

IBM Tivoli Storage Manager for Windows Backup-Archive
Clients
Version 6.3

Installation and User's Guide



IBM Tivoli Storage Manager for Windows Backup-Archive
Clients
Version 6.3

Installation and User's Guide



Note:

Before using this information and the product it supports, read the information in "Notices" on page 655.

This edition applies to Version 6.3 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, and 5608-E07), and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters. This edition replaces SC23-9792-04.

Changes since the previous edition are marked with a vertical bar (|) in the left margin. Ensure that you are using the correct edition for the level of the product.

© **Copyright IBM Corporation 1993, 2011.**

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

Contents

Tables xi

About this publication xiii

| | |
|--|------|
| Who should read this publication. | xiii |
| Publications | xiii |
| Tivoli Storage Manager publications | xiv |
| Support information | xvi |
| Getting technical training | xvi |
| Searching knowledge bases | xvii |
| Contacting IBM Software Support | xix |
| Conventions used in this publication | xx |
| Reading syntax diagrams. | xx |

New in version 6.3 xxiii

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

| | |
|--|----|
| Upgrading the backup-archive client | 1 |
| Upgrade path for clients and servers | 1 |
| Considerations for migrating between processor architectures x32 and x64 | 2 |
| Additional upgrade information | 3 |
| Upgrading Open File Support or online image | 3 |
| Automatic backup-archive client deployment | 4 |
| Client environment requirements | 5 |
| Windows client environment requirements | 6 |
| NDMP support requirements (Extended Edition only) | 8 |
| Installation requirements for backing up and archiving Tivoli Storage Manager FastBack client data | 8 |
| Client configuration wizard for Tivoli Storage Manager FastBack | 9 |
| Installing the Tivoli Storage Manager Windows client | 9 |
| Silent installation | 11 |
| Upgrading, modifying, reinstalling, or uninstalling the Tivoli Storage Manager Windows client | 15 |
| Troubleshooting problems during installation | 16 |

Chapter 2. Configure the Tivoli Storage Manager client 19

| | |
|--|----|
| Creating and modifying the client options file | 19 |
| Create a shared directory options file | 21 |
| Creating multiple options files | 21 |
| Environment variables. | 22 |
| Configuring the Web client | 22 |
| Configuring the client scheduler | 24 |
| Starting the client scheduler | 26 |
| Scheduling events using the GUI | 26 |
| Configuring Tivoli Storage Manager client/server communication across a firewall | 27 |
| Configuring Tivoli Storage Manager client/server communication with Secure Sockets Layer | 29 |
| Certificate Authorities root certificates | 32 |
| Configure your system for journal-based backup | 33 |

| | |
|---|----|
| Configuring the journal engine service | 33 |
| Client-side data deduplication | 40 |
| Configuring the client for data deduplication | 44 |
| Excluding files from data deduplication | 46 |
| Configuring the client to back up and archive Tivoli Storage Manager FastBack data. | 47 |
| Configuring Tivoli Storage Manager to protect FastBack client data | 48 |
| Configuring the backup-archive client in a cluster server environment. | 50 |
| Protecting data in MSCS and VCS clusters (Windows Server 2003) | 50 |
| Protecting data in MSCS clusters (Windows Server 2008) | 56 |
| Configure the Web client in a cluster environment | 57 |
| Frequently asked questions | 63 |
| Configuring online-image backup support | 65 |
| Configuring Open File Support. | 66 |
| Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups | 67 |
| Register your workstation with a server | 68 |
| Closed registration | 69 |
| Open registration | 69 |
| Creating an include-exclude list | 69 |
| Include-exclude options | 71 |
| Determine compression and encryption processing | 78 |
| Preview include-exclude list files | 79 |
| Include and exclude option processing | 80 |
| Processing rules when using UNC names | 82 |

Chapter 3. Getting started 83

| | |
|--|----|
| Backup-archive client operations and security rights | 83 |
| Backup Operators group operations | 84 |
| Considerations before you start using a Backup Operators group account | 85 |
| Permissions required to restore files that use adaptive subfile backup | 85 |
| Permissions required to back up, archive, restore or retrieve files on cluster resources | 86 |
| Tivoli Storage Manager client authentication | 86 |
| User account control | 87 |
| Starting a Java GUI session | 87 |
| Tivoli Storage Manager password | 88 |
| Setup wizard | 88 |
| Starting a command-line session | 88 |
| Using batch mode | 89 |
| Issuing a series of commands by using interactive mode. | 89 |
| Displaying Euro characters in a command-line prompt. | 90 |
| Use options on the DSMC command | 90 |
| Specifying input strings that contain blank spaces or quotation marks. | 91 |

| | |
|--|-----------|
| Starting a web client session | 91 |
| User privileges | 92 |
| Start the client scheduler automatically | 93 |
| Changing your password | 93 |
| Sorting file lists using the Tivoli Storage Manager GUI | 94 |
| Displaying online help | 95 |
| Ending a session | 95 |
| Online forums | 96 |
| Other sources of online help | 96 |
| Chapter 4. Backing up your data | 97 |
| Planning your backups (Windows) | 97 |
| Understanding which files are backed up | 98 |
| Open file support for backup operations | 99 |
| Backing up data using the GUI | 101 |
| Specifying drives in your domain | 102 |
| Backing up data using the command line | 103 |
| Deleting backup data | 105 |
| When to back up and when to archive files | 105 |
| Pre-backup considerations (Windows) | 106 |
| LAN-free data movement | 106 |
| Unicode file spaces (Windows) | 107 |
| Incremental backups on memory-constrained systems | 108 |
| Incremental backups on systems with a large number of files | 108 |
| Performing a backup with limited bandwidth | 109 |
| Control processing with an include-exclude list | 110 |
| Data encryption during backup or archive operations | 110 |
| Maximum file size for operations | 111 |
| How Tivoli Storage Manager handles long user and group names | 111 |
| Incremental, selective, or incremental-by-date backups (Windows) | 112 |
| Full and partial incremental backup | 112 |
| Incremental-by-date backup | 115 |
| Comparing incremental-by-date, journal-based, and NetApp snapshot difference to full incremental and partial incremental backups | 116 |
| Selective backup | 118 |
| Backing up files from one or more file spaces for a group backup (Windows) | 118 |
| Backing up data with client-node proxy support (Windows) | 119 |
| Enabling multiple node operations from the GUI | 120 |
| Setting up encryption | 120 |
| Scheduling backups with client-node proxy support | 120 |
| Associate a local snapshot with a server file space (Windows) | 122 |
| Backing up Windows system state | 122 |
| Backing up Automated System Recovery files | 124 |
| Backing up in preparation for Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008) | 125 |
| Creating a Tivoli Storage Manager options file for Automated System Recovery (Windows Vista, Windows 7, or Windows Server 2008) | 125 |

| | |
|---|-----|
| Backing up the boot drive and system drive for Automated System Recovery (Windows Vista, Windows 7, or Windows Server 2008) | 126 |
| Backing up in preparation for Automated System Recovery (Windows Server 2003) | 127 |
| Creating an Automated System Recovery diskette (Windows Server 2003) | 127 |
| Backing up boot drive and system drive for Automated System Recovery (Windows Server 2003) | 128 |
| Backing up user profiles for ASR recovery (Windows Server 2003) | 128 |
| Image backup | 129 |
| Performing prerequisite tasks before creating an image backup | 130 |
| Utilizing image backups to perform file system incremental backups | 131 |
| Performing an image backup using the GUI | 133 |
| Performing an image backup using the command line | 134 |
| Back up NAS file systems using Network Data Management Protocol | 135 |
| Backing up NAS file systems with the Web client GUI using NDMP protocol | 136 |
| Back up NAS file systems using the command line | 137 |
| Methods for backing up and recovering data on NAS file servers when using CIFS | 139 |
| Support for CDP Persistent Storage Manager | 140 |
| Backing up VMware virtual machine data | 141 |
| Scenario: Backing up your virtual machines (Windows) | 141 |
| Preparing the environment for VMware backup processing | 143 |
| Running full virtual machine backups | 145 |
| Backing up virtual machines on a Windows Hyper-V system | 146 |
| Hyper-V backup support limitations | 147 |
| Back up and archive Tivoli Storage Manager FastBack data | 148 |
| Net Appliance CIFS share definition backups | 149 |
| Display backup processing status | 149 |
| Backup (Windows): Additional considerations | 152 |
| Open files | 152 |
| Ambiguous file space names in file specifications | 152 |
| Management classes | 153 |
| Deleted file systems | 153 |
| Removable media backup | 154 |
| Fixed drives | 155 |
| NTFS file spaces | 155 |
| Universal Naming Convention names | 155 |
| Microsoft Dfs file protection methods | 157 |

| | |
|---|------------|
| Chapter 5. Restoring your data | 161 |
| Duplicate file names | 161 |
| Universal Naming Convention names restore | 162 |
| Active or inactive backup restore | 162 |
| Restoring files and directories | 163 |
| Restoring data using the GUI | 163 |

| | |
|--|-----|
| Examples for restoring data using the command line | 163 |
| Plan for temporary space to restore migrated files | 167 |
| Restoring Windows system state | 167 |
| Restoring Automated System Recovery files | 169 |
| Restoring the operating system when the computer is working | 170 |
| Recovering a computer when the operating system is not working (Windows Vista, Windows 7, and Windows Server 2008) | 170 |
| Creating a bootable WinPE CD (Windows Vista, Windows 7, and Windows Server 2008). | 170 |
| Restoring the operating system with Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008) | 172 |
| Recovering a computer when the operating system is not working (Windows Server 2003) | 173 |
| Hardware and software requirements for ASR recovery (Windows Server 2003) | 174 |
| Creating a Tivoli Storage Manager client installation CD for ASR recovery (Windows Server 2003) | 175 |
| Restoring the operating system with Automated System Recovery (Windows Server 2003) | 175 |
| Microsoft Dfs tree and file restore | 177 |
| Restoring an image | 177 |
| Restoring an image using the GUI | 179 |
| Restoring an image using the command line | 180 |
| Restore data from a backup set | 180 |
| Restore backup sets: considerations and restrictions | 182 |
| Backup set restore | 183 |
| Restoring backup sets using the GUI | 184 |
| Backup set restores using the client command-line interface | 185 |
| Restore Net Appliance CIFS shares | 186 |
| Restoring data from a VMware backup | 186 |
| Restoring file-level VM backups | 187 |
| Restoring full VM backups for VMware Consolidated Backup (VCB) | 188 |
| Restoring full VM backups | 190 |
| Restore Windows individual Active Directory objects | 191 |
| Reanimate tombstone objects or restoring from a system state backup | 191 |
| Restoring Active Directory objects using the GUI and command line | 192 |
| Restrictions and limitations when restoring Active Directory objects | 193 |
| Preserve attributes in tombstone objects | 194 |
| Modifying the client acceptor and agent services to use the Web client | 195 |
| Authorizing another user to restore or retrieve your files | 195 |
| Restoring or retrieving files from another client node | 196 |
| Restoring or retrieving your files to another workstation | 197 |
| Deleting file spaces | 198 |
| Restoring data to a point in time | 199 |

| | |
|--|-----|
| Restore NAS file systems | 200 |
| Restoring NAS file systems using the Web client | 200 |
| Restoring NAS files and directories using the Web client | 201 |
| Options and commands to restore NAS file systems from the command line | 203 |

Chapter 6. Archive and retrieve your data (Windows) 205

| | |
|---|-----|
| Archive files | 205 |
| Snapshot backup or archive with open file support | 206 |
| Archiving data using the GUI | 206 |
| Archive data examples using the command line | 207 |
| Archiving data with client node proxy | 208 |
| Deleting archive data | 210 |
| Retrieve archives | 210 |
| Retrieving archives using the GUI | 211 |
| Retrieve archives using the command line | 211 |

Chapter 7. Tivoli Storage Manager scheduler overview 213

| | |
|---|-----|
| Examples: Blank spaces in file names in schedule definitions | 214 |
| Preferential start times for certain nodes | 214 |
| Scheduler processing options | 215 |
| Return codes from preschedulecmd and postschedulecmd scripts | 216 |
| Client-acceptor scheduler services versus the traditional scheduler services | 217 |
| Set the client scheduler process to run as a background task and start automatically at startup | 217 |
| Examples: Display information about scheduled work | 219 |
| Display information about completed work | 221 |
| Examples: Event logs | 222 |
| Specify scheduling options | 225 |
| Enable or disable scheduled commands | 225 |
| Change processing options used by the scheduler service | 226 |
| Manage multiple schedule requirements on one system | 227 |

Chapter 8. Client return codes 229

Chapter 9. Storage management policies 231

| | |
|--|-----|
| Policy domains and policy sets | 231 |
| Management classes and copy groups | 232 |
| Display information about management classes and copy groups | 233 |
| Copy group name attribute | 233 |
| Copy type attribute | 233 |
| Copy frequency attribute | 234 |
| Versions data exists attribute | 234 |
| Versions data deleted attribute | 234 |
| Retain extra versions attribute | 234 |
| Retain only version attribute | 235 |
| Copy serialization attribute | 235 |

| | |
|--|-----|
| Copy mode attribute | 236 |
| Copy destination attribute | 236 |
| Retain versions attribute. | 236 |
| Deduplicate data attribute | 237 |
| Select a management class for files | 237 |
| Assign a management class to files | 237 |
| Override the management class for archived files | 238 |
| Select a management class for directories | 239 |
| Bind management classes to files | 239 |
| Rebind backup versions of files | 240 |
| Retention grace period | 240 |
| Event-based policy retention protection. | 241 |
| Archive files on a data retention server. | 242 |

Chapter 10. Using the Client Service Configuration Utility 243

| | |
|--|-----|
| Install the backup-archive scheduler service | 243 |
| Using the Client Service Configuration Utility (Windows) | 243 |
| Use the Dsmcutil command | 246 |
| Dsmcutil commands: Required options and examples | 247 |
| Dsmcutil valid options | 254 |

Chapter 11. Windows Server 2003 ASR supplemental information (Windows). 259

| | |
|--|-----|
| Windows Server 2003 ASR questions and answers | 259 |
| Windows Server 2003 ASR diagnostic tips | 262 |
| Windows Server 2003 ASR in depth | 263 |
| Windows Server 2003 ASR information in the Windows Registry. | 266 |
| Windows Server 2003 ASR Reference information | 267 |

Chapter 12. Processing options 269

| | |
|---|-----|
| Processing options overview | 269 |
| Communication options | 270 |
| TCP/IP options | 270 |
| Named Pipes option | 271 |
| Shared memory options | 271 |
| Server and node options. | 271 |
| Node options | 272 |
| Backup and archive processing options. | 272 |
| Restore and retrieve processing options. | 282 |
| Scheduling options | 285 |
| Format and language options | 286 |
| Command processing options | 286 |
| Authorization options | 287 |
| Error processing options. | 287 |
| Transaction processing options | 288 |
| Web client options. | 289 |
| Diagnostics options | 289 |
| Using options with commands | 289 |
| Entering options with a command | 290 |
| Initial command-line-only options | 295 |
| Client options that can be set by the Tivoli Storage Manager server | 296 |
| Client options reference | 298 |
| Adlocation | 298 |
| Archmc. | 298 |

| | |
|---|-----|
| Asnodename | 299 |
| Asrmode | 301 |
| Auditlogging | 301 |
| Auditlogname | 304 |
| Autodeploy | 306 |
| Autofsrename | 307 |
| Backmc. | 309 |
| Backupsetname | 310 |
| Casesensitiveaware | 311 |
| Changingretries | 312 |
| Class | 313 |
| Clientview | 313 |
| Clusterdisksonly | 314 |
| Clusternode. | 317 |
| Collocatebyfilespec. | 319 |
| Commmethod | 320 |
| Commrestartduration. | 321 |
| Commrestartinterval. | 321 |
| Compressalways. | 322 |
| Compression | 323 |
| Computername | 324 |
| Console | 325 |
| Createnewbase | 326 |
| Datacenter | 328 |
| Datastore. | 329 |
| Dateformat | 329 |
| Dedupcachepath. | 331 |
| Dedupcachesize. | 332 |
| Deduplication | 332 |
| Deletefiles. | 334 |
| Description. | 334 |
| Detail. | 335 |
| Dfsbackupmntpnt | 336 |
| Diffsnapshot | 337 |
| Dirmc | 338 |
| Dirsonly | 340 |
| Disablenqr | 340 |
| Diskbuffsize | 341 |
| Diskcachelocation | 342 |
| Domain. | 343 |
| Domain.image | 346 |
| Domain.nas | 346 |
| Domain.vmfile | 348 |
| Domain.vmfull | 350 |
| Enable8dot3namesupport. | 352 |
| Enablearchiveretentionprotection | 353 |
| Enabledupcache | 354 |
| Enablelanfree | 355 |
| Encryptiontype. | 356 |
| Encryptkey | 357 |
| Errorlogmax. | 359 |
| Errorlogname | 361 |
| Errorlogretention | 361 |
| Determine compression and encryption processing | 363 |
| Exclude options | 364 |
| Fbranch | 369 |
| Fbclientname | 370 |
| Fbpolicyname | 371 |
| Fbreposlocation | 373 |
| Fbserver | 374 |

| | | | |
|----------------------------------|-----|-----------------------------|-----|
| Fbvolumentname | 375 | Revokeremoteaccess | 445 |
| Filelist | 377 | Runasservice | 446 |
| Filename | 379 | Schedcmddisabled | 447 |
| Filesonly | 381 | Schedcmdexception | 448 |
| Fromdate | 381 | Schedlogmax | 449 |
| Fromnode | 382 | Schedlogname | 450 |
| Fromtime | 383 | Schedlogretention | 451 |
| Frsprimaryrestore | 384 | Schedmode | 453 |
| Groupname | 384 | Schedrestretrdisabled | 454 |
| Guitreeviewafterbackup | 385 | Scrolllines | 455 |
| Host | 386 | Scrollprompt | 456 |
| Httpport | 386 | Sessioninitiation | 457 |
| Hsmreparsetag | 386 | Shmport | 459 |
| Ieobjtype | 388 | Showmembers | 460 |
| Ifnewer | 389 | Skipmissingsyswfiles | 460 |
| Imagegapsize | 389 | Skipntpermissions | 461 |
| Imagetofile | 390 | Skipntsecuritycrc | 462 |
| Inactive | 391 | Snapdiff | 463 |
| Incl excl | 392 | Snapshotcachelocation | 467 |
| Include options | 393 | Snapshotcachesize | 468 |
| Incrbydate | 399 | Snapshotfsidleretries | 469 |
| Incremental | 400 | Snapshotfsidlewait | 470 |
| Incrthreshold | 401 | Snapshotproviderfs | 472 |
| Journalpipe | 402 | Snapshotproviderimage | 473 |
| Lanfreecommmethod | 402 | Snapshotroot | 474 |
| Lanfreeshmport | 404 | Srvoptsetencryptiondisabled | 477 |
| Lanfreetcppport | 404 | Srvprepostscheddisabled | 478 |
| Lanfreetcpsserveraddress | 405 | Srvprepostsnapdisabled | 479 |
| Language | 406 | Ssl | 480 |
| Latest | 407 | Sslfipsmode | 481 |
| Localbackupset | 408 | Stagingdirectory | 481 |
| Managedservices | 409 | Subdir | 482 |
| Maxcmdretries | 410 | Subfilebackup | 484 |
| Memoryefficientbackup | 411 | Subfilecachepath | 485 |
| Mode | 412 | Subfilecachesize | 486 |
| Monitor | 414 | Tapeprompt | 486 |
| Namedpipename | 415 | Tcpadminport | 488 |
| Nasnodename | 416 | Tcpbuffsize | 488 |
| Nodename | 417 | Tcpcadaddress | 489 |
| Nojournal | 418 | Tcpclientaddress | 490 |
| Noprompt | 419 | Tcpclientport | 491 |
| Numberformat | 419 | Tcpnodeelay | 491 |
| Optfile | 421 | Tcpport | 492 |
| Password | 422 | Tcpserveraddress | 493 |
| Passwordaccess | 423 | Tcpwindowsize | 494 |
| Pick | 424 | Timeformat | 495 |
| Pitdate | 425 | Toc | 496 |
| Pittime | 426 | Todate | 497 |
| Postschedulecmd/Postnschedulecmd | 427 | Totime | 498 |
| Postsnapshotcmd | 428 | Txnbytelimit | 499 |
| Preschedulecmd/Prenschedulecmd | 429 | Type | 500 |
| Preserveleastaccessdate | 431 | Usedirectory | 501 |
| Preservepath | 432 | V2archive | 502 |
| Presnapshotcmd | 434 | Verbose | 503 |
| Queryschedperiod | 436 | Verifyimage | 504 |
| Querysummary | 437 | Virtualfsname | 504 |
| Quiet | 438 | Virtualnodename | 505 |
| Replace | 439 | Vmbackdir | 506 |
| Resetarchiveattribute | 440 | Vmbacknodelete | 507 |
| Resourceutilization | 442 | Vmbackuptype | 508 |
| Retryperiod | 444 | Vmmc | 509 |

| | |
|--------------------------------------|-----|
| Vmbackvcbtransport | 510 |
| Vmchost | 511 |
| Vmcpw | 512 |
| Vmctlmc | 513 |
| Vmcuser | 513 |
| VmList | 514 |
| Vmprocessvmwithindependent | 515 |
| Vmprocessvmwithprdm | 516 |
| Vmvstortransport | 517 |
| Vssusesystemprovider | 518 |
| Webports | 519 |

Chapter 13. Using commands 521

| | |
|--|-----|
| Start and end a client command session | 524 |
| Process commands in batch mode | 524 |
| Process commands in interactive mode | 525 |
| Enter client command names, options, and parameters | 526 |
| Command name | 526 |
| Options | 526 |
| Parameters | 527 |
| File specification syntax | 527 |
| Wildcard characters | 528 |
| Client commands reference | 529 |
| Archive | 529 |
| Open file support | 531 |
| Archive FastBack | 532 |
| Backup ASR | 535 |
| Backup FastBack | 536 |
| Backup Group | 538 |
| Backup Image | 540 |
| Offline and online image backup | 543 |
| Utilizing image backup to perform file system incremental backup | 544 |
| Backup NAS | 545 |
| Backup Systemstate | 547 |
| Backup VM | 549 |
| Cancel Process | 552 |
| Cancel Restore | 552 |
| Delete Access | 553 |
| Delete Archive | 554 |
| Delete Backup | 556 |
| Delete Filespace | 558 |
| Delete Group | 559 |
| Expire | 561 |
| Help | 562 |
| Incremental | 564 |
| Open file support | 567 |
| Journal-based backup | 568 |
| Backing up NTFS volume mount points | 569 |
| Back up Microsoft Dfs root | 570 |
| Incremental-by-Date | 571 |
| Adaptive subfile backups | 571 |
| Associate a local snapshot with a server file space | 571 |
| Loop | 572 |
| Macro | 573 |
| Monitor Process | 573 |
| Preview Archive | 574 |
| Preview Backup | 575 |
| Query Access | 576 |

| | |
|---|-----|
| Query Adobjects | 577 |
| Query Archive | 578 |
| Query ASR | 581 |
| Query Backup | 581 |
| Query NAS file system images | 585 |
| Query Backupset | 585 |
| Query Backupset without the backupsetname parameter | 586 |
| Query Filespace | 588 |
| Query NAS file spaces | 590 |
| Query Group | 590 |
| Query Image | 591 |
| Query Inclxcl | 593 |
| Query Mgmtclass | 594 |
| Query Node | 595 |
| Query Options | 596 |
| Query Restore | 597 |
| Query Schedule | 598 |
| Query Session | 598 |
| Query Systeminfo | 599 |
| Query Systemstate | 601 |
| Query VM | 602 |
| Restart Restore | 604 |
| Restore | 605 |
| Restoring NTFS volume mount points | 610 |
| Restore Microsoft Dfs junctions | 611 |
| Restore active files | 611 |
| Universal Naming Convention restores | 611 |
| Restore from file spaces that are not Unicode-enabled | 612 |
| Restore named streams | 612 |
| Restore sparse files | 612 |
| Restore Adobjects | 612 |
| Restore ASR | 614 |
| Restore Backupset | 615 |
| Restore backup sets: considerations and restrictions | 618 |
| Restore backup sets in a SAN environment | 619 |
| Restore Backupset without the backupsetname parameter | 620 |
| Restore Group | 622 |
| Restore Image | 624 |
| Restore NAS | 627 |
| Restore Systemstate | 629 |
| Restore VM | 631 |
| Retrieve | 635 |
| Retrieve archives from file spaces that are not Unicode-enabled | 638 |
| Retrieve named streams | 638 |
| Retrieve sparse files | 639 |
| Schedule | 639 |
| Selective | 640 |
| Open file support | 643 |
| Adaptive subfile backups | 644 |
| Associate a local snapshot with a server file space | 644 |
| Set Access | 644 |
| Set Event | 646 |
| Set Password | 648 |

| | | | |
|--|------------|---------------------------|------------|
| Appendix. Accessibility features for the Tivoli Storage Manager product family. | 653 | Trademarks | 657 |
| Notices | 655 | Glossary | 659 |
| | | Index | 681 |

Tables

| | | | |
|--|-----|---|-----|
| 1. Tivoli Storage Manager server publications | xiv | 44. Format and language options | 286 |
| 2. Tivoli Storage Manager storage agent publications | xiv | 45. Command processing options | 287 |
| 3. Tivoli Storage Manager client publications | xv | 46. Authorization options | 287 |
| 4. Tivoli Storage Manager data protection publications | xv | 47. Error processing options | 287 |
| 5. IBM Tivoli Storage Manager troubleshooting and tuning publications | xv | 48. Transaction processing options | 288 |
| 6. Windows client communication methods | 7 | 49. Web client options | 289 |
| 7. Supported features on Windows platforms | 8 | 50. Diagnostics options | 289 |
| 8. Stoppable services | 15 | 51. Client command options | 291 |
| 9. Data deduplication settings: Client and server | 43 | 52. Options that are valid on the initial command line only | 296 |
| 10. Options for excluding file spaces and directories | 72 | 53. Options that can be set by the Tivoli Storage Manager server | 297 |
| 11. Options for controlling processing using include and exclude statements | 73 | 54. Clusternode and clusterdiskonly combinations | 317 |
| 12. Wildcard and other special characters | 76 | 55. Interaction of domain definitions from several sources | 345 |
| 13. Specifying a drive specification using wildcards | 77 | 56. Options for controlling compression and encryption processing | 363 |
| 14. Using wildcard characters with include and exclude patterns | 78 | 57. System services components and corresponding keywords | 366 |
| 15. Options for controlling compression and encryption processing | 79 | 58. Other optional parameters | 397 |
| 16. UNC name patterns and DOS patterns | 82 | 59. Incremental command: Related options | 464 |
| 17. Required user security rights for Tivoli Storage Manager backup and restore services | 83 | 60. Commands | 521 |
| 18. Working with your files using the Tivoli Storage Manager GUI | 94 | 61. Wildcard characters | 529 |
| 19. Planning your backups | 97 | 62. Archive command: Related options | 530 |
| 20. Command line backup examples | 103 | 63. Archive FastBack command: Related options | 533 |
| 21. Maximum file size | 111 | 64. Backup FastBack command: Related options | 537 |
| 22. Comparing incremental image backup methods | 133 | 65. Backup Group command: Related options | 540 |
| 23. NAS options and commands | 137 | 66. Backup Image command: Related options | 542 |
| 24. Client command line informational messages | 150 | 67. Backup NAS command: Related options | 547 |
| 25. UNC examples | 156 | 68. Delete Archive command: Related options | 555 |
| 26. Command-line restore examples | 163 | 69. Delete Backup command: Related options | 557 |
| 27. Backup set GUI restore restrictions | 181 | 70. Delete Filespace command: Related options | 559 |
| 28. Backup set command-line restore restrictions | 182 | 71. Delete Group command: Related options | 560 |
| 29. NAS options and commands | 203 | 72. Expire command: Related options | 562 |
| 30. Command-line archive examples | 207 | 73. Incremental command: Related options | 565 |
| 31. Command line examples of retrieving archives | 211 | 74. Query Adobjects command: Related options | 577 |
| 32. Sample classic query schedule output | 220 | 75. Query Archive command: Related options | 579 |
| 33. Sample enhanced query schedule output | 221 | 76. Query ASR command: Related options | 581 |
| 34. Client return codes and their meanings | 229 | 77. Query Backup command: Related options | 583 |
| 35. Default attribute values in the standard management class | 233 | 78. Query Backupset command: Related options | 586 |
| 36. TCP/IP options | 270 | 79. Query Backupset command: Related options | 587 |
| 37. Named Pipes communication option | 271 | 80. Query Filespace command: Related options | 589 |
| 38. Shared memory communication options | 271 | 81. Query Group command: Related options | 591 |
| 39. Server and node options | 271 | 82. Query Image command: Related options | 592 |
| 40. Node options | 272 | 83. Query Mgmtclass command: Related options | 595 |
| 41. Backup and archive processing options | 272 | 84. Query Node command: Related options | 596 |
| 42. Restore and retrieve processing options | 282 | 85. Query Options command: Related options | 596 |
| 43. Scheduling options | 285 | 86. Query Systeminfo command: Related options | 600 |
| | | 87. Query Systemstate command: Related options | 601 |
| | | 88. Query VM command: Related options | 603 |
| | | 89. Restore command: Related options | 607 |
| | | 90. Restore Adobjects command: Related options | 613 |
| | | 91. Restore ASR command: Related options | 615 |
| | | 92. Restore Backupset command: Related options | 617 |
| | | 93. Restore Group command: Related options | 623 |

| | | | |
|--|-----|---|-----|
| 94. Restore Image command: Related options | 626 | 97. Restore VM command: Related options | 633 |
| 95. Restore NAS command: Related options | 628 | 98. Retrieve command: Related options | 636 |
| 96. Restore Systemstate command: Related options | 630 | 99. Schedule command: Related options | 640 |
| | | 100. Selective command: Related options | 641 |

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 permits users to back up and archive files from their workstations or file servers to storage, and restore and retrieve backup versions and archived copies of files to their local workstations.

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

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

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.

Related information

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

Chapter 2, “Configure the Tivoli Storage Manager client,” on page 19

Publications

Publications for the IBM Tivoli Storage Manager family of products are available online. The IBM 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 Tivoli.

To search all publications, go to the Tivoli Storage Manager information center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>.

You can download PDF versions of publications from the Tivoli Storage Manager information center or from the IBM Publications Center at <http://www.ibm.com/shop/publications/order/>.

Go to Tivoli Documentation Central to find information centers that contain official product documentation for current and previous versions of Tivoli products, including the Tivoli Storage Manager product family. You can find Tivoli Documentation Central at <https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Home>.

You can also order some related publications from the IBM Publications Center website. The website provides information about ordering publications from countries other than the United States. In the United States, you can order publications by calling 1-800-879-2755.

Tivoli Storage Manager publications

The following tables list the publications that make up the Tivoli Storage Manager library.

Table 1. Tivoli Storage Manager server publications

| Publication title | Order number |
|--|---------------------|
| <i>IBM Tivoli Storage Manager for AIX Installation Guide</i> | GC23-9781 |
| <i>IBM Tivoli Storage Manager for AIX Administrator's Guide</i> | SC23-9769 |
| <i>IBM Tivoli Storage Manager for AIX Administrator's Reference</i> | SC23-9775 |
| <i>IBM Tivoli Storage Manager for HP-UX Installation Guide</i> | GC23-9782 |
| <i>IBM Tivoli Storage Manager for HP-UX Administrator's Guide</i> | SC23-9770 |
| <i>IBM Tivoli Storage Manager for HP-UX Administrator's Reference</i> | SC23-9776 |
| <i>IBM Tivoli Storage Manager for Linux Installation Guide</i> | GC23-9783 |
| <i>IBM Tivoli Storage Manager for Linux Administrator's Guide</i> | SC23-9771 |
| <i>IBM Tivoli Storage Manager for Linux Administrator's Reference</i> | SC23-9777 |
| <i>IBM Tivoli Storage Manager for Oracle Solaris Installation Guide</i> | GC23-9784 |
| <i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Guide</i> | SC23-9772 |
| <i>IBM Tivoli Storage Manager for Oracle Solaris Administrator's Reference</i> | SC23-9778 |
| <i>IBM Tivoli Storage Manager for Windows Installation Guide</i> | GC23-9785 |
| <i>IBM Tivoli Storage Manager for Windows Administrator's Guide</i> | SC23-9773 |
| <i>IBM Tivoli Storage Manager for Windows Administrator's Reference</i> | SC23-9779 |
| <i>IBM Tivoli Storage Manager for z/OS Media Installation and User's Guide</i> | SC27-4018 |
| <i>IBM Tivoli Storage Manager Upgrade and Migration Guide for V5 Servers</i> | GC27-4017 |
| <i>IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack®</i> | SC27-2828 |

Table 2. Tivoli Storage Manager storage agent publications

| Publication title | Order number |
|--|---------------------|
| <i>IBM Tivoli Storage Manager for SAN for AIX Storage Agent User's Guide</i> | SC23-9797 |

Table 2. Tivoli Storage Manager storage agent publications (continued)

| Publication title | Order number |
|---|---------------------|
| <i>IBM Tivoli Storage Manager for SAN for HP-UX Storage Agent User's Guide</i> | SC23-9798 |
| <i>IBM Tivoli Storage Manager for SAN for Linux Storage Agent User's Guide</i> | SC23-9799 |
| <i>IBM Tivoli Storage Manager for SAN for Oracle Solaris Storage Agent User's Guide</i> | SC23-9800 |
| <i>IBM Tivoli Storage Manager for SAN for Windows Storage Agent User's Guide</i> | SC23-9553 |

Table 3. Tivoli Storage Manager client publications

| Publication title | Order number |
|--|---------------------|
| <i>IBM Tivoli Storage Manager for UNIX and Linux: Backup-Archive Clients Installation and User's Guide</i> | SC23-9791 |
| <i>IBM Tivoli Storage Manager for Windows: Backup-Archive Clients Installation and User's Guide</i> | SC23-9792 |
| <i>IBM Tivoli Storage Manager Using the Application Programming Interface</i> | SC23-9793 |
| <i>IBM Tivoli Storage Manager for Space Management for UNIX and Linux: User's Guide</i> | SC23-9794 |
| <i>IBM Tivoli Storage Manager HSM for Windows Administration Guide</i> | SC23-9795 |

Table 4. Tivoli Storage Manager data protection publications

| Publication title | Order number |
|--|---------------------|
| <i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Installation and User's Guide</i> | GC27-4010 |
| <i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for UNIX and Linux Installation and User's Guide</i> | SC27-4019 |
| <i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle for Windows Installation and User's Guide</i> | SC27-4020 |
| <i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Installation and User's Guide</i> | GC27-4009 |
| <i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino® UNIX and Linux Installation and User's Guide</i> | SC27-4021 |
| <i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino for Windows Installation and User's Guide</i> | SC27-4022 |
| <i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for DB2</i> | SC33-6341 |
| <i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Installation and User's Guide for Oracle</i> | SC33-6340 |
| <i>IBM Tivoli Storage Manager for Virtual Environments Installation and User's Guide</i> | SC27-2898 |
| <i>IBM Tivoli Storage Manager for Microsoft SharePoint Guide</i> | N/A |

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications

| Publication title | Order number |
|---|---------------------|
| <i>IBM Tivoli Storage Manager Problem Determination Guide</i> | GC23-9789 |

Table 5. IBM Tivoli Storage Manager troubleshooting and tuning publications (continued)

| Publication title | Order number |
|--|--------------|
| <i>IBM Tivoli Storage Manager Performance Tuning Guide</i> | GC23-9788 |
| <i>IBM Tivoli Storage Manager Client Messages and Application Programming Interface Return Codes</i> | SC27-2878 |
| <i>IBM Tivoli Storage Manager Server Messages and Error Codes</i> | SC27-2877 |
| <i>IBM Tivoli Storage Manager for Mail: Data Protection for Microsoft Exchange Server Messages</i> | GC27-4011 |
| <i>IBM Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server Messages</i> | GC27-4012 |
| <i>IBM Tivoli Storage Manager for Databases: Data Protection for Oracle Messages</i> | SC27-4014 |
| <i>IBM Tivoli Storage Manager for Mail: Data Protection for Lotus Domino Messages</i> | SC27-4015 |
| <i>IBM Tivoli Storage Manager for Enterprise Resource Planning: Data Protection for SAP Messages</i> | SC27-4016 |

Note: You can find information about IBM System Storage® Archive Manager at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/c_complydataretention_ovr.html.

Support information

You can find support information for IBM products from various sources.

Start at the IBM Support Portal: <http://www.ibm.com/support/entry/portal/>. You can select the products that you are interested in and search for a wide variety of relevant information.

Getting technical training

Information about Tivoli technical training courses is available online.

Go to the following websites to sign up for training, ask questions, and interact with others who use IBM storage products.

Tivoli software training and certification

Choose from instructor led, online classroom training, self-paced Web classes, Tivoli certification preparation, and other training options at <http://www.ibm.com/software/tivoli/education/>

Tivoli Support Technical Exchange

Technical experts share their knowledge and answer your questions in webcasts at http://www.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html.

Storage Management community

Interact with others who use IBM storage management products at <http://www.ibm.com/developerworks/servicemanagement/sm/index.html>

Global Tivoli User Community

Share information and learn from other Tivoli users throughout the world at <http://www.tivoli-ug.org/>.

IBM Education Assistant

View short "how to" recordings designed to help you use IBM software products more effectively at <http://publib.boulder.ibm.com/infocenter/ieduasst/tivv1r0/index.jsp>

Searching knowledge bases

If you have a problem with your Tivoli Storage Manager family product, there are several knowledge bases that you can search.

Begin by searching the Tivoli Storage Manager Information Center at <http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3>. From this website, you can search the current Tivoli Storage Manager documentation.

Searching the Internet

If you cannot find an answer to your question in the IBM Tivoli Storage Manager information center, search the Internet for the information that might help you resolve your problem.

To search multiple Internet resources, go to the IBM support website at <http://www.ibm.com/support/entry/portal/>.

You can search for information without signing in. Sign in using your IBM ID and password if you want to customize the site based on your product usage and information needs. If you do not already have an IBM ID and password, click **Sign in** at the top of the page and follow the instructions to register.

From the support website, you can search various resources including:

- IBM technotes
- IBM downloads
- IBM Redbooks® publications
- IBM Authorized Program Analysis Reports (APARs)

Select the product and click **Downloads** to search the APAR list.

If you still cannot find a solution to the problem, you can search forums and newsgroups on the Internet for the latest information that might help you find problem resolution.

An independent user discussion list, ADSM-L, is hosted by Marist College. You can subscribe by sending an email to listserv@vm.marist.edu. The body of the message must contain the following text: `SUBSCRIBE ADSM-L your_first_name your_family_name`.

To share your experiences and learn from others in the Tivoli Storage Manager and Tivoli Storage FlashCopy Manager user communities, go to the following wikis:

Tivoli Storage Manager wiki

<http://www.ibm.com/developerworks/wikis/display/tivolistoragemanager>

Tivoli Storage FlashCopy Manager wiki

[https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Storage FlashCopy Manager](https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli%20Storage%20FlashCopy%20Manager)

Using IBM Support Assistant

IBM Support Assistant is a complimentary software product that can help you with problem determination. It is available for some Tivoli Storage Manager and Tivoli Storage FlashCopy Manager products.

To learn about which products are supported, go to the IBM Support Assistant download web page at <http://www.ibm.com/software/support/isa/download.html>.

IBM Support Assistant helps you gather support information when you must open a problem management record (PMR), which you can then use to track the problem. The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

You can find more information at the IBM Support Assistant website:

<http://www.ibm.com/software/support/isa/>

You can also install the stand-alone IBM Support Assistant application on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products that you use. Find add-ons for specific products at <http://www.ibm.com/support/docview.wss?uid=swg27012689>.

Finding product fixes

A product fix to resolve your problem might be available from the IBM software support website.

You can determine what fixes are available by checking the IBM software support website at <http://www.ibm.com/support/entry/portal/>.

- If you previously customized the site based on your product usage:
 1. Click the link for your product, or a component for which you want to find a fix.
 2. Click **Downloads**, and then click **Fixes by version**.
- If you have not customized the site based on your product usage, click **Downloads** and search for your product.

Receiving notification of product fixes

You can receive notifications about fixes, flashes, upgrades, and other news about IBM products.

To sign up to receive notifications about IBM products, follow these steps:

1. From the support page at <http://www.ibm.com/support/entry/portal/>, click **Sign in to create, manage, or view your subscriptions** in the **Notifications** pane.
2. Sign in using your IBM ID and password. If you do not have an ID and password, click **register now** and complete the registration process.
3. Click **Manage all my subscriptions** in the **Notifications** pane.
4. Click the **Subscribe** tab and then click **Tivoli**.
5. Select the products for which you want to receive notifications and click **Continue**.

6. Specify your notification preferences and click **Submit**.

Contacting IBM Software Support

You can contact IBM Software Support if you have an active IBM subscription and support contract and if you are authorized to submit problems to IBM.

To obtain help from IBM Software Support, complete the following steps:

1. Ensure that you have completed the following prerequisites:
 - a. Set up a subscription and support contract.
 - b. Determine the business impact of your problem.
 - c. Describe your problem and gather background information.
2. Follow the instructions in “Submitting the problem to IBM Software Support” on page xx.

Setting up a subscription and support contract

Set up a subscription and support contract. The type of contract that you need depends on the type of product you have.

For IBM distributed software products (including, but not limited to, IBM Tivoli, Lotus®, and Rational® products, as well as IBM DB2® and IBM WebSphere® products that run on Microsoft Windows or on operating systems such as AIX or Linux), enroll in IBM Passport Advantage® in one of the following ways:

- **Online:** Go to the Passport Advantage website at <http://www.ibm.com/software/lotus/passportadvantage/>, click **How to enroll**, and follow the instructions.
- **By telephone:** You can call 1-800-IBMSERV (1-800-426-7378) in the United States. For the telephone number to call in your country, go to the IBM Software Support Handbook web page at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you must understand and assess the business impact of the problem you are reporting.

| | |
|-------------------|--|
| Severity 1 | Critical business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution. |
| Severity 2 | Significant business impact: The program is usable but is severely limited. |
| Severity 3 | Some business impact: The program is usable with less significant features (not critical to operations) unavailable. |
| Severity 4 | Minimal business impact: The problem causes little impact on operations, or a reasonable circumvention to the problem has been implemented. |

Describing the problem and gathering background information

When explaining a problem to IBM, it is helpful to be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?

- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can the problem be re-created? If so, what steps led to the failure?
- Have any changes been made to the system? For example, hardware, operating system, networking software, and so on.
- Are you using a workaround for this problem? If so, be prepared to explain it when you report the problem.

Submitting the problem to IBM Software Support

You can submit the problem to IBM Software Support online or by telephone.

Online

Go to the IBM Software Support website at [http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_\(general\)](http://www.ibm.com/support/entry/portal/Open_service_request/Software/Software_support_(general)). Sign in to access IBM Service Requests and enter your information into the problem submission tool.

By telephone

For the telephone number to call in your country, go to the IBM Software Support Handbook at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html> and click **Contacts**.

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 $\longrightarrow\blacktriangleleft$ 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.



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.



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



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.



Optional choices

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



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



Repeatable choices

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

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



Defaults

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

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



New in version 6.3

Read about new features and other changes to the IBM Tivoli Storage Manager 6.3 Backup-Archive Client.

The Backup-Archive Client can be upgraded (rpm -U) and freshened (rpm -F) on Linux systems

Version 6.3.0 rpm packages can be upgraded using the rpm -U command or freshened using the rpm -F command. You are no longer required to uninstall the files before performing an upgrade. These upgrade options apply only to rpm packages delivered with Tivoli Storage Manager backup-archive clients that are version 6.3.0 and newer.

Journal based backup support added for Linux clients

You can use journal-based backup on Linux backup-archive clients to determine which files are eligible for backup before an incremental backup operation begins. By not querying the server and then comparing the query results with the client file system, journal-based backup can potentially reduce the amount of time that it takes to incrementally back up files on a Linux system. All local file systems are supported by journal-based backup.

Client performance monitor feature

The Tivoli Storage Manager client performance monitor allows you to gather and review performance data about backup, archive, and restore operations for a Tivoli Storage Manager client. The data is collected by the Tivoli Storage Manager API can be used to detect performance bottlenecks for client operations. The client performance monitor is launched from the Tivoli Storage Manager Administration Center. See *Using the Application Programming Interface* for the API options that must be set.

Linux path length has been extended

The maximum path length that is supported by the Backup/Archive client on Linux has been changed to 4096 to match the PATH_MAX supported by the operating system.

Automatic client deployment support for UNIX and Linux clients

The Tivoli Storage Manager server administrator can automatically deploy a client to a UNIX or Linux client workstation. The following UNIX or Linux clients can be deployed from the Tivoli Storage Manager server:

- AIX®
- HP-UX
- Linux
- Mac OS
- Solaris

64-bit clients on Linux, Solaris, and Mac

The Tivoli Storage Manager Backup-Archive Client and Tivoli Storage Manager API for Linux (x86_64), Mac (Intel), and Solaris SPARC are now 64-bit applications. The 32-bit Tivoli Storage Manager API is no longer shipped.

Cluster configuration wizard (Windows)

A client configuration wizard is available to simplify the client

| configuration for processing Microsoft Cluster Server (MSCS) cluster disk
| resources and propagate the configuration to other nodes in the cluster
| group.

| **vStorage API supports off host backups from a vStorage backup server**

| You can now perform full VM backups and full VM incremental backups
| and restores from a Tivoli Storage Manager agent node on a Linux
| computer.

| **The filelist option can now be used with the snapshotroot option**

| In Tivoli Storage Manager version 6.3, you can now use the filelist
| option with the snapshotroot option on incremental, selective, or archive
| operations. Using these two options, together, was not allowed in previous
| versions.

| **Two new options for backing up VMware virtual machines that contain
| independent disks, or RDM disks that are provisioned in physical compatibility
| mode (pRDM)**

| Before this release, attempts to back up a VMware virtual machine failed if
| the virtual machine had one or more independent disks, or one or more
| raw device mapping (RDM) disks that are provisioned in physical
| compatibility mode (pRDM). Independent and pRDM disks are still
| omitted from backups, but two new options (VMPROCESSVMWITHINDEPENDENT
| and VMPROCESSVMWITHPRDM) have been added to specify whether the
| Backup-Archive Client should back up the standard and vRDM disks in
| the virtual machine. The defaults are VMPROCESSVMWITHPRDM NO and
| VMPROCESSVMWITHINDEPENDENT NO. Both options are valid for backups from
| Windows or Linux clients. The options can be specified in the client
| options file (dsm.opt) or in the system options file (dsm.sys), on the
| command line, or in a client options set on the server. Neither option
| applies to backups of Hyper-V systems.

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

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

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

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 6.3 from a previous version.

Upgrade path for clients and servers

As part of a migration plan from Tivoli Storage Manager Version 6.1 or 6.2 to Tivoli Storage Manager Version 6.3, Tivoli Storage Manager clients and servers can be upgraded at different times. This configuration is supported as long as the older version is supported.

To help prevent disruption to your backup and archive activities during the migration, follow these guidelines:

- A Tivoli Storage Manager Version 6.1 or 6.2 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 6.3 server.
- A Tivoli Storage Manager Version 6.3 client can perform backup, restore, archive, retrieve, and query functions to a Tivoli Storage Manager Version 6.1 or 6.2 server.
- If you back up or archive data from a Tivoli Storage Manager Version 6.1 or 6.2 client to any Tivoli Storage Manager server, you can restore or retrieve that data using a Tivoli Storage Manager Version 6.3 client.
- If you back up or archive data from a Tivoli Storage Manager Version 6.3 client, you cannot restore or retrieve that data by using a Tivoli Storage Manager Version 6.2 or earlier client.
- All administrative clients can administer Tivoli Storage Manager Version 6.1, 6.2, and 6.3 servers.
- NAS backups performed on a Tivoli Storage Manager Version 6.3 server using the **BACKUP NODE** server command can only be restored using the **RESTORE NODE** server command or a Tivoli Storage Manager Version 6.3 client.

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

Additional upgrade information

This section explains some additional information that you need to know when upgrading your Tivoli Storage Manager 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.

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

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

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

Upgrading Open File Support or online image

The Open File Support (OFS) and online image installation features have been replaced with a single Logical Volume Snapshot Agent (LVSA) feature.

This feature is selected by default unless you are upgrading from an installation where OFS or online image was being used.

You can install LVSA for use during both online image and OFS operations, but those features are not automatically enabled. To enable OFS, the `snapshotproviderfs` option must be set in the `dsm.opt` file, and to enable online image, the `snapshotproviderimage` option must be set in the `dsm.opt` file.

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

Use the Setup wizard to select **NONE**, **VSS**, or **LVSA** for each of the OFS and online image functions. If **LVSA** is selected and it is not already installed on your system, it is installed at this time.

If you are migrating from a previous version of the Tivoli Storage Manager client where you were using the LVSA for OFS or online image, and you decide during the installation to continue to use the LVSA, then you do not need to explicitly set the `snapshotproviderfs` or `snapshotproviderimage` options. Since you do not need

to set these options, it is easier to install the new client on a large number of systems, because the dsm.opt file will not need to be updated to continue to use the OFS or online image functions.

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.

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.

Start the client deployment process as a post-schedule command of the scheduler. By default, the `autodeploy` option is enabled, and the client workstation is restarted if required. The `autodeploy` option can be set in several places. It can be set by the server in the Administration Center, or it can be set in the options file that is associated with the client scheduler or client acceptor daemon. The user at the client computer can override any other `autodeploy` option setting by setting the `autodeploy` option in the `dsm.opt` file on the client computer. The deployment manager detects options files associated with the scheduler or client acceptor daemon, with the following limitation: If there are multiple scheduler or client acceptor daemon processes running on the same machine, at the same time, and the processes use different options files, the deployment manager uses the `autodeploy` setting in one of the options files.

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 cause is likely due to 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 backup or 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 Administrators Guide* 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 clients version 6.2 and later.

The deployment manager writes log and trace data for a deployment operation to the client's disk. The default location of the logs is in the following directory on the client's disk: C:\Program Files\Tivoli\TSM\IBM_ANR_WIN\Vxxx\log\; where xxx represents the version, release, modification, and fix pack information for the deployed Backup-Archive Client.

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 306

Client environment requirements

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

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

- “NDMP support requirements (Extended Edition only)” on page 8

For current information concerning the client environment prerequisites for all Tivoli Storage Manager supported client platforms, go to the Web site at:
<http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>

Windows client environment requirements

This section contains client environment information, Tivoli Storage Manager client components, and hardware and software requirements for the Windows Vista, Windows Server 2003, and Windows Server 2008 platforms.

For current information concerning the client environment prerequisites for all Tivoli Storage Manager supported client platforms, go to the Web site at:
<http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>

Windows client installable components

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

You can install the following components with Tivoli Storage Manager Version 6.2 or later:

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

System requirements for Windows clients

The Tivoli Storage Manager Windows client requires a minimum amount of disk space for installation and a supported operating system.

Disk space requirements

The following disk space is required to install the Windows client:

- 2.3 GB of free disk space is required to install the 32-bit Windows client, and all language packs.
- 2.9 GB of free disk space is required to install the 64-bit Windows client, and all language packs.

If you select a custom installation, the online installation instructions display how much space is required for the selected installation options, however, that number might not be accurate. The online installation instructions also displays (correctly) how much space is available on the hard disk.

The Java client GUI (graphical user interface) and web GUI require a minimum screen resolution of 600 by 800 pixels.

Software requirements

The Windows client requires a one of the following operating systems:

- Windows Server 2003 (all 32 bit and 64 bit editions)

- Windows Server 2003 R2 (all 32-bit and 64-bit editions)
- Windows Vista (all 32-bit and 64-bit editions)
- Windows 7 (all 32-bit and 64-bit editions)
- Windows Server 2008 (all 32-bit and 64-bit editions)
- Windows Server 2008 R2 (all 64-bit editions)

Restriction: Support for Windows Server 2003 R2 is the same as for Windows Server 2003. Thus the client does not provide support for the new features of Windows Server 2003 R2.

For GPFS™ support (Windows Server 2008 64-bit only), you must also install GPFS version 3.

The supported operating systems are listed in a technote at the following URL: <http://www.ibm.com/support/docview.wss?rs=663&context=SSGSG7&uid=swg21243309>.

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 Tivoli Storage Manager Version 6.2, or later, Windows client:

Table 6. 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) | AIX, HP-UX, Linux, Solaris, Windows, z/OS® |
| Named Pipes | Named Pipes (Standard with all supported Windows platforms) | Windows |
| Shared Memory | TCP/IP (Standard with all supported Windows platforms) | Windows |

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.

The following table shows the supported and unsupported features on the various Windows platforms.

Note:

1. LVSA is not signed, therefore the LVSA driver is included with the Tivoli Storage Manager Client package, but it is not installed by default. If your organization does not permit using unsigned drivers, you can use VSS.
2. ASR from WinPE environment is not supported on 64-bit architectures with Extensible Firmware Interface (EFI).

Table 7. Supported features on Windows platforms

| Features | 2003 32-bit | 2003 64-bit x64 | Vista 32-bit | Vista 64-bit x64 | Windows 7 32-bit | Windows 7 64 bit x64 | 2008 32-bit | 2008 64-bit x64 |
|--|-------------|-----------------|--------------|----------------------|------------------|----------------------|-------------|----------------------|
| Journal-based backup | yes | yes | yes | yes | yes | yes | yes | yes |
| Online image backup | yes | yes | yes | yes | yes | yes | yes | yes |
| Offline image backup | yes | yes | yes | yes | yes | yes | yes | yes |
| System state support with Volume Shadowcopy Services (VSS) | yes | yes | yes | yes | yes | yes | yes | yes |
| LAN-free operations | yes | yes | no | no | no | no | yes | yes |
| Automated System Recovery (ASR) | yes | yes | yes | BIOS: yes EFI: no | yes | BIOS: yes EFI: no | yes | BIOS: yes EFI: no |
| Open File Support (OFS) | yes | yes | yes | yes | yes | yes | yes | yes |

Windows supported file systems

This topic lists the types of file systems that are supported on the Tivoli Storage Manager Windows Backup-Archive Client.

The Tivoli Storage Manager Windows Backup-Archive Client supports the following types of file systems:

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

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

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 Version 6.1.3.x (where x is 1 or higher) or Version 6.2 or later
- Tivoli Storage Manager server Version 6.1.3 or higher
- Tivoli Storage Manager Administration Center Version 6.1.3
 - Required only if you want to use integrated Tivoli Storage Manager FastBack - administration.

The Tivoli Storage Manager FastBack environment must be running. For information about installing and setting up Tivoli Storage Manager FastBack, see *Tivoli Storage Manager FastBack Installation and User's Guide*, at: <http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6>

For information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack, see *IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack*.

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

- Install Tivoli Storage Manager 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 Tivoli Storage Manager client on a workstation where the FastBack DR Hub is installed. In this case, the prerequisites are: the FastBack Disaster Recovery Hub setup, the FastBack shell, and the FastBack mount.
- Install Tivoli Storage Manager 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 47

Client configuration wizard for Tivoli Storage Manager FastBack

The IBM Tivoli Storage Manager 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 Tivoli Storage Manager 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 V5.5.1 or later versions.

Related concepts

“Configuring Tivoli Storage Manager to protect FastBack client data” on page 48

Installing the Tivoli Storage Manager Windows client

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

If the Logical Volume Snapshot Agent (LVSA) component is going to be installed, you must restart the system to install or update the `tsmlvsa.sys` filter driver.

You can install the backup-archive clients on the following Windows platforms:

- Windows Vista
- Windows 7
- Windows Server 2003
- Windows Server 2003 R2
- Windows Server 2008
- Windows Server 2008 R2

Follow these steps to install the software on your Windows system.

1. Insert the DVD that contains the Tivoli Storage Manager Windows client into your DVD drive. If you have autorun enabled, the installation dialog should start when the DVD loads. If the installation dialog does not start automatically, you can start it manually. Select **Run** from the Start menu and at the prompt, type: `x:\setup` where `x` is your DVD drive. Click **OK**.
2. Follow the instructions displayed on the screen. If files from a previous version are in the installation directory, the Windows installer presents these options: **Modify**, **Repair**, and **Remove**. To install a new version of the product, first remove the currently installed version using the **Remove** option. To add a component that was not initially installed, select the **Modify** option.

Installation setup types:

There are two setup types:

- Choosing **Typical** installs the minimum necessary to provide normal backup and archive functions. This includes the Backup-Archive Client, the API Runtime files, and the Web Client.
- Choosing **Custom** takes you to the Custom Setup window. From this window, you can click on any program feature icon to modify that feature if it is not mandatory for installation. You can select from the following program features:
 - Backup-Archive Client
 - Backup-Archive Client GUI Files (Mandatory; cannot be deselected)
 - Backup-Archive Client Web Files (Mandatory; cannot be deselected)
 - Client API Runtime Files (Mandatory; cannot be deselected)
 - Client API SDK Files (Optional; not enabled by default)
 - Administrative Client Command Line Files (Optional; not enabled by default)
 - Logical Volume Snapshot Agent (LVSA) (Optional; not enabled by default)
 - VMware backup tool (Optional; not enabled by default)
 - The Tivoli Storage Manager client now makes use of language packs for non-English language support. Each supported language has its own installation package that must be installed in order to use Tivoli Storage Manager in a supported, non-English language. The Tivoli Storage Manager client is a prerequisite for installing a Tivoli Storage Manager Client Language Pack.

Note:

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

2. The Backup-Archive Client component includes the client scheduler files.
3. The installer displays the exact amount of disk space that is required for each program feature. Ensure that there is enough disk space on the destination drive for the files you choose to install. The installation program will not install to a destination drive with insufficient disk space.
4. If you do not have a dsm.opt file, a setup wizard is launched automatically when you start the GUI. The setup wizard can help you configure an initial options file.

Related tasks

“Creating and modifying the client options file” on page 19

“Starting a web client session” on page 91

“Starting the client scheduler” on page 26

Silent installation

The Tivoli Storage Backup-Archive Client installation program supports silent, unattended installations.

Note: The backup-archive client requires the Microsoft Visual C++ 2010 redistributable package as a setup prerequisite. The graphical installation program installs this package for you. If you are silently installing the client using MSIEXEC, you must separately install the Microsoft Visual C++ 2010 redistributable package. The package can be installed before or after the silent installation of the client is performed, but it must be installed before you use the backup-archive client.

Use the following executable files to install the C++ 2010 redistributable package. In these paths, *dir* represents the drive and directory in which the .exe files exist. These files require elevated privileges to install them. See below for information about how to install software as the administrator.

Windows 32-bit clients

dir\ISSetupPrerequisites\{D4408760-7D16-4F02-B7C1-E68612F5FBFD}\
(contains vcredist_x86.exe)

Windows 64-bit clients

dir\ISSetupPrerequisites\{270b0954-35ca-4324-bbc6-ba5db9072dad}\
(contains vcredist_x86.exe)

dir\ISSetupPrerequisites\{7f66a156-bc3b-479d-9703-65db354235cc} (contains
vcredist_x64.exe)

To perform a silent installation of the C++ redistributable packages, use the following commands. Enter each command on a single line. These files require elevated privileges to install them. See below for information about how to install software as the administrator.

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

```
vcredist_x86.exe /q /c:"msiexec /i vcredist.msi /qn /!v*  
%temp%\vcredist_x86.log"
```

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

```
vcredist_x86.exe /q /c:"msiexec /i vcredist.msi /qn /!v*  
%temp%\vcredist_x86.log"  
  
vcredist_x64.exe /q /c:"msiexec /i vcredist.msi /qn /!v*  
%temp%\vcredist_x64.log"
```

Silent installation on Windows 2008, Windows Vista, and Windows 7 systems

You cannot use the user account control feature in Windows Vista, Windows 7, and Windows 2008 for silent installations. To install the product using the silent method:

- Run the installation from a manageability tool such as Tivoli.
- Run the installation from a command prompt:
 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 using the command prompt window.
- Disable the user account control before installing the product. To disable the user account, complete the following steps:
 1. Type MSCONFIG in the start menu **Start Search** text box.
 2. Click the **Tools** tab. Scroll down and click **Disable UAC**.
 3. Click **Launch**. When the command is done, close the window.
 4. Close the command prompt window and restart your system to apply the changes.

Enable user account control by selecting the **Enable UAC** line, and then click **Launch**.

Silent client installation on Windows 32-bit systems

When you place a customized version of the **msiexec** command (which invokes the Microsoft Software Installer) in a script or batch file, you can perform installations on multiple Windows systems. The following is a sample command to install the backup-archive command-line client, client GUI, web client, API, Administrative command-line client, and LVSA. You might need to customize this example to run correctly on your system. While the command is physically spread across multiple lines on this page, it is entered 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="d:\program files\tivoli\tsm"
ADDLOCAL="BackupArchiveGUI,BackupArchiveWeb,ApiRuntime,
AdministrativeCmd,LVSA"
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. The language pack installation is the same for both

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

Silent language pack installation (32 and 64-bit clients)

You must install client features before you install a language pack. Install a

language pack with a command like the following. Although the command is physically split across multiple lines on this page, it is one single command.

```
msiexec /i
"Z:\tsm_images\TSM_BA_Client\IBM Tivoli Storage Manager Client - French.msi"
RebootYesNo="No" REBOOT="Suppress" ALLUSERS=1
INSTALLDIR="d:\program files\tivoli\tsm"
ADDLOCAL="LanguageFiles" TRANSFORMS=1036.mst /qn /l*v "c:\log.txt"
```

The command parameters are as follows:

msiexec

Invokes the Microsoft Software Installer (MSI) program.

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

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

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

RebootYesNo="No" REBOOT="Suppress"

Under certain conditions, a system reboot might be necessary for the installation to complete successfully. This option causes the installation program to not reboot the system if circumstances would otherwise cause the reboot to occur. While this option is convenient, 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 would be required is if the installation was an upgrade to an existing Tivoli Storage Manager client, and the installation was performed while the client programs were running. Therefore, shut down all Tivoli Storage Manager client programs and services before performing the installation.

ALLUSERS=1

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

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

Specifies the destination path. If you have already installed this product or a previous version of this product on your workstation, 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 backup-archive features are available as shown in the following table:

| Windows 32 bit client features | Windows 64 bit client features | Feature description |
|--------------------------------|--------------------------------|---------------------------|
| BackupArchiveGUI | BackupArchiveGUI | Graphical user interface |
| BackupArchive web | BackupArchive web | 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 |
| LVSA | LVSA | Online image backup and OFS |
| VMwareBackupTools | VMwareBackupTools | VMware backup tools |

ADDLOCAL="LanguageFiles"

After you install client features, you can install a language pack. You must install client features before you install a language pack.

| Windows 32 bit client features | Windows 64 bit client features | Feature description |
|--------------------------------|--------------------------------|-------------------------|
| LanguageFiles | LanguageFiles | Language-specific files |

TRANSFORMS=1033.mst

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

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

/qn Specifies to perform the installation silently.

/!v "c:\log.txt"

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

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

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

- If you want to install a predefined dsm.opt file, place a copy of the dsm.opt file in the ... \CONFIG directory located within the install image, for example:
tsm_images\TSM_BA_Client\baclient\Program Files\Tivoli\TSM\config
- The file must be named dsm.opt.

- The install program 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 install program does not copy over an existing dsm.opt file.
 - dsm.opt exists in the ..\CONFIG directory of the install image, as described earlier.

Upgrading, modifying, reinstalling, or uninstalling the Tivoli Storage Manager Windows client

The Windows Backup-Archive Client can be upgraded, modified, reinstalled, or uninstalled.

If you have previously installed the client and you want to reinstall it into a different directory, uninstall the older client files before you install the new client files.

Wait for any in-process Backup-Archive Client tasks to complete before you uninstall or modify the backup-archive Windows client.

Stop any running services before reinstalling or modifying the client. Use the Windows control panel or service manager commands to display the services that are running. Stop the services shown in Table 8 using either the control panel or a command prompt window.

Table 8. Stoppable services

| Control panel display name | Command line procedure |
|----------------------------|------------------------|
| TSM Journal Service | net stop tsmjbbd.exe |
| TSM Client Acceptor | net stop dsmcad.exe |
| TSM Client Scheduler | net stop dsmcsvc.exe |
| TSM Remote Client Agent | net stop dsagent.exe |

To remove any of these services without uninstalling the client, use the **Help me configure** options in **Utilities > Setup wizard** in the GUI (dsm.exe). You can also use the setup wizard to remove configuration information for online image support and open file support.

The following files are not overwritten if you upgrade or reinstall a client by installing it into the same directory used by the previous installation:

- dsm.opt
- dsmerror.log
- dsmsched.log
- dsmwebcl.log

The configuration settings from the previous version, are in the dsm.opt file, and are preserved for use with the new client software.

During a reinstallation, you are prompted to see if you want the installer to overwrite existing exe, ini, mst, dll, xml, and .properties files. Select **Yes to all** to overwrite these existing files to ensure that you have the current libraries and initialization files.

Follow these steps to uninstall or modify the IBM Tivoli Storage Manager Backup-Archive Client:

1. Open the control panel feature used to add, modify, or remove programs. Click **Start > Control Panel > Programs and Features**.
2. Select the IBM Tivoli Storage Manager client from the list of installed programs.
3. To uninstall an existing client, click **Uninstall**.
4. To fix missing or corrupted files, shortcuts, and registry entries, select the **Repair** option.
5. To remove individual components, select **Change** and then click **Next**. Next to the installable features that you want to uninstall, select **This feature will not be available**.
6. Complete and close the setup wizard.
7. Reboot the computer if you are prompted to do so. Rebooting removes locked files that must be removed before you upgrade or reinstall the client. Rebooting also updates locked files that must be preserved and reused when you reinstall or upgrade the client.

Troubleshooting problems during installation

If you are upgrading from a previous version of Tivoli Storage Manager and there are Tivoli Storage Manager services running (for example, Client Acceptor or Scheduler), you might see an error during 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 running the Tivoli Storage Manager services and retry the installation process.

Installing LVSA using the setup wizard

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

Removing online image backup or open file support during the upgrade

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

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

DiskSpace=8000 ;DiskSpace requirement in KB
to
DiskSpace=0000 ;DiskSpace requirement in KB

Chapter 2. Configure the Tivoli Storage Manager client

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

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

There are some *required* and *optional* configuration tasks. The *required* configuration tasks include the following:

- “Creating and modifying the client options file”
- “Register your workstation with a server” on page 68

Optional configuration tasks include the following:

- “Create a shared directory options file” on page 21
- “Creating multiple options files” on page 21
- “Environment variables” on page 22
- “Configuring the Web client” on page 22
- “Configuring the client scheduler” on page 24
- “Configuring the journal engine service” on page 33
- “Configuring online-image backup support” on page 65
- “Configuring Open File Support” on page 66
- “Creating an include-exclude list” on page 69

Creating and modifying the client options file

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

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

Use the client options file to identify the Tivoli Storage Manager server to contact for services, and communication protocol necessary to establish communications with that server. The client options file can also include authorization options, backup and archive processing options, scheduling options, and other options that govern the behavior of the client. When you run the Tivoli Storage Manager client GUI, if it does not detect an options file, it automatically launches the configuration wizard to help you create one. If you want to manually start the wizard later, perhaps to modify your options file, follow these steps:

1. Launch the GUI by either double-clicking the **TSM Backup Client** icon on your desktop, or going to the Windows task bar and clicking **Start → Programs → Tivoli Storage Manager → Backup Archive GUI**.

2. From the **Utilities** menu, select **Setup Wizard**. This launches the Tivoli Storage Manager Client Configuration Wizard.
3. On the first panel of the Tivoli Storage Manager Client Configuration Wizard, make sure **Help me configure the TSM Backup Archive Client** is selected.
4. Follow the prompts to finish creating your options file. Click the **Help** button if you need assistance completing the wizard.

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

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

```
COMMMethod TCPIP
```

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

```
TCPServeraddress dsmhost.endicott.ibm.com
```

You can also specify what you want to call your workstation by using the `nodename` option. When you register with the backup server during setup, this is the name the server uses for your workstation. If you do not specify a node name, the server identifies your workstation by its host name. You can type **hostname** at a command prompt to see the host name of your workstation.

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

Your backup server can override your options by using server defined and enforced client option sets.

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

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

Related concepts

“Client options reference” on page 298

“Communication options” on page 270

Chapter 12, “Processing options,” on page 269

“Register your workstation with a server” on page 68

Related reference

“Passwordaccess” on page 423

Create a shared directory options file

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

Windows Vista, Windows 7, Windows Server 2003, Windows Server 2008, and Windows Server 2008 R2 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.

For more information see *IBM Tivoli Storage Manager for Windows Installation Guide*, SC32-0138.

Creating multiple options files

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

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

Then do one of the following:

1. From a command prompt, use the **copy** command to replace the `dsm.opt` file with the appropriate options file before you begin a backup or archive session. For example: `copy a.opt dsm.opt`, then `dsm`.
2. Use the `optfile` option on the command line to specify the options file you want to use. For example: `dsm -optfile=b.opt`.
3. Define the `DSM_CONFIG` environment variable to specify the options file to use: `SET DSM_CONFIG=C:\Program Files\Tivoli\TSM\baclient\b.opt`.

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

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

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

Environment variables

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

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

- You want to invoke Tivoli Storage Manager from a directory other than the directory where Tivoli Storage Manager 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 Tivoli Storage Manager client message file dsc*.txt is located:

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

Related reference

“Optfile” on page 421

Configuring the Web client

This task guides you through the steps to configure the Web client.

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

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

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

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

1. If *MANAGEDServices schedule* is set, specify *MANAGEDServices webclient schedule* in the *dsm.opt* file.
2. Ensure that you specify *passwordaccess generate* in the client options file (*dsm.opt*).
3. Install the Client Acceptor Service by entering the following command, where *nodename* and *password* are your Tivoli Storage Manager node name and password.

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

TSM CAD is an example name, you can use any name you want. The default name is TSM Client Acceptor.

4. Install the Remote Client Agent Service by entering the following command, where *nodename* and *password* are your Tivoli Storage Manager node name and password.

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

TSM AGENT is an example name, you can use any name as long as it is different from the CAD name. The default name is TSM Remote Client Agent. The */partnername* option value must match the name of the CAD service. The default is name is TSM Client Acceptor.

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

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

The options applicable only to the **dsmcad** program are `httpport`, `managementservices`, and `webports`. You can also use options such as `optfile` or `errorlogname`. You can use the `managementservices` option to specify whether the Tivoli Storage Manager client acceptor daemon also manages the Tivoli Storage Manager scheduler.

The Administration Center, which is a component available through the Tivoli Integrated Portal, allows you to securely hyperlink to the client workstation. The web client GUI is launched without the administrator needing to sign on again. In order for this to be fully secured, SSL needs to be enabled for the Tivoli Integrated Portal server.

All Web client messages are written to the Web client log file, `dsmwebcl.log`. Error messages are written to the error log file `dsmerror.log`, or the file you specify with the `errorlogname` option. The `dsmwebcl.log` and `dsmerror.log` files reside in the directory you specify with the `DSM_LOG` environment variable or

in the current working directory. Set this environment variable to the directory where the log files should reside. The root directory is *not* a valid value for DSM_LOG. Specify a directory other than the root directory.

6. To access the Web client, enter the following URL from any supported browser, where *your_machine_name* is the host name of the workstation running the Web client.:

```
http://your_machine_name:1581
```

Port 1581 is the default port number. You can set a different port number using the `httpport` option.

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

The options applicable only to the `dsmcad` program are `httpport`, `manageservices`, and `webports`. You can also use options such as `optfile` or `errorlogname`. You can use the `manageservices` option to specify whether the Tivoli Storage Manager client acceptor daemon also manages the Tivoli Storage Manager scheduler.

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

After installing and configuring the Web client on your workstation you can use the Web client to perform backup, archive, restore, and retrieve operations.

Related concepts

“Scheduling options” on page 285

Related tasks

“Starting a web client session” on page 91

Related reference

“Httpport” on page 386

“Passwordaccess” on page 423

Configuring the client scheduler

This task guides you through the steps to configure the client scheduler.

Your Tivoli Storage Manager administrator can schedule Tivoli Storage Manager to perform tasks automatically.

For example, you can automatically back up files at the end of each day or archive some of your files every Friday. This procedure, known as *central scheduling*, is a cooperative effort between the server and your client node. Your administrator associates clients with one or more schedules that are part of the policy domain maintained in the server database. The Tivoli Storage Manager administrator defines central scheduling on the server and you start the client scheduler on your workstation. Once you start the client scheduler, further intervention is not necessary.

With client scheduling, you can perform the following tasks:

- Display information about available schedules.
- Display information about work that the schedule has completed.
- Modify scheduling options in the client options file (dsm.opt).

The Tivoli Storage Manager Client Acceptor service (CAD) can manage the scheduler. In this case, the CAD serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either run immediately or the scheduler exits. The CAD restarts the scheduler when it is time to run the scheduled event. This reduces the number of background processes on your workstation and resolves memory retention problems that can occur when running the scheduler service without CAD management.

You cannot use the `dsmcad` for scheduling when you set the `sessioninitiation` option to *serveronly*.

Use the `managedservices` option in your client options file (dsm.opt) to specify whether the CAD manages the scheduler.

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

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

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

Note:

- If the `sessioninitiation` option is set to *serveronly* in your client options file (dsm.opt), the client setup wizard and scheduler service might be unable to initiate authentication with the Tivoli Storage Manager server. To avoid this problem, ensure that the **Contact the TSM Server to validate password** checkbox on the TSM Authentication page is unchecked.
- In the Service login options window, select the **Automatically when Windows boots** option to have the service started automatically when Windows starts so that your schedules are run.
- You can also use the **Scheduler Service Configuration** utility. The **Scheduler Service Configuration** utility must be run from an account that belongs to the Administrator/Domain Administrator group. You can start multiple client scheduler services on your system.

- If you are having the CAD manage your schedule, start the CAD service, but do not start the scheduler service. The scheduler service is started and stopped automatically by the CAD service as needed. If you are not managing the schedule with the CAD, then you can use the Services Control Panel or the `net start` command to start the Scheduler service.
- If you include files for encryption, ensure that the `encryptkey save` option is set in the options file. This option is set by selecting **Save Encryption Key Password Locally** on the Authorization panel in the preference editor. Setting this option enables unattended scheduled services. If the encryption key has not been previously saved, you must perform an attended backup of at least one file so that you get the encryption prompt and you can save the key.

Related concepts

“Enable or disable scheduled commands” on page 225

“Scheduling options” on page 285

Related tasks

“Configuring the Web client” on page 22

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Related reference

“Managedservices” on page 409

“Sessioninitiation” on page 457

Starting the client scheduler

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

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 12, “Processing options,” on page 269

Scheduling events using the GUI

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

1. From the Tivoli Storage Manager main GUI window, select **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 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.

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.

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 `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.
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 Tivoli Storage Manager scheduler using the setup wizard, and the Tivoli Storage Manager server is behind a firewall, the node password will not get stored on the client 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.
4. When installing Tivoli Storage Manager scheduler using the setup wizard or `dsmcutil`, and the Tivoli Storage Manager server is behind a firewall, the node password will not get stored on the client 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.

5. When configuring the Tivoli Storage Manager 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.

6. The Tivoli Storage Manager client cannot prompt for the encryption key password in scheduler mode. If you are using Tivoli Storage Manager data encryption, you must run an initial interactive backup once to set up the encryption key by opening the TCP/IP connection from the client 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 Tivoli Storage Manager 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, consider the following:

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

Related tasks

“Configuring the client scheduler” on page 24

Related reference

“Sessioninitiation” on page 457

“Tcpadminport” on page 488

“Tcpport” on page 492

“Webports” on page 519

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.

The following client components support SSL:

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

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

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. You can do this using the GSKit command-line utility, gsk8capicmd or gsk8capicmd_64.
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 root certificate of the CA. You can do this using the GSKit command-line utility, gsk8capicmd or gsk8capicmd_64.

Tip: After you have completed this step, if the server gets a new certificate signed by the same CA, the client does not need to import the root certificate again.

3. If you are recovering the Tivoli Storage Manager 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** command is provided by Global Security Kit (GSKit). Tivoli Storage Manager automatically installs GSKit in \Program Files\IBM\gsk8. However, if GSKit has been installed prior to Tivoli Storage Manager installation, it is possible that it is in some other location. You might have to obtain the GSKit location from the following registry key:

```
HKLM\SOFTWARE\IBM\gsk8\CurrentVersion\InstallPath
```

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

1. Open a command window and change the directory to your Tivoli Storage Manager client directory, for example: cd "c:\Program Files\Tivoli\TSM\baclient"
2. Add the GSKit binary path and library path to the PATH environment variable, for example:


```
set PATH=C:\Program Files\IBM\gsk8\bin;  
C:\Program Files\IBM\gsk8\lib;%PATH%
```

If you are configuring SSL on the Tivoli Storage Manager client for the first time, you need to 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 -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 when you specify the `SSLTCPPOINT` or `SSLTCPADMINPORT` server option in the `dsmserv.sys` file, the certificate file is created when you restart the server with these options set. Tivoli Storage Manager version 6.3 servers (and newer version) generate files named `cert256.arm` and `cert.arm`. Tivoli Storage Manager servers older than version 6.3 only generate 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 local-client key database. Then, to enable trust for the certificate, run the following command:

```
gsk8capicmd -cert -add -db dsmcert.kdb -pw <password>
            -label "TSM server <servername> self-signed key"
            -file <path_to_cert256.arm> -format ascii -trust enable
```

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. See “Certificate Authorities root certificates” on page 32 for the preinstalled list of the root certificates of the external Certificate Authorities.

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 local-client key database. Then, to enable trust for the certificate, run the following command:

```
gsk8capicmd -cert -add -db dsmcert.kdb -pw <password>
            -label "XYZ Certificate Authority" -file <path to CA root certificate>
            -format ascii -trust enable
```

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 Tivoli Storage Manager 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 Tivoli Storage Manager client directory, you must copy `dsmcert.kdb` and `dsmcert.sth` into that directory.

4. On a 64-bit platform, GSKit is installed in \Program files\IBM\gsk8_64. The GSKit location can vary, so check the registry key. For a 64-bit platform, use the gsk8capicmd_64 command.
5. Due to system performance concerns, use SSL only for sessions where it is needed. Consider adding additional processor resources on the Tivoli Storage Manager server system to manage the increased requirements.
6. In order for a client to connect to a server using Transport Layer Security (TLS) version 1.2, the certificate's signature algorithm needs to be SHA-1 or stronger. If you are using a self-signed certificate, you must use the cert256.arm certificate. Your TSM 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 has been 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. This is normal behavior that you can correct by specifying the correct port number.

Related reference

“`Ssl`” on page 480

“`Sslfipsmode`” on page 481

Certificate Authorities root certificates

When the Tivoli Storage Manager Backup-Archive client is delivered, it includes a list of root certificates for a number of common Certificate Authorities.

Here is the 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.

Related tasks

“Configuring the journal engine service”

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.

Journal-based backup is enabled by installing and configuring the Tivoli Storage Manager journal service. The Tivoli Storage Manager journal service can be installed with the GUI Setup wizard or with the **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`.

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

1. From the main window, open the **Utilities** menu and select **Setup Wizard**.
2. Select the **Help me configure the TSM Journal Engine** check box.
3. Select the task you want to perform. You can install a new 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.

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.
```
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 mentioned above.

Note:

1. Each journal service instance is associated to only one Tivoli Storage Manager client node name. Changing the association requires a restart of the journal service to recognize the new association.
2. 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 114

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\jnlServer`. For example, in `dsm.opt`:

```
JournalPipe \\.\pipe\jnlServer1
```

Under `tsmjbbd.ini` [JournalSettings] stanza:

```
JournalPipe=\\.\pipe\jnlServer1
```

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

NlsRepos

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

```
NlsRepos=dscenu.txt
```

ErrorLog

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

```
ErrorLog=jbberror.log
```

JournalDir

Specifies the directory where journal database files are stored and written. The default directory is the journal service installation directory. You can specify different journal locations for each file system being journaled. This is useful when running in a clustered environment because the location of the journal must be accessible by each 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:\tsmjjournal`), 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.

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:\excludedir\*
c:\dir1\excludefile
*.*\*.tmp
```

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

JournaledFileSystemSettings stanza

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

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

Syntax for *JournaledFileSystemSettings* stanza:
`[JournaledFileSystemSettings]`

Syntax for stanza settings:
`JournaledFileSystemSetting=value`

You can specify the following *JournaledFileSystemSettings* values:

DirNotifyBufferSize

Specifies the size of the buffer to record change notifications for a particular journal file system. You might need to increase this value for 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.

JournaledFileSystems

Specifies a space delimited list of file systems to journal. Full file system specifications and Windows junctions are supported. There is no default value. You must specify at least one journaled file system for the journal service to run. Journaled file systems can be added or removed online without having to restart the service. For example:

```
JournaledFileSystems=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

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

[JournalFileSystemSettings.f:\]
```

```

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

```

Related concepts

“Overriding stanzas”

Overriding stanzas

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

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

Syntax for JournaledFileSystemSettings stanza:

```
[JournaledFileSystemSettings.fs]
```

Syntax for stanza settings:

```
JournaledFileSystemSetting=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 on Tivoli Storage Manager: *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.

- 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 being 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 non-deduplicated storage pool, client extents are reassembled into contiguous files.

Server-side data deduplication offers more protection against data loss. 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 Tivoli Storage Manager 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

The following prerequisites apply to client-side data deduplication:

- 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 IBM Tivoli Storage Manager 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.
- Files on the client must not be excluded from client-side data deduplication. By default, all files are included. You can optionally exclude specific files from data deduplication.
- Files on the client must not be encrypted. Encrypted files and files from encrypted file systems cannot be deduplicated.
- Files must be more than 2 KB. Files that are 2 KB or less are not deduplicated.

The server can limit the maximum transaction size for data deduplication by setting the CLIENTDEDUPTXNLIMIT option on the server. See the *Administrator's Guide* for details.

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 9.

Table 9. Data deduplication settings: Client and server

| Value of the client DEDUPLICATION option | Setting on the server | Data deduplication location |
|--|------------------------------------|-----------------------------|
| Yes | On either the server or the client | Client |
| Yes | On the server only | Server |
| No | On either the server or the client | Server |
| No | On the server only | Server |

Encrypted files

The Tivoli Storage Manager server and the backup-archive client cannot deduplicate encrypted files. If an encrypted file is encountered during data deduplication processing, the file is not deduplicated, and a message is logged.

Tip: You do not have to process encrypted files separately from files that are eligible for client-side data deduplication. Both types of files can be processed in the same operation. However, they are sent to the server in different transactions.

As a security precaution, you can take one or more of the following steps:

- Enable storage-device encryption together with client-side data deduplication.
- Use client-side data deduplication only for nodes that are secure.
- If you are uncertain about network security, enable Secure Sockets Layer (SSL).
- If you do not want certain objects (for example, image objects) to be processed by client-side data deduplication, you can exclude them on the client. If an object is excluded from client-side data deduplication and it is sent to a storage pool that is set up for data deduplication, the object is deduplicated on server.
- Use the **SET DEDUPVERIFICATIONLEVEL** command to detect possible security attacks on the server during client-side data deduplication. Using this command, you can specify a percentage of client extents for the server to verify. If the server detects a possible security attack, a message is displayed.

Related tasks

“Configuring the client for data deduplication”

Related reference

“Deduplication” on page 332

“Exclude options” on page 364

“Dedupcachepath” on page 331

“Dedupcachesize” on page 332

“Enablededupcache” on page 354

“Ieobjtype” on page 388

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 configure your client to use data deduplication, ensure that the requirements listed in “Client-side data deduplication” on page 40 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. See *Tivoli Storage Manager Administrator's Guide* for details.

Follow these steps to enable data deduplication on the client:

1. Add the deduplication yes option to the `dsm.opt` file.
2. You can also set this option in the GUI:
 - a. Select **Edit > Client Preferences**
 - b. Click the **Deduplication** tab
 - c. Select the **Enable Deduplication** check box
 - d. Click **OK** to save your selections and close the Preferences Editor

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 using these steps:

1. 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 adding the ENABLEDEDUPCACHE NO option to the dsm.opt file.

You can also set this option in the Tivoli Storage Manager backup-archive client preference editor by clearing the **Enable Deduplication** check box.

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 40

Related reference

“Deduplication” on page 332

“Enablededupcache” on page 354

“Exclude options” on page 364

Excluding files from data deduplication

You can exclude a file from data deduplication during backup or archive processing.

You can exclude only files for archive data deduplication. You can exclude files, images, system state objects, and ASR for backup data deduplication.

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.

You can also exclude files from data deduplication processing in the `dsm.opt` file:

1. Add the deduplication yes option
2. Exclude client-side data deduplication for image backup of drive E: by adding `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 40

Related reference

“Deduplication” on page 332

“Enablededupcache” on page 354

“Exclude options” on page 364

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, there are some configuration tasks that you must complete.

First ensure that you have configured the backup-archive client and that you installed the Tivoli Storage Manager FastBack client.

Install the Tivoli Storage Manager FastBack client using the information in *Tivoli Storage Manager FastBack Installation and User's Guide*, here: <http://publib.boulder.ibm.com/infocenter/tsmfbinf/v6>.

After you install the Tivoli Storage Manager FastBack client, complete the following tasks. You can also use the Client Configuration wizard for Tivoli Storage Manager FastBack.

1. Register a Tivoli Storage Manager node for each Tivoli Storage Manager FastBack client where Tivoli Storage Manager backs up or archives data. The node name must be the short host name of the FastBack client.

This is a one-time configuration performed once for each Tivoli Storage Manager 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.

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 Tivoli Storage Manager 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 Tivoli Storage Manager client connects. Run the `set password -type=fastback` command once for each repository where the Tivoli Storage Manager client is expected to connect.

The credentials that are stored depends on these configurations:

- Tivoli Storage Manager client on a FastBack server
- Tivoli Storage Manager client on a FastBack DR Hub
- Tivoli Storage Manager client on a dedicated proxy workstation

See *IBM Tivoli Storage Manager Integration Guide for Tivoli Storage Manager FastBack* for information about integrating Tivoli Storage Manager and Tivoli Storage Manager FastBack.

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 Tivoli Storage Manager to protect FastBack client data”

Related reference

“Set Password” on page 648

Configuring Tivoli Storage Manager 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 Tivoli Storage Manager backup-archive client is properly configured with a Tivoli Storage Manager server. Also make sure that the Tivoli Storage Manager client acceptor service (`dsmcad.exe`) is running. You can use the Tivoli Storage Manager Client Configuration wizard in the Tivoli Storage Manager Java 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 Tivoli Storage Manager 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 Tivoli Storage Manager backup-archive client and FastBack server or DR 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 graphical user interface:

1. Ensure that the Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager Client Acceptor service is running.
For example: http://<machine_name_or_ip_address>:1585
3. From the Tivoli Storage Manager 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.

Tivoli Storage Manager in a VCS environment is supported on Windows 2003 and 2008.

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

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

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

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

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

| 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 421

Protecting data in MSCS and VCS clusters (Windows Server 2003)

| Installing and configuring the backup-archive client on Windows Server 2003
| systems in an MSCS or VCS environment requires you to install the
| backup-archive client on the cluster nodes and configure the client to back up the
| drives in a cluster disk group.

Install the backup-archive client on the cluster nodes

Install the backup-archive client software on a local disk on each cluster node. The executable files should reside in the same location on each local drive.

The default location for the backup-archive client is in C:\Program Files\tivoli\tsm\baclient.

Configuring the backup-archive client to process local nodes

To back up the local (non-clustered) drives and process system state information, in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment, the Tivoli Storage Manager Client Scheduler Service should use the following combination of the options: CLUSTERNODE NO (default) and CLUSTERDISKONLY YES (default).

The `clusternode` and `clusterdiskonly` options should be invoked before MSCS or VCS is started, because although the Tivoli Storage Manager Client Scheduler Service is configured to back up local drives, the scheduler depends on the cluster service. The Tivoli Storage Manager Scheduler communicates with the cluster service (MSCS or VCS) to build the cluster disks map, get the cluster name, etc.

To ensure that the MSCS or VCS loads before the Tivoli Storage Manager Client Scheduler service, a dependency should be added (for MSCS service or VCS service) to the Tivoli Storage Manager Client Scheduler service.

Follow these steps to add the service dependency:

1. Start **regedt32.exe**
2. Locate the following subkey in the registry: HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services*TSM Scheduler Service*, (where *TSM Scheduler Service* is the name of the Tivoli Storage Manager Scheduler)
3. Double-click the **DependOnService** value
4. In the **Value data** box, type `ClusSvc` for the MSCS environment or `HAD` for the VCS environment, and then click **OK**
5. Click **File > Exit** to close the registry editor
6. Restart your computer

Note: After disabling or uninstalling the MSCS or VCS, the dependency should be removed.

You can edit your `dsm.opt` file on each local node to process local disk drives using the following options:

nodename

If no value is specified, the backup-archive client uses the local workstation name.

domain If you do not specify a value for this option, the backup-archive client processes all local drives that are not owned by the cluster.

clusternode

Do not specify this option when processing local drives.

You can configure the Tivoli Storage Manager Backup-Archive Scheduler Service to back up the local cluster nodes.

Related reference

“Clusternode” on page 317

“Domain” on page 343

“Nodename” on page 417

Configure the backup-archive client to process cluster disk resources

Ensure that the backup-archive client manages each cluster group that contains physical disk resources as a unique node.

This ensures that the backup-archive client correctly manages all disk resources, regardless of which cluster node owns the resource at the time of backup.

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 **mcs-cluster** contains the following groups and resources:

- **group-a** - Contains physical disk **q:** (quorum), and physical disk **r:**
Note: VCS does not have quorum disk.
- **group-b** - Contains physical disk **s:**, and physical disk **t:**

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

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

Step 2: Configure the client options file:

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

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

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

nodename

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

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

See “Frequently asked questions” on page 63 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, Tivoli Storage Manager:

1. Checks for a cluster environment (MSCS or VCS).

2. Uses the cluster name instead node name for file space naming and encryption. This allows the use of one password file for all nodes in the cluster.
3. Builds a list of shared volumes and works only with shared volumes. Back up of local volumes not permitted if the `clusternode` option set to `yes`.

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 called `main.cf`, which is located in the path pointed by `%VCS_HOME%conf/config` on each node in the cluster. If this file is corrupted, the cluster configuration is also corrupted. Be very careful when handling this file. The `VCS_HOME` environment variable points to the directory where VCS is installed on the node.

passwordaccess

Specify the generate value.

managedservices

(Optional). Specifies whether the Tivoli Storage Manager Client Acceptor service (CAD) manages the scheduler, the Web client, or both. The examples in this appendix assume that the CAD manages both the Web client and the scheduler for each cluster group. To specify that the CAD manages both the Web client and the scheduler, enter the following option in the `dsm.opt` file for each cluster group:

```
managedservices webclient schedule
```

errorlogname

Specify a unique error log name.

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

schedlogname

Specify a unique schedule log name.

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

Related reference

“Clusternode” on page 317

“Domain” on page 343

“Errorlogname” on page 361

“Managedservices” on page 409

“Nodename” on page 417

“Passwordaccess” on page 423

“Schedlogname” on page 450

Step 3: Configure the scheduler service:

Configure a Tivoli Storage Manager backup-archive scheduler service for each cluster group by using the `dsmcutil` command.

Configure the scheduler service on each node in the cluster group. The scheduler service must have the same name on each node and be available for failover (moved to the other nodes in the cluster).

To install the Tivoli Storage Manager Scheduler Service for **group-a** from workstation **node-1**, ensure that **node-1** currently owns **group-a** and issue the following command:

```
dsmcutil install SCHEDuler /name:"tsm scheduler service: group-a"  
/clientdir:"c:\Program Files\tivoli\tsm\baclient" /optfile:q:\tsm\  
dsm.opt /node:mscs-cluster-group-a /password:nodepassword  
/validate:yes /autostart:no /startnow:no /clusternode:yes  
/clustername:mscs-cluster
```

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

Note: For VCS, the default value of the /autostart option is no.

Using the Microsoft Cluster Administrator utility or VCS Configuration Editor, move **group-a** to **node-2**. From **node-2**, issue the same **dsmcutil** command (above) to install the service on **node-2**. Repeat this procedure for each cluster group.

Related tasks

“Frequently asked questions” on page 63

Step 4: Create a generic service resource for failover:

Creating a generic service resource allows the Tivoli Storage Manager services for a cluster group to correctly failover between the two nodes of the cluster.

MSCS:

To add a Generic Service resource to each cluster group managed by the backup-archive client, use the Microsoft Cluster Administrator utility.

1. Select the **group-a** cluster group under the MSCS-Cluster\Groups folder and select **File > New > Resource** from the dropdown menu.
2. In the New Resource dialog, enter a unique name in the **Name** field. For example: TSM SCHEDULE SERVICE for GROUP-A. Enter a description in the **Description** field. For example: TSM cluster schedule for cluster group GROUP-A. Select the **Generic Service** resource type in the **Resource Type** field. Enter the group name in the **Group** field. Press **Enter**.
3. In the Possible Owner dialog, ensure that all cluster nodes appear as possible owners. Press **Enter** to continue.
4. In the Dependencies dialog add all physical disk resources as Resource Dependencies. Press **Enter** to continue.
5. In the Generic Service Parameters dialog, enter the client acceptor service name that you specify in step 3 of “Configure the Web client in a cluster environment” on page 57, in the **Service Name** field. Leave the **Startup Parameters** field blank. Press **Enter**.
6. In the Registry Replication dialog, add the registry key corresponding to the Tivoli Storage Manager node name and server name. The format for this key is: `HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\nodename\TSM server instance name`, where *nodename* is the name of your Tivoli Storage Manager node, and *server instance name* is the name of the Tivoli Storage Manager server that the node connects to. For example, if the node name is **mscs-cluster-group-a** and the Tivoli Storage Manager server name is **tsmsv1**, then enter the following registry key in the Registry Replication dialog:

HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\mscs-cluster-group-a\tsmv1. This entry matches an existing key in the Windows registry.

7. Select the new resource from the Microsoft Cluster Administrator utility, and click **File** and then **Bring Online** from the dropdown menu.

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

The new scheduler service is now associated with the cluster group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

Note: If you manually change the Tivoli Storage Manager password, you must stop the scheduler service, regenerate the password, and restart the scheduler service. The password can be generated by running the command:

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```

Important: The generic service resource should remain online during any password changes. Password corruption occurs if the generic service resource is offline while the password is changed.

VCS:

VCS provides a separate utility to create or modify your existing configuration.

This utility is called the VCS Configuration Editor. By default, this utility is installed when VCS Cluster Manager is installed. You can also manually edit the VCS configuration file `main.cf` file; refer to the appropriate VCS documentation for more information.

To add a Generic Service resource to each cluster group managed by the backup-archive client, use the VCS Configuration Editor, as follows:

1. Open the VCS Configuration Editor. You are prompted with the Build a new configuration or modify existing configuration window which provides the following options: **New Config** - If you select this option you are prompted for the path for the `types.cf` file, and **Open Existing Config** - If you select this option, the configuration window opens. Click on the 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. In the Registry Replication dialog box, 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 **tsmv1**, enter the following registry key in the Registry Replication

dialog box: HKEY_LOCAL_MACHINE\SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\mscs-cluster-group-a\tsmv1. This entry matches an existing key in the Windows registry.

7. Select the **ServiceName** attribute and enter the name of scheduler or CAD service that you want to make High-Availability.
8. When you are finished, close the window. The Configuration window prompts you whether to save the configuration; click **Yes**.

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

1. Open the VCS Configuration Editor. You are prompted with the Build a new configuration or modify existing configuration window which provides the following options: **New Config** - If you select this option you are prompted for the path for the types.cf file, and **Open Existing Config** - If you select this option, the configuration window opens. Click the resource group that 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**.

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

The new scheduler service is now associated with the cluster group. If the group is moved (failed) to the other nodes in the cluster, the service should correctly fail over between the cluster nodes and notify both cluster nodes of automatic password changes.

Note: If you manually change the Tivoli Storage Manager password, you must stop the scheduler service, regenerate the password, and restart the scheduler service. The password can be generated by running the command:

```
dsmc query session -optfile="q:\tsm\dsm.opt"
```

Protecting data in MSCS clusters (Windows Server 2008)

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 Windows Server 2008 or Windows Server 2008 R2 as their operating system.

Configuring cluster protection (Windows 2008)

Use the Tivoli Storage Manager cluster wizard to configure the Backup-Archive Client to protect cluster resources. The wizard gathers the information needed so the Backup-Archive Client can protect cluster resources, and log on to the server.

Before you run the Tivoli Storage Manager 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 or the web-accessible administration center command-line interface and register the node 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.

You only need to run the wizard on 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.

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. Note that you cannot use the wizard to backup 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 “Install the backup-archive client on the cluster nodes” on page 50 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:

First you must identify the cluster groups to be managed.

Perform the steps under step 1 of “Configure the backup-archive client to process cluster disk resources” on page 52.

Step 2: Configure the client options file:

Now you must configure the client options file.

Perform the steps under step 2 of “Configure the backup-archive client to process cluster disk resources” on page 52.

In addition, specify the following options in the `dsm.opt` file for each cluster group:

httpport

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

Note: It is preferable to specify a different log name using the `schedlogname` option in the `dsm.opt` file for each cluster group.

Related reference

“Httpport” on page 386

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 the above 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 63 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

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

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

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

MSCS:

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

Follow these steps to add an IP address resource:

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**.

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.

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

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.

MSCS:

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

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.

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 63 for information on what to do if a generic service resource for the cluster group fails because the CAD service has been removed.

VCS:

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

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**.

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" on page 63 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.

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

Alternatively, you can point your browser at the network name and httpport. For example, if you used a network name of **cluster1groupa** and specified an 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.

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

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

1. Right-click on the desktop and select **New > Shortcut**.
2. In the window that appears, find the path to the `dsm.exe` executable (located by default in directory `C:\program files\tivoli\tsm\baclient\`). 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"`.

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 Tivoli Storage Manager cluster scheduler service resource for backup?

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

1. The cluster drive resource, and any related resource shares, must exist and reside within the designated cluster group as defined in the Microsoft Cluster Administrator utility or VCS Configuration Editor. The designated cluster group must already contain the cluster scheduler service resource for which this new drive 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 CAD service has been removed and now the generic service resource for the cluster group is failing. How can this be corrected?

A: The CAD can be used to control the scheduler, the Web client, or both for Tivoli Storage Manager in a cluster environment. If the CAD is removed without updating the generic cluster resource, the resource fails. To correct this:

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

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

Configuring online-image backup support

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

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

To configure online image backup, perform the following steps:

1. Select **Utilities** → **Setup Wizard** from the Tivoli Storage Manager GUI main window. The Client Configuration Wizard panel appears.
2. Select **Help me configure Online Image Support** and click the **Next** button. The TSM Online Image Support Wizard panel appears.
3. Select the Snapshot Provider you want to use and then click the **Next** button. The providers you can select from are **VSS** or **LVSA**. To disable online image support, select **None**.
 - a. If you select **None** and the LVSA is installed, and the LVSA is not being used for OFS, the **Uninstall the Logical Volume Snapshot Agent (LVSA)** checkbox is enabled. If selected, the LVSA storage filter driver is uninstalled.
4. Click the **Finish** button to complete the setup.
5. Complete each panel in the wizard and click the **Next** button to continue. To return to a previous panel, click the **Back** button. To display help information for a panel, click the **Help** button.
6. If you selected as the provider in 3, restart the system if prompted.

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

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

Related concepts

“Client options reference” on page 298

“Image backup” on page 129

Configuring Open File Support

The Open File Support (OFS) feature can be selected at installation time or installed later using the Tivoli Storage Manager client GUI setup wizard.

By default, the OFS feature is not installed.

The install program or the setup wizard performs all the steps necessary to install, set up, and enable the OFS feature so that the next backup or archive operation attempts to take advantage of the open file support. You can also use the install program or setup wizard to install or remove the OFS feature. If the LVSA is chosen as the snapshot provider, the installation or removal of this feature might require a workstation reboot.

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

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

To configure OFS, perform the following steps:

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

The following are some things to consider when you configure Open File Support:

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

The LVSA driver is a non-signed driver. It is included with the Tivoli Storage Manager client package but is not installed by default. If your organization has a strict policy which does not allow non-signed drivers to be installed, then you can use VSS as your online image provider. Offline image backups can be performed without installing the LVSA driver.

Related concepts

“Client options reference” on page 298

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 Tivoli Storage Manager client.

1. Define a new user on your NetApp file server 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 difference incremental, until their password is changed. Users in groups that do not have these capabilities cannot log in to the storage system. Refer to NetApp documentation for details on defining a user id and a password on the NetApp file server.

2. Export the NetApp volumes and consider the following recommended settings:

Tip: See the NetApp documentation for details on exporting the NetApp volumes for use with Windows.

- Map the NetApp volumes using CIFS.
 - Ensure the NetApp volumes have the NTFS security setting.
3. Set the user ID, and password on Tivoli Storage Manager client for the user ID that you created in step 1 on page 67 using the following steps:
 - a. Log on as the user with read/write access to the CIFS share.
 - b. From the Tivoli Storage Manager 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 hostname of your NetApp file server.

snapdiff_user

This value is the user ID that you created in step 1 on page 67.

newPassword

This value is the password for the user ID that you created in step 1 on page 67.

Related reference

"Snapdiff" on page 463

"Createnewbase" on page 326

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.

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 server.

The first time you start a session, Tivoli Storage Manager prompts you for information necessary to register your workstation with the server identified in your client options file. You need to supply your node name, a password, and contact information.

When you use open registration:

- Your client node is assigned to a policy domain named **standard**.
- You can delete archived copies of files from server storage, but not backup versions of files.

If necessary, your Tivoli Storage Manager administrator can change these defaults later.

Creating an include-exclude list

If you do not create an include-exclude list, Tivoli Storage Manager considers all files for backup services and uses the default management class for backup and archive services.

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

Specify 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.

You can use the following methods to create an include-exclude list or specify an include-exclude file:

1. You can add include-exclude statements in the Tivoli Storage Manager or Web client directory tree. The online help provides detailed instructions.
2. Open the **Edit** menu and select **Client Preferences**. In the Preferences dialog, click the **Include/Exclude** section. You can specify an INCLEXCL file using the Preferences editor. However, you cannot create the INCLEXCL file using the Preferences editor.
3. 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. Tivoli Storage Manager evaluates all `exclude.dir` statements *first* (regardless of their position within the include-exclude list), and removes the excluded directories and files from the list of objects available for processing. All other include-exclude statements are processed from the bottom of the list up. Therefore, it is important to enter all your include-exclude statements in the proper order. For example, in the following include-exclude list the `includefile.txt` file *is not* backed up:

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

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

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

5. Save the file and close it. For Mac OS X, ensure that you save the file as plain text encoded as Unicode (UTF-8 or UTF-16). Do not add the `.txt` extension.
6. Restart your Tivoli Storage Manager client and the scheduler and client acceptor services to enable your include-exclude list.

Related concepts

“System files to exclude” on page 74

Chapter 9, “Storage management policies,” on page 231

Related reference

“Incl`excl`” on page 392

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.

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

Table 10 on page 72 lists the options you can use to exclude file spaces and directories from processing.

Table 10. Options for excluding file spaces and directories

| Option | Description |
|---|---|
| exclude.dir "Exclude options" on page 364 | <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: <ul style="list-style-type: none"> Exclude directory <code>C:\MyPrograms\Traverse</code> and its files and subdirectories: <pre>exclude.dir c:\MyPrograms\Traverse</pre> 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. <pre>exclude.dir c:\MyPrograms\Traverse*</pre> Exclude all directories whose names begin with <code>temp</code>, and are located within directory <code>x:\documents and settings</code> and its subdirectories, where <code>x:</code> is any drive. <pre>exclude.dir "x:\documents and settings\...\temp"</pre> Exclude all directories whose names begin with <code>temp</code>, regardless of the drive or directory in which they reside: <pre>exclude.dir temp*</pre> The following example is invalid because it ends with a directory delimiter: <pre>exclude.dir c:\MyPrograms\Traverse\</pre> <ul style="list-style-type: none"> 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. <pre>exclude x:* exclude.dir x:*</pre> 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.sanjose.ibm.com.

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

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

Here is an example of the new format using UNC:

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

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

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

Table 11 lists the options that you can use to control processing with include and exclude statements.

Table 11. 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 364 |
| 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 393 |
| include.backup | | |
| include.file | | |
| include.fs | Use this option to set options on a file space-by-file space basis. | "Include options" on page 393 |
| Archive processing | | |

Table 11. 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 364 |
| 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 393 |
| 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 <code>exclude.dir</code> statements. This option is for all Windows clients. | “Exclude options” on page 364 |
| 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 <code>exclude.image</code> . This option is valid for all Windows clients. | “Exclude options” on page 364 |
| include.fs.nas | Use the <code>include.fs.nas</code> option to bind a management class to Network Attached Storage (NAS) file systems. To specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, use the <code>toc</code> option with the <code>include.fs.nas</code> option in your client options file (<code>dsm.opt</code>). See “Toc” on page 496 for more information. This option is valid for all Windows clients. | “Include options” on page 393 |
| 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 platforms. | “Include options” on page 393 |
| System state processing | | |
| exclude.systemservice | Excludes individual system components from system state backup. Input can be the keyword or component name to be excluded from backup. This option is valid for Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 <i>only</i> . | “Exclude options” on page 364 |
| include.systemstate | Assigns management classes for backup of the Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 system state. The default is to bind the system state object to the default management class. | “Include options” on page 393 |

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".

Note:

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.

Tivoli Storage Manager 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 528

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 12. Wildcard and other special characters

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

Table 12. Wildcard and other special characters (continued)

| Character | Function |
|-----------|---|
| * | The match-all character. For example: <ul style="list-style-type: none"> The pattern <code>ab*</code>, matches <code>ab</code>, <code>abb</code>, <code>abxxx</code>, but does not match <code>a</code>, <code>b</code>, <code>aa</code>, <code>bb</code>. The pattern <code>ab*rs</code>, matches <code>abrs</code>, <code>abtrs</code>, <code>abrsrs</code>, but does not match <code>ars</code>, or <code>aabrs</code>, <code>abrss</code>. The pattern <code>ab*ef*rs</code>, matches <code>abefrs</code>, <code>abefghrs</code>, but does not match <code>abefr</code>, <code>abers</code>. The pattern <code>abcd.*</code>, matches <code>abcd.c</code>, <code>abcd.txt</code>, but does not match <code>abcd</code>, <code>abdc</code>, or <code>abcdtxt</code>. |
| \... | The match- <i>n</i> character matches zero or more directories. |
| [| The open character-class character begins the enumeration of a character class. For example: <code>xxx[abc]</code> matches <code>xxxa</code> , <code>xxxb</code> , or <code>xxxc</code> . |
| - | The character-class range includes characters from the first character to the last character specified. For example: <code>xxx[a-z]</code> matches <code>xxxa</code> , <code>xxxb</code> , <code>xxxc</code> , ... <code>xxxz</code> . This format should not be used to specify remote drives in an <i>exclude</i> statement. |
| \ | The literal escape character. When used within a character class, it treats the next character literally. When used outside of a character class, it is not treated in this way. For example, if you want to include the <code>]</code> in a character class, enter <code>[...\.]</code> . The escape character removes the usual meaning of <code>]</code> as the close character-class character. |
|] | The close character-class character ends the enumeration of a character class. |
| : | 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: <code>d:\direct\file.nam</code> |

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 13 shows how to correctly specify shared drives.

Table 13. Specifying a drive specification using wildcards

| Incorrect | Correct |
|--|--|
| <code>\\remote*\...*.*</code> | <code>\\remote*\$\...*.*</code> |
| <code>\\remote\?\...*.*</code> | <code>\\remote\?\$\...*.*</code> |
| <code>\\remote*\...\pagefile.sys</code> | <code>\\remote*\$\...\pagefile.sys</code> |

Related concepts

“Wildcard characters” on page 528

Examples using wildcards with include and exclude patterns

Tivoli Storage Manager 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 14 shows how to use wildcard characters to include or exclude files.

Table 14. 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. | <code>exclude ?:*.*bak</code> <code>include d:\dev*.*bak</code> |
| Exclude all files in any directory named "tmp" and its subdirectories, <i>except</i> for the file <code>d:\tmp\save.fil</code> . | <code>exclude ?:\...\tmp\...*</code> <code>include d:\tmp\save.fil</code> |
| Exclude any <code>.obj</code> file for backup in any directory on the c: e: f: and g: drives. | <code>exclude [ce-g]:\...*.obj</code> The c: e: f: and g: drives are local or removable. |
| Exclude the <code>.obj</code> files found in the root directory in the d: drive <i>only</i> . | <code>exclude d:*.obj</code> |
| Exclude any file that resides under the tmp directory found on any drive. | <code>exclude ?:\tmp\...*</code> |
| Exclude the <code>c:\mydir\test1</code> directory and any files and subdirectories under it. | <code>exclude.dir c:\mydir\test1</code> |
| Exclude all directories under the <code>\mydir</code> directory with names beginning with <code>test</code> . | <code>exclude.dir c:\mydir\test*</code> |
| Exclude all directories directly under the <code>\mydir</code> directory with names beginning with <code>test</code> , on any drive. | <code>exclude.dir ?:\mydir\test*</code> |
| Exclude the raw logical volume from image backup. | <code>exclude.image c:*</code> |
| Exclude all directories and files on the local drives, except the c: drive. | <code>exclude [abd-z]:\...*</code> <code>exclude.dir [abd-z]:\...*</code> |

Related concepts

“Examples using wildcards with include and exclude patterns”

Related reference

“Exclude options” on page 364

Determine compression and encryption processing

Tivoli Storage Manager evaluates `exclude.dir` and any other include-exclude options controlling backup and archive processing, and then the decision is made about which files undergo compression and encryption processing.

The following options determine which files undergo compression and encryption processing.

Table 15. Options for controlling compression and encryption processing

| Option | Description | Page |
|---|---|-------------------------------|
| Compression processing | | |
| exclude.compression | Excludes files from compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives. | “Exclude options” on page 364 |
| include.compression | Includes files for compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives. | “Include options” on page 393 |
| Encryption processing | | |
| exclude.encrypt | Excludes files from encryption processing. | “Exclude options” on page 364 |
| include.encrypt | Includes files for encryption processing. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. Important: The <code>include.encrypt</code> option is the only way to enable encryption on the Backup-Archive client. If no <code>include.encrypt</code> statements are used encryption will not occur. | “Include options” on page 393 |
| 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 364 |
| 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 393 |

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 Tivoli Storage Manager client Java GUI directory tree shows detailed information of included and excluded objects. The directory tree windows in the Tivoli Storage Manager Client Java GUI allow you to select files and directories to include or exclude. You should use this **preview** command to make sure that you include and exclude the correct files. The following is a sample scenario for using the include-exclude preview function.

For example, follow these steps to back up the files on your `/Users/home` file space:

1. Bring up the Tivoli Storage Manager Client Java GUI and open the Backup tree. You can see all of the directories and files that have been excluded by your options file and other sources.
2. Scroll down the tree and notice that all of the `*.o` files in your `/Volumes/home/mary/myobjdir` are backed up.

3. You don't 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 574

"Preview Backup" on page 575

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, Tivoli Storage Manager 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, Tivoli Storage Manager evaluates the include-exclude list from the bottom up and stops when it finds an include or exclude statement that matches the file it is processing. The order in which the include and exclude options are entered therefore affects which files are included and excluded.

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:

3. If a match *is not* found, files are implicitly included and backed up.
4. When a file is backed up, it is bound to the default management class unless it matched an `include` statement that specified a different management class name, in which case the file is bound to that management class.

The following examples demonstrate *bottom up* processing.

Example 1

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

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

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

1. Rule 3 (the last statement defined) is checked *first* because of bottom-up processing. The pattern `c:\foo\junk*.obj` does not match the file name that is being processed.
2. Processing moves to Rule 2 and checks. This time, pattern `c:\foo\...*.obj` matches the file name that is being processed. Processing stops, the option is checked, and it is 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 71

Chapter 12, “Processing options,” on page 269

Related reference

“Exclude options” on page 364

“Query Includexcl” on page 593

Processing rules when using UNC names

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

Tivoli Storage Manager uses the rules described in “Include and exclude option processing” on page 80. Also the rules in “Explicit use of UNC names for remote drives” apply.

Explicit use of UNC names for remote drives

Tivoli Storage Manager recognizes explicit use of UNC names for remote drives.

For example, as shown in Table 16, 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 16. UNC name patterns and DOS patterns

| UNC name pattern | DOS pattern |
|--|----------------------------------|
| <code>\\remote\c\$\include\file.out</code> | <code>r:\include\file.out</code> |
| <code>\\remote\c\$\...\file.out</code> | <code>r:\...\file.out</code> |
| <code>\\remote\share4\exclude*</code> | <code>s:\exclude*</code> |
| <code>\\remote\share2\...\?.out</code> | <code>t:\...\?.out</code> |

Conversion of DOS pathnames for fixed and remote drives

Tivoli Storage Manager 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\sharex\file.txt
           \\remote\share\file[a-z].txt
matches    \\remote\share\fileg.txt
```

Chapter 3. Getting started

This topic describes some tasks that you need to do before you use the Tivoli Storage Manager Backup-Archive Client.

Perform these tasks before using Tivoli Storage Manager:

- “Starting a Java GUI session” on page 87
- “Starting a command-line session” on page 88
- “Starting a web client session” on page 91
- “Start the client scheduler automatically” on page 93
- “Changing your password” on page 93
- “Sorting file lists using the Tivoli Storage Manager GUI” on page 94
- “Displaying online help” on page 95
- “Ending a session” on page 95

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 17 summarizes the user security rights needed for backup and restore.

Table 17. Required user security rights for Tivoli Storage Manager backup and restore services

| Operating system | Account | What can I back up and restore? |
|----------------------------|--|---|
| Windows 2008 | Member of Administrators group or member of Backup Operators group | <ul style="list-style-type: none">• All file and directory objects• System state data (Backup Operators group cannot back up ASR writer data and cannot restore system state data) |
| Windows 2003 | Member of Administrators group or member of Backup Operators group | <ul style="list-style-type: none">• All file and directory objects• System state data |
| Windows Vista or Windows 7 | Member of Administrators group or member of Backup Operators group | <ul style="list-style-type: none">• All file and directory objects• System state data (Backup Operators group cannot back up ASR writer data and cannot restore system state data) |

Table 17. Required user security rights for Tivoli Storage Manager backup and restore services (continued)

| Operating system | Account | What can I back up and restore? |
|--|--------------------|---|
| Windows 2003, Windows Vista, or Windows 7 | All other accounts | <ul style="list-style-type: none"> • All file and directory objects to which the account has access • System Access Control Lists (SACL) used for auditing entries is not backed up or restored |
| <p>The account must have the "Backup files and directories" and "Restore files and directories" security settings to back up these objects:</p> <ul style="list-style-type: none"> • File objects • Directory objects • System state data with Windows 2003, Windows Vista, or Windows 7 <p>By default, the Administrators and Backup Operators groups are granted these security settings.</p> | | |

By default, Tivoli Storage Manager client services run under the local system account. However, the local system account does not have access to network mapped drives and does not have the same permissions and logon properties as a user that is logged 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.

Here is a list of the Tivoli Storage Manager Backup-Archive client operations that a member of the Backup Operators can do:

- Back up and restore files (see Table 17 on page 83)
- Back up system state

On Windows Vista, Windows 7, and Windows 2008, you must be a member of the Administrators group to back up ASR writer data.

- Restore system state (Windows 2003 only)

On Windows Vista, Windows 7, and Windows 2008, you must be a member of the Administrators group to restore system state.

- Start the TSM scheduler service

Here is a list of the Tivoli Storage Manager Backup-Archive client operations that a member of the Backup Operators cannot do:

- Start any other services (TSM client acceptor, TSM remote client agent, and TSM 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 Tivoli Storage Manager Backup-Archive client with an Administrators group account you might not be able to launch the client because you cannot open the log files (for example `dsmerror.log`). To alleviate this problem, you can grant the Backup Operators group Read and Write permissions to the log files or the directories containing these log files.
- If you have been creating Windows 2003, Windows Vista, or Windows 7 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 109

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. See the appropriate *Tivoli Storage Manager Administrator's Reference* for more information about the QUERY ADMIN command.

Related reference

"Revokeremoteaccess" on page 445

User account control

Starting with Microsoft Windows Vista and Windows 2008, the security architecture was updated with a new User Account Control (UAC) component.

The UAC component changes how user accounts are configured and used, and how applications use system-level privileges.

As a result, when you start a Tivoli Storage Manager client on those operating systems, the Windows Security dialog box is displayed. Depending on how you logged on, the form of the dialog might differ:

- Members of the Administrators group are prompted for confirmation to start the Tivoli Storage Manager client.
- Members of the Backup Operators group are prompted to provide the password for a Backup Operators group account or for an administrator account.

In addition, members of the Backup Operators group must have the following access privileges to the Tivoli Storage Manager client root directory:

- Modify
- Read and Execute
- List Folder Contents
- Read
- Write

Starting a Java GUI session

This section explains the different methods to start a GUI session.

Start a GUI session using one of the following methods:

- Click the Windows **Start** button and select **Programs** → **Tivoli Storage Manager** → **Backup Archive GUI**.
- Click the Windows **Start** button and select **Run**, then enter the full path to the backup client executable file (**dsm.exe**).
- To start the non-Java GUI, click **Start** > **Run**, then enter the full path to the **dsmmfc.exe** file.
- On the command line, change directory to the Tivoli Storage Manager installation directory and enter **dsm**.

- On Windows Vista or Windows 7, when you log on as an administrator, you are running as a standard user. If you open a command prompt without running as an administrator, you are asked for consent to run **dsm**.

Tivoli Storage Manager locates the client user-options file (dsm.opt) and starts with the options specified in that file.

Related concepts

Chapter 2, “Configure the Tivoli Storage Manager client,” on page 19

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 93

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 Client Command Line**

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 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 83

“Options in interactive mode” on page 526

“Start and end a client command session” on page 524

Chapter 13, “Using commands,” on page 521

Using batch mode

Use *batch* mode to enter a single client command. When you use batch mode, you must precede the command with **dsmc**.

For example, to 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, Tivoli Storage Manager 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 423

Issuing a series of commands by using interactive mode

Use *interactive* mode when you want to issue a series of commands.

Tivoli Storage Manager establishes the connection to the server 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 Vista or 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, Tivoli Storage Manager 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).

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.

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.

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 382

“Set Access” on page 644

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 installing the web client on your workstation you can use the web client to perform backup, archive, restore, and retrieve operations from any browser that is running Java Runtime Environment (JRE) 5 or 6.

The web client facilitates the use of assistive devices for users with disabilities and contains improved keyboard navigation. The native look and feel of the platform running the browser is preserved.

The web client uses most desktop font and color settings when it is run in browsers on Windows platforms.

Refer to the software requirements topic for your platform to find the browsers where the web client runs.

For Microsoft Internet Explorer browsers, you must enable the **Scripting of Java applets**. This setting is also enabled by default. You can verify this by following these steps:

1. Open the **Tools** menu and select **Internet Options**
2. From the Internet Options dialog, select the **Security** tab.
3. Click the web content zone in which you use the IBM Tivoli Storage Manager web client and then click the **Custom Level** button.
4. In the Security Settings dialog, ensure that **Enable** is selected under the **Scripting of Java applets** setting.

If your browser does not have the correct JRE level, the web client notifies you and if possible, tries to automatically install the correct JRE for you.

For Microsoft Internet Explorer on Windows platforms, the web client automatically downloads and installs JRE 5 or 6.

You can download and install JRE 5 or 6 (unless otherwise noted), manually from the appropriate URL:

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

To use the web client, specify the URL of the client workstation running the web client in your web browser. You must also specify the HTTP port number defined on the client workstation; the default is 1581. For example:

<http://myhost.mycompany.com:1581>

Note: Entering a different URL or pressing the browser **Back** button during an operation disconnects the web client and causes the current operation to end.

Note: Back up and restore activities running in conjunction with a NAS box continues after the web client disconnects.

Related tasks

“Configuring the Web client” on page 22

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 86

Related reference

“Revokeremoteaccess” on page 445

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

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Changing your password

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

Tivoli Storage Manager prompts you for the password if one is required. Contact your Tivoli Storage Manager administrator if you do not know your password.

Important: The password discussed in this topic is different than the password used for encrypting files.

To change your password from the GUI:

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

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

For UNIX, Linux, and Windows clients:

```
dsmc set password
```

Then, enter your old and new passwords when prompted.

A Tivoli Storage Manager password can be up to 63 characters. Valid characters are:

| Character | Description |
|-----------|--|
| A–Z | Any letter; A through Z, upper or lower case |
| 0–9 | Any number; 0 through 9 |
| + | Plus |
| . | Period |
| _ | Underscore |
| - | Hyphen |
| & | Ampersand |

A password is not case sensitive.

Related concepts

“Start the client scheduler automatically” on page 93

Related reference

“Password” on page 422

“Set Password” on page 648

Sorting file lists using the Tivoli Storage Manager GUI

This topic explains how to work with your files using the Tivoli Storage Manager GUI.

Table 18. Working with your files using the Tivoli Storage Manager GUI

| Task | Procedure |
|---|--|
| Displaying files | To display files in a directory, click the folder icon next to the directory name. The files appear in the File List box on the right. |
| Sorting the file list | <ul style="list-style-type: none"> 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 Tivoli Storage Manager GUIs.
2. Using the Tivoli Storage Manager 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 391

“Pick” on page 424

Displaying online help

You can display online help in any of the following ways: On the Tivoli Storage Manager client GUIs, from the Web client, or from the **dsmc** command line.

- **On Tivoli Storage Manager client GUIs:**
 - Open the **Help** menu.
 - 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 562

Ending a session

You can end a Tivoli Storage Manager client session from the Tivoli Storage Manager client GUI or from the **dsmc** command line.

- **From the Tivoli Storage Manager client GUI main hub 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. Tivoli Storage Manager 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 Tivoli Storage Manager 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 572

Online forums

To participate in user discussions of Tivoli Storage Manager you can subscribe to the ADSM-L list server.

This is a user forum maintained by Marist College. While not officially supported by IBM, Tivoli Storage Manager developers and other IBM support staff also participate on an informal, best-effort basis. Because this is not an official IBM support channel, you should contact IBM Technical Support if you require a response specifically from IBM. Otherwise there is no guarantee that IBM will respond to your question on the list server.

You can subscribe by sending a note to the following e-mail address:

`listserv@vm.marist.edu`

The body of the message must contain the following:

`SUBSCRIBE ADSM-L yourfirstname yourlastname`

The list server will send you a response asking you to confirm the subscription request. Once you confirm your subscription request, the list server will send you further instructions. You will then be able to post messages to the list server by sending e-mail to:

`ADSM-L@vm.marist.edu`

If at a later time you want to unsubscribe from ADSM-L, you can send a note to the following e-mail address:

`listserv@vm.marist.edu`

The body of the message must contain the following:

`SIGNOFF ADSM-L`

You can also read and search the ADSM-L archives, join discussion forums, and access other resources at the following URL:

<http://www.adsm.org>

Other sources of online help

An anonymous FTP server (<ftp://public.dhe.ibm.com/storage>) is available where you can find maintenance and other Tivoli Storage Manager-related materials.

Three other anonymous servers are unofficially maintained by non-IBM volunteers.

The following are the anonymous servers:

`ftp.rz.uni-karlsruhe.de` (mirror - Germany)

`ftp.wu-wien.ac.at` (mirror - Austria)

`ftp.cac.psu.edu` (mirror - Pennsylvania)

You can get maintenance information from the Tivoli Storage Manager support page at: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager

Chapter 4. Backing up your data

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

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

- Estimate
- Preferences editor
- Setup wizard

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

Tivoli Storage Manager provides backup and archive services for all files residing in the following types of file systems: File Allocation Table (FAT), NT File System (NTFS), and FAT 32.

The following is a list of primary backup tasks.

- “Planning your backups (Windows)”
- “Pre-backup considerations (Windows)” on page 106
- “Incremental, selective, or incremental-by-date backups (Windows)” on page 112
- “Deleting backup data” on page 105
- “Backing up files from one or more file spaces for a group backup (Windows)” on page 118
- “Backing up Windows system state” on page 122
- “Backing up Automated System Recovery files” on page 124
- “Image backup” on page 129
- “Back up NAS file systems using Network Data Management Protocol” on page 135
- “Preparing the environment for VMware backup processing” on page 143
- “Backing up virtual machines on a Windows Hyper-V system” on page 146
- “Net Appliance CIFS share definition backups” on page 149

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 19. 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 105 for more information. |
|--------------------------|--|
-

Table 19. Planning your backups (continued)

| | |
|--------------------------|---|
| <input type="checkbox"/> | See “Pre-backup considerations (Windows)” on page 106 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 110 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 112• “Backing up files from one or more file spaces for a group backup (Windows)” on page 118• “Backing up Windows system state” on page 122• “Backing up Automated System Recovery files” on page 124• “Image backup” on page 129• “Back up NAS file systems using Network Data Management Protocol” on page 135 |
| <input type="checkbox"/> | For additional backup considerations, see “Backup (Windows): Additional considerations” on page 152. |

Understanding which files are backed up

When you request a backup, Tivoli Storage Manager backs up a file if some requirements are met.

The following are the requirements that must be met for Tivoli Storage Manager to back up a file.

- The selected management class contains a backup copy group.
- The file meets the serialization requirements defined in the backup copy group. If serialization is *static* or *shared static*, and the file changes during backup, the file will not be backed up.
- The file meets the mode requirements defined in the backup copy group. If the mode is *modified*, the file must have changed since the last backup. If the mode is *absolute*, the file can be backed up even if it does not change.
- The file meets the frequency requirements defined in the backup copy group. The specified minimum number of days since the last backup must elapse before a file is backed up.
- 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 performing a system state backup. You can back up the system state only as a single entity because of dependencies among the system state components. You can not 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 231

“Management classes and copy groups” on page 232

Related tasks

“Backing up Windows system state” on page 122

Open file support for backup operations

There are two snapshot providers that can be used for open file support: LVSA and VSS.

VSS is the preferred provider, since it utilizes the strategic snapshot solution from Microsoft. The LVSA is provided for customers migrating from prior versions of Tivoli Storage Manager, where LVSA was utilized, and as an alternative to VSS if there are issues with using VSS.

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

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

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

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

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

OFS should not be used for backing up locked Windows system files, such as the Windows system state (Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7). The Tivoli Storage Manager client has advanced features for backing up data contained within these files. The backup of the system data that is contained in these files requires additional processing and must be backed up in a group to allow for a successful restore. These files are excluded from the Tivoli Storage Manager file level backup.

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

features. For a current list of Data Protection clients go to this Web site:
<http://www.ibm.com/software/tivoli/products/storage-mgr/product-links.html>

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

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

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

For information about Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client Open File Support** document, which is Technote number 1248971, at: http://www.ibm.com/support/entry/portal/Overview/Software/Tivoli/Tivoli_Storage_Manager

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

To control an open file support operation you can specify these additional options in your *dsm.opt* file or as values of the *include.fs* option: *snapshotproviderfs*, *presnapshotcmd* and *postsnapshotcmd*. Additionally, when the LVSA is the snapshot provider, the following additional options can be specified: *snapshotcachelocation*, *snapshotcachesize*, *snapshotfsidleretries*, and *snapshotfsidlewait*.

Note:

1. You can use the *include.fs* option to set snapshot options on a per file system basis.
2. For the LVSA, use the *snapshotcachelocation* option to relocate the cache if necessary. You can specify a snapshotcachelocation for a specific drive using the *include.fs* option.

3. Open file support is provided for both backup and archive. For backup, this includes incremental, incremental by date, selective, incremental image, adaptive subfile backup, and journal-based backup.
4. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
5. For the LVSA, if the client is unable to create a snapshot; for example, if the `snapshotfsid` is not met or the `snapshotlocation` option is not valid, failover to a non-OFS backup occurs on the volumes where this problem exists.
6. To enable OFS support in a cluster environment, all workstations in the cluster must have OFS configured. If LVSA is being used as the snapshot provider, ensure it is installed and configured on each workstation in the cluster. You can also use VSS by setting the `snapshotproviderfs` option.

Related concepts

“Upgrading Open File Support or online image” on page 3

Chapter 12, “Processing options,” on page 269

Related tasks

“Backing up Windows system state” on page 122

“Configuring Open File Support” on page 66

Backing up data using the GUI

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

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

To perform a GUI backup, use the following steps:

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

The following are some items to consider when using 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 used is either a default selected for you, or one that you assign to the file using an *include* option in the include-exclude options list. Select **Utilities** → **View Policy Information** from the backup-archive client or Web client GUI to view the backup policies defined by the Tivoli Storage Manager server for your client node. Select **Edit** → **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 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**.

Related concepts

Chapter 9, “Storage management policies,” on page 231

Related tasks

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Specifying drives in your domain

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

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

You can exclude any domain (including the `systemobject` domain) in your default domain from backup processing using the **Backup** tab in the Preferences editor. You can also exclude drives or the `systemobject` domain by specifying the dash (-) operator before the drive or the `systemobject` domain. For example, in the following option Tivoli Storage Manager 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 Tivoli Storage Manager GUI to perform these backup functions.

Related reference

"Domain" on page 343

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.

Table 20. 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 564 for more information about the incremental command. See "Full and partial incremental backup" on page 112 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 343 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 343 for more information. For Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 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 343 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 482 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 474 for more information. |

Incremental-by-date backup

Table 20. Command line backup examples (continued)

| Task | Command | Considerations |
|---|--|---|
| Perform an incremental-by-date backup of your default client domain. | <code>dsmc incremental -incrbydate</code> | Use the <code>incrbydate</code> 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 399 for more information about the <code>incrbydate</code> option. |
| <i>Selective backups</i> | | |
| Back up all files in the <code>d:\proj</code> directory. | <code>dsmc selective d:\proj\</code> | 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 640 for more information about the selective command. |
| Back up the <code>d:\proj</code> directory and all its subdirectories. | <code>dsmc selective d:\proj\ -subdir=yes</code> | See “Subdir” on page 482 for more information about the <code>subdir</code> option. |
| Back up the <code>d:\h1.doc</code> and <code>d:\test.doc</code> files. | <code>dsmc selective d:\h1.doc d:\test.doc</code> | You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the <code>filelist</code> option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. See “Filelist” on page 377 for more information. |
| Back up a list of files in the <code>c:</code> drive. | <code>dsmc selective -filelist=c:\filelist.txt</code> | Use the <code>filelist</code> option to process a list of files. See “Filelist” on page 377 for more information. |
| Assuming that you initiated a snapshot of the <code>C:</code> drive and mounted the snapshot as the logical volume <code>\\florence\c\$\snapshots\snapshot.0</code> , run a selective backup of the <code>c:\dir1\sub1</code> directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name <code>C:</code> . | <code>dsmc selective c:\dir1\sub1* -subdir=yes snapshot=\\florence\c\$\snapshots\snapshot.0</code> | See “Snapshotroot” on page 474 for more information. |

Related concepts

“Backup (Windows): Additional considerations” on page 152
Chapter 13, “Using commands,” on page 521

Deleting backup data

If your administrator has given you authority, you can delete individual backup copies from the Tivoli Storage Manager server without deleting the entire file space.

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

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:

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**.

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

Related reference

“Delete Backup” on page 556

When to back up and when to archive files

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

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

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

- The number of versions
- The number of days each additional backup copy is kept

- What happens to backup data versions when the file is deleted on the client system

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

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

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

Backups protect against file damage or loss that could occur through accidental deletion, corruption, disk crashes, and so forth. The server maintains one or more backup versions for each file that you back up. Older versions are deleted as newer versions are made. The number of backup versions the server maintains is set by your administrator.

Archive copies are saved for long-term storage. Your administrator can limit how long archive copies are kept. The server can store an unlimited number of archive versions of a file. Archives are useful if you need to go back to a particular version of your files, or you want to delete a file from your workstation and retrieve it at a later time, if necessary. For example, you might need to save spreadsheets for tax purposes, but because you are not using them, you do not want to leave them on your workstation.

Related concepts

“Restore data from a backup set” on page 180

Pre-backup considerations (Windows)

This section discusses some items to consider 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 Tivoli Storage Manager 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 Managed System for SAN Storage Agent on the client workstation.

The Tivoli Storage Manager Managed System for SAN Storage Agent is a separate product and it is included on its own DVD, with the Tivoli Storage Manager for SAN product.

For more information about installing and configuring the Tivoli Storage Manager Managed System for SAN Storage Agent, refer to the IBM Tivoli Storage Manager Storage Agent User's Guides.

LAN-free data movement options

This topic describes the options that you can use to enable LAN-free data movement, after you have installed and configured the Tivoli Storage Manager Managed System for SAN feature on the client workstation.

The following are the options that you can use 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.

lanfreetcport

Specifies the TCP/IP port number where the Storage Agent is listening.

lanfreetcpserveraddress

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

Related reference

“Enablelanfree” on page 355

“Lanfreecommmethod” on page 402

“Lanfreeshmport” on page 404

“Lanfreetcport” on page 404

“Lanfreetcpserveraddress” on page 405

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 393

Related reference

“Autofsrename” on page 307

“Detail” on page 335

“Query Filespace” on page 588

“Restore” on page 605

“Retrieve” on page 635

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 Tivoli Storage Manager to process only one directory at a time, which reduces memory consumption but increases backup time. When you specify `yes`, Tivoli Storage Manager analyzes only one directory at a time for backup consideration. If performance remains poor, check your communication buffer settings and the communication link between your system and the Tivoli Storage Manager server. If your system is not memory constrained, setting the `memoryefficientbackup` option to `yes` degrades your backup performance.

Related reference

“Memoryefficientbackup” on page 411

Incremental backups on systems with a large number of files

The client can use very large amounts of memory to perform incremental backup operations, especially on file systems that contain large numbers of files.

The term *memory* as used here is the addressable memory available to the client process. Addressable memory is a combination of physical RAM and virtual memory.

On average, the client uses approximately 300 bytes of memory per object (file or directory). Thus for a file system with one million files and directories, the Tivoli Storage Manager client requires, on average, approximately 300 MB of memory. The exact amount of memory 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 Tivoli Storage Manager backup client memory requirement.

The maximum number of files can be determined by dividing the maximum amount of memory available to a process by the average amount of memory needed per object.

The total memory requirement can be reduced by any of the following methods:

- Use the client option `memoryefficientbackup diskcachemethod`. This choice reduces the use of memory to a minimum at the expense of performance and a significant increase in disk space required for the backup. The file description data from the server is stored in a disk-resident temporary database, not in memory. As directories on the workstation are scanned, the database is

consulted to determine whether to back up, update, or expire each object. At the completion of the backup, the database file is deleted.

- Use the client option `memoryefficientbackup yes`. The average memory used by the client then becomes 300 bytes times the number of directories plus 300 bytes per file in the directory being processed. Note that for file systems with large numbers (millions) of directories, the client still might not be able to allocate enough memory to perform incremental backup with `memoryefficientbackup yes`.
- If the client option `resourceutilization` is set to a value greater than 4, and there are multiple file systems being backed up, then reducing `resourceutilization` to 4 or lower limits the process to incremental backup of a single file system at a time. This reduces the memory requirement. If backup of multiple file systems in parallel is required for performance reasons, and the combined memory requirements exceed the process limits, then multiple instances of the backup client can be used to back up multiple file systems in parallel. For example, if you want to back up two file systems at the same time but their memory requirements exceed the limits of a single process, then start one instance of the client to back up one of the file systems, and start a second instance of the client to back up the other file system.
- Use the `- incrbydate` client option to perform an "incremental-by-date" backup.
- Use the Tivoli Storage Manager client `exclude.dir` option to prevent the client from traversing and backing up directories that do not need to be backed up.
- Reduce the number of files per file system by spreading the data across multiple file systems.

Related reference

"`Snapdiff`" on page 463

"`Exclude options`" on page 364

"`Incrbydate`" on page 399

"`Memoryefficientbackup`" on page 411

"`Resourceutilization`" on page 442

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*.

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

Perform the following steps to use adaptive subfile backup:

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

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 368

“Determine compression and encryption processing” on page 78

Related reference

“Exclude options” on page 364

“Subfilebackup” on page 484

“Subfilecachepath” on page 485

“Subfilecachesize” on page 486

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 containing other files you want to exclude.

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

Tivoli Storage Manager uses *management classes* to determine how to manage your backups on the server. Every time you back up a file, the file is assigned a management class. The management class is either a default chosen for you, or one you assign to the file using the `include` option in the include-exclude list. If you assign a management class, it must contain a backup copy group for the file to be backed up.

You can also add include-exclude statements in the Tivoli Storage Manager Client GUI directory tree. The **preview** command allows you to see the resultant effects of the currently defined include-exclude list without need of performing an actual backup of data.

Related tasks

“Creating an include-exclude list” on page 69

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Related reference

“**Preview Backup**” on page 575

Data encryption during backup or archive operations

For the strongest possible encryption, use 128-bit Advanced Encryption Standard (AES) data encryption, with the `encryptiontype` option.

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.

To encrypt file data, you must select an encryption key password, which Tivoli Storage Manager uses to generate the encryption key for encrypting and decrypting the file data. You can specify whether to save the encryption key password in the Windows Registry by using the `encryptkey` option.

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

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

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

Related concepts

“Backup (Windows): Additional considerations” on page 152

Related reference

“Encryptiontype” on page 356

“Encryptkey” on page 357

“Exclude options” on page 364

“Include options” on page 393

Maximum file size for operations

The table in this topic specifies the maximum file sizes for backup, restore, archive, and retrieve operations.

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

Table 21. Maximum file size

| File system | Max file size (in bytes) |
|--------------------|---------------------------------|
| FAT16 | 2 147 483 647 (2 GB) |
| FAT32 | 4 294 967 295 (4 GB) |
| NTFS | 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 above outlined rule.

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

Your administrator might set up schedules to automatically perform backups. The following sections discuss how to back up files without a schedule.

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

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

Related tasks

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Full and partial incremental backup

An incremental backup backs up only new and changed files. Whether you perform a full or partial incremental backup depends on what you select for backup.

If you select entire drives, you are performing a full incremental backup. If you select a directory tree or individual files, you are performing a partial incremental backup.

The first time you run a full incremental backup, Tivoli Storage Manager backs up all the files and directories that you specify. This process can take a long time if the number of files is large, or if one or more very large files must be backed up. Subsequent full incremental backups only back up new and changed files. This allows the backup server to maintain current versions of your files, without having to waste time or space by backing up files that already exist in server storage.

Depending on your storage management policies, the server might keep more than one version of your files in storage. The most recently backed up files are active backup versions. Older copies of your backed up files are inactive versions. However, if you delete a file from your workstation, the next full incremental backup causes the active backup version of the file to become inactive. If you need

to restore a file you have deleted, and if a full incremental backup has been run since you deleted the file, then you need to restore an inactive version of the file (assuming that a version still exists on the server). The number of inactive versions maintained by the server and how long they are retained is governed by the management policies defined by your server administrator. The purpose of the active versions is to represent which files existed on your file system at the time of the backup.

To perform a full or partial incremental backup using the client Java GUI, select the incremental (complete) option from the pull-down menu at the top of the screen on the backup window, or use the **incremental** command in the command-line interface. Specify file systems, directory trees, or individual files to include in the backup.

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

- Back up new files.
- Update the attributes on the server. Changes to read-only, hidden and compressed file attributes causes Tivoli Storage Manager to send only attribute changes to the server. Changes to sparse file, reparse point, and encrypted file attributes cause Tivoli Storage Manager to send the entire file to the server. The archive attribute is not examined by Tivoli Storage Manager in determining changed files.
- Back up files whose contents changed since the last backup.
- Changes include *any* of the following:
 - File size
 - Date or time of last modification
 - Extended Attributes
 - Access Control List
 - Sparse, reparse point or encrypted file attributes.
 - NTFS file security descriptors. These are the Owner Security Identifier (SID), Group SID, Discretionary Access Control List (ACL), and System ACL.
- Expire backup versions of files on the server that do not have corresponding files on the workstation. The result is that files which no longer exist on your workstation will not have active backup versions on the server. However, inactive versions 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 rebound. Objects for which only inactive backup versions exist are not rebound.

Important: Each directory is also backed up under the following circumstances:

- It has not been backed up previously
- Its permissions have changed since the last backup
- Its Access Control List has changed since the last backup
- Its Extended Attributes have changed since the last backup
- Its time stamp has changed since the last backup

You might be able to use the `preserveLastAccessDate` option during a backup or archive operation to specify whether to reset the last access date to its original value following a backup or archive operation. By default, the Tivoli Storage

Manager client *will not* reset the last access date of any backed up or archived files to their original value prior to the backup or archive operation.

Directories are counted in the number of objects backed up. To exclude directories and their contents from backup, use the `exclude.dir` option.

Related concepts

Chapter 9, “Storage management policies,” on page 231

Related reference

“Exclude options” on page 364

“PreserveLastAccessDate” on page 431

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

Related tasks

“Configuring the journal engine service” on page 33

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 K) 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 `snapdiff` option to invoke the snapshot difference backup from NetApp when running a full-volume incremental backup. Using this option reduces memory usage and is faster.

Consider the following restrictions when running a full-volume incremental backup using the `snapdiff` option, to ensure that data is backed up when it should be.

- A file is excluded due to an exclude rule in the include-exclude file. Tivoli Storage Manager 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. Tivoli Storage Manager does not inspect every file on the volume during backup.
- If you used the **`dsmc delete backup`** command to explicitly delete a file from the Tivoli Storage Manager inventory, NetApp cannot detect that a file was manually deleted from Tivoli Storage Manager. 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 Tivoli Storage Manager to back it up again.
- Policy changes such as changing the policy from **`mode=modified`** to **`mode=absolute`** are not detected.
- The entire file space is deleted from the Tivoli Storage Manager inventory. This action causes the `snapdiff` option to create a new snapshot to use as the source, and a full incremental backup to be run.

The NetApp software determines what is a changed object, not Tivoli Storage Manager.

If you run a full volume backup of an NFS-mounted or a CIFS-mapped NetApp or N-Series volume, all the snapshots under the snapshot directory might also be backed up. To avoid this situation, you can 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.

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 101

Related reference

“Selective” on page 640

Backing up files from one or more file spaces for a group backup (Windows)

You can use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

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

- All objects in the group are assigned to the same management class. 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 using the `mode` option.

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 180

Related reference

“**Backup Group**” on page 538

“Include options” on page 393

“Mode” on page 412

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.

This is useful when the workstation responsible for performing the backup can change over time, such as with a cluster.

The `asnodename` option also allows data to be restored from a different system than the one which performed the backup.

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

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

Use the `asnodename` option with the appropriate command to back up, archive, restore, and retrieve data under the target node name on the Tivoli Storage Manager server. This support is only available with Tivoli Storage Manager Version 5.3 and higher server and client.

To enable this option, follow these steps:

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

Suggestions:

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

Restrictions enforced within a proxied session:

- You cannot perform a system state object backup or restore.

- If a backup set contains system state object, it will not be restored within a proxied session.
- You cannot access another node (either from GUI drop down or use of the `fromnode` option).
- You cannot use the `clusternode` option.
- You cannot perform NAS backup or restore.

Related reference

“Asnodename” on page 299

Enabling multiple node operations from the GUI

This topic lists the steps that you must follow to enable multiple node operation from the GUI.

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the `QUERY PROXYNODE` administrative client command.
2. Select **Edit** → **Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the target node.
4. Click **Apply** and then **OK** to close the preferences window.

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

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

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

Setting up encryption

This topic lists the steps that you must follow to set up encryption with the `encryptkey` option.

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

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

1. Specify `encryptkey=prompt` in the options file.
2. Ensure that users of the agent nodes in the multiple node environment are using the same encryption key.

Important:

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

Scheduling backups with client-node proxy support

Multiple nodes can be used to perform backup operations using the scheduler.

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='-asnodel=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: (1) `register node NODE-A`, (2) `register node NODE-B`, (3) `register node NODE-C`.
- The administrator grants proxy authority to each agent node, by issuing the following commands: (1) `grant proxynode target=NODE-Z agent=NODE-A`, (2) `grant proxynode target=NODE-Z agent=NODE-B`, (3) `grant proxynode target=NODE-Z agent=NODE-C`.
- The administrator defines the schedules, by issuing the following commands: (1) `define schedule standard proxy1 description="NODE-A proxy schedule" action=incremental options="-asnodel=NODE-Z" objects=C: startdate=05/21/2005 starttime=01:00`, (2) `define schedule standard proxy2 description="NODE-B proxy schedule" action=incremental options="-asnodel=NODE-Z" objects=D: startdate=05/21/2005 starttime=01:00`, (3) `define schedule standard proxy3 description="NODE-C proxy schedule" action=incremental options="-asnodel=NODE-Z" objects=E: startdate=05/21/2005 starttime=01:00`.

Note: Place the `asnodel` option in the schedule definition only. Do not place it in the client options file, on the command line, or in any other location.

For more information about the server scheduler commands, see *IBM Tivoli Storage Manager Administrator Command Reference*.

Start the schedules by either configuring a Tivoli Storage Manager schedule service, or by using the following client command: `dsmc sched`

You can also use the client acceptor daemon, 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.

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 474

Backing up Windows system state

Tivoli Storage Manager uses VSS to back up all system state components as a single object, to provide a consistent point-in-time snapshot of the system state. System state consists of all startable system state and system services components.

Tivoli Storage Manager supports the Microsoft Volume Shadowcopy Service (VSS) on Windows Server 2003, Windows Server 2008, Windows 7, and Windows Vista.

The system state backup consists of data from several VSS writers. When the files that belong to the System Writer change, an incremental backup is used for these files. Using incremental backup on the System Writer files reduces the amount of time it takes to back up the system state. A full backup is used for all other system state data.

By default, a progressive incremental backup is used for the System Writer files in the system state. If you want to run a full backup of all system state data, you can specify `include.systemstate mc_name` in the client options file (`dsm.opt`), where `mc_name` is the name of the management class with copy group mode absolute.

Startable system state components include the following:

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

System services components include the following:

- Background Intelligent Transfer Service (BITS)
- Event logs (2003 only)
- Removable Storage Management Database (RSM)
- Cluster Database (cluster node only)
- Remote Storage Service
- Terminal Server Licensing
- Windows Management Instrumentation (WMI)
- Internet Information Services (IIS) metabase
- Dynamic Host Configuration Protocol (DHCP)
- Windows Internet Name Service (WINS)

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

You must be a member of the Administrators or Backup Operators group to back up system state information.

Important:

1. Your Windows client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.

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.

You can use the `include.systemstate` option to assign management classes for back up of the system state. The default is to bind the system state object to the default management class.

You can use the `exclude.systemservice` option in your client options file (`dsm.opt`) to exclude individual system services components from backup. Input can be the keyword or component name to be excluded from backup.

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

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

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

On Windows Vista, Windows 7, and Windows Server 2008 clients, **backup systemstate** also backs up ASR data.

Related tasks

“Restoring Windows system state” on page 167

Related reference

“Backup Systemstate” on page 547

“Domain” on page 343

“Exclude options” on page 364

“Include options” on page 393

“Query Systemstate” on page 601

“Restore Systemstate” on page 629

Backing up Automated System Recovery files

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.

You must be a member of the Administrators or Backup Operators group to back up ASR files.

On Windows Vista, Windows 7, and Windows Server 2008, the Tivoli Storage Manager backup-archive client backs up ASR data when the backup-archive client backs up the Windows system state. Use the **backup systemstate** command to back up ASR files on these operating systems. On Windows Server 2003, backing up ASR files is a separate task from backing up system state.

Back up the ASR files by typing the following command: `dsmc backup asr`.

Tivoli Storage Manager 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

“Backing up in preparation for Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008)”

“Backing up in preparation for Automated System Recovery (Windows Server 2003)” on page 127

Related tasks

“Creating an Automated System Recovery diskette (Windows Server 2003)” on page 127

“Restoring Automated System Recovery files” on page 169

Related reference

“Backup ASR” on page 535

“Backup Systemstate” on page 547

Backing up in preparation for Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008)

Specific backups and media are required for Windows Vista, Windows 7, and Windows Server 2008 Automated System Recovery (ASR).

Note that ASR data is not backed up, and hence cannot be restored, on 64-bit architectures with Extensible Firmware Interface (EFI).

Creating a Tivoli Storage Manager options file for Automated System Recovery (Windows Vista, Windows 7, or Windows Server 2008)

Before you can recover a Windows Vista, Windows 7, or Windows Server 2008 computer using Automated System Recovery (ASR), you must create an options file. The options file is unique for each computer.

This task assumes that you created a generic bootable WinPE CD or DVD. A generic bootable WinPE CD does not contain the Tivoli Storage Manager 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.

1. Find a copy of the Tivoli Storage Manager 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 dsm.opt specify errorlogname x:\dsmerror.log.

Note: This example uses x: because in WinPE mode, the default system drive is x:.

- b. Enter the Tivoli Storage Manager node name with the nodename option.
 - c. Optional: If you plan to restore the system state from files stored on a Tivoli Storage Manager server, enter the Tivoli Storage Manager 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 Tivoli Storage Manager 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. The other ways to gain access to the server are the following:

- 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 tell you 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.

You created an options file that contains Tivoli Storage Manager configuration information that is unique for each computer. This information complements the generic bootable WinPE CD.

Related tasks

“Creating a bootable WinPE CD (Windows Vista, Windows 7, and Windows Server 2008)” on page 170

Backing up the boot drive and system drive for Automated System Recovery (Windows Vista, Windows 7, or Windows Server 2008)

Before you can recover your Windows Vista, Windows 7, or Windows Server 2008 computer using ASR (Automated System Recovery), you must have a complete backup of the boot drive and system drive.

1. Perform a full incremental backup of your system and boot drives. Assuming that your system and boot files are on drive c:, 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 112

Related tasks

“Backing up Windows system state” on page 122

Backing up in preparation for Automated System Recovery (Windows Server 2003)

Specific backups and media are required for Windows Server 2003 system recovery with ASR (Automated System Recovery).

In addition to the following tasks, consider using security policy administrative tools to enable the recovery console to access all paths and removable media. You can use the Windows Recovery Console to facilitate easier problem determination in the event of recovery errors. You must set these permissions prior to an Automated System Recovery situation.

Creating an Automated System Recovery diskette (Windows Server 2003)

To recover your Windows Server 2003 system using ASR (Automated System Recovery), you must create an Automated System Recovery diskette.

Create a recovery diskette any time there is a hardware configuration change or service pack update to your workstation.

Always create an ASR diskette after you back up ASR files. Store the ASR diskette in a safe location for future use.

You can also create the ASR diskette for another Tivoli Storage Manager client node if you know the password for the other node, and if the ASR files on the other node were backed up to a Tivoli Storage Manager server. To create the diskette for another node with the client GUI, set the `virtualnodename` option in your client options file (`dsm.opt`) to the node name of the workstation for which you are creating the ASR diskette prior to creating the ASR diskette. If you use the client command line interface, specify the other node name with the `virtualnodename` option like this: `dsmc restore asr a:\`

`-virtualnodename=dagordon`

1. Back up the ASR files by typing the following command: `dsmc backup asr`. The Windows Automated System Recovery window is displayed by the Windows API that creates the ASR files. ASR data is generated on your computer and backed up to the Tivoli Storage Manager server.
2. Copy the ASR files to the ASR diskette

Client command line interface

- a. Insert a blank, formatted diskette into the drive (assumed to be drive A for this example), then issue the following TSM command:
`dsmc restore asr a:\`
- b. Use the Windows **label** command to make the diskette label TSMASR. The volume label for the diskette must be TSMASR for recovery to succeed. For example: `label a: TSMASR`

client GUI

- a. Select **Utilities** → **Create ASR Diskette** from the GUI main window. The Create Automated System Recovery Diskette window appears.
- b. Specify the diskette drive for creating the ASR diskette, insert a diskette into the drive, and click **Finish**.

Tivoli Storage Manager copies the required ASR files to the diskette and labels the diskette volume TSMASR. If you forget to back up the ASR files prior to creating the diskette, the following error is shown:

No ASR filespace was found on server. Make sure you have backed up ASR prior to performing this operation.

3. Modify TSMASR.CMD on the ASR diskette.

Place the two commands **NTBACKUP** and **PAUSE** following the **dsmc restore systemstate** command, as in the following example:

```
:NETWORK
dsmc restore \\TSMWIN2K3\C$\ C:\ -subdir=yes
-preservepath=complete -replace=all -tapeprompt=no
-nodename=NTBACKUP -optfile=tsmasr.opt
dsmc restore systemstate -asrmode=yes -nodename=NTBACKUP
-computername=TSMWIN2K3 -optfile=tsmasr.opt
NTBACKUP
PAUSE
```

4. Modify TSMASR.OPT on the ASR diskette.

Remove the **passwordaccess generate** statement, if it is present.

Backing up boot drive and system drive for Automated System Recovery (Windows Server 2003)

To recover your Windows Server 2003 system using ASR (Automated System Recovery), you must have a complete backup of the boot drive and system drive.

You must back up the ASR files using the backup asr command before performing an incremental backup of the system and boot drives. Backing up the ASR files generates the ntdll.asr and smss.asr files and places them in the windows\repair directory. These files must be present in the incremental backup of your system and boot drives in order for Automated System Recovery to succeed. To verify that these files have been backed up, use the following command (assume that your Windows system is installed on the c: drive): QUERY BACKUP C:\WINDOWS\REPAIR*.ASR

1. Perform a complete incremental backup of your system and boot drives. The boot drive is the partition containing the boot files such as ntldr and boot.ini. The system drive is the partition containing your WINDOWS system directory. Assuming that your system and boot files are on drive c:, enter the following command: dsmc incremental c:
2. 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 with these commands to display file level detail.

Related tasks

“Backing up Windows system state” on page 122

Backing up user profiles for ASR recovery (Windows Server 2003)

Before recovering your Windows Server 2003 user profiles using ASR (Automated System Recovery), you must back up user profiles.

Use Ntbackup.exe to back up Windows user profiles to local storage. The Tivoli Storage Manager backup-archive client can then back up that local file.

1. Create a script which lists all of the instances of NTUSER.DAT and UsrClass.dat into a Ntbackup.exe backup script. Then run Ntbackup.exe to back up the files into local storage. The backup script is created in Unicode format which is required by Ntbackup.exe. This example command called c:\ntexport.cmd can create an Ntbackup.exe backup script called c:\tsmback.bks. c:\tsmback.bks backs up all instances of NTUSER.DAT and UsrClass.dat on the C: drive. If the user profiles are located on a different drive, modify this example to reflect the appropriate drive location:

```
REM list all ntuser.dat files on the C: drive (both hidden and
REM not hidden REM and put them in a file called tsmback.bks
cmd /u /c dir c:\ntuser.dat /s /b /ah > c:\tsmback.bks
cmd /u /c dir c:\ntuser.dat /s /b >> c:\tsmback.bks
cmd /u /c dir c:\usrclass.dat /s /b >> c:\tsmback.bks
cmd /u /c dir c:\usrclass.dat /s /b /ah >> c:\tsmback.bks
REM call ntbackup to backup the contents of tsmback.bks
ntbackup backup @c:\tsmback.bks /j "TSM Test Export" /f "c:\ntbackup.bkf
```

2. Run c:\ntexport.cmd. The command creates the file c:\ntbackup.bkf.
3. Back up the file c:\ntbackup.bkf using the command: DSMC SELECTIVE c:\ntbackup.bkf
4. Verify that the file c:\ntbackup.bkf was backed-up: DSMC QUERY BACKUP c:\ntbackup.bkf

Related tasks

“Restoring user profiles during ASR recovery (Windows Server 2003)” on page 177

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 FAT, FAT32, NTFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks used by the file system or smaller than the *imagegapsize* parameter is backed up.

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

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

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

An image backup provides the following benefits:

- Backs up file systems containing a large number of files faster than a full file system incremental backup.

- Improves the speed with which Tivoli Storage Manager restores file systems containing many small files.
- Conserves resources on the server during backups since only one entry is required for the image.
- Provides a point-in-time picture of your logical volume, which might be useful if your enterprise needs to recall that information.
- Restores a corrupt file system or raw logical volume. Data is restored to the same state it was when the last logical volume backup was performed.

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

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

If the Tivoli Storage Manager online image support is configured, Tivoli Storage Manager performs an online image backup, during which the volume is available to other system applications. The snapshot provider, as specified by the `snapshotproviderimage` option, maintains a consistent image of a volume during online image backup.

If you use VSS, you do not need to install LVSA.

You can use the `snapshotproviderimage` option with the **backup image** command or the `include.image` option to specify whether to perform an offline or online image backup.

Related tasks

“Configuring online-image backup support” on page 65

Related reference

“`Snapshotproviderimage`” on page 473

Performing prerequisite tasks before creating an image backup

This topic lists some items to consider before you perform an image backup.

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

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

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

Related concepts

Chapter 9, “Storage management policies,” on page 231

Related reference

“Exclude options” on page 364

“Include options” on page 393

“Snapshotproviderimage” on page 473

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.

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:
 - 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 101

“Performing an image backup using the GUI” on page 133

“Restoring an image using the GUI” on page 179

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.

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.

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 information

“Restoring an image using the GUI” on page 179

“Performing an image backup using the GUI”

Related tasks

“Performing an image backup using the GUI”

“Restoring an image using the GUI” on page 179

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 22. 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 Tivoli Storage Manager image feature is configured, you can create an image backup where the real volume is available to other system applications.

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, Tivoli Storage Manager honors the `snapshotproviderimage` setting in your client options file (`dsm.opt`). If the Tivoli Storage Manager online image support is configured, Tivoli Storage Manager performs an online image backup, during which the volume is available to other system applications.

To create an image backup of your file system or raw logical volume, perform the following steps:

1. Click on the **Backup** button in the Tivoli Storage Manager main window. The Backup window appears.
2. Expand the directory tree and select the objects you want to back up. To back up a raw logical volume, locate and expand the **RAW** directory tree object.
3. Click **Backup**. The Backup **Task List** window displays the backup processing status. The Backup Report window displays a detailed status report.
 - 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 during the current session *only*.
- Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files. To determine the actual stored image size, select **View** → **File Details**. The actual stored image size is noted in the Stored Size field.
- To modify specific backup options, click the **Options** button. The options you select are effective during the current session *only*.
- Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files. To determine the actual stored image size, select **View** → **File Details**. The actual stored image size is noted in the Stored Size field.

Related reference

“Snapshotproviderimage” on page 473

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 540

“Mode” on page 412

“Restore Image” on page 624

“Snapshotproviderimage” on page 473

Back up NAS file systems using Network Data Management Protocol

Through support of Network Data Management Protocol (NDMP), Tivoli Storage Manager Windows, AIX, and Solaris backup-archive clients can efficiently back up and restore network attached storage (NAS) file system images to automated tape drives or libraries that are locally attached to Network Appliance and EMC Celerra NAS file servers.

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

Incremental backup can also be used to back up NAS file system snapshots. See the **incremental** command and `snapshotroot`, `snaptiff`, `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 8

Related reference

“Diffsnapshot” on page 337

“Incremental” on page 564

“Snapdiff” on page 463

“Snapshotroot” on page 474

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 the *authentication=on* must be specified at the server.

You are always prompted for a user ID and password. To display NAS nodes and perform NAS functions, you must enter an authorized administrative user ID and password. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

You can use the *toc* option with the *include.fs.nas* option in the client options file to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use Tivoli Storage Manager for 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

“Restore NAS file systems” on page 200

Related tasks

“Configuring the Web client” on page 22

Related reference

“Toc” on page 496

Back up NAS file systems using the command line

You can use the command line to back up NAS file system images.

Table 23 lists the commands and options that you can use to back up NAS file system images from the command line.

Table 23. NAS options and commands

| Option or command | Definition | Page |
|-----------------------------|--|-------------------------------|
| <code>domain.nas</code> | Use the <code>domain.nas</code> option to specify the volumes to include in your default domain for NAS backups. | “Domain.nas” on page 346 |
| <code>exclude.fs.nas</code> | Use the <code>exclude.fs.nas</code> option to exclude file systems on the NAS file server from an image backup when used with the backup nas command. | “Exclude options” on page 364 |
| | This option is valid for all Windows clients. | |

Table 23. NAS options and commands (continued)

| Option or command | Definition | Page |
|--------------------------------|--|--------------------------------|
| <code>include.fs.nas</code> | Use the <code>include.fs.nas</code> option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the <code>toc</code> option with the <code>include.fs.nas</code> option in your client options file. This option is valid for all Windows clients. | "Include options" on page 393 |
| <code>query node</code> | Use the query node command to display all the nodes for which a particular administrative user ID has authority to perform operations. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using. | "Query Node" on page 595 |
| <code>backup nas</code> | 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 545 |
| <code>toc</code> | Use the <code>toc</code> option with the backup nas command or the <code>include.fs.nas</code> option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. | "Toc" on page 496 |
| <code>monitor process</code> | 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 573 |
| <code>cancel process</code> | 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 552 |
| <code>query backup</code> | Use the query backup command with the <code>class</code> option to display information about file system images backed up for a NAS file server. | "Query Backup" on page 581 |
| <code>query filesystem</code> | Use the query filesystem command with the <code>class</code> option to display a list of file spaces belonging to a NAS node. | "Query Filespace" on page 588 |
| <code>delete filesystem</code> | Use the delete filesystem command with the <code>class</code> 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 558 |

A NAS file system specification uses the following conventions:

- NAS nodes represent a new node type. The NAS node name uniquely identifies a NAS file server and its data to Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}. Do not use brace delimiters in the option file.

Note: When you initiate a NAS backup operation using the 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 and other necessary tasks before data movement occurs.

Related reference

“Toc” on page 496

Methods for backing up and recovering data on NAS file servers when using CIFS

Tivoli Storage Manager can process network attached storage (NAS) file-server data that is accessed using the Common Internet File System (CIFS).

You can use the following methods to back up and recover data on NAS devices:

- Use a Tivoli Storage Manager backup-archive client to back up and restore data, by using CIFS or NFS to access files from the backup-archive client. Data can be stored on the Tivoli Storage Manager server with file-level granularity, 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 Tivoli Storage Manager client accesses individual files. The method requires that data flow through the Tivoli Storage Manager client. This method also requires that the data flows through the Tivoli Storage Manager server unless a LAN-free configuration is used.
- Use the `snapi ff` option to mitigate the performance problems of CIFS backup. This option stores data with file-level granularity using progressive incremental backup for NFS or CIFS.
- Use a Tivoli Storage Manager backup-archive client running on the NAS device, provided that you can use external programs with the NAS operating system. This method decreases processor usage of CIFS or NFS. Data can be stored on the Tivoli Storage Manager server with file-level granularity using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool. This method requires that data flow through the Tivoli Storage Manager 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 Tivoli Storage Manager. File systems are backed up as full images (all files) or differential images (all files that have changed since the last full backup). Backed up images are stored on a tape device that is accessed by the NAS file server. This method provides high performance because there is no

data flow through a Tivoli Storage Manager client or server. Data that is backed up to Tivoli Storage Manager using NDMP cannot be migrated, reclaimed, or backed up to a copy storage pool.

These are some limitations of NAS file server data when using CIFS:

- Performance degradation occurs because data is being accessed remotely.
- The mapped drives appear to Tivoli Storage Manager as NTFS file systems, but they might not have the full functionality of 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.

Tip: Use NDMP with Tivoli Storage Manager on a NAS file server to back up and restore volumes instead of backing up and restoring the volumes using remote mapped drives.

Related reference

“Snapdiff” on page 463

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 including the IBM TotalStorage NAS 200, 300, and 300G.

You can use the Tivoli Storage Manager backup-archive client to back up the persistent images (PI) of a volume produced by PSM. You must first ensure that the volume has a label. You can then use PSM to schedule or create a persistent image with a specific image name, such as `snapshot.daily`, and set the number of images to save to 1. PSM overwrites the PI as needed and you can use Tivoli Storage Manager to incrementally back up the PI. In this case, Tivoli Storage Manager only backs up the files that changed between snapshots. One advantage of backing up a PSM PI rather than the actual volume, is that there are no open files in the PI.

Consider the following items before using Persistent Storage Manager:

- By default, a PSM schedule uses a variable name (`snapshot.%i`) and keeps a number of images.

Important: Do not use Tivoli Storage Manager with PSM in this way. Tivoli Storage Manager sees each image as unique and makes a complete copy of each.

- Tivoli Storage Manager requires that the volume used to make the PI has a label. If the volume does not have a label, Tivoli Storage Manager will not back up its PI.
- You use the Tivoli Storage Manager image backup function to back up the original volume used to create the PI, but you cannot use the backup image function to back up the PI.
- To avoid backing up unnecessary files when backing up PSM, include the following entries in your client option file (`dsm.opt`):

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

Backing up VMware virtual machine data

You can use the Backup-Archive Client to backup VMware virtual machines.

When the Backup-Archive Client is installed on a Windows computer, it can be used to perform the following backup and restore operations for VMware virtual machines:

- Full VM backup
- Full VM incremental backup (available only with Tivoli Storage Manager for Virtual Environments)
- File-level backup, using the VMware vStorage APIs and the VMware vStorage backup server
- Item recovery (file, folder) from a full VM backup (available only with Tivoli Storage Manager for Virtual Environments Data Protection for VMware Recovery Agent)
- Full VM restore

A full VM backup copies the entire virtual machine, including the virtual disks and the virtual machine configuration file. A full VM incremental backup backs up only the files that have changed since the last full VM backup. Repeatedly creating full VM incremental backups creates large files. Periodically perform a full VM backup to reduce the disk space needed to store the backed up guest.

VMware consolidated backup (VCB) has been deprecated. Use the vStorage APIs to perform full VM backups. The vStorage APIs provide the following benefits:

- Less storage is needed on the vStorage backup server because with the vStorage APIs, a temporary directory to contain the backup is not needed on the vStorage backup server.
- Backups created using the vStorage APIs do not require the VMware converter tool to restore them.

Related reference

“**Backup VM**” on page 549

“Vmmc” on page 509

“**Query VM**” on page 602

“**Restore VM**” on page 631

“Vmchost” on page 511

“Vmcpw” on page 512

“Vmcuser” on page 513

“Domain.vmfile” on page 348

“Domain.vmfull” on page 350

Scenario: Backing up your virtual machines (Windows)

This scenario guides you through an example configuration of VMware guests, Tivoli Storage Manager and the vStorage backup server to protect a VMware environment.

In this scenario, you define a full VMware backup of the guests that runs once a week and a daily, file-level backup of the same VMware guests. This configuration ensures that there are frequent backups of the VMware guests while limiting the required size and time to store them.

After you have configured your VMware backups using this scenario, you can customize it to better match the configuration of your environment.

1. Configure the vStorage backup server. See “Preparing the environment for VMware backup processing” on page 143.
2. Define a full VM backup for each of the VMware guests. See “Running full virtual machine backups” on page 145.
3. Schedule the full VM backup that runs once a week.
4. Optional: If you have multiple VMware guests, separate the pool of VMware guests into groups to reduce backup time using the following steps:

Note: Placing VMware guests in to separate pools reduces backup times because each backup is run in parallel on the vStorage backup server.

- a. Create a group for each VMware ESX server.
 - b. Create a Tivoli Storage Manager node name for each group.
 - c. Create a client options file for each group.
 - d. Create a scheduler service for each group.
5. Define a file-level backup of each of the VMware guests.
 6. Optional: Exclude files that you do not want backed up from file-level backup using the following steps:

Important: If you are running file-level backups from the VMware guests, their exclude list must match the exclude list on the proxy backup server.

- a. From the virtual machine, open the Tivoli Storage Manager client `dsm.opt` file.
- b. Add exclude statements for the directories and files that you are not backing up, using the following command:

```
exclude "\\*\X$\Directory\File"
```

where *X* is the drive letter, *Directory* is the directory name, and *File* is the file name.

For more information about excluding files using the `dsm.opt` file, see *Client include exclude lists* in the *Tivoli Storage Manager Administrator's Guide*.

7. Schedule the file-level backups to run six days per week. The full backup runs on the seventh day.
8. Optional: Create a pre-freeze and post-thaw VMware script for each of the virtual machines using the following steps:

Tip: Use the scripts to stop a database, before the VMware server takes the snapshot, and then starting it again after the snapshot is complete.

- a. Create the script files in the following file name and location:
 - `C:\Windows\pre-freeze-script.bat`
 - `C:\Windows\post-thaw-script.bat`
 - b. Add the VMware commands to run application-consistent backup in virtual machines.
9. Optional: Specify compression and deduplication to reduce the backup size of the VMware backups using the following steps:
 - a. From the VMware guests, open the Tivoli Storage Manager client `dsm.opt` file.
 - b. Enable compression for each group on the proxy server by adding the following option to the group's `dsm.opt` file: `compression yes`.

Tip: You can only use the compression option with the VSTORE full vm backup if the backup is being saved to a storage pool that has been enabled for deduplication. See “Determine compression and encryption processing” on page 78 and “Compression, encryption, and adaptive subfile backup processing” on page 398 for more information.

- c. Enable deduplication by adding the following option to the client `dsm.opt` file: `deduplication yes`.

Preparing the environment for VMware backup processing

Use the following steps to prepare the VMware environment to be backed up. The vStorage backup server can run either a Windows or Linux client.

1. Configure your storage environment so that it can be backed up using 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: If you will be running a direct SAN backup, zone the SAN and configure the disk subsystem host mappings so that all VMware ESX servers and the vStorage backup server access the same disk volumes.
2. Configure the vStorage backup server using the following steps:
 - a. When the backup-archive client runs on a vStorage backup server, this client configuration is referred to as the Tivoli Storage Manager *data mover node* (previously called the backup proxy server). If you configure the data mover node to directly access the storage volumes, turn off automatic drive letter assignment. If it is not turned off, the client on the data mover node might corrupt the Raw Data Mapping (RDM) of the virtual disks and backups will fail.

A Tivoli Storage Manager data mover node typically uses the SAN to back up and restore data. If the Tivoli Storage Manager data mover node is on a Windows Server 2008 or Windows Server 2008 R2 computer, you must set the Windows SAN policy to OnlineAll if you plan to use the SAN to restore data. 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
```

Additionally, if the backup-archive client is installed in a virtual machine on a Windows Server 2008 or Windows Server 2008 R2 computer, the SAN policy on that computer must also be set to OnlineAll if you want to use the hotadd transport to restore data from dynamically added disks.

In either configuration (where the client uses the san or hotadd transport), if the Windows SAN policy is not set to OnlineAll, restore operations fail, and the following message is displayed:

```
ANS9365E VMware vStorage API error.
TSM function name: vddksdk Write
TSM file : vmvddkdisk.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

See the “Vmvstortransport” on page 517 option for a description of the option’s defaults, and how you can override them, if necessary.

- b. On Windows systems: Install the Tivoli Storage Manager client on the vStorage backup server. At the custom setup page of the installation wizard, select **VMware Backup Tools**.

Important: If you are moving the backup data using a LAN-free backup, the SAN must have separate connections for tape and disk.

3. Modify Tivoli Storage Manager using the following steps:
 - a. Access the administrative command line of Tivoli Storage Manager client.
 - b. From the Tivoli Storage Manager client of the vStorage backup server, enter 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 it.

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 Tivoli Storage Manager client of the vStorage backup server, enter the following command to register the node:

```
register node my_guest_name my_password
```

where *my_guest_name* is the full name of the VMware guest that you are backing up.

- d. Repeat step 3c for each VMware guest that you are backing up.
 - e. Grant proxy access for each of the VMware guests by entering the following command:

```
grant proxynode agent=my_server_name target=my_guest_name
```

When the command completes, the following message is returned:

```
GRANT PROXYNODE: success.  
Node my_server_name is granted proxy authority to node my_guest_name.
```

- f. Repeat step 3e for each guest that you are backing up.
4. If you use the **BACKUP VM** command on a virtual machine that has volumes that are mounted to a directory instead of a drive letter, the file might not be stored to the correct location because the mount point does 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:

Tip: The drive letter error is caused because the mount points for a VMware guest running Microsoft Windows do not have a drive letter assignment. When you use the VMware vStorage APIs for Data Protection, it creates a filespace name that includes a number assignment. The filespace names that were created for the mount point do not correspond to the actual mount points of the backed up file.

- a. To restore the backed up files to the original location, map the drive or assign a drive letter to the mount point from the VMware guest 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 reference

“Backup VM” on page 549

“Query VM” on page 602

“Restore VM” on page 631

“Vmchost” on page 511

“Vmcpw” on page 512

“Vmcuser” on page 513

“Vmlist” on page 514

“Vmvstortransport” on page 517

Running full virtual machine backups

A full virtual machine (VM) backup is a backup of an entire virtual machine snapshot. It is like a Tivoli Storage Manager image backup.

The backup environment must be configured as described in “Preparing the environment for VMware backup processing” on page 143. The vStorage backup server must run a Windows or Linux client.

1. Configure the Backup-Archive Client on the vStorage backup server using the following steps:
 - a. From the welcome page of the Tivoli Storage Manager client, 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 Servers or the host name of the Virtual Center. If you specify the Virtual Center host name, 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 specified in the **Host** field.
 - g. Optional: If you want to override the default management class for full VM 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 will be stored.
 - i. Click **OK** to save your changes.
2. Verify that your system is configured correctly by running a backup of one of the virtual machines using the following steps:
 - a. At the command line of the vStorage backup server, enter the command:


```
dsmc backup vm my_vm_name
```

where *my_vm_name* is the name of your virtual machine.
 - b. When the command ends, verify that it completed without errors. The following message indicates that the command completed successfully:

```
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
```

3. Verify that you can restore the virtual machines files using the following steps:
 - a. At the command-line interface of the vStorage backup server, enter the 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 there are any restore processing failures, view the Tivoli Storage Manager error log for more information.

Tip: The log file is saved to `c:\Program Files\Tivoli\TSM\baclient\dsmerror.log`

Related reference

“Backup VM” on page 549

“Vmbackuptype” on page 508

Backing up virtual machines on a Windows Hyper-V system

Use Hyper-V on Windows Server 2008 systems to host virtual machines that allow other operating systems to run independently on one physical workstation.

You can use the Tivoli Storage Manager backup-archive client for these tasks on Hyper-V systems:

- Back up and restore snapshots of Hyper-V virtual machines using the Microsoft Volume Shadow Copy Services (VSS) interface. You can restore a virtual machine using the files created by VSS.
- Restore either individual virtual machines or a group of virtual machines running on a Hyper-V server for disaster recovery purposes.
- Back up data without stopping the virtual machine or stopping any running applications within the virtual machine.

Restriction: You can only back up data this way if the operating system of the virtual machine supports VSS. For further details about this restriction, refer to: [http://technet.microsoft.com/en-us/library/dd252619\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd252619(WS.10).aspx)

- Add backup and restore support for guest operating systems without installing a Tivoli Storage Manager client on the guest VM.

For both disaster recovery and long-term data backup support, perform these steps:

1. Install the Tivoli Storage Manager client on the Hyper-V host Windows 2008 operating system (for example, the parent partition).
2. On the Tivoli Storage Manager Server, define a single node for the backup of the Hyper-V machine.
3. Back up a running virtual machine:
 - From the Tivoli Storage Manager GUI **Backup** panel, 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` on the 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 drives (VHD) are configured to support VSS):
 - The virtual machine that is running pauses
 - A snapshot is taken
 - The virtual machine resumes
 - b. If Microsoft clustering is enabled, you might need to set the `clusterdisksonly` option to **NO** to back up Hyper-V VMs. This action is necessary if any VM files or supporting VM files reside on disks that are not clustered.
4. Restore a virtual machine:
- From the Tivoli Storage Manager GUI **Restore** panel, expand the Hyper-V VMs container, and select the virtual machine to restore.
 - You can also enter: `dsmc restore vm1 -vmlist=vm1 -vmbackuptype=HyperVFull` on the Tivoli Storage Manager command line, where `vm1` is the name of the virtual machine.

Important: When a restore operation is complete, the existing virtual machine is stopped (if it is running) and all existing files (for example, VHD) are deleted. When the backup snapshot is restored, all the virtual machines that existed when the files were backed up are recreated. This includes any Hyper-V snapshots that existed when the files were backed up.

You can set up Tivoli Storage Manager scheduled backups on a regular basis (for example, once a week). Create the schedules as MACRO type schedules with the macro file located in the Tivoli Storage Manager Backup-Archive Client directory of the Hyper-V machine. The contents have the `dsmc` command, for example: `dsmc backup vm -vmbackuptype=HyperVFull`.

Related reference

- “Backup VM” on page 549
- “Query VM” on page 602
- “Restore VM” on page 631
- “Clusterdisksonly” on page 314

Hyper-V backup support limitations

Due to the tight integration of Microsoft Failover Clustering with Cluster Shared Volumes and Hyper-V, the following limitations apply when using Tivoli Storage Manager.

Microsoft Hyper-V support is provided for the following configurations:

- A dedicated Hyper-V system or
- A Hyper-V configuration that is part of Microsoft Failover Clustering or
- A Hyper-V configuration that is part of Microsoft Failover Clustering with Cluster Shared Volumes

VM backups can only be performed by a backup-archive client installed on the cluster node that is the current owner of the VM. To see which VMs are owned by the node, navigate the Windows menu items: **ServerManager->Roles->Hyper-V Manager**.

A VM should be restored only by a backup-archive client installed on the cluster node that is the current owner of the VM. To determine the VM configurations of the cluster, navigate the Windows menu items: **ServerManager > Features > Failover Cluster Manager > cluster name > Services and Applications**. Use the **Services and Applications** window to change which node is the owner of the VM before you restore the VM.

VMs currently owned by different nodes which have files on a common Cluster Shared Volume cannot be backed up simultaneously. The second backup will fail as the shared volume will be owned by the first node. To determine which node owns the volume, navigate the Windows menu items: **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).

If your cluster environment typically has all VMs owned by a single node, and other nodes only provide failover support, the following applies:

- Consider using CLUSTERNODE YES for backup from the active node and restore to failover nodes (if needed).
- Consider running the Tivoli Storage Manager scheduler as a failover service to periodically backup all VMs. For example: `dsmc backup vm * -vmbackuptype=hypervfull`. A failover will start the scheduler on the failover node.

If your cluster environment typically has some VMs owned by one node, and other VMS owned by other nodes, the following applies:

- Consider using the “Asnodename” on page 299 option to allow for backup from multiple nodes and restore to failover nodes (if needed).
- Consider running the Tivoli Storage Manager scheduler as active on each node to periodically backup the VMs owned by each node. The schedules need to 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 which move over on the next backup.

Back up and archive Tivoli Storage Manager FastBack data

Use IBM 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 47

Related reference

“Fbclientname” on page 370

“Fbpolicyname” on page 371

“Fbvolumename” on page 375

Net Appliance CIFS share definition backups

Network Appliance (NetApp) CIFS share definitions include share permissions set on the file server.

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 186

Related reference

“Snapdiff” on page 463

Display backup processing status

During a backup, by default Tivoli Storage Manager displays the status of each file it attempts to back up.

Tivoli Storage Manager reports the 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 24 on page 150 lists some informational messages and meanings.

Table 24. Client command line informational messages

| Informational message | Meaning |
|------------------------------------|--|
| Directory--> | Indicates the directory that you back up. |
| Updating--> | Indicates that only the file meta data is sent, not the data itself. |
| Expiring--> | Indicates an object (file or directory) on the server that no longer exists on the client is expired and made inactive on the server. |
| Subfile--> | Indicates that a file is processed using adaptive subfile backup. The file size indicated in this message is the size of the entire file, not the amount of data actually sent. |
| Total number of objects inspected: | As indicated. When using journal-based backup (AIX and Windows only), the number of objects inspected might be less than the number of objects backed up. When using snapshot difference incremental backup, the number of objects inspected is zero. The number is zero because the client performs an incremental backup of the files that NetApp reported as changed. The client does not scan the volume looking for files that have changed. |
| Total number of objects backed up: | As indicated. |
| 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 239 for more information. |
| Total number of objects deleted: | This is a count of the objects deleted from the client workstation after being successfully archived on the server. The count is zero for all backup commands. |
| Total number of objects expired: | See the section about full and partial incremental backup for more information. |
| Total number of objects failed: | Objects can fail for several reasons. Check the dsmdir.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. |
| 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. |
| 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 dsmc.exe INCREMENTAL C:\Users* -su=yes |
| 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 24. 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: | The average rate at which Tivoli Storage Manager and the network transfer data between the client and the server. This is calculated by dividing the total number of bytes transferred by the time that elapses from the beginning to the end of the process. Both Tivoli Storage Manager processing and network time are included in the aggregate transfer rate. Therefore, the aggregate transfer rate is lower than the network transfer rate. Note: On occasion, the aggregate data transfer rate might be higher than the network data transfer rate. This is because the backup-archive client can have multiple simultaneous sessions with the backup server. If you set the resourceutilization option, Tivoli Storage Manager attempts to improve performance and load balancing by using multiple sessions when it backs up a volume or other set of files. When multiple sessions are open during backup, the data transfer time represents the sum of the times reported by all sessions. In this case, aggregate data transfer time is incorrectly reported as higher. However, when running with a single session, the aggregate data transfer rate should always be reported as lower than the network data transfer rate. |
| 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: $== (1 - (10240/102400)) \times 100 == 90\%$. |
| 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: | 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: $(1 - (25 \text{ MB}/100 \text{ MB})) \times 100 == 75\%$. 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. |
| 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. |
| Total objects deduplicated: | The number of files that have been deduplicated. |

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.
 - Tivoli Storage Manager can back up the file even if it is open and changes during the backup. This is only useful if the file 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 Tivoli Storage Manager attempts to back up open files, this is not always possible. Some files are open exclusively for the application that opened them. If Tivoli Storage Manager encounters such a file, it cannot read it for backup purposes. If you are aware of such file types in your environment, you should exclude them from backup to avoid seeing error messages in the log file.

Related concepts

“Display information about management classes and copy groups” on page 233

“Select a management class for files” on page 237

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 Tivoli Storage Manager 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 231

Related tasks

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

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 inactivate 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 231

Removable media backup

Tivoli Storage Manager backs up your removable media (such as tapes, cartridges or diskettes) based on the drive label, not the drive letter.

If a drive has no label, the backup does not occur. This use of drive labels permits you to perform such tasks as backing up different diskettes from the a: drive.

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

Because Tivoli Storage Manager uses the labels to manage backups and archives of your removable media, you occasionally need to use those labels to locate data when using commands. For example, if you try to restore a file on diskette or 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, Tivoli Storage Manager can locate your files if you use the file space name based on the original drive label. A mismatch between a label and a file space name might occur if you label your drives again, or if you access Tivoli Storage Manager from a different workstation than the one from which you backed up the files. If you have not relabeled the drive, and you are at the same workstation where the file was backed up, then you can use the drive letter as a shorthand version of the file space name (drive label).

Fixed drives

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

NTFS file spaces

When you back up files on New Technology File System (NTFS) partitions, Tivoli Storage Manager 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 an NTF file-space name that is mixed or lowercase, 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.

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: Universal Naming Convention backup

You can back up shared files in a network through the use of a UNC name. This topic shows some examples of backing up UNC name files.

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 25. 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 69

Related reference

“Domain” on page 343

Microsoft Dfs file protection methods

There are some methods that you can use to protect the data in your Microsoft Dfs environment.

Here are the methods you should use to protect your Microsoft Dfs data:

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 Tivoli Storage Manager back up and restores to be coordinated among several workstations.

Note:

1. This feature applies to Windows Server 2003.
2. 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 Tivoli Storage Manager can be used to restore Dfs links.

You can use the `dfsbackupmntpnt` option to specify whether Tivoli Storage Manager sees a Dfs mount point as a Microsoft Dfs junction or as a directory.

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, Tivoli Storage Manager creates a directory under the Dfs root using the same name as the junction point and restores the data in that directory.

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

Dfs root

```
\\almaden\snj64test
```

Dfs link1

\\almaden\snj64test\tools

Dfs link2

\\almaden\snj64test\trees

Backup procedure:

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

```
dsmc inc \\almaden\snj64test
```

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

```
dsmc inc \\almaden\snj64test\tools
```

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

```
dsmc inc \\almaden\snj64test\trees
```

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

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

Restore procedure:

1. Manually recreate shares at target 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 \\almaden\snj64test\tools\* -sub=yes
```

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

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

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

Tivoli Storage Manager limitations:

- Tivoli Storage Manager does not restore root of Dfs. To recreate the Dfs tree, manually create the Dfs root first, then start restore to recreate the links.
- Tivoli Storage Manager can back up the Dfs tree (both domain based Dfs and stand alone Dfs) hosted on local workstation only. You cannot back up Dfs if the Dfs host server is not your local workstation.
- Tivoli Storage Manager cannot recreate shared folders on restore. For example, if you delete the junction and the shared folder the junction points to, restoring the Dfs root 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, then Tivoli Storage Manager does not display the replica data.
- If a Dfs root is added or modified, Tivoli Storage Manager will not back it up. You must specify the Dfs root in the `domain` option in the client options file (`dsm.opt`) regardless of whether `DOMAIN ALL-LOCAL` is specified.

Related reference

“Dfsbackupmtpnt” on page 336

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 you want to restore using file specification (file path, name, and extension), a directory list, or a subdirectory path to a directory and its subdirectories.

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

- Estimate
- Preferences editor

Note: When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.

The following are the primary restore tasks:

- “Restoring files and directories” on page 163
- “Restoring Windows system state” on page 167
- “Restoring Automated System Recovery files” on page 169
- “Microsoft Dfs tree and file restore” on page 177
- “Restoring an image” on page 177
- “Restore data from a backup set” on page 180
- “Restoring data to a point in time” on page 199
- “Restore NAS file systems” on page 200
- “Authorizing another user to restore or retrieve your files” on page 195
- “Restoring or retrieving files from another client node” on page 196
- “Restoring or retrieving your files to another workstation” on page 197
- “Deleting file spaces” on page 198
- “Restoring data from a VMware backup” on page 186

Related tasks

“Starting a web client session” on page 91

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 *abcdefghijk.doc* might originally have existed in the directory as *abcdefghijk.doc* and *abcdef~2.doc*. During the restore, if *abcdefghijk.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 439

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 391

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 using the GUI

This section lists the steps to follow to restore files and directories using the client GUI.

1. Click **Restore** on the main window. The Restore window appears.
2. Expand the directory tree by clicking the plus (+) sign or the folder icon next to an object in the tree. Select the object that you want to restore. To search or filter files, click the **Search** icon from the tool bar.
3. To modify specific restore options, click the **Options** button. Any options you change are effective during the current session *only*.
4. Click **Restore**. The Restore Destination window appears. Enter the appropriate information.
5. Click **Restore**. The Restore Task List window displays the processing status.

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 26. Command-line restore examples

| Task | Command | Considerations |
|--|---|---|
| Restore the most recent backup version of the <code>c:\doc\h1.doc</code> 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 <code>latest</code> 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 407 for more information. |

Table 26. Command-line restore examples (continued)

| Task | Command | Considerations |
|---|---|--|
| 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 424 and "Inactive" on page 391 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 <code>subdir</code> option to restore directory attributes/permissions. See "Subdir" on page 482 for more information about the <code>subdir</code> 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 425 and "Pittime" on page 426 for more information about the <code>pitdate</code> and <code>pittime</code> options. |
| Restore the c:\doc\h2.doc file to its original directory on the workstation, named <i>star</i> . | <p><code>dsmc restore c:\doc\h2.doc \\star\c\$\</code></p> <p>To restore the file to <i>star</i> which has been renamed <i>meteor</i>, 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 (<i>meteor</i>, 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. |
| Restore a file that was originally backed up from the diskette 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 377 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 622 for more information. |

Related concepts

Chapter 13, “Using commands,” on page 521

Related reference

“Restore” on page 605

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.

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:

The no query restore process is outlined below.

1. The client tells the server that a no query restore is going to be performed and provides the server with details about file spaces, directories, and files.
2. The server sorts the data using an internal sort table which minimizes tape mounts.
3. The data to be restored is sent to the client. File and directory objects stored on disk are sent immediately since sorting for such data is not required before restoring it.
4. You can use multiple sessions to restore the data. If the data resides on multiple tapes, there are multiple mount points available at the server. The combination of using the `resourceutilization` option and `MAXNUMMP` allows multiple sessions.

When you enter an unrestricted wildcard source file specification on the **restore** command and do not specify any of the options: `inactive`, `latest`, `pick`, `fromdate`, `todate`, or `volinformation`, the client uses a *no query restore* method for restoring files and directories from the server. This method is called *no query restore* because instead of querying the server for each object to be restored, a single restore request is sent to the server. In this case, the server returns the files and directories to the client without further action by the client. The client merely accepts the data coming from the server and restores it to the destination named on the **restore** command.

Using the command-line client, an example of an unrestricted wildcard command would be:

```
c:\mydocs\2004\*
```

An example of a restricted wildcard file specification would be:

```
c:\mydocs\2004\sales.*
```

Restartable restore process:

If the restore process stops because of a power outage or network failure, the server records the point at which this occurred.

This record is known to the client as a *restartable restore*. It is possible to have more than one restartable restore session. Use the **query restore** command 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 442

“Restore” on page 605

Plan for temporary space to restore migrated files

When you restore migrated files, you also need temporary space for the resident version of these files.

If you do not specify enough space for the migrated files and the resident versions of the files combined, the operation might fail. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about the **restoremigstate** option.

Restoring Windows system state

Tivoli Storage Manager supports the Microsoft Volume Shadowcopy Service (VSS) on Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7. Tivoli Storage Manager uses VSS to restore the system state.

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. On Windows Server 2003, most of the system

state data is SFP files, so a lot of free space is required. If there is absolutely no extra space on the system drive, contact your service representative.

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.

On Windows Vista, Windows Server 2008, and Windows 7, you must be a member of the Administrators group to restore Windows system state information. On Windows Server 2003, you must be a member of the Administrators or Backup Operators group to restore Windows system state information.

On Windows Vista, Windows 7, and Windows Server 2008 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:

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.

On the command line, use the **restore systemstate** command to restore a backup of system state.

Considerations:

- If you are upgrading from a Windows XP workstation to a Windows Vista workstation, you cannot restore the Windows XP system objects that were backed up to the server.
- Your Windows client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
- If Active Directory is installed, you must be in Active Directory restore mode.
- See 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 operating system is not working (Windows Server 2003)” on page 173

“Recovering a computer when the operating system is not working (Windows Vista, Windows 7, and Windows Server 2008)” on page 170

Related tasks

“Restoring the operating system when the computer is working” on page 170

Related reference

“Restore Systemstate” on page 629

Restoring Automated System Recovery files

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.

You must be a member of the Administrators or Backup Operators group to back up ASR files.

On Windows Vista, Windows 7, and Windows Server 2008, the Tivoli Storage Manager backup-archive client restores ASR data when the backup-archive client restores the Windows system state. Use the **restore systemstate** command to restore ASR files on these operating systems. On Windows XP and Windows Server 2003, restoring the ASR files is a separate task from restoring system state.

To restore the ASR files for Windows XP and Windows Server 2003, use the Tivoli Storage Manager GUI as follows.

1. Click **Restore** in the GUI main window. The Restore window appears.
2. Expand the directory tree by clicking the plus sign + or the folder icon next to an object in the tree.
3. Click the selection box next to the Automated System Recovery node or click on the ASR selection box from the file list.

When you select Automated System Recovery from the directory tree, the results are displayed in the file list. Any inactive version of a particular file is indicated by the icon marked with an X next to the selection box. If you restore inactive ASR files in ASR recovery mode, you must ensure that the system state is restored to the same point in time as the ASR files. This requires that you must manually modify the `tmsmr.cmd` file because the default command parameters in the `tmsmr.cmd` file do restore the system to the last active backup. The restore commands in the `tmsmr.cmd` file do not specify inactive point-in-time parameters.

4. Click **Restore**. Enter values for the fields. You can choose to restore the ASR files to their original destination on the system drive (\adm.sys\ASR) or to a different location, such as portable media.
5. Click **Restore**. The Restore Task List window displays the processing status.

Related concepts

"Recovering a computer when the operating system is not working (Windows Server 2003)" on page 173

"Recovering a computer when the operating system is not working (Windows Vista, Windows 7, and Windows Server 2008)"

Related reference

"Restore ASR" on page 614

Restoring the operating system when the computer is working

If your computer is working, you can restore the operating system from backed up files.

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.

1. Restore the system drive. For example: `dsmc restore c:* -sub=yes -rep=all`.
2. Restore system state. For example: `dsmc restore systemstate`.

Related concepts

"Recovering a computer when the operating system is not working (Windows Server 2003)" on page 173

Recovering a computer when the operating system is not working (Windows Vista, Windows 7, and Windows Server 2008)

If the computer has a catastrophic hardware or software failure, you can recover a Windows Vista, Windows 7, or Windows Server 2008 operating system with Automated System Recovery (ASR).

Note that ASR data is not backed up, and hence cannot be restored, on 64-bit architectures with Extensible Firmware Interface (EFI).

Related concepts

"Recovering a computer when the operating system is not working (Windows Server 2003)" on page 173

Related tasks

"Restoring the operating system when the computer is working"

Creating a bootable WinPE CD (Windows Vista, Windows 7, and Windows Server 2008)

Before you can recover a Windows Vista, Windows 7, and Windows Server 2008 computer using ASR (Automated System Recovery), you must create a bootable Windows Preinstallation Environment (WinPE) CD or DVD.

You must create a technician computer as described in Microsoft Windows Automated Installation Kit (AIK) documentation. Use the latest version of AIK.

Create a bootable WinPE CD any time that there is a hardware configuration change or service pack update to your workstation.

The bootable WinPE CD contains a Tivoli Storage Manager client installation package, and a minimum operating system.

In this task, you create a generic bootable WinPE CD. Use a generic bootable WinPE CD to recover all computers of the same operating system and architecture. The generic bootable WinPE CD does not contain the Tivoli Storage Manager options file (dsm.opt), because the options file is unique for each computer. You must provide the Tivoli Storage Manager options file for each computer on separate media.

1. Prepare the minimal operating system.
 - a. On the technician computer, click **Start > All Programs**, and as appropriate for your operating system, click **Windows OPK** or **Windows AIK**. Right-click **Deployment Tools Command Prompt**, and then click **Run as administrator**. By default, all tools are installed at `C:\Program Files\kit name\Tools`, where *kit name* is **Windows OPK** or **Windows AIK**.
 - b. At the command prompt, run the `Copype.cmd` script. The script prepares the minimum operating system. The script requires two arguments: hardware architecture and destination location. The hardware architecture value can be `x86`, or `amd64`. The destination value is the path to a local folder. For example:

For x86 architecture:
`copype.cmd x86 c:\winpe_x86`

For x64 architecture:
`copype.cmd amd64 c:\winpe_x64`
 - c. Copy the base image (`Winpe.wim`) to the `\Winpe_x86\ISO\sources` folder and rename the file to `Boot.wim`. For example:

For x86 architecture:
`copy c:\winpe_x86\winpe.wim c:\winpe_x86\ISO\sources\boot.wim`

For x64 architecture:
`copy c:\winpe_x64\winpe.wim c:\winpe_x64\ISO\sources\boot.wim`
2. Configure the Tivoli Storage Manager client installation package.
 - a. Unpack the Tivoli Storage Manager client installer files into the folder of your choice.
 - 1) Run the installer executable file.
 - 2) Indicate where to save the files. The default directory is `c:\tsm_images`.
 - 3) When prompted to choose the setup language, cancel the wizard. The files are unpacked but the client is not installed.
 - b. In a command prompt window in the directory where the client installer is unpacked, run the Tivoli Storage Manager preparation script **`dsmc-pe.cmd`** or **`dsmc-pe-x64.cmd`**. If the bootable WinPE CD is for a 32-bit computer, run **`dsmc-pe.cmd -set`**. If the bootable WinPE CD is for a 64-bit computer, run **`dsmc-pe-x64.cmd -set`**. The script prepares the Tivoli Storage Manager client to run in the WinPE environment. For example, open a command prompt window and run the command `inc:\tsm_images\TSM_BA_Client\`.
 - c. Copy Tivoli Storage Manager files to the WinPE location. For example:
`xcopy c:\tsm_images\TSM_BA_Client* c:\winpe_x64\iso\ /e /y`
3. Create the WinPE ISO image. At a command prompt, enter this command:

```
oscdimg -n -o -bC:\winpe_x86\etfsboot.com
C:\winpe_x86\ISO C:\winpe_x86\winpe_x86.iso
```

4. Burn the ISO image to a CD or DVD.

The bootable WinPE CD contains a minimum operating system and Tivoli Storage Manager client. Use this CD or DVD when restoring a computer with ASR recovery.

Restoring the operating system with Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008)

You can restore the operating system of a Windows Vista, Windows 7, or Windows Server 2008 computer with Automated System Recovery (ASR).

Ensure that you have the bootable WinPE CD that contains a Tivoli Storage Manager client. Ensure that you have portable media with a Tivoli Storage Manager options file (dsm.opt) that is configured for the computer to be recovered.

1. Start the computer with the bootable WinPE CD or DVD.
2. From a command prompt, set environment variables DSM_DIR and DSM_CONFIG. Set DSM_CONFIG to the location of the Tivoli Storage Manager client options file (dsm.opt). Set DSM_DIR to the location of the Tivoli Storage Manager client on the bootable WinPE CD or DVD. For example:

For x86 architecture:

```
SET DSM_CONFIG=e:\dsm.opt
SET DSM_DIR="x:\PROGRAM files\Tivoli\TSM\baclient"
```

For x64 architecture:

```
SET DSM_CONFIG=e:\dsm.opt
SET DSM_DIR="x:\PROGRAM files 64\Tivoli\TSM\baclient"

SET DSM_CONFIG=e:\dsm.opt
SET DSM_DIR="x:\TSM_ba_client_6.2.2\PROGRAM files\Tivoli\TSM\baclient"
```

3. Restore the system state with the Tivoli Storage Manager client.

For example, to restore the system state from a Tivoli Storage Manager server that is identified in the options file, issue a command like the following one:

```
dsmc restore systemstate
```

In WinPE mode, the default is asrmode=yes.

The backup copies on the Tivoli Storage Manager server require password authentication. You can provide the password in several ways:

- a. Before issuing the **restore systemstate** command, import Tivoli Storage Manager registry entries that you saved from the computer before it failed. Saving registry entries is an optional step in the task to create a Tivoli Storage Manager client options file for ASR recovery.
- b. Provide the password in the dsm.opt options file with the password option.
- c. Enter the password when the client connects with the Tivoli Storage Manager server.

If you do not have the password, the Tivoli Storage Manager administrator can change the password for your node.

If you restore the system state from a backup set on a CD or DVD, issue a command like the following one:

```
dsmc restore backupset systemstate
-backupsetname=c:\recovery\backupsetfile.name -loc=file
```

In WinPE mode, the default is `asrmode=yes`.

You can access the backup copies on the backup set without password authentication. After the system state restore is completed, Tivoli Storage Manager client creates and formats critical partitions for the computer. A *critical partition* is the partition that is critical for the operating system to function. For example, if the operating system is installed on the C: drive, then the partition containing the C: drive is critical.

4. Mount and restore the system drive. In the Windows Preinstallation Environment (WinPE), the Tivoli Storage Manager client restores ASR data by default.
 - a. Mount the system drive with the system **MOUNTVOL** command. For example, if C:\ is a mount point and the globally unique identifier (GUID) of the system drive is \\?\Volume{8d587604-54e6-11df-803c-0014a4db5a82}\, enter a command like the following one:

```
MOUNTVOL C:\ \\?\Volume{8d587604-54e6-11df-803c-0014a4db5a82}\
```
 - b. Restore the files on the system drive. For example, if the system drive is C:\, issue a command like the following one:

```
dsmc restore \\computer_name\c$ \* c:\ -su=yes -repl=yes
```
5. Copy C:\Windows\BOOT\PCAT\bootmgr file to a hidden partition. Specify the destination with the GUID of the hidden partition. For example, if the GUID of the hidden partition is \\?\Volume{8d587603-54e6-11df-803c-0014a4db5a82}\, issue this command:

```
copy C:\Windows\BOOT\PCAT\bootmgr \\?\Volume{8d587603-54e6-11df-803c-0014a4db5a82}\
```

In Windows Vista, and in subsequent releases of Microsoft Windows, the operating system installation creates a hidden bootable partition to store boot data. The hidden partition is different from the partition for the system drive. The hidden partition is restored when ASR data is restored during the system state restore. Copy the bootmgr file to the hidden partition. The hidden partition does not have a drive letter mapping, so you must use GUID format for the destination drive in the copy command. You can distinguish the hidden partition from other partitions because it has a size of 100 MB.

6. Restart the computer to complete the recovery process.

You can now restore other volumes.

Related tasks

“Creating a bootable WinPE CD (Windows Vista, Windows 7, and Windows Server 2008)” on page 170

“Creating a Tivoli Storage Manager options file for Automated System Recovery (Windows Vista, Windows 7, or Windows Server 2008)” on page 125

Related reference

“Restore” on page 605

“Restore Systemstate” on page 629

Recovering a computer when the operating system is not working (Windows Server 2003)

If the computer has a catastrophic hardware or software failure, you can recover a Windows Server 2003 computer with Automated System Recovery (ASR).

The functions described in this section are also documented in the Tivoli field guide, *Using Microsoft Windows Automated System Recovery (ASR) to Recover*

Windows XP and Windows 2003 Systems with the IBM Tivoli Storage Manager Backup-Archive Client for Windows, available at the following Web site:
http://www.ibm.com/software/sysmgmt/products/support/Field_Guides.html

Related concepts

“Recovering a computer when the operating system is not working (Windows Vista, Windows 7, and Windows Server 2008)” on page 170

Related tasks

“Restoring the operating system when the computer is working” on page 170

Hardware and software requirements for ASR recovery (Windows Server 2003)

Restoring your Windows Server 2003 system in ASR recovery mode requires specific hardware and software.

Hardware requirements for Automated System Recovery (Windows Server 2003)

Specific hardware is required for Windows Server 2003 Automated System Recovery (ASR).

- The hardware configuration of the target system must be identical to that of the original system, with the exception of the hard disks, video cards, and network interface cards.

Note: If there is a change in the video cards or networking cards, they must be manually configured.

- The target system must have the same number of disks as the original system.
- The size of each disk on the target system must be greater than or equal to the size of the corresponding disk on the original system. Note that disks that appear to be the same size might not have the same number of sectors due to geometry differences. Choose a disk on the target system which has identical disk geometry as the original system or choose a larger disk on the target system.
- All disks must have 512 byte sectors.
- All disks (and only such disks) that are available to ASR on which to perform a recovery must be accessible.
- The ASR files must be accessible from a local floppy drive.
- The floppy and DVD drives cannot be external PC-card drives.
- The original system must be an x86 workstation.
- The disks on the original system must have 512 byte sectors.
- All basic volumes in the original system must be simple. Non-simple volumes are not supported. If a volume is mirrored, it must be a dynamic volume, not a basic volume.

Software requirements for Automated System Recovery (Windows Server 2003)

Specific software is required for Windows Server 2003 Automated System Recovery (ASR).

System software

The operating system installation DVD is required. Network starting of the operating system is not supported for ASR.

ASR diskette

The ASR diskette is created from the ASR files. The label must be TSMASR.

Tivoli Storage Manager backup-archive client installation CD

The Tivoli Storage Manager client installation package must be at the same or higher level than the client that was used to create the backups for the system being recovered.

Boot drive and system drive backup copies

You must have a complete backup of your boot drive and system drive. The backup must contain NTDLL.ASR and SMSS.ASR in the WINDOWS\REPAIR directory. These files are generated by the **backup asr** command before backing up the boot drive and system drive.

Tivoli Storage Manager password

If you restore boot drive and system drive backup copies from a Tivoli Storage Manager server, you are prompted for your Tivoli Storage Manager user ID and password.

Your Tivoli Storage Manager node password is requested during ASR, even if you set the passwordaccess option to *generate*. During ASR, the generated Tivoli Storage Manager password is unavailable in the Windows registry of the temporary operating system image installed by Windows. If you do not know your Tivoli Storage Manager node password, ask your Tivoli Storage Manager administrator to update the node password and tell you the new password. After ASR, you can use the **set password** command to change the password to another value.

If you restore backup copies from a backup set on local media, you are not prompted for your Tivoli Storage Manager user ID and password.

Creating a Tivoli Storage Manager client installation CD for ASR recovery (Windows Server 2003)

To recover your Windows Server 2003 system using ASR (Automated System Recovery), you must create a special Tivoli Storage Manager client installation package.

Prepare the Tivoli Storage Manager client installation package in package-for-the-web format burned onto a CD. This is the single self-extracting installation package that IBM makes available with FTP. The volume label of the CD must be TSMCLI. The client package must be renamed to TSMCLI.EXE and placed in the root directory of the CD. Follow these steps to prepare the client package:

1. Download the Tivoli Storage Manager client package.
2. Rename the package to TSMCLI.EXE. Ensure that the TSMCLI.EXE file path is designated as the root directory (\) in your CD burning software.
3. Using the method described by your CD burning software, specify TSMCLI as the volume label.
4. Burn the CD using the method described by your CD burning software.

Restoring the operating system with Automated System Recovery (Windows Server 2003)

If the computer is not working, you can recover a Windows Server 2003 computer with Automated System Recovery (ASR).

1. Insert the Windows Server 2003 operating system installation CD into the CD drive.
2. Restart your computer. Press the appropriate key to start from the CD. You might need to configure your computer BIOS to enable start from the CD.
3. Press **F6** if you need to install any vendor-acquired SCSI or RAID drivers.
4. Press **F2** to enter the ASR recovery mode during text mode section of Setup.
5. Insert the ASR recovery diskette (TSMASR) into the diskette drive when prompted for the Windows ASR recovery diskette. The following tasks are performed automatically: (1) Windows reads the `asrnpn.sif` file from the diskette, and (2) Windows reformats the boot volume and might reformat other partitions on the disk if the partition layout is different than the original system. This partitioning and format process can cause several restarts.
6. Insert the Tivoli Storage Manager installation CD (TSMCLI) into the CD drive when prompted. The Tivoli Storage Manager client package (TSMCLI.EXE) is copied to the `%SystemRoot%\temp` directory.
7. Insert the ASR recovery diskette (TSMASR) again when prompted. The `tmasr.opt`, `tmasr.cmd`, and `waitforevent.exe` files are copied to the `%SystemRoot%\temp` directory.
8. Remove the diskette when prompted, prior to system restart. The system restarts into GUI-mode setup.
9. Insert the operating system CD into the drive when prompted. The prompt message is unclear; it reads as follows:

Setup was unable [to] access files needed to continue. This could be caused by an error on the CD Media or the CD is no longer present in the drive.

After setup is completed, the `tmasr.cmd` file runs. The Tivoli Storage Manager client is silently installed.

You are prompted to choose whether you want to restore from a network connected Tivoli Storage Manager server (press 1), a local backup set from file (press 2), or a local backup set from tape (press 3).
10. Choose the location of your backup copies.

Tivoli Storage Manager prompts for your ID and password if you are restoring from the Tivoli Storage Manager server. If you are restoring from the local backup set, you are prompted for the path to the local backup set, for example `r:\backupsetfile1.ost`. You are prompted for subsequent volumes if your backup set occupies more than one volume.

Tivoli Storage Manager commands are issued to restore the boot drive and system drive.

Tivoli Storage Manager commands are issued to restore system state.

You modified **TSMASR.CMD** on the ASR recovery diskette so that **NTBACKUP** is called after restoring system state. The recovery continues with the Windows backup and restore wizard.

After the system state is restored, the following message is displayed: Press any key to continue. . . . Do not press any key to continue, but wait for the Windows backup and restore wizard to start. Complete the wizard to restore the user profiles.

You can now restore other volumes.

Restoring user profiles during ASR recovery (Windows Server 2003)

Recover user profiles during a Windows Server 2003 system recovery.

You created the Automated System Recovery (ASR) recovery diskette. You inserted the **NTBACKUP** command in the **TSMASR.CMD** command script on the ASR recovery diskette. You saved user profile data in `c:\ntbackup.bkf` and backed up `c:\ntbackup.bkf`.

This task occurs during ASR recovery, after system state is restored. The following steps complete the Windows backup and restore wizard to restore user profiles.

1. In the welcome panel, click **Advanced Mode** and click **Next**.
2. Click **Restore Wizard (Advanced)**, then click **Next**.
3. In the file tree, select `c:\ntbackup.bkf` and click **Next**. `ntbackup.bkf` contains the user profiles data.
4. In the Completing the Restore Wizard panel, click **Advanced** to specify additional options.
5. In the restore destination panel select **Original location**.
6. In the restore method panel select **Replace existing files**.
7. In the Advanced Restore Options panel, leave the default settings (all options are set).
8. In the Completing the Restore Wizard panel, click **Finish**. The user profiles are restored.
9. When the restore completes, close the NTBACKUP application.
10. In the command prompt window with the message Press any key to continue..., press a key. ASR continues. The command file finishes and the computer restarts.
11. Remove the ASR recovery diskette. The system, including user profiles, is restored.

You can now restore other volumes.

Related tasks

“Creating an Automated System Recovery diskette (Windows Server 2003)” on page 127

“Backing up user profiles for ASR recovery (Windows Server 2003)” on page 128

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 157

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, Tivoli Storage Manager 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 390

“Verifyimage” on page 504

Restoring an image using the GUI

You can use the GUI to restore an image of your file system or raw logical volume.

Follow these steps to restore an image of your file system or raw logical volume:

1. Click **Restore** from the main window. The Restore window appears.
2. Expand the directory tree.
3. Locate the object in the tree named **Image** and expand it. Click the selection box next to the image you want to restore. You can obtain detailed information about the object by highlighting the object and selecting **View** → **File Details...** from the main window or click the **View File details** button.
4. **(Optional)** To perform an incremental image restore, click the **Options** button to open the Restore Options window and select the **Image plus incremental directories and files** option. If you want to delete inactive files from your local file system, select the **Delete inactive files from local** check box. Click the **OK** button.
5. Click **Restore**. The Restore Destination window appears. The image can be restored to the volume with the drive letter or mount point from which it was originally backed up. Alternatively, a different volume can be chosen for the restore location.
6. Click the **Restore** button to begin the restore. The **Task List** window appears showing the progress of the restore. The Restore Report window displays a detailed status report.

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 390

“Verifyimage” on page 504

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 27. 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. |

Table 27. Backup set GUI restore restrictions (continued)

| Data type in the backup set | Backup set location | | |
|-----------------------------|--|---|---|
| | Local (location=file or location=tape) | Tivoli Storage Manager Server (TOC available) | Tivoli Storage Manager Server (TOC not available) |
| system state | 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 28. 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 408
- “Query Backupset” on page 585
- “Query Image” on page 591
- “Restore” on page 605
- “Restore Backupset” on page 615
- “Restore Image” on page 624
- “Restore Systemstate” on page 629

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 Tivoli Storage Manager 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 information

“Localbackupset” on page 408

“Restore” on page 605

“Restore Image” on page 624

“Restore Backupset” on page 615

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.

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 commands below to restore data from backup sets. Use the appropriate command for the location of the backup set and the data in the backup set. (See the table in the parent topic)

restore
restore backupset
restore image

restore systemobject

restore systemstate

Related reference

“Query Backupset” on page 585

“Query Image” on page 591

“Restore” on page 605

“Restore Backupset” on page 615

“Restore Image” on page 624

“Restore Systemstate” on page 629

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

“Net Appliance CIFS share definition backups” on page 149

Restoring data from a VMware backup

There are several options for restoring files to a VMware guest. The method that you use depends on the version of Tivoli Storage Manager client software you are using to run the restore and the way the VMware files were backed up.

File-level restore

A file level restore is useful for restoring an individual file that might have been lost or damaged. Use this option when it is not practical to restore an

entire VMware image. As best practice, do not restore entire VMware guests if you used file level-backups because the restore command does not recreate the Windows system states.

You can only use the file-level restore option if you have previously performed a file-level backup of the VMware guests. You cannot use file-level restore to restore individual files from a full VM backup of a VMware guest. The files are backed up like there was a Tivoli Storage Manager client installed on each VMware guest.

You can restore a VMware file that has been backed up using the following methods:

- Restore the files to the vStorage backup server that backed up the VMware guest. Use this method if you have not installed a Tivoli Storage Manager client on the VMware guest you are restoring the files to.
- Restore the files from a Tivoli Storage Manager client that was installed on the VMware guest.

Full VMware backup

A full VMware backup is used to restore an entire VMware guest. For this method, the VMware guest is replaced with the restored image or a new VMware guest is created with that restored image. The full VMware restore is useful because you are restoring all of the VMware files and the Windows system states. If you have Tivoli Storage Manager for Virtual Environments Data Protection for VMware Recovery Agent, you can restore individual files.

You can restore a Full VMware backup using the following methods:

- If the VMware client is using Tivoli Storage Manager versions earlier than 6.2.2, you must restore the full VMware backup using VMware Consolidated Backup.
- If the VMware client is using Tivoli Storage Manager version 6.2.2 or higher, you must restore the full VMware backup using the vStorage API.

Note: Tivoli Storage Manager version 6.2.2 or higher can restore full VMware backups that were made with the pre-6.2.2 version of Tivoli Storage Manager.

Related tasks

“Restoring file-level VM backups”

“Restoring full VM backups for VMware Consolidated Backup (VCB)” on page 188

“Restoring full VM backups” on page 190

Restoring file-level VM backups

A file level restore is useful for restoring an individual file that might have been lost or damaged.

Consider the following prerequisites before you start this procedure:

- You can only use the file-level restore option if you have setup a file-level backup of the VMware guests. You cannot use this procedure to restore files that were part of a full backup. To restore a full VM backup, see “Restoring full VM backups for VMware Consolidated Backup (VCB)” on page 188 or “Restoring full VM backups” on page 190

- If you are restoring a file from the off-host backup server onto the VMware guest, the off-host backup server must be configured as a proxy. See “Preparing the environment for VMware backup processing” on page 143 for more information.

Important: Restoring a file-level backup to a VMware guest uses the same command syntax that you use to restore a file to a Tivoli Storage Manager client. Do not use the restore VM command for this scenario because it is specific to full-VM restores.

1. From the off-host backup server, query the Tivoli Storage Manager server for available backups by entering the following command:

```
dsmc q MyFilePath -asnode=myVMName
```

where *MyFilePath* is the location of the backups.

2. Restore the full VMware backup that was listed in step 1 by entering one of the following commands:

- If you are restoring with a backup-archive client on the VMware guest from which the file was backed up, enter the following command:

```
dsmc restore myFileName
```

- If you are restoring with a backup-archive client on the off-host backup server (the backup proxy server), you can restore the file to off-host backup server. Then copy the file to the VMware guest. To restore the file to the off-host backup server, enter the following command:

```
dsmc restore myFileName fileLocation -asnode=myVMName
```

See the following parameters for the values that are specific to your environment:

myFileName

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.

fileLocation

Specifies the path and file name where the file is stored on the off-host backup server.

myVMName

Specifies the Tivoli Storage Manager node name of the VMware guest from which the file was backed up. The off-host backup server restores the file to the off-host backup server as a proxy for the VMware guest node.

Related concepts

“Restoring data from a VMware backup” on page 186

Related tasks

“Restoring full VM backups for VMware Consolidated Backup (VCB)”

“Restoring full VM backups” on page 190

Restoring full VM backups for VMware Consolidated Backup (VCB)

You can restore a full VMware backup to recreate all of the files of a VMware guest. Use this procedure to restore full VM backups that were created by older (Tivoli Storage Manager client 6.2.0 or earlier) clients, using VMware Consolidated Backup.

To restore a full VMware backup using this version of the client, see “Restoring full VM backups” on page 190.

1. If the full-VM that you are restoring will overwrite the existing VMware guest, delete the existing full-VM guest before you restore.

Tip: If you restore the full VM backup to a new VMware guest, you do not need to delete the original.

2. Query the Virtual Machine for full VMware backups using the following steps:
 - a. From the off-host backup server, enter the following command:

```
dsmc q vm *
```

The command lists the backups like the following screen:

| # | 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. View the output to identify the VMware that you are restoring.
3. Restore the full VMware backup using the **REStore VM** command. For example, to restore a virtual machine from a given point in time, use the `-pitdate` and `-pittime` options.

```
dsmc RESTORE VM vmname destination -pitdate=date -pittime=hh:mm:ss
```

where:

vmname

This is the display name of the virtual machine that you want to restore.

destination

Is the directory to restore the vmdk file to.

`-pitdate`

The date that the backup was created.

`-pittime`

The time the backup was created.

4. When the restore completes, you are prompted with the following message. 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))
```

Note: If you enter N, the command line returns without opening the VMware Converter. However, the image must be converted before it can be restored. You can open the Converter tool from the Windows Start menu.

5. Use the VMware vCenter Converter tool to convert the restored VCB image back into a virtual machine on a VMware server.
 - a. From the Converter tool, click **Convert Machine**.
 - b. In the **Virtual machine file** field, enter the directory where the .vmx file was restored to.

Note: The .vmx file was restored to the directory specified by the `vmbackdir` option of the `restore vm` command.

- c. 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 to use it.

Related concepts

“Restoring data from a VMware backup” on page 186

Related tasks

“Restoring full VM backups for VMware Consolidated Backup (VCB)” on page 188

“Restoring full VM backups”

Restoring full VM backups

You can restore a full VMware backup to recreate all of the files of a VMware guest.

You can restore the backup files directly to the VMware server. The restore procedure is different than using the VMware Consolidated Method tools that required you to restore the files to the off-host backup server and use the VMware converter tool before you could restore it to the VMware server.

1. If the full VM that you are restoring will replace the existing VMware guest, delete the existing VM guest before you restore.

Tip: If you restore the VMware backup to a new VMware guest, you do not need to delete the original.

2. Query the Virtual Machine for VMware backups using the following steps:
 - a. From the off-host backup server, enter the following command:

```
dsmc q vm *
```

The command lists the backups like the following screen:

| # | 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. View the output to identify the VMware that you are restoring.

3. Restore the full VMware backup using the **REStore VM** command. For example, in the following command, the virtual machine is restored and given a new name:

```
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 to use it.

Related concepts

“Restoring data from a VMware backup” on page 186

Related tasks

“Restoring file-level VM backups” on page 187

“Restoring full VM backups” on page 190

“Restoring full VM backups for VMware Consolidated Backup (VCB)” on page 188

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 2003 and Windows Server 2008 client, use the **restore adobjects** command to restore local, deleted Active Directory objects (tombstone objects). On the Windows Server 2008 server client, 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 167

Related reference

“Restore Adobjects” on page 612

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.

Important: Tombstone object reanimation is the only method available on the Windows Server 2003 domain controller.

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 2008 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 storing Active Directory objects** before performing the restore.

Related concepts

“Preserve attributes in tombstone objects” on page 194

Chapter 5, “Restoring your data,” on page 161

“Restrictions and limitations when restoring Active Directory objects” on page 193

Related tasks

“Restoring Windows system state” on page 167

Related reference

“Restore Aobjects” on page 612

Restoring Active Directory objects using the GUI and command line

To restore individual Active Directory objects, you must run the Tivoli Storage Manager 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**. For Windows Server 2008, 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 (only on Windows Server 2008):
 - 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 appears showing 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 577

“Restore Adobjects” on page 612

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 Tivoli Storage

Manager, 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. Tivoli Storage Manager 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 Tivoli Storage Manager command line using the object GUID, in which case the command-line client only reanimates the container object and not its children. In the Tivoli Storage Manager 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, Tivoli Storage Manager 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, Tivoli Storage Manager recreates the container using the original name at the time of the backup. When restoring a tombstone object, Tivoli Storage Manager 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 161

Related reference

“Restore Aobjects” on page 612

“Stagingdirectory” on page 481

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 161

Related reference

“Restore Aobjects” on page 612

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 Aobjects” on page 612

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.

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 Tivoli Storage Manager. 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:

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.

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 tasks

“Creating an Automated System Recovery diskette (Windows Server 2003)” on page 127

Related reference

“Delete Access” on page 553

“Query Access” on page 576

“Set Access” on page 644

Restoring or retrieving files from another client node

After users grant you access to their files on the server, you can restore or retrieve those files to your local system.

You can display 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:

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**.

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 filesystem** command to display a list of file spaces. For example, to display a list of the file spaces of cougar, enter:

```
dsmc query filesystem -fromnode=cougar
```

Important: The Tivoli Storage Manager 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 Tivoli Storage Manager 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 382

“Restore” on page 605

“Retrieve” on page 635

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 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 disk, 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 195

“Restoring or retrieving files from another client node” on page 196

Deleting file spaces

If your Tivoli Storage Manager administrator grants you authority, you can delete entire file spaces from the server.

You cannot delete individual backup 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:

1. From the main window, click **Utilities** → **Delete Filespaces**.
2. Select the file spaces you want to delete.
3. Click **Delete**. Tivoli Storage Manager prompts you for confirmation before deleting the file space.

You can also delete a file space using the **delete filespace** command. Use the `class` option with the **delete filespace** command to delete NAS file spaces.

Related reference

“Class” on page 313

“Delete Filespace” on page 558

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.

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 client notifies the server when files are deleted from a client file space or directory. Selective and incremental-by-date backups do not notify the server about deleted files. Run incremental backups at a frequency consistent with possible restore requirements.

If you request a point-in-time restore with a date and time that is prior to the oldest version maintained by the Tivoli Storage Manager server, the object is not restored to your system. Files which were deleted from your workstation prior to the point-in-time specified will not be restored.

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, Tivoli Storage Manager might not be able to restore all objects to the point-in-time you specify.
2. If you delete a file or directory, the next time you run an incremental backup, the active backup version becomes inactive and the oldest versions that exceed the number specified by the *versions data deleted* attribute of the management class are deleted.

When performing a point-in-time restore, consider the following:

- Tivoli Storage Manager restores file versions from the most recent backup before the specified point-in-time date. Ensure the point-in-time that you specify is not the same as the date and time this backup was performed.
- If the date and time you specify for the object you are trying to restore is earlier than the oldest version that exists on the server, Tivoli Storage Manager cannot restore that object.
- Point-in-time restore restores files deleted from the client workstation after the point-in-time date but not files deleted before this date.
- Tivoli Storage Manager cannot restore a file created after the point-in-time date and time. When a point-in-time restore runs, files that were created on the client after the point-in-time date are not deleted.

To perform a point-in-time restore using the client GUI, use the following steps:

1. Click the **Restore** button in the main window. The Restore window appears.

2. Click the **Point-in-Time** button from the Restore window. The Point in Time Restore window appears.
3. Select the **Use a Point-in-Time Date** selection box. Select the date and time and click **OK**. The point in time that you specified appears in the Point in Time display field in the Restore window.
4. Display the objects you want to restore. You can search for an object by name, filter the directory tree, or work with the directories in the directory tree.
5. Click the selection boxes next to the objects you want to restore.
6. Click the **Restore** button. The Restore Destination window is displayed. Enter the appropriate information.
7. Click the **Restore** button to start the restore. The Restore Task List window displays the restore processing status.

Note: If there are no backup versions of a directory for the point-in-time you specify, files within that directory are not restorable from the GUI. However, you can restore these files from the command line.

You can start point-in-time restore from the command-line client using the `pitdate` and `pittime` options with the **query backup** and **restore** commands. For example, when you use the `pitdate` and `pittime` options with the **query backup** command, you establish the point-in-time for which file information is returned. When you use `pitdate` and `pittime` with the **restore** command, the date and time values you specify establish the point-in-time for which files are returned. If you specify `pitdate` without a `pittime` value, `pittime` defaults to 23:59:59. If you specify `pittime` without a `pitdate` value, it is ignored.

Related concepts

Chapter 9, “Storage management policies,” on page 231

Related reference

“**Backup Image**” on page 540

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 tasks

“Configuring the Web client” on page 22

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.

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.

Considerations:

- Workstation and remote (NAS) backups are mutually exclusive in a Restore window. After selecting an item for restore, the next item you select must be of the same type (either NAS or non NAS).
- Details will not appear in the right-frame of the Restore window for NAS nodes or images. To view information about a NAS image, highlight the NAS image and select **View** → **File Details** from the menu.
- To delete NAS file spaces, select **Utilities** → **Delete Filespaces**. You can delete both workstation and remote objects.

Restoring NAS files and directories using the Web client

You can use the `toc` option with the `include.fs.nas` option in your client options file to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup.

If you save TOC information, you can use Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the `TOCDESTINATION` attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation. If you do not save TOC information, you can still restore individual files or directory

trees using the RESTORE NODE server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

To restore NAS files and directories:

1. Click the **Restore** from the main window. The Restore window appears.
2. Expand the directory tree if necessary. To expand a node in the tree, click the plus sign (+) next to an object in the tree. 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.

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 496

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 29. 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 authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the Web client. | “ Query Node ” on page 595 |
| query backup | Use the query backup command with the <i>class</i> option to display information about file system images backed up for a NAS file server. | “ Query Backup ” on page 581 |
| query filesystem | Use the query filesystem command with the <i>class</i> option to display a list of file spaces belonging to a NAS node. | “ Query Filespace ” on page 588 |
| restore nas | Restores the image of a file system belonging to a Network Attached Storage (NAS) file server. | “ Restore NAS ” on page 627 |
| 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 573 |
| 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 552 |
| delete filesystem | Use the delete filesystem with the <i>class</i> option to display a list of file spaces belonging to a NAS node so that you can choose one to delete. | “ Delete Filespace ” on page 558 |

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)

Archiving and retrieving files is similar to backing up and restoring files. This section discusses the primary archive and retrieve tasks.

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

All client archive and retrieve procedures in this topic also apply to the Web client, except the following:

- Preferences editor
- Setup wizard

The following are the primary archive and retrieve tasks:

- “Archiving data using the GUI” on page 206
- “Archive data examples using the command line” on page 207
- “Deleting archive data” on page 210
- “Retrieving archives using the GUI” on page 211
- “Retrieve archives using the command line” on page 211

Related concepts

“When to back up and when to archive files” on page 105

Related tasks

“Starting a web client session” on page 91

Archive files

To archive files, select the files you want to archive. You can select the files by name or description, or select them from a directory tree.

Your administrator might have set up schedules to automatically archive certain files on your workstation. The following sections discuss how to archive files without using a schedule.

requires you to assign an archive description for all archived files. An archive description identifies data through a meaningful description that you can use later to identify files and directories. You can enter as many as 254 characters to describe your archived data. If you do not enter a description, Tivoli Storage Manager assigns the following default archive description:

Archive Date: mm/dd/yyyy

where mm/dd/yyyy is the current date.

When you select the archive function from the backup-archive GUI, Tivoli Storage Manager displays a list of all previously used archive descriptions. You can assign these archive descriptions to future archives.

Incremental backup might recall migrated files, while selective backup and archive always recall migrated files, if you do not use the `skipmigrated` option. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about using the `skipmigrated` option.

Related tasks

“Set the client scheduler process to run as a background task and start automatically at startup” on page 217

Snapshot backup or archive with open file support

If open file support has been configured, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or "in use") by other applications.

The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the archive. You can set the `snapshotproviderfs` parameter of the `include.fs` option to `none` to specify which drives do not use open file support.

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

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 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
3. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not installed.
4. To enable open file support in a cluster environment all workstations in the cluster should have the OFS feature configured.
5. When using the open file support feature with either VSS or LVSA, 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 Tivoli Storage Manager Open File Support restrictions and issues, search for the **TSM Client v5.2 Open File Support** document under the **Storage Management** product category at the following Web site:

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

Related concepts

Chapter 12, “Processing options,” on page 269

Related tasks

“Configuring Open File Support” on page 66

Archiving data using the GUI

You can archive specific files or entire directories from a directory tree. You can also assign a unique description for each group of files you archive (archive package).

To archive your files:

1. Click the **Archive** button in the GUI main window. The Archive window appears.
2. Expand the directory tree by clicking the plus sign (+) or the folder icon next to an object in the tree. To search or filter files, click the **Search** icon from the tool bar.
3. Enter a description, accept the default description, or select an existing description for your archive package in the **Description** pull down box.
4. To modify specific archive options, click the **Options** button. Any options you change are effective during the current session *only*. To estimate the transfer time for your archive selections click the **Estimate** button.
5. Click the **Archive** button. The Archive Status window displays the progress of the archive.

Archive data examples 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 30. 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 334. |
| 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 377. |
| 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 377. |
| 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 334. |
| 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 482. |

Table 30. Command-line archive examples (continued)

| Task | Command | Considerations |
|--|---|---|
| Use the <code>v2archive</code> option with the archive command to archive only files in the <code>c:\relx\dir1</code> directory. | <code>dsmc archive c:\relx\dir1\ -v2archive</code> | Tivoli Storage Manager archives only files in the <code>c:\relx\dir1</code> directory. Directories that exist in the path are not processed. For more information about the <code>v2archive</code> option, see “V2archive” on page 502. |
| Use the <code>archmc</code> 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 <code>archmc</code> option, see “Archmc” on page 298. For more information about management classes, see Chapter 9, “Storage management policies,” on page 231. |
| Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as the logical volume <code>\\florence\c\$\snapshots\ snapshot.0</code> , archive the <code>c:\dir1\sub1</code> directory tree from the local snapshot and manage it on the Tivoli Storage Manager server under the file space name C: | <code>dsmc archive c:\dir1\sub1\ -subdir=yes -snapshotroot=\ florence\c\$\snapshots\snapshot.0</code> | See “Snapshotroot” on page 474 for more information. |

Related reference

“Archive” on page 529

Associating a local snapshot with a server file space

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, in conjunction 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 only manage data created by a volume snapshot.

Related reference

“Snapshotroot” on page 474

Archiving data with client node proxy

Archives of multiple nodes that share storage can be consolidated to a common target node name on the Tivoli Storage Manager server.

This is useful when the 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. This support is only available with Tivoli Storage Manager Version 5.3 and higher.

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:

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

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

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

Follow these steps to enable multinode operation from the GUI:

1. Verify that the client node has proxy authority to a target node (or authorized to act as the target node) using the **QUERY PROXYNODE** administrative client command.
2. Select **Edit** → **Preferences** to open the preferences window.
3. Select the **General** tab and fill in the **As Node Name** field with the name of the proxy authorized target node.
4. Click **Apply** and then **OK** to close the preferences window.

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

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

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

Tip:

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

- Do not use target nodes as traditional nodes. Use them only for multiple node processing.

Restrictions enforced within a proxied session:

- You cannot perform a system object or system state backup or restore.
- You cannot access another node (either from GUI drop down or use of the fromnode option).
- You cannot use the clusternode option.
- You cannot perform NAS backup or restore.

Related reference

“Asnodename” on page 299

Deleting archive data

If your administrator has given you authority, you can delete individual archive copies from the Tivoli Storage Manager server without deleting the entire file space.

To determine if you have this authority, select **File** → **Connection Information** from the Tivoli Storage Manager GUI or Web client main menu. Your authority status is provided in the **Delete Archive Files** field.

To delete archive copies from the Tivoli Storage Manager GUI or Web client:

1. Select **Delete Archive Data** from the **Utilities** menu. The Archive Delete window appears.
2. Expand the Directory tree by clicking the plus sign (+) or folder icon next to the object you want to expand. Objects on the tree are grouped by archive package description.
3. Click the **Delete** button. Tivoli Storage Manager prompts you for confirmation before deleting archive selections. The Archive Delete Task List window displays the progress of the delete.

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

Related reference

“Delete Archive” on page 554

Retrieve archives

Select the *retrieve* function to recover an archive copy of a file or a directory.

Note: When retrieving a directory, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because retrieves the directories first, then adds the files to the directories.

You can also retrieve archive copies from the directory tree, filter the directory tree, and retrieve archive copies of files owned by someone else. To do any of these, click the **Retrieve** button on the Tivoli Storage Manager main window and follow the directions provided in the task help of the GUI.

Important: When you retrieve a file without any specifications, and more than one version of the archive copy exists on the server, Tivoli Storage Manager retrieves

all of the copies. After the first copy is retrieved, the second copy 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 161

Retrieving archives using the GUI

This topic guides you through the steps to retrieve your archived files using the Tivoli Storage Manager GUI.

1. Click the **Retrieve** button on the GUI main window. The Retrieve window appears.
2. Expand the directory tree by clicking the plus sign (+) or the folder icon next to an object you want to expand. To search or filter files, click the **Search** icon from the tool bar.
3. Enter your search criteria in the Find Files (Retrieve) window.
4. Click the **Search** button. The Matching Files (Retrieve) window appears.
5. Click the selection boxes next to the files you want to retrieve and close the Matching Files (Retrieve) window.
6. Enter your filter criteria in the Find Files (Retrieve) window.
7. Click the **Filter** button. The Retrieve window displays the filtered files.
8. Click the selection boxes next to the filtered files or directories you want to retrieve.
9. To modify specific retrieve options, click the **Options** button located to the right of the **Estimate** button. Any options you change are effective during the current session *only*.
10. Click **Retrieve**. The Retrieve Destination window appears. You can retrieve files to a directory or drive other than the one from where they were originally archived. You can also select how much of the parent directory structure is recreated at the retrieve location.
11. Click **Retrieve**. The Retrieve Status window displays the processing status.

Retrieve archives using the command line

You *retrieve* a file when you want to return an archive copy from the server to your workstation. This topic shows some examples of how to retrieve archived files using the command line.

You can retrieve a single file, a group of files, or all the files in a directory or subdirectory. When you retrieve a file, Tivoli Storage Manager 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 31. Command line examples of retrieving archives

| Task | Command | Considerations |
|---|---|--|
| Retrieve the c:\doc\h2.doc file to its original directory. | dsmc retrieve c:\doc\h2.doc | If you do not specify a destination, the files are retrieved to their original location. |
| Retrieve the c:\doc\h2.doc file under a new name and directory. | dsmc retrieve c:\doc\h2.doc c:\proj2\h3.doc | None |

Table 31. Command line examples of retrieving archives (continued)

| Task | Command | Considerations |
|--|---|--|
| Retrieve all files archived with a specific description to a directory named retr1 at a new location | <code>dsmc retrieve c:* d:\retr1\ -sub=yes -desc="My first archive"</code> | None |
| Retrieve all files from the c:\projecta directory that end with the characters .bak to the c:\projectn directory. | <code>dsmc retrieve c:\projecta*.bak c:\projectn</code> | None |
| Use the pick option display a list of archives from which you can select files to retrieve. | <code>dsmc retrieve c:\project* -pick</code> | See "Pick" on page 424 for more information about the <i>pick</i> option. |
| Retrieve a file originally archived from the diskette, workathome on the a: drive, to a diskette in the a: drive labeled extra . | <code>dsmc retrieve {workathome}\doc\ h2.doc a:\doc\h2.doc</code> | If you are retrieving a file to a disk that has a different label other than the disk from which the file was archived, use the file space name (label) of the archive disk rather than the drive letter. |
| Retrieve the c:\doc\h2.doc file to its original directory on the workstation, named <i>star</i> . | <code>dsmc retrieve c:\doc\h2.doc \\star\c\$\</code> To retrieve the file to <i>star</i> which has been renamed <i>meteor</i> , enter: <code>dsmc retrieve \\star\c\$\ doc\h2.doc \\meteor\c\$\</code> You could also enter: <code>dsmc retrieve \\star\c\$\ doc\h2.doc c:\</code> This example is valid because if the workstation name is not included in the specification, the local workstation is assumed (<i>meteor</i> , in this case). | For the purposes of this manual, the workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This is true even if you are retrieving to the same physical workstation, but the workstation has a new name. |

Related reference

"Retrieve" on page 635

Chapter 7. Tivoli Storage Manager scheduler overview

The Tivoli Storage Manager central scheduler allows client operations to occur automatically at specified times.

In order to understand scheduling with Tivoli Storage Manager, several terms need to be defined:

schedule definition

A definition on the Tivoli Storage Manager server which specifies critical properties of the automated activity including the type of action, the time the action should take place, and how frequently the action takes place. There are numerous other properties which can be set (see the appropriate *Tivoli Storage Manager Administrator's Reference Guide* for a detailed description of the **define schedule**.)

schedule association

An assignment to a specific schedule definition for a client node. Multiple schedule associations allow single schedule definitions to be used by many client nodes. Because schedule definitions are included with specific policy domains, it is only possible for nodes defined to a certain policy domain to be associated with schedules defined in that domain.

scheduled event

A specific occurrence of when a schedule is executed 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 the "Tivoli Storage Manager for" application.

The schedule *startup window* indicates the acceptable time period for a scheduled event to start. The startup window is defined by these schedule definition parameters: *startdate*, *starttime*, *durunits*, and *duration*. The *startdate* and *starttime* options define the beginning of the startup window for the very first scheduled event. The beginning of the startup windows for subsequent scheduled events vary depending on the *period* and *perunit* values of the schedule definition. The *duration* of the schedule window defines the length of the startup window. The schedule action is required to start within the startup window. To illustrate, consider the results of the following schedule definition:

```
define schedule standard test1 action=incremental starttime=12:00:00 period=1
perunits=hour dur=30 duru=minutes
```

| Event | Window start | Window end | Actual start (just an example, times 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 defining or updating a schedule *objects* parameter with file specifications that contain blank spaces, use quotation marks around each file specification that contains blanks, then single quotation marks around all of the specifications.

The following are some examples.

```
objects="C:\Program Files\MyApp\Some file.doc"
objects="C:\Program Files\MyApp\Some file.doc" "C:\Another file.txt"
C:\noblanks.txt'
objects="'E:\My Directory With Blank Spaces\''
```

This ensures that C:\Program Files\MyApp\Some file.doc is treated as a single file name, as opposed to two separate files (C:\Program Files\MyApp\Some, and file.doc).

You can also refer to the *objects* parameter information for the DEFINE SCHEDULE and UPDATE SCHEDULE commands in the appropriate *IBM Tivoli Storage Manager Administrator's Reference*.

Related concepts

“Specifying input strings that contain blank spaces or quotation marks” on page 91

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

This topic lists the processing options that impact the behavior of the client scheduler.

On the Tivoli Storage Manager client, you can define most of these options in the client options file. However, some of these options can be set globally on the Tivoli Storage Manager server for all Tivoli Storage Manager clients. The *Tivoli Storage Manager Administrator's Guide* provides detailed information on all topics described in the section.

This table shows which options are defined by the client and server, and which options are overridden by the server.

| Option | Client defined | Server defined | Server global override |
|--------------------------------------|----------------|--|------------------------------|
| managedservices | X | | |
| maxcmdretries | X | | SET MAXCMDRETRIES command |
| maxschedsessions | | X | |
| postschedulecmd, postnschedulecmd | X | | |
| preschedulecmd, prenschedulecmd | X | | |
| queryschedperiod | X | | SET QUERYSCHEDPERIOD command |
| randomize | | X | |
| retryperiod | X | | SET RETRYPERIOD command |
| schedcmddisabled | X | | |
| schedlogname | X | | |
| schedlogretention | X | | |
| schedmode | X | | SET SCHEDMODES command |
| sessioninitiation | X | X | UPDATE NODE command |
| tcpclientaddress | X | X (also defined on server when sessioninit= <i>serveronly</i> as part of the node definition) | |
| tcpclientport | X | X (also defined on server when sessioninit= <i>serveronly</i> as part of the node definition) | |

Client defined options are defined in the `dsm.opt` file. The Tivoli Storage Manager server can also define some options in a client options set, or as part of the options parameter of the schedule definition. The Tivoli Storage Manager server can also set some options globally for all clients. By default, the client setting for these

options is honored. If the global override on the Tivoli Storage Manager server is set, the client setting for the option is ignored. Defining client options as part of the schedule definition is useful if you want to use specific options for a scheduled action that differ from the option settings normally used by the client node, or are different for each schedule the node 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*.

Return codes from `preschedulecmd` and `postschedulecmd` scripts

This topic lists the return codes that you might see when you use the `preschedulecmd` and `postschedulecmd` options.

- If the command specified by the `preschedulecmd` option ends with a nonzero return code, Tivoli Storage Manager considers the command to have failed. In this case, neither the scheduled event nor any `postschedulecmd` or `postnschedulecmd` command runs. 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 have failed. The administrative **query event** command with `format=detailed` option shows that the event completed with return code 8, unless the scheduled operation completed with a higher return code, in which case the higher return code 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 interpreting the return code from a command, Tivoli Storage Manager considers 0 to mean success, and anything else to mean failure. While this behavior is widely accepted in the industry, it is not 100% guaranteed. For example, the developer of the `widget.exe` command might exit with return code 3, if `widget.exe` ran successfully. Therefore, it is possible that the `preschedulecmd` or `postschedulecmd` command might end with a nonzero return code and be successful. To prevent Tivoli Storage Manager from treating such commands as failed, you should wrap these commands in a script, and code the script so that it interprets the command return codes correctly. The script should exit with return code 0 if the command was successful; otherwise it should exit with a nonzero return code. The logic for a script running `widget.exe` might look like this:

```
run 'widget.exe'  
  if lastcc == 3  
    exit 0  
  else  
    exit 1
```

Related reference

“Postschedulecmd/Postnschedulecmd” on page 427

“Preschedulecmd/Prenschedulecmd” on page 429

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.

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.
- You must restart the scheduler process for updated Tivoli Storage Manager options to take effect.

Tip: Restart the traditional scheduler periodically to free system resources previously used by system calls.

Set 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 which starts automatically when your system is started.

This is true whether you use the client acceptor daemon (CAD) 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 Tivoli Storage Manager client command such as `dsmc query session`. Note that for testing purposes, you can always start the scheduler in the foreground by running `dsmc sched` from a command prompt (without a `managedservices` stanza set).

On Windows platforms, the scheduler and client acceptor daemon run as services. You can create and manage these services using either the setup wizard or the Tivoli Storage Manager Client Service Configuration Utility, `dsmcutil.exe`.

- To start the setup wizard, select **Utilities > Setup Wizard** in the Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager client scheduler can be managed by the client acceptor daemon. When setting up Tivoli Storage Manager scheduler services to run with client acceptor daemon management, two services must be created: the scheduler service and the client acceptor daemon service. When installing the client acceptor daemon service with `dsmcutil.exe`, use the `/cadschedname:` parameter to identify which Tivoli Storage Manager scheduler service the client acceptor daemon should manage. If you use the setup wizard to install the scheduler, you can select the **Use the CAD to manage the scheduler** checkbox, which automatically creates both services and associates them.

Using the Client Service Configuration Utility, you can use either of the following methods:

client acceptor daemon-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 daemon and associate scheduler service with the client acceptor daemon:

```
dsmcutil inst CAD /name:"TSM Client Acceptor" /cadschedname:"TSM Client Scheduler" /node:tsmclient1 /password:secret /autostart:yes
```
5. Manually start the client acceptor daemon service:

```
net start "TSM Client Acceptor"
```

Traditional 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 Tivoli Storage Manager client scheduler service on Windows, you should set the services to automatically recover from a failure, as follows:

- Bring up 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 daemon 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 as **Take No Action** for first, second, and subsequent failures. The same recovery settings can also be defined to increase the reliability of the TSM Journal Service.

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
```

Tivoli Storage Manager displays detailed information about all scheduled work for your client node. Table 32 on page 220 displays sample classic **query schedule** output.

Table 32. 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, Tivoli Storage Manager runs the incremental backup on your default domain. The schedule has no expiration date.

To more accurately determine the status of scheduled events, the **query schedule** output for an enhanced schedule, on Tivoli Storage Manager Version 5.3 client and above, includes new fields. These fields are always displayed, even if it is a classic schedule or a Tivoli Storage Manager Version 5.3 client session with a pre-Version 5.3 server, but the new fields are blank. Note that for a down-level (prior to Tivoli Storage Manager Version 5.3) client, the server reports the period as indefinite and the day of week as an illegal day. Table 33 on page 221 displays sample enhanced **query schedule** output.

Table 33. 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 225

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

4097: Informational message

Example 1:

```
Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:57 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
TSM 515 Scheduler halted.
```

Example 2:

```
Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:57 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Scheduler Terminated, service ending.
```

Example 3:

```
Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
TSM Client Scheduler 'TSM 515 Scheduler'
Started.
```

Example 4:

```
Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/31/2002
Time: 8:29:56 AM
```

User: DILE\Administrator
Computer: MIKEDILE
Description:
Starting Scheduler.

Example 5:

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4097
Date: 10/30/2002
Time: 8:06:09 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Incremental backup of volume '\\MIKEDILE\C\$'

4098: Warning message

Example 1:

Event Type: Warning
Event Source: AdsmClientService
Event Category: None
Event ID: 4098
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Error Initializing TSM Api, unable to verify
Registry Password, see dsierror.log.

Example 2:

Event Type: Warning
Event Source: AdsmClientService
Event Category: None
Event ID: 4098
Date: 9/20/2002
Time: 6:20:10 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
ANS1802E Incremental backup of '\\mikedile\
c\$' finished with 3 failure

4099: Error message

Example 1:

Event Type: Error
Event Source: AdsmClientService
Event Category: None
Event ID: 4099
Date: 9/17/2002
Time: 6:53:13 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Scheduler exited with a result code of 4.

Example 2:

Event Type: Error
Event Source: AdsmClientService
Event Category: None
Event ID: 4099
Date: 9/17/2002
Time: 6:27:19 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
ANS4987E Error processing '\\mikedile\e\$\
tsm520c\client\winnt\mak \dsmwin32.ncb':
the object is in use by another process

4100: Scheduler command message

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4100
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Next Scheduled Event Obtained from Server
SNJEDS1 (MVS):

Schedule Name: NIGHTLY_BACKUP
Action: Incremental
Objects: (none)
Options: (none)
Server Window Start: 19:00:00 on 10/31/2002

4101: Backup/Archive statistics for scheduler Displays backup and archive statistics, which might be useful in determining the success or failure of a command.

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4101
Date: 10/30/2002
Time: 8:29:21 PM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Backup/Archive Statistics for Schedule Backup
NIGHTLY_BACKUP :

Total number of objects inspected: 158,688
Total number of objects backed up: 2,486
Total number of objects updated: 0
Total number of objects rebound: 0
Total number of objects deleted: 0
Total number of objects expired: 12
Total number of objects failed: 0
Total number of bytes transferred: 1.15 GB
Data transfer time: 104.35 sec
Network data transfer rate: 11,564.84 KB/sec
Aggregate data transfer rate: 866.99 KB/sec
Objects compressed by: 100%
Elapsed processing time: 00:23:11

4102: Restore/Retrieve statistics for scheduler Similar to 4101, except that it is for scheduled restore and retrieve commands.

4103: Backup-Archive Client Service Startup Parameters

Event Type: Information
Event Source: AdsmClientService
Event Category: None
Event ID: 4103
Date: 10/31/2002
Time: 8:29:56 AM
User: DILE\Administrator
Computer: MIKEDILE
Description:
Backup/Archive Client Service Startup
Parameters:

Service Name : TSM 515 Scheduler
Last Update : Oct 14 2002
Client PTF Level : 5.1.5.2
Service Directory : D:\Program Files\
Tivoli\TSM515\baclient
Client Options File : E:\users\mikedile\
logfiles\dsm.opt
Client Node : MIKEDILE
Comm Method : (default or obtained from
client options file)
Server : (default or obtained from client
options file)
Port : (default or obtained from client
options file)
Schedule Log : E:\users\mikedile\logfiles\
dsmsched.log
Error Log : E:\users\mikedile\logfiles\
dsmerror.log
MS Cluster Mode : (default or obtained
from client options file)

Journal Based Backup Service:

4097: Informational message

4098: Warning message
4099: Error message
4100: Journal Based Backup service file
monitor parameters
4101: Journal Based Backup service database
parameters
4102: Journal Based Backup Service
configuration parameters

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 285

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 447

“`Schedrestretrdisabled`” on page 454

“`Srvprepostscheddisabled`” on page 478

“`Srvprepostsnapdisabled`” on page 479

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 client scheduler” on page 24.

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.

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 Tivoli Storage Manager client option file or Tivoli Storage Manager server override options.

The disadvantages of using multiple schedule processes:

- A unique file space for each node name on the Tivoli Storage Manager server is created.
- When restoring the data, you must use the same node name associated with the backup.

You must create a separate service for each schedule process. If you are using the client acceptor daemon to manage the scheduler, a client acceptor daemon service and schedule service are required for each schedule. The following is an example of setting up two schedule processes to be managed by the client acceptor daemon:

```
dsmcutil inst /name:"TSM Client Scheduler1" /optfile:"c:\tsm\dsm.opt1"  
/node:tsmcli_sched1 /password:secret /autostart:no /startnow:no
```

```
dsmcutil inst CAD /name:"TSM Client Acceptor1" /optfile:"c:\tsm\dsm.opt1"
```

```
/cadschedname:"TSM Client Scheduler1" /node:tsmcli_sched1 /password:secret  
/autostart:yes
```

```
dsmcutil inst /name:"TSM Client Scheduler2" /optfile:"c:\tsm\dsm.opt2"  
/node:tsmcli_sched2 /password:secret /autostart:no /startnow:no
```

```
dsmcutil inst CAD /name:"TSM Client Acceptor2" /optfile:"c:\tsm\dsm.opt2"  
/cadschedname:"TSM Client Scheduler2" /node:tsmcli_sched2 /password:secret  
/autostart:yes
```

Unique option files are required for each schedule instance, and must be identified at the time of service creation:

Option file #1 (c:\tsm\dsm.opt1)

```
tcps          tsmserve1.storage.sanjose.ibm.com  
nodename      tsmcli_sched1  
passwordaccess generate  
schedlogname  c:\tsm\dsm\sched1.log  
errorlogname  c:\tsm\dsm\error1.log  
schedmode     prompted  
tcpclientport 1507  
domain        h:  
managedservices schedule
```

Option file #2 (c:\tsm\dsm.opt2)

```
tcps          tsmserve1.storage.sanjose.ibm.com  
nodename      tsmcli_sched2  
passwordaccess generate  
schedlogname  c:\tsm\dsm\sched2.log  
errorlogname  c:\tsm\dsm\error2.log  
schedmode     prompted  
tcpclientport 1508  
domain        i:  
managedservices schedule
```

Related concepts

“Change processing options used by the scheduler service” on page 226

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.

The exception to the above rules is warning or error messages that individual files could not be processed. For such a skipped file, the return code is 4. For cases where the return code is not 0, you can examine the dserror.log file (and, for scheduled events, the dsmsched.log file).

For a description of the return codes and their meanings, see the following table.

Table 34. Client return codes and their meanings

| Code | Explanation |
|------|--|
| 0 | All operations completed successfully. |
| 4 | The operation completed successfully, but some files were not processed. There were no other errors or warnings. This return code is very common. Files are not processed for various reasons. The most common reasons are: <ul style="list-style-type: none">• The file satisfies an entry in an exclude list.• The file was in use by another application and could not be accessed by the client.• The file changed during the operation to an extent prohibited by the copy serialization attribute. See “Copy serialization attribute” on page 235. |
| 8 | The operation completed with at least one warning message. For scheduled events, the status is Completed. Review dserror.log (and dsmsched.log for scheduled events) to determine what warning messages were issued and to assess their impact on the operation. |
| 12 | The operation completed with at least one error message (except for error messages for skipped files). For scheduled events, the status is Failed. Review the dserror.log file (and dsmsched.log file for scheduled events) to determine what error messages were issued and to assess their impact on the operation. As a general rule, this return code means that the error was severe enough to prevent the successful completion of the operation. For example, an error that prevents an entire drive from being processed yields return code 12. |

Table 34. Client return codes and their meanings (continued)

| Code | Explanation |
|--------------|---|
| <i>other</i> | <p>For scheduled operations where the scheduled action is <code>COMMAND</code>, the return code is the return code from the command that was executed. If the return code is 0, the status of the scheduled operation is <i>Completed</i>. If the return code is nonzero, then the status is <i>Failed</i>.</p> <p>Some commands might issue a nonzero return code to indicate success. For these commands, you can avoid a <i>Failed</i> status by <i>wrapping</i> the command in a script that invokes the command, interprets the results, and exits with return code 0 if the command was successful (the script should exit with a nonzero return code if the command failed). Then ask your Tivoli Storage manager server administrator to modify the schedule definition to invoke your script instead of the command.</p> |

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

```
selective c:\MyTools\* -subdir=yes
incremental c:\MyPrograms\TestDriver\* -subdir=yes
archive e:\TSM510C\* -subdir=yes
```

If the first command completes with return code 0; the second command completed with return code 8; and the third command completed with return code 4, the return code for the macro is 8.

See the *IBM Tivoli Storage Manager Administrator's Reference* for your operating system for more information about QUERY EVENT.

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 Tivoli Storage Manager associates your data 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 237

“Retention grace period” on page 240

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 35 shows the default values for the backup and archive copy groups in the standard management class.

Table 35. 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 is zero (0) and mode is *modified*, a file or directory is backed up *only if* it changed since the last incremental backup. If frequency is zero (0) and mode is *absolute*, a file is backed up every time you run an incremental backup against it. This attribute is not checked for selective backups.

For archive copy groups, copy frequency is always CMD (command). There is no restriction on how often you archive a file.

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. Tivoli Storage Manager 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. Tivoli Storage Manager attempts to perform a backup or archive as many as four additional times, depending on the value specified on the `changingretries` option in your 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. Tivoli Storage Manager 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.

Only for Windows 2003 and XP: If open file support is configured, Tivoli Storage Manager performs a snapshot backup of files that are open or locked by other applications. The snapshot allows the backup to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to the file system are not included in the backup.

Related concepts

“Open file support for backup operations” on page 99

Related tasks

“Configuring Open File Support” on page 66

Related reference

“Snapshotproviderimage” on page 473

Copy mode attribute

The *copy mode* attribute determines whether a file or directory is considered for incremental backup regardless of whether it changed or not since the last backup.

Tivoli Storage Manager does not check the mode for selective backups.

The value for this parameter can be one of the following:

- **Modified.** The file is considered for incremental backup *only if* it has changed since the last backup. A file is considered changed if any of the following are true:
 - The date or time of the last modification is different.
 - The file size is different.
 - The file attributes, with the exception of archive, are different.
 - If only the file metadata changes (such as access permissions), but the file data does not change, Tivoli Storage Manager might back up only the metadata.
- **Absolute.** The file is considered for incremental backup regardless of whether it changed since the last backup. For archive copy groups, the mode is always *absolute*, indicating that a file is archived regardless of whether it changed since the last archive request.

Copy destination 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 40

Related reference

“Deduplication” on page 332

“Enablededupcache” on page 354

“Exclude options” on page 364

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. Tivoli Storage Manager does not check those attributes 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 235

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 excl` 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 **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 example below demonstrates how to use a management class:

```
exclude ?:\...\*.sno
include c:\winter\...\*.ice mcweekly
include c:\winter\december\*.ice mcdaily
include c:\winter\january\*.ice mcmonthly
include c:\winter\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, **mcmonthly**.
3. Any file with an extension of `ice` in the **december** directory is assigned to the management class, **mcdaily**.
4. Any other files with an extension of `ice` in any directory under **winter** are assigned to the management class, **mcweekly**.
5. Any file with an extension of `sno` in any directory is excluded from backup. The exception to this rule is `white.sno` in the `february` directory in the `winter` directory.

To specify your own default management class `mgmt_class_name` for files that are not explicitly included, put the statement below at the top of your include list:

```
include ?:\...\* mgmt_class_name
```

Related reference

“Dirmc” on page 338

“Include options” on page 393

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" (RETONLY) setting meets your backup requirements for directories, it might not be necessary to take any action to associate directories with that management class. Tivoli Storage Manager does it automatically when it backs up your directories.

If there is more than one management class with the longest RETONLY setting, 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 Tivoli Storage Manager client updates the insert date on the archived directory to ensure that the directory is not expired before the files under it.

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 a Tivoli Storage Manager GUI, any previous copies of the file that you archived remain bound to the management class specified when you archived them.

If a file is deleted on the client system then that inactive objects of the file are not rebound.

Refer to the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information on how files and directories are associated with management classes.

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.

Refer to the *IBM Tivoli Storage Manager Administrator's Guide* for your operating system, for more information on how files and directories are associated with management classes.

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 Tivoli Storage Manager manages a file using the backup retention grace period, it does not create any new backup versions of the file. All existing backup versions of the file expire 30 days (or the number of days specified in your policy domain) from the day they are marked inactive.

Archive copies are never rebound because each archive operation creates a different archive copy. Archive copies remain bound to the management class name specified when the user archived them. If the management class to which an archive copy is bound no longer exists or no longer contains an archive copy group, the server uses the default management class. If you later change or replace the default management class, the server uses the updated default management class to manage the archive copy. If the default management class does not contain an archive copy group, the server uses the archive retention grace period specified for the policy domain.

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.

Using the Tivoli Storage Manager copy group value `RETINIT=CREATE` starts the data retention period when the file is archived. Using the copy group value `RETINIT=EVENT` starts the data retention period when the server is notified that the event has occurred.

The following example demonstrates this concept:

The user has two files, `create.file` and `event.file`. The user has available two management classes; `CREATE`, with `RETINIT=CREATE`, and `EVENT`, with `RETINIT=EVENT`. Both management classes have a 60-day retention period. The user, on the same day, archives both files:

```
dsmc archive create.file -archmc=CREATE
dsmc archive event.file -archmc=EVENT
```

Ten days later, the user issues the `set event -type=hold` command for the `create.file` file, so the file cannot be deleted. On the same day the user issues the `set event -type=activate` for the `event.file` file. At this time, `create.file` has 50 days left on its retention period, and `event.file` has 60 days. If no other action is taken, `create.file` 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).

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

Related reference

“Set Event” on page 646

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. Using the Client Service Configuration Utility

The following Tivoli Storage Manager services can be installed during installation or by using the Tivoli Storage Manager Client Service Configuration Utility: (1) Backup-Archive Scheduler Service, (2) Client Acceptor Service (CAD), (3) Remote Client Agent Service, and (4) Tivoli Storage Manager Journal Engine Service.

For more information about using the Tivoli Storage Manager Client Service Configuration Utility to install the Client Acceptor Service and Remote Client Agent Service, see related information about using the `Dsmcutil` command.

Related concepts

“Use the `Dsmcutil` command” on page 246

Install the backup-archive scheduler service

You can use either the Tivoli Storage Manager GUI or the TSM Client Service Configuration Utility to install the Tivoli Storage Manager scheduler.

- From the Tivoli Storage Manager GUI click **Utilities** and then click **Setup Wizard**. Select the **Help me configure the TSM Client Scheduler** option.
- If you have an account that belongs to the Administrator/Domain Administrator group, you can use the **TSM Client Service Configuration Utility** to configure client services on both local and remote Windows 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 CAD to manage the scheduler.

This example illustrates the use of the Tivoli Storage Manager scheduler.

When the Tivoli Storage Manager backup-archive client is registered with a Tivoli Storage Manager server, the procedure involves the following steps:

1. **On the server:**
 - a. Define a schedule for the policy domain to which the backup-archive client is registered.
 - b. Associate the backup-archive client node to the defined schedule.
2. **On the backup-archive client:**
 - a. Install the Tivoli Storage Manager scheduler client as a Windows service for the backup-archive client.
 - b. Start the scheduler service installed for the backup-archive client.

Examples: Automating backups

This section provides examples of how to automate your backups.

The example below uses the following assumptions:

- The Tivoli Storage Manager backup-archive client is registered to a Tivoli Storage Manager server with a node name of **mars** and a password of **marpswd** 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:

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. Tivoli Storage Manager displays this message: ANR2500I Schedule WNT_DAILY_INCR defined in policy domain BACLIWNT.
2. To associate the backup-archive client to this schedule, issue the following command: define association bacliwnt wnt_daily_incr mars. Tivoli Storage Manager displays this message: ANR2510I Node MARS associated with schedule WNT_DAILY_INCR in policy domain BACLIWNT. At this point, a schedule that performs an incremental backup is defined on the Tivoli Storage Manager server. The schedule starts around 9:00 pm. The schedule is re-executed once a day and can start on any day of the week. If you want to confirm that the schedule and association are set correctly, you can use the **Query Schedule** command.

On the backup-archive client:

This example assumes that you installed the Tivoli Storage Manager backup-archive client in the c:\program files\tivoli\tsm\baclient directory. It also assumes that the options files in each of these directories are updated so that the communication parameters point to the Tivoli Storage Manager server.

1. Login using an account with administrative privileges.
2. Open a command prompt window and issue the following command: cd /d "c:\program files\tivoli\tsm\baclient". If the path contains a space, for example c:\program files\tivoli\tsm\baclient, enclose the name in double 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 *shedlogname* option in the options file, dsm.opt.
3. Output from scheduled commands is sent to the log file. After scheduled work is performed, check the log to ensure the work completed successfully. When a scheduled command is processed the schedule log might contain the following entry: Scheduled event *eventname* completed successfully. This is merely an indication that Tivoli Storage Manager successfully issued the scheduled command associated with the *eventname*. No attempt is made to determine the success or failure of the command. You should assess the success or failure of the command by evaluating the return code from the scheduled command in the schedule log. The schedule log entry for the return code of the command is prefaced with the following text: Finished command. Return code is:.

Related tasks

"Dsmcutil valid options" on page 254

Related reference

"Query Schedule" on page 598

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.

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:

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 247

Related tasks

"Dsmcutil valid options" on page 254

Creating a new scheduler and associating a CAD to manage the scheduler

This section provides step-by-step instructions to create a new scheduler and associate a CAD to manage the scheduler.

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 installing 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"`.

Note: The CAD and scheduler starts, as described above. 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.


Use the Dsmcutil command

The Tivoli Storage Manager Client Service Configuration Utility, **dsmcutil**, allows Tivoli Storage Manager Client Services installation and configuration on local and remote Windows workstations.

You can install the following Tivoli Storage Manager services under a double-byte character set (DBCS) path:

- Backup-Archive Scheduler Service
- Client Acceptor Service (CAD)
- Remote Client Agent Service
- Journal Engine Service
- VMware Tools Service

The Client Service Configuration Utility must be run from an account that belongs to the Administrator/Domain Administrator group. The syntax of the utility is:



Note:

1. For Windows Vista, you must right click on the command prompt and select "Run as administrator".
2. Options that you specify with `dsmcutil` commands override those in your options file (`dsm.opt`).

The account running 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 documented here, the minimum abbreviation is in upper case. Upper case letters denote the shortest acceptable truncation. If an item appears entirely in upper case letters, it cannot be truncated. You can type the item in any combination of upper case or lower case letters.

Related concepts

Chapter 2, "Configure the Tivoli Storage Manager client," on page 19

Dsmcutil commands: Required options and examples

The `INSTall` command installs and configures a Tivoli Storage Manager Service.

`INSTall Scheduler`

Installs and configures the Tivoli Storage Manager Scheduler Service.

These are the required `INSTall` command options:

- `/name:servicename`
- `/password:password`
- `/clusternode:Yes|No` (required if running the Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS)).
- `/clustername:clustername` (required if running the MSCS or VCS).

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.

The `/clientdir:clientdir` option can also be used, the default is the current directory.

The following files must exist in the directory specified for `/clientdir:clientdir`:

- `dsmcsvc.exe`
- `dscenu.txt`
- `dsm.opt`
- `dsmntapi.dll`
- `tsmutil1.dll`

Note: If the service is being installed on a remote 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 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 Tivoli Storage Manager server is contacted to verify that the specified node and password are valid. When the password is validated it is generated (encrypted) into the Windows Registry:

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service"  
/node:ALPHA1 /password:nodepw /autostart:yes
```

Task Install a scheduler service named **TSM Central Scheduler Service** on remote 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.sanjose.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 command.

Command:

```
dsmcutil install scheduler /name:"TSM Central Scheduler Service:  
group-a" /clientdir:c:\tsm\baclient /optfile:q:\tsm\baclient\  
dsm.opt /node:mcs-cluster-group-a /password:n  
/validate:no /autostart:yes /startnow:yes  
/clusternode:yes /clustername:mcs-cluster
```

INSTAll CAD

Installs and configures the Client Acceptor Service (CAD). Required options are:

- */name:servicename*
- */node:nodename*
- */password:password*

Other valid options are:

- */optfile:optionsfile*

- */httpport:httpport*
- */webports:webports*

Task Install a Client Acceptor Service called TSM CAD. TSM CAD uses a node called **test** to connect to the Tivoli Storage Manager server. Use the options file `c:\tsm\baclient\dsm.opt` to connect to Tivoli Storage Manager.

Command:

```
dsmcutil install cad /name:"TSM CAD" /node:test
/password:test /optfile:c:\tsm\baclient\dsm.opt
```

INSTAll Journal

Installs a journaling engine service on all Windows clients. 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 Tivoli Storage Manager 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:servicename*
- */node:nodename*
- */password:password*
- */partnername:partner service name*

Other valid options are:

- */optfile:optionsfile*

Task Install a Remote Client Agent Service called TSM AGENT. TSM AGENT uses a node called **test** to connect to the Tivoli Storage Manager server. The options file `c:\tsm\baclient\dsm.opt` is used to connect to Tivoli Storage Manager. The partner CAD service is TSM CAD.

Command:

```
dsmcutil install remoteagent /name:"TSM AGENT" /node:test
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:
"TSM CAD"
```

Note: Both the Remote Client Agent Service and the Client Acceptor Service must be installed to run the Web client. The Client Acceptor Service must be

installed *before* the Remote Client Agent Service. Use the */partnername:* option to specify the name of the partner Client Acceptor Service.

INSTall vmwaretools

Installs and configures the VMware tools service. There are no required options.

Command:

```
dsmcutil install vmwaretools
```

REMove

Remove an installed Client Service. The required option is */name:servername*.

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 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:servername*, and the registry values to update. Other valid options are:

- */clientdir:clientdir*
- */optfile:optionsfile*
- */eventlogging:[Yes|No]*
- */node:nodename*
- */autostart:[Yes|No]*
- */clusternode:Yes|No* (required if running the MSCS or VCS).
- */clustername:clustername* (required if running the MSCS or VCS).

Task Update the client directory and options file for the specified scheduler service. All required client service files must reside in the specified directory.

Note: The communication options specified with the **dsmcutil** command here take precedence over those specified in the client options file.

Command:

```
dsmcutil update /name:"TSM Central Scheduler Service"  
/clientdir:c:\tsmptf3\baclient /optfile:c:\tsmptf3\baclient\dsm.opt
```

Task Update the specified scheduler service to use the TCP/IP protocol to connect to a Tivoli Storage Manager server at the specified host name on the specified port.

Command:

```
dsmcutil update /name:"TSM Central Scheduler Service"  
/commsserver:nt1.sanjose.com /commport:1521 /commmethod:  
tcpip
```

UPDate CAD

Updates Client Acceptor Service registry values. The required option for this command is */name:servername*, and the registry values to update. Other valid options are:

- */node:nodename*
- */password:password*
- */optfile:optionsfile*
- */httpport:httpport*
- */webports:webports*
- */cadschedname:schedulername*

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:servername*, and the registry values to update. Other valid options are:

- */node:nodename*
- */password:password*
- */optfile:optionsfile*
- */partnername:partner service name*

Task Update a Remote Client Agent Service called TSM AGENT. TSM AGENT uses a node called **test** to connect to the Tivoli Storage Manager server. The options file *c:\tsm\baclient\dsm.opt* is used to connect to Tivoli Storage Manager. The partner CAD service is TSM CAD.

Command:

```
dsmcutil update remoteagent /name:"TSM AGENT" /node:test  
/password:test /optfile:c:\tsm\baclient\dsm.opt /partnername:  
"TSM CAD"
```

Query Scheduler

Query Scheduler Service registry values. Required options are: */name:servername*. Other valid options are:

- */machine:machinename*
- */clientdir*
- */optfile*

- */eventlogging*
- */node*
- */commmethod*
- */commport*
- */commserver*
- */errorlog*
- */schedlog*

Note: Do not specify a value for the non-required options. Tivoli Storage Manager returns option registry values for the scheduler service you specify.

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:servername*. Other valid options are:

- */machine:machinename*
- */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 (32-bit). There are no valid options for this command.

Task Query the journaling engine service, **TSM Journal Service**.

Command:

```
dsmcutil query journal
```

Query REMOTEAgent

Queries Remote Client Agent Service registry values. The required option for this command is */name:servername*. Other valid options are:

- */machine:machinename*

- */node*
- */optfile*
- */partnername*
- */clientdir*

Note: Do not specify a value for these options.

Task Query registry settings for the specified Remote Client Agent Service.

Command:

```
dsmcutil query remoteagent /name:"TSM AGENT"
```

List

Lists installed Client Services. There are no required options.

Task Locate and list the installed Tivoli Storage Manager client services on the local workstation.

Command:

```
dsmcutil list
```

Task List the installed Tivoli Storage Manager 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: servicename* option.

Task Start the journaling engine service, **TSM Journal Service**.

Command:

```
dsmcutil start /name:"TSM Journal Service"
```

STOP

Use the **Stop** command to stop a client service. The **Stop** command requires the */name: servicename* option.

Task Stop the journaling engine service, **TSM Journal Service**.

Command:

```
dsmcutil stop /name:"TSM Journal Service"
```

UPDATEPW

Generate encrypted Tivoli Storage Manager registry password. The **UPDATEPW** command requires the */node: nodename* and */password: password* options. Optionally, you can use the following options:

- */validate:[Yes|No]*
- */clusternode:Yes|No* (required if running the MSCS or VCS).
- */clustername:clustername* (required if running the MSCS or VCS).
- */force:[Yes|No]*

The password is validated with the Tivoli Storage Manager server if the */validate:Yes* option is specified. The password is updated on the server if you specify the */updateonserver:Yes* option. If you specify this option, you must specify the current password with the */oldpassword* option.

Task Update the encrypted registry password for the specified node. Validate and update the password on the specified Tivoli Storage Manager server which resides on the specified TCP/IP hostname and port:

Command:

```
dsmcutil updatepw /node:alpha1 /commMethod:tcpip  
/commServer:alpha1.sanjose.ibm.com /commPort:1500  
/password:newpw /oldpassword:oldpw /updateonserver:yes
```

Related concepts

“Journal-based backup” on page 114

Related tasks

“Dsmcutil valid options”

Related reference

“Incremental” on page 564

Dsmcutil valid options

This section lists the valid **dsmcutil** options that you can specify to use the scheduler service.

/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 */clustername* (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 Tivoli Storage Manager options file. This is the options file the specified Client Service uses to connect to Tivoli Storage Manager. The utility also uses the file to connect to Tivoli Storage Manager to validate and update passwords. Note that although this option overrides the default option file in the current directory (*dsm.opt*), the Tivoli Storage Manager API requires that a default option file exists in the

current directory. UNC names are not allowed if the local system account is set to logon. The default is the `dsm.opt` file in the `/clientdir` directory.

/partnername:partner service name

This option is used when installing a Remote Client Agent Service to specify the partner Client Acceptor Service.

/password:password

The Tivoli Storage Manager password which is generated (encrypted) into the Windows Registry.

/schedlog:schedlog

Specifies the fully qualified name of the client schedule log.

/srcdir:pathname

Use this option in conjunction with the */copyfiles* option to specify the fully qualified source path to copy the service installation to another location prior to installing the service.

/startnow:[Yes | No]

Specifies whether `dsmcutil` starts the specified service after executing the command; the default is *Yes*. If you specify *No*, you must start the service manually using the services control panel applet, or the NET START **name of the Tivoli Storage Manager service**.

/updateonserver:[Yes | No]

Specifies whether the specified password is updated on the Tivoli Storage Manager server. Requires using the */oldpassword* option.

/validate:[Yes | No]

Specifies whether to perform validation when displaying or updating the encrypted registry password. The default is *Yes*.

/webports: webports

Specifies the TCP/IP port number used by the Tivoli Storage Manager Client Acceptor service and the Tivoli Storage Manager Web client agent service for communications with the Web GUI.

Chapter 11. Windows Server 2003 ASR supplemental information (Windows)

This section provides supplemental information for using Microsoft Windows Automated System Recovery (ASR) to recover Windows 2003 systems with the IBM Tivoli Storage Manager backup-archive client for Windows.

Requirement: You must install TSM dlls and VC90 CRT as a side-by-side assembly to run within WinPE. Use the script DSMC-PE.CMD to set up the environment. Run the script from the network share where the client package is extracted.

Note: The scope of this section does not include disaster recovery planning.

To set up the environment, issue the following command:

```
DSMC-PE SET
```

To reset the changes made by the DSMC-PE SET command, issue the following command:

```
DSMC-PE CLEAN
```

The functions and the information described in this section are also documented in the Tivoli field guide entitled *Using Microsoft Windows Automated System Recovery (ASR) to Recover Windows XP and Windows 2003 Systems with the IBM Tivoli Storage Manager Backup-Archive Client for Windows*, which is available at the following Web site:

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

Related concepts

“Windows Server 2003 ASR Reference information” on page 267

Windows Server 2003 ASR questions and answers

This section provides some questions and answers about Windows Server 2003 system recovery.

Q: What if I did not create the ASR diskette prior to needing it for ASR?

A: You can also create the ASR recovery diskette for another node if you are granted access to the ASR file space by the other node.

Q: What do I do if the smss.ASR and ntdll.ASR files are not contained in my boot drive backup and I consequently encounter an error in the ASR recovery process?

A: If you have another workstation running the same maintenance level of Windows as the system you want to recover, you can copy the smss.ASR and ntdll.ASR files from the \windows\repair directory of that workstation to a diskette. Re-run the ASR recovery process to the point where the Tivoli Storage Manager ASR recovery script prompts you for network or backup set recovery. Use the minimized Tivoli Storage Manager command window to copy the files to the workstation being recovered. For example:

```
copy a:\*.asr c:\windows\repair
```

Q: How do I create a local backup set containing the information required for ASR recovery?

A: While you can use a single command to place all the required file spaces for recovery in a single file space, it is preferable for performance reasons to use separate backup sets. For Windows 2003, use one backup set for the system drive, and one backup set for system state. Use the following Tivoli Storage Manager server command to create backup sets on the server:

Syntax:

```
generate backupset <nodename> <prefix> "system state" devclass=<file based device class> nametype=unicode
```

```
generate backupset <nodename> <prefix> "\\<machinename>\<systemdrive>" devclass=<file based device class> nametype=unicode
```

Examples:

```
generate backupset nodea nodeabackup "system state" devc=file nametype=unicode
```

```
generate backupset nodea nodeabackup "\\<machinea>\c$" devc=file nametype=unicode
```

To identify the files associated with the backup set, enter the following command:

```
q volhistory type=backupset
```

Copy the files to a DVD using DVD burning software. You can rename the files for convenience.

Q: What files are placed on the ASR recovery diskette, and what do they do?

A: The following files are placed on the ASR diskette:

asr.sif System information file. Provides disk layout information and a list of files to be copied to the system for Tivoli Storage Manager to perform ASR recovery.

asrnpn.sif

Plug and Play device information used by Windows during an ASR operation.

tsmasr.cmd

Command file that installs the Tivoli Storage Manager client and invokes the commands to restore system state and system and boot drives.

waitforevent.exe

Program that detects the completion of the Tivoli Storage Manager client installation.

tsmasr.opt

A copy of the Tivoli Storage Manager client options file that matches the Tivoli Storage Manager environment in effect during the ASR backup.

Q: After performing the ASR restore, the user account from which the backup was performed has missing shortcuts in the start menu and has lost desktop customizations. What causes this to happen?

A: The likely cause is that the user profile of this account was in use during the backup or restore operation.

Q: Why does the `dsmc restore` command seem to run more slowly in ASR mode than when I run it on a fully installed system? The output scrolls slowly on the console. How can I make restore run faster?

A: In the GUI setup mode phase of ASR, the default video driver installed by Windows does not provide support for hardware acceleration. To improve performance, you can do the following:

1. Minimize the command window after you respond to the prompts for network or local backup set restore. If you are performing local backup set restore, you must maximize the command window to respond to prompts for additional restore volumes.
2. Maximize the diagnostic command prompt, and move that window to cover most of the restore command window. This reduces the effort required by the video subsystem to draw the changing text on the restore command window.
3. Use the `quiet` option in the `tsmasr.opt` options file to reduce the amount of feedback displayed on the restore console. You would edit `tsmasr.opt` on the ASR recovery diskette prior to running ASR. The disadvantage of using the `quiet` option is that you will not see any restore progress information except for final restore statistics.

Q: How do you use ASR in conjunction with systems that are running in a cluster?

A: Restoring clustered systems using ASR with Tivoli Storage Manager is similar to the methods you would use in a non-ASR environment. There are two scenarios: (1) restoring a cluster node and (2) restoring an entire cluster when no nodes are running in the cluster.

Restoring a cluster node

This is the most common restore scenario which you would use when an individual node of a cluster fails. The local node must be restored but the shared disks have failed over to another node. The cluster database is backed up by Tivoli Storage Manager as part of **systemstate**. The **restore systemstate clusterdb** command issued during a single system ASR insures the cluster database is restored correctly. No additional steps are required.

When the node reboots, the workstation should join the cluster and begin to work. Since the shared disks have failed over to another node, they appear to be inaccessible to the node.

Restoring an entire cluster

If there are no nodes running in the cluster, and an ASR is performed on the first node, the quorum information is placed on the node disk in the repair directory. The quorum information is not copied to the quorum disk during the ASR process. Use the `clustrest.exe` utility from the Windows Resource Kit to copy the quorum information to the quorum disk following the ASR restore. Reboot the workstation and the single node should resume normal cluster operations. Follow the **Restoring a cluster node** procedure described above for remaining cluster nodes that need to be recovered.

Related tasks

"Creating an Automated System Recovery diskette (Windows Server 2003)" on page 127

Windows Server 2003 ASR diagnostic tips

This section provides some tips to help in your Windows Server 2003 ASR diagnosis.

1. The `tsmasr.cmd` file starts a secondary command window (minimized) for debugging. *Do not* close or interrupt the main `tsmasr.cmd` command window; when this command shell instance is closed, the ASR process exits and the system reboots. You can use this window to retry Tivoli Storage Manager commands issued from `tsmasr.cmd` if there is a failure, such as unavailability of the Tivoli Storage Manager server.
2. You can use tracing. Use `TRACEFLAGS WIN2K` and `SYSTEMOBJECT`. You can modify the `tsmasr.opt` client options file on the floppy disk using Notepad, to include entries for `TRACEFILE` and `TRACEFLAGS`. The `TRACEFILE` can be directed to the floppy if it is not too large, or to a file in a directory that is accessible from the recovery console, such as `%SystemRoot%\temp` or `%SystemRoot%\temp\~tmasr\baclient`.
3. You can modify the `tmasr.cmd` file on the floppy disk using Notepad to include additional diagnostic commands.
4. To improve restore performance, you can edit the `tmasr.opt` client options file and insert the `quiet` option. This significantly reduces the feedback written to the console window during restore; only final statistics are displayed. However, this can inhibit diagnosis in the event of a restore failure.
5. If an ASR recovery failure occurs and you cannot boot the system, you can reboot into Windows recovery console and examine the following files in the `%SystemRoot%\temp\~tmasr\baclient` directory:

dsmerror.log

In this file, you might find entries such as: `ReadPswdFromRegistry(): RegOpenPathEx(): Win32 RC=2`. These are expected when Tivoli Storage Manager attempts to read the generated password from the Windows Registry, and the password is not present.

tmasrdiag.txt

The presence of this file indicates that ASR processing was completed. The file contains the output from the query backup `%SystemRoot%\repair*.asr` command. Ensure that the output of this command shows entries for the `ntd11.ASR` and `smss.ASR` files. These files are generated by the Tivoli Storage Manager **backup asr** command and should be included in the incremental backup of the system drive. If these files are not present after Tivoli Storage Manager restores the system drive, the ASR restore fails. In this case, the `asr.err` file contains error text such as the following:

```
Setup was unable to start the recovery application
"C:\windows\temp\tmasr.cmd /sifpath=C:\windows\repair\asr.sif".
The error code returned was 0x2.
```

6. The setup log from the silent installation of the Tivoli Storage Manager client is found in the `%SystemRoot%\repair\tmasrsetup.log` file.
7. There is limited network functionality available in ASR mode. Simple FTP is available, and basic TCP/IP connectivity can be used in a DHCP environment. NetBios services are not available, so you cannot access files on network shares.

8. There are two lines in the `tsmasr.cmd` file preceded by a REM Diagnostic line. The first line would be used to circumvent a problem with the silent install. The pause would stop ASR processing so that you could use the second command window to perform diagnostics. The second line directs the output of the query backup `%SystemRoot%\repair*.asr` command to a file on the floppy disk, rather than the default (preceding) command, which directs the output to a local file.
9. The **waitforevent.exe** command in the `tsmasr.cmd` file detects the completion of the asynchronous silent installation and pauses the execution of the `tsmasr.cmd` file until the install is complete, or 10 minutes (600) seconds have transpired. If the timeout occurs, the Timed out message displays in the command window. If the Timed out message appears, you should examine the `tsmasrsetup.log` file in the Windows repair directory for installation errors. The Timed out message can also be caused by installing a Tivoli Storage Manager client package prior to Version 5.2.0 in ASR mode. You must have a Version 5.2.0 or higher client to have the required support for ASR recovery.
10. When the ASR process encounters an error during GUI mode setup, the contents of the `asr.err` file appear on the screen in Notepad. Unfortunately, the error codes are not documented by Microsoft. However, you can determine what program was running:

asr_fmt.exe

The ASR format utility for volumes.

asr_pfu.exe

The ASR protected file utility.

asr_ldm.exe

The ASR utility for Logical Disk Manager.

tsmasr.cmd

The Tivoli Storage Manager command file for ASR recovery.

Windows Server 2003 ASR in depth

ASR.SIF: The `asr.sif` file is a Unicode file created by the Tivoli Storage Manager **backup asr** command.

Some portions of the file are generated by the operating system and some are created by Tivoli Storage Manager. This section examines those portions of `ASR.SIF`, shown in **bold** type, that are created by Tivoli Storage Manager in the following example.

Note that the line numbers are for reference in this discussion and are not part of the actual file.

```

1. ;
2. ; Microsoft Windows Automated System Recovery State Information File
3. ;

4. [VERSION]
5. Signature="$Windows NT$"
6. ASR-Version="1.0"
7. Provider="IBM Tivoli Storage Manager"

8. [SYSTEMS]
9. 1="DILEXP", "x86", "5.1", "C:\WINDOWS", 1, 0x00010100, "480 0 -60 0-10-0-5 2:00:00.0 0-4-0-1 2:00:00.0", "Pacific Standard Time", "Pacific Standard Time"

10. [BUSES]
11. 1=1,1

```

12. 2=1,3
13. [DISKS.MBR]
14. 1=1,1,0,0x7a4c5b3b,512,63,255,8467200
15. 2=1,2,1,0x9d7c45ca,512,63,255,12594960
16. [DISKS.GPT]
17. [PARTITIONS.MBR]
18. 1=1,0,0,,0x00,0x05,0x05,16065,3887730,0x0
19. 2=1,1,0,"\\??\Volume{76ae2548-13a2-11d7-902a-002035c2c2a2}"",0x00,0x06,0x06,3903795,3887730,0x0
20. 3=1,4,0,"\\??\Volume{e93d26b4-132e-11d7-aa59-806d6172696f}"",0x00,0x06,0x06,16128,3887667,0x0
21. 4=2,0,3,"\\??\Volume{e93d26b2-132e-11d7-aa59-806d6172696f}"",0x80,0x07,0x07,63,8193087,0x1000
22. 5=2,1,0,,0x00,0x05,0x05,8193150,4385745,0x0
23. 6=2,4,0,"\\??\Volume{e93d26b3-132e-11d7-aa59-806d6172696f}"",0x00,0x0b,0x0b,8193213,4385682,0x0
24. [PARTITIONS.GPT]
25. [COMMANDS]
26. 1=1,3000,0,"%SystemRoot%\system32\asr_fmt.exe","/restore"
27. 2=1,4990,1,"%SystemRoot%\system32\asr_pfu.exe","/restore"
28. 3=1,2000,1,"%SystemRoot%\system32\asr_ldm.exe","/restore"
29. [ASRFMT.FIXEDVOLUMES]
30. 1=1,"\\??\Volume{e93d26b3-132e-11d7-aa59-806d6172696f}"", "\DosDevices\D:", "FAT32", "PROGRAMS", 0x4000
31. 2=1,"\\??\Volume{76ae2548-13a2-11d7-902a-002035c2c2a2}"", "\DosDevices\F:", "FAT", "GHOST2", 0x8000
32. 3=1,"\\??\Volume{e93d26b4-132e-11d7-aa59-806d6172696f}"", "\DosDevices\E:", "FAT", "GHOST1", 0x8000
33. 4=1,"\\??\Volume{e93d26b2-132e-11d7-aa59-806d6172696f}"", "\DosDevices\C:", "NTFS", "", 0x1000
34. [ASRFMT.REMOVABLEMEDIA]
35. 1=1,"\\Device\CdRom0", "\\??\Volume{e93d26b1-132e-11d7-aa59-806d6172696f}"", "\DosDevices\R:"
36. 2=1,"\\Device\Floppy0", "\\??\Volume{e93d26b0-132e-11d7-aa59-806d6172696f}"", "\DosDevices\A:"
37. [ASRLDM.VOLUMESTATE]
38. 0, 34, 1, 125, "LDM Configuration", "Microsoft0"
39. [INSTALLFILES]
40. 1=1, "TSMASR", "\\Device\Floppy0", "tsmasr.opt", "%SystemRoot%\temp\tsmasr.opt", "TSM Options", 0x00000001
41. 2=1, "TSMASR", "\\Device\Floppy0", "tsmasr.cmd", "%SystemRoot%\temp\tsmasr.cmd", "TSM installation batch file", 0x00000000
42. 3=1, "TSMASR", "\\Device\Floppy0", "waitforevent.exe", "%SystemRoot%\temp\waitforevent.exe", "TSM utility file", 0x00000000
43. 4=1, "TSMCLI", "\\Device\CdRom0", "tsmcli.exe", "%SystemRoot%\temp\tsmcli.exe", "TSM Setup Executable", 0x00000001
44. [IBM.TIVOLISTORAGEMANAGER]
45. "IBM Tivoli Storage Manager"

Line 7:

Distinguishes this ASR file from one created by other backup products, such as NTBackup.

Line 29:

In the [COMMANDS] stanza, Tivoli Storage Manager tells ASR to execute tsmasr.cmd.

Lines 41-44:

In the [INSTALLFILES] stanza, Tivoli Storage Manager tells ASR the source and destination of files to be copied during text mode setup. The "TSMASR" and "TSMCLI" entries are the volume labels of the removable media for the source files. The \Device\ entries specify where ASR looks for the removable media (floppy or DVD). The next two entries on each line specify the source and destination of the copy. Next is a comment describing the entry. Finally, there is a flag value. 0x00000001 means that ASR always prompts for the media to be mounted. 0x00000000 means *do not prompt*.

Lines 45-46:

This is a vendor-provided stanza reserved for Tivoli Storage Manager use. Currently, Tivoli Storage Manager does not exploit this stanza area.

TSMASR.CMD

The following example of the tsmasr.cmd file is from a Windows workstation. The line numbers are for reference in this discussion and are not part of the actual file.

```
1. @start "IBM Tivoli Storage Manager" /MIN cmd.exe
2. @echo off
3. @title IBM Tivoli Storage Manager
4. %SystemRoot%\temp\tsmcli.exe -s -e setup.exe -a /s
   /v"INSTALLDIR="%SystemRoot%\temp\~\tsmasr\" ADDLOCAL="Client\"
   TRANSFORM="C:\tsm_images\TSM_BA_CLIENT\1033.mst\" /qn /l*v
   \"C:\WINDOWS\repair\tsmasrsetup.log\" RebootYesNo="No\" REBOOT="Suppress\"
   ALLUSERS=1 "
5. %SystemRoot%\temp\waitforevent.exe 600
6. REM Diagnostics
7. REM pause
8. pushd %SystemRoot%\temp\~\tsmasr\baclient
9. copy %SystemRoot%\temp\tsmasr.opt >nul
10. @echo Do you wish to perform Automated Service Recovery (ASR) restore from the TSM
    server or from a local backup set (from file or tape)?
11. @echo 1. Restore from the TSM server.
12. @echo 2. Restore from a local backup set (from file).
13. @echo 3. Restore from a local backup set (from tape).
14. :choice
15. set /P CH=
16. if /I "%CH%"=="1" goto NETWORK
17. if /I "%CH%"=="2" goto BACKUPSET_FILE
18. if /I "%CH%"=="3" goto BACKUPSET_TAPE
19. goto choice
20. :NETWORK
21. dsmc restore \\DILEXP\C$ C:\ -subdir=yes -preservepath=complete -replace=all -
    tapeprompt=no -nodename=DILEXPNEW -optfile=tsmasr.opt
22. dsmc restore systemobject -asrmode=yes -nodename=DILEXPNEW -computername=DILEXP -
    optfile=tsmasr.opt
23. REM Diagnostics
24. dsmc query backup \\dilexp\c$\WINDOWS\repair\*.asr -asrmode=yes -
    nodename=DILEXPNEW -optfile=tsmasr.opt >tsmasrdiag.txt
25. REM dsmc query backup \\dilexp\c$\WINDOWS\repair\*.asr -asrmode=yes -
    nodename=DILEXPNEW -optfile=tsmasr.opt >a:\tsmasrdiag.txt
26. @GOTO DONE
27. :BACKUPSET_FILE
28. dsmc restore backupset TSMASR \\DILEXP\C$ C:\ -asrmode=yes -subdir=yes -
    preservepath=complete -replace=all -tapeprompt=no -nodename=DILEXPNEW -
    location=file -optfile=tsmasr.opt
29. dsmc restore backupset TSMASR systemobject -asrmode=yes -nodename=DILEXPNEW -
    location=file -replace=all -computername=DILEXP -optfile=tsmasr.opt
30. @GOTO DONE
31. :BACKUPSET_TAPE
32. dsmc restore backupset TSMASR \\DILEXP\C$ C:\ -asrmode=yes -subdir=yes -
    preservepath=complete -replace=all -tapeprompt=no -nodename=DILEXPNEW -
```

```

location=tape -optfile=tsmasr.opt
33. dsmc restore backupset TSMASR systemobject -asrmode=yes -nodename=DILEXPNEW -
location=tape -replace=all -computername=DILEXP -optfile=tsmasr.opt
34. :DONE
35. @echo on

```

Line 1:

Starts a minimized command prompt window which you can use for diagnostic purposes during ASR restore.

Line 4-5:

Performs silent install of the Tivoli Storage Manager client into the %SystemRoot%\temp\~tsmasr directory. The installation progress is logged to the Windows repair directory in file tsmasrsetup.log. The waitforevent command detects the completion of the asynchronous silent installation and pauses the execution of the tsmasr.cmd file until the install is complete, or 10 minutes (600) seconds have transpired. If the timeout occurs then Timed out message is displayed in the command window.

Line 9:

The original client options file (stored as tsmasr.opt) is copied to the baclient directory.

Lines 21-22:

These Tivoli Storage Manager commands are issued to restore the system drive and system state (Windows 2003) from a network-connected Tivoli Storage Manager server. Using the asrmode and computername options in the **restore systemstate** command allow system state to be restored to a different workstation. The name of the workstation in ASR mode is *machinename*.

Lines 24-25:

These lines are for diagnostic use. The presence of the tsmasrdiag.txt file in the baclient directory indicates that the ASR process completed. The tsmasrdiag.txt file contains the output from query backup %SystemRoot%\repair*.asr command. Ensure that the output of this command shows entries for the ntdll.ASR and smss.ASR files.

Lines 28-29 and 32-33:

These Tivoli Storage Manager commands are issued to perform local backup set restore (from tape or file) of the system drive; system state (Windows 2003). Note the special TSMASR backup set designator. This designator in combination with the asrmode=*yes* option, signal the restore process to prompt the user for the first volume of the backup set.

Windows Server 2003 ASR information in the Windows Registry

| This section provides some information about where to find Windows Server 2003
| ASR information in the registry.

The ASR commands and location are found in:

HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows
NT\CurrentVersion\Asr\Commands

The files not to be backed up by ASR are found in:

HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\
BackupRestore\FilesNotToBackup

The registry keys not to be restored during ASR recovery are found in:
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Control\
BackupRestore\AsrKeysNotToRestore

Windows Server 2003 ASR Reference information

This section lists some Microsoft knowledge base articles about Automated System Recovery that are instructive.

- 299044: How to Install Additional Files During Automated System Recovery
- 314470: Definition of System Partition and Boot Partition
- 314686: ASR Cannot Restore System Disk to Replacement Disk of Identical Size
- 314688: A "Logical Disk Manager ASR Utility Error" Message When You Use ASR to Restore Disks That Are in a RAID Set
- 316484: "The Files for the Recovery Diskette Could Not Be Created" Error Message Occurs When You Use Ntbackup for Automated System Recovery

The following publication supplements the recovery information provided in **Disaster Recovery Strategies with Tivoli Storage Management - SG-24-6844**:

<http://www.redbooks.ibm.com/abstracts/sg246844.html>

Chapter 12. Processing options

You can use defaults for processing options or you can tailor the processing options to meet your specific needs. This topic provides an overview of processing options and an options reference section that provides detailed information about each option.

Related concepts

“Using options with commands” on page 289

Related information

Reading syntax diagrams

Processing options overview

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

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

- 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

Tivoli Storage Manager also includes a group of client command options that you can enter *only* on the command line with specific commands. You can override some of the options in your options file by entering them with appropriate backup-archive commands.

Note: Some of the processing options used by the Tivoli Storage Manager central scheduler are defined in the Windows registry when the schedule services are configured. These options can also be specified in the client options file. When the scheduler runs as a service, processing options specified in the registry override the same options specified in the client options file.

Related concepts

“Entering options with a command” on page 290

Related tasks

“Creating and modifying the client options file” on page 19

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 320

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 36. TCP/IP options

| Option | Description |
|--|---|
| <code>httpport</code> “Httpport” on page 386 | Specifies a TCP/IP port address for the Tivoli Storage Manager Web client. |
| <code>lanfreetcpport</code> “Lanfreetcpport” on page 404 | Specifies the TCP/IP port number where the Tivoli Storage Manager storage agent is listening. |
| <code>lanfreetcpserveraddress</code> “Lanfreetcpserveraddress” on page 405 | Specifies the TCP/IP address for the Tivoli Storage Manager storage agent. |
| <code>tcpbuffsize</code> “Tcpbuffsize” on page 488 | Specifies the size, in kilobytes, of the Tivoli Storage Manager internal TCP/IP communication buffer. |
| <code>tcpnodelay</code> “Tcpnodelay” on page 491 | Specifies whether the server or client disables the delay of sending successive small packets on the network. |
| <code>tcpadminport</code> “Tcpadminport” on page 488 | 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> “Tpcadaddress” on page 489 | Specifies a TCP/IP address for <code>dsmcad</code> . |
| <code>tcpport</code> “Tcpport” on page 492 | Specifies the TCP/IP port address for a Tivoli Storage Manager server. |

Table 36. TCP/IP options (continued)

| Option | Description |
|--|---|
| tcpserveraddress "Tcpserveraddress" on page 493 | Specifies the TCP/IP address for a Tivoli Storage Manager server. |
| tcpwindowsize "Tcpwindowsize" on page 494 | Specifies the size, in kilobytes, of the TCP/IP sliding window for your client node. |
| webports "Webports" on page 519 | 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 37. Named Pipes communication option

| Option | Description |
|--|--|
| namedpipename "Namedpipename" on page 415 | 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 38. Shared memory communication options

| Option | Description |
|--|--|
| lanfreeshmport "Lanfreeshmport" on page 404 | 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 459 | Specifies the unique number that is used by the client and the server to identify shared memory area used for communications. |

Server and node options

Authorized User

This topic provides information about the server to contact for backup-archive services, and the client node to request backup-archive services.

Table 39. Server and node options

| Option | Description |
|--|---|
| asnodename "Asnodename" on page 299 | Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel. |
| nasnodename "Nasnodename" on page 416 | Specifies the node name for the NAS file server when processing NAS file systems. |
| nodename "Nodename" on page 417 | Use the nodename option in your dsm.sys file to identify your workstation to the server to establish communications. |

Table 39. Server and node options (continued)

| Option | Description |
|--|---|
| virtualnodename "Virtualnodename" on page 505 | The virtualnodename option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation. |

Node options

This topic explains the node options that you can use to request backup-archive services.

Table 40. Node options

| Option | Description |
|--|---|
| asnodename "Asnodename" on page 299 | Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel. |
| clusterdiskonly "Clusterdiskonly" on page 314 | Specifies whether the Tivoli Storage Manager Backup-Archive client allows the backup of only clustered disks, when running in a Microsoft Cluster Server (MSCS) or VERITAS Cluster Server (VCS) environment and using the <i>clusternode=yes</i> processing option. |
| clusternode "Clusternode" on page 317 | Specifies whether Tivoli Storage Manager manages cluster drives in a Microsoft Cluster Server (MSCS) or Veritas Cluster Server (VCS) environment. |
| nasnodename "Nasnodename" on page 416 | Specifies the node name for the NAS file server when processing NAS file systems. |
| nodename "Nodename" on page 417 | Use the nodename option in your client options file dsm.opt to identify your workstation to the server to establish communications. |
| virtualnodename "Virtualnodename" on page 505 | The virtualnodename option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation. |

Backup and archive processing options

This topic describes the options that you can use to control some aspects of backup and archive processing.

Table 41. Backup and archive processing options

| Option | Description |
|--|---|
| archmc "Archmc" on page 298 | 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 299 | Use the asnodename option to allow agent nodes to back up or restore data on behalf of another node (the target node). This enables concurrent operations from multiple nodes to store data to the same target node and file space in parallel. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|--|
| autofsrename "Autofsrename" on page 307 | 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. |
| backmc "Backmc" on page 309 | Specifies the Tivoli Storage Manager management class to apply to the backup fastback subcommand for retention purposes. |
| changingretries "Changingretries" on page 312 | Specifies the number of times a back up or archive action is retried on a file that is in use. |
| class "Class" on page 313 | 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 322 | 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 323 | 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 326 | The createnewbase option creates a base snapshot and uses it as a source to run a full incremental. This ensures that the backup of any files that might have been skipped during the snapshot difference incremental. |
| deduplication "Deduplication" on page 332 | Specifies whether to eliminate redundant data on the client side when transferring data to the Tivoli Storage Manager server, during backup or archive processing. |
| dedupcachepath "Dedupcachepath" on page 331 | 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 332 | Determines the maximum size of the data deduplication cache file. |
| domain.vmfile "Domain.vmfile" on page 348 | Specifies the virtual machines to include in your backup vm file-level VMware Consolidated Backup operations. |
| domain.vmfull "Domain.vmfull" on page 350 | Specifies the virtual machines to include in your backup vm full VM image VMware Consolidated Backup operations. |

|
|
|
|

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|---|
| enablededupcache "Enablededupcache" on page 354 | Specifies whether you want to enable client-side data deduplication cache, so that Tivoli Storage Manager gets the changed data from the cache. |
| deletefiles "Deletefiles" on page 334 | Use the deletefiles option with the archive command to delete files from your workstation after you archive them. You can also use this option with the restore image command and the incremental option to delete files from the restored image if they were deleted after the image was created. |
| description "Description" on page 334 | The description option assigns or specifies a description for files when performing archive, delete, retrieve, query archive, or query backupset operations. |
| detail "Detail" on page 335 | 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 337 | Use the diffsnapshot option to determine whether Tivoli Storage Manager creates a differential snapshot. |
| dfsbackupmntpnt "Dfsbackupmntpnt" on page 336 | Specifies whether Tivoli Storage Manager sees a Dfs (NTFS or FAT) as a junction or a directory. This option is valid for Windows 2003 server <i>only</i> . |
| dirmc "Dirmc" on page 338 | 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 340 | Backs up, restores, archives, retrieves, or queries directories <i>only</i> . |
| diskcachelocation "Diskcachelocation" on page 342 | 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 343 | Specifies the drives to include in your default client domain for an incremental backup. |
| domain.image "Domain.image" on page 346 | Specifies the file systems and raw logical volumes that you want to include in your client domain for an image backup. This option is for Windows 32-bit clients <i>only</i> . |
| domain.nas "Domain.nas" on page 346 | Specifies the volumes to include in your default domain for NAS image backups. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|--|
| enablearchiveretentionprotection "Enablearchiveretentionprotection" on page 353 | Allows the client to connect to a data retention server. |
| enablelanfree "Enablelanfree" on page 355 | 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. |
| encryptiontype "Encryptiontype" on page 356 | Allows you to select AES 128 bit data encryption, providing a stronger form of data encryption than DES 56 bit data encryption. |
| encryptkey "Encryptkey" on page 357 | Specifies whether to save the encryption key password locally when performing a backup-archive operation or whether to prompt for the encryption key password. |
| exclude.archive "Exclude options" on page 364 | Excludes a file or a group of files that match the pattern from archive services <i>only</i> . |
| exclude.compression "Exclude options" on page 364 | 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 364 | Excludes a directory, its files, and all its subdirectories and their files from backup processing. |
| exclude.encrypt "Exclude options" on page 364 | Excludes specified files from encryption processing. |
| exclude.fs.nas "Exclude options" on page 364 | 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 364 | 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 364 | Excludes files from adaptive subfile backup processing. This option does not apply to archive processing. |
| exclude.systemservice "Exclude options" on page 364 | Excludes individual system services components from backup services. Input can be the keyword or component name to be excluded from back up. |
| fbbranch "Fbbranch" on page 369 | Specifies the branch ID of the remote FastBack server to back up or archive. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|--|
| fbclientname "Fbclientname" on page 370 | Specifies the name of one or more FastBack clients to back up from the backup proxy. |
| fbpolicyname "Fbpolicyname" on page 371 | Specifies the name of one or more Tivoli Storage Manager FastBack policies that you want to back up from the backup proxy. |
| fbrepositlocation "Fbrepositlocation" on page 373 | Specifies the location of the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager client proxy to connect to issue MOUNT DUMP , MOUNT ADD and MOUNT DEL commands. |
| fbserver "Fbserver" on page 374 | Specifies host name of the FastBack server workstation or the FastBack DR Hub workstation that owns the repository specified by the fbrepositlocation option. |
| fbvolumename "Fbvolumename" on page 375 | Specifies the name of one or more Tivoli Storage Manager FastBack volumes to back up from the backup proxy. |
| filelist "Filelist" on page 377 | Specifies a list of files to be processed for the command. Tivoli Storage Manager opens the designated filelist and processes the files listed within according to the command. |
| filesonly "Filesonly" on page 381 | Backs up, restores, retrieves, or queries files <i>only</i> . |
| groupname "Groupname" on page 384 | 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 385 | Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after an operation finishes. |
| ieobjtype "Ieobjtype" on page 388 | Specifies an object type for a client-side data deduplication operation. This option is used with the <code>include.dedup</code> and <code>exclude.dedup</code> options. |
| imagegapsize "Imagegapsize" on page 389 | 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 392 | Specifies the path and file name of an include-exclude options file. |
| "Include options" on page 393 | Use these options to include files or assign management classes for backup processing. |
| include include.backup include.file | |
| include.archive "Include options" on page 393 | Includes files or assigns management classes for archive processing. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|---|
| include.compression "Include options" on page 393 | 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 393 | Includes the specified files for encryption processing. By default, Tivoli Storage Manager does not perform encryption processing. |
| include.fs "Include options" on page 393 | Use the include.fs option to specify processing options for a given 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 393 | Use the include.fs.nas option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the toc option with the include.fs.nas option in your client options file (dsm.opt). See "Toc" on page 496 for more information. |
| include.image "Include options" on page 393 | 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 393 | Includes files for adaptive subfile backup processing. This option does not apply to archive processing. |
| include.systemstate "Include options" on page 393 | Assigns management classes for back up of the Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 system state. The default is to bind the system object to the default management class. |
| incrbydate "Incrbydate" on page 399 | Use with the incremental command to request an incremental backup by date. |
| incremental "Incremental" on page 400 | 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. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|---|
| incrthreshold "Incrthreshold" on page 401 | 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 411 | Specifies a memory-saving backup algorithm for incremental backups when used with the incremental command. |
| mode "Mode" on page 412 | 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 containing a list of files from one or more file space origins. backup vm To specify whether to perform a full or incremental backup of a VMware virtual machine when vmbackuptype=fullvm and when you have installed IBM Tivoli Storage Manager for Virtual Environments. |
| monitor "Monitor" on page 414 | Specifies whether you want to monitor an image backup of file systems belonging to a Network Attached Storage (NAS) file server. |
| noprompt "Noprompt" on page 419 | suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands. |
| nojournal "Nojournal" on page 418 | 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 421 | Specifies the client options file you want to use when you start a Tivoli Storage Manager session. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|--|
| postsnapshotcmd "Postsnapshotcmd" on page 428 | 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. |
| preservelastaccessdate "Preservelastaccessdate" on page 431 | Use this option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following a backup or archive operation. By default, the Tivoli Storage Manager client 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 434 | 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 440 | Specifies whether Tivoli Storage Manager resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server. This option is valid for all Windows clients. |
| skipntpermissions "Skipntpermissions" on page 461 | Specifies whether to back up, archive, retrieve, or restore Windows security information. |
| skipntsecuritycrc "Skipntsecuritycrc" on page 462 | Specifies whether to compute the security CRC for permission comparison during subsequent backups. Use this option on all Windows clients. |
| snapdiff "Snapdiff" on page 463 | Specifies an incremental backup of the files reported as changed by NetApp, instead of scanning the volume looking for files that have changed. Use this option with a NAS full volume incremental backup. |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | During a snapshot-based backup or open file support operation, use this option to specify the location where the LVSA places the cache files. This option is only valid if the LVSA is installed and configured for open file or online image support. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|--|
| snapshotcachesize "Snapshotcachesize" on page 468 | Use this option to determine the size of the cache file containing the original data for blocks that change during the operation. The value is the percent of the used space that is expected to change on the volume on which the snapshot is taken. On a volume with high file system write activity, this value can be increased to remove the condition where the snapshot cache fills up during the operation. The snapshot cache location is specified with the snapshotcachelocation option. This option is only valid if the LVSA is installed and configured for open file or online image support. |
| snapshotfsidleretries "Snapshotfsidleretries" on page 469 | During an online image backup or open file support operation, use this option to specify the number of additional times the LVSA tries to achieve the snapshot file system idle wait time before the online image backup or open file support operation fails. This option is only valid if the LVSA is installed and configured for open file or online image support. |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | During an online image backup or open file support operation, use this option to specify the amount of time that must pass in which there is no write activity (read activity is ignored) on a volume before a snapshot can be taken. This option is only valid if the LVSA is installed and configured for open file or online image support. |
| snapshotproviderfs "Snapshotproviderfs" on page 472 | Use the snapshotproviderfs option to enable snapshot-based file backup and archive operations, and to specify a snapshot provider. |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | Use the snapshotproviderimage option to enable snapshot-based online image backup, and to specify a snapshot provider. |
| snapshotroot "Snapshotroot" on page 474 | Use the snapshotroot option with the incremental , selective , or archive commands in conjunction with a 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 482 | Specifies whether to include subdirectories of a named directory. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|--|---|
| subfilebackup "Subfilebackup" on page 484 | Specifies whether Tivoli Storage Manager uses adaptive subfile backup. |
| subfilecachepath "Subfilecachepath" on page 485 | Specifies the path where the client cache resides for adaptive subfile backup processing. |
| subfilecachesize "Subfilecachesize" on page 486 | Specifies the client cache size for adaptive subfile backup. |
| tapeprompt "Tapeprompt" on page 486 | Specifies whether you want Tivoli Storage Manager to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice. |
| toc "Toc" on page 496 | Use the toc option with the backup nas command or the <code>include.fs.nas</code> option to specify whether Tivoli Storage Manager saves Table of Contents (TOC) information for each file system backup. If you save TOC information, you can use the query toc server command to determine the contents of a file system backup in conjunction with the restore node server command to restore individual files or directory trees. You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. |
| type "Type" on page 500 | Use the type option with the query node command to specify the type of node to query. |
| v2archive "V2archive" on page 502 | Use the v2archive option with the archive command to archive only files to the server. Tivoli Storage Manager will not process directories that exist in the path of the source file specification. |
| virtualfsname "Virtualfsname" on page 504 (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 511 | 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 VMware Consolidated Backup commands are directed. |
| vmcpw "Vmcpw" on page 512 | Used with the backup VM , restore VM , or query VM commands to specify the password of the VirtualCenter or ESX user specified with the <code>vmcuser</code> option. |

Table 41. Backup and archive processing options (continued)

| Option | Description |
|----------------------------------|---|
| vmcuser "Vmcuser" on page 513 | Used with the backup VM , restore VM , or query VM commands to specify the user name of the VMware VirtualCenter or ESX server where the VMware Consolidated Backup commands are directed. |
| vmllist "Vmlist" on page 514 | 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. |

The following options are Backup-Archive Client options that apply only to HSM for Windows migrated files. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about these options:

- Restorecheckstubaccess
- Restoremigstate
- Skipmigrated

Restore and retrieve processing options

This topic describes the options that you can use for restoring and retrieving.

Table 42 lists the restore and retrieve processing options that are available.

Table 42. Restore and retrieve processing options

| Option | Description |
|---|---|
| asrmode "Asrmode" on page 301 | 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 generated in the asr.sif file by the backup asr command only. This option should not be used outside the context of ASR recovery mode. |
| backupsetname "Backupsetname" on page 310 | The backupsetname option specifies either the name of the backup set, or the name of the file or tape device containing the backup set. This option is used in conjunction with the location option. |
| computername "Computername" on page 324 | Use the computername option to specify the name of the computer for which you are performing system state recovery in Automated System Recovery (ASR) mode. This option is used in the context of restore commands generated in the asr.sif file by the backup asr command only. This option should not be used outside the context of ASR recovery mode. |
| dirsonly "Dirsonly" on page 340 | Qualifies the operation (backup, archive, restore, retrieve) to process directories alone. |
| disablenqr "Disablenqr" on page 340 | Specifies whether the Tivoli Storage Manager Backup-Archive Client can use the no-query restore method for restoring files and directories from the server. |

Table 42. Restore and retrieve processing options (continued)

| Option | Description |
|---|--|
| filelist "Filelist" on page 377 | Specifies a file containing a list of files to be processed by the specified command. |
| filesonly "Filesonly" on page 381 | Qualifies the operation (backup, archive, restore, retrieve) to process files alone. |
| fromdate "Fromdate" on page 381 | 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 382 | 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. |
| fromtime "Fromtime" on page 383 | 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. |
| frsprimaryrestore "Frsprimaryrestore" on page 384 | Specifies whether Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost. For example, in a disaster recovery situation when all systems containing the SYSVOL system object are lost, it is necessary to restore the first SYSVOL in this manner. |
| guitreeviewafterbackup "Guitreeviewafterbackup" on page 385 | Specifies whether the client GUI is returned to the Backup, Restore, Archive, or Retrieve window after a successful operation completes. |
| ifnewer "Ifnewer" on page 389 | Replaces an existing file with the latest backup version only if the backup version is newer than the existing file. |
| imagetofile "Imagetofile" on page 390 | 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 391 | Displays a list of active and inactive files when used with the pick option. |
| latest "Latest" on page 407 | Restores the most recent backup version of a file whether it is active or inactive. |
| localbackupset "Localbackupset" on page 408 | Specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the server to restore a local backup set on a stand-alone workstation. |
| monitor "Monitor" on page 414 | Specifies whether you want to monitor an image restore of one or more file systems belonging to a network-attached storage (NAS) file server. |
| noprompt "Noprompt" on page 419 | suppresses the confirmation prompt that is presented by the delete group , delete archive , expire , restore image , and set event commands. |
| optfile "Optfile" on page 421 | Specifies the client options file you want to use when you start a Tivoli Storage Manager session. |

Table 42. Restore and retrieve processing options (continued)

| Option | Description |
|--|---|
| pick "Pick" on page 424 | 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 425 | 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 426 | 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 432 | 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 439 | Specifies whether to overwrite an existing file, or to prompt you for your selection when you restore or retrieve files. |
| showmembers "Showmembers" on page 460 (does not apply to Mac OS X) | Displays all members of a group. |
| subdir "Subdir" on page 482 | Specifies whether you want to include subdirectories of a named directory. |
| tapeprompt "Tapeprompt" on page 486 | Specifies whether you want Tivoli Storage Manager to wait for a tape required for a restore or retrieve to be mounted, or to prompt you for your choice. |
| todate "Todate" on page 497 | 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 498 | 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 500 | Use the type option with the query node command to specify the type of node to query. |
| verifyimage "Verifyimage" on page 504 | 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. |

The following options are Backup-Archive Client options that apply only to HSM for Windows migrated files. Refer to *IBM Tivoli Storage Manager for HSM for Windows Administration Guide* for details about these options:

- Restorecheckstubaccess
- Restoremigstate
- Skipmigrated

Scheduling options

This topic discusses the options that you can use to regulate central scheduling. Tivoli Storage Manager uses scheduling options only when the Scheduler is running.

Table 43 lists the scheduling options that are available.

Table 43. Scheduling options

| Option | Description |
|--|---|
| managementservices “Managementservices” on page 409 | Specifies whether the Storage Manager client acceptor service manages the Web client, the scheduler, or both. |
| maxcmdretries “Maxcmdretries” on page 410 | Specifies the maximum number of times the client scheduler attempts to process a scheduled command that fails. |
| postschedulecmd/postnschedulecmd “Postschedulecmd/Postnschedulecmd” on page 427 | Specifies a command to process after running a schedule. |
| preschedulecmd/preschedulecmd “Preschedulecmd/Preschedulecmd” on page 429 | Specifies a command to process before running a schedule. |
| querschedperiod “Querschedperiod” on page 436 | Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work. |
| retryperiod “Retryperiod” on page 444 | 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 446 | 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 447 | Specifies whether to disable the scheduling of generic commands specified by your Tivoli Storage Manager administrator. |
| schedlogmax “Schedlogmax” on page 449 | Specifies the maximum size of the scheduler log and Web client log, in megabytes. |
| schedlogname “Schedlogname” on page 450 | Specifies the path and file name where you want to store schedule log information. |
| schedlogretention “Schedlogretention” on page 451 | 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 453 | Specifies which schedule mode to use, <i>polling</i> or <i>prompted</i> . |
| schedrestretrdisabled “Schedrestretrdisabled” on page 454 | Specifies whether to prevent the Tivoli Storage Manager Server administrator from executing restore or retrieve schedule operations. |

Table 43. Scheduling options (continued)

| Option | Description |
|--|--|
| sessioninitiation "Sessioninitiation" on page 457 | 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 478 | 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 479 | 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 490 | 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 453) for details. |
| tcpclientport "Tcpclientport" on page 491 | 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 453) for details. |

Format and language options

Format and language options allow you to select different formats for date, time, numbers, and for different languages if you have the appropriate client language pack installed for that language.

Table 44. Format and language options

| Option | Description |
|---|--|
| dateformat "Dateformat" on page 329 | Specifies the format for displaying dates. |
| language "Language" on page 406 | Specifies the language used for messages. |
| numberformat "Numberformat" on page 419 | Specifies the format for displaying numbers. |
| timeformat "Timeformat" on page 495 | Specifies the format for displaying time. |

Command processing options

This topic explains the options that you can use with the Tivoli Storage Manager commands.

Command processing options allow you to control some of the formatting of data on your terminal screen.

Table 45. Command processing options

| Option | Description |
|---|--|
| quiet "Quiet" on page 438 | 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 455 | 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 456 | Specifies whether you want Tivoli Storage Manager to stop and wait after displaying the number of lines of information you specified with the scrolllines option, or scroll through and stop at the end of the information list. |
| verbose "Verbose" on page 503 | 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 46 lists the authorization options that are available.

Table 46. Authorization options

| Option | Description |
|---|--|
| autodeploy "Autodeploy" on page 306 | Specifies whether you want to enable or disable an automatic deployment of the client if a restart is required. |
| password "Password" on page 422 | Specifies a Tivoli Storage Manager password. |
| passwordaccess "Passwordaccess" on page 423 | 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 445 | 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 Tivoli Storage Manager treats the entries in the log file.

Table 47 lists the error processing options that are available.

Table 47. Error processing options

| Option | Description |
|---|---|
| errorlogmax "Errorlogmax" on page 359 | Specifies the maximum size of the error log, in megabytes. |
| errorlogname "Errorlogname" on page 361 | Specifies the fully qualified path and file name of the file where you want to store information about errors that occur during processing. |

Table 47. Error processing options (continued)

| Option | Description |
|--|---|
| errorlogretention "Errorlogretention" on page 361 | 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 Tivoli Storage Manager processes transactions between the client and server.

Table 48 lists the transaction processing options that are available.

Table 48. Transaction processing options

| Option | Description |
|--|---|
| collocatebyfilespec "Collocatebyfilespec" on page 319 | Specifies that you want the Tivoli Storage Manager client to use only one server session to send objects generated from one file specification. Setting the <code>collocatebyfilespec</code> option to <code>yes</code> 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 321 | 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 321 | 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 341 | Specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files. |
| largecommbuffers "Diskbuffsize" on page 341 | This option has been replaced by the <code>diskbuffsize</code> option. At this time, <code>largecommbuffers</code> is still accepted by the Tivoli Storage Manager client in order to ease the transition to the new option. However, the value specified by <code>largecommbuffers</code> is ignored in favor of the <code>diskbuffsize</code> setting. Important: Discontinue the use of <code>largecommbuffers</code> because future releases of Tivoli Storage Manager might not accept this option. |
| resourceutilization "Resourceutilization" on page 442 | Use the <code>resourceutilization</code> option in your client options file <code>dsm.opt</code> to regulate the level of resources the Tivoli Storage Manager server and client can use during processing. |
| txnbytelimit "Txnbytelimit" on page 499 | Specifies the number of kilobytes the client program buffers before it sends a transaction to the server. |
| usedirectory "Usedirectory" on page 501 | Provides a convenient way to simplify client communication configuration by overriding <code>commmethod</code> 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

This topic describes the options for the Tivoli Storage Manager Web Client.

Table 49 lists the Web client options that are available.

Table 49. Web client options

| Option | Description |
|--|--|
| httpport "Httpport" on page 386 | Specifies a TCP/IP port address for the Web client. |
| managerservices "Managerservices" on page 409 | Specifies whether the Storage Manager client acceptor service manages the Web client, the scheduler, or both. |
| revokeremoteaccess "Revokeremoteaccess" on page 445 | Restricts administrator access on a client workstation through the Web client. |
| webports "Webports" on page 519 | 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 50 lists the diagnostics options that are available.

Table 50. Diagnostics options

| Option | Description |
|---------------------------------|---|
| console "Console" on page 325 | Use the console option with the query systeminfo command to output system information to the console. |
| filename "Filename" on page 379 | 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 599

Using options with commands

You can override some of the options in your client options file (dsm.opt) file by entering them with appropriate Tivoli Storage Manager commands.

Tivoli Storage Manager processes options in the following order (precedence):

1. Options defined on the server with server-enforced client options. The server overrides client values.
2. Options entered locally on the command line.
3. Options defined on the server for a schedule using the options parameters.
4. Options entered locally in the options file.

5. Options received from the server with client option sets not set as forced by the server. The server *does not* override client values if not forced.
6. Default option values.

Tivoli Storage Manager also includes a group of client command options that you can enter *only* on the command line with specific commands. For a complete list of command-line options, a description, and where to go for more information, see Table 51 on page 291.

Entering options with a command

This topic lists 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. There should be no 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 submitting the command-line arguments to the Tivoli Storage Manager client application. In such cases, 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 entering 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 set in the preferences file. However, when you use the domain option with the **incremental** command, it adds to the domain specified in your client options file rather than overriding the current value.
- The maximum number of bytes for a file name and file path combined is 6255. However, the file name itself cannot exceed 255 bytes and the path leading 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 51 on page 291 lists client command options that you can enter only on the command line with specific commands.

Table 51. Client command options

| Command option | Description | Commands |
|---|---|--|
| archmc "Archmc" on page 298 | 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 |
| asrmode "Asrmode" on page 301 | 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 generated in the asr.sif file by the backup asr command only. This option should not be used outside the context of ASR recovery mode. This option is valid for the Windows 2003 .NET client only. | restore asr |
| class "Class" on page 313 | Specifies whether to display a list of NAS objects or client objects when using the following commands: | query backup delete filesystem query filesystem |
| computername "Computername" on page 324 | Use the computername option to specify the name of the computer for which you are performing system state recovery in Automated System Recovery (ASR) mode. Use the computername option with restore commands generated in the asr.sif file by the backup asr command only. Do not use this option outside the context of ASR recovery mode. This option is valid for Windows Server 2003 clients only. | restore systemstate restore backupset |
| console "Console" on page 325 | Use the console option with the query systeminfo command to output system information to the console. | query systeminfo |
| deletefiles "Deletefiles" on page 334 | 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 334 | Assigns or specifies a description for files when performing archive, delete, retrieve, or query archive operations. | archive delete archive query archive query backupset retrieve |
| detail "Detail" on page 335 | Displays management class, file space, backup, and archive information depending on the command with which it is used. | delete filesystem query archive query backup query filesystem query mgmtclass |
| dirsonly "Dirsonly" on page 340 | Backs up, restores, archives, retrieves, or queries directories <i>only</i> . | archive incremental query archive query backup restore restore backupset retrieve selective |

Table 51. Client command options (continued)

| Command option | Description | Commands |
|-----------------------------------|---|--|
| filelist "Filelist" on page 377 | Specifies a list of files to be processed for the command. Tivoli Storage Manager opens the designated filelist and processes the files listed within according to the command. | archive backup group delete archive delete backup expire incremental query archive query backup restore retrieve selective |
| filename "Filename" on page 379 | 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 381 | Backs up, restores, retrieves, or queries files <i>only</i> . | archive incremental query archive query backup restore restore backupset retrieve selective |
| fromdate "Fromdate" on page 381 | 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 382 | 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 |
| fromtime "Fromtime" on page 383 | 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 384 | Specifies the fully qualified name for a group. | backup group |

Table 51. Client command options (continued)

| Command option | Description | Commands |
|---------------------------------------|--|--|
| ifnewer “Ifnewer” on page 389 | Replaces existing files with the latest backup version only if the backup version is newer than the existing version. | restore restore backupset restore group retrieve |
| imagetofile “Imagetofile” on page 390 | 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. | restore image |
| inactive “Inactive” on page 391 | 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 399 | Requests an incremental backup by date. | incremental |
| incremental “Incremental” on page 400 | Applies changes to the base image using information from incremental backups made after the original image backup. | restore image |
| latest “Latest” on page 407 | Restores the most recent backup version of a file whether it is active or inactive. | restore restore group |
| mode “Mode” on page 412 | 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 containing a list of files from one or more file space origins. | backup group backup nas backup image restore nas |
| monitor “Monitor” on page 414 | Specifies whether you want to monitor an image backup or restore of one or more file systems belonging to a Network Attached Storage (NAS) file server. | backup nas restore nas |
| nojournal “Nojournal” on page 418 | 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 |

Table 51. Client command options (continued)

| Command option | Description | Commands |
|---|---|---|
| noprompt "Noprompt" on page 419 | 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 421 | Specifies the client options file you want to use when you start a Tivoli Storage Manager session. | dsmc.exe |
| pick "Pick" on page 424 | 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 425 | 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 426 | 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 |
| preservepath "Preservepath" on page 432 | 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 446 | 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 |

Table 51. Client command options (continued)

| Command option | Description | Commands |
|---|---|--|
| showmembers “Showmembers” on page 460 | Displays all members of a group. | query group query systemstate restore group |
| todate “Todate” on page 497 | 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 498 | 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 500 | Use the type option with the query node command to specify the type of node to query. | query node |
| v2archive “V2archive” on page 502 | Use the v2archive option with the archive command to archive only files to the server. Tivoli Storage Manager will not process directories that exist in the path of the source file specification. | archive |
| verifyimage “Verifyimage” on page 504 | 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. This option is valid for Windows 32-bit platforms <i>only</i> . | restore image |
| virtualfsname “Virtualfsname” on page 504 | Specifies the name of the virtual file space for the group on which you want to perform the operation. | backup group |

Initial command-line-only options

There is a subset of client options that are valid on the initial command line only. Many of these options establish the runtime environment, such as the `commethod` and `optfile` options. Options in this category are not valid in interactive, macro, or scheduler modes. They generate an error and cause processing to stop.

Table 52 on page 296 lists the options that are valid only on the initial command line.

Table 52. 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 | queryschedperiod |
| computername | resourceutilization |
| deduplication | retryperiod |
| diskbuffsize | runasservice |
| editor | shedlogmax |
| enablededupcache | shedlogname |
| enablelanfree | shedlogretention |
| errorlogmax | shedmode |
| errorlogname | sessioninitiation |
| errorlogretention | subfilebackup |
| incrthreshold | subfilecachepath |
| lanfreecommmethod | subfilecachesize |
| lanfreeshmport | tcpbuffsize |
| lanfreetcpport | 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 53 on page 297 lists the options that can be set by the server.

Table 53. 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 311
- "Changingretries" on page 312
- "Collocatebyfilespec" on page 319
- "Compressalways" on page 322
- "Compression" on page 323
- "Dirmc" on page 338
- "Disablenqr" on page 340
- "Diskcachelocation" on page 342
- "Domain" on page 343
- "Domain.image" on page 346
- "Domain.nas" on page 346
- "Encryptiontype" on page 356
- "Encryptkey" on page 357
- "Exclude options" on page 364
- "Incllexcl" on page 392
- "Include options" on page 393
- "Managedservices" on page 409
- MAXCANDPROCS
- "Maxcmdretries" on page 410
- MAXMIGRATORS
- "Memoryefficientbackup" on page 411
- "Postschedulecmd/Postnschedulecmd" on page 427
- "Postsnapshotcmd" on page 428
- "Preschedulecmd/Prenschedulecmd" on page 429
- "Preservelastaccessdate" on page 431
- "Presnapshotcmd" on page 434
- "Queryschedperiod" on page 436
- "Quiet" on page 438
- "Resetarchiveattribute" on page 440
- "Resourceutilization" on page 442
- "Retryperiod" on page 444
- "Schedmode" on page 453
- "Scrolllines" on page 455
- "Scrollprompt" on page 456
- "Snapshotcachelocation" on page 467
- "Snapshotcachesize" on page 468
- "Snapshotfsidleretries" on page 469
- "Snapshotfsidlewait" on page 470
- "Snapshotproviderfs" on page 472
- "Snapshotproviderimage" on page 473
- "Stagingdirectory" on page 481
- "Subdir" on page 482
- "Subfilebackup" on page 484
- "Subfilecachepath" on page 485
- "Subfilecachesize" on page 486
- "Tapeprompt" on page 486
- "Txnbytelimit" on page 499
- "Verbose" on page 503
- "Vmchost" on page 511
- "Vmcpw" on page 512
- "Vmcuser" on page 513
- "Vmllist" on page 514
- "Vmprocessvmwithindependent" on page 515
- "Vmprocessvmwithprdm" on page 516

Note:

1. See Tivoli Storage Manager for Space Management for UNIX and Linux, client command line interface help or user's guide.
2. See *Tivoli Storage Manager for Mail, Data Protection for Microsoft Exchange Server Installation and User's Guide*.

Related information

For information about creating client options sets on the Tivoli Storage Manager server, see the Tivoli Storage Manager Administrator's Guide for your server platform.

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.

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 Windows Server 2003 and Windows Server 2008 clients only. The Tivoli Storage Manager client API does not support this option.

Syntax



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 Windows Server 2008 only.

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 client API does not support this option.

Syntax

►►—ARCHMC =—*managementclass*—►►

Parameters

managementclass

Specifies an available management class in the active policy set of your policy domain. This management class overrides the default management class and any `include` statements for the files and directories you are archiving.

Examples

Command line:

```
dsmc archive -archmc=ret2yrs c:\plan\proj1\budget.jan\*
```

Asnodename

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 has been granted authority by a Tivoli Storage Manager administrator to perform client operations on behalf of a target node, which is another client node. The administrator uses the **grant proxynode** command on the Tivoli Storage Manager server to grant this authority.

When you back up or archive data without the `asnodename` option, the backed up data is stored in a file space on the Tivoli Storage Manager server that matches the Tivoli Storage Manager nodename of the drive on which the original data exists. For example, if a computer with a nodename of "TAURUS" backs up data on its F: drive, then the data is stored in a server file space named `\\taurus\fs`; node Taurus owns the file space.

When you use the `asnodename` option to back up data on behalf of a target node, the data is stored in a file space owned by the target node. For example, if a computer having a Tivoli Storage Manager nodename of TAURUS backs up data on its F: drive using `-asnodename=scorpio`, the data is stored in a file space named `\\scorpio\fs`, and the files are owned by the 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 TAURUS uses the option `-asnodename=scorpio` to restore data from `\\scorpio\f$`, Tivoli Storage Manager attempts to restore the data to the F: drive on a computer named SCORPIO. This might not produce the expected result, especially if the SCORPIO computer does not exist or if you do not have access to it.

To restore files, specify a destination for the restored files. In the following example, the restore command is entered on the TAURUS node. The command restores all files and subdirectories from the `Users\taurus\education` directory in the `\\scorpio\f$` file space to the local F: drive.

```
dsmc restore \\scorpio\f$\users\taurus\education\* f:\ -subdir=yes
-asnodename=scorpio
```

Restrictions:

- You cannot use the `asnodename` option with `-fromnode` and you cannot perform NAS backup using `asnodename`.
- `asnodename` can be used for clustered systems, although no specific cluster software is supported.
- You cannot use the `asnodename` option with system state components.
- When backing up an image, you must identify the volume by using the UNC name. You cannot use the drive letter.
- Tivoli Storage Manager does not track the originating system for the files. Because the `asnodename` option can cause data to be stored in the same file space name, regardless of the system from which the backup was taken, you will need to make sure you know which files should be restored to each system.

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 scorpio
```

Command line:

```
dsmc incr f:\* -subdir=yes -asnodename=scorpio
```

This command backs up the entire F disk to the server file space called \\scorpio\fs\$.

This option is not valid in interactive mode, but it can be defined in the options portion of a schedule definition.

Asrmode

Use the **asrmode** option with the **restore** 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 Windows Server 2003 clients.

This option is valid for Windows Vista, Windows7, and Windows Server 2008 clients only when these clients are running in Windows Preinstallation Environment (WinPE). Note that ASR data is not backed up, and hence cannot be restored, on 64-bit architectures with Extensible Firmware Interface (EFI).

Syntax



Parameters

No Specifies that Tivoli Storage Manager does not perform the restore operation in system ASR recovery mode. This is the default for Windows Server 2003 clients.

Yes Specifies that Tivoli Storage Manager performs the restore operation in ASR recovery mode. This is the default for Windows Vista, Windows 7, and Windows Server 2008 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 **auditlogging** 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 (dmsched.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 Tivoli Storage Manager to store audit log information.

If you specify a file name only, the file is stored in your current directory. The default is the installation directory with a file name of `dsmaudit.log`. The `dsmaudit.log` file *cannot* be a symbolic link.

Examples

Run an incremental backup with audit logging enabled.

Options file:

Store the audit log in a non-default path.

```
auditlogname c:\mypath\myaudit.log
```

Command line:

Back up a list of files using the maximum level of auditing, which would enable a separate application, such as a Perl script, to verify the results:

```
dsmc i -filelist=file.lst -auditlogging=full  
-auditlogname="c:\program files\tivoli\tsm\baclient\  
temp_audit001.log"
```

Sample output

The following is a sample execution and output file:

```
C:\Program Files\Tivoli\TSM\baclient>dsmc i  
c:\test\* -sub=yes -auditlogging=full  
IBM Tivoli Storage Manager  
Command Line Backup/Archive Client Interface  
Client Version 5, Release 5, Level 0.0  
Client date/time: 04/21/2007 15:52:25  
(c) Copyright by IBM Corporation and other(s)  
1990, 2007. All Rights Reserved.  
  
Node Name: PATMOS  
Session established with server PATMOS_5331: Windows  
Server Version 5, Release 3, Level 3.1  
Server date/time: 04/21/2007 15:52:26  
Last access: 04/21/2007 15:52:06
```

```
Incremental backup of volume 'c:\test\*'
```

```

Normal File-->          1,048,576 \\patmos\c$\test
  \dir1\file1 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
  \dir1\file2 [Sent]
Normal File-->           1,024 \\patmos\c$\test
  \dir1\file3 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
  \dir2\file1 [Sent]
Normal File-->          1,048,576 \\patmos\c$\test
  \dir2\file2 [Sent]
Normal File-->           1,024 \\patmos\c$\test
  \dir2\file3 [Sent]
Successful incremental backup of '\\patmos\c$\test\*'

```

```

Total number of objects inspected:      12
Total number of objects backed up:      6
Total number of objects updated:        0
Total number of objects rebound:       0
Total number of objects deleted:        0
Total number of objects expired:        0
Total number of objects failed:         0
Total number of subfile objects:        0
Total number of bytes transferred:     400.85 KB
Data transfer time:                     0.00 sec
Network data transfer rate:             0.00 KB/sec
Aggregate data transfer rate:           382.85 KB/sec
Objects compressed by:                  91%
Subfile objects reduced by:             0%
Elapsed processing time:                 00:00:01
ANS1900I Return code is 0.
ANS1901I Highest return code was 0.

```

The following are the audit log contents:

```

04/21/2007 15:52:25 ANS1650I Command:
  i c:\test\*
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\dir1
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\dir2
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file1
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file2
04/21/2007 15:52:26 ANS1661I Unchanged:
  \\patmos\c$\test\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir1\file3
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file1
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file2
04/21/2007 15:52:26 ANS1651I Backed Up:
  \\patmos\c$\test\dir2\file3

```

Related information

For more information about the audit logging facility refer to “Audit logging” on page 301.

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 4

Autofsrename

The `autofsrename` option renames an existing file space that is not Unicode-enabled on the Tivoli Storage Manager server so that a Unicode-enabled file space with the original name can be created for the current operation.

When you specify `autofsrename yes` in your client options file, and the server value of `autofsrename` is set to `client`, Tivoli Storage Manager generates a unique name by appending `_OLD` to the file space name you specify in the current operation. For example, Tivoli Storage Manager renames the file space `\\your-node-name\h$` to `\\your-node-name\h$_OLD`. If the new file space name is too long, the suffix replaces the last characters of the file space name, as follows:

```
\\your-node-name_0LD
```

If the new file space name already exists on the server, Tivoli Storage Manager renames the new file space `\\your-node-name_0LDx`, where `x` is a unique number.

Tivoli Storage Manager creates new Unicode-enabled file spaces that contain only the data specified in the current operation. For example, to archive files from your `H:` disk named `\\your-node\h$`, issue the following archive command:

```
arc h:\logs\*.log
```

Before the archive takes place, the server renames the file space to `\\your-node\h$_OLD`. The archive places the data specified in the current operation into the Unicode-enabled file space named `\\your-node\h$`. The new Unicode-enabled file space now contains only the `\logs` directory and the `*.log` files specified in the operation. Tivoli Storage Manager stores all subsequent full and partial incremental, selective backup, and archive data in the new Unicode-enabled file spaces.

Renamed file spaces remain on the server as stabilized file spaces. *These file spaces contain all the original data, which you can restore as long as they remain on the server.*

Note: When an existing file space is renamed during Unicode conversion, any access rules defined for the file space remain applicable to the original file space. New access rules must be defined to apply to the new Unicode file space.

After installation, perform a full incremental backup and rename all existing file spaces that are not Unicode-enabled and back up the files and directories within them under the new Unicode-enabled file spaces. This operation requires increased processing time and storage on the server.

File spaces that are not Unicode-enabled can be viewed in the character set of the locale from which Tivoli Storage Manager backed up the files. A workstation running in a different locale might be unable to view or restore from these file spaces. Unicode-enabled file spaces that are backed up in one locale are visible in all other locales, provided that the workstation has the proper fonts installed.

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

Supported Clients

This option is valid for all Windows clients. The server can define the `autofsrename` option and override the `autofsrename` setting on the client. The Tivoli Storage Manager client API does not support this option.

Options File

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

Syntax



Parameters

Yes

Specifies that Tivoli Storage Manager automatically renames all file spaces that are not Unicode-enabled in the current backup or archive operation.

No

Specifies that Tivoli Storage Manager does not rename file spaces that are not Unicode-enabled in the current backup or archive operation.

Prompt

Specifies that Tivoli Storage Manager prompts you whether to rename the file spaces that are not Unicode-enabled in the current operation. This is the default.

Considerations:

- This option applies *only* when the server sets the `autofsrename` option to client.
- When the client scheduler is running, the default behavior is to not prompt you. The next interactive session prompts you to rename the file space.

- The client prompts you *only* one time per file space. If you specify no at the prompt, the client cannot rename the file spaces later. However, the Tivoli Storage Manager administrator can rename the file spaces on the server.
- When backing up files to a file space that is not Unicode-enabled, the Unicode-enabled client skips the files and directories with names containing characters from a code page that is different from the current locale.
- If files and directories with names containing characters from a code page other than the current locale were previously backed up with a client that was not Unicode-enabled, they might be expired. The Unicode-enabled client expires these files if you do not migrate the file space to a Unicode-enabled file space. You can back up and archive these files to a Unicode-enabled file space.

Examples

Options file:

```
autofsrename yes
```

Related information

“Restore from file spaces that are not Unicode-enabled” on page 612

Backmc

The **backmc** option specifies the Tivoli Storage Manager 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 client 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 180

Casesensitiveaware

The `casesensitiveaware` option specifies whether the Windows Backup-Archive client should filter out file and directory objects which have names conflicting in case only.

NTFS volumes are case sensitive and allow case sensitive file names to be stored. Although the Windows operating system is not case sensitive, applications such as Windows Services for UNIX (SFU) exploit POSIX and allow case sensitive file names. SFU is typically included with Windows operating systems such as Windows Powered OS and Windows Storage Server 2003. These operating systems are typically deployed on hardware (for example, NAS hardware) which is acting as a dedicated file server in a heterogeneous environment.

If there are UNIX clients storing files on NTFS volumes in these Windows file server environments, 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 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 Tivoli Storage Manager Windows Backup-Archive client cannot guarantee the restore integrity of such objects. When a name casing conflict arises, the Tivoli Storage Manager Backup-Archive client can only guarantee the restore integrity of the first file in an alphabetical sort. On an ASCII-based operating system such as Windows, this means that capital letters come first alphabetically before their lower-case counterparts, so `'MySwork.xls'` would alphabetically precede `'mywork.xls'`.

In this example, if the `casesensitiveaware` option is used, only `'MyWork.xls'` 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 being 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 Tivoli Storage Manager should attempt to identify object names which differ in casing only and filter out objects which have casing conflicts and cannot be guaranteed to be restored properly.

no Specifies that the Tivoli Storage Manager should not attempt to identify object names which differ in casing only. This is the default.

Changingretries

The `changingretries` option specifies how many additional times you want the client to attempt to back up or archive a file that is in use. Use this option with the **archive**, **incremental**, and **selective** commands.

This option is applied only when *serialization*, an attribute in a management class copy group, is *shared static* or *shared dynamic*.

With *shared static* serialization, if a file is open during an operation, the operation repeats the number of times that you specify. If the file is open during each attempt, the operation does not complete.

With *shared dynamic* serialization, if a file is open during an operation, the operation repeats the number of times that you specify. The backup or archive occurs during the last attempt whether the file is open or not. Open file support can be used to back up files that are locked or in use.

Supported Clients

This option is valid for all Windows clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab, **Number of retries if file is in use** field of the Preferences editor.

Syntax

►►—CHAngingretries— *numberretries* —————►►

Parameters

numberretries

Specifies the number of times a backup or archive operation is attempted if the file is in use. The range of values is zero through 4; the default is 4.

Examples

Options file:

```
changingretries 3
```

Command line:

```
-cha=3
```

Class

The `class` option specifies whether to display a list of NAS or client objects when using the `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 client API does not support this option.

Options File

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

Syntax



Parameters

client

Specifies that you want to display a list of file spaces for a client node. This is the default.

nas

Specifies that you want to display a list of file spaces for a NAS node.

Examples

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

Command line:

```
q backup -nasnodename=nodename -class=nas
```

Clientview

The `clientview` option is available to users who have upgraded from the Tivoli Storage Manager Express backup client to the Tivoli Storage Manager enterprise client.

You must be connected to a Tivoli Storage Manager Version 5.4 or higher server to use this option. The `clientview` option allows you to choose either the express view or the standard view of the client graphical user interface (GUI).

Supported Clients

This option is valid for all Windows clients.

Options File


Place this option in the `dsm.opt` file. To switch to the Express view:

1. In the IBM Tivoli Storage Manager window, select **Edit > 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 IBM Tivoli Storage Manager window, click **Modify Settings**.
2. From the **General** tab of the Preferences editor, in the **Client View** field, click **Standard**.
3. Click **OK** to save your change.

Syntax



Parameters

standard

Specifies that the standard, or enterprise, view of the Tivoli Storage Manager client GUI should be used. The standard view contains the advanced features of the Tivoli Storage Manager Backup-Archive Client GUI. This is the default.

express

Specifies that the express view of the Tivoli Storage Manager client GUI should be used. The express view contains the same features as the Express backup client GUI.

Clusterdiskonly

The `clusterdiskonly` option specifies whether the Tivoli Storage Manager Backup-Archive client allows the backup of only clustered disks in specific environments.

Tivoli Storage Manager allows the backup of only clustered disks when running in the following environments:

- A Microsoft Cluster Server (MSCS)
- Failover Clustering on Windows Server 2008
- Failover Clustering on Windows Server 2008 R2
- VERITAS Cluster Server (VCS) environment, when the `clusternode` processing option is set to `yes`

The Tivoli Storage Manager Backup-Archive client previously only allowed backups and restores of data on clustered drives that were mounted as a drive letter.

It is becoming increasingly common to find clustered drives mounted as volume mount points. Windows 2003 Server and products such as IBM® Tivoli SANergy® allow the user to surpass the 26-drive-letter limitation by allowing volume mount points to be defined on a clustered server. Tivoli Storage Manager can protect data on cluster disks mounted as drive letters on Windows Server 2003 or Windows

Server 2008 and 2008 R2. Tivoli Storage Manager can also protect data on cluster disks mount as volume mount points with these considerations:

Windows Server 2003:

TSM cannot determine whether a volume using a volume mount point is a cluster volume. You must set `clusterdisksonly=no` and manually specify the cluster volumes in a `DOMAIN` statement.

Windows Server 2008:

The TSM client can automatically determine whether a volume using a volume mount point is a cluster volume. For Windows 2008, set `clusterdisksonly=no` to allow TSM to automatically determine which cluster volumes to process.

Note that when the `clusterdisksonly` option is specified as `yes`, the Tivoli Storage Manager Backup-Archive client continues to segregate local drives from cluster drives when evaluating the `ALL-LOCAL` domain option. When the `clusterdisksonly` option is specified as `no`, you must explicitly define the backup domains. When the `clusterdisksonly` option is specified as `no`, the Tivoli Storage Manager Backup-Archive client also bypasses enumeration of cluster resources in order to determine which resources represent cluster drives. This can be beneficial if initialization of the Backup-Archive client exhibits long delays.

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax



Parameters

Yes

Specifies that the Tivoli Storage Manager only allows the processing of cluster drives. On Windows Server 2003, the TSM client cannot determine which volumes using volume mount points are cluster drives. This option should not be used on Windows 2003 when cluster volumes using volume mount points are in use. Yes is the default.

No Specifies that the Tivoli Storage Manager allows the processing of any disk when the `clusternode` processing option is set to `yes`.

Examples

Scenario 1: Back up a node which manages the local (non-clustered) drives and the system state information

This is the node that is basically dedicated to the restoration of the physical system should there be a hardware failure. There are no clustered drives which appear as volume mount points (for example, a product like IBM Tivoli SANergy[®] is not being exploited).

Options file:

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY YES (default)
DOMAIN ALL-LOCAL (default)
EXCLUDE c:\...\file.txt
```

Scenario 1b: Back up a node which manages the local (non-clustered) drives and the system state information and bypass enumeration of cluster resources

This is a scenario similar to scenario 1 which can be deployed if the Tivoli Storage Manager Backup-Archive client takes an inappropriate amount of time during startup processing. During initialization of the Backup-Archive client, all of the cluster resources are enumerated to determine which resources represent cluster disk devices. This processing can be skipped by coding the `clusterdisksonly` parameter to `no`.

Options file:

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY NO
DOMAIN C: D: (local drives must be explicitly enumerated)
EXCLUDE c:\...\file.txt
```

Scenario 2: Back up a node which manages the local (non-clustered) drives and the system state information in a clustered environment using volume mount points as cluster resources

This node is dedicated to the restoration of the physical system if there is a hardware failure. There are clustered drives which appear as volume mount points in this environment (for example, IBM Tivoli SANergy, Windows server 2003). Also, be sure to remove any volume mount points from the incremental processing domain.

Options file:

```
CLUSTERNODE NO (default)
CLUSTERDISKSONLY YES (default)
DOMAIN ALL-LOCAL
EXCLUDE c:\...\file.txt
```

Scenario 3: Back up a node which manages the clustered drives within a cluster resource group and bypass enumeration of cluster resources

This is a scenario which can be deployed if the Tivoli Storage Manager Backup-Archive client takes an inappropriate amount of time during startup processing. During initialization of the Backup-Archive client, all of the cluster resources are enumerated to determine which resources represent cluster disk devices. This processing can be skipped by coding the `clusterdisksonly` parameter to `no`.

Options file:

```
CLUSTERNODE YES
CLUSTERDISKSONLY NO
DOMAIN f: g:
EXCLUDE f:\...\file.txt
```

Scenario 4: Back up a node which manages the clustered drives within a cluster resource group, using volume mount points as cluster resources

In this scenario, it is assumed that the node is responsible for backing up a cluster resource group that has two drives, `f:` and `f:\mnt`. There are clustered drives which appear as volume mount points (IBM Tivoli SANergy, Windows 2003 server). Ensure that you define the incremental processing domain as only the volumes within a cluster resource group and set `clusterdisksonly` to `yes` (only for Windows 2003 server). If you

have multiple cluster resource groups, assign a unique Tivoli Storage Manager node to manage each cluster resource group.

Options file (Windows Server 2008)

```
CLUSTERNODE YES
CLUSTERDISKONLY YES
DOMAIN f: f:\mnt
EXCLUDE f:\mnt\...\file.txt
```

Options file (Windows Server 2003)

```
CLUSTERNODE YES
CLUSTERDISKONLY NO
DOMAIN f: f:\mnt
EXCLUDE f:\mnt\...\file.txt
```

Table 54 shows the list of `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 <code>clusterdiskonly</code> is set to <code>yes</code> , the cluster disk map is built. This combination is used for backing up local drives. |
| yes | yes | This is the basic way to run in a cluster node to back up cluster disks; the cluster disk map is built. For Windows Server 2003 systems, this option combination will not process cluster volumes which use volume mount points. |
| yes | no | For Windows 2003 server, set these options to back up volumes mounted as mount points. For Windows Server 2008 and newer Windows systems, you only need to specify <code>clusternode=yes</code> unless you need to bypass cluster volume enumeration for performance reasons. |

Clusternode

The `clusternode` option specifies how the Tivoli Storage Manager Backup-Archive Client manages cluster drives.

Tivoli Storage Manager manages clustered drives in the following environments:

- A Microsoft Cluster Server (MSCS)
- Failover Clustering on Windows Server 2008
- Failover Clustering on Windows Server 2008 R2
- VERITAS Cluster Server (VCS)

When the `clusternode` option is set to *yes*, only shared cluster drives are available for backup and archive processing. When you set the `clusternode` option to *yes*, the node name defaults to the cluster name.

To back up local drives or Windows Server 2003 systemstate, you must set the `clusternode` option to *no*.

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

Use the `optfile` option to properly call the correct (cluster) `dsm.opt` for all Tivoli Storage Manager programs to ensure proper Tivoli Storage Manager functionality for cluster related operations. See the `OPTFILE` topic for more information.

Supported Clients

This option is valid for Windows Server 2003 and Windows Server 2008.

Options File

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

Syntax



Parameters

Yes

Specifies that you want Tivoli Storage Manager to manage cluster drives in the following environments:

- A MSCS
- Failover Clustering on Windows Server 2008
- Failover Clustering on Windows Server 2008 R2
- 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 421

Collocatebyfilespec

Use the `collocatebyfilespec` option to specify whether the Tivoli Storage Manager client uses only one server session to send objects generated from one file specification.

Setting the `collocatebyfilespec` option to *yes* attempts to eliminate interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Considerations:

- Use the `collocatebyfilespec` option only if the storage pool is going directly to tape. If you use this option going to a disk storage pool, you could affect some load balancing, and therefore, performance.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the Tivoli Storage Manager client to use only one server session to send objects generated from one file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape, unless another tape is required for more capacity. Restore performance can increase as a result.

No Specifies that the Tivoli Storage Manager client can (depending on the execution dynamics and on the setting of the `resourceutilization` option of 3 or higher), use more than one server session to send the files from one file specification. This is the default.

Backup performance might increase as a result. If the files are backed up to tape, files 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.

Commethod

The commethod 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 Version 4 or Version 6 should be used, depending on the system configuration and the results of a domain name service lookup. A valid DNS environment must be available.

NAMedpipes

The interprocess communication method that permits message data streams to pass between a client and a server. Use this communication method with a Tivoli Storage Manager server running on the same workstation as the client.

SHAREmem

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:

```
commethod tcpip
```

Use only TCP/IP Version 4

```
commethod V6Tcpip
```

Use both TCP/IP Version 4 and Version 6, depending on how the system is configured, and the results of a domain name service lookup.

Note: The `dsmc schedule` command cannot be used when both `SCHEDMODE prompt` and `commmethod V6Tcpi` are specified.

Command line:

```
-comm=tcpip  
-comm=V6Tcpi
```

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

Commrestartduration

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.

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

►►—`COMPRESSAlways`—

| |
|-----|
| Yes |
| No |

 —————►►

Parameters

Yes

File compression continues even if the file grows as a result of compression. This is the default.

No Backup-archive client objects are resent uncompressed if they grow during compression. API behavior depends on the application. Application backups might fail.

Examples

Options file:

```
compressalways yes
```

Command line:

```
-compressa=no
```

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

Compression

The compression option compresses files *before* you send them to the server.

Compressing your files reduces data storage for backup versions and archive copies of your files. It can, however, affect Tivoli Storage Manager throughput. A fast processor on a slow network connection benefits from compression, but a slow processor on a fast network connection does not.

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

The **backup image** command uses the compression option value specified in the dsm.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.

Tivoli Storage Manager backs up a sparse file as a regular file if client compression is off. Set the compression option to *yes* to enable file compression when backing up sparse files to minimize network transaction time and maximize server storage space.

If you set the compressalways option to *yes*, compression continues even if the file size increases. To stop compression if the file size grows, and resend the file uncompressed, set the compressalways option to *no*.

If you set the compression option to *yes*, you can control compression processing in the following ways:

- Use the `exclude.compression` option in your client options file (dsm.opt) to exclude specific files or groups of files from compression processing.
- 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.

Parameters

computername

Specifies the name of the computer for which you are restoring the system state.

Examples

Command line:

```
restore systemstate -asrmode=yes -computername=originaljoe
```

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

Related information

“Recovering a computer when the operating system is not working (Windows Server 2003)” on page 173

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 379

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 463 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 support only 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 **creatnewbase=migrate** option.


Supported Clients

This option is valid for the following clients:

- All Windows clients except IA 64 bit

Enter the **creatnewbase** option on the command line. Specify this option with the **snappdiff** option.

Syntax



Parameters

no Specifies that a snapshot difference incremental is run. If Tivoli Storage Manager 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 using it to run a scan-based incremental backup. Use this option to backup 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.

It is recommended that you do not set **creatnewbase=yes** for any schedule that runs a daily snapshot difference backup. Instead, create a separate, monthly schedule that has the **creatnewbase=yes** option.

IGNore

Specifies that a snapshot difference incremental backup is run when Tivoli Storage Manager 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 tells 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 Tivoli Storage Manager 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 runs. The migrate option is different from the yes option because the migrate option creates a base snapshot only when Tivoli Storage Manager 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 information

“Snapdiff” on page 463

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 67

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 Tivoli Storage Manager supports, the Web client uses the date format for US English.
2. When you change the date format and use the `schedlogretention` option to prune the schedule log, Tivoli Storage Manager removes all entries in the schedule log with a different date format when pruning the log. When you change the date format and use the `errorlogretention` option to prune the error log, Tivoli Storage Manager removes all entries in the error log with a different date when pruning the log. When changing the date format, copy the schedule log and error log if you want to preserve log entries that contain a different date format.

You can use the dateformat option with the following commands.

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query asr**
- **query backup**
- **query filespace**
- **query image**
- **query 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`, `frontime`, `today`, `fromdate`, and `pittime`.

For example, if you specify the `timeformat` option as `TIMEFORMAT 4`, the value that you provide on the `frontime` 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 client 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 Tivoli Storage Manager client or API installation directory.

Examples

Options file:

```
dedupcachepath c:\logs\dedup\
```

Command line:

Does not apply.

Related reference

“Enablededupcache” on page 354

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 client 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. Tivoli Storage Manager 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 Tivoli Storage Manager 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. Refer to the *Administrator's Guide* for details.
- The file size must be larger than 2 K.


Supported Clients

This option is valid for all clients. This option is also valid on the Tivoli Storage Manager client 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 393

"Exclude options" on page 364

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 client API does not support this option.

Syntax

▶—DELEtefiles—▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc archive c:\foo\*.c -deletefiles
dsmc rest image c: -incre -deletefiles
```

Description

The `description` option assigns or specifies a description for files when performing archive, delete archive, retrieve, query archive, or query backupset.

For example, if you want to archive a file named `budget.jan` and assign to it the description **2002 Budget for Proj 1**, you would enter:

```
dsmc archive -des="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 client API does not support this option.

Syntax

►—DEscription =— —*description*—◄

Parameters

description

Assigns a description to the file you are archiving. If you do not specify a description with the **archive** command, the default is Archive Date:*x*, where *x* is the current system date. Note that the date is always 10 characters long. If your date format uses a two digit year, there 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, and archive information depending on the command with which it is used.

Use the `detail` option with the **query mgmtclass** command to display detailed information about each management class in your active policy set. If you do not use the `detail` option, only the management class name and a brief description are displayed on the screen. If you specify the `detail` option, information about attributes in each copy group contained in each management class is displayed on the screen. A management class can contain a backup copy group, an archive copy group, both, or neither.

A Unicode-enabled file space might not display correctly if the server cannot display the Unicode name. In this case, use the file space identifier (fsID) of the file space to identify these file spaces on the server. Use the `detail` option with the **delete filespace** and **query filespace** commands to determine the fsID of a file space. The fsID also appears in the file information dialog in the 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

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 following commands:

- **delete filesystem**
- **query adobjects**
- **query archive**
- **query backup**
- **query filesystem**
- **query inclexcl**
- **query mgmtclass**
- `query systemstate`

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶—DETail—————▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query mgmtclass -detail
dsmc query filesystem -detail
dsmc query backup file1 -detail
dsmc query systemstate -detail
```

Dfsbackupmntpnt

The `dfsbackupmntpnt` option specifies whether Tivoli Storage Manager views a Microsoft Dfs junction residing on an NTFS or FAT drive as a junction or a directory.

If Tivoli Storage Manager views Microsoft Dfs junction as a junction, only the metadata of the junction is backed up or archived. The subtree under the junction point is not backed up or archived.

This option is effective only when you back up or archive a Microsoft Dfs root and is ignored when you back up or archive a Microsoft Dfs junction. To restore a Dfs tree, the root of the tree must already exist.

Supported Clients

This option is valid for Windows Server 2003 *only*.

Options File

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

Syntax



Parameters

Yes

Specifies that Tivoli Storage Manager views all Microsoft Dfs junctions as *junctions* and backs up only the name of any mounted junction it encounters during a backup operation. This is the default.

No

Specifies that Tivoli Storage Manager views all Microsoft Dfs junctions as *directories* and backs up the contents of files and subdirectories of any junction it encounters during a backup operation.

Examples

Options file:

```
dfsbackupmntpnt no
```

Command line:

Does not apply.

Related information

For more information on backing up a Dfs root, see “Microsoft Dfs file protection methods” on page 157.

Diffsnapshot

The `diffsnapshot` option controls whether Tivoli Storage Manager should create the differential snapshot when running a snapshot difference incremental backup .

If the differential snapshot is not created by Tivoli Storage Manager, 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 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 Tivoli Storage Manager are deleted by Tivoli Storage Manager 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 Tivoli Storage Manager to use the most recent snapshot that was created. Whatever method is used to create named snapshots, snapshot names differing only by case will not work properly with the `snappdiff` option. Snapshots that are created by Tivoli Storage Manager will not have the casing problem. Snapshots created by methods outside of Tivoli Storage Manager are never deleted by Tivoli Storage Manager.

Supported Clients

This option is valid for the following clients:

- 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 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 an 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 information

“[Snapdiff](#)” on page 463

“[Createnewbase](#)” on page 326

“[Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups](#)” on page 67

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 237 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 client API does not support this option.

Syntax

►►—Dirsonly—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc query backup -dirsonly c:*
```

Disablenqr

The `disablenqr` option specifies whether the Tivoli Storage Manager Backup-Archive Client can use the "no query restore" method for restoring files and directories from the server.

If you set the `disablenqr` option to *no* (the default), the client can use the "no query restore" process.

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 "no query restore" method.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

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

Syntax



Parameters

No Specifies that Tivoli Storage Manager can use the "no query restore" method. This is the default.

Yes Specifies that the client uses only the "standard restore" method. The "no query restore" method is not allowed.

Examples

Options file:

```
disablenqr yes
```

Command line

```
-disablenqr=yes
```

Diskbuffsize

The `diskbuffsize` option specifies the maximum disk I/O buffer size (in kilobytes) that the client can use when reading files. 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 is set to *diskcachemethod*. The default location is to create the disk cache file in the root of the file space being processed.

Examples

Options file:

```
diskcachelocation c:\temp  
diskcachelocation c:\tivoli\data
```

Command line:

Does not apply.

See “Include options” on page 393 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 invoke the **incremental** command without a file specification.

Tivoli Storage Manager uses the domain value in the following situations to determine which drives to process during an incremental backup:

- When you run an incremental backup 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 Tivoli Storage Manager backup-archive client GUI

There are several places where the `domain` option can be defined:

- In an options file
- On the command line, when entered with a Tivoli Storage Manager command
- In a client option set, defined on the Tivoli Storage Manager server with the **define clientopt** command
- As an option on a scheduled command, defined on the Tivoli Storage Manager 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, Tivoli Storage Manager 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 you invoke the incremental command with a file specification, Tivoli Storage Manager 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 client 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 domain objects. This parameter is the default.

The **all-local** domain includes the system state for Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7.

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 Windows Server 2003, Windows Server 2008, Windows Vista, and Windows 7 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 are 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 (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

The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab → **Domain for Backup** box of the Preferences editor.

Syntax



Parameters

domain

Defines the file systems or raw logical volumes to include in your default client image domain.

Examples

Options file:

```
domain.image d: e: f: domain.image f:\mnt\raw\rawmnt1  
f:\mnt\fs\fsmnt1
```

Command line:

Does not apply.

Domain.nas

The `domain.nas` option specifies the volumes to include in your NAS image backups.

You can specify *all-nas* to include all the mounted file systems on the NAS file server, except those you exclude with the `exclude.fs.nas` option.

Tivoli Storage Manager uses your domain for NAS image backups when you run a **backup nas** command and you do not specify which volumes to process.

When you 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, Tivoli Storage Manager adds volumes that you specify on the command line to the volumes defined in your dsm.opt file. For example, if you enter `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, Tivoli Storage Manager backs up the `vol/vol0`, `vol/vol1`, and `vol/vol2` volumes on node `nas1`.

If you set the `domain.nas` option to *all-nas* in the dsm.opt file, Tivoli Storage Manager backs up all mounted volumes on the NAS file server. When performing a backup, if you use a file specification and set the `domain.nas` option to *all-nas* in the dsm.opt file, *all-nas* takes precedence.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

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

Syntax



Parameters

domain

Defines the volumes you want to process. You cannot exclude volumes by specifying the dash (-) operator.

all-nas

Processes all mounted volumes on the NAS file server, except those you exclude with the `exclude.fs.nas` option. This is the default. If there is no `domain.nas` statement in the dsm.opt file and no volumes specified on the command line, Tivoli Storage Manager backs up all mounted volumes on the NAS server.

Examples

Options file:

```
domain.nas nas1/vol/vol0 nas1/vol/vol1
domain.nas all-nas
```

Command line:

Does not apply.

Domain.vmfile

The `domain.vmfile` option specifies the virtual machines to include in your file-level off-host backup of VMware virtual machine data operations.

This option is not valid for Hyper-V backups.

You can specify `all-windows` to include all virtual machines that are defined to the Virtual Center or to the ESX Server that is specified on the `vmhost` option. VMware file-level backups are only available for Windows-type guest operating systems from a Windows backup proxy.

The virtual machines 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 `-vmlist` option specified.

When you specify `all-windows`, `vmhost`, or `vmfolder` on the `domain.vmfile` option, you enable automatic inclusion of newly-created virtual machines on subsequent backups.


Supported Clients

This option is valid for Windows clients that are configured as an off-host backup of VMware virtual machine data proxy. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`), or on the command line. You can set this option on the **VM Backup** tab of the Preferences editor.

Syntax



Parameters

Important: Multiple keywords must be separated by a semicolon. There cannot be any spaces following the semicolons.

You can use these keywords to easily add groups of virtual machines:

- all-windows
- vmhost
- vmfolder

Then Tivoli Storage Manager detects the new virtual machines that were added to the VMware environment after the initial setup.

vmname

Defines the virtual machine host name you want to process or a list of virtual machine host names separated by commas.

vm=vmname

VM keyword specifies that the next set of values are virtual machine host names. VM is the default and it is not required.

all-windows

Process all virtual machines that are defined to the Virtual Center or the ESX server that is specified on the VMCHOST option, that have a guest operating system type of Windows. All-windows 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 VMCHOST option. The virtual machines also must be running on the ESX server that is specified by the host name. Host name can include multiple ESX servers separated by commas. Allows for automatic inclusion of newly added virtual machines.

vmfolder=foldername

Process all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the VMCHOST that is defined as residing in the VMware folder specified by the folder name. Folder name can include multiple VMware folders separated by commas. Allows for automatic inclusion of newly added virtual machines.

You can exclude individual virtual machines from the domain by specifying the exclude operator (-) before the vmname. You cannot use any domain exclusion operator (-) in front of any domain keywords.

-vm=vmname

-VM keyword specifies that the next set of values are virtual machine host names to exclude from the domain.

Examples

Options file:

```
domain.vmf file all-windows
```

File-level backup of all virtual machines running on Windows operating systems.

```
domain.vmf file "vmfolder=TSMSvt,TSMdev"
```

File-level VM backup of all virtual machines defined to VM folder TSMSvt and TSMdev.

```
domain.vmf file "vmfolder=TSMSvt;VM=testvm1,testvm2"
```

File-level VM backup of all virtual machines defined to VM folders TSMSvt and TSMdev and virtual machines testvm1 and testvm2.

```
domain.vmf file "all-windows;-vm=machine1,machine2"
```

Process all Windows virtual machines for file-level **backup vm** command, except for the virtual machines named machine1 and machine 2.

```
-domain.vmf file="vm1,vm2"
```

Process the virtual machines vm1 and vm2 during a file-level backup.

Command line:

`-domain.vmfule="vm1,vm2"`

Domain.vmfule

The `domain.vmfule` option specifies the virtual machines to include in your **backup vm** full VM image off-host backup of VMware virtual machine data operations.

You can specify ALL-VM to include 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 specified with the `domain.vmfule` option are only processed when the **backup vm** command is issued without a virtual machine, a list of virtual machines, or the `-vmlist` option specified.

Important: Full VM backups can take significantly longer to complete than file-level backups. Use multiple Tivoli Storage Manager scheduler instances to process the total set of virtual machines in subgroups that allow the backups to complete within the backup window. Each subgroup is backed up by a separate Tivoli Storage Manager scheduler instance. Each scheduler instance uses a unique client options file that defines unique `nodename` and `domain.vmfule` options. This approach enables the processing of multiple virtual machine backups in parallel from a single backup proxy.

The following example shows how VMware hosts can be used as the subgroup dividing point by specifying a single host to be processed by each scheduler instance. Each of the following options must be specified in a different client options file to be picked up by a unique Tivoli Storage Manager scheduler instance.

```
DOMAIN.VMFULL VMHOST=esxhost1.storage.usca.ibm.com
DOMAIN.VMFULL VMHOST=esxhost2.storage.usca.ibm.com
DOMAIN.VMFULL VMHOST=esxhost3.storage.usca.ibm.com
```


Supported Clients

This option is valid for Windows and Linux clients that are configured as a vStorage backup server. The server can also define this option.

Options File

Set this option in the client options, using the command line, or by using the **VM Backup** tab of the Preferences editor.

Syntax



Parameters

Important: Multiple keywords must be separated by a semicolon. There cannot be any spaces following the semicolons.

You can use these keywords to easily add groups of virtual machines:

- all-vm
- all-windows
- vmhost
- vmfolder

Then Tivoli Storage Manager detects the new virtual machines that were added to the VMware environment after the initial setup.

vmname

Defines the virtual machine host name that you want to process or a list of virtual machine host names separated by commas.

vm=vmname

VM keyword specifies that the next set of values are virtual machine host names. VM is the default and it is not required.

all-vm

Process all virtual machines that are defined to the Virtual Center or to the ESX server that is specified on the VMCHOST option.

all-windows

Process 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

Process 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 be running on the ESX server that is specified by the host name. Host name can include multiple ESX servers separated by commas.

vmfolder=foldername

Process 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 be defined as residing in the VMware folder that is specified by the folder name. Folder name can include multiple VMware folders separated by commas.

You can exclude individual virtual machines from the domain by specifying the exclude operator (-) before the vmname. You cannot use any domain exclusion operator (-) in front of any domain keyword.

-vm=vmname

-VM keyword specifies that the next set of values are virtual machine host names to exclude from the domain.

Examples

Options file:

```
Perform a full VM backup of all virtual machines.  
domain.vmfull all-vm
```

```
Perform a full VM backup of all virtual machines,
```

excluding virtual machine testvm1 and testvm2.
domain.vmfull "all-vm;-VM=testvm1,testvm2"

Perform a full VM backup of all virtual machines defined to
VM folders SanJoseTivoliLab and TucsonTivoliLab.
domain.vmfull vmfolder=SanJoseTivoliLab,TucsonTivoliLab

Perform a full VM backup of all virtual machines running on
ESX servers bovar, doomzoo, and kepler.
domain.vmfull vmhost=brovar.ibm.com,doomzoo.ibm.com,kepler.ibm.com

Command line:

Exclude virtual machines named testvm1 and testvm2 from a
full VM backup.

domain.vmfull "all-vm;-VM=testvm1,testvm2"

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 a large number of files with short 8.3 names on a large number of Windows systems, this can impact the size of the Tivoli Storage Manager server database.

Important: Consult with your Tivoli Storage Manager server administrator before using 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 need to 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 client 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 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 Tivoli Storage Manager 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 332

“Dedupcachepath” on page 331

“Dedupcachesize” on page 332

Enablelanfree

The `enablelanfree` option specifies whether to enable an available LAN-free path to a storage area network (SAN) attached storage device.

A LAN-free path allows backup, restore, archive, and retrieve processing between the Tivoli Storage Manager client and the SAN-attached storage device.

To support LAN-free data movement you must install and configure the Tivoli Storage Manager Managed System for SAN Storage Agent on the client workstation. For more information, refer to *IBM Tivoli Storage Manager for Windows Storage Agent User's Guide*.

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 615 for more information.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **General** tab → **Enable Lanfree** check box of the Preferences editor.

Syntax



Parameters

Yes

Specifies that you want to enable an available LAN-free path to a SAN-attached storage device.

No Specifies that you do not want to enable a LAN-free path to a SAN-attached storage device. This is the default.

Examples

Options file:

```
enablelanfree yes
```

Command line:

```
-enablelanfree=yes
```

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

Related information

To specify a communication protocol between the Tivoli Storage Manager client and Storage Agent, see “`Lanfreecommmethod`” on page 402.

Encryptiontype

The `encryptiontype` option allows you to choose the algorithm for data encryption.

The `encryptiontype` option allows you to use AES 128-bit data encryption, providing a stronger form of data encryption than DES 56-bit data encryption. The encryption type only affects backup and archive operations. The data that you include is stored in encrypted form, and encryption does not affect the amount of data sent or received. During restore and retrieve operations, the encrypted data 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 data encryption provides a stronger form of data encryption than DES 56-bit data encryption. This is the default.

DES56

DES 56-bit data encryption.

Examples

Options file:

```
encryptiontype des56
```

Command line:

Does not apply.

Encryptkey

The Tivoli Storage Manager client supports the option to encrypt files being backed up or archived to the Tivoli Storage Manager server. This option is enabled with the `include.encrypt` option.

All files matching the pattern on the `include.encrypt` specification 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 Tivoli Storage Manager API has an alternate way of specifying `encryptkey=generate`; the previous `enableclientencryptkey=yes` option can also be specified to request generate encryption processing.
2. The `enableclientencryptkey=yes` API option is still supported, so it is possible when using the API to specify two conflicting options. For example, `enableclientencryptkey=yes` and `encryptkey=prompt` or `encryptkey=save`.
3. When conflicting values are specified, the Tivoli Storage Manager API 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 Tivoli Storage Manager client's password file. A prompt is issued for an initial encryption key password, and after the initial prompt, the saved encryption key password in the password file is used for the backups and archives of files matching the `include.encrypt` specification. 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 is the default.

Restrictions:

- This option can only be used when `passwordaccess generate` is also specified.
- The root user or a Tivoli Storage Manager authorized user must specify the initial encryption key password.

prompt

The management of the encryption key password is provided by the user. The user is prompted for the encryption key password when the Tivoli Storage Manager client begins a backup or archive. A prompt for the same 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.

Restriction: This option can only be used by the root user or a Tivoli Storage Manager authorized user.

generate

An encryption key password is dynamically generated when the Tivoli Storage Manager client begins a backup or archive. This generated key password is used for the backups of files matching the `include.encrypt` specification. The generated key password, in an encrypted form, is kept on the Tivoli Storage Manager server. The key password is returned to the Tivoli Storage Manager client to enable the file to be decrypted on restore and retrieve operations.

Examples

Options file:

encryptkey prompt

Command line:

Does not apply.

Errorlogmax

The `errorlogmax` option specifies the maximum size of the error log, in megabytes. 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, Tivoli Storage Manager attempts to create it.

Examples

Options file:

```
errorlogname c:\temp\dsmerror.log
```

Command line:

```
-errorlogname=c:\temp\dsmerror.log
```

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

The location of the log file specified using the Client Service Configuration Utility or the client configuration wizard overrides the location specified in the client options file (dsm.opt).

Errorlogretention

The errorlogretention option specifies how many days to maintain error log entries before pruning, and whether to save the pruned entries in other files.

The error log is pruned when the first error is written to the log after a Tivoli Storage Manager session is started. If the only session you run is the client scheduler, and you run it twenty-four hours a day, the error log might not be pruned according to your expectations. Stop the session and start it again to allow the scheduler to prune the error log.

If you change from 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 `dsmerlog.pru` log file.

Syntax



Parameters

N or *days*

Specifies how long to wait before pruning the error log.

N Do not prune the error log. This permits the error log to grow indefinitely. This is the default.

days

The number of days to keep log file entries before pruning the log. The range of values is zero through 9999.

D or S

Specifies whether to save the pruned entries. Enter a space or comma to separate this parameter from the previous one.

D Discard the error log entries when you prune the log. This is the default.

S Save the error log entries when you prune the log.

The pruned entries are copied from the error log to the dsmerlog.pru file located in the same directory as the dsmerror.log file.

Examples

Options file:

Prune log entries from the dsmerror.log file that are older than 365 days and save the pruned entries in dsmerlog.pru.errorlogretention 365 S

Command line:

-errorlogr=365,S

Options file:

Prune log entries from the dsmerror.log file that are older than 365 days and do not save the pruned entries.errorlogretention 365 D

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

Determine compression and encryption processing

Tivoli Storage Manager evaluates `exclude.dir` and any other include-exclude options controlling backup and archive processing, and then the decision is made about which files undergo compression and encryption processing.

The following options determine which files undergo compression and encryption processing.

Table 56. Options for controlling compression and encryption processing

| Option | Description | Page |
|----------------------------------|---|-------------------------------|
| Compression processing | | |
| <code>exclude.compression</code> | Excludes files from compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives. | "Exclude options" on page 364 |
| <code>include.compression</code> | Includes files for compression processing if <code>compression=yes</code> is specified. This option applies to backups and archives. | "Include options" on page 393 |
| Encryption processing | | |
| <code>exclude.encrypt</code> | Excludes files from encryption processing. | "Exclude options" on page 364 |

Table 56. Options for controlling compression and encryption processing (continued)

| Option | Description | Page |
|---|---|-------------------------------|
| 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 393 |
| 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" |
| 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 393 |

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 inclexcl option.
4. As with other include-exclude statements, you can use the inclexcl option to specify a file that can be in Unicode format, containing exclude statements with file names in Unicode.

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, Tivoli Storage Manager 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 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.

exclude.systemservice

Excludes individual system services components from system state backup. Input can be the keyword or component name to be excluded from processing. If you specify the component name and there are spaces, enclose the name in quotation marks. Table 57 shows valid system services components and their corresponding keywords.

If you are specifying a component that is not listed in Table 57, you can specify the component name. This option is valid for Windows Server 2003 and Windows Vista *only*.

Table 57. 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.systemservice eventlog
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 75 for examples of statements using UNC file names.

See “System files to exclude” on page 74 for a list of files that you should always exclude.

“Inclxcl” on page 392

See “Include and exclude groups of files with wildcard characters” on page 76 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 Tivoli Storage Manager compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.
- You must set the `compression` option to *yes* to enable compression processing. If you do not specify the `compression` option or you set the `compression` option to *no*, Tivoli Storage Manager does not perform compression processing.

If you set the compression option to *yes* and no `exclude.compression` statements exist, Tivoli Storage Manager considers all files for compression processing.

- Tivoli Storage Manager processes `exclude.dir` and other include-exclude statements first. Tivoli Storage Manager then considers any `exclude.compression` statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.compression c:\test\file.txt
include c:\test\file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the `c:\test\file.txt` file is a candidate for backup, but is not a candidate for compression processing.

- Include-exclude compression processing is valid for backup and archive processing *only*. The `exclude.compression` option does not affect whether files are excluded from backup or archive processing, only whether they are excluded from compression processing.

Related reference

“Compression” on page 323

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*, Tivoli Storage Manager does not perform adaptive subfile backup processing.

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

- Tivoli Storage Manager processes `exclude.dir` and other include-exclude statements first. Tivoli Storage Manager then considers any `exclude.subfile` statements. For example, consider the following include-exclude list:

```
exclude c:\test\*.*
exclude.subfile c:\test\file.txt
include c:\test\file.txt
```

Tivoli Storage Manager examines the statements (reading from bottom to top) and determines that the `c:\test\file.txt` file is a candidate for backup, but is not a candidate for adaptive subfile backup.

- Include-exclude adaptive subfile backup processing is valid for backup and restore processing *only*. The `exclude.subfile` option does not affect whether files are excluded from backup or restore processing, only whether they are excluded from adaptive subfile backup processing.
- As with other include-exclude statements, you can specify `exclude.subfile` statements in Unicode using the `inclxcl` option.

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 392

“Subfilebackup” on page 484

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 Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the exclude statement applies. If you do not specify a NAS node name, the file system identified applies to all NAS file servers.
- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: `/vol/vol0`.

For example, to exclude `/vol/vol1` from backup services on all NAS nodes, specify the following exclude statement:

```
exclude.fs.nas */vol/vol1
```

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

►►—FBBranch=*branch_ID*—►►

Parameters

branch_ID

Specifies the FastBack server branch ID. The value is part of the disaster recovery configuration of the FastBack server.

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 Tivoli Storage Manager backup-archive client. To back up or archive volumes from a Linux FastBack client, use the Linux Tivoli Storage Manager 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 Tivoli Storage Manager by the Windows Tivoli Storage Manager 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



Notes:

- 1 Specify up to 10 FastBack policy names.

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 Tivoli Storage Manager 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 Tivoli Storage Manager backup-archive client is installed on a DR Hub server or a FastBack server workstation. When the Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager 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  
-fbrepositon=\myFbServerMachine.company.com\Rep
```

Command line:

The Tivoli Storage Manager 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  
-fbrepositon=\myDrHubMachine.company.com\Rep  
-fbbranch=branch1
```

Fbvolumename

Use the `fbvolumename` option with the **backup fastback** or **archive fastback** commands.

The `fbvolumename` option is the name of one or more comma-separated Tivoli Storage Manager FastBack volumes to back up or archive from the backup proxy. Values for the `fbvolumename` option are not valid if more than one FastBack client is specified in the `fbclientname` option.

If you specify multiple FastBack clients in the `fbclientname` option, you cannot specify values for the `fbvolumename` option.

Supported Clients

This option is valid for 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



Notes:

- 1 Specify up to 10 FastBack volume names.

Parameters

volume_name

Specifies the name of the Tivoli Storage Manager FastBack volumes. You can specify up to 10 FastBack volume names.

Important:

When specifying the **archive fastback** or **backup fastback** command:

1. *At least* one FBpolicyName is always *required*.
2. You can specify up to 10 values for FBPolicyName, if no values are specified for both FBClientName *and* FBVolumeName.
3. When you specify a FBClientName value, there must be *only one* value for FBPolicyName.
4. You can specify up to 10 values for FBClientName if *only one* PolicyName is specified, and *no* values for FBVolumeName are specified.
5. When you specify the FBVolumeName option, you can have *only one* FBPolicy, and *only one* FBClientName specified. You must specify exactly one FBClientName. It cannot be omitted.
6. You can specify multiple FBVolumeNames if condition 5 is satisfied.

Examples

Command line:

```
dsmc backup fastback -fbpolicyname=Policy1 -fbclientname=client1  
-fbvolumename=c:,f: -fbserver=myFbServer  
-fbrepositlocation=\\myFbServer.company.com\REP
```

Backs up volumes c: and f: from FastBack client Client1, found in policy Policy1.

Command line:

```
dsmc archive fastback -fbpolicyname=Policy1 -fbclientname=client1  
-fbvolumename=c:,f: -fbserver=myFbServer  
-fbrepositlocation=\\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 Tivoli Storage Manager 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, Tivoli Storage Manager ignores all other file specifications on the command line.

The files (entries) listed in the `filelist` must adhere to the following rules:

- Each entry must be a fully-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).
- Paths must not contain wildcard characters. Do not include * or ? in a path.
- 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 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"
```

You can use the `filelist` option during an open file support operation. In this case, Tivoli Storage Manager processes the entries in the filelist from the virtual volume instead of the real volume.

If an entry in the filelist indicates a directory, only that directory 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. Tivoli Storage Manager skips any entries in the filelist that are not valid files or directories. Tivoli Storage Manager logs errors and processing continues to the next entry.

Use file specifications with the **restore** and **retrieve** commands to denote the destination for the restored filelist entries. For example, in the 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: Tivoli Storage Manager might back up a directory twice if the following conditions exist:

- The filelist contains an entry for the directory
- The filelist contains one or more entries for files within that directory
- No backup of the directory exists

For example, your filelist includes the entries `c:\dir0\myfile` and `c:\dir0`. If the `\dir0` directory does not exist on the server, the `c:\dir0` directory is sent to the server a second time.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►►—FILEList =— —*filelistspec*—

Parameters

filelistspec

Specifies the location and name of the file that contains the list of files to process with the command.

Note: When you specify the `filelist` option on the command line, the `subdir` option is ignored.

Examples

Command line:

```
sel -filelist=c:\avi\filelist.txt
```

Filename

Use the `filename` option with the **query systeminfo** command to specify a file name in which to store information.

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 - 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

```
►► 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 325

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 client API does not support this option.

Syntax

▶▶—FILESOnly—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incremental -filesonly
```

Fromdate

Use the `fromdate` option with the `fromtime` option to specify a date and time from which you want to search for backups or archives during a restore, retrieve, or query operation.

Files that were backed up or archived before this date and time are not included, although older directories might be included, if necessary, to restore or retrieve the files.

Use the `fromdate` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**

- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶ FROMDate =— *date*—————▶▶

Parameters

date

Specifies the date from which you want to search for backup copies or archived files. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc query backup -fromdate=12/11/2003 c:\Windows\Program
Files\*.exe
```

Fromnode

The `fromnode` option permits one node to perform commands for another node. A user on another node must use the **set access** command to permit you to query, restore, or retrieve files 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 Tivoli Storage Manager 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 Tivoli Storage Manager 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.

Tivoli Storage Manager ignores this option if you do not specify the `fromdate` option.

Use the `fromtime` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►► FROMTime =— —*time*—————►►

Parameters

time

Specifies a beginning time on a specific date from which you want to search for backed up or archived files. If you do not specify a time, the time defaults to 00:00:00. Specify the time in the format you selected with the `timeformat` option.

When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `totime` options.

Examples

Command line:

```
dsmc q b -timeformat=4 -fromt=11:59AM -fromd=06/30/2003 -tot=11:59PM  
-tod=06/30/2003 c:\*
```

Frsprimaryrestore

The `frsprimaryrestore` option specifies whether Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost.

For example, in a disaster recovery situation when all systems containing the SYSVOL system object are lost, it is necessary to restore the first SYSVOL in this manner. In this case, set the `frsprimaryrestore` option to *yes* in your client options file (`dsm.opt`). If you are restoring a single system into an existing active directory tree, this option is not necessary.

Supported Clients

This option is valid for Windows Server 2003 clients only.

Options File

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

Syntax



Parameters

Yes

Specifies that Tivoli Storage Manager allows the primary restoration of the SYSVOL system object if all replication partners have been lost.

No Specifies that Tivoli Storage Manager *does not* allow the primary restoration of the SYSVOL system object if all replication partners have been lost. This is the default.

Examples

Options file:

```
frsprimaryrestore yes
```

Command line:

Does not apply.

Groupname

Use the `groupname` option with the **backup group** command to specify the name for a group. You can only perform operations on new groups or the current active version of the group.

Supported Clients

This option is valid for all Windows clients.

Syntax

►►—GROUPName =— —name—►►

Parameters

name

Specifies the name of the group which 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

►►—GUITREEViewafterbackup No Yes—►►

Parameters

No Returns you to the Tivoli Storage Manager main window after a successful operation completes. This is the default.

Yes

Returns you to the Backup, Restore, Archive, or Retrieve window after a successful operation completes.

Examples

Options file:

```
gui treeview after backup 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 client 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 → **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:

```
hsmreparse tag 22
```

Specify an HSM reparse tag in hexadecimal format:

```
hsmreparse tag 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 client 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 364

“Include options” on page 393

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 client API does not support this option.

Syntax

▶—IFNewer—————▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore -ifnewer d:\logs\*.log
```

Imagegapsize

Use the `imagegapsize` option with the **backup image** command, in the 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 k in size is not backed up. Gaps that are exactly 10 k is backed up. Empty regions that are exactly 10 k and that are smaller than 10 k is backed up, even though they do not contain data. However, an empty region that is smaller than 10 k 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 client API does not support this option.

Options File

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

Syntax

▶▶—IMAGEgapsize— *size*—————▶▶

Parameters

size

Specifies the minimum size of empty regions in a formatted logical volume that should be skipped during an image backup. You can specify k (kilobytes) m (megabytes) or g (gigabytes) qualifiers with the value. Without a qualifier, the value is interpreted in kilobytes. Valid values are 0 through 4294967295 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 will *not* be backed up. For LAN-based and LAN-free image backup the default value is 32 k.

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 client API does not support this option.

Syntax

▶—IMAGETOfile—▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image d: e:\diskD.img -imagnetofile
```

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 client API does not support this option.

Syntax

▶▶—INActive—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore -inactive c:\id\projecta\ -pick
```

Inclexc1

The `inclexc1` option specifies the path and file name of an include-exclude options file.

Multiple `inclexc1` 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 `inclexc1` 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

▶▶—INCLExc1— *filespec*—▶▶

Parameters

filespec

Specifies the path and file name of *one* include-exclude options file.

Examples

Options file:

```
inclexc1 c:\dsm\backup.exc1
```

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 69.

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.

1. Open TextEdit. Click **Format > Make PlainText**.
2. Enter your include and exclude statements.
3. Click **File** and then click **Save As**.
4. From **PlainText Encoding**, select **Unicode (UTF-8)** or **Unicode (UTF-16)**, specify the file and target directory, and then save the file. Do not add the .txt extension.
5. Place an `incl excl` option specifying the include-exclude file you just created in your `dsm.sys` file.
6. Restart the Tivoli Storage Manager client.

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 Tivoli Storage Manager 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, Tivoli Storage Manager uses the default management class in the active policy set of your policy domain. Use the **query mgmtclass** command to display information about the management classes available in your active policy set.

You can include remotely accessed files by specifying Universal Naming Convention (UNC) names in your include statement.

Remember: Tivoli Storage Manager compares the files it processes against the patterns specified in the include-exclude statements, reading from the bottom to the top of the options file.

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, **Define Include-Exclude Options** section of the Preferences editor.

Syntax



include, include.backup, include.file

Use these options to include files or assign management classes for backup processing.

The `include` option affects archive and backup processing. If you want to assign different management classes for archive and backup processing, always specify `include.archive` and `include.backup` with their own management classes. In 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, Tivoli Storage Manager does not perform encryption processing.

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 does not occur.

include.fs

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

To control an open file support operation with LVSA, you can specify these additional options in your `dsm.opt` file or as values of the `include.fs` option: `snapshotproviderfs`, `presnapshotcmd`, `postsnapshotcmd`. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: `snapshotcachelocation`, `snapshotcachesize`, `snapshotfsidleretries`, `snapshotfsidlewait`.

To control how Tivoli Storage Manager processes your file space for incremental backup, you can specify these additional options in your `dsm.opt` file or as values of the `include.fs` option: `diskcachelocation` and `memoryefficientbackup`. If these options appear both in the options file and an `include.fs` option, the `include.fs` values are used for the specified file space in place of any values in an option file or on the command line.

Examples:

Options file or include/exclude file:

```
include.fs d: memoryefficientbackup=diskcachem
diskcachelocation=e:\temp
include.fs e: memoryefficientbackup=diskcachem
diskcachelocation=c:\temp
```

include.fs.nas

Use the `include.fs.nas` option to bind a management class to Network Attached Storage (NAS) file systems. You can also specify whether Tivoli Storage Manager saves Table of Contents (TOC) information during a NAS file system image backup, using the `toc` option with the `include.fs.nas` option in your client options file (`dsm.opt`).

include.image

Includes a file space or logical volume, or assigns a management class when used with the **backup image** command. The **backup image** command ignores all other include options.

By default, Tivoli Storage Manager performs an offline image backup. To enable and control an online image operation, you can specify these options in your `dsm.opt` file or as values of the `include.image` option: `snapshotproviderimage`, `presnapshotcmd`, `postsnapshotcmd`. Also, if you selected the LVSA as your snapshot provider, the following additional options can be set: `snapshotcachelocation`, `snapshotcachesize`, `snapshotfsidleretries`, `snapshotfsidlewait`

include.subfile

Includes files for adaptive subfile backup processing. This option does not apply to archive processing. This option is valid for all Windows clients.

include.systemstate

Assigns management classes for backup of the system state. The default is to bind the system state object to the default management class.

During system state backup, a progressive incremental backup is used for the System Writer files in the system state. If you want to run a full backup of all system state data, you can specify `include.systemstate mc_name` in the client options file (`dsm.opt`), where `mc_name` is the name of the management class with copy group mode absolute.

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 using `include.systemstate` the only valid pattern is ALL (all types of system state objects). By default, Tivoli Storage Manager binds all objects to the default management class.

optional_parameter

mgmtclassname

Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

Use the following command to bind a management class to a group:

```
include c:\myvfs\mygroup mymgmtclass
```

Table 58. Other optional parameters

| optional_parameter | Use with option |
|--|-----------------|
| ieobjtype "Ieobjtype" on page 388 | include.dedup |
| memoryefficientbackup "Memoryefficientbackup" on page 411 | include.fs |
| diskcachelocation "Diskcachelocation" on page 342 | include.fs |
| postsnapshotcmd "Postsnapshotcmd" on page 428 | include.image |
| presnapshotcmd "Presnapshotcmd" on page 434 | include.image |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | include.image |
| snapshotcachesize "Snapshotcachesize" on page 468 | include.image |
| snapshotfsidleretries "Snapshotfsidleretries" on page 469 | include.image |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | include.image |
| snapshotproviderfs "Snapshotproviderfs" on page 472 | include.image |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | include.image |

Examples

Options file:

```
include c:\proj\text\devel.*
include c:\proj\text\* textfiles
include ?:* managall
include WAS_ND_NDNODE mgmtclass
include WAS_APPNODE mgmtclass
include.backup c:\win98\system\* mybackupclass
include.archive c:\win98\system\* myarchiveclass
include.encrypt c:\win98\proj\gordon\*
include.compress c:\test\file.txt
include.subfile c:\test\file.txt
include.image h: MGMTCLASSNAME
    snapshotproviderimage=lvsa
    snapshotcachesize=40
include.image x:
    snapshotproviderimage=none
include.image y:
    snapshotproviderimage=vss
include.image z: MGMTCLASSNAME
    snapshotproviderimage=none
include.fs c:
    snapshotproviderfs=lvsa snapshotcachesize=40
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
```

Command line:

Does not apply.

Related information

“Exclude files with UNC names” on page 75

“Configuring Open File Support” on page 66

See Chapter 12, “Processing options,” on page 269 for more information about the `snapshotcachelocation`, `snapshotcachesize`, `snapshotfsidleretries`, and `snapshotfsidlewait` options.

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*, Tivoli Storage Manager 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*, Tivoli Storage Manager does not perform adaptive subfile backup processing.
- Tivoli Storage Manager processes `exclude.dir` and other include-exclude statements first. Tivoli Storage Manager then considers any `include.compression`, `include.encrypt` and `include.subfile` statements. For example, consider the following include-exclude list:

```
exclude c:\test\file.txt
include.compression c:\test\file.txt
include.encrypt c:\test\file.txt
include.subfile c:\test\file.txt
```

Tivoli Storage Manager examines the `exclude c:\test\file.txt` statement first and determines that `c:\test\file.txt` is excluded from backup processing and is, therefore, not a candidate for compression, encryption, or adaptive subfile backup processing.

- Include-exclude compression and encryption processing is valid for backup and archive processing *only*.
- Include-exclude adaptive subfile backup processing is valid for backup and restore processing *only*.
- As with other include-exclude statements, you can use the `inclxcl` option to specify a file that is in Unicode format, which contains `include.compression`, `include.encrypt` and `include.subfile` statements specifying Unicode files. See “`Inclxcl`” on page 392 for more information.

Related reference

“Compression” on page 323

“Subfilebackup” on page 484

Processing NAS file systems

Use the `include.fs.nas` option to bind a management class to NAS file systems and to control if Table of Contents information is saved for the file system backup.

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 Tivoli Storage Manager. You can prefix the NAS node name to the file specification to specify the file server to which the include statement applies. If you do not specify a NAS node name, the file system you specify applies to all NAS file servers.
- Regardless of the client platform, NAS file system specifications use the forward slash (/) separator, as in this example: /vol/vol0.
- NAS file system designations on the command line require brace delimiters {} around the file system names, such as: {/vol/vol0}. Do not use brace delimiters in the option file.

Use the following syntax:

► *pattern*— *mgmtclassname*— *toc=value* ◄

Parameters

pattern

Specifies the objects to include for backup services, to assign a specific management class, or to control TOC creation. You can use wildcards in the pattern.

mgmtclassname

Specifies the name of the management class to assign to the objects. If a management class is not specified, the default management class is used.

toc=value

See “Toc” on page 496 for more information.

For example, to assign a management class to the /vol/vol1 file system of a NAS node called netappsj, specify the following include statement:

```
include.fs.nas netappsj/vol/vol1 nasMgmtClass toc=yes
```

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 Tivoli Storage Manager client, but before the incremental-by-date backup completes, are not backed up and will not be backed up in future incremental-by-date backups, unless the files are modified again. For this reason, a regular incremental backup should be run periodically, without specifying the `incrbydate` option.

An incremental-by-date updates the date and time of the last incremental at the server. If you perform an incremental-by-date on only part of a file system, the date of the last full incremental is not updated and the next incremental-by-date backs up these files again.

Both full incrementals and incrementals-by-date back up new and changed files. An incremental-by-date takes less time to process than a full incremental and requires less memory. However, unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files because:

- It does not expire backup versions of files that are deleted from the workstation.

- It does not rebind backup versions to a new management class if the management class has changed.
- It does not back up files with attributes that have changed, such as NTFS security information, unless the modification dates and times have also changed.
- It ignores the copy group frequency attribute of management classes.

Note: If you have limited time during the week to perform backups, but extra time on weekends, you can maintain current server storage of your workstation files by performing an incremental backup with the `incrbydate` option on weekdays and a full incremental backup on weekends.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—INCRbydate—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:
`dsmc incremental -incrbydate`

Incremental

Use the `incremental` option with the **restore image** command to ensure that any changes that were made to the base image are also applied to the restored image.

If you also use the `deletefiles` option, changes include the deletion of files and directories that were in the original image but later deleted from the workstation.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—INCRemental—▶▶

Examples

Command line:
`res i d: -incremental`

Incrthreshold

The `incrthreshold` option specifies the threshold value for the number of directories in any journaled file space that might have active objects on the server, but no equivalent object on the workstation.

When a Windows client deletes a file or directory with a long name, it sometimes reports this using a compressed name. After the object is deleted, the compressed name might be reused and the deletion notice can no longer identify a unique object. During a journaled incremental backup of a file space, this can result in the *no active version* response from the server resulting in an unsuccessful expire for an object.

The `incrthreshold` option allows you to specify what to do when this condition arises:

- If you set the `incrthreshold` option to `0` (the default), Tivoli Storage Manager takes no action. The primary consequence is that, during a restore of such a directory, these objects might be inadvertently restored. When the next non-journaled incremental is run on this directory, Tivoli Storage Manager expires all objects in the directory that exist on the server but not on the workstation.
- If you specify a value greater than zero, Tivoli Storage Manager saves 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`, Tivoli Storage Manager performs a full incremental backup of these directories whenever a *no active version* response is received during a full file space journaled incremental backup.

Supported Clients

This option is for all Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab → **Threshold for non-journal incremental backup** field of the Preferences editor.

Syntax

►►—`INCRThreshold—numberdirectories`—◄◄

Parameters

numberdirectories

Specifies the threshold value for the number of directories in any journaled file space that might contain active files that should be expired. When this threshold is reached during a full file space journaled incremental, the client initiates an incremental backup on each such directory at the completion of the journaled backup. The range of values is 0 through 2,000,000,000; the default is 0.

Examples

Options file:

```
incrthreshold 1
```

Command line:

```
-increthreshold=1
```

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

Related information

See “Incremental” on page 564 for more information about journaled backups.

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\jnlService1
```

Syntax

►—JOURNALPipe—pipename—►

Parameters

pipename

Specify the name of the pipe the client attaches to when performing a journal-based backup.

Examples

Options file:

```
JOURNALP jbbpipe1
```

Command line:

This option cannot be set on the command line.

Lanfrecommethod

The `lanfrecommethod` 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 `lanfrecommethod 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 your Tivoli Storage Manager client:

TCPip

The Transmission Control Protocol/Internet Protocol (TCP/IP) communication method.

Use the `lanfreecpport` option to specify the TCP/IP port number where the Storage Agent is listening.

V6Tcpip

Indicates that either TCP/IP Version 4 or Version 6 should be used, depending on the system configuration and results of a domain name service lookup. The only time this is not true is when **dsmc schedule** is used and `shedmode` 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 `lanfreecpport` option if you want to use the `NAMedpipes` communication method for LAN-free communication.

SHAREmem

Use the shared memory communication method when the client and Storage Agent are running on the same system. Shared memory provides better performance than the TCP/IP protocol.

Examples

Options file:

```
lanfreecommmethod tcp
```

Use only TCP/IP Version 4

```
lanfreecommmethod V6Tcpip
```

Use both TCP/IP Version 4 or Version 6, depending on how the system is configured and the results of a domain name service lookup.

Command line:

```
-lanfreec=tcp
```

```
-lanfreec=V6Tcpip
```

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

Related information

“Lanfreeshmport”

“Lanfreetcport”

Lanfreeshmport

Use the `lanfreeshmport` option when `lanfreecommmethod=SHAREdmem` is specified for communication between the Tivoli Storage Manager client and Storage Agent. This enables processing between the client and the SAN-attached storage device.

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; the default is 1.

Examples

Options file:

```
lanfrees 1520
```

Command line:

```
-lanfrees=1520
```

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

Related information

“Lanfreecommmethod” on page 402

Lanfreetcport

The `lanfreetcport` option specifies the TCP/IP port number where the Tivoli Storage Manager Storage Agent is listening.

Use this option when you specify `lanfreecommmethod=TCPIP` for communication between the Tivoli Storage Manager client and Storage Agent. Do not specify the `lanfreetcport` option if you want to use the NAMEDpipes communication method for LAN-free communication.

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 `lanfreetcpport` value must match Storage Agent `tcpport` value for communications with the Storage Agent (virtual server). The client `tcpport` value must match the server `tcpport` value for communications with the actual server.

Examples

Options file:

```
lanfreetcp 1520
```

Command line:

```
-lanfreetcp=1520
```

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

Related information

“Lanfreecommmethod” on page 402

Lanfreetcpserveraddress

The `lanfreetcpserveraddress` option specifies the TCP/IP address for a Tivoli Storage Manager Storage Agent.

Use this option when you specify `lanfreecommmethod=TCPIP` or `V6TCPIP` for communication between the Tivoli Storage Manager client and Storage Agent.

Overriding the default for this option is useful when configuring LAN-free in an environment where the client and storage agent are running on different systems. You can obtain this Storage Agent address from your administrator.

Supported Clients

This option is valid for all Windows clients, except the 64-bit Windows Server 2003.

Options File

Place this option in the client system-options file.

Syntax

▶—LANFREETCPServeraddress— *stagent_address*—▶

Parameters

stagent_address

Specifies a 1 to 64 character TCP/IP address for a server. Specify a TCP/IP domain name or a numeric IP address. The numeric IP address can be either a TCP/IP v4 or TCP/IP v6 address. The default value is 127.0.0.1 (localhost).

Examples

Options file:

```
LANFREETCPServeraddress stagent.bayamon.ibm.com
```

```
LANFREETCPServeraddress 192.168.1.50
```

Command line:

Does not apply.

Language

The language option specifies the national language in which to present client messages.

You can use US English (ENU) with all clients.

The Tivoli Storage Manager client automatically detects the language of the system locale and displays Tivoli Storage Manager for that language. For example, a supported operating system on a Windows system running a French locale displays Tivoli Storage Manager in French by default, without specifying the language option; provided that the French language pack is installed. If Tivoli Storage Manager cannot load the French message catalog, it defaults to the US English language pack. If the client is running on an unsupported language/locale combination, such as French/Canada or Spanish/Mexico, Tivoli Storage Manager defaults to US English. You can override the default language by specifying the language option.

Note: The language option does not affect the Web client. The Web client displays in the language associated with the locale of the browser. If the browser is running in a locale that Tivoli Storage Manager does not support, the Web client displays in US English.

Supported Clients

This option is valid for all clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Regional Settings** tab, **Language** drop-down list of the Preferences editor.

Syntax

►—LANGUage— *language*—►

Parameters

language

Specifies the language you want to use. The available languages include:

- ENU (English, United States)
- PTB (Brazilian Portuguese)
- CHS (Chinese, Simplified)
- CHT (Chinese, Traditional)
- FRA (Standard French)
- DEU (Standard German)
- ITA (Standard Italian)
- JPN (Japanese)
- KOR (Korean)
- ESP (Standard Spanish)
- CSY (Czech)
- HUN (Hungarian)
- PLK (Polish)
- RUS (Russian)

Examples

Options file:

```
language enu
```

Command line:

Does not apply.

Latest

Use the `latest` option 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 client API does not support this option.

Syntax

►—LATest—►

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore c:\devel\projecta\ -latest
```

Localbackupset

The `localbackupset` option specifies whether the Tivoli Storage Manager GUI bypasses initial logon with the Tivoli Storage Manager server to restore a local backup set on a standalone workstation.

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 615

Managedservices

The `managedservices` option specifies whether the Tivoli Storage Manager client acceptor service manages the scheduler, the Web client, or both.

Restriction: You cannot use the `dsmcad` for scheduling when you set the `sessioninitiation` option to **serveronly**.

The client acceptor daemon serves as an external timer for the scheduler. When the scheduler is started, it queries the server for the next scheduled event. The event is either executed immediately or the scheduler exits. The client acceptor daemon restarts the scheduler when it is time to execute the scheduled event.

Note:

1. If you set the `schedmode` option to **prompt**, the server prompts the client acceptor daemon when it is time to run the schedule. The scheduler connects to and disconnects from the server when the client acceptor daemon is first started.

The `dsmc schedule` command cannot be used when both `schedmode prompt` and `commethod V6Tcpip` are specified.

2. 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 client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **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 `managementservices` option in your client options file (`dsm.opt`).

Task Specify that the client acceptor daemon manages only the Web client.

```
managementservices webclient
```

Task Specify that the client acceptor daemon manages only the scheduler.

```
managementservices schedule
```

Task Specify that the client acceptor daemon manages both the Web client and the scheduler.

```
managementservices schedule webclient
```

Note: The order in which these values are specified is not important.

Command line:

Does not apply.

Related information

“Passwordaccess” on page 423

See “Configuring the client scheduler” on page 24 for instructions to set up the client acceptor daemon to manage the scheduler.

“Sessioninitiation” on page 457

Maxcmdretries

The `maxcmdretries` option specifies the maximum number of times the client scheduler (on your workstation) attempts to process a scheduled command that fails.

The command retry starts *only if* the client scheduler has not yet backed up a file, never connected to the server, or failed before backing up a file. This option is only used when the scheduler is running.

Your Tivoli Storage Manager administrator can also set this option. If your Tivoli Storage Manager administrator specifies a value for this option, that value overrides what you specify in the client options file *after* your client node successfully contacts the server.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, **Maximum command retries** field of the Preferences editor.

Syntax

▶—MAXCMDRetries— —*maxcmdretries*—▶

Parameters

maxcmdretries

Specifies the number of times the client scheduler can attempt to process a scheduled command that fails. The range of values is zero through 9999; the default is 2.

Examples

Options file:

```
maxcmdr 4
```

Command line:

```
-maxcmdretries=4
```

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

Memoryefficientbackup

The `memoryefficientbackup` option specifies the memory-conserving algorithm to use for processing full file space backups.

One method backs up one directory at a time, using less memory. The other method uses much less memory, but requires more disk space.

Use the `memoryefficientbackup` option with the **incremental** command when your workstation is memory constrained. You can also use this option as a parameter to the `include.fs` option in order to select the algorithm Tivoli Storage Manager uses on a per-filespace basis.

Use `memoryefficientbackup=diskcachemethod` for any file space that has too many files for Tivoli Storage Manager to complete the incremental backup with either the default setting, `memoryefficientbackup=no`, or with `memoryefficientbackup=yes`. The disk cache file created by the initial disk cache incremental backup can require up to 5 gigabytes of disk space for each million files or directories being backed up.

The 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, **Use memory-saving algorithm** check box of the Preferences editor.

Syntax



Parameters

No Your client node uses the faster, more memory-intensive method when processing incremental backups. This is the default.

Yes

Your client node uses the method that requires less memory when processing incremental backups.

Diskcachemethod

Your client node uses the method that requires much less memory but more disk space when processing incremental backups for full file systems.

Examples

Options file:

```
memoryefficientbackup yes
memoryefficientbackup diskcachem
```

Command line:

```
-memoryef=no
```

Related information

“Include options” on page 393

Mode

Use the mode option with the **backup group**, **backup image**, **backup nas** and **backup vm**, commands to specify whether you want to perform a full or differential group backup containing a list of files from one or more file spaces.

You can use the mode option with the following 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

To specify whether to perform a full or incremental backup of a VMware virtual machine when `vmbackuptype=fullvm` and when you have installed IBM Tivoli Storage Manager for Virtual Environments.

The mode option has no effect on a raw logical device backup.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

Syntax

For image backup of client file systems (all Windows clients).



For image backup of NAS file systems (all Windows clients).



For group backups (all Windows clients).



For full VM backup of VMware virtual machines when you have installed IBM Tivoli Storage Manager for Virtual Environments. .



Parameters

selective

Specifies that you want to perform a full (selective) image backup. This is the default for image backup of client file systems.

incremental

Specifies that you want to back up only new and changed files after the most recent image backup (full or incremental). If an image backup has not yet been made, then the first backup is a full image backup, regardless of the mode option.

Deleted files are not marked inactive on the server.

full

Specifies that you want to perform a full backup of NAS, virtual machines, or group objects. This is the default for group backups.

differential

This is the default for NAS objects. Specifies that you want to perform a NAS or group backup of files that changed since the last full backup. If there is no copy of a full image stored on the Tivoli Storage Manager server, a full backup occurs. If a full image exists, whether it is restorable, or expired and being maintained because of dependent differential images, specifying **MODE=differential** sends a differential image backup. If a full image is sent during a differential backup, it is reflected as a full image using the **QUERY NASBACKUP** server command.

A full image can be eligible for expiration based on versioning or retention (verexists reextra), but still be maintained on the Tivoli Storage Manager server to allow for restoring dependent differential images. A full image that is eligible for expiration *cannot* be selected for restore, so it is not displayed using the **QUERY NASBACKUP** server command. The differential image backups that depend on an "expired" full image can be restored.

Examples

Task Perform a full VM incremental backup of a virtual machine named **vm1**.

```
dsmc backup vm vm1 -mode=incremental
```

Task Perform the NAS image backup of the entire file system.

```
dsmc backup nas -mode=full -nasnodename=nas1 {/vol/vol0} {/vol/vol1}
```

Task Back up the 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=incremental
```

Task Perform a full backup of all the files in filelist c:\dir1\filelist1 to the virtual file space name \virtfs containing the group leader c:\group1 file.

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

Related reference

“Backup VM” on page 549

“Backup Group” on page 538

“Backup Image” on page 540

“Backup NAS” on page 545

Monitor

The **monitor** option specifies whether to monitor an image backup or restore of file systems belonging to a Network Attached Storage (NAS) file server.

If you specify **monitor=yes**, Tivoli Storage Manager monitors the current NAS image backup or restore operation and displays processing information on your screen. This is the default.

If you specify **monitor=no**, Tivoli Storage Manager does not monitor the current NAS image backup or restore operation and is available to process the next command.

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. Tivoli Storage Manager 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 filespace** command to interactively delete NAS file spaces from server storage.

Use the `nasnodename` option to identify the NAS file server. Place the `nasnodename` option in your client 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, Tivoli Storage Manager prompts for the password assigned to the node you specify, if a password is required.

If you want to restore or retrieve files from the server while you are working from a different workstation, use the *virtualnodename* option. You can also use the *asnodename* option, if it is set up by the administrator.

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.

- 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, **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 505

Nojournal

Use the **nojournal** option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

Journal-based incremental backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Attribute changes to an object require a 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 client API does not support this option.

Syntax

▶▶—NOJournal—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc incr c: -nojournal
```

Related information

See “Journal-based backup” on page 568 for more information about journal-based incremental backup.

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 client API does not support this option.

Syntax

▶▶—NOPrompt—————▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc delete archive -noprompt c:\home\project\*
```

Numberformat

The `numberformat` option specifies the format you want to use to display numbers.

Use this option if you want to change the default number format for the language of the message repository you are using.

By default, the backup-archive and administrative clients obtain format information from the locale definition in effect at the time the client is called. Consult the documentation on your local system for details about setting up your locale definition.

Note: The `numberformat` option does not affect the Web client. The Web client uses the number format for the locale that the browser is running in. If the browser is not running in a supported locale, the Web client uses the number format for US English.

You can use the `numberformat` option with the following commands:

- **delete archive**
- **delete backup**
- **expire**
- **query archive**
- **query asr**
- **query backup**
- **query image**
- **query nas**
- **query 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 Tivoli Storage Manager 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, Tivoli Storage Manager assumes the file name specified is located in the current working directory. The default is `dsm.opt`.

Examples

Command line:

`dsmc query session -optfile=myopts.opt`

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

Password

The password option specifies a Tivoli Storage Manager password.

If you do not specify this option and your administrator has set authentication to *On*, you are prompted for a password when you start a Tivoli Storage Manager session.

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 Tivoli Storage Manager clients are directed to a different Tivoli Storage Manager server, all clients must re-authenticate with the server because the stored encrypted password must be regenerated.

The password option is ignored when the passwordaccess option is set to *generate*.

Supported Clients

This option is valid for all clients.

Options File

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

Syntax

►►—PASsword— —password—►►

Parameters

password

Specifies a 1 to 63 character password. A password is not case-sensitive. Valid characters include:

Characters

Description

| | |
|-----|---|
| A–Z | Any letter, A through Z, uppercase or lowercase |
| 0–9 | Any number, 0 through 9 |
| + | Plus |
| . | Period |
| _ | Underscore |
| - | Hyphen |
| & | Ampersand |

To keep your client node password secure, enter commands without the password and wait for Tivoli Storage Manager to prompt you for the password.

API applications must supply the password when a session is initiated. The application is responsible for obtaining the password.

generate

Encrypts and stores your password locally and generates a new password when the old password expires. The new password is randomly generated by the Tivoli Storage Manager client

A password prompt is displayed when registering a workstation with a server using open registration or if your administrator changes your password manually.

Examples

Options file:

```
passwordaccess generate
```

Command line:

Does not apply.

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 client API does not support this option.

Syntax

►—Pick—◄

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore c:\project\* -pick -inactive
```

Pitdate

Use the `pitdate` option with the `pitime` 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 client API does not support this option.

Syntax

►—PITDate =— —date—◄

Parameters

date

Specifies the appropriate date. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc restore -pitdate=08/01/2003 c:\myfiles\
```

Pittime

Use the `pittime` option with the `pitdate` option to establish a point in time 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 client API does not support this option.

Syntax

►►—PITtime =— —*time*—►►

Parameters

time

Specifies a time on a specified date. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the `timeformat` option.

When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `tottime` options.

Examples

Command line:

```
dsmc query backup -pitt=06:00:00 -pitd=08/01/2003 c:\myfiles\
```

Postschedulecmd/Postnschedulecmd

The postschedulecmd/postnschedulecmd option specifies a command that the client program processes after it runs a schedule.

The client program waits for the command to complete before it continues with other processing.

If you do not want to wait, specify postnschedulecmd.

Note:

1. If the postschedulecmd command does not complete with return code 0, the client reports that the scheduled event completed with return code 8 (unless the scheduled operation encounters a more severe error yielding a higher return code). If you do not want the postschedulecmd command to be governed by this rule, you can create a script or batch file that invokes the command and exits with return code 0. Then configure postschedulecmd to invoke the script or batch file. The return code for the postnschedulecmd command is not tracked, and does not influence the return code of the scheduled event.
2. The server can also define the postschedulecmd option (and the postnschedulecmd option).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax

```
▶▶ [ POSTSCHEDULECMD | POSTNSCHEDULECMD ] — "cmdstring" —▶▶
```

Parameters

"cmdstring"

Specifies the command to process. You can enter a command to be run after a schedule with this option. Use only one postschedulecmd option.

Specify the command string just as you would enter it from the operating system command prompt. If the command string contains any blank spaces, enclose the command string in single quotation marks. For example:

```
'net stop someservice'
```

Use a blank, or null, string for *cmdstring* if you want to prevent any commands from running that the Tivoli Storage Manager server administrator uses for

postschedulecmd or preschedulecmd. If you specify a blank or null string on *either* option, it prevents the administrator from using a command on *both* options.

If your administrator uses a blank or null string on the postschedulecmd option, you cannot run a post-schedule command.

Examples

Options file:

```
posts startdb.cmd
posts 'rename c:\myapp\logfile.log logfile.new'
posts 'net start "simple service"'
posts 'rename "c:\myapp\log file.log" "log file.new"'
posts '"C:\Program Files\MyTools\runreport.bat"
log1.in log2.in'
```

Command line:

```
-postschedulecmd="'restart database'"
```

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

Postsnapshotcmd

The postsnapshotcmd option allows you to run operating system shell commands or scripts after the Tivoli Storage Manager 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.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (dsm.opt). You can 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 393

“Srvprepostscheddisabled” on page 478

Preschedulecmd/Prenschedulecmd

The `preschedulecmd` option specifies a command that the client program processes before it runs a schedule.

The client program waits for the command to complete before it starts the schedule. If you do not want it to wait, specify `prenschedulecmd`.

Note:

1. Successful completion of the `preschedulecmd` command is considered to be a prerequisite to running the scheduled operation. If the `preschedulecmd` command does not complete with return code 0, the scheduled operation and any `postschedulecmd` and `postnschedulecmd` commands will not run. The client reports that the scheduled event failed, and the return code is 12. If you do not want the `preschedulecmd` command to be governed by this rule, you can create a script or batch file that invokes the command and exits with return code 0.

Then configure `preschedulecmd` to invoke the script or batch file. The return code for the `preschedulecmd` command is not tracked, and does not influence the return code of the scheduled event.

2. The server can also define the `preschedulecmd` option (and the `preschedulecmd` option).

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax



Parameters

cmdstring

Specifies the command to process. Use only one `preschedulecmd` option. You can enter a command to be executed before a schedule using this option.

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.

PreserveLastAccessDate

Use the `preserveLastAccessDate` option during a backup or archive operation to specify whether to reset the last access date of any specified files to their original value following the backup or archive operation.

By default, the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This option requires extra processing time during backup and archive for each file that is sent to the Tivoli Storage Manager server, so it should only be enabled when necessary.

Any application that touches a file can implicitly cause the last access date of the file to change to the time that the application touches it. This is a function of the file system, not the application. Because of this, when the client backs up or archives a file, it can trigger an update to the last access date of the file. This can cause problems for other applications such as Storage Resource Management (SRM), whose processing relies on accurate last access dates.

If you have enabled open file support, the last access date for files is always preserved regardless of the setting for `preserveLastAccessDate`. When using open file support, leave this option disabled.

Use this option with the **incremental**, **selective**, or **archive** commands.

Note:

1. This option only applies to files; it does not apply to directories.
2. Resetting the last access date incurs additional overhead that impact backup and archive performance. The last access date should be reset only if you are using other another application, such as a Storage Resource Management (SRM) that relies on accurate last access dates.
3. The last access date cannot be preserved on files which are write protected either by the read-only attribute or by a restrictive NTFS security permission.
4. You cannot reset the last access date of read-only files. The `preserveLastAccessDate` option ignores read-only files and does not change their date.

Supported Clients

This option is valid for all Windows clients.

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 Specifies that the Tivoli Storage Manager client *will not* reset the last access date of any backed up or archived files to their original value following the backup or archive operation. This is the default.

Yes

Specifies that the Tivoli Storage Manager client *will* reset the last access date of any backed up or archived files to their original value following the backup or archive operation.

Examples

Options file:

```
preservelastaccessdate yes
```

Command line:

```
dsmc incr c: e: f: -preservelastaccessdate=yes
```

Preservepath

The `preservepath` option specifies how much of the source path to reproduce as part of the target directory path when you restore or retrieve files to a new location.

Use the `-subdir=yes` option to include the entire subtree of the source directory (directories and files below the lowest-level source directory) as source to be restored. If a required target directory does not exist, it is created. If a target file has the same name as a source file, it is overwritten. Use the `-replace=prompt` option to have Tivoli Storage Manager prompt you before files are overwritten.


Use the `preservepath` option with the following commands:

- **restore**
- **restore backupset**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Syntax



Parameters

Subtree

Creates the lowest-level source directory as a subdirectory of the target directory. Files from the source directory are stored in the new subdirectory. This is the default.

Complete

Restores the entire path, starting from the root, into the specified directory. The entire path includes all the directories *except* the file space name.

NOBase

Restores the contents of the source directory without the lowest level, or base directory, into the specified destination directory.

NONE

Restores all selected source files to the target directory. No part of the source path at or above the source directory is reproduced at the target.

If you specify `SUBDIR=yes`, Tivoli Storage Manager restores all files in the source directories to the single target directory.

Examples

Command line:

For the examples below, assume the server file space contains the following backup copies:

```
c:\h1\m1\file.a
c:\h1\m1\file.b
c:\h1\m1\11\file.x
c:\h1\m1\11\file.y
```

This command:

```
dsmc res backupset my.backupset.file /fs/h1/m1/ /u/ann/ -su=yes
creates a local backupset file named "my.backupset.file".
```

Restores these directories and files:

```
c:\ann\h1\m1\file.a
c:\ann\h1\m1\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=nobase.
```

Restores these directories and files:

```
c:\ann\file.a
c:\ann\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=subtree.
```

Restores these directories and files:

```
c:\ann\m1\file.a
c:\ann\m1\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -preser=none.
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=  
complete
```

Restores these directories and files:

```
c:\ann\h1\m1\file.a  
c:\ann\h1\m1\file.b  
c:\ann\h1\m1\l1\file.x  
c:\ann\h1\m1\l1\file.y
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=nobase.
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b  
c:\ann\l1\file.x  
c:\ann\l1\file.y
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=subtree.
```

Restores these directories and files:

```
c:\ann\m1\file.a  
c:\ann\m1\file.b  
c:\ann\m1\l1\file.x  
c:\ann\m1\l1\file.y
```

This command:

```
dsmc res c:\h1\m1\ c:\ann\ -su=yes -preser=none.
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y
```

This command:

```
dsmc res backupset c:\h1\m1\ c:\ann\ -su=yes  
-preser=nobase -loc=file
```

Restores these directories and files:

```
c:\ann\file.a  
c:\ann\file.b  
c:\ann\file.x  
c:\ann\file.y
```

Presnapshotcmd

The `presnapshotcmd` option allows you to run operating system commands before the Tivoli Storage Manager client starts a snapshot.

This allows you to quiesce an application before the Tivoli Storage Manager 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.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set 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 393

"Srvprepostscheddisabled" on page 478

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 client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Scheduler** tab, **Query Schedule Interval** field of the Preferences editor.

Syntax

►►—QUERYSCHeDperiod— *—hours—*—————►►

Parameters

hours

Specifies the number of hours the client scheduler waits between attempts to contact the server for scheduled work. The range of values is 1 through 9999; the default is 12.

Examples

Options file:

```
querysch 6
```

Command line:`-queryschedperiod=8`

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

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 client 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 client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

▶—QUIET—▶

Parameters

There are no parameters for this option.

Examples

Options file:

```
quiet
```

Command line:

```
-quiet
```

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

Replace

The `replace` option specifies whether to overwrite existing files on your workstation, or to prompt you for your selection when you restore or retrieve files.

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, Tivoli Storage Manager skips files without prompting during a scheduled operation.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). 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. Tivoli Storage Manager 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.

Resetarchiveattribute

Use the `resetarchiveattribute` option to specify whether Tivoli Storage Manager resets the Windows archive attribute on files that are successfully backed up to a Tivoli Storage Manager server.

Tivoli Storage Manager also resets the archive attribute during incremental backups if it is determined that there is already an active object on the Tivoli Storage Manager server. The `resetarchiveattribute` option is useful in conjunction with applications, such as the IBM Tivoli Storage Resource Manager, as a simple way to report on the backup status of files.

The Windows archive attribute is used to indicate that a file has changed since the last backup. After Tivoli Storage Manager resets the archive attribute, the Windows operating system turns the attribute back to *ON* after the file has been modified. Tivoli Storage Manager does not use the Windows archive attribute to determine if a file is a candidate for incremental backup, but only manipulates this attribute for reporting purposes. Tivoli Storage Manager uses a much more sophisticated method to determine candidacy for incremental backup.

There are several applications which manipulate or examine the Windows archive attribute. Be aware of the ramifications of using the `resetarchiveattribute` option in conjunction with these products.

If you set the `resetarchiveattribute` option to *yes*, after a file has been successfully backed up to the Tivoli Storage Manager server, the Tivoli Storage Manager client resets the Windows archive attribute on the local file system:

- The Windows archive attribute is reset during incremental and selective backups after the file has been successfully committed to the Tivoli Storage Manager database. This attribute is not reset for archive, or image operations.
- The Windows archive attribute is not reset when processing system objects or system state objects.
- The Windows archive attribute is not reset for directory entries.

In addition, in order for the local file system to reflect the current active object inventory on the Tivoli Storage Manager server, the `resetarchiveattribute` option instructs the Tivoli Storage Manager client to reset the Windows archive attribute on the local file system if it is determined during incremental backup that a valid, active backup copy of the file already exists on the Tivoli Storage Manager server. This behavior is not displayed in the following cases:

- Incremental backup operations which do not examine the stored client attributes on the Tivoli Storage Manager server, such as journal-based backup or incremental-by-date processing.
- Files that are not examined during an incremental backup operation because they are excluded from backup processing.

The Tivoli Storage Manager client does not guarantee the accuracy of the current setting of the Windows archive attribute. For example, if the `resetarchiveattribute` option is set to *yes* and a file examined by a reporting product indicates that the Windows archive attribute is *OFF* for a particular file, this does not necessarily mean that a valid, active backup copy of the file exists on the Tivoli Storage Manager server. Factors that could contribute to this type of situation include:

- An independent software vendor product is manipulating the Windows archive attribute
- A file space was deleted from the Tivoli Storage Manager server
- A backup tape was lost or destroyed

There should be no significant performance degradation when using the `resetarchiveattribute` option. The `resetarchiveattribute` option does not affect restore processing.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

This option is valid in the client options file (dsm.opt) or server client options set. You can set this option on the **Backup** tab of the Preferences editor.

Syntax

►►—RESETARCHIVEATTRIBUTE =

| |
|-----|
| No |
| Yes |

—►►

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 112

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 client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab, **Resource Utilization** field of the Preferences editor.

Syntax

►►—RESOURceutilization— *number* —►►

Parameters

number

Specifies the level of resources the Tivoli Storage Manager server and client can use during processing. The range of values that you can specify is 1 through 10.

Examples

Options file:

```
resourceutilization 7
```

Command line:

```
-resourceutilization=7
```

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

Regulating backup and archive sessions

When you request a backup or archive, the client can use more than one session to the server.

The default is to use a maximum of two sessions; one to query the server and one to send file data. The client can use only one server session if you specify a `resourceutilization` setting of 1.

A client can use more than the default number of sessions when connecting to a server that is Version 3.7 or higher. For example, `resourceutilization=10` permits up to eight sessions with the server. Multiple sessions can be used for querying the server and sending file data.

Multiple query sessions are used when you specify multiple file specifications with a backup or archive command. For example, if you enter 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 or not the second session starts depends on how long it takes to query the server about files backed up on file space A. The client might also try to read data from the file system and send it to the server on multiple sessions.

Note: During a backup operation, if you enter multiple file specifications, the result might be that files from one file specification are stored on multiple tapes and interspersed with files from different file specifications. This can decrease restore performance. Setting the `collocatebyfilespec` option to *yes* eliminates interspersing of files from different file specifications, by limiting the client to one server session per file specification. Therefore, if you store the data to tape, files for each file specification are stored together on one tape (unless another tape is required for more capacity).

Related reference

“`Collocatebyfilespec`” on page 319

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 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 reside on 4 tapes and some on disk, you could use up to 5 sessions during the restore.
2. The Tivoli Storage Manager server can set the maximum number of mount points a node can use on the server using the MAXNUMMP parameter. If the resourceutilization option value exceeds the value of the MAXNUMMP on the server for a node, the backup can fail with an *Unknown System Error* message.
3. You can get a multi-session restore from your single **restore** command, and from a single volume on the server, if that volume is device class FILE.

For example, if the data you want to restore is on 5 different tape volumes, the maximum number of mount points is 5 for your node, and resourceutilization is set to 3, then 3 sessions 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 to the number of restore sessions allowed for 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.

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 client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Scheduler** tab, **Retry period** field of the Preferences editor.

Syntax

►►—RETRYPeriod— —*minutes*——————►►

Parameters

minutes

Specifies the number of minutes the client scheduler waits between attempts to contact the server, or to process a scheduled command that fails. The range of values is 1 through 9999; the default is 20.

Examples

Options file:

```
retryp 10
```

Command line:

```
-retryperiod=10
```

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

Revokeremoteaccess

The `revokeremoteaccess` option restricts an administrator with client access privilege from accessing a client workstation that is running the Web client.

This option does not restrict administrators with client owner, system, or policy privilege from accessing your workstation through the Web client.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Web Client** tab of the Preferences editor.

Syntax



Parameters

None

Does not revoke access to administrators who have client access authority for the client. This is the default.

Access

Revokes access to administrators who have client access authority for the client.

Examples

Options file:

```
revokeremoteaccess none
```

Command line:

Does not apply.

Runasservice

The runasservice option forces the client command process to continue running, even if the account that started the client logs off.

Use this option with the **AT** command and the **dsmc sched** command when you schedule client command batch jobs. The runasservice option is *not* valid in any options file (dsm.opt or tsmasr.opt).

Important: Use the scheduler service when running Tivoli Storage Manager services unattended. Set runasservice=*yes* only to schedule Tivoli Storage Manager commands using the Windows **AT** command. Setting runasservice=*yes* might interfere with other interactive uses of the Tivoli Storage Manager client.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

Syntax



Parameters

No Does not force the client command process to continue running, even if the account that started the client logs off. This is the default.

Yes

Forces the client command process to continue running, even if the account that started the client logs off.

commands by the server using the `action=command` option on the DEFINE SCHEDULE server command. This is the default.

Examples

Options file:

```
schedcmddisabled no
```

Command line:

Does not apply.

Related information

“Query Schedule” on page 598

Schedcmexception

The `schedcmexception` 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 `schedcmexception` 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 exactly.

Examples

Options file:

```
schedcmddisabled yes
```

```
schedcmdexception "start dir c: /s"  
schedcmdexception "start echo hello, world!"
```

Examples:

Disable the scheduling of commands by the server `action=command` option on the DEFINE SCHEDULE server command, except for two schedules defined on the Tivoli Storage Manager Server:

Related information

"Schedcmddisabled" on page 447

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 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 > 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, Tivoli Storage Manager attempts to create it.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Scheduler** tab, **Schedule Log** button of the Preferences editor.

Important: Set the `DSM_LOG` environment variable to name a directory where the log is to be placed. The directory specified must have permissions which allow write access from the account under which the client is run.

Syntax

▶—SCHEDLOGName— *filespec* —▶

Parameters

filespec

Specifies the path and file name where you want to store schedule log information when processing scheduled work. If any part of the path you specify does not exist, Tivoli Storage Manager attempts to create it.

If you specify a file name only, the file is stored in your current directory. The default is the current working directory with a file name of `dsmsched.log`.

Examples

Options file:

```
schedlogname c:\mydir\schedlog.jan
```

Command line:

```
-schedlogn=c:\mydir\schedlog.jan
```

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

Related information

See “Errorlogname” on page 361 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, **Schedule Mode** section of the Preferences editor.

Syntax



Parameters

POLLing

The client scheduler queries the server for scheduled work at prescribed time intervals. This is the default. You can set the time intervals using the `querschedperiod` option.

PRompted

The client scheduler waits for the server to contact your client node when scheduled work needs to be done.

Note:

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 306

“Tcpclientaddress” on page 490

“Tcpclientport” on page 491

Schedrestretrdisabled

The `schedrestretrdisabled` option specifies whether to disable the execution of restore or retrieve schedule operations.

Supported Clients

This option is valid for all clients. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (dsm.opt) for the scheduler. You can set this option on the **Scheduler** tab, **Schedule Command** button of the Preferences editor.

Syntax



Parameters

No Specifies that Tivoli Storage Manager does not disable the execution of restore and retrieve schedule operations. This is the default.

Yes

Specifies that Tivoli Storage Manager disable the execution of restore and retrieve schedule operations.

Examples

Options file:

```
schedrestretrdisabled yes
```

Command line:

Does not apply.

Scrolllines

The scrolllines option specifies the number of lines of information that are displayed on your screen at one time.

Use this option when you set the scrollprompt option to *Yes*.

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 client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Command Line** tab, **Number of lines to display on screen** field of the Preferences editor.

Syntax

►—SCROLLLines— *number*—►

Parameters

number

Specifies the number of lines of information that are displayed on your screen at one time. The range of values is 1 through 80; the default is 20.

Examples

Options file:

```
scrolllines 25
```

Command line:

```
-scroll=25
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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 Tivoli Storage Manager to stop and wait after displaying the number of lines of information you specified with the scrolllines option, or scroll through and stop at the end of the information list.

You can use the scrollprompt option with the following commands only:

- **delete filespace**
- **query archive**
- **query backup**
- **query backupset**
- **query filespace**
- **query 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 client API does not support this option.

Options File

Place this option in the client user-options file (dsm.opt). You can set this option on the **Command Line** tab, **Pause after displaying the following number of lines** field of the Preferences editor.

Syntax



Parameters

No Scrolls to the end of the list and stops. This is the default.

Yes

Stops and waits after displaying the number of lines you specified with the `scrolllines` option. The following prompt is displayed at the bottom of the screen:

Press 'Q' to quit, 'C' to continuous scroll, or 'Enter' to continue.

Examples

Options file:

```
scrollprompt yes
```

Command line:

```
-scrollp=yes
```

This option is valid on the initial command line and in interactive mode. If you use this option in interactive mode, it affects only the command with which it is specified. When that command completes, the value reverts to the value at the beginning of the interactive session. This 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, it is unnecessary to open any ports on the firewall. If you set the `sessioninitiation` option to **serveronly**, the client will not attempt to contact the server. All sessions must be initiated by server prompted scheduling on the port defined on the client with the `tcpclientport` option. The `sessioninitiation` option only affects the behavior of the client scheduler running in the prompted mode. If you set the `sessioninitiation` option to **serveronly**, with the exception of client acceptor daemon-managed schedulers, the command-line client, 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

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.

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 client scheduler” on page 24

“`Tcpclientport`” on page 491

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 with previous versions of Tivoli Storage Manager to complete.
- 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 Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows Server 2003 R2, Windows Vista, and Windows 7 clients only.

Options File

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

Syntax



Parameters

Yes

Specifies that you want the Tivoli Storage Manager client to 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 Tivoli Storage Manager 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 547

Skipntpermissions

The skipntpermissions option bypasses processing of NTFS security information.

Select this option for incremental backups, selective backups, archives, retrieves and restores. Use this option with the **archive**, **incremental**, **restore**, **retrieve**, and **selective** commands.

Supported Clients

This option is valid for all Windows clients.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **General** tab of the Preferences editor.

Syntax



Parameters

No If you specify *No*, the NTFS security information is backed up, restored, archived or retrieved. This is the default.

Yes

If you specify *Yes*, the NTFS security information is not backed up, restored, archived or retrieved.

Examples

Options file:

```
skipntp yes
```

Command line:

```
-skipntp=yes
```

Skipntsecuritycrc

The `skipntsecuritycrc` option controls the computation of the security cyclic redundancy check (CRC) for a comparison of NTFS security information during an incremental or selective backup, archive, restore, or retrieve operation.

If you set the `skipntsecuritycrc` option to *no* (the default), performance might be slower because the program must retrieve all the security descriptors.

Use this option with the following commands:

- **archive**
- **incremental**
- **restore**
- **retrieve**
- **selective**

Supported Clients

This option is valid for all Windows clients.

Options File

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

Syntax



Parameters

No If you specify *No*, the security CRC is generated during a backup. This is the default.

Yes

If you specify *Yes*, the security CRC is not generated during a backup. All the permissions are backed up, but the program 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` 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 Tivoli Storage Manager client cannot determine their mount points programmatically.

You must configure a user ID and password on the Tivoli Storage Manager 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 67.

Use this option with an incremental backup of a NAS file server volume instead of a simple incremental or incremental with `snapshotroot` whenever the NAS file server is running ONTAP 7.3.0, or later, for performance reasons. Do not use the `snapdiff` and `snapshotroot` options together.

The first time you perform an incremental backup with the snapshot difference option, a snapshot is created (the base snapshot) and a traditional incremental backup is run using this snapshot as the source. The name of the snapshot that is created is recorded in the Tivoli Storage Manager database.

The second time an incremental backup is run with this option, a newer snapshot is either created, or an existing one is used to find the differences between these two snapshots. The second snapshot is called the `diffsnapshot`. Tivoli Storage Manager then incrementally backs up the files reported as changed by NetApp to the Tivoli Storage Manager server. The filesystem that you select for snapshot difference processing must be mounted to the root of the volume. You cannot use the `snapdiff` option for any filesystem that is not mounted to the root of the volume. After backing up data using the `snapdiff` 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`. Tivoli Storage Manager does not delete the snapshot if it was not created by Tivoli Storage Manager. You can also run a snapshot-difference incremental-backup by specifying the `-diffsnapshot=latest` option.

For NAS and N-Series file servers that are running ONTAP 7.3.0, or later, you can use the `createnewbase` option to backup any files that might have been skipped due to 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 have not changed the include-exclude file, but you

have removed the rule that excluded the file. The NetApp API only detects file changes between two snapshots, not changes to the include-exclude file.

- If you have added an include statement to the option file, that include option does not take effect unless NetApp detects that the file has changed. This is because Tivoli Storage Manager does not inspect each file on the volume during backup.
- You have used the **dsmc delete backup** command to explicitly delete a file from the Tivoli Storage Manager inventory. NetApp does not detect that a file has been manually deleted from Tivoli Storage Manager. Therefore, the file remains unprotected in Tivoli Storage Manager storage until it is changed on the volume and the change is detected by NetApp, signalling Tivoli Storage Manager 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 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 new base snapshot and uses it as a source to run a full incremental backup. NetApp controls what constitutes a changed object.

Note: In the list of options used by the traditional **incremental** command, the last column shows the interaction of each option with the **snaptiff** option. The definitions of *valid*, *not valid*, and *no effect* are:

Valid Processing runs normally when the option is used.

Not valid

If the option is used with the **snaptiff** option, an error message is generated.

No effect

The option can be used, but it is ignored.

Table 59. Incremental command: Related options

| Option | Where specified | With snaptiff |
|---|---|----------------------|
| asnodename "Asnodename" on page 299 | Client options file (dsm.opt) or command line. | Valid |
| autofsrename "Autofsrename" on page 307 | Client options file (dsm.opt) only. | No effect |
| changingretries "Changingretries" on page 312 | Client options file (dsm.opt) or command line. | No effect |
| compressalways "Compressalways" on page 322 | Client options file (dsm.opt) or command line. | Valid |
| compression "Compression" on page 323 | Client options file (dsm.opt) or command line. | Valid |
| createnewbase "Createnewbase" on page 326 | Command line only. | Valid |
| diffsnapshot "Diffsnapshot" on page 337 | Command line only. | Valid |
| dirsonly "Dirsonly" on page 340 | Command line only. | Valid |
| domain "Domain" on page 343 | Client options file (dsm.opt) or command line only. | Valid |

Table 59. Incremental command: Related options (continued)

| Option | Where specified | With snapdiff |
|---|--|---|
| enablelanfree "Enablelanfree" on page 355 | Client options file (dsm.opt) or command line. | Valid |
| encryptiontype "Encryptiontype" on page 356 | Client options file (dsm.opt). | Valid |
| encryptkey "Encryptkey" on page 357 | Client options file (dsm.opt). | Valid |
| exclude.fs.nas "Exclude options" on page 364 | Client options file (dsm.opt). | No effect |
| filelist "Filelist" on page 377 | Command line only. | Not valid |
| filesonly "Filesonly" on page 381 | Command line only. | Valid |
| include.fs.nas "Include options" on page 393 | Client options file (dsm.opt) or command line. | No effect |
| inlexcl "Inlexcl" on page 392 | Client options file (dsm.opt). | Valid, but only when a file change is detected by NetApp. |
| incrbydate "Incrbydate" on page 399 | Command line only. | Not valid |
| memoryefficientbackup "Memoryefficientbackup" on page 411 | Client options file (dsm.opt), server, or command line. | No effect |
| monitor "Monitor" on page 414 | Command line only. | Not valid |
| nojournal "Nojournal" on page 418 | Command line only. | Not valid |
| postsnapshotcmd "Postsnapshotcmd" on page 428 | Client options file (dsm.opt) or with the include.fs option. | Valid |
| preservelastaccessdate "Preservelastaccessdate" on page 431 | Client options file (dsm.opt) or command line. | Valid |
| presnapshotcmd "Presnapshotcmd" on page 434 | Client options file (dsm.opt) or with the include.fs option. | Valid |
| resetarchiveattribute "Resetarchiveattribute" on page 440 | Client options file (dsm.opt). | Valid |
| skipntpermissions "Skipntpermissions" on page 461 | Client options file (dsm.opt) or command line. | Valid |
| skipntsecuritycrc "Skipntsecuritycrc" on page 462 | Client options file (dsm.opt) or command line. | Valid |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | Client options file (dsm.opt) or with the include.fs option. | No effect |
| snapshotcachesize "Snapshotcachesize" on page 468 | Client options file (dsm.opt) or with the include.fs option. | No effect |
| snapshotfsidleretries "Snapshotfsidleretries" on page 469 | Client options file (dsm.opt) or with the include.fs option. | No effect |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | Client options file (dsm.opt) or with the include.fs option. | No effect |

Table 59. Incremental command: Related options (continued)

| Option | Where specified | With snapdiff |
|---|---|---------------|
| snapshotproviderfs "Snapshotproviderfs" on page 472 | Client options file (dsm.opt) or with the include.fs option. | Not valid |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | Client options file (dsm.opt) or with the include.image option. | Not valid |
| snapshotroot "Snapshotroot" on page 474 | Command line only. | Not valid |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. | Not valid |
| subfilebackup "Subfilebackup" on page 484 | Client options file (dsm.opt) or command line. | Not valid |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. | Valid |
| toc "Toc" on page 496 | Command line only. | Not valid |
| virtualfsname "Virtualfsname" on page 504 | Command line only. | Not valid |

Supported Clients

This option is valid for the following clients:

- All Windows clients except IA 64 bit

Syntax

►►—SNAPDiff—◄◄

Parameters

There are no parameters for this option.

Examples

Command line:

Perform a snapdiff incremental backup from a snapshot 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 snapdiff incremental backup from a snapshot 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:

Perform 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:
```

Perform a snapshot difference 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 have made some include or exclude changes:

```
dsmc incremental -snapdiff -createnewbase=yes h:
```

Related information

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 67

“Diffsnapshot” on page 337

“Set Password” on page 648

“Configuring NetApp and Tivoli Storage Manager for snapshot difference incremental backups” on page 67

Snapshotcachelocation

Use the `snapshotcachelocation` option during an online image backup or open file support operation to specify where the Tivoli Storage Manager Logical Volume Snapshot Agent (LVSA) stores the snapshot cache.

The snapshot cache stores original copies of blocks that change during the course of the backup operation.

Important: Use the default snapshot cache location `x:\tsmlvsa`, where `x:` is the drive being backed up. Specify a location on a different drive if the drive being backed up is expected to change significantly during the course of the backup operation and is unlikely to have enough free space to cache the original data blocks. If the `snapshotprovider` option is VSS, this option is ignored.

For an online image backup, use the `snapshotcachelocation` option with the **backup image** command, in the `dsm.opt` file, or with the `include.image` option.

For open file support operations, you can specify the `snapshotcachelocation` option in your `dsm.opt` file or with the `include.fs` option.

Supported Clients

This option is valid for all supported Windows clients. The server can also define this option.

Options File

Place this option in the client options file. You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

Parameters

path

Specifies the location where the LVSA places the cache files. The path must specify a formatted volume. With the single drive enhancements of LVSA in Tivoli Storage Manager Version 5.3, cache files can be placed on the same volume that is being backed up. The default location is to create the OBF files on the volume on which the snapshot is taken. Snapshot cache location cannot be placed on a GPFS drive.

Examples

Options file:

```
snapshotcachelocation c:\temp\
```

Command line:

```
-snapshotcachelocation=c:\temp\
```

Related information

See “Include options” on page 393 for more information about `include.fs`.

Snapshotcachesize

Use the `snapshotcachesize` option to specify an appropriate size to create the snapshot.

The size estimation is needed for the cache file containing the original data for blocks that change during an online image backup or open file support operation

On a volume with high file system write activity, this value can be increased to remove the condition where the snapshot cache fills up during the operation. The snapshot cache location is specified with the `snapshotcachelocation` option. This option is only valid if OFS or online image backup has been configured and the LVSA has been selected as the snapshot provider.

For online image backups, use the `snapshotcachesize` option with the **backup image** command, the `include.image` option, or in your `dsm.opt` file.

For open file support operations, use the `snapshotcachesize` option in your `dsm.opt` file or with the `include.fs` option.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

▶—SNAPSHOTCACHESize— *size* —▶

Parameters

size

Specifies the maximum size of the snapshot cache file as a percentage of the used space of the volume for which the snapshot is taken. The range of values is 1 through 100. The default is 1. Note that Tivoli Storage Manager always uses a minimum of 500 MB.

If the snapshot cache size exceeds the amount of free space on the volume where the snapshot cache file resides, the client truncates the snapshot cache file size to 90 percent of the available free space and writes the following message in `dsmerror.log`:

Snapshot cache size adjusted to accommodate insufficient space

Examples

Options file:

```
snapshotcachesize 40
```

Command line:

```
-snapshotcachesize=40
```

Related information

See “Include options” on page 393 for more information about `include.fs`.

Snapshotfsidleretries

Use the `snapshotfsidleretries` option to specify the number of additional times the LVSA should try to achieve the snapshot file system idle wait time before the online image backup or open file support operation fails.

This option is only valid if OFS or online image backup has been configured, and the LVSA has been selected as the snapshot provider.

While the snapshot provider gives a point-in-time image of the volume, it does not guarantee that there is no disk related activity while the snapshot is taken. This might lead to an image backup that requires you to run `chkdsk` before the image backup is accessed. A message indicates that a `chkdsk` is necessary when you access the volume after the restore. In the case of OFS, a file might not be able to be read. If the file cannot be read, a warning of the file read problem appears in the `dsmerror.log` as a failed backup file, and the file is backed up during the next incremental backup.

For an online image backup, use the `snapshotfsidleretries` option with the **backup image** command, the `include.image` option, or in your `dsm.opt` file. Place the `include.image` statement containing the `snapshotfsidleretries` value in your `dsm.opt` file.

For open file support operations, use the `snapshotfsidleretries` option in your `dsm.opt` file or the with the `include.fs` option.

If the value you specified with the `snapshotfsidlewait` option cannot be satisfied in the number of retries you specify with the `snapshotfsidleretries` option, the snapshot and the image backup fails. If this occurs, you must adjust the `snapshotfsidlewait` and `snapshotfsidleretries` values until the snapshot is successful. You can also use the `presnapshotcmd` and `postsnapshotcmd` options to temporarily quiesce applications until the snapshot starts. You can also use the Preferences editor to specify these options.

Supported Clients

This option is valid for the Windows 32-bit clients *only*. The Tivoli Storage Manager client API does not support this option. 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 **Image-Snapshot** tab of the Preferences editor.

Syntax

▶—SNAPSHOTFSIDLERetries— *—retrynumber—*▶

Parameters

retrynumber

Specifies the number of additional times the LVSA should try to achieve the snapshot file system idle wait time before the online image backup operation fails. The range of values is zero through 99; the default is 99.

Examples

Options file:

```
include.image h: MYMC snapshotfsidleretries=2
```

Command line:

```
-snapshotfsidleretries=2
```

Related information

“Include options” on page 393

“Postsnapshotcmd” on page 428

“Presnapshotcmd” on page 434

“Snapshotfsidlewait”

Snapshotfsidlewait

Use the `snapshotfsidlewait` option during an online image backup or open file support operation to specify the amount of time that must pass in which there is no write activity (read activity is ignored) on a volume before a snapshot can be taken. This allows the disk to reach an idle state.

This option is only valid if OFS or online image backup has been configured, and the LVSA has been selected as the snapshot provider.

While the snapshot provider gives a point-in-time image of the volume, it does not guarantee that there is no disk related activity while the snapshot is taken. This might lead to an image backup that requires you to run `chkdsk` before the image backup is accessed. A message indicates that a `chkdsk` is necessary when you access the volume after the restore. In the case of OFS, a file might not be able to be read. If the file cannot be read, a warning of the file read problem appears in the `dsmerror.log` as a failed backup file, and the file is backed up during the next incremental backup. The snapshot file system idle wait period can decrease the likelihood of creating a snapshot while disk activity is occurring on the volume by delaying the creation of the snapshot until the disk is idle (with no write activity) for a specified amount of time. For example, if you specify a value of 2 seconds, a period of 2 seconds must pass without write activity on a volume before a snapshot can occur.

This option is only valid if the LVSA is installed and configured on your system allowing you to perform an online image backup or open file support operation.

For an online image backup, you can use the `snapshotfsidlewait` with the **backup image** command, the `include.image` option, or in your `dsm.opt` file when performing an online image backup. Place the `include.image` statement containing the `snapshotfsidlewait` value in your `dsm.opt` file.

For open file support operations, use the `snapshotfsidlewait` option in your `dsm.opt` file or the with the `include.fs` option.

Supported Clients

This option is valid for the Windows 32-bit clients *only*. The Tivoli Storage Manager client API does not support this option. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax

▶▶—SNAPSHOTFSIDLEwait— *maximum wait time* —————▶▶
 └ *minimum wait time* ─┘

Parameters

maximum wait time

Specifies the maximum amount of time that must pass during which there is no write activity on a volume before a snapshot can be taken. You can specify **ms** (milliseconds) or **s** (seconds) qualifiers with the value. If you do not specify a qualifier, the value you specify is interpreted as seconds. The range of values is zero through 999; the default value is 2 s. A value of zero (0) indicates that the snapshot occurs immediately without waiting for a period of inactivity. You can specify a *maximum wait time* value *without* specifying a *minimum wait time* value.

minimum wait time

Specifies the minimum amount of time that must pass during which there is no write activity on a volume before a snapshot can be taken. You can specify

ms (milliseconds) or **s** (seconds) qualifiers with the value. If you do not specify a qualifier, the value you specify is interpreted as seconds. The range of values is zero through 999; the default value is 50 **ms**.

If you set `snapshotfsidlewait` to 6,1 and `snapshotfsidleretries` to 5, Tivoli Storage Manager waits for i/o inactivity for 5 seconds, then retries 5, 4, 3, 2, and finally 1. The maximum value, minimum value, and the value of `snapshotfsidleretries` are used to determine the value to decrease by on each retry. This allows the longer wait times to be attempted first but still allows snapshots to occur if the activity is too high. If `snapshotfsidleretries` is set to zero (0), no retries are attempted, even if a minimum value is specified.

If you specify both maximum and minimum wait times, the following rules apply:

Note:

1. You must specify the *maximum wait time*.
2. If you do not specify a *minimum wait time*, the *maximum wait time* is used until the value of `snapshotfsidleretries` is met.
3. The *maximum wait time* must be greater than or equal to the *minimum wait time* or you receive an invalid option error.
4. The *maximum wait time* qualifier can be different than the *minimum wait time* qualifier. For example, `SNAPSHOTFSIDLEWait 4s,500ms` is valid.
5. You can use either a comma or a space to separate the *maximum wait time* and *minimum wait time*.

Examples

Options file:

```
SNAPSHOTFSIDLEWait 5
SNAPSHOTFSIDLEWait 5,1
SNAPSHOTFSIDLEWait 4S,500MS
SNAPSHOTFSIDLEWait 7S,1
SNAPSHOTFSIDLEWait 7,1S
SNAPSHOTFSIDLEWait 7,500MS
```

```
Include.Fs
include.fs f: snapshotproviderfs=lvsa
snapshotfsidlewait=4s,500ms
```

```
Include.Image
include.image f: snapshotproviderimage=lvsa
snapshotfsidlewait=7s,1s
```

Command line:

```
-snapshotfsidlewait=2s
-snapshotfsidlewait=4s,500ms
```

Related information

“Include options” on page 393

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:

LVSA

Specifies that the LVSA should be used to provide OFS support.

Note: LVSA must be installed before it can be specified. Install LVSA during the install process or use the setup wizard.

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

Refer to “Configuring Open File Support” on page 66 for information about installing LVSA.

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:

LVSA

Specifies that the LVSA should be used to provide online image support.

Note: LVSA must be installed before it can be specified. Install LVSA during the install process or use the setup wizard.

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

Refer to “Configuring Open File Support” on page 66 for information about installing LVSA.

Snapshotroot

Use the `snapshotroot` option with the **incremental**, **selective**, or **archive** commands in conjunction with a 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 back up specification (source) and the *snapshotroot* value can be a network share mounted file specification. For example, the snapshotroot option can be used to back up a network share file system hosted on a network-attached storage (NAS) that supports snapshot.

In the following example, c:\snapshots\snapshot.0 is network share mounted from a NAS file server and \\florance\c\$ represents the snapshot created at the NAS file server.

```
dsmc incr \\florance\C$ -snapshotroot=c:\shapshots
\snapshot.0
```

You can also specify a directory with the snapshotroot option when backing up each fileset as a separate file space.

The snapshotroot option does not provide any facilities to take a volume snapshot, only to manage data created by a volume snapshot.

For example, consider an application that takes a snapshot of the c: drive and mounts it as the NTFS junction point \\florence\c\$\snapshots\snapshot.0. If you back up this data using the following command, a unique file space 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 alternate 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 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). 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 lymgmtclass
exclude \\florence\c$\mydocs\*.doc
```

The following include-exclude statements are not valid because they contain the name of the snapshot:

```
include \\florence\c$\snapshots\snapshot.1\dir1\...\
*.txt lymgmtclass
exclude \\florence\c$\mydocs\*.doc
```

You must use the `snapshotroot` option in conjunction with a single file specification for a incremental, selective, or archive operation. You cannot specify multiple file specifications or no file specifications. For example, these commands are valid:

```
dsmc incr c: -snapshotroot=\\florence\c$\snapshots\snapshot.0
dsmc incr c:\dir1\* -snapshotroot=\\florence\c$\snapshots\
snapshot.0
```

The following command is invalid because it contains two file specifications:

```
dsmc incr c:\dir1\* e:\dir1\* -snapshotroot=\\florence\c$\
snapshots\snapshot.0
```

The following command is invalid because it contains no file specification:

```
dsmc incr -snapshotroot=\\florence\c$\snapshots\snapshot.0
```

Note:

1. Ensure that the `snapshotroot` references a snapshot of the correct volume. Ensure that `snapshotroot` refers to the root of the snapshot. If these rules are not followed, unintended results such as files expiring incorrectly can result.
2. If you specify the `filelist` option and the `snapshotroot` option, all files specified in the `filelist` option are assumed to be located 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**, **backup systemstate**, etc.
4. You cannot use the `snapshotroot` option with the `snappdiff` option.
5. If you specify the `snapshotroot` option in conjunction with the Tivoli Storage Manager Open File Support feature provided by the Logical Volume Snapshot Agent (LVSA), the Tivoli Storage Manager Open File Support feature will not be used for the operation. Use the Tivoli Storage Manager LVSA to provide open file support instead of using the `snapshotroot` option in conjunction with another snapshot provider.

Important: Use the `snapshotroot` option with caution in conjunction with the Tivoli Storage Manager journal-based backup feature. Since there is no coordination between the Tivoli Storage Manager journal and the vendor-acquired snapshot provider, unwanted behavior can occur with journal notifications received after the snapshot occurs. For example, files might not be backed up, or they might be backed up redundantly to the Tivoli Storage Manager server. If you want to use the Tivoli Storage Manager journal-based

backup feature in conjunction with a snapshot provider, use it in conjunction with the Tivoli Storage Manager LVSA which provides full integration with the journal-based backup feature.

6. You can use the `snapshotroot` option in conjunction with the `preschedulecmd` and `postschedulecmd` options, or in an automated script that you execute with the Tivoli Storage Manager 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 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 a 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 a Tivoli Storage Manager server:

- `encryptkey generate`
- `exclude.encrypt`
- `include.encrypt`

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

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

Syntax

►►—SRVOPTSETENCryptiondisabled *no* *yes*—◄◄

Parameters

yes

Tivoli Storage Manager client ignores the values of the listed encryption options in a client options set from a Tivoli Storage Manager server.

no

Tivoli Storage Manager client processes the setting of the listed encryption options in a client options set from a 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 Tivoli Storage Manager client node.

Supported Clients

This option is valid for all Tivoli Storage Manager clients that use the Tivoli Storage Manager client scheduler. The server cannot define this option.

Options File

Place this option in the client options file (`dsm.opt`) for the scheduler. You can set this option on the **Schedule Command** button of the Preferences editor.

Syntax



Parameters

No Specifies that Tivoli Storage Manager allows pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled operations. If a pre-schedule or a post-schedule command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that Tivoli Storage Manager prevents pre-schedule and post-schedule commands defined by the Tivoli Storage Manager administrator to execute on

the client system, when performing scheduled operations. If a pre-schedule or a post-schedule command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator will *not* override the corresponding command defined in the client option file. This option can be used in conjunction with the `schedcmddisabled` and `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 Tivoli Storage Manager client node.

Supported Clients

This option is valid for Windows clients that support the image snapshot backup command. The server cannot define this option. The Tivoli Storage Manager API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`) for the scheduler. You can set this option on the **Image-Snapshot** tab of the Preferences editor.

Syntax



Parameters

No Specifies that Tivoli Storage Manager allows pre-snapshot and post-snapshot commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled image snapshot backup operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator overrides the corresponding command defined in the client option file. This is the default.

Yes

Specifies that Tivoli Storage Manager allows pre-snapshot and post-snapshot commands defined by the Tivoli Storage Manager administrator to execute on the client system, when performing scheduled image snapshot backup

operations. If a pre-snapshot or a post-snapshot command is defined by both the client and the Tivoli Storage Manager administrator, the command defined by the administrator will *not* override the corresponding command defined in the client option file. This option can be used in conjunction with the `schedcmddisabled` and `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 Windows clients.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Communications** tab of the Preferences editor.

Syntax



Parameters

No Specifies that the Tivoli Storage Manager client disables SSL. *No* is the default.

Yes

Specifies that the Tivoli Storage Manager client enables SSL.

To enable SSL, specify `ssl=yes` and change the value of the `tcpport` option. Changing the value of the `tcpport` option is needed, 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 29.

Sslfipsmode

The `sslfipsmode` option specifies if 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 except for Macintosh.

Options File

Set this option in the client options file. You cannot specify it as a command line parameter (using `dsmc`) 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 IBM Tivoli Storage Manager 6.3 and newer versions. Set this client option to `no` if the client uses SSL to connect to a IBM Tivoli Storage Manager that is not at version 6.3, 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 IBM Tivoli Storage Manager version 6.3 (or newer).

Example

To enable SSL FIPS mode on the client:

```
SSLFIPSMODE yes
```

Stagingdirectory

The `stagingdirectory` option defines the location where the client can keep any data it generates in order 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.

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 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, Tivoli Storage Manager checks for the existence of the USER environment variables in the following order, and uses the first path found:

1. The path specified by the TMP user variable
2. The path specified by the TMP system variable
3. The path specified by the TEMP user variable
4. The path specified by the TEMP system variable
5. The Windows system directory

Examples

Options file:

```
stagingdirectory c:\tsmdata
```

Command line:

```
-stagingdir="e:\tsmdata"
```

Related information

“Query Aobjects” on page 577

“Restore Aobjects” on page 612

“Diffsnapshot” on page 337

“Snapdiff” on page 463

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, Tivoli Storage Manager 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 client API does not support this option.

Options File

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

Syntax



Parameters

No Subdirectories are not processed. This is the default.

Yes

Subdirectories are processed. Because the client program searches all subdirectories of a directory that is being processed, processing can take longer to complete. Specify *Yes* only when necessary.

Note: If you use the `preservepath` option in addition to `subdir=yes`, it can affect which subdirectories are processed.

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 432

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. Ask your Tivoli Storage Manager administrator to set the proper values.
2. You can use open file support in conjunction with subfile backup operations.
3. The subfilebackup option does not work correctly for migrated files. If you use a combination of subfilebackup and non-subfilebackup for migrated files, your data might be corrupted on the server.

Supported Clients

This option is valid for all Windows clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Backup** tab of the Preferences editor.

Syntax



Parameters

No Specifies that Tivoli Storage Manager does not use adaptive subfile backup. This is the default.

Yes

Specifies that Tivoli Storage Manager uses adaptive subfile backup.

Examples

Options file:

```
subfileb yes
```

Command line:

```
-subfileb=yes
```

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

Subfilecachepath

The `subfilecachepath` option specifies the path where the client cache resides for adaptive subfile backup processing.

If you do not specify a path, Tivoli Storage Manager creates a path called `\cache` under the directory where the Tivoli Storage Manager executables reside.

All directories and subdirectories in the path name you specify with the `subfilecachepath` option must exist. For example, if you specify `c:\temp\cache`, the `c:\temp` directory must already exist. You should exclude the `subfilecache` directory from backup.

Supported Clients

This option is valid for all Windows clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab of the Preferences editor.

Syntax

```
▶▶—SUBFILECACHEPath— —path_name—————▶▶
```

Parameters

path_name

Specifies the path name where the client cache resides for adaptive subfile backup processing.

Examples

Options file:

```
subfilecachep c:\temp\cc_tsm
```

Command line:

```
-subfilecachep=c:\temp\cc_tsm
```

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

Subfilecachesize

The `subfilecachesize` option specifies the client cache size for adaptive subfile backup.

If the cache size is too small, base files for some files will not be cached and subfile processing will not apply for them. However, setting the value too large can take up more disk space than can be spared. The files maintained in the cache should closely reflect the files used on a regular basis.

Supported Clients

This option is valid for all Windows clients. The server can also define this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Backup** tab of the Preferences editor.

Syntax

►►—SUBFILECACHESize— —size—◄◄

Parameters

size

Specifies the size, in megabytes, of the client cache for adaptive subfile backup processing. The range is 1 through 1024 (1 GB); the default is 10.

Examples

Options file:

```
subfilecaches 10
```

Command line:

```
-subfilecaches=10
```

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

Tapeprompt

The `tapeprompt` option specifies whether you want Tivoli Storage Manager to wait for a tape mount if it is required for a backup, archive, restore, or retrieve process, or to be prompted for a choice.

In the Tivoli Storage Manager GUI, the Media Mount dialog can display the **Information Not Available** value in the Device and Volume Label fields if you perform a standard (also known as classic) restore or retrieve operation. This value means that this information is only available for no query restore or retrieve operations; not a standard restore or retrieve operation. The Device field displays the name of the device on which to mount the media needed to process an object. The Volume Label field displays the name of the volume needed to process an object.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the **Communication** tab, **Buffer Size** field of the Preferences editor.

Syntax

►►—TCPBuffsize— *size*—

Parameters

size

Specifies the size, in kilobytes, that you want to use for the internal TCP/IP communication buffer. The range of values is 1 through 512; the default is 32.

Depending on the operating system communication settings, your system might not accept all values in the range of 1 through 512.

Examples

Options file:

```
tcpb 32
```

Command line:

```
-tcpbuffsize=32
```

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

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 client API does not support this option.

Options File

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

Syntax

►►—TCPCADAddress— *cad_address*—

Parameters

cad_address

Specifies a TCP/IP Internet domain name or a numeric IP address. You can specify `TCPip` or `V6Tcpip`. This option uses TCP/IP Version 4 if `TCPip` is specified and uses Version 6 if `V6Tcpip` is specified. You can only use IPv6 addresses if you specified the `commethod V6Tcpip` option.

Examples

Options file:

```
tcpcada dsmclnt.sanjose.ibm.com
```

Command line:

```
-tcpcadaddress=128.33.10.249  
-tcpcadaddress=khozt.mycompany.mydomain.com  
-tcpcadaddress=2002:92b:111:221:128:33:10:249
```

This option is valid only on the initial command line of the dsmcad program. It is not valid with other dsm modules.

Related information

See “Commmethod” on page 320 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 option.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (dsm.opt). You can set this option on the Scheduler tab, **Your TCP/IP address** field of the Preferences editor.

Syntax

```
▶▶—TCPCLIENTAddress— —client_address—▶▶
```

Parameters

client_address

Specifies the TCP/IP address you want the server to use to contact your client node. Specify a TCP/IP Internet domain name or a numeric IP address. The numeric IP address can be either a TCP/IPv4 or TCP/IPv6 address. You can only use IPv6 addresses if you specified the commethod *V6Tcip* option.

Examples

Command line:

```
-tcpclientaddress=128.33.10.249
-tcpclientaddress=khozt.mycompany.mydomain.com
-tcpclientaddress=2002:92b:111:221:128:33:10:249
```

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

Tcpclientport

The `tcpclientport` option specifies a TCP/IP port number for the server to contact the client when the server begins the server prompted scheduled operation.

Use this option only if you specify the *prompted* parameter with the `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 client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Scheduler** tab, **Your TCP/IP port** field of the Preferences editor.

Syntax

▶—TCPCLIENTPort— *client_port_address*—▶

Parameters

client_port_address

Specifies the TCP/IP port address you want the server to use to contact your client node. The range of values is 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.

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, **Server Address** field of the Preferences editor.

If this option is not specified, the client attempts to contact a Tivoli Storage Manager server running on the same computer as the Tivoli Storage Manager 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 V6Tcipip` option.

Examples

Options file:

```
tcps dsmchost.endicott.ibm.com
```

Command line:

```
-tcpsserveraddress=129.33.24.99
```

```
-tcpsserveraddress=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.

Tcpwindow size

Use the `tcpwindow size` option to specify, in kilobytes, the size you want to use for the TCP/IP sliding window for your client node.

The sending host cannot send more data until it receives an acknowledgment and a TCP receive window update. Each TCP packet contains the advertised TCP receive window on the connection. A larger window allows the sender to continue sending data and can improve communication performance.

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

►►—TCPWindow size— *—window_size—*—————►►

Parameters

window_size

Specifies the size, in kilobytes, to use for your client node TCP/IP sliding window. The range of values is 0 through 2048. A value of 0 allows Tivoli Storage Manager to use the operating system default TCP window size. Values from 1 to 2048 indicate that the window size is in the range of 1KB to 2MB. The default is 63. If you specify a value less than 1, the TCP window size defaults to 1. If you specify a value greater than 2048, the TCP window size defaults to 2048.

Note:

1. The TCP window acts as a buffer on the network. It is not related to the `tcpbuffsize` option, or to the send and receive buffers allocated in client or server memory.
2. A window size larger than the buffer space on the network adapter might degrade throughput due to resending packets that were lost on the adapter.
3. Depending on the operating system communication settings, your system might not accept all values in the range of values.
4. Windows 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:

```
tcpwindow size 63
```

Command line:

```
-tcpw=63
```

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

Timeformat

The `timeformat` option specifies the format in which you want to display 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 Tivoli Storage Manager supports, the Web client uses the time format for US English.

You can use the `timeformat` option with the following commands:

- `delete archive`
- `delete backup`
- `expire`
- `query archive`
- `query asr`
- `query backup`
- `query filespace`
- `query image`
- `query nas`
- `query 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 options file (`dsm.opt`). You can set this option on the **Regional Settings** tab, **Time Format** field of the Preferences editor.

Syntax

►►—`TIMEformat`— *format_number*—►►

Parameters

format_number

Displays time in one of the formats listed below. Select the format number that

corresponds to the format you want to use. When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `totime` options.

- 1 23:00:00
- 2 23,00,00
- 3 23.00.00
- 4 12:00:00 A/P
- 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 Tivoli Storage Manager saves table of contents (TOC) information for each file system backup.

You should consider the following when deciding whether you want to save TOC information:

- If you save TOC information, you can use the `QUERY TOC` server command to determine the contents of a file system backup in conjunction with the `RESTORE NODE` server command to restore individual files or directory trees.
- You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore.
- Creation of a TOC requires that you define the `TOCDESTINATION` attribute in the backup copy group for the management class to which this backup image is bound. Note that TOC creation requires additional processing, network resources, storage pool space, and possibly a mount point during the backup operation.

- If you do not save TOC information, you can still restore individual files or directory trees using the RESTORE NODE server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

Supported Clients

This option is valid for all Windows clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place the `include.fs.nas` statement containing the `toc` value in the client options file (`dsm.opt`).

Syntax



Parameters

Yes

Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. However, the backup fails if an error occurs during creation of the TOC.

No Specifies that Tivoli Storage Manager does not save TOC information during a NAS file system image backup.

Preferred

Specifies that Tivoli Storage Manager saves TOC information during a NAS file system image backup. The backup does not fail if an error occurs during creation of the TOC. This is the default.

Note: If the mode option is set to *differential* and you set the `toc` option to *preferred* or *yes*, but the last full image does not have a TOC, Tivoli Storage Manager performs a full image backup and creates a TOC.

Examples

Options file:

```
include.fs.nas netappsj/vol/vol0 homemgmtclass toc=yes
```

Command line:

```
backup nas -nasnodename=netappsj {/vol/vol0} -toc=yes
```

Todate

Use the `todate` option with the `totime` option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation

Use the `todate` and `totime` options with the `fromtime` and `fromdate` options to request a list of backed up or archived files within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2002 and 11:59 PM on July 30, 2002.

Use the `todate` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**
- **retrieve**

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►► `-TODate =— —date—` ◀◀

Parameters

date

Specifies an ending date. Enter the date in the format you selected with the `dateformat` option.

When you include `dateformat` with a command, it must precede the `fromdate`, `pitdate`, and `todate` options.

Examples

Command line:

```
dsmc restore -todate=12/11/2003 c:\myfiles\
```

Totime

Use the `totime` option with the `todate` option to specify an ending date and time to which you want to search for backups or archives during a restore, retrieve, or query operation. Tivoli Storage Manager ignores this option if you do not specify the `todate` option.

Use the `totime` and `todate` options with the `fromtime` and `fromdate` options to request a list of files that were backed up within a period of time. For example, you might request a list of files that were backed up between 6:00 AM on July 1, 2003 and 11:59 PM on July 30, 2003.

Use the `totime` option with the following commands:

- **delete backup**
- **query archive**
- **query backup**
- **restore**
- **restore group**

- retrieve

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

►—TOTime =— —*time*—◄

Parameters

time

Specifies an ending time. If you do not specify a time, the time defaults to 23:59:59. Specify the time in the format you selected with the `timeformat` option.

When you include the `timeformat` option in a command, it must precede the `fromtime`, `pittime`, and `totime` options.

Examples

Command line:

```
dsmc query backup -totime=23:59:00 -todate=06/30/2003 c:\mybackups\
```

Txnbytelimit

The `txnbytelimit` option specifies the number of kilobytes the client program buffers before it sends a transaction to the server.

A *transaction* is the unit of work exchanged between the client and server. 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 **Transaction Buffer Size** field of 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 client API does not support this option.

Syntax

►►—Type =

| |
|---------------|
| <i>any</i> |
| <i>nas</i> |
| <i>server</i> |
| <i>client</i> |

—►►

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 client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Communication** tab of the Preferences editor.

Syntax



Parameters

Yes

Specifies that the client 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.

V2archive

Use the `v2archive` option with the **archive** command to archive only files to the server.

Tivoli Storage Manager will not process directories that exist in the path of the source file specification.

This option differs from the `filesonly` option in that the `filesonly` option archives the directories that exist in the path of the source file specification.

The `v2archive` and `dirsonly` options are mutually exclusive and an error message is displayed if you use both options in the same **archive** command.

If you use this option, you might want to consider the following:

- You might experience performance problems when retrieving large amounts of data archived with this option.
- You might want to use this option only if you are concerned about expiration performance on a server that already contains extremely large amounts of archived data.
- If there are multiple files with the same name for the `v2archive` option, the files are archived multiple times, with their directory structure. The `v2archive` option archives only the files.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Syntax

▶—V2archive—▶

Parameters

There are no parameters for this option.

Examples

This command:

```
dsmc archive c:\relx\dir1\ -v2archive -su=y
```

Archives these files:

```
c:\relx\dir1\file1
c:\relx\dir1\file2
c:\relx\dir1\file3
c:\relx\dir1\dir2\file4
c:\relx\dir1\dir2\file5
```

Note: Tivoli Storage Manager does not archive `c:\relx\dir1` and `c:\relx\dir1\dir2`.

Verbose

The verbose option specifies that you want to display detailed processing information on your screen. This is the default.

When you run the **incremental**, **selective**, or **archive** commands, information is displayed about each file that is backed up. Use the quiet option if you do not want to display this information.

The following behavior applies when using the verbose and quiet options:

- If the server specifies either the quiet or verbose option in the server client option set, the server settings override the client values, even if **force** is set to *no* on the server.
- If you specify quiet in your `dsm.opt` file, and you specify `-verbose` on the command line, `-verbose` prevails.
- If you specify both `-quiet` and `-verbose` on the same command, the last option encountered during options processing prevails. If you specify `-quiet -verbose`, `-verbose` prevails. If you specify `-verbose -quiet`, `-quiet` prevails.

Supported Clients

This option is valid for all clients. The server can also define this option. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). You can set this option on the **Command Line** tab, **Do not display process information on screen** checkbox of the Preferences editor.

Syntax

▶▶—VErbose—▶▶

Parameters

There are no parameters for this option.

Examples

Options file:

```
verbose
```

Command line:

```
-verbose
```

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

Verifyimage

Use the `verifyimage` option with the **restore image** command to specify that you want to enable detection of bad sectors on the destination target volume.

If bad sectors are detected on the target volume, Tivoli Storage Manager issues a warning message on the console and in the error log.

Supported Clients

This option is valid for Windows 32-bit clients *only*. The Tivoli Storage Manager client API does not support this option.

Syntax

▶▶—VERIFYImage—▶▶

Parameters

There are no parameters for this option.

Examples

Command line:

```
dsmc restore image d: -verifyimage
```

Virtualfsname

Use the *virtualfsname* option with the **backup group** command to specify the name of the virtual file space for the group on which you want to perform the operation. The *virtualfsname* cannot be the same as an existing file space name.

Supported Clients

This option is valid for all Windows clients.

Syntax

▶▶—VIRTUALF\$name =— —*fsname*—▶▶

Parameters

fsname

Specifies the name of the container for the group on which you want to perform the operation.

Examples

Command line:

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

Virtualnodename

The `virtualnodename` option specifies the node name of your workstation when you want to restore or retrieve files to a different workstation.

When you use the `virtualnodename` option in your client options file, or with a command:

- You must specify the name you specified with the `nodename` option in your client options file (`dsm.opt`). This name should be different from the name returned by the **hostname** command on your workstation.
- Tivoli Storage Manager prompts for the password assigned to the node you specify, if a password is required. If you enter the correct password, you have access to all backups and archives that originated from the specified node.

When connecting to a server, the client must identify itself to the server. This login identification is determined in the following ways:

- If the `nodename` and `virtualnodename` options are not specified, or a virtual node name is not specified on the command line, the default login ID is the name returned by the **hostname** command.
- If the `nodename` option is specified, the name specified with the `nodename` option overrides the name returned by the **hostname** command.
- If the `virtualnodename` option is specified, or a virtual node name is specified on a command line, it cannot be the same name as the name returned by the **hostname** command.

When the virtual node name is accepted by the server, a password is required (assuming authentication is on), even if the `passwordaccess` option is *generate*. The password is not stored in the Windows Registry even when the `passwordaccess` option is set to *generate*. Once a connection to the server is established, then access is permitted to any file backed up using this login ID.

For the Java GUI, the `virtualnodename` option only works if you are root user. If you need to use the `virtualnodename` option with a non-root user account, use the command-line client (or native GUI, for platforms that still have one).

Note: The Tivoli Storage Manager 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 Tivoli Storage Manager 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.

Vmbackdir

Use the `vmbackdir` option with the **backup VM** command if you want to change where full VM control information is stored on the vStorage backup server.

Supported Clients

The directory specified by this option must be located 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.

Options File

Set this option in the client options file, or specify it on the command line as an option for the **backup vm** command.

Syntax

▶—VMBACKDir—directory—▶

Parameters

directory

Specifies the path where the full VM backup control information is stored on the vStorage backup server.

The default is `\mnt\tsmvmbbackup\fullvm\`

Examples

Options file:

```
VMBACKD c:\mnt\tsmvmbbackup\
```

Command line:

```
dsmc backup vm -VMBACKUPT=fullvm -VMBACKD=G:\san_temp\
```


Command line:

```
dsmc backup vm -vmlist=vm1 -VMBACKUPT=fullvm -VMBACKNOD=yes
```

Vmbackuptype

Use the `vmbackuptype` option with the **backup VM**, **query VM**, or **restore VM** commands to specify the type of virtual machine backup to complete.

You can specify a VMware file-level backup, a VMware full VM backup, or a Hyper-V full VM backup.

Supported Clients

This command is valid for Windows that are configured as a vStorage backup server or as a Hyper-V setup. 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***File*

Specifies a file-level backup is performed.

FULLvm

Specifies that a full VM backup is performed.

HYPERVFULL

Required for Hyper-V systems. Specifies a full virtual machine backup of all virtual machines that are defined on the Hyper-V host machine.

If the `vmbackuptype=hypervfull` option is specified, most of the other options associated with backing up files from a vStorage backup server are not allowed, or they are ignored (for example: `VMCHost`, `VMCUser`, `VMCPw`, `VMBackdir`, `VMFULLNODELETE`).

Examples

Perform a file-level virtual-machine backup of `vm1.example.com`.

```
dsmc backup vm vm1 -VMBACKUPT=File
```

Perform a full-VM image backup of virtual machine `vm1.example.com` and `vm2.example.com`.

```
dsmc backup vm vm1,vm2 -VMBACKUPT=FULLvm
```

Options file:

```
VMBACKUPT full
```

Command line:

```
dsmc backup vm vm1 -VMBACKUPT=file -vmchost=virtctr
-vmcuser=virctr_admin -vmcpw=xxxxx
```

Perform 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
```

Perform a file-level virtual-machine backup of `vm1.example.com` and `vm2.example.com` using the VMware ESX Server machine `vm serv1.example.com`, to the Tivoli Storage Manager server, using machine names `vm1` and `vm2`.

```
dsmc backup vm -domain.vmfile=vm1[xxxxxxxxx]
```

Perform a file-level virtual-machine backup of `vm1` with the IP address of VM, using the VMware ESX Server machine `vm serv1.example.com`, to the Tivoli Storage Manager server with machine name `vm1`.

```
dsmc backup vm -VMBACKUPT=full -vmhost=virtctr  
-vmcuser=virtctr_admin -vmcpw=xxxxx
```

Perform 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"
```

Perform a full virtual-machine backup of Hyper-V virtual machines named "VM 1" and "VM 2", to the Tivoli Storage Manager server.

Vmmc

Use the `vmmc` option to store the VMware backups with a management class other than the default management class. This option is only valid if `vmbackuptype=fullvm`

You can specify this option with the **backup VM** command, client options file `dsm.opt`, or the client option that is set by the server.

Supported Clients

This command 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 `dsm.sys`), or on the command line.

Syntax

```
▶▶—VMMC—class_name—————▶▶
```

Parameters

class_name

Specifies a management class that applies to the VMware backup. If you do not set this option, the default management class of the node is used.

Examples

Command line:

```
dsmc backup vm "myVirtualMachine" -vmc=myManagementClass
```

Runs a backup of the virtual machine named myVirtualMachine and saves it according to myManagementClass.

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.

This option is valid for Windows 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, 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. See the VMware document *Virtual Machine Backup Guide* that is specific for your version of ESX software for a description of the hotadd transport and the setup that is required. The *Virtual Machine Backup Guide* is available 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 508

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.

Putting the password in the client options file is a security risk because the file is not encrypted. Use the client preferences editor to save the encrypted password locally. Click **Edit** > **Client Preferences** > **Authentication** > **Save encryption key password locally**.

As an alternative to the preferences editor, you store the password locally and encrypted, 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 648

Vmctlmc

This option specifies the management class to use when backing up VMware control files. The management class that you specify here overrides the default management class, which would otherwise be used to store control files during a VMware backup.

To optimize full VM incremental backup and file recovery using Tivoli Data Protection for VMware, control files should be stored in a disk-based storage pool on the Tivoli Storage Manager server. The management class you specify can allow smaller control files to be stored in a different storage pool than the larger data files that are part of a VMware backup.

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 identified in the following technote: <https://www.ibm.com/support/docview.wss?uid=swg21497028>

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
```

Vmlist

Use the `vmlist` option with the **backup VM**, **query VM**, or **restore VM** commands to specify the host name or list of host names of the virtual machine to back up.

This option is only valid for Hyper-V backups.

Multiple values can be specified (see the following examples).

Supported Clients

This command is valid for Windows that are configured as an off-host backup of VMware virtual machine data proxy. 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— [ —[tsmnodename]— ]▶▶
```

Parameters

vmname

Specifies the virtual machine name to back up or a list of virtual machine names.

*tsmnode*name

This optional parameter can be used in situations where the name of the virtual machine and Tivoli Storage Manager node name do not match, or where DNS name resolution is not available. The option defaults to the hostname portion of the value specified with *vmname*.

Examples

Options file:

```
VMList vm1,vm2,vm3
VMList vm3.domain.com,vm8.domain.com
VMList 9.22.100.85[vm4],9.22.100.90[vm5]
VMList 9.22.100.99,9.22.100.101
VMList vm3[tsmnodeX]
```

Command line:

```
-VMList=vm1,vm2,9.88.200.4
```

Command line:

```
-vmlist='myvm1 - 2k8 x64,myvm2 - SLES10 sp2,almadenvm1_win vista'
```

Command line:

```
-vmlist="myvm1 - 2k8 x64,myvm2 - SLES10 sp2,almadenvm1_win vista"
```

Command line:

```
dsmc backup vm -vmbackuptype=HYPERVFULL -vmlist="vm1,vm2,vm3"
```

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, Tivoli Storage Manager 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 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”

Vmprocessvmwithprdm

Use this option to control whether full VMware virtual machine backups are processed if the 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, Tivoli Storage Manager 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 recreate 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 be 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 515

Vmvstortransport

The `vmvstortransport` option specifies the preferred transports order (hierarchy) to use when backing up 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 Tivoli Storage Manager 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 backup or restore virtual machine files using VADP.

Examples

If the SAN is not available, do not transport backups or restores over the LAN
VMVSTORTRANSPORT san

The Backup-Archive Client is running in a virtual machine, but do not use the hotadd transport
VMVSTORTRANSPORT nbdssl:nbd

Use the LAN transport, even if nbdssl is available, to obtain better performance
VMVSTORTRANSPORT nbd

The SAN transport is preferred, but use nbd when the SAN is not available, and do not use nbdssl or hotadd
VMVSTORTRANSPORT san:nbd

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 client 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.

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. The port value TCP/IP assigns is in the range of 1024 through 5000.

Supported Clients

This option is valid for all clients. The Tivoli Storage Manager client API does not support this option.

Options File

Place this option in the client options file (`dsm.opt`). 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 CAD Port** fields.

Syntax

```
▶▶—WEBPorts— —cadport— —agentport—▶▶
```

Parameters

cadport

Specifies the *required* Tivoli Storage Manager client acceptor service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

agentport

Specifies the *required* Tivoli Storage Manager Web client agent service port number. The range of values is 1000 through 32767. If a value is not specified, the default, zero (0), causes TCP/IP to randomly assign a free port number.

Examples

Options file:

```
webports 2123 2124
```

Command line:

```
webports 2123, 2124
```

|

Chapter 13. Using commands

Tivoli Storage Manager provides a command-line interface (CLI) that you can use as an alternative to the graphical user interface (GUI). This 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 524
- “Enter client command names, options, and parameters” on page 526
- “Wildcard characters” on page 528

The following table provides an alphabetical list of the commands and a brief description.

Table 60. Commands

| Command | Description |
|--|---|
| archive “Archive” on page 529 | Archives files from a workstation to Tivoli Storage Manager storage. |
| archive fastback “Archive FastBack” on page 532 | Archives volumes specified by the <code>fbpolicyname</code> , <code>fbclientname</code> and <code>fbvolumename</code> options for long term retention. |
| backup asr “Backup ASR” on page 535 | Generates Automated System Recovery (ASR) files in the <code>adsm.sys\ASR</code> staging directory and backs them up to the ASR file space on the server. This command is only valid for Windows Server 2003 clients. |
| backup fastback “Backup FastBack” on page 536 | Backs up volumes specified by the <code>fbpolicyname</code> , <code>fbclientname</code> and <code>fbvolumename</code> options for long term retention. |
| backup group “Backup Group” on page 538 | 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 540 | Creates an image backup of one or more file systems or logical volumes that you specify. |
| backup nas “Backup NAS” on page 545 | 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 547 | 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 vm “Backup VM” on page 549 | Backs up virtual machines specified in the <code>vm list</code> option. |
| cancel process “Cancel Process” on page 552 | 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 552 | 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 553 | 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 554 | Deletes archived files from Tivoli Storage Manager server storage. |
| delete backup “Delete Backup” on page 556 | Deletes active and inactive backup files from Tivoli Storage Manager server storage. |
| delete filespace “Delete Filespace” on page 558 | Deletes file spaces in Tivoli Storage Manager server storage. |
| delete group “Delete Group” on page 559 | Deletes a group backup on the Tivoli Storage Manager server. |
| expire “Expire” on page 561 | Inactivates backup objects that you specify in the file specification or with the <code>filelist</code> option. |
| help “Help” on page 562 | Displays a Table of Contents of help topics for the command-line client. |
| incremental “Incremental” on page 564 | 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 572 | Starts an interactive command session. |
| macro “Macro” on page 573 | Executes commands within a macro file that you specify. |
| monitor process “Monitor Process” on page 573 | 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 574 | Simulates an archive command without sending data to the server. |
| preview backup “Preview Backup” on page 575 | Simulates a backup command without sending data to the server. |
| query access “Query Access” on page 576 | Displays a list of current authorization rules. |
| query adobjects “Query Adobjects” on page 577 | Displays a list of current authorization rules. |
| query archive “Query Archive” on page 578 | Displays a list of archived files. |
| query asr “Query ASR” on page 581 | Displays information about the active backup of the Automated System Recovery files on the Tivoli Storage Manager server. This command is valid for all supported Windows clients. |
| query backup “Query Backup” on page 581 | Displays a list of backup versions. |
| query backupset “Query Backupset” on page 585 | 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. |

Table 60. Commands (continued)

| Command | Description |
|--|---|
| query filepace “Query Filespace” on page 588 | 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 590 | Displays information about group backups and their members. |
| query image “Query Image” on page 591 | Displays information about image backups. |
| query inclexcl “Query Inclexcl” on page 593 | 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 594 | Displays information about available management classes. |
| query node “Query Node” on page 595 | Displays all the nodes for which an administrative user ID has authority to perform operations. |
| query options “Query Options” on page 596 | Displays all or part of your options and their current settings. |
| query restore “Query Restore” on page 597 | Displays a list of your restartable restore sessions in the server database. |
| query schedule “Query Schedule” on page 598 | Displays information about scheduled events for your node. |
| query session “Query Session” on page 598 | 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 599 | Gathers Tivoli Storage Manager system information and outputs this information to a file or the console. |
| query systemstate “Query Systemstate” on page 601 | 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 vm “Query VM” on page 602 | Verifies the successful backups of the virtual machines from the vStorage backup server. |
| restart restore “Restart Restore” on page 604 | Displays a list of restartable restore sessions from which you can one to restart. |
| restore “Restore” on page 605 | Restores copies of backup versions of your files from a Tivoli Storage Manager server. |
| restore adobjects “Restore Adobjects” on page 612 | Restores individual Active Directory objects from the local Active Directory Deleted Objects container. |
| restore asr “Restore ASR” on page 614 | Restores the active version of the Automated System Recovery (ASR) files from the Tivoli Storage Manager server. This command is valid for Windows Server 2003 clients only. |
| restore backupset “Restore Backupset” on page 615 | 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 622 | Restores specific members or all members of a group backup. |
| restore image “Restore Image” on page 624 | Restores a file system or raw volume image backup. |

Table 60. Commands (continued)

| Command | Description |
|--|--|
| <code>restore nas</code> “Restore NAS” on page 627 | Restores the image of a file system belonging to a Network Attached Storage (NAS) file server. |
| <code>restore systemstate</code> “Restore Systemstate” on page 629 | Restores a backup of the system state. This command is valid for all supported Windows clients. |
| <code>restore vm</code> “Restore VM” on page 631 | Restores a full VM backup, and returns the full VM backup files to the <code>vmbackdir</code> directory on the vStorage backup server. |
| <code>retrieve</code> “Retrieve” on page 635 | Retrieves copies of archived files from the Tivoli Storage Manager server. |
| <code>schedule</code> “Schedule” on page 639 | Starts the client scheduler on the workstation. |
| <code>selective</code> “Selective” on page 640 | Backs up selected files. |
| <code>set access</code> “Set Access” on page 644 | 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. |
| <code>set event</code> “Set Event” on page 646 | Allows you to specify the circumstances for when archived data is deleted. |
| <code>set password</code> “Set Password” on page 648 | 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 information

Reading syntax diagrams

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. Tivoli Storage Manager processes the command and returns to the command prompt.

Use interactive mode when you want to enter a *series* of commands. Since Tivoli Storage Manager establishes connection to the server only once for interactive mode, a series of commands can be processed more quickly. Tivoli Storage Manager processes the commands and returns to the `tsm>` prompt.

Process commands in batch mode

Some options are valid *only* on the initial command line and not in interactive mode. These options generally affect the operation of the entire session.

For example, the command **dsmc query session -errorlogname=myerror.log** is accepted and it does name the error log. However, it is accepted simply because it appears in the initial command, even though the option is not valid for the query command.

There are also some options that are always valid on the initial command line as well as on individual commands in interactive mode. Therefore, certain options are accepted on the initial command line even though they have no effect on the command being entered. For example, **dsmc query session -subdir=yes** is a valid command, but in this case the *-subdir* option has no effect on the command that was entered.

When you enter a *single* command in batch mode, precede it with the executable program name, **dsmc**. For example, to process the **incremental** command in batch mode, you would enter:

```
dsmc incremental
```

Tivoli Storage Manager prompts you each time you enter a command if the *passwordaccess* option is set to *prompt* and authentication on the server is set to *On*. Type your password and press Enter.

You can also enter your password using the password option with a command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc incremental -password=secret
```

If you set the *passwordaccess* option to *generate* in your *dsm.opt* file, you do not need to specify the password with the command. Tivoli Storage Manager *only* prompts you for your password if you are registering your workstation with a server or manually changing your password.

Related concepts

Chapter 12, "Processing options," on page 269

Process commands in interactive mode

Use the *interactive* mode (or *loop* mode) to enter a series of commands.

Enter **dsmc** on the command line and press Enter. When the **tsm>** command prompt appears, type the command name and press Enter. *Do not* precede each command with the executable program name, **dsmc**. Alternatively, you can enter **dsmc loop** on the command line to start a client command session in interactive mode. **Loop** is the default command for **dsmc**.

If a password is required, Tivoli Storage Manager prompts you before you enter the first command.

Type your password and press Enter.

You can also enter your password using the password option with the **loop** command, but your password appears on the screen. For example, if your password is *secret*, enter:

```
dsmc loop -password=secret
```

To end an interactive session, enter **quit** at the prompt.

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 298

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, Tivoli Storage Manager uses the default value. For example, the **restore** command includes a required parameter, **sourcefilespec**, that specifies the path and file name in storage that you want to restore. The optional parameter, **destinationfilespec**, specifies the path where you want to place the restored files. If you do not specify the **destinationfilespec**, by default Tivoli Storage Manager restores the files to the original source path. If you want to restore the files to a *different* directory, enter a value for **destinationfilespec**.

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 below 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 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 either VSS or LVSA, the Tivoli Storage Manager 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 above. The snapshot volume name can be up to 1024 bytes.

- When entering the **sourcefilespec**, if the directory name ends with \, then * is implied.

When entering a **destinationfilespec**, if the name ends with /, then it is considered a directory, otherwise it is considered a file.

The following example illustrates these two rules. Even though `mydir` and `yourdir` are directories, the command will fail because `/*` is implied after `mydir`, and `yourdir` is considered a file:

```
restore c:\home\mydir\ c:\away\yourdir
```

The following example illustrates the second rule. Even though `mydir` and `yourdir` are directories, the command will fail because `mydir` and `yourdir` are considered files:

```
restore c:\home\mydir c:\away\yourdir
```

- If a file specification does not begin with a file space name (an opening directory delimiter), the file specification is assumed to be a subdirectory of the current working directory. The client appends the destination 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 the file specification contains spaces, it must be enclosed in quotation marks. For example:

```
dsmc sel "x:\dir one\file1"
```

However, when the filespec is a directory with a trailing backslash, an extra backslash (\) must be added. If an extra backslash is not added, the filespec will not be processed correctly, and the operation might cause unexpected results.

The following example is incorrect:

```
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 directories 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 377

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 | abdefjrs | 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 Tivoli Storage Manager 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 in conjunction with a 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.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

filespec

Specifies the path and name of the file you want to archive. You can use wildcards to specify groups of files or all the files in a directory. Separate file specifications with a space. You can also use the `filelist` option to process a list of files. The Tivoli Storage Manager client opens the file that you specify with this option and processes the list of files within according to the specific command.

You can specify as many file specifications as available resources or other operating system limits permit.

Table 62. Archive command: Related options

| Option | Where to use |
|-------------------------------------|--|
| <code>archmc</code> | Command line only. |
| <code>autofsrename</code> | Client options file (<code>dsm.opt</code>) only. |
| <code>changingretries</code> | Client options file (<code>dsm.opt</code>) or command line. |
| <code>compressalways</code> | Client options file (<code>dsm.opt</code>) or command line. |
| <code>compression</code> | Client options file (<code>dsm.opt</code>) or command line. |
| <code>deletefiles</code> | Command line only. |
| <code>description</code> | Command line only. |
| <code>dironly</code> | Command line only. |
| <code>encryptiontype</code> | Client options file (<code>dsm.opt</code>). |
| <code>encryptkey</code> | Client options file (<code>dsm.opt</code>). |
| <code>filelist</code> | Command line only. |
| <code>filesonly</code> | Command line only. |
| <code>postsnapshotcmd</code> | Client options file (<code>dsm.opt</code>) or with the <code>include.fs</code> option. |
| <code>preservelastaccessdate</code> | Client options file (<code>dsm.opt</code>) or command line. |
| <code>presnapshotcmd</code> | Client options file (<code>dsm.opt</code>) or with the <code>include.fs</code> option. |
| <code>skipntpermissions</code> | Client options file (<code>dsm.opt</code>) or command line. |

Table 62. Archive command: Related options (continued)

| Option | Where to use |
|-----------------------|--|
| skipntsecuritycrc | Client options file (dsm.opt) or command line. |
| snapshotcachelocation | Client options file (dsm.opt) or with the include.fs option. |
| snapshotcachesize | Client options file (dsm.opt) or include.fs option. |
| snapshotfsidleretries | Client options file (dsm.opt) or with include.fs option. |
| snapshotfsidlewait | Client options file (dsm.opt) or with include.fs option. |
| snapshotroot | Command line only. |
| subdir | Client options file (dsm.opt) or command line. |
| tapeprompt | Client options file (dsm.opt) or command line. |
| v2archive | Command line only. |

Examples

Task Archive a single file 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 information

“Configuring Open File Support” on page 66

“Include options” on page 393

Open file support

If open file support has been configured, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or "in use") by other applications.

The snapshot allows the archive to be taken from a point-in-time copy that matches the file system at the time the snapshot is taken. Subsequent changes to

the file system are not included in the archive. You can set the **snapshotproviderfs** parameter of the `include.fs` option to `none` to specify which drives do not use open file support.

To control an open file support operation with LVSA, you can specify these additional options in your `dsm.opt` file or as values of the `include.fs` option: `snapshotcachelocation`, `snapshotcachesize`, `snapshotfsidleretries`, `snapshotfsidlewait`, `presnapshotcmd`, `postsnapshotcmd`.

You can also use VSS for open file support. If you use VSS, you do not need to install LVSA.

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Use the `snapshotcachelocation` option to relocate the cache if necessary. You can specify a snapshot cache location for a specific drive using the `include.fs` option.
3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if LVSA was not installed.
5. To enable open file support in a cluster environment, all systems in the cluster should have LVSA 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 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 Tivoli Storage Manager by the Windows Tivoli Storage Manager client.

If a policy specification contains both Windows and Linux FastBack clients, only the Linux volumes will be backed up or archived to Tivoli Storage Manager by the Linux Tivoli Storage Manager 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 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.
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 371 | Command line and scheduler. |
| fbserver "fbserver" on page 374 | Command line and scheduler. |
| fbclientname "fbclientname" on page 370 | Command line and scheduler. |
| fbvolumename "fbvolumename" on page 375 | Command line and scheduler. |

Table 63. Archive FastBack command: Related options (continued)

| Option | Where to use |
|---|-----------------------------|
| fbreposlocation "Fbreposlocation" on page 373 | Command line and scheduler. |
| fbbranch "Fbbranch" on page 369 | Command line and scheduler. |
| archmc "Archmc" on page 298 | Command line and scheduler. |

Examples

Command line:

The Tivoli Storage Manager 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 Tivoli Storage Manager 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 FBBbranch option

Command line:

The Tivoli Storage Manager client is installed on a dedicated proxy machine with Tivoli Storage Manager FastBack administrative command line and FastBack mount. The Tivoli Storage Manager 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 Tivoli Storage Manager client is installed on a dedicated proxy machine with Tivoli Storage Manager FastBack administrative command line and FastBack mount. The Tivoli Storage Manager 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"
```

Backup ASR

Use the **backup asr** command to generate Automated System Recovery (ASR) files in the adsm.sys\ASR staging directory and back them up to the ASR file space on the Tivoli Storage Manager server.

This command generates the asr.sif file that contains instructions required for Tivoli Storage Manager to restore the files necessary during ASR recovery mode.

Note: ASR recovery requires a valid system state backup, and this command does not perform that backup automatically. For Windows Server 2003, use the **backup systemstate** command to create the asr.sif file before you use the **Backup ASR** command to back up that file on the server.

Supported Clients

This command is valid for Windows Server 2003 clients only.

Syntax

▶▶ Backup ASR ◀◀

Parameters

There are no parameters for this command.

Examples

Task Back up the ASR files to the Tivoli Storage Manager server.

Command: backup asr

Related information

Backup FastBack

Use the **backup fastback** command to back up Tivoli Storage Manager FastBack volumes specified by the `fbpolicyname`, `fbclientname` and `fbvolumename` options for long-term retention.

Before using this command, configure the client to back up and archive Tivoli Storage Manager FastBack data. Also, before you issue this command, at least one snapshot should exist in the Tivoli Storage Manager FastBack repository for the Tivoli Storage Manager FastBack policy being archived or backed up.

If a policy specification contains both Windows and Linux FastBack clients, only the Windows volumes will be backed up or archived to Tivoli Storage Manager by the Windows Tivoli Storage Manager 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 371 | Command line and scheduler. |
| fbserver "Fbserver" on page 374 | Command line and scheduler. |
| fbclientname "Fbclientname" on page 370 | Command line and scheduler. |
| fbvolumename "Fbvolumename" on page 375 | Command line and scheduler. |
| fbreposlocation "Fbreposlocation" on page 373 | Command line and scheduler. |
| fbbranch "Fbbranch" on page 369 | Command line and scheduler. |
| backmc "Backmc" on page 309 | Command line and scheduler. |

Examples

Command line:

The Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager client is installed on a dedicated proxy machine with FastBack administrative command line and FastBack mount. The Tivoli Storage Manager 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  
-fbrepositlocation=\\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 Tivoli Storage Manager client is installed on a dedicated proxy machine with FastBack administrative command line and FastBack mount. The Tivoli Storage Manager 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  
-fbrepositlocation=\\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"
```

Backup Group

Use the **backup group** command to create and back up a group containing a list of files from one or more file space origins to a virtual file space on the Tivoli Storage Manager server.

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

- All objects in the group are assigned to the same management class.
- 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 backup are assigned as members of the differential group backup. This data is not resent to the server, reducing backup time.

The **backup group** command requires the following options:

filelist

Specifies a list of files to add to a new group.

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 Tivoli Storage Manager server.
4. Unless you are running Mac OS X, use the **delete group** command to delete a specific group backup from the Tivoli Storage Manager server.
5. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager 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 377 | Command line only. |
| groupname "Groupname" on page 384 | Command line only. |
| mode "Mode" on page 412 | Command line only. |
| snapshotproviderfs "Snapshotproviderfs" on page 472 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | Client options file (dsm.opt) or with include.image option. |
| virtualfsname "Virtualfsname" on page 504 | 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 393

"Query Group" on page 590

"Restore Group" on page 622

"Delete Group" on page 559

"Query Filespace" on page 588

Backup Image

The **backup image** command creates an image backup of one or more volumes on your system.

These volumes can be formatted FAT, FAT32, NTFS, or unformatted RAW volumes. If a volume is NTFS-formatted, only those blocks used by the file system 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 with *mode=incremental*.

The client backs up the files whose modification date and time (client time) is later than the date and time of the last incremental backup of the file system on which the file resides (server time).

If the server time is ahead of the client time, an incremental-by-date operation or image backup with *mode=incremental* skips 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 running the Tivoli Storage Manager client must have administrator authority to successfully perform any type of image backup.
2. The API must be installed to use the **backup image** command.

Use the `include.image` option to include a file system or logical volume for image backup, or to specify volume-specific options for image backup.

The **backup image** command uses the `compression` option.

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 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 point assigned or a drive letter assigned. 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 322 | Client options file (dsm.opt) or command line. |
| compression "Compression" on page 323 | Client options file or command line. |
| imagegapsize "Imagegapsize" on page 389 | Use with the backup image command, the include.image option, or in the options file. |
| mode "Mode" on page 412 | Command line only. |
| postsnapshotcmd "Postsnapshotcmd" on page 428 | Use with the backup image command, the include.image option , or in the options file. |
| presnapshotcmd "Presnapshotcmd" on page 434 | Use with the backup image command, the include.image option , or in the options file. |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | Use with the backup image command, the include.image option , or in the options file. |
| snapshotcachesize "Snapshotcachesize" on page 468 | Use with the backup image command, the include.image option , or in the options file. |
| snapshotfsidretries "Snapshotfsidretries" on page 469 | Use with the backup image command, the include.image option , or in the options file. |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | Use with the backup image command, the include.image option , or in the options file. |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | 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.

Command: dsmc backup image m:\mnt\myntfs

Task Back up the h: drive using an image incremental backup that backs up only new and changed files after the last full image backup.

Command: dsmc backup image h: -mode=incremental

Task Perform an offline image backup of the f: drive.

Command: dsmc backup image f: -snapshotproviderimage=none

Task Perform an online image backup of the f: drive.

Command:

dsmc backup image f: -snapshotproviderimage=LVSA

Task Back up the f: drive which is mapped to a volume which has not been formatted with a file system.

Command: `dsmc backup image f:`

Related information

“Imagegapsize” on page 389

“Snapshotproviderimage” on page 473

“Configuring Open File Support” on page 66

“Image backup” on page 129 to install and configure the LVSA

“Snapshotcachesize” on page 468

“Mode” on page 412

“Comparing methods 1 and 2” on page 133 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, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or "in use") by other applications. To control an open file support operation with LVSA, you can specify these additional options in your `dsm.opt` file or as values of the `include.fs` option: `snapshotcachelocation`, `snapshotcachesize`, `snapshotfsidleretries`, `snapshotfsidlewait`, `snapshotproviderfs`, `presnapshotcmd`, `postsnapshotcmd`.

You can also use VSS for open file support. If you use VSS, you do not need to install LVSA.

The following considerations apply to offline and online image backups:

- If you create an image of the system drive, you cannot restore it to the original location. Restore of any image requires that the client have an exclusive lock of the volume you are restoring to, so the system drive cannot be restored since Tivoli Storage Manager is unable to lock the system drive. You can restore an image backup of the system drive to an alternate location.
- Because of different system component configurations, the system image not be consistent across components (such as Active Directory). Some of these components can be configured to use different volumes where parts are installed on the system drive and others to non-system volumes.
- 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. Tivoli Storage Manager will not back up a volume without a mount point or drive letter.
- If bad disk sectors are detected on the source drive during a LAN-free or LAN-based image backup, data corruption occur. In this case, bad sectors are

skipped when sending image data to the Tivoli Storage Manager server. If bad disk sectors are detected during the image backup, a warning message is issued after the image backup completes.

Utilizing image backup to perform file system incremental backup

There are two methods of utilizing image backups to perform efficient incremental backups of your file system. These backup methods allow you to perform point-in-time restore of your file systems and improve backup and restore performance.

You can perform the backup only on formatted volumes; not on raw logical volumes. You can 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, Tivoli Storage Manager ignores the `deletefiles` option when the `image+image incremental` technique of backing up has been used. The restore will include files that were deleted after the last full image backup plus the latest versions of files added or changed after the last image backup.

Note: You should perform full image backups periodically in the following cases. 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 belonging to a Network Attached Storage (NAS) file server, otherwise known as NDMP Backup. Tivoli Storage Manager prompts you for an 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 Tivoli Storage Manager 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 in conjunction with the **RESTORE NODE** server command to restore individual files or directory trees.

You can also use the Tivoli Storage Manager Web client to examine the entire file system tree and select files and directories to restore. Creation of a TOC requires that you define the **tocdestination** attribute in the backup copy group for the management class to which this backup image is bound. 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, 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 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 authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS backup processing.


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, Tivoli Storage Manager processes all of the file systems 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 412 | Command line only. |
| monitor "Monitor" on page 414 | Command line only. |
| nasnodename "Nasnodename" on page 416 | Client options file (dsm.opt) or command line. |
| toc "Toc" on page 496 | Command line or with the include.fs.nas 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=netappsj {/vol/vol0} -toc=yes

Related information

"Nasnodename" on page 416

"Toc" on page 496

"Mode" on page 412

"Monitor" on page 414

"Cancel Process" on page 552

"Domain.nas" on page 346

Backup Systemstate

Use the **backup systemstate** command to back up all startable system state and system services components as a single object, to provide a consistent point-in-time snapshot of the system state.

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

The list of startable system state and system services components is dynamic and can change depending on service pack and operating system features installed. Tivoli Storage Manager allows for the dynamic discovery and backup of these components.

The system state backup consists of data from several VSS writers. When the files that belong to the System Writer change, an incremental backup is used for these files. Using incremental backup on the System Writer files reduces the amount of time it takes to back up the system state. A full backup is used for all other system state data.

During system state backup, a progressive incremental backup is used for the System Writer files in the system state. If you want to run a full backup of all system state data, you can specify `include.systemstate mc_name` in the client options file (`dsm.opt`), where `mc_name` is the name of the management class with copy group mode absolute.

On Windows Vista, Windows 7, and Windows Server 2008 clients running BIOS, **backup systemstate** also backs up ASR data. ASR data is not backed up on 64-bit architectures with Extensible Firmware Interface (EFI).

Note:

1. To back up Windows system state, your client must be connected to a Tivoli Storage Manager Version 5.2.0 or higher server.
2. The system and boot files component of system state is backed up only if a member (file) of that component has changed since the last backup. If a member changes, the entire group of files that comprise that component are backed up.
3. The Tivoli Storage Manager Windows client does not allow the backup of any individual component.
4. Use the `include.systemstate` option in your client options file (`dsm.opt`) to assign management classes for backup of system state. The default is to bind the system state object to the default management class.
For example: `include.systemstate ALL BASVT2`.
5. Use the **query systemstate** command to display information about a backup of the system state on the Tivoli Storage Manager server.
6. 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 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

“Backing up in preparation for Automated System Recovery (Windows Vista, Windows 7, and Windows Server 2008)” on page 125

“Backing up in preparation for Automated System Recovery (Windows Server 2003)” on page 127

“Query Systemstate” on page 601

“Restore Systemstate” on page 629

Backup VM

Use the **backup vm** command to start a full backup of a virtual machine (backup a vmdk).

Use the **backup vm** command to back up VMware or Hyper-V virtual machines.

Issue this command from the Backup-Archive Client on the vStorage backup server for the ESX server that you want to back up.

VMware backup

On VMware, one or more virtual machines are processed using the IBM Tivoli Storage Manager data mover node. The data mover node is the name given to a configuration where the Backup-Archive Client runs on the vStorage backup server. The backup files are stored as if they were backed up from within the virtual machines. You must configure the VMware virtual machine before using this command, see “Preparing the environment for VMware backup processing” on page 143.

You can run a full VM backup which stores a backup copy of complete virtual disk images and configuration information for a virtual machine. Full VM backup enables a complete restore of the complete virtual machine, although it takes more time and more IBM Tivoli Storage Manager server space.

If you install IBM Tivoli Storage Manager for Virtual Environments, you can also specify the **-mode** option to run an incremental or full backup when **backuptype=fullvm**.

A Full VM backup uses VMware Change Block Tracking (CBT) support to enable content-aware (used-block only) backups. VMware CBT requires an ESX 4.0 or later host (with virtual hardware 7). You cannot perform Full VM backups on virtual machines that do not support CBT.

For ESX 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 **Backup** window in the backup-archive GUI, or use the `dsmc show vm all client` command to view the current CBT status. CBT status has 3 values:

Off Indicates the CBT configuration parameter (**ctkEnabled**) has not been 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 it is enabled in the virtual machine's configuration parameters (`ctkEnabled=true`). The client enables CBT when you perform your first full VM backup of a virtual machine. After the first full VM backup is performed, only the changed blocks on the disk are backed up or restored. Once the status has been changed to on, it remains on, even if the virtual machine is deleted from the Tivoli Storage Manager server. To turn it off, you must set the **ctkEnabled** configuration parameter to false. In the vSphere client, right-click the virtual machine that you want to change the CBT status for. Click **Edit Settings > Options > General > Configuration Parameters** and set the **ctkEnabled** configuration parameter to false.

Tip: You can only use the compression option with backups if the backup is being saved to a storage pool that has been enabled for deduplication. See “Determine compression and encryption processing” on page 78 for more information. See also “Compression, encryption, and adaptive subfile backup processing” on page 398 for more information.

Attention: IBM Tivoli Storage Manager Backup-Archive Client cannot back up vCenter virtual machine templates. Also, virtual machines that are deployed as vApps are not included in full VM backups.

File-level VM backup provides individual file restore within the virtual machine and incremental backup, although it does not have an easy full machine recovery procedure. You might want to use a combination of file level VM backups with periodic full VM backups for optimal results.

Hyper-V backup

For Hyper-V backup, use the `vm list` option to specify the name of the virtual machine, or a list of virtual machine names to back up. If `-vm list=?` is specified, then a listing of all the Hyper-V virtual machines defined on the Hyper-V server is printed to the terminal.

For Hyper-V backup, you must specify `HyperVFull` on the `vm backup type` option.

Supported Clients

This command is valid only on supported Windows clients that are installed on a vStorage backup server to protect Hyper-V or VMware virtual machines.

Syntax

```
► Backup VM [VMNAME] [options] ►
```

Parameters

VMNAME

Specify the name of the virtual machine. The name can be the DNS host name or the virtual machine display name. If you do not specify VMNAME, you can identify the virtual machine with one of the following options:

- domain.vmfile
- domain.vmfull
- vmlist

vStorage APIs for data protection examples

```
dsmc backup vm vm1 -vmbackuptype=file
```

```
dsmc backup vm vm3,vm4 -vmbackuptype=fullvm
```

```
dsmc backup vm vmlocal -vmbackuptype=fullvm
```

To run a file-level virtual machine backup of vm1.example.com using the VMware VirtualCenter machine virtctr.example.com, see the following example:

```
dsmc backup vm vm1 -vmbackuptype=file -vmhost=virtctr
```

Hyper-V examples

```
dsmc backup vm -vmbackuptype=hypervfull vm1
```

Full virtual machine backup of virtual machine named vm1.

```
dsmc backup vm -vmbackuptype=hypervfull -vmlist=vm1,vm2
```

Full virtual machine backup of virtual machines named vm1 and vm2.

```
dsmc backup vm -vmbackuptype=hypervfull
```

Backs up all virtual machines defined on the Hyper-V host machine.

```
dsmc backup vm -vmbackuptype=hypervfull -vmlist=?
```

Prints out a list of all virtual machines defined on the Hyper-V host machine.

Related reference

“Query VM” on page 602

“Restore VM” on page 631

“Domain.vmfull” on page 350

“Domain.vmfile” on page 348

“Mode” on page 412

“Vmbackdir” on page 506

“Vmbacknodelete” on page 507

“Vmbackuptype” on page 508

“Vmchost” on page 511

“Vmcpw” on page 512

“Vmouser” on page 513

“Vmmc” on page 509

“Vmlist” on page 514

“Vmprocessvmwithindependent” on page 515

“Vmprocessvmwithprdm” on page 516

“Vmvstortransport” on page 517

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. Tivoli Storage Manager prompts you for an 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 only cancel one restartable restore session at a time. Run the **cancel restore** command again to cancel additional restores. To restart restartable restore sessions, use the **restart restore** command.

Use the **cancel restore** command under the following circumstances:

- You cannot back up files affected by the restartable restore.
- You want to cancel restartable restore sessions.
- Restartable restore sessions lock the file space so that files cannot be moved off of the 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.

Command:

cancel restore

Delete Access

The **delete access** command deletes authorization rules for files that are stored on the server.

When you delete an authorization rule, you revoke user access to any files or images 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 you want to delete.

Command:

delete access

See the following screen example:

| Index | Type | Node | Owner | Path |
|-------|---------|-------|--------|-----------------------|
| 1 | Backup | node1 | daisy | c:\dev\proja\list.c |
| 2 | Archive | node3 | marm | c:\fin\budg\depta.jan |
| 3 | Backup | node4 | susie | c:\plan\exp\deptc.feb |
| 4 | Archive | node5 | susies | c:\mfg\invn\parta.wip |

Enter Index of rule(s) to delete, or quit to cancel:

To delete the authorization rules that allow *marm* and *susies* to access your files, type **2 4** or **2,4**, then press Enter.

Delete Archive

The **delete archive** command deletes archived files from Tivoli Storage Manager server storage. Your administrator must give you the authority to delete archived files.

Important: When you delete archived files, you cannot retrieve them. Verify that the files are obsolete before you delete them.

Supported Clients

This command is valid for all clients.

Syntax

```
►► Delete ARchive [ —options ] [ —filespec ] —filespec —{—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 Tivoli Storage Manager client opens the file 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 has changed, or if you are deleting files 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 file space name 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 68. Delete Archive command: Related options

| Option | Where to use |
|---|--|
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| description "Description" on page 334 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| noprompt "Noprompt" on page 419 | Command line only. |
| numberformat "Numberformat" on page 419 | Client options file (dsm.opt) or command line. |
| pick "Pick" on page 424 | Command line only. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |

Examples

Task Delete files from file space abc in the proj directory.

Command: delete archive {"abc"}\proj*

Task Delete a file named budget.

Command: delete archive c:\plan\proj1\budget.jan

Task Delete all files archived from the c:\plan\proj1 directory with a file extension of .txt.

Command: delete archive c:\plan\proj1*.txt

Task Delete files archived from the c:\project directory 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.

Command: delete archive c:\project* -pick

Task Delete selected files from the group of files archived with the description "Monthly Budgets 1999" located in c:\projects and its subdirectories.

Command: delete ar c:\projects* -description="Monthly Budgets 1999" -pick -subdir=yes

Related information

"Filelist" on page 377

Delete Backup

The **delete backup** command deletes backup files from Tivoli Storage Manager server storage. Your administrator must give you authority to delete backup files.


You can delete your backed up files from the Tivoli Storage Manager client node workstation if the function is enabled on the definition of the Tivoli Storage Manager node. When you do this, the Tivoli Storage Manager server takes all of the backed up files that meet the 'filespec' and 'deltype' specified and deactivate 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. 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.

Supported Clients

This command is valid for all clients.

Syntax



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. Separate file specifications with a space. You can also use the `filelist` option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command.

Note: If you indicate *filespace name*, do not include a drive letter in the file specification.

When using `-deltype=inactive` or `-deltype=active`, use wildcard characters to specify a group of files or all files in a directory.

When using `-deltype=all`, specify a fully-wildcarded directory.

objtype

Specifies the type of backup delete you want to perform. You can specify either of the following values:

FILE

Specifies that you want to delete directories and files. This is the default.

IMAGE

Specifies that you want to delete an image backup. Specifies that you want to delete an image backup. Objtype=image is not supported on Mac OS X.

delttype

Specifies the deletion type. Specify one of the following values:

ACTIVE

Delete only active file objects. Directory objects are not deleted. This is the default.

Note: If there are any inactive objects, then after the active object is deleted, the most current inactive object is changed from inactive to active.

To delete all versions of a file, first issue the **delete backup** command with `-delttype=inactive`, then issue the command again with `-delttype=active`.

INACTIVE

Delete only inactive file objects. Directory objects are not deleted.

ALL

Delete all active and inactive objects below a given directory, including all subdirectories and their files.

Note: The parent directory of the deleted files and subdirectories is not deleted. If you specify `delttype=ALL`, you cannot use the `pick` option because `delttype=ALL` and the `pick` option are mutually exclusive.

Table 69. Delete Backup command: Related options

| Option | Where to use |
|--|--|
| description "Description" on page 334 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| fromdate "Fromdate" on page 381 | Command line, and in the GUI find function. |
| fromtime "Fromtime" on page 383 | Command line, and in the GUI find function. |
| noprompt "Noprompt" on page 419 | Command line only. |
| pick "Pick" on page 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line, and in the GUI find function. |
| pittime "Pittime" on page 426 | Command line, and in the GUI find function. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |

Table 69. Delete Backup command: Related options (continued)

| Option | Where to use |
|-----------------------------|---|
| todate "Todate" on page 497 | Command line, and in the GUI find function. |
| totime "Totime" on page 498 | 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 names ending in .txt backed up from the c:\plan\proj1 directory and its subdirectories.

Command: delete backup c:\plan\proj1*.txt -deltype=inactive -subdir=yes

Task Delete selected active files backed up from the c:\project directory. Use the -pick option to display a list of backup copies that match the file specification. From the list, you can select which versions to delete.

Command: delete backup c:\project* -pick

Task Delete all active and inactive versions of files and subdirectories in c:\user\myproject.

Command: delete backup c:\user\myproject* -deltype=all

Note: The backup versions of directory object c:\user\myproject are not deleted.

Related information

"Filelist" on page 377

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 70. Delete Filespace command: Related options

| Option | Where to use |
|--|---|
| <code>class</code> “Class” on page 313 | Command line only. |
| <code>detail</code> “Detail” on page 335 | Command line only. |
| <code>nasnodename</code> “Nasnodename” on page 416 | Client options file or command line. |
| <code>scrolllines</code> “Scrolllines” on page 455 | Client options file or command line. |
| <code>scrollprompt</code> “Scrollprompt” on page 456 | 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 416

“Class” on page 313

Delete Group

Use the **delete group** command to delete a group backup on the Tivoli Storage Manager server.

After deleting a group, the group leader (`virtualfsname`) remains on the Tivoli Storage Manager server. It contains no members (file or directories) but is reported in a subsequent **query filesystem** command. There are no files listed if the `showmembers` option is added. Deleting a group does not remove the file space that

it resides in because there might be other groups in it. Use **delete filespace** if you want to remove the file space and all the data it contains.


Note:

1. Use the `inactive` option to display both active and inactive group backup versions. By default, Tivoli Storage Manager only 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, Tivoli Storage Manager prompts you for confirmation before deleting the group backup. Using this option can speed up the delete procedure. However, it also increases the danger of accidentally deleting a group backup version that you want to save. Use this option with caution.
4. Use the **query filespace** command to display virtual file space names for your node that are stored on the Tivoli Storage Manager 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 71. Delete Group command: Related options

| Option | Where to use |
|--|--------------------|
| <code>inactive</code> “Inactive” on page 391 | Command line only. |
| <code>noprompt</code> “Noprompt” on page 419 | Command line only. |
| <code>pick</code> “Pick” on page 424 | Command line only. |
| <code>pitdate</code> “Pitdate” on page 425 | Command line only. |
| <code>pittime</code> “Pittime” on page 426 | 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 391

“Pick” on page 424

“Noprompt” on page 419

“Query Filespace” on page 588

Expire

The **expire** command inactivates the backup objects you specify in the file specification or with the `filelist` option.

When working in interactive mode, a prompt notifies you before files are expired.

The **expire** command does not remove workstation files. If you expire a file or directory that still exists on your workstation, the file or directory is backed up again during the next incremental backup unless you exclude the object from backup processing.

If you expire a directory that contains active files, those files will not appear in a subsequent query from the GUI. However, these files are displayed on the command line, if you specify the proper query with a wildcard character for the directory.

Note: Because the **expire** command changes the server picture of the client file system without actually changing the client file system, the **expire** command is not allowed on files located on a file system monitored by the Tivoli Storage Manager journal service.

Supported Clients

This command is valid for all clients.

Syntax

```

▶▶—EXPire ————— —filespec————▶▶
      └──options──┘

```

Parameters

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.

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 this case the combination should be unique. Non-unique abbreviations result in the display of the first section of the entire help file matching 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 might 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 only 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 have been 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 less than 72 characters causes sentences that are 72 characters wide to wrap to the next line. This can cause the displayed help text to begin somewhere within the section rather than at the beginning. The undisplayed lines can be viewed by using the scrolling function of the terminal to move up.

Examples

Task Display the table of contents of the help topics.

Command: `dsmc help`

Task Display the information in help topic 2.1.2

Command: `dsmc help 2.1.2`

Task Display help information on the **archive** command.

Command: `dsmc help archive`

Task Display help information on message ANS1036.

Command: `dsmc help 1036`

Incremental

The **incremental** command backs up all new or changed data in the locations that you specify, unless you exclude them from backup services.

You can back up all new or changed files or directories in the default client domain or from file systems, directories, or files.

To incrementally back up selected files or directories, enter a file specification in the command. If you do not enter a file specification, the default is to back up files or directories in the default domain.

The following attributes in the management class assigned to the file or directory affect whether the data is actually backed up:

Frequency

The number of days that must elapse between successive backups for the file. The **frequency** attribute applies only to a full incremental backup.

This management class attribute is ignored during a journal-based backup.

Mode Permits you to back up only files that changed since the last backup (*modified*) or back up the files whether they changed or not (*absolute*).

Serialization

Permits or denies backup of files or directories according to the following values:

- **static**: In order to be backed up, data must not be modified during backup or archive.
- **shared static**: If data in the file or directory changes during each of the allowed attempts to back up or archive it, it is not backed up or archived. The value of the `changingretries` option determines how many attempts are made. The default is 4.
- **dynamic**: The object is backed up or archived on the first attempt whether or not data changes during the process.
- **shared dynamic**: The object is backed up or archived on the last attempt, even if data changes during the process.

Using the `include` option in an include-exclude list, you can override the default management class for a file or group of files.

You can perform either a *full incremental* backup or an *incremental by date* backup. The default is a full incremental backup.

If you are journaling a file system and the journal is valid, the full incremental backup performs a journal-based backup. More than one journal-based backup session can be started, but only one journal-based backup session can proceed. All other journal-based backup sessions that need access to the same filespace must wait until the current journal-based backup session has completed before the next session can proceed. 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 *selective* backup that backs up only the files, directories or empty directories that you specify regardless of whether they have changed.


A full incremental backs up all files or directories that are new or have changed since the last incremental backup. During a full incremental backup, the client queries the server or the journal database. Tivoli Storage Manager uses this information to:

- Back up new files or directories.
- Back up files or directories whose contents have changed.
- Mark inactive backup versions on the server for files or directories that are deleted from the workstation.
- Rebind backup versions to management classes if the management class assignments change.

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 Tivoli Storage Manager client opens the file 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 default domain or the domain specified as an option is backed up.

If you specify a file system, all new and changed files are backed up. In addition, the last incremental date for the file space is updated on the server. If you specify a file or directory, the last incremental date is not updated. This means the file or directory might be backed up again if a later backup is performed using the `incrbydate` option. If you specify a file system, specify the file system without a trailing slash.

Table 73. Incremental command: Related options

| Option | Where to use |
|--|--|
| <code>autofsrename</code> "Autofsrename" on page 307 | Client options file (<code>dsm.opt</code>) only. |
| <code>changingretries</code> "Changingretries" on page 312 | Client options file (<code>dsm.opt</code>) or command line. |
| <code>compressalways</code> "Compressalways" on page 322 | Client options file (<code>dsm.opt</code>) or command line. |
| <code>compression</code> "Compression" on page 323 | Client options file (<code>dsm.opt</code>) or command line. |
| <code>diffsnapshot</code> "Diffsnapshot" on page 337 | Command line only. |
| <code>dirsonly</code> "Dirsonly" on page 340 | Command line only. |
| <code>domain</code> "Domain" on page 343 | Client options file (<code>dsm.opt</code>) or command line only. |

Table 73. Incremental command: Related options (continued)

| Option | Where to use |
|---|---|
| encryptiontype "Encryptiontype" on page 356 | Client options file (dsm.opt). |
| encryptkey "Encryptkey" on page 357 | Client options file (dsm.opt). |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| incrbydate "Incrbydate" on page 399 | Command line only. |
| memoryefficientbackup "Memoryefficientbackup" on page 411 | Client user-options file (dsm.opt), server, or command line. |
| nojournal "Nojournal" on page 418 | Command line only. |
| postsnapshotcmd "Postsnapshotcmd" on page 428 | Client options file (dsm.opt) or with include.fs option. |
| preservelastaccessdate "Preservelastaccessdate" on page 431 | Client options file (dsm.opt) or command line. |
| presnapshotcmd "Presnapshotcmd" on page 434 | Client options file (dsm.opt) or with include.fs option. |
| resetarchiveattribute "Resetarchiveattribute" on page 440 | Client options file (dsm.opt). |
| skipntpermissions "Skipntpermissions" on page 461 | Client options file (dsm.opt) or command line. |
| skipntsecuritycrc "Skipntsecuritycrc" on page 462 | Client options file (dsm.opt) or command line. |
| snapdiff " Snapdiff " on page 463 | Command line only. |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotcachesize "Snapshotcachesize" on page 468 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotfsidleretries "Snapshotfsidleretries" on page 469 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotproviderfs "Snapshotproviderfs" on page 472 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotproviderimage "Snapshotproviderimage" on page 473 | Client options file (dsm.opt) or with the include.image option. |
| snapshotroot "Snapshotroot" on page 474 | Command line only. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |

Examples

Task Run an incremental backup of the default client domain specified in your client options file (dsm.opt).

Command: Incremental

Task Run an incremental backup of the c:, d:, and e: drives.

Command: i c: d: e:

Task Run an incremental backup of the \home\ngai directory and its contents on the current drive.

Command: i \home\ngai\

Task Assuming that you initiated a snapshot of the C: drive and mounted the snapshot as \\florence\c\$\snapshots\snapshot.0, run an incremental backup of all files and directories under the local snapshot and manage them on the Tivoli Storage Manager server under the file space name c:.

Command: dsmc inc c: -snapshotroot=\\florence\c\$\snapshots\snapshot.0

Task Run a snapdiff 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.

Task Run a snapdiff 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.

Command: incremental -snapdiff H:

Task Run a snapdiff 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 the operation occurs using the latest snapshot (the active snapshot) for volume H:.

Command: incremental -snapdiff H: -diffsnapshot=LATEST

Related information

“Journal-based backup” on page 568

“Selective” on page 640

“Include options” on page 393

“Incrthreshold” on page 401

“Dfsbackupmntpnt” on page 336

Chapter 12, “Processing options,” on page 269 for adaptive subfile backups.

Open file support

If open file support has been configured, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or “in use”) by other applications.

To control an open file support operation with LVSA, you can specify these additional options in your dsm.opt file or as values of the include.fs option: snapshotcachelocation, snapshotcachesize, snapshotfsidleretries, snapshotfsidlewait, snapshotproviderfs, presnapshotcmd, postsnapshotcmd.

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

snapshotcachelocation, snapshotcachesize, snapshotfsidleretries, snapshotfsidlewait, presnapshotcmd, postsnapshotcmd.

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Use the `snapshotcachelocation` option to relocate the cache if necessary. You can specify a `snapshotcachelocation` for a specific drive using the `include.fs` option.
3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
5. To enable open file support in a cluster environment all systems in the cluster should have the OFS feature configured.

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 which are being monitored by the journal engine service.

Tivoli Storage Manager does not use the journaling facility inherent in Windows NTFS file systems or any other journaled file system.

The *journal engine service* records changes to an object or its attributes in a journal database. During a journal-based backup, the client obtains a list of files that are eligible for backup from the journal database. Performing backups on a regular basis maintains the size of the journal.

Journal-based backup can increase backup performance. With journal-based backup, the client does not scan the local file system or obtain information from the server to determine which files to process. Journal-based backup also reduces network traffic between the client and server.

Tivoli Storage Manager filters the list based on the current include-exclude list. Tivoli Storage Manager processes, expires, and updates the resulting files according to policy constraints, such as serialization. The management-class copy frequency attribute is ignored during journal-based backup.

The journal engine service excludes specific system files (pagefile, registry, etc.) from having changes recorded in the journal. Because changes to these files are not journaled, Tivoli Storage Manager does not back up these files. See the journal service configuration file `tsmjbbd.ini` located in the Tivoli Storage Manager installation directory for specific system files that are excluded.

To support journal-based backup, you must install the journaling engine service. Install this service by using the **dsmcutil** command or the GUI Setup wizard.

If the file specification on the **incremental** command is a file space, Tivoli Storage Manager processes any journal entries for that file space. Tivoli Storage Manager

processes directories and wildcarded file specifications in the same way. Tivoli Storage Manager uses the domain list if you do not specify a file specification.

Note: Journal-based backup might not fall back to the traditional incremental backup if the policy domain of your node is changed on the server, depending on when the policy set within the domain was last updated and the date of the last incremental. In this case, you must force a full traditional incremental backup to rebind the files to the new domain. Use the `nojournal` option with the **incremental** command to specify that you want to perform a traditional full incremental backup, instead of the default journal-based backup.

When a user deletes a file with a long name, the Windows operating system might supply a short, or compressed, name to the journal engine service. After the object is deleted, the compressed name can be reused and the deletion notice might no longer identify a unique object. During a journaled incremental backup, the attempt to expire the file fails because the compressed name is not known to the server. When this occurs, a record is placed in the journal indicating that the current directory is not exactly represented at the server. Use the `incrthreshold` option to specify what action is taken when this occurs.

Under the following conditions, the journal database is considered invalid and the client reverts to the traditional full incremental backup:

- A journaled file space name has changed.
- The client node name has changed.
- The client contacts a different server to do the backup.
- Policy changes have occurred (new policy set activation).
- The journal is corrupt (out of space conditions, disk error).
- The journal service is not running.
- The journal service is stopped or started for any reason, including system reboot.

Journal-based backup differs from the traditional full incremental backup in the following ways:

- Tivoli Storage Manager does not enforce non-default copy frequencies (other than 0).
- Attribute changes to an object require a back up of the entire object.

You can use the `nojournal` option with the **incremental** command to perform a traditional full incremental backup instead of the default journal-based backup.

Multiple journal-based backup sessions are possible.

Backing up NTFS 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 `C:` backs up only the junction (`C:\mount`), and not the data under `C:\mount`.

Related concepts

“Restoring NTFS volume mount points” on page 610

“Restoring data on NTFS mounted volumes” on page 610

“Backing up data on NTFS mounted volumes”

Backing up data on NTFS 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 F: in the domain. In this case, duplicate backups can be avoided by configuring the domain to back up C:\mount or F:, but not both.

To back up the data on the mounted volume, run an incremental backup of the mount point using the incremental command:

```
dsmc incremental c:\mount
```

You can also add C:\mount to the DOMAIN option in order 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, C: and the data on the volume 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 C:.
- 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 volume mount points” on page 610

“Restoring data on NTFS mounted volumes” on page 610

“Backing up NTFS volume mount points” on page 569

Back up Microsoft Dfs root

If you perform an incremental backup of Microsoft Dfs root with `dfsbackupmntpnt=yes` specified, Tivoli Storage Manager backs up only the junction points, *not* the subtree under the junctions.

If you want to traverse the Dfs tree and back up the files and subdirectories of any junction it encounters, specify the `dfsbackupmntpnt=no` option. If you want to backup both the Dfs tree structure and the data contained in the Dfs tree you must run two backups: one with `dfsbackupmntpnt=yes` and one with `dfsbackupmntpnt=no`.

This option has no effect if you are backing up individual junctions. The *exclude.dir* option behavior for Dfs junctions is same as for mounted virtual volumes.

Note: If a Dfs root is added or modified, Tivoli Storage Manager will not back it up. You must specify the Dfs root in the `domain` option in the client options file (`dsm.opt`) regardless of whether `DOMAIN ALL-LOCAL` is specified.

Incremental-by-Date

An incremental-by-date backup backs up new and changed files with a modification date later than the date of the last incremental backup stored at the server, unless the files are excluded from backup by an **exclude** statement.

If an incremental-by-date is performed on only part of a file system, the date of the last full incremental is not updated, and the next incremental-by-date will back up these files again. Use the **query filespace** command to determine the date and time of the last incremental backup of the entire file system.

To perform an incremental-by-date backup, use the `incrbydate` option with the **incremental** command.

Unlike a full incremental, an incremental-by-date does not maintain current server storage of *all* your workstation files 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 Tivoli Storage Manager attempted to transfer during *all* command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

Adaptive subfile backups

If you plan to perform a backup over a network device with limited bandwidth, such as a modem, you can reduce the network traffic by using the `subfilebackup` option.

If you plan to use this option, ensure that you specify the `subfilebackup`, `subfilecachepath` and `subfilecachesize` options during the initial backup of your file. The options specify whether to perform an adaptive subfile backup along with the path and size of the data to be transferred.

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 you enter on the initial command line overrides 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 Tivoli Storage Manager has finished processing, enter `QQ` on the Tivoli Storage Manager console. In many cases, but not all, this interrupts the command.

Related information

Chapter 12, "Processing options," on page 269 for options that you cannot use in interactive mode.

Macro

The **macro** command executes a series of commands that you specify in a macro file.

By including the **macro** command within a macro file, you can nest as many as ten levels of commands.

Comment lines are not supported within the macro file that you specify for the **macro** command.

Supported Clients

This command is valid for all clients.

Syntax

►► `Macro` `—macroname` ◀◀

Parameters

macroname

Specifies the fully qualified name of the file containing the commands.

Examples

The following is an example of how to use the **macro** command.

Task Selectively back up files in the following directories:

- `c:\devel\project\proja`
- `c:\devel\project\projb`
- `c:\devel\project\projc`

Command: `macro backabc.mac`

where `backabc.mac` contains the following statements:

```
selective c:\devel\project\proja\*.*
selective c:\devel\project\projb\*.*
selective c:\devel\project\projc\*.*
```

Monitor Process

The **monitor process** command displays a list of current NAS (if NDMP support is enabled) image backup and restore processes for which the administrative user has authority. Tivoli Storage Manager prompts you for an 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-delimited text file that can be imported into a spreadsheet program. The preview contains information such as whether the file is excluded or included. If the file is excluded, the pattern, or reason, 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 – 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 – 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 you have given access to backup versions or archive copies of specific files.

Tivoli Storage Manager displays a list of authorization rules that you defined with the **set access** command or with *Node Access List* on the graphical user interface (GUI) Utilities menu.

The following information is included.

- Authority you gave a user to restore backup versions or retrieve archive copies.
- The node name of the user to whom you gave authorization.
- The files to which the user has access.

Supported Clients

This command is valid for all clients.

Syntax

▶▶—Query Access—▶▶

Parameters

There are no parameters for this command.

Examples

Task Display a list of users who have access to your files.

Command: query access

Query 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 2008 clients, Active Directory object information can also be displayed from full system-state backups on the server.

Supported Clients

This command is valid for the Windows Server 2003 and Windows Server 2008 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\#".

Tivoli Storage Manager cannot display any object names that contain an asterisk (*) as part of the name.

Table 74. Query Adobjects command: Related options

| Option | Where to use |
|--|--|
| adlocation "Adlocation" on page 298 | Command line only. |
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| detail "Detail" on page 335 | Command line only. |
| pitdate (option is ignored when adlocation is not specified) "Pitdate" on page 425 | Command line only. |
| pittime (option is ignored when adlocation is not specified) "Pittime" on page 426 | Command line only. |

Table 74. Query Aobjects command: Related options (continued)

| Option | Where to use |
|--|--|
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |

Examples

Task Query all local deleted objects.

Command: query aobjects

Task Query all local deleted objects for a user with the name starting with Fred.

Command: query aobjects "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 aobjects "cn=Users,DC=bryan,DC=test,DC=ibm,DC=com" -adloc=server

Task Query all local deleted objects for organizational unit testou.

Command: query aobjects "ou=testou"

Task Query the local deleted object with a GUID of E079130D-3451-4C69-8349-31747E26C75B.

Command: query aobjects {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
- Encryption type
- Client-side data deduplication
- Retention initiation
- Whether or not the file is on hold

Size Archive Date - Time File - Expires on - Description

```
-----
219 B 09/03/2009 09:32:13 \\halley\m$\tsm620c.0901fa\debug\bin\
winnt_unicode\dsm.opt 09/03/2010
Archive Date: 09/03/2009
RetInit:STARTED Obj
```

Held:NO
 Modified: 09/02/2009 19:43:00 Created: 09/01/2009 15:31:23
 Compressed: NO Encryption Type: None
 Client-deduplicated: NO

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. This is the name on the workstation drive from which the file was archived. The following example is valid for specifying a UNC name: {'\\machine\C\$'}.

Use the *filespace-name* if the name has 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* 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 75. Query Archive command: Related options

| Option | Where to use |
|---------------------------------------|--|
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| description "Description" on page 334 | Command line only. |
| detail "Detail" on page 335 | Command line only. |
| dirsonly "Dirsonly" on page 340 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| fromdate "Fromdate" on page 381 | Command line only. |

Table 75. Query Archive command: Related options (continued)

| Option | Where to use |
|---|--|
| fromnode "Fromnode" on page 382 | Command line only. |
| frontime "Frontime" on page 383 | Command line only. |
| numberformat "Numberformat" on page 419 | Client options file (dsm.opt) or command line. |
| querysummary "Querysummary" on page 437 | Command line only. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |
| todate "Todate" on page 497 | Command line only. |
| totime "Totime" on page 498 | 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 containing a file extension of .dev. Use the dateformat and timeformat options.
- Command:** q ar -date=5 -time=4 c:\proj*.dev
- Task** Last week you changed the label of your c: drive to **store** and archived some files. Yesterday you changed the label to **dev** and archived more files. Display a list of all the files you archived in the c:\proj directory when the label was **store**.

Command: q ar {store}\proj*

Task Last week you archived files from a diskette labeled **docs**. Display a list of all the files you archived.

Command: q ar {docs}*

Query ASR

The **query ASR** command displays information about the active backup of the Automated System Recovery (ASR) files on the Tivoli Storage Manager server.

Supported Clients

This command is valid for Windows 2003 clients.

Syntax

►► Query ASR —options— ►►

Parameters

Table 76. Query ASR command: Related options

| Option | Where to use |
|---|---|
| dateformat "Dateformat" on page 329 | Client option file (dsm.opt) or command line. |
| inactive "Inactive" on page 391 | Command line only. |
| numberformat "Numberformat" on page 419 | Client option file (dsm.opt) or command line. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| timeformat "Timeformat" on page 495 | Client option file (dsm.opt) or command line. |

Examples

Task Display information about the active backup of the ASR object on the server.

Command: query asr

Query Backup

The **query backup** command displays a list of backup versions of your files that are stored on the Tivoli Storage Manager, 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 ten characters of the management class name appear.

If you use the **detail** option with the **query backup** command, the client displays the following additional information:

- Last modification date
- Creation date
- Compression
- Encryption type
- Client-side data deduplication

Supported Clients

This command is valid for all clients.

Syntax

```

▶▶ Query Backup [ -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. Do not use wild cards 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 2003, Windows Vista, Windows 2008, and Windows 7 system state. This value is supported only when used from a backup-archive client running on Windows 2003, Windows Vista, Windows 2008, or Windows 7.

{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 file space name 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 310 | Command line only. |
| class "Class" on page 313 | Command line only. |
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| detail "Detail" on page 335 | Command line only. |
| dirsonly "Dirsonly" on page 340 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| fromdate "Fromdate" on page 381 | Command line only. |
| fromowner "Fromnode" on page 382 | Command line only. |
| frontime "Frontime" on page 383 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| nasnodename "Nasnodename" on page 416 | Client options file (dsm.opt) or command line. |
| numberformat "Numberformat" on page 419 | Client options file (dsm.opt) or command line. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| querysummary "Querysummary" on page 437 | Command line only. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |

Table 77. Query Backup command: Related options (continued)

| Option | Where to use |
|-----------------------------|--------------------|
| todate "Todate" on page 497 | Command line only. |
| totime "Totime" on page 498 | 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.

Command: q b {"abc"}\proj*.*

Task Display a list of all active and inactive backup versions that were backed up from the c:\proj directory.

Command: q backup -ina c:\proj*

Task Display a list of all your backups in the c:\dir1 directory. Use the detail option to display the last modification date and the creation date of each file.

Command: q backup -detail c:\dir1*

Task Display a list of all active and inactive backup versions that were backed up from the c:\proj directory. Use the dateformat and timeformat options to reformat the dates and times.

Command: q b -date=5 -time=4 -ina c:\proj*

Task Last week you backed up files from a diskette labeled docs. Display a list of those files.

Command: q b {docs}*

Task Query file system images from the nas2 NAS file server.

Command: query backup -nasnodename=nas2 -class=nas

Task Display a list of all files from your c: drive that are contained in the backup set weekly_accounting_data.32145678.

Command: query backup c:* -subdir=yes
-backupsetname=weekly_accounting_data.32145678

Task Display a list of all the files that are contained in the backup set weekly_accounting.32145678.

Command: query filespace -backupsetname=weekly_accounting.32145678

Task Display information about all the active and inactive backup versions of the system state on the server.

Command: query backup -ina systemstate

Related information

"Restore data from a backup set" on page 180

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. Tivoli Storage Manager 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 313

“Nasnodename” on page 416

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 a 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 the three choices below.

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 Tivoli Storage Manager 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 Tivoli Storage Manager searches for the backup set from the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set from a local file.

tape Specifies that Tivoli Storage Manager 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 334 | Command line only. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | 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 called `monthly_financial_data` from the Tivoli Storage Manager server.

Command: `query backupset -backupsetname=monthly_financial_data.12345678`

Task Query the backup set in the file: `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 180

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 Tivoli Storage Manager 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 Tivoli Storage Manager 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 Tivoli Storage Manager searches for the backup set from the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set from a local file.

tape Specifies that Tivoli Storage Manager 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 334 | Command line only. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |

Table 79. Query Backupset command: Related options (continued)

| Option | Where to use |
|--|--|
| scrollprompt "Scrollprompt" on page 456 | 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 filec:\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 180

Query Filespace

The **query filespace** command displays a list of file spaces for a node 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. You can also specify a single file space name to query.

A *file space* is a logical space on the server that contains files you backed up or archived. Tivoli Storage Manager assigns a separate file space on the server for each node at your workstation from which you back up or archive files.

Tivoli Storage Manager assigns a separate file space on the server for each file system at your workstation from which you back up or archive files. The file space name is the same as the file system name.

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

```

>> Query Filespace [ _filespace_ ] [ _options_ ]

```

Parameters

*filespace*name

Specifies an optional character string which can include wildcards. Use this argument to specify a subset of file spaces. The default is to display all file spaces.

Table 80. Query Filespace command: Related options

| Option | Where to use |
|--|--|
| backupsetname "Backupsetname" on page 310 | Command line only. |
| class "Class" on page 313 | Command line only. |
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| detail "Detail" on page 335 | Command line only. |
| fromnode "Fromnode" on page 382 | Command line only. |
| nasnodename "Nasnodename" on page 416 | Client options file (dsm.opt) or command line. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | 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 filespace -date=5 -time=4
```

Query a file space from the `nas2` NAS file server.

```
query filespace -nasnodename=nas2 -class=nas
```

Display the `\\florence\c$` file space.

```
query filespace \\florence\c$
```

Display all of the file space names on the server with a file space name ending in '\$' belonging to system named `florence`.

```
query filespace \\florence\*$
```

Display file spaces in the backup set named `monthly_accounting.23456789`.

```
query filespace -backupsetname=monthly_accounting.23456789
```

Related information

"Restore data from a backup set" on page 180

“Nasnodename” on page 416

“Class” on page 313

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, Tivoli Storage Manager prompts for an administrator ID.

Place the `nasnodename` option in your client options file (`dsm.opt`). The value in the client options file is the default, but this value can be overridden on the command line. If the `nasnodename` option is not specified in the client options file, it must be specified on the command line when processing NAS file systems.

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.

OK

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 option is valid for all clients.

Syntax

► Query Group — *filespec* — *options* ►

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 382 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| showmembers "Showmembers" on page 460 (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 588

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 - This is the volume size which was backed up.
- Stored Size - This is the actual image size stored on the server. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
- File system type
- Backup date and time
- Management class assigned to image backup
- Whether the image backup is an active or inactive copy
- The image name

Note: The Tivoli Storage Manager API must be installed to use the **query image** command.

Supported Clients

This command is valid for all Windows clients.

Syntax

```

>> Query Image [ -options ] [ -logicalvolumename ] [ -filespacename ] >>

```

Parameters

logicalvolumename

The name of a logical volume you want to query. You must specify the exact name of the image. You cannot use wildcards. The default is all active images (unless restricted by one or more options).

filespacename

Specifies the file system name that you want to query.

Omitting *logicalvolumename* and *filespacename* causes all images to be displayed.

Table 82. Query Image command: Related options

| Option | Where to use |
|--|---|
| backupsetname "Backupsetname" on page 310 | Command line only. |
| dateformat "Dateformat" on page 329 | Client option file (dsm.opt) or command line. |
| fromnode "Fromnode" on page 382 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| numberformat "Numberformat" on page 419 | Client option file (dsm.opt) or command line. |
| pitdate "Pitdate" on page 425 | Command line only. |

Table 82. Query Image command: Related options (continued)

| Option | Where to use |
|---|--|
| pittime "Pittime" on page 426 | Command line only. |
| scrolllines "Scrolllines" on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | 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 180

Query Incllexcl

The **query incllexcl** command displays a list of include-exclude statements in the order in which they are processed during backup and archive operations. The list displays the type of option, the scope of the option (archive, all, etc.), and the name of the source file.

You can test the validity of patterns you want to use in your include-exclude list before you actually insert them in your options file. See the *test pattern* explanation below.

Use the detail option to display the management class that is associated with an include-exclude statement.

Supported Clients

This command is valid for all clients.

Syntax

```

>>> Query INCLexcl [test pattern] [-DETail]

```

Parameters

test pattern

Use for testing the validity of patterns you want to use in your include-exclude list. When you use a test pattern with this command, the following occurs:

- The internal include-exclude list is *not* displayed
- The pattern is processed as if it had come from an include-exclude statement, including all the usual error checking
- The pattern is displayed as it would appear in the include-exclude list

If the test pattern has no errors, the compiled pattern result is the same as the test pattern.

-DETail

Display the management class that is associated with the include-exclude statement.

Examples

Here is an example include-exclude list for client-side data deduplication:

```
| Exclude Dedup *...\file2 dsm.opt
| IEOBJTYPE: FILE
| Exclude Dedup g:\*\* dsm.opt
| IEOBJTYPE: IMAGE
| Exclude Dedup *...\ALL dsm.opt
| IEOBJTYPE: SYSTEMSTATE
| Exclude Dedup *...\ALL dsm.opt
| IEOBJTYPE: ASR
```

Task Display a list of include-exclude statements.

Command: query incl excl

Task Display a list of include-exclude statements. Display the management class that is associated with each statement.

Command: query incl excl -detail

Task Test the validity of this pattern: ..\?x?*.log

Command: query incl excl ..\?x?*.log

Note: Some system files are excluded explicitly by Tivoli Storage Manager. You can use the **query incl excl** command to display a list of these files.

Query Mgmtclass

The **query mgmtclass** command displays information about the management classes available in your active policy set.

Your administrator defines management classes that contain attributes controlling whether a file is eligible for backup or archive services. Management classes also determine how backups and archives are managed on the server.

Your active policy set contains a default management class; it can contain any number of additional management classes. You can assign specific management classes to files using **include** options that are located in the client options file (dsm.opt). If you do not assign a management class to a file, Tivoli Storage Manager uses the default management class.

When you archive files, you can override the assigned management class by using the `archmc` option.

Supported Clients

This command is valid for all clients.

Syntax

```
►►—Query Mgmtclass [—options]►►
```

Parameters

Table 83. Query Mgmtclass command: Related options

| Option | Where to use |
|---------------------------------|--------------------|
| detail “Detail” on page 335 | Command line only. |
| fromnode “Fromnode” on page 382 | 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. Tivoli Storage Manager prompts you for an administrator ID.

The authorized administrative user ID should have at least client owner authority over the client workstation node they are using either from the command line or from the web.

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*.

Supported Clients

This command is valid for all clients.

Syntax

►► Query Node —options— —————►►

Parameters

Table 84. Query Node command: Related options

| Option | Where to use |
|--|--|
| type “Type” on page 500 | Command line only. |
| scrolllines “Scrolllines” on page 455 | Client options file (dsm.opt) or command line. |
| scrollprompt “Scrollprompt” on page 456 | Client options file (dsm.opt) or command line. |

Examples

Task Display all NAS nodes.

Command: query node -type=nas

Related information

“Type” on page 500

Query Options

Use the **query options** command to display all or part of your options and their current settings relevant to the command-line client.

Supported Clients

This command is valid for all clients.

Syntax

►► Query Options —options— —pattern— —————►►

Parameters

pattern

An optional character string which can include wildcards. Use this argument to specify a subset of options. The default is to display all options.

Table 85. Query Options command: Related options

| Option | Where to use |
|--|--|
| scrolllines “Scrolllines” on page 455 | Client options file (dsm.opt) or command line. |

Table 85. Query Options command: Related options (continued)

| Option | Where to use |
|---|--|
| scrollprompt "Scrollprompt" on page 456 | Client options file (dsm.opt) or command line. |

Examples

Task Display all options and their values.

Command: query options

Task Display only options beginning with *comm*.

Command: query options comm*

Task Display the value of the replace option.

Command: query options replace

Query Restore

The **query restore** command displays a list of your restartable restore sessions in the server database. The list contains these fields: owner, replace, subdir, preservepath, source, and destination.

A restartable restore session is created when a wildcard restore command fails because of network outage, client failure, server outage, or a similar problem. When such a failure occurs, the file space is locked on the server and its files cannot be moved off the 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 scheduled for your node. Your administrator can set up schedules to perform automatic backups and archives for you. To plan your work, use this command to determine when the next scheduled events occur.

The enhanced **query schedule** command on a Tivoli Storage Manager Version 5.3 and above client reports new parameters. The **query schedule** command prior to Tivoli Storage Manager Version 5.3 is referred to as classic. Refer to the Administrator's Guide and Administrator's Reference for your operating system for more information on the classic and enhanced commands.

Supported Clients

This command is valid for all clients.

Syntax

▶▶—Query Schedule—▶▶

Parameters

There are no parameters for this command.

Examples

Task Display your scheduled events.

Command: query schedule

Query Session

The **query session** command displays information about your session, including the current node name, when the session was established, server information, and server connection information.

Supported Clients

This command is valid for all clients.

Syntax

▶▶—Query SSession—▶▶

Parameters

There are no parameters for this command.

Examples

Task Display your session information.

Command: query session

A sample **query session** display follows:

```
Server Name.....: HALLEY_SERVER1
Server Type.....: Windows
Archive Retain Protect...: "No"
Server Version.....: Ver. 6, Rel. 2, Lev. 0.0
Last Access Date.....: 09/03/2009 09:08:13
Delete Backup Files.....: "No"
Delete Archive Files.....: "Yes"
Deduplication.....: "Server Only"

Node Name.....: HALLEY
User Name.....:
```

The possible client-side deduplication values are:

- None
 - This is displayed when connected to 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 assist in diagnosing problems, although users who are familiar with the concepts addressed by this information might also find it useful.

If you use the *console* option, no special formatting of the output is performed to accommodate screen height or width. Therefore, the console output can be difficult to read due to length and line-wrapping. If the console output is difficult to read, use the *filename* option with the **query systeminfo** command to allow the output to be written to a file that can subsequently be submitted to IBM support.

Supported Clients

This command is valid for all 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 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 inlexcl** 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.
- 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 325 | Command line only. |
| filename "Filename" on page 379 | 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 379

“Console” on page 325

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 Tivoli Storage Manager Windows client supports standard and detailed format.

Supported Clients

This command is valid for the Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows Server 2003 R2, Windows Vista, and Windows 7 clients only.

Syntax

►► Query SYSTEMState [options] ►►

Parameters

Table 87. Query Systemstate command: Related options

| Option | Where to use |
|--|---|
| backupsetname “Backupsetname” on page 310 | Command line only. |
| dateformat “Dateformat” on page 329 | Client option file (dsm.opt) or command line. |
| inactive “Inactive” on page 391 | Command line only. |
| numberformat “Numberformat” on page 419 | Client option file (dsm.opt) or command line. |
| pitdate “Pitdate” on page 425 | Command line only. |
| pittime “Pittime” on page 426 | Command line only. |

Table 87. Query Systemstate command: Related options (continued)

| Option | Where to use |
|--|---|
| showmembers "Showmembers" on page 460 | Command line only. |
| timeformat "Timeformat" on page 495 | Client option file (dsm.opt) or command line. |
| detail "Detail" on page 335 | 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 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 related to Active Directory in the output.

Query VM

Use the **query VM** command to verify the successful backups of the virtual machines.

Supported Clients

This command is valid only on supported Windows and Linux clients that are installed on a vStorage backup server for a Hyper-V or VMware virtual machine.

Syntax

```
►► Query VM — vmname — options ►►
```

Parameters

vmname

Specifies the virtual machine host name that you want to process.

You can also specify the wildcard character "*" to show all vm backups on the Tivoli Storage Manager server.

For Hyper-V backup, if the `vmbackuptype=hypervfull` option is missing, the `vm list` option values are used. If neither the `vmname` parameter or the `vm list` option is specified, all virtual machines are shown. Wildcards are not allowed for the `vmname` parameter.

Table 88. Query VM command: Related options

| Option | Where to use |
|--|--------------------------------------|
| <code>inactive</code> "Inactive" on page 391 Valid for <code>vmbackuptype=fullvm</code> Valid for <code>vmbackuptype=hypervfull</code> | Command line. |
| <code>pitdate</code> "Pitdate" on page 425 Valid for <code>vmbackuptype=fullvm</code> Valid for <code>vmbackuptype=hypervfull</code> | Command line. |
| <code>pittime</code> "Pittime" on page 426 Valid for <code>vmbackuptype=fullvm</code> Valid for <code>vmbackuptype=hypervfull</code> | Command line. |
| <code>vmbackuptype</code> "Vmbackuptype" on page 508 | Command line or client options file. |
| <code>vmchost</code> "Vmchost" on page 511 | Command line or client options file. |
| <code>vmcpw</code> "Vmcpw" on page 512 | Command line or client options file. |
| <code>vmcuser</code> "Vmcuser" on page 513 | Command line or client options file. |

Examples

The following are samples of using the **query VM** command with full VM and file-level VM.

Full VM

```

tsm> QUERY VM * -INA -vmbackuptype=FULLVM
#  Backup Date  Mgmt Class  Type  A/I  Virtual
   |            |            |      |    |  Machine
   |            |            |      |    |
1  09/26/2008  DEFAULT    VMFULL A    vm1
   | 11:34:14
2  09/26/2008  DEFAULT    VMFULL A    vm2
   | 11:34:37
3  09/26/2008  DEFAULT    VMFULL A    vm3
   | 11:34:49
4  09/26/2008  DEFAULT    VMFULL I    vm2
   | 12:35:08
5  09/26/2008  DEFAULT    VMFULL I    vm3
   | 12:35:18
    
```

File Level

```

tsm> QUERY VM * -vmbackuptype=FILE

Filespace Query for Virtual Machine 'vm1'

Accessing as node: vm1
#   Last Incr Date      Type   File Space Name
--
1   09/25/2008 12:19:23 NTFS   \\vm1\c$   fsID: 2
2   09/25/2008 12:19:25 NTFS   \\vm1\e$   fsID: 3
3   09/25/2008 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   09/25/2008 10:49:10 NTFS   \\vm2\c$   fsID: 2
2   09/25/2008 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   09/25/2008 10:49:22 NTFS   \\vm3\c$   fsID: 2
2   09/25/2008 10:49:21 NTFS   \\vm3\h$   fsID: 1

```

Hyper-V full

```

dsmc query vm -vmbackuptype=hypervfull

#   Backup Date   Mgmt Class   Type   A/I   Virtual Machine
   -----   -
1   07/24/2009     DEFAULT     Full   A     Windows Vista
2   07/25/2009     DEFAULT     Full   A     Windows 2003
3   08/04/2009     DEFAULT     Full   A     Windows 2003

```

Related tasks

“Preparing the environment for VMware backup processing” on page 143

Restart Restore

The **restart restore** command displays a list of your restartable restore sessions in the server database.

You can only restart one restartable restore session at a time. Run the **restart restore** command again to restart additional restores.

The restarted restore uses the same options you used in the failed restore. The restarted restore continues from the point at which the restore previously failed.

To cancel restartable restore sessions, use the **cancel restore** command. Use the **restart restore** command when:

- Restartable restore sessions lock the file space at the server so that files cannot be moved off the sequential volumes of the server.
- You cannot back up files affected by the restartable restore.

Options from the failed session supersede new or changed options for the restarted session.

Supported Clients

This command is valid for all 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. Tivoli Storage Manager uses the `preservepath` option with the `subtree` value as the default for restoring files.

Note:

1. When restoring a directory, its modification date and time is set to the date and time of the restore, not to the date and time the directory had when it was backed up. This is because Tivoli Storage Manager restores the directories first, then adds the files to the directories.
2. An error 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 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 should occur, perform any of the following actions to correct it:

- Restore the file with the short file name to a different location.
- Stop the restore and change the name of the existing file.
- Disable the short file name support on Windows.
- Do not use file names that would conflict with the short file naming convention; for example, do not use ABCDEF~1.DOC.

If you set the `subdir` option to *yes* when restoring a specific path and file, Tivoli Storage Manager recursively restores *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories.

See the Microsoft Knowledge Base article Q121007, *How to Disable the 8.3 Name Creation on NTFS Partitions*, for more information.


If the **restore** command is retried because of a communication failure or session loss, the transfer statistics display the bytes Tivoli Storage Manager attempted to

transfer across *all* command attempts. Therefore, the statistics for bytes transferred might not match file statistics, such as those for file size.

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-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 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 backed up from another node with drive labels that are different from yours.

Note: You must specify a mixed or lowercase NTFS file space name 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, Tivoli Storage Manager restores the files to the original source path.

When entering the *destinationfilespec*, consider the following:

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory. If you are restoring a single file, you can optionally end the specification with a file name if you want to give the restored file a new name.
- If the *sourcefilespec* is wildcarded or *subdir=yes* is specified, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

Note: If the destination path or any part of it does not exist, Tivoli Storage Manager 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 three choices below.

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 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 Tivoli Storage Manager 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 Tivoli Storage Manager searches for the backup set from the server. This is the default.

file Specifies that Tivoli Storage Manager searches for the backup set from a local file.

tape Specifies that Tivoli Storage Manager searches for the backup set from a local tape device.

Table 89. Restore command: Related options

| Option | Where to use |
|-------------------------------------|--|
| asrmode "Asrmode" on page 301 | Command line only. |
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |

Table 89. Restore command: Related options (continued)

| Option | Where to use |
|--|--|
| dironly "Dironly" on page 340 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| fromdate "Fromdate" on page 381 | Command line only. |
| fromnode "Fromnode" on page 382 | Command line only. |
| frontime "Frontime" on page 383 | Command line only. |
| ifnewer "Ifnewer" on page 389 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| latest "Latest" on page 407 | Command line only. |
| numberformat "Numberformat" on page 419 | 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 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| preservepath "Preservepath" on page 432 | Command line only. |
| replace "Replace" on page 439 | Client options file (dsm.opt) or command line. |
| skipntpermissions "Skipntpermissions" on page 461 | Client options file (dsm.opt) or command line. |
| skipntsecuritycrc "Skipntsecuritycrc" on page 462 | Client options file (dsm.opt) or command line. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |

Table 89. Restore command: Related options (continued)

| Option | Where to use |
|-----------------------------|--------------------|
| today "Today" on page 497 | Command line only. |
| totime "Totime" on page 498 | Command line only. |

Examples

Task Restore a single file named `budget.fin`.

Command: `restore c:\devel\projecta\budget.fin`

Task Restore a single file named `budget.fin` which resides in the current directory.

Command: `restore file budget.fin`

Task Restore files from the abc NTFS file space `proj` directory.

Command: `rest {"abc"}\proj*.*`

Task Restore all files with a file extension of `.c` from the `c:\devel\projecta` directory.

Command: `rest c:\devel\projecta*.c`

Task Restore all files with an extension of `.c` from the `\devel\projecta` directory located in the `winnt` file space.

Command: `rest {winnt}\devel\projecta*.c`

Task Restore all files with a file extension of `.c` from the `c:\devel\projecta` directory to the `c:\newdevel\projectn\projecta` directory. If the `projectn` or `projectn\projecta` directory does not exist, it is created.

Command: `restore c:\devel\projecta*.c c:\newdevel\projectn\`

Task Restore files in the `c:\project` directory. Use the `pick` and `inactive` options to select active and inactive backup versions.

Command: `restore c:\project* -pi -ina`

Task Restore all files in the `c:\mydir` directory to their state as of 1:00 PM on August 17, 2002.

Command: `restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\`

Task Restore a file from the renamed file space `\\your-node\h$_OLD` to its original location. Enter both the source and destination as follows:

Command: `res \\your-node\h$_OLD\docs\myresume.doc h:\docs\`

Task Restore all files in the `c:\mydir` directory to their state as of 1:00 PM on August 17, 2002.

Command: `restore -pitd=8/17/2002 -pitt=13:00:00 c:\mydir\`

Task Restore a single file named `budget.fin` contained within the backup set `daily_backup_data.12345678`.

Command: `restore c:\projecta\budget.fin
-backupsetname=daily_backup_data.12345678 -location=server`

Related information

“Restore data from a backup set” on page 180

“Preservepath” on page 432

Restoring NTFS 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, Tivoli Storage Manager 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 volume mount points” on page 569

“Backing up data on NTFS mounted volumes” on page 570

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 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 exist on C:\mount:

- C:\mount\projects\2009plan.doc
- C:\mount\projects\2010plan.doc
- C:\mount\master_list.xls

If the restore command (shown above) 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 volume mount points” on page 610

“Backing up NTFS volume mount points” on page 569

“Backing up data on NTFS mounted volumes” on page 570

Restore Microsoft Dfs junctions

To restore Microsoft Dfs junctions, you must restore Microsoft Dfs root.

If you select the junction point itself, Tivoli Storage Manager restores data under junction, but not the junction itself. If you select a junction point that no longer exists under Dfs root, Tivoli Storage Manager creates a local directory under Dfs root with the same name as the junction before restoring data.

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:

```
se1 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

Tivoli Storage Manager 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 file system, set the Tivoli Storage Manager server communication time out value (`idletimeout`) to the maximum value of 255 to avoid client session timeout.

Tivoli Storage Manager is restricted to restoring sparse files that are less than 4 gigabytes in size.

The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user performing the restore has a disk quota (e.g., belongs to the Backup Operator Group), Tivoli Storage Manager will not restore any data that exceeds the disk quota of the restore user and will display a "Disk Full" message.
- If the user performing the restore does not have a disk quota (e.g., belongs to the Administrator Group), Tivoli Storage Manager will restore all data and transfer ownership of those files which exceed the disk quota of the original owner to the user 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.

On Windows Server 2008 clients, individual Active Directory objects can also be restored from full system-state backups on the Tivoli Storage Manager server.

Supported Clients

This command is valid for Windows Server 2003 and Windows Server 2008 clients.

Syntax

```
►► Restore ADOBJects sourcepathspec options ►►
```

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:

```
\  
#  
+  
=  
<  
>
```

For example, "cn=test#" is entered as "cn=test\#".

Tivoli Storage Manager cannot display any object names that contain an asterisk (*) as part of the name.

Do not use wildcards when you specify a distinguished name.

Note:

Table 90. Restore Aobjects command: Related options

| Option | Where to use |
|--|--|
| adlocation "Adlocation" on page 298 | Command line only. |
| dateformat (the option is ignored when adlocation is not specified) "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |
| pitdate (the option is ignored when adlocation is not specified) "Pitdate" on page 425 | Command line only. |
| pittime (the option is ignored when adlocation is not specified) "Pittime" on page 426 | Command line only. |
| replace "Replace" on page 439 | Client options file (dsm.opt) or command line. |

Table 90. Restore Adobjects command: Related options (continued)

| Option | Where to use |
|---|--|
| timeformat (the option is ignored when adlocation is not specified) "Timeformat" on page 495 | 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 ASR

Use the **restore asr** command to restore the Automated System Recovery (ASR) files to a specified location.

By default, the client restores these files to the adsm.sys\ASR staging directory.

You can create an ASR recovery diskette by specifying the diskette drive as the restore location or by copying the files from the adsm.sys\ASR staging directory to a diskette.

Supported Clients

This command is valid for Windows Server 2003 clients only.

You can restore ASR files for Windows Vista, Windows 7, and Windows Server 2008 with the restore systemstate command.

Syntax

```

>> REStore ASR [—destinationfilespec] [—options]
  
```

Parameters

destinationfilespec

Specifies the path where you want to place the restored files. If you do not specify a destination, the client restores the files to the adsm.sys\ASR directory on the system drive root directory.

Table 91. Restore ASR command: Related options

| Option | Where to use |
|---------------------------------|--------------------|
| inactive "Inactive" on page 391 | Command line only. |
| pick "Pick" on page 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |

Examples

Task Restore the ASR files from the Tivoli Storage Manager server to the adsm.sys\ASR directory on the system drive root directory. Use the pick and inactive options to select from a list of both active and inactive backup versions.

Command: restore asr -pick -inactive

Related information

"Creating an Automated System Recovery diskette (Windows Server 2003)" on page 127


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 backed up from another node with drive labels that are different from yours.

Note: You must specify a mixed or lowercase NTFS file space name 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, Tivoli Storage Manager restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, Tivoli Storage Manager assumes the last name is a file name and reports an error. If you are restoring a single file, you can optionally end the destination file specification with a file name if you want to give the restored file a new name. 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 three choices below.

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 native device driver, not the device driver 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 Tivoli Storage Manager client searches for backup sets on the Tivoli Storage Manager server. If you specify the location parameter, the value must be one of the three choices below.

server Specifies that the backup set is on the Tivoli Storage Manager server. This is the default.

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 92. Restore Backupset command: Related options

| Option | Where to use |
|---|--|
| dironly "Dironly" on page 340 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| ifnewer "Ifnewer" on page 389 | Command line only. |
| preservepath "Preservepath" on page 432 | Command line only. |
| quiet "Quiet" on page 438 | Client options file (dsm.opt) or command line. |
| replace "Replace" on page 439 | Client options file (dsm.opt) or command line. |
| skiptpermissions "Skiptpermissions" on page 461 | Client options file (dsm.opt) or command line. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |

Examples

Task Restore the entire backup set called `monthly_financial_data.87654321` from the server.

Command:

```
dsmc restore backupset
  -backupsetname=monthly_financial_data.87654321
  -loc=server
```

Task Restore the entire backup set from the `\\.\tape0` device.

Command:

```
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.

Command:

```
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.

Command:

```
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.

Command:

```
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.

Command:

```
dsmc restore backupset m:\budget\salary.xls
-backupsetname=c:\budget\weekly_budget_data.ost
-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.

Command:

```
dsmc restore backupset systemstate
-BACKUPSetname=c:\recovery\backupsetfile.name
-loc=file -replace=all
```

Related information

“Restore data from a backup set” on page 180

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 Tivoli Storage Manager 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 information

“Localbackupset” on page 408

“Restore” on page 605

“Restore Image” on page 624

“Restore Backupset” on page 615

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. Tivoli Storage Manager 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. Prior to the introduction of the **backupsetname** parameter, the Tivoli Storage Manager client restored backup sets with a different syntax. The previous syntax is currently 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 legacy syntax of **restore backupset**.

{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 backed up from another node with drive labels that are different from yours.

Note: You must specify a mixed or lowercase NTFS file space name 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, Tivoli Storage Manager restores the files to the original source path. If you are restoring more than one file, you must end the file specification with a directory delimiter (/), otherwise, Tivoli Storage Manager assumes the last name is a file name and reports an error. If you are restoring a single file, you can optionally end the destination file specification with a file name if you want to give the restored file a new name. When the **sourcefilespec** does not exist on the target workstation, you must specify **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 native device driver, not the device driver 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 Tivoli Storage Manager client searches for backup sets on the Tivoli Storage Manager server. If you specify the location parameter, the value must be one of the three choices below.

server Specifies that the backup set is on the Tivoli Storage Manager server. This is the default.

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.

Command: `dsmc restore backupset monthly_financial_data.87654321 -loc=server`

- Task** Restore the entire backup set from the \\.\tape0 device.
Command: 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.
Command: 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.
Command: 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.
Command: 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.
Command: 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.
Command: dsmc restore backupset c:\recovery\backupsetfile.name systemstate -loc=file -replace=all

Related information

“Restore data from a backup set” on page 180

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



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 93. Restore Group command: Related options

| Option | Where to use |
|--|--|
| backupsetname "Backupsetname" on page 310 | Command line only. |
| fromdate "Fromdate" on page 381 | Command line only. |
| fromnode "Fromnode" on page 382 | Command line only. |
| fromtime "Fromtime" on page 383 | Command line only. |
| ifnewer "Ifnewer" on page 389 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| latest "Latest" on page 407 | Command line only. |
| pick "Pick" on page 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| preservepath "Preservepath" on page 432 | Command line only. |
| replace "Replace" on page 439 | Client options file (dsm.opt) or command line. |
| showmembers "Showmembers" on page 460 | Command line only. |
| skiptpermissions "Skiptpermissions" on page 461 | Client options file (dsm.opt) or command line. |
| skiptsecuritycrc "Skiptsecuritycrc" on page 462 | Client options file (dsm.opt) or command line. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |

Table 93. Restore Group command: Related options (continued)

| Option | Where to use |
|-------------------------------------|--|
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |
| todate "Todate" on page 497 | Command line only. |
| totime "Totime" on page 498 | 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 615

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 running the Tivoli Storage Manager client *must* have administrator authority to successfully perform any type of image restore.
2. If you use HSM for Windows or HSM for UNIX, and you restore a file system image backup and plan to run reconciliation, you must restore the files that have been backed up after the image backup. Otherwise, migrated files that have been created after the image backup expire from the HSM archive storage on the Tivoli Storage Manager server.

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

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 file system to a FAT32 volume or vice-versa.
- The destination volume to which you restore must exist and be the same size or larger than the source volume.
- The physical layout of the target volume (striped, mirrored, etc.) can differ.
- The target volume is overwritten with data contained in the image backup.
- It is not necessary to format a target volume prior to restoring an image backup containing a file system.
- Tivoli Storage Manager requires an exclusive lock to destination volume you are restoring. The client 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 backed up by the client:
 - Image Size
 - Stored Size - This is the actual image size stored on the server. Because image backup allows you to back up only used blocks in a file system, the stored image size on the Tivoli Storage Manager server could be smaller than the volume size. For online image backups, the stored image can be larger than the file system based on the size of the cache files.
 - File system type
 - Backup date and time
 - Management class assigned to image backup
 - Whether the image backup is an active or inactive copy
 - The image name
- If for some reason a restored image is corrupted, you should run **chkdsk** to check for and repair any bad sectors or data inconsistencies (unless the restored volume is RAW).

Supported Clients

This command is valid for all Windows clients.

Syntax

```
► REStore Image [ options ] [ sourcefilespec ]
                  [ "sourcefilespec" ]
► [ destinationfilespec ]
```

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 file system to a FAT32 volume or vice-versa.

Table 94. Restore Image command: Related options

| Option | Where to use |
|--|---|
| backupsetname "Backupsetname" on page 310 | Command line only. |
| dateformat "Dateformat" on page 329 | Client option file (dsm.opt) or command line. |
| deletefiles "Deletefiles" on page 334 | Command line only. |
| fromnode "Fromnode" on page 382 | Command line only. |
| imagetofile "Imagetofile" on page 390 | Command line only. |
| inactive "Inactive" on page 391 | Command line only. |
| incremental "Incremental" on page 400 | Command line only. |
| noprompt "Noprompt" on page 419 | Command line only. |
| pick "Pick" on page 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| timeformat "Timeformat" on page 495 | Client option file (dsm.opt) or command line. |
| verifyimage "Verifyimage" on page 504 | 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 located on

a dynamic disk the lost space can be recovered by increasing the size of the volume. This 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 recorded at the server. The changes include deletion of files.

Command: `dsmc restore image h: -incremental -deletefiles`

Task Restore the d: drive to its original location. Use the `verifyimage` option to enable detection of bad sectors on the target volume.

Command: `dsmc restore image d: -verifyimage`

Task If bad sectors present on the target volume, use the `imagetofile` option to restore the d: drive to the e:\diskD.img file to avoid data corruption.

Command: `dsmc restore image d: e:\diskD.img -imagetofile`

Task Restore the e: drive from the backup set `weekly_backup_data.12345678` to its original location.

Command: `restore image e:
-backupsetname=weekly_backup_data.12345678`

Related information

“Verifyimage” on page 504

“Imagetofile” on page 390

Restore NAS

The **restore nas** command restores the image of a file system belonging to a Network Attached Storage (NAS) file server. When using an interactive command-line session with a non-administrative ID, Tivoli Storage Manager prompts for an administrator ID.

The NAS file server performs the outboard data movement. A server process performs the restore.

If you used the `toc` option with the **backup nas** command or the `include.fs.nas` option to save Table of Contents (TOC) information for each file system backup, you can use the **QUERY TOC** server command to determine the contents of a file system backup in conjunction with the **RESTORE NODE** server command to restore individual files or directory trees. You can also use the Web client to examine the entire file system tree and select files and directories to restore. If you do not save TOC information, you can still restore individual files or directory trees using the **RESTORE NODE** server command, provided that you know the fully qualified name of each file or directory and the image in which that object was backed up.

Use the `nasnodename` option to specify the node name for the NAS file server. 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 owned by the NAS node you specify. From this list you can select one or more images to restore. If you select multiple images to restore using the `pick` option, do not use the `monitor` option or you will serialize the restores. To start multiple restore processes simultaneously when restoring multiple images, do not specify `monitor=yes`.

Use the `monitor` option to specify whether you want to monitor a NAS file system image restore and display processing information on your screen.

Use the **monitor process** command to display a list of current restore processes for all NAS nodes for which your administrative user ID has authority. The authorized administrative user ID should have at least client owner authority over both the NAS node and the client workstation node they are using either from command line or from the web.

Use the **cancel process** command to stop NAS restore processing.

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 specifying the *sourcefilespec*.

destinationfilespec

Specifies the name of an existing mounted file system on the NAS device over which you want to restore the image. This parameter is optional. The default is the original location of the file system on the NAS device.

Table 95. Restore NAS command: Related options

| Option | Where to use |
|-------------------------------------|--|
| dateformat "Dateformat" on page 329 | Client option file (<code>dsm.opt</code>) or command line. |
| inactive "Inactive" on page 391 | Command line only. |

Table 95. Restore NAS command: Related options (continued)

| Option | Where to use |
|---|--|
| mode "Mode" on page 412 | Command line only. |
| monitor "Monitor" on page 414 | Command line only. |
| nasnodename "Nasnodename" on page 416 | Client options file (dsm.opt) or command line. |
| numberformat "Numberformat" on page 419 | Client option file (dsm.opt) or command line. |
| pick "Pick" on page 424 | Command line only. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| timeformat "Timeformat" on page 495 | 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 416

"Monitor" on page 414

"Cancel Process" on page 552

Restore Systemstate

Use the **restore systemstate** command to restore the backup of the system state stored on a Tivoli Storage Manager server, or inside a backup set from a Tivoli Storage Manager server, when the `backupsetname` option is specified.


The Tivoli Storage Manager Windows client does not support the current commands for restoring individual startable system state components.

On Windows Vista, Windows 7, and Windows Server 2008 clients running BIOS, **restore systemstate** also restores ASR files when running in WinPE (Windows Preinstallation Environment). The backup-archive client does not back up, and hence cannot restore, ASR data on 64-bit architectures with Extensible Firmware Interface (EFI).

Supported Clients

This command is valid for all supported Windows clients.

Syntax



Parameters

Table 96. Restore Systemstate command: Related options

| Option | Where to use |
|---|---|
| asrmode "Asrmode" on page 301 | Command line only. |
| backupsetname "Backupsetname" on page 310 | Command line only. |
| dateformat "Dateformat" on page 329 | Client option file (dsm.opt) or command line. |
| numberformat "Numberformat" on page 419 | Client option file (dsm.opt) or command line. |
| pitdate "Pitdate" on page 425 | Command line only. |
| pittime "Pittime" on page 426 | Command line only. |
| replace "Replace" on page 439 | Client option file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | 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 operating system is not working (Windows Vista, Windows 7, and Windows Server 2008)” on page 170

“Recovering a computer when the operating system is not working (Windows Server 2003)” on page 173

Restore VM

Use the **restore vm** command to fully restore a virtual machine (restore a vmdk).

Use the **restore vm** command to restore a full VM backup (a vmdk restore) of a VMware or Hyper-V virtual machine.

Issue this command from the Backup-Archive Client on the vStorage backup server for the ESX server that you want to restore. The client sends the command to the IBM Tivoli Storage Manager server, and the server sends the backed up disk to the ESX server using the transport method specified that you specify in the client options file.

Restoring a full VM backup creates a new virtual machine; the configuration 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.

Full VM backups created using VCB must be restored using VCB. 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 have performed a file-level back up of the virtual machine’s files, you can restore the backed-up files to the virtual machine using the command line interface or the Java GUI.


If you have the backup-archive client installed on a separate machine that is configured as a vStorage backup server, you can restore the full VM backups to the ESX server that they came from, or to a different server. To restore to a different VM, use the `-host` option.

The version 6.3 backup-archive client can restore a VMware virtual machine that was backed up by a version 6.2.0 (or earlier) backup-archive client.

Supported Clients

This command is valid only on supported Windows clients that are installed on a vStorage backup server for a Hyper-V or VMware virtual machine.

Syntax



Parameters

Note: Any parameter that contains spaces must be enclosed in quotation (" ") marks.

sourcevmspec

Specifies the name of the virtual machine that was backed up.

VMName

Specifies the new name for the virtual machine after it is restored (if you do not want to use the name specified by sourcevmspec).

When restoring Hyper-V systems, wildcards in the VMName are not supported and will cause errors if they are used.

DATACENTER

Specifies the name of the datacenter to restore the VM to as it is defined in the vSphere vCenter. If the datacenter is contained in a folder, you must specify the -datacenter option when you restore the virtual machine and include the folder structure of the datacenter in the datacenter name.

-datacenter=folder_name/datacenter_name

When you restore a virtual machine 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 datacenter. Without a folder name to help locate the original datacenter, the restore fails.

HOST

Specifies the domain name of the ESX host server to restore to as it is defined in the vSphere vCenter.

DATASTORE

Specifies the data store to restore the virtual machine to. The data store can be on a SAN, NAS, or iSCSI device. You can specify only one data store when restoring a virtual machine. If you do not specify a datastore parameter, the virtual machine's vmdk file is restored to the data store it was on when the backup was created.

destinationfilespec

Specifies the location where VCB Full VM image files are restored. If this option is not specified the vmbackdir option is used. This parameter is for Hyper-V and VMware VCB restores only.

Table 97. Restore VM command: Related options

| Option | Where to use |
|--|---|
| datacenter "Datacenter" on page 328 | Command line or options file. |
| datastore "Datastore" on page 329 | Command line or options file. |
| host "Host" on page 386 | Command line or options file. |
| inactive "Inactive" on page 391 | Command line |
| pick "Pick" on page 424 | Command line |
| pitdate "Pitdate" on page 425 | Command line |
| pittime "Pittime" on page 426 | Command line |
| vmbackdir "Vmbackdir" on page 506 | Command line or client options file. (Not supported for Hyper-V backup) |
| vmbackuptype "Vmbackuptype" on page 508 | Command line or client options file. You cannot specify this option if you specify a vmname |
| vmvstortransport "Vmvstortransport" on page 517 | Command line or client options file. |
| vmlist "Vmlist" on page 514 | Command line or client options file. You cannot specify this option if you specify a vmname |

Examples

Task Restore the most recent backup version of myVM to its original name. Use the VMware management interface to delete the original virtual machine, before you restore it using this syntax.

```
dsmc restore vm myvm
```

Task Restore the most recent backup version of *myvm* to a new machine that is created with the name "Test Machine", with the restore target for datacenter, 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 datacenter named *mydatacenter*. The datacenter 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"

```

Example:

```
dsmc RESTORE VM vm2 G:\restdir
```

Files are restored to:

```
G:\restdir\fullvm\vm2\RESTORE_DATE_2009_06_30[11-22-11]\
```

Example:

```
dsmc RESTORE VM vm2
```

With the vmbckdir option set to the default, the C:\mnt\tsmvmbackup file is restored to:

```
C:\mnt\tsmvmbackup\fullvm\vm2\RESTORE_DATE_2009_06_30[11-22-11]\
```

Example:

```
dsmc restore vm myvm -vmname="Lab1" -host=esx1.us.ibm.com
```

List all the active and inactive objects that are a backup type of hypervfull. The selected objects are restored.

All the Hyper-V backup objects are displayed on the first page. After selecting the objects that you want to restore, the second page lists the backup details for the active and inactive objects.

This example shows how the information is displayed across two pages when using the **RESTORE VM** command with the backup type of hypervfull.

TSM Scrollable PICK Window - Restore VSS Filespace

| # | S | L | Writer Name | Component Name | Component Caption |
|----|---|---|----------------------|----------------------|-------------------|
| 1. | Y | 0 | Microsoft Hyper-V VS | 7C8A3FF8-2788-406A-B | sirius RH 5.4 |
| 2. | Y | 0 | Microsoft Hyper-V VS | AD277AC8-13FF-47FF-A | drnovm1 - 2k8 x64 |
| 3. | Y | 0 | Microsoft Hyper-V VS | DA11B31E-22D1-4E31-8 | Meanvm1 |

<U>=Up <D>=Down <T>=Top =Bottom <R#>=Right <L#>=Left <O>=Ok
<C>=Cancel pick

After selecting #3 and 'O' for Ok, the second page is displayed:

TSM Scrollable PICK Window - Restore VSS Filespace

| # | S | L | Backup Date/Time | Size | Dest/Type | A/I | Component Name |
|----|---|---|---------------------|---------|------------|-----|----------------|
| 1. | Y | 1 | 01/06/2010 00:41:43 | 2.80 GB | TSM - Full | I | Meanvm1 |
| 2. | Y | 1 | 01/06/2010 16:17:30 | 2.84 GB | TSM - Full | A | Meanvm1 |

<U>=Up <D>=Down <T>=Top =Bottom <R#>=Right <L#>=Left <E>=Expand
<O>=Ok <C>=Cancel pick

Related tasks

“Preparing the environment for VMware backup processing” on page 143

Retrieve

The **retrieve** command obtains copies of archived files from the Tivoli Storage Manager server. You can retrieve specific files or entire directories.

Use the **description** option to specify the descriptions assigned to the files you want to retrieve.

Use the **pick** option to display a list of your archives from which you can select an archive to retrieve.

Retrieve the files to the same directory from which they were archived, or to a different directory. Tivoli Storage Manager uses the **preservepath** option with the **subtree** value as the default for restoring files.

Note:

1. When retrieving a directory, its modification date and time is set to the date and time of the retrieve, not to the date and time the directory had when it was archived. This is because Tivoli Storage Manager retrieves the directories first, then adds the files to the directories.
2. An error 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 should occur, perform any of the following actions to correct it:

- Retrieve the file with the short file name you specified to a different location.
- Stop the retrieval, and change the name of the existing file.
- Disable the short file name support on Windows.
- Do not use file names that conflict with the short file naming convention. For example, do not use ABCDEF~1.DOC.

If you are using client Version 3.1.0.5 or later, the workstation name is part of the file name. Therefore, if you archive files on one workstation and you want to retrieve them to another workstation, you must specify a destination. This is true even if you are retrieving to the same physical workstation, but the workstation has a new name. For example, to retrieve the `c:\doc\h2.doc` file to its original directory on the workstation, named *star*, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

To retrieve the file to *star*, which has now been renamed *meteor*, you would enter:

```
dsmc retrieve c:\doc\h2.doc \\meteor\c$\
```

You could also enter:


```
dsmc retrieve c:\doc\h2.doc \\star\c$\
```

You can enter the command in either of the preceding ways, because if the workstation name is not included in the specification, the local workstation is assumed (*meteor*, in this case).

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*, 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 retrieve. This name is the drive label on the workstation drive from which the files were archived.

Use the file space name if the drive label name has changed, or if you are retrieving files archived from another node with drive label names that are different from yours.

Note: You must specify a mixed or lowercase NTFS file space name enclosed in 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 reside. If you do not specify a destination, Tivoli Storage Manager restores the files to the original source path.

When entering the *destinationfilespec* consider the following:

- If the *sourcefilespec* names a single file, the *destinationfilespec* can be a file or a directory.
- If the *sourcefilespec* is wildcarded or you specify the `subdir=yes` option, the *destinationfilespec* must be a directory and end with a directory delimiter (\).

Note: If the destination path or any part of it does not exist, Tivoli Storage Manager creates it.

Table 98. Retrieve command: Related options

| Option | Where to use |
|--|--|
| dateformat "Dateformat" on page 329 | Client options file (dsm.opt) or command line. |

Table 98. Retrieve command: Related options (continued)

| Option | Where to use |
|--|--|
| description "Description" on page 334 | Command line only. |
| dirsonly "Dirsonly" on page 340 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only |
| fromdate "Fromdate" on page 381 | Command line only |
| fromnode "Fromnode" on page 382 | Command line only. |
| fromtime "Fromtime" on page 383 | Command line only |
| ifnewer "Ifnewer" on page 389 | Command line only |
| pick "Pick" on page 424 | Command line only. |
| preservepath "Preservepath" on page 432 | Command line only. |
| replace "Replace" on page 439 | Client options file (dsm.opt) or command line. |
| skiptpermissions "Skiptpermissions" on page 461 | Client options file (dsm.opt) or command line |
| skiptsecuritycrc "Skiptsecuritycrc" on page 462 | Client options file (dsm.opt) or command line |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |
| timeformat "Timeformat" on page 495 | Client options file (dsm.opt) or command line. |
| todate "Todate" on page 497 | Command line only. |
| totime "Totime" on page 498 | Command line only. |

Examples

Task Retrieve a single file named `budget.fin`.

Command: `ret c:\devel\projecta\budget.fin`

Task Retrieve all files with an extension of `.c` from the `c:\devel\projecta` directory.

Command: `ret c:\devel\projecta*.c`

- Task** Retrieve all files with a file extension of .c from the \devel\projecta directory on the winnt file space.
Command: ret {winnt}\devel\projecta*.c
- Task** Retrieve all files in the c:\devel directory.
Command: ret c:\devel*
- Task** Retrieve files from the abc file space proj directory.
Command: ret {abc}\proj*.*
- Task** Retrieve all files with a file extension of .c from the c:\devel\projecta directory to the c:\newdevel\projectn\projecta directory. If the \projectn or the \projectn\projecta directory does not exist, it is created.
Command: ret c:\devel\projecta*.c c:\newdevel\projectn\
- Task** Retrieve files in the c:\project directory. Use the pick option.
Command: ret c:\project* -pick
- Task** Retrieve a file from the renamed file space \\your-node\h\$_OLD to its original location. Enter both the source and destination as follows:
Command: ret \\your-node\h\$_OLD\docs\myresume.doc h:\docs\

Related information

“Client options reference” on page 298

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

Tivoli Storage Manager 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 file system, set the Server communication time out value (idletimeout) to the maximum value of 255 to avoid client session timeout.

The following issues apply if more data is restored than the Microsoft disk quota allows:

- If the user performing the retrieve has a disk quota (e.g., belongs to the Backup Operator Group), Tivoli Storage Manager will not retrieve any data that exceeds the disk quota of the retrieve user and will display a "Disk Full" message.
- If the user performing the retrieve does not have a disk quota (e.g., belongs to the Administrator Group), Tivoli Storage Manager will retrieve all data and transfer ownership of those files which exceed the disk quota of the original owner to the user 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 99. Schedule command: Related options

| Option | Where to use |
|--|--|
| maxcmdretries "Maxcmdretries" on page 410 | Client options file (dsm.opt) or command line. |
| password "Password" on page 422 | client options file (dsm.opt) |
| queryschedperiod "Queryschedperiod" on page 436 | Client options file (dsm.opt) or command line. |
| retryperiod "Retryperiod" on page 444 | Client options file (dsm.opt) or command line. |
| schelogname "Schelogname" on page 450 | Client options file (dsm.opt) or command line. |
| schedmode "Schedmode" on page 453 | Client options file (dsm.opt) or command line. |
| sessioninitiation "Sessioninitiation" on page 457 | Client options file (dsm.opt) or command line. |
| tcpclientport "Tcpclientport" on page 491 | 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 regarding scheduled work are sent to the `dsmsched.log` file or to the file you specify with the `schelogname` option in your client options file (`dsm.opt`). If you do not specify a directory path with the file name in the `schelogname` 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 lose 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. This might result in more than one copy of the same file on the server. If this occurs, you might not have as many different down-level versions of the file on the server as you intended. Your version limit might consist of identical files. To avoid this, use the **incremental** command to back up only new and changed files.

You can selectively back up single files *or* directories. You can also use wildcard characters to back up groups of related files.

If you set the `subdir` option to *yes* when backing up a specific path and file, Tivoli Storage Manager recursively backs up *all* subdirectories under that path, and any instances of the specified file that exist under *any* of those subdirectories.

During a selective backup, a directory path might be backed up, even if the specific file that was targeted for backup is not found. For example 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 Tivoli Storage Manager attempts to transfer during *all* command attempts. Therefore, the statistics for bytes transferred might not match the file statistics, such as those for file size.

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. You can specify as many file specifications as available resources or other operating system limits permit. Separate file specifications with a space. You can also use the `filelist` option to process a list of files. The Tivoli Storage Manager client opens the file you specify with this option and processes the list of files within according to the specific command. When backing up a file system, specify the file system with a trailing slash.

Table 100. Selective command: Related options

| Option | Where to use |
|---|---|
| <code>changingretries</code> "Changingretries" on page 312 | Client options file (<code>dsm.opt</code>) or command line. |
| <code>compressalways</code> "Compressalways" on page 322 | Client options file (<code>dsm.opt</code>) or command line. |

Table 100. Selective command: Related options (continued)

| Option | Where to use |
|--|--|
| compression "Compression" on page 323 | Client options file (dsm.opt) or command line. |
| dironly "Dironly" on page 340 | Command line only. |
| filelist "Filelist" on page 377 | Command line only. |
| filesonly "Filesonly" on page 381 | Command line only. |
| postsnapshotcmd "Postsnapshotcmd" on page 428 | Client options file (dsm.opt) or with the include.fs option. |
| preservelastaccessdate "Preservelastaccessdate" on page 431 | Client options file (dsm.opt) or command line. |
| presnapshotcmd "Presnapshotcmd" on page 434 | Client options file (dsm.opt) or with the include.fs option. |
| skipntpermissions "Skipntpermissions" on page 461 | Client options file (dsm.opt) or command line. |
| skipntsecuritycrc "Skipntsecuritycrc" on page 462 | Client options file (dsm.opt) or command line. |
| snapshotcachelocation "Snapshotcachelocation" on page 467 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotcachesize "Snapshotcachesize" on page 468 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotfsidleretries "Snapshotfsidleretries" on page 469 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotfsidlewait "Snapshotfsidlewait" on page 470 | Client options file (dsm.opt) or with include.fs option. |
| snapshotproviderfs "Snapshotproviderfs" on page 472 | Client options file (dsm.opt) or with the include.fs option. |
| snapshotroot "Snapshotroot" on page 474 | Command line only. |
| subdir "Subdir" on page 482 | Client options file (dsm.opt) or command line. |
| tapeprompt "Tapeprompt" on page 486 | Client options file (dsm.opt) or command line. |

Examples

Task Back up the proj_a.dev file in the c:\devel directory.

Command: `sel c:\devel\proja.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 307

“Include options” on page 393

“Subfilebackup” on page 484

Open file support

If open file support has been configured, Tivoli Storage Manager performs a snapshot backup or archive of files that are locked (or "in use") by other applications.

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

Note:

1. You can use the `include.fs` option to set snapshot options on a per file system basis.
2. Use the `snapshotcachelocation` option to relocate the cache if necessary. You can specify a `snapshotcachelocation` for a specific drive using the `include.fs` option.
3. Open file support is only available for local fixed volumes (mounted to either drive letters or volume mount points) formatted with FAT, FAT32 or NTFS file systems. This support includes SAN-attached volumes that meet these requirements.
4. If the client is unable to create a snapshot, failover to non-OFS backup occurs; the same backup support that would be done if the OFS feature was not configured.
5. To enable open file support in a cluster environment all systems in the cluster should have the OFS feature configured.

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.

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 might 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 filesystem.


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\*\* * * *
```

{filespace}

Specifies the file space name (enclosed in braces) on the server that contains the files to which you are giving access. This name is the drive label name on the workstation drive from which the file was backed up or archived. Use the file space name if the drive label name has changed.

image-fs

The name of the image file system to be shared. This can be specified as an asterisk (*) to allow access to all images owned by the user granting access.

node

Specifies the client node of the user to whom you are giving access. Use wildcards to give access to more than one node with similar node names. Use an asterisk (*) to give access to all nodes.

user

This is an optional parameter that restricts access to the named user at the specified node. This parameter is applicable for UNIX clients only. Windows clients can omit this parameter. If this parameter is specified on a Windows client, it is ignored.

Examples

Task Give the user at node_2 authority to restore all files with an extension of .c from the c:\devel\proja directory.

Command: set access backup c:\devel\proja*.c node_2

Task Give the user at node_3 authority to retrieve all files in the c:\devel directory, but do not permit access to files in subdirectories of c:\devel, such as c:\devel\proj.

Command: set access archive c:\devel* node_3

Task Give all nodes whose names end with bldgb the authority to restore all backup versions from all directories on the d: drive. The d: drive has the file space name of project.

Command: set ac b {project}** *bldgb

Set Event

The **set event** command allows you to specify the circumstances for when archived data is deleted.

You can use the **set event** command to:

- Prevent the deletion of data at the end of its assigned retention period (Deletion hold)
- Allow the expiration to take place, as defined by the archive copy group (Release a deletion hold)

- Start the expiration clock running when a particular event occurs (Notify the server that an event has occurred)

Objects affected can be specified with a standard Tivoli Storage Manager filespec (including wildcards), a list of files whose names are in the file specified using the `filelist` option, or a group of archived files with the description specified with the `description` option.

Note: When only a `<filespec>` is used, *all* archived copies of files or folders matching the filespec are affected. If you only want to affect certain versions of a file, use the `-pick` option and select from the displayed list.


Interaction with down-level servers:

If the `set event` command is issued when the client is connected to a server that does not support event-based policy (previous to Tivoli Storage Manager 5.2.2), the command is rejected with an error message indicating that the current server does not support event-based policy.

Supported Clients

This command is valid for all clients.

Syntax



Parameters

TYPE=

Specifies the event type setting. This parameter must be specified.

hold

Prevents the object from being deleted regardless of expiration policy.

release

Allows normal event-controlled expiration to take place.

activateretention

Signals the server that the controlling event has occurred and starts the expiration clock running.

-pick

Provides a list of objects from which the user can select to apply the event.

The following options can also be used and serve their usual purpose:

- `Dateformat`
- `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 329

“Numberformat” on page 419

“Noprompt” on page 419

“Subdir” on page 482

“Timeformat” on page 495

Set Password

The **set password** command changes the Tivoli Storage Manager password for your workstation.

If you omit the old and new passwords when you enter the **set password** command, you are prompted once for the old password and twice for the new password.


A password is not case-sensitive, and it can be as many as 63 characters. Valid characters are:

- a-z** Any letter, a through z, upper or lower-case
- 0-9** Any number, 0 through 9
- +** Plus
- .** Period
- _** Underscore
- Hyphen
- &** Ampersand

Supported Clients

This command is valid for all clients.

Syntax



Parameters

oldpw

Specifies the current password for your workstation.

newpw

Specifies the new password for your workstation.

TYPE

Specifies if this password is for the workstation or file server. Use **TYPE=TSM** to specify the Tivoli Storage Manager password for your workstation. Use **TYPE=FILER** to specify that this password is for snapshot difference operations on a file server. The default is TSM.

For **TYPE=FILER** you must specify the file server name, the user ID, and the password. For example: `dsmc set password -type=filer myfiler filerid filerpasswd.`

If you are using the **dsmc** command in batch mode, and your password contains any special characters (for example, `pa$password`), you must enclose the password in single quotation marks to avoid truncation (for example, `'pa'`).

Note: When you specify **TYPE=FILER**, the password is stored in the `TSM.PWD` file or the Windows registry without validating that the password is valid. Passwords that are stored with the **TYPE=FILER** parameter can be shared

between Tivoli Storage Manager nodes. For example, a password stored by NODEA can be used by NODEB. Only one set of credentials is stored per file server.

You can also use TYPE=VM to set the password. Here are some examples: `dsmc SET PASSWORD -type=VM Vchost VCuserid Vcpassword`, `dsmc SET PASSWORD -type=VM vchost.us.example.com Administrator secret`.

Use the Preferences editor to set the `vmchost`, `vmcuser`, and `vmcpw` options. You can also use the **set password** command to save the `vmcuser` and `vmcpw` options and then set the `vmchost` option, for example: `set password -type=VM vmchost vmcuser vmcpw`.

When using the **set password** TYPE=VM option, make sure that the value specified for `vmchost` is the same as the value that is used on the **set password** command. When using TYPE=VM, the `vmcuser` and `vmcpw` options are not necessary, for example:

```
dsmc SET PASSWORD -type=VM archon.example.com Administrator secret
VMCHOST archon.example.com
```

Use the TYPE=FastBack parameter to store the Tivoli Storage Manager FastBack credentials that are required for mounting and dismounting the FastBack volumes to back up or archive to the Tivoli Storage Manager server.

Important: You must define the user credentials that are required to mount and unmount FastBack volumes from a repository to the Tivoli Storage Manager client before you enter the Tivoli Storage Manager FastBack subcommand. Use the `reposlocation` option to define the credentials.

Here is a brief description of the various configurations and credentials that you need:

- If the Tivoli Storage Manager client is installed on the Tivoli Storage Manager FastBack DR Hub workstation, the client attempts to connect to a local repository. You must manually define the Tivoli Storage Manager FastBack admin user ID and the corresponding password using this command:

```
dsmc set password -type=fastback fbserver fbadminuserid fbadminpass
```

In the previous command, the `fbserver` is the short host name of the Tivoli Storage Manager FastBack Server workstation.

You must issue the `set password` command once for each FastBack Server branch repository on the FastBack DR Hub machine that the Tivoli Storage Manager Backup-Archive Client is expected to connect to.

Follow these steps for each of the network share repositories where the Tivoli Storage Manager client is connected:

1. Configure the repository for remote network access from FastBack Manager. Refer to the *Tivoli Storage Manager FastBack Installation and User's Guide* for details.

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 Tivoli Storage Manager 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 the user ID on the Tivoli Storage Manager client workstation that has read access to the network share repository.

Domain is the domain name of the user ID.

Networkaccesspassword is the password of the network access user ID.

These credentials are retrieved based on the short host name that you specify with the `fbserver` option.

otherserver

Specifies the file server host name if you specify `TYPE=FILER`. The file server host name must be the fully qualified DNS name and cannot be the IP address in dotted-decimal format or the short local-host name. This parameter only applies if you specify `TYPE=FILER`.

otheruserid

Specifies the user ID to be used for the entity named by the **otherserver** parameter. This parameter only applies if you specify `TYPE=FILER`.

otherpassword

Specifies the password to be used for the entity named by the **otherserver** parameter. This parameter only applies if you specify `TYPE=FILER`.

Examples

The following examples use the **set password** command.

Task Change your password from `osecret` to `nsecret`.

```
Command: set password osecret nsecret
```

Task Set up a user ID and password for the root user on the file server `myFiler.example.com`.

```
Command: 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`.

```
Command: 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.

```
Command: 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 Tivoli Storage Manager Backup-Archive Client is expected to connect.
2. For network share repositories the `dsmc set password -type=fastback` command should be in this format: `dsmc set password -type=fastback myFBServer domainName:userId password`.

3. For a FastBack server or a Tivoli Storage Manager FastBack Hub, the user ID and password specified must have Tivoli Storage Manager 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 Tivoli Storage Manager Backup-Archive Client is expected to connect to.
4. The server name specified, `myFBServer`, in the above example, must match the name to be specified on the `-fbserver` option specified on the `backup fastback` or `archive fastback` command.

Related reference

“Snapdiff” on page 463

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 following list includes the major accessibility features in the Tivoli Storage Manager family of products:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

The Tivoli Storage Manager Information Center, and its related publications, are accessibility-enabled. The accessibility features of the information center are described at http://publib.boulder.ibm.com/infocenter/tsminfo/v6r3/topic/com.ibm.help.ic.doc/iehs36_accessibility.html.

Keyboard navigation

On Windows, the Tivoli Storage Manager product family follows Microsoft conventions for all keyboard navigation and access. Drag and Drop support is managed using the Microsoft Windows Accessibility option known as MouseKeys. For more information about MouseKeys and other Windows accessibility options, please refer to the Windows online help (keyword: MouseKeys).

On other operating systems, these products follow the operating-system conventions for keyboard navigation and access.

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 for more 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
Armonk, NY 10504-1785
U.S.A.*

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
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who 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
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information 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 is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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 <http://www.ibm.com/legal/copytrade.shtml>.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary includes terms and definitions for IBM Tivoli Storage Manager and IBM Tivoli Storage FlashCopy Manager products.

To view glossaries for other IBM products, go to <http://www.ibm.com/software/globalization/terminology/>.

The following cross-references are used in this glossary:

- *See* refers the reader from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- *See also* refers the reader to a related or contrasting term.

A

absolute mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup even if the file has not changed since the last backup. See also *modified mode*.

access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights. For example, an access control list is associated with a file that identifies the users who can access that file and their access rights.

access mode

An attribute of a storage pool or a storage volume that specifies whether the server can write to or read from the storage pool or storage volume. The access mode can be read/write, read-only, or unavailable. Volumes in primary storage pools can also have an access mode of destroyed. Volumes in copy storage pools can also have an access mode of offsite.

acknowledgment

The transmission of acknowledgment characters as a positive response to a data transmission.

ACL See *access control list*.

activate

To validate the contents of a policy set and then make it the active policy set.

active-data pool

A named set of storage pool volumes that contain only active versions of client backup data.

active file system

A file system to which space management has been added. With space management, tasks for an active file system include automatic migration, reconciliation, selective migration, and recall. Contrast with *inactive file system*.

active policy set

The activated policy set that contains the policy rules in use by all client nodes that are assigned to the policy domain. See also *policy domain* and *policy set*.

active version

The most recent backup copy of a file stored. The active version of a file cannot be deleted until a backup process detects that the user has either replaced the file with a newer version or has deleted the file from the file server or workstation. Contrast with *inactive version*.

activity log

A log that records normal activity messages that are generated by the server. These messages include information about server and client operations, such as the start time of sessions or device I/O errors.

adaptive subfile backup

A type of backup that sends only changed portions of a file to the server, instead of sending the entire file. Adaptive subfile backup reduces network traffic and increases the speed of the backup.

administrative client

A program that runs on a file server, workstation, or mainframe that administrators use to control and monitor the Tivoli Storage Manager server. Contrast with *backup-archive client*.

administrative command schedule

A database record that describes the

planned processing of an administrative command during a specific time period. See also *client schedule*.

administrative privilege class

See *privilege class*.

administrative session

A period of time during which an administrator user ID communicates with a server to perform administrative tasks. Contrast with *client node session*.

administrator

A user who is registered to the server as an administrator, and who is authorized to perform tasks and issue commands through the assignment of an administrative privilege class.

Advanced Program-to-Program Communication (APPC)

An implementation of the SNA LU 6.2 protocol that allows interconnected systems to communicate and share the processing of programs.

agent node

A client node that has been granted proxy authority to perform operations on behalf of another client node, which is the target node.

aggregate

An object, stored in one or more storage pools, consisting of a group of logical files that are packaged together. See also *logical file* and *physical file*.

aggregate data transfer rate

A performance statistic that indicates the average number of bytes that were transferred per second while processing a given operation.

APPC See *Advanced Program-to-Program Communication*.

application client

A program that is installed on a system to protect an application. The Tivoli Storage Manager server provides backup services to an application client.

archive

To copy programs, data, or files to other storage media, usually for long-term storage or security. Contrast with *retrieve*.

archive copy

A file or group of files that was archived to server storage.

archive copy group

A policy object containing attributes that control the generation, destination, and expiration of archived files.

archive-retention grace period

The number of days that the storage manager retains an archived file when the server is unable to rebind the file to an appropriate management class. See also *bind*.

association

(1) The defined relationship between a client node and a client schedule. An association identifies the name of a schedule, the name of the policy domain to which the schedule belongs, and the name of a client node that performs scheduled operations.

(2) On a configuration manager, the defined relationship between a profile and an object such as a policy domain. Profile associations define the configuration information that is distributed to a managed server when it subscribes to the profile.

audit

To check for logical inconsistencies between information that the server has and the actual condition of the system. The storage manager can audit information about items such as volumes, libraries, and licenses. For example, when a storage manager audits a volume, the server checks for inconsistencies between information about backed-up or archived files that are stored in the database and the actual data that are associated with each backup version or archive copy in server storage.

authentication

The process of checking a user's password before permitting user access to the Tivoli Storage Manager server. Authentication can be turned on or off by an administrator with system privilege.

authentication rule

A specification that another user can use to either restore or retrieve files from storage.

authority

The right to access objects, resources, or functions. See also *privilege class*.

authorization rule

A specification that permits another user to either restore or retrieve a user's files from storage.

authorized user

A user who has administrative authority for the Tivoli Storage Manager client on a workstation. This user changes passwords, performs open registrations, and deletes file spaces.

AutoFS

See *automounted file system*.

automatic detection

A feature that detects, reports, and updates the serial number of a drive or library in the database when the path from the local server is defined.

automatic migration

The process that is used to automatically move files from a local file system to storage, based on options and settings that are chosen by a root user on a workstation. See also *threshold migration* and *demand migration*.

automatic reconciliation

The process that is used to reconcile file systems at regular intervals. The intervals are set by a user with root user authority. See also *reconciliation*.

automounted file system (AutoFS)

A file system that is managed by an automounter daemon. The automounter daemon monitors a specified directory path, and automatically mounts the file system to access data.

B**backup-archive client**

A program that runs on a workstation or file server and provides a means for users to back up, archive, restore, and retrieve files. Contrast with *administrative client*.

backup copy group

A policy object containing attributes that control the generation, destination, and expiration of backup versions of files. A backup copy group belongs to a management class.

backup-retention grace period

The number of days the storage manager retains a backup version after the server is unable to rebind the file to an appropriate management class.

backup set

A portable, consolidated group of active versions of backup files that are generated for a backup-archive client.

backup set collection

A group of backup sets that are created at the same time and which have the same backup set name, volume names, description, and device classes. The server identifies each backup set in the collection by its node name, backup set name, and file type.

backup version

A file or directory that a client node backed up to server storage. More than one backup version can exist in server storage, but only one backup version is the active version. See also *active version* and *inactive version*.

bind To associate all versions of a file with a management class name. See *rebind*.

bindery

A database that consists of three system files for a NetWare server. The files contain user IDs and user restrictions.

C

cache To place a duplicate copy of a file on random access media when the server migrates a file to another storage pool in the hierarchy.

cache file

A snapshot of a logical volume created by Logical Volume Snapshot Agent. Blocks are saved immediately before they are modified during the image backup and their logical extents are saved in the cache files.

CAD See *client acceptor*.

central scheduler

A function that permits an administrator to schedule client operations and administrative commands. The operations can be scheduled to occur periodically or on a specific date. See *client schedule* and *administrative command schedule*.

client A software program or computer that requests services from a server.

client acceptor

An HTTP service that serves the applet for the web client to web browsers. On Windows systems, the client acceptor is installed and run as a service. On AIX, UNIX, and Linux systems, the client acceptor is run as a daemon, and is also called the *client acceptor daemon (CAD)*.

client acceptor daemon (CAD)

See *client acceptor*.

client domain

The set of drives, file systems, or volumes that the user selects to back up or archive data, using the backup-archive client.

client node

A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

client node session

A session in which a client node communicates with a server to perform backup, restore, archive, retrieve, migrate, or recall requests. Contrast with *administrative session*.

client options file

An editable file that identifies the server and communication method, and provides the configuration for backup, archive, hierarchical storage management, and scheduling.

client option set

A group of options that are defined on the server and used on client nodes in conjunction with client options files.

client-polling scheduling mode

A method of operation in which the client queries the server for work. Contrast with *server-prompted scheduling mode*.

client schedule

A database record that describes the planned processing of a client operation during a specific time period. The client operation can be a backup, archive, restore, or retrieve operation, a client operating system command, or a macro. See also *administrative command schedule*.

client/server

Pertaining to the model of interaction in

distributed data processing in which a program on one computer sends a request to a program on another computer and awaits a response. The requesting program is called a client; the answering program is called a server.

client system-options file

A file, used on AIX, UNIX, or Linux system clients, containing a set of processing options that identify the servers to be contacted for services. This file also specifies communication methods and options for backup, archive, hierarchical storage management, and scheduling. This file is also called the *dsm.sys* file. See also *client user-options file*.

client user-options file

A file that contains the set of processing options that the clients on the system use. The set can include options that determine the server that the client contacts, and options that affect backup operations, archive operations, hierarchical storage management operations, and scheduled operations. This file is also called the *dsm.opt* file. For AIX, UNIX, or Linux systems, see also *client system-options file*.

closed registration

A registration process in which only an administrator can register workstations as client nodes with the server. Contrast with *open registration*.

collocation

The process of keeping all data belonging to a single-client file space, a single client node, or a group of client nodes on a minimal number of sequential-access volumes within a storage pool. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored.

collocation group

A user-defined group of client nodes whose data is stored on a minimal number of volumes through the process of collocation.

commit point

A point in time when data is considered consistent.

Common Programming Interface for Communications (CPI-C)

A call-level interface that provides a consistent application programming interface (API) for applications that use program-to-program communications. CPI-C uses LU 6.2 architecture to create a set of interprogram services that can establish and end a conversation, send and receive data, exchange control information, and notify a partner program of errors.

communication method

The method by which a client and server exchange information. See also *Transmission Control Protocol/Internet Protocol*.

communication protocol

A set of defined interfaces that permit computers to communicate with each other.

compression

A function that removes repetitive characters, spaces, or strings of characters from the data being processed and replaces the repetitive characters with control characters. Compression reduces the amount of storage space that is required for the data.

configuration manager

A server that distributes configuration information, such as policies and schedules, to managed servers according to their profiles. Configuration information can include policy and schedules. See also *managed server* and *profile*.

conversation

A connection between two programs over a session that allows them to communicate with each other while processing a transaction.

copy backup

A full backup in which the transaction log files are not deleted so that backup procedures that use incremental or differential backups are not disrupted.

copy group

A policy object containing attributes that control how backup versions or archive copies are generated, where backup versions or archive copies are initially

located, and when backup versions or archive copies expire. A copy group belongs to a management class. See also *archive copy group*, *backup copy group*, *backup version*, and *management class*.

copy storage pool

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class (for space-managed files). See also *primary storage pool* and *destination*.

CPI-C See *Common Programming Interface for Communications*.

D**daemon**

A program that runs unattended to perform continuous or periodic functions, such as network control.

damaged file

A physical file in which Tivoli Storage Manager has detected read errors.

data access control mode

A mode that controls whether a command can access a migrated file, see a migrated file as zero-length, or receive an input/output error if it attempts to access a migrated file. See also *execution mode*.

database backup series

One full backup of the database, plus up to 32 incremental backups made since that full backup. Each full backup that is run starts a new database backup series. A number identifies each backup series.

database snapshot

A complete backup of the entire database to media that can be taken off-site. When a database snapshot is created, the current database backup series is not interrupted. A database snapshot cannot have incremental database backups associated with it. See also *database backup series*. Contrast with *full backup*.

data deduplication

A method of reducing storage needs by eliminating redundant data. Only one instance of the data is retained on storage

media. Other instances of the same data are replaced with a pointer to the retained instance.

data manager server

A server that collects metadata information for client inventory and manages transactions for the storage agent over the local area network. The data manager server informs the storage agent with applicable library attributes and the target volume identifier.

data mover

A device that moves data on behalf of the server. A network-attached storage (NAS) file server is a data mover.

data storage-management application-programming interface (DSMAPI)

A set of functions and semantics that can monitor events on files, and manage and maintain the data in a file. In an HSM environment, a DSMAPI uses events to notify data management applications about operations on files, stores arbitrary attribute information with a file, supports managed regions in a file, and uses DSMAPI access rights to control access to a file object.

default management class

A management class that is assigned to a policy set. This class is used to govern backed up or archived files when a file is not explicitly associated with a specific management class through the include-exclude list.

deduplication

See *data deduplication*.

demand migration

The process that is used to respond to an out-of-space condition on a file system for which hierarchical storage management (HSM) is active. Files are migrated to server storage until space usage drops to the low threshold that was set for the file system. If the high threshold and low threshold are the same, one file is migrated.

desktop client

The group of backup-archive clients that includes clients on Microsoft Windows, Apple, and Novell NetWare operating systems.

destination

A copy group or management class attribute that specifies the primary storage pool to which a client file will be backed up, archived, or migrated.

device class

A named set of characteristics that are applied to a group of storage devices. Each device class has a unique name and represents a device type of disk, file, optical disk, or tape.

device configuration file

(1) For a server, a file that contains information about defined device classes, and, on some servers, defined libraries and drives. The information is a copy of the device configuration information in the database.

(2) For a storage agent, a file that contains the name and password of the storage agent, and information about the server that is managing the SAN-attached libraries and drives that the storage agent uses.

device driver

A program that provides an interface between a specific device and the application program that uses the device.

disaster recovery manager (DRM)

A function that assists in preparing and using a disaster recovery plan file for the server.

disaster recovery plan

A file that is created by the disaster recovery manager (DRM) that contains information about how to recover computer systems if a disaster occurs and scripts that can be run to perform some recovery tasks. The file includes information about the software and hardware that is used by the server, and the location of recovery media.

domain

A grouping of client nodes with one or more policy sets, which manage data or storage resources for the client nodes. See *policy domain* or *client domain*.

DRM See *disaster recovery manager*.

DSMAPI

See *data storage-management application-programming interface*.

dynamic serialization

A type of copy serialization in which a file or folder is backed up or archived on the first attempt regardless of whether it changes during a backup or archive.

E

EA See *extended attribute*.

EB See *exabyte*.

EFS See *Encrypted File System*.

Encrypted File System (EFS)

A file system that uses file system-level encryption.

enterprise configuration

A method of setting up servers so that the administrator can distribute the configuration of one of the servers to the other servers, using server-to-server communication. See also *configuration manager*, *managed server*, *profile*, and *subscription*.

enterprise logging

The process of sending events from a Tivoli Storage Manager server to a designated event server. The event server routes the events to designated receivers, such as to a user exit. See also *event*.

error log

A data set or file that is used to record error information about a product or system.

estimated capacity

The available space, in megabytes, of a storage pool.

event (1) An administrative command or a client operation that is scheduled to be run using Tivoli Storage Manager scheduling.

(2) A message that an Tivoli Storage Manager server or client issues. Messages can be logged using Tivoli Storage Manager event logging.

event record

A database record that describes actual status and results for events.

event server

A server to which other servers can send events for logging. The event server routes the events to any receivers that are enabled for the sending server's events.

exabyte (EB)

For processor storage, real and virtual storage, and channel volume, 1 152 921 504 606 846 976 bytes. For disk storage capacity and communications volume, 1 000 000 000 000 000 000 bytes.

exclude

The process of identifying files in an include-exclude list. This process prevents the files from being backed up or migrated whenever a user or schedule enters an incremental or selective backup operation. A file can be excluded from backup and space management, backup only, or space management only.

exclude-include list

See *include-exclude list*.

execution mode

A mode that controls the space-management related behavior of commands that run under the **dsmmode** command.

expiration

The process by which files, data sets, or objects are identified for deletion because their expiration date or retention period has passed.

expiring file

A migrated or premigrated file that has been marked for expiration and removal from storage. If a stub file or an original copy of a premigrated file is deleted from a local file system, or if the original copy of a premigrated file is updated, the corresponding migrated or premigrated file is marked for expiration the next time reconciliation is run.

extend

To increase the portion of available space that can be used to store database or recovery log information.

extended attribute (EA)

Names or value pairs that are associated with files or directories. There are three classes of extended attributes: user attributes, system attributes, and trusted attributes.

extent The part of a file that is created during the data-deduplication process. Extents are compared with other file extents to identify duplicates.

external library

A type of library that is provided by Tivoli Storage Manager that permits LAN-free data movement for StorageTek libraries that are managed by Automated Cartridge System Library Software (ACSL). To activate this function, the Tivoli Storage Manager library type must be EXTERNAL.

F**file access time**

On AIX, UNIX, or Linux systems, the time when the file was last accessed.

file age

For migration prioritization purposes, the number of days since a file was last accessed.

file device type

A device type that specifies the use of sequential access files on disk storage as volumes.

file server

A dedicated computer and its peripheral storage devices that are connected to a local area network that stores programs and files that are shared by users on the network.

file space

A logical space in server storage that contains a group of files that have been backed up or archived by a client node, from a single logical partition, file system, or virtual mount point. Client nodes can restore, retrieve, or delete their file spaces from server storage. In server storage, files belonging to a single file space are not necessarily stored together.

file space ID (FSID)

A unique numeric identifier that the server assigns to a file space when it is stored in server storage.

file state

The space management mode of a file that resides in a file system to which space management has been added. A file can be in one of three states: resident, premigrated, or migrated. See also *resident file*, *premigrated file*, and *migrated file*.

file system migrator (FSM)

A kernel extension that intercepts all file system operations and provides any space

management support that is required. If no space management support is required, the operation is passed to the operating system, which performs its normal functions. The file system migrator is mounted over a file system when space management is added to the file system.

file system state

The storage management mode of a file system that resides on a workstation on which the hierarchical storage management (HSM) client is installed. A file system can be in one of these states: native, active, inactive, or global inactive.

frequency

A copy group attribute that specifies the minimum interval, in days, between incremental backups.

FSID See *file space ID*.

FSM See *file system migrator*.

full backup

The process of backing up the entire server database. A full backup begins a new database backup series. See also *database backup series* and *incremental backup*. Contrast with *database snapshot*.

fuzzy backup

A backup version of a file that might not accurately reflect what is currently in the file because the file was backed up at the same time as it was being modified.

fuzzy copy

A backup version or archive copy of a file that might not accurately reflect the original contents of the file because it was backed up or archived the file while the file was being modified. See also *backup version* and *archive copy*.

G**General Parallel File System**

A high-performance shared-disk file system that can provide data access from nodes in a cluster environment.

gigabyte (GB)

In decimal notation, 1 073 741 824 when referring to memory capacity; in all other cases, it is defined as 1 000 000 000.

global inactive state

The state of all file systems to which

space management has been added when space management is globally deactivated for a client node. When space management is globally deactivated, hierarchical storage management (HSM) cannot perform migration, recall, or reconciliation. However, a root user can update space management settings and add space management to additional file systems. Users can access resident and premigrated files.

Globally Unique Identifier (GUID)

An algorithmically determined number that uniquely identifies an entity within a system.

GPFS See *General Parallel File System*.

GPFS node set

A mounted, defined group of GPFS file systems.

group backup

The backup of a group containing a list of files from one or more file space origins.

GUID See *Globally Unique Identifier*.

H

hierarchical storage management (HSM)

A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

hierarchical storage management (HSM) client

A client program that works with the Tivoli Storage Manager server to provide hierarchical storage management (HSM) for a system. See also *hierarchical storage management* and *space manager client*.

HSM See *hierarchical storage management*.

HSM client

See *hierarchical storage management client*.

I

ILM See *information lifecycle management*.

image A file system or raw logical volume that is backed up as a single object.

image backup

A backup of a full file system or raw logical volume as a single object.

inactive file system

A file system for which space management has been deactivated. Contrast with *active file system*.

inactive version

A backup version of a file that is either not the most recent backup version, or that is a backup version of a file that no longer exists on the client system. Inactive backup versions are eligible for expiration processing according to the management class assigned to the file. Contrast with *active version*.

include-exclude file

A file containing statements to determine the files to back up and the associated management classes to use for backup or archive. See also *include-exclude list*.

include-exclude list

A list of options that include or exclude selected files for backup. An exclude option identifies files that should not be backed up. An include option identifies files that are exempt from the exclusion rules or assigns a management class to a file or a group of files for backup or archive services.

incremental backup

(1) A copy of all database data that has changed since the most recent successful full backup operation. An incremental backup is also known as a *cumulative backup image* because each incremental backup includes the contents of the previous incremental backup.

(2) The process of backing up information in the database that is new or changed since the last full backup. Contrast with *full backup*. See also *database backup series*.

(3) For Data Protection for Microsoft Exchange Server, a backup in which the transaction logs are backed up and then cleared.

individual mailbox restore

See *mailbox restore*.

information lifecycle management (ILM)

GPFS policy-based file management for storage pools and file sets.

inode The internal structure that describes the individual files on AIX, UNIX, or Linux systems. An inode contains the node, type, owner, and location of a file.

inode number
A number specifying a particular inode file in the file system.

IP address
A unique address for a device or logical unit on a network that uses the IP standard.

J

job file
A generated file that contains configuration information for a migration job. The file is XML format and can be created and edited in the hierarchical storage management (HSM) client for Windows client graphical user interface.

journal-based backup
A method for backing up Windows clients and AIX clients that exploits the change notification mechanism in a file to improve incremental backup performance by reducing the need to fully scan the file system.

journal daemon
On AIX, UNIX, or Linux systems, a program that tracks change activity for files residing in file systems.

journal service
In Microsoft Windows, a program that tracks change activity for files residing in file systems.

K

kilobyte (KB)
For processor storage, real and virtual storage, and channel volume, 210 or 1 024 bytes. For disk storage capacity and communications volume, 1 000 bytes.

L

LAN See *local area network*.

LAN-free data movement
The movement of client data between a client system and a storage device on a storage area network (SAN), bypassing the local area network. This process is also referred to as *LAN-free data transfer*.

LAN-free data transfer
See *LAN-free data movement*.

leader data
Bytes of data, from the beginning of a migrated file, that are stored in the file's corresponding stub file on the local file system. The amount of leader data that is stored in a stub file depends on the stub size that is specified.

library
(1) A repository for demountable recorded media, such as magnetic disks and magnetic tapes.
(2) A collection of one or more drives, and possibly robotic devices (depending on the library type), which can be used to access storage volumes.

library client
A server that uses server-to-server communication to access a library that is managed by another storage management server. See also *library manager*.

library manager
A server that controls device operations when multiple storage management servers share a storage device. See also *library client*.

local (1) Pertaining to a device, file, or system that is accessed directly from a user system, without the use of a communication line.
(2) For HSM products, pertaining to the destination of migrated files that are being moved.

local area network (LAN)
A network that connects several devices in a limited area (such as a single building or campus) and that can be connected to a larger network.

local shadow volumes
Data that is stored on shadow volumes localized to a disk storage subsystem.

LOFS See *loopback virtual file system*.

logical file
A file that is stored in one or more server storage pools, either by itself or as part of an aggregate. See also *aggregate* and *physical file*.

logical occupancy
The space that is used by logical files in a

storage pool. This space does not include the unused space created when logical files are deleted from aggregate files, so it might be less than the physical occupancy.

logical unit (LU)

An access point through which a user or application program accesses the Systems Network Architecture (SNA) network to communicate with another user or application program.

logical unit number (LUN)

In the Small Computer System Interface (SCSI) standard, a unique identifier that is used to differentiate devices, each of which is a logical unit (LU).

logical volume

A portion of a physical volume that contains a file system.

logical volume backup

A backup of a file system or logical volume as a single object.

Logical Volume Snapshot Agent (LVSA)

Software that can act as the snapshot provider for creating a snapshot of a logical volume during an online image backup.

loopback virtual file system (LOFS)

A file system that is created by mounting a directory over another local directory, also known as mount-over-mount. A LOFS can also be generated using an automounter.

LU See *logical unit*.

LUN See *logical unit number*.

LVSA See *Logical Volume Snapshot Agent*.

M

macro file

A file that contains one or more storage manager administrative commands, which can be run only from an administrative client using the MACRO command. Contrast with *Tivoli Storage Manager command script*.

mailbox restore

A function that restores Microsoft Exchange Server data (from IBM Data Protection for Microsoft Exchange backups) at the mailbox level or mailbox-item level.

managed object

In Tivoli Storage Manager, a definition in the database of a managed server that was distributed to the managed server by a configuration manager. When a managed server subscribes to a profile, all objects that are associated with that profile become managed objects in the database of the managed server. In general, a managed object cannot be modified locally on the managed server. Objects can include policy, schedules, client option sets, server scripts, administrator registrations, server definitions, and server group definitions.

managed server

A Tivoli Storage Manager server that receives configuration information from a configuration manager using a subscription to one or more profiles. Configuration information can include definitions of objects such as policy and schedules. See also *configuration manager*, *subscription*, and *profile*.

management class

A policy object that users can bind to each file to specify how the server manages the file. The management class can contain a backup copy group, an archive copy group, and space management attributes. See also *copy group*, *space manager client*, *bind*, and *rebind*.

maximum transmission unit

The largest possible unit of data that can be sent on a given physical medium in a single frame. For example, the maximum transmission unit for Ethernet is 1500 bytes.

MB See *megabyte*.

media server

In a z/OS environment, a program that provides access to z/OS disk and tape storage for Tivoli Storage Manager servers that run on operating systems other than z/OS.

megabyte (MB)

(1) 1 048 576 bytes (2 to the 20th power) when used in this publication.

(2) For processor storage, real and virtual storage, and channel volume, 2 to the power of 20 or 1 048 576 bits. For disk

storage capacity and communications volume, 1 000 000 bits.

metadata

Data that describes the characteristics of data; descriptive data.

migrate

To move data from one storage location to another. In Tivoli Storage Manager products, migrating can mean moving data from a client node to server storage, or moving data from one storage pool to the next storage pool defined in the server storage hierarchy. In both cases the movement is controlled by policy, such as thresholds that are set. See also *migration threshold*.

migrated file

A file that has been copied from a local file system to Tivoli Storage Manager storage. For HSM clients on UNIX or Linux systems, the file is replaced with a stub file on the local file system. On Windows systems, creation of the stub file is optional. See also *stub file* and *resident file*. For HSM clients on UNIX or Linux systems, contrast with *premigrated file*.

migrate-on-close recall mode

A mode that causes a migrated file to be recalled back to its originating file system temporarily. Contrast with *normal recall mode* and *read-without-recall recall mode*.

migration job

A specification of files to migrate, and actions to perform on the original files after migration. See also *job file*.

migration threshold

High and low capacities for storage pools or file systems, expressed as percentages, at which migration is set to start and stop.

mirroring

The process of writing the same data to multiple locations at the same time. Mirroring data protects against data loss within the recovery log.

mode

A copy group attribute that specifies whether to back up a file that has not been modified since the last time the file was backed up. See *modified mode* and *absolute mode*.

modified mode

In storage management, a backup copy-group mode that specifies that a file is considered for incremental backup only if it has changed since the last backup. A file is considered a changed file if the date, size, owner, or permissions of the file have changed. See also *absolute mode*.

mount limit

The maximum number of volumes that can be simultaneously accessed from the same device class. The mount limit determines the maximum number of mount points. See also *mount point*.

mount point

On the Tivoli Storage Manager server, a logical drive through which volumes in a sequential access device class are accessed. For removable-media device types, such as tape, a mount point is a logical drive that is associated with a physical drive. For the file device type, a mount point is a logical drive that is associated with an I/O stream. The number of mount points for a device class is defined by the value of the mount limit attribute for that device class. See also *mount limit*.

mount retention period

The maximum number of minutes that the server retains a mounted sequential-access media volume that is not being used before it dismounts the sequential-access media volume.

mount wait period

The maximum number of minutes that the server waits for a sequential-access volume mount request to be satisfied before canceling the request.

MTU See *maximum transmission unit*.

N**Nagle algorithm**

An algorithm that reduces congestion of TCP/IP networks by combining smaller packets and sending them together.

named pipe

A type of interprocess communication that permits message data streams to pass between peer processes, such as between a client and a server.

NAS See *network-attached storage*.

NAS node

A client node that is a network-attached storage (NAS) file server. Data for the NAS node is transferred by a NAS file server that is controlled by the network data management protocol (NDMP). A NAS node is also called a NAS file server node.

native file system

A file system that is locally added to the file server and is not added for space management. The hierarchical storage manager (HSM) client does not provide space management services to the file system.

native format

A format of data that is written to a storage pool directly by the Tivoli Storage Manager server. Contrast with *non-native data format*.

NDMP

See *Network Data Management Protocol*.

NetBIOS

See *Network Basic Input/Output System*.

network-attached storage (NAS) file server

A dedicated storage device with an operating system that is optimized for file-serving functions. A NAS file server can have the characteristics of both a node and a data mover.

Network Basic Input/Output System (NetBIOS)

A standard interface to networks and personal computers that is used on local area networks to provide message, print-server, and file-server functions. Application programs that use NetBIOS do not have to handle the details of LAN data link control (DLC) protocols.

Network Data Management Protocol (NDMP)

A protocol that allows a network storage-management application to control the backup and recovery of an NDMP-compliant file server, without installing vendor-acquired software on that file server.

network data-transfer rate

A rate that is calculated by dividing the total number of bytes that are transferred by the data transfer time. For example, this rate can be the time that is spent transferring data over a network.

node A file server or workstation on which the backup-archive client program has been installed, and which has been registered to the server.

node name

A unique name that is used to identify a workstation, file server, or PC to the server.

node privilege class

A privilege class that gives an administrator the authority to remotely access backup-archive clients for a specific client node or for all clients in a policy domain. See also *privilege class*.

non-native data format

A format of data that is written to a storage pool that differs from the format that the server uses for operations.

normal recall mode

A mode that causes a migrated file to be copied back to its originating file system when it is accessed.

O**offline volume backup**

A backup in which the volume is locked so that no other system applications can access it during the backup operation.

online volume backup

A backup in which the volume is available to other system applications during the backup operation.

open registration

A registration process in which users can register their workstations as client nodes with the server. Contrast with *closed registration*.

operator privilege class

A privilege class that gives an administrator the authority to disable or halt the server, enable the server, cancel server processes, and manage removable media. See also *privilege class*.

options file

A file that contains processing options. On Windows and NetWare systems, the file is called *dsm.opt*. On AIX, UNIX, Linux, and Mac OS X systems, the file is called *dsm.sys*.

originating file system

The file system from which a file was

migrated. When a file is recalled using normal or migrate-on-close recall mode, it is always returned to its originating file system.

orphaned stub file

A file for which no migrated file can be found on the Tivoli Storage Manager server that the client node is contacting for space management services. For example, a stub file can be orphaned when the client system-options file is modified to contact a server that is different than the one to which the file was migrated.

out-of-space protection mode

A mode that controls whether the program intercepts out-of-space conditions. See also *execution mode*.

P

pacing

In SNA, a technique by which the receiving system controls the rate of transmission of the sending system to prevent overrun.

packet In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole.

page A defined unit of space on a storage medium or within a database volume.

partial-file recall mode

A recall mode that causes the hierarchical storage management (HSM) function to read just a portion of a migrated file from storage, as requested by the application accessing the file.

password generation

A process that creates and stores a new password in an encrypted password file when the old password expires. Automatic generation of a password prevents password prompting. Password generation can be set in the options file (passwordaccess option). See also *options file*.

path An object that defines a one-to-one relationship between a source and a destination. Using the path, the source accesses the destination. Data can flow from the source to the destination, and back. An example of a source is a data

mover (such as a network-attached storage [NAS] file server), and an example of a destination is a tape drive.

pattern-matching character

See *wildcard character*.

physical file

A file that is stored in one or more storage pools, consisting of either a single logical file, or a group of logical files that are packaged together as an aggregate. See also *aggregate* and *logical file*.

physical occupancy

The amount of space that is used by physical files in a storage pool. This space includes the unused space that is created when logical files are deleted from aggregates. See also *physical file*, *logical file*, and *logical occupancy*.

plug-in

A self-contained software component that modifies (adds, or changes) the function in a particular system. When a plug-in is added to a system, the foundation of the original system remains intact.

policy domain

A grouping of policy users with one or more policy sets, which manage data or storage resources for the users. The users are client nodes that are associated with the policy domain.

policy privilege class

A privilege class that gives an administrator the authority to manage policy objects, register client nodes, and schedule client operations for client nodes. Authority can be restricted to certain policy domains. See also *privilege class*.

policy set

A group of rules in a policy domain. The rules specify how data or storage resources are automatically managed for client nodes in the policy domain. Rules can be contained in management classes. See also *active policy set* and *management class*.

premigrated file

A file that has been copied to Tivoli Storage Manager storage, but has not been replaced with a stub file on the local file system. An identical copy of the file resides both on the local file system and

in Tivoli Storage Manager storage. Premigrated files occur on UNIX and Linux file systems to which space management has been added. Contrast with *migrated file* and *resident file*.

premigrated files database

A database that contains information about each file that has been premigrated to Tivoli Storage Manager storage. The database is stored in a hidden directory named `.SpaceMan` in each file system to which space management has been added.

premigration

The process of copying files that are eligible for migration to Tivoli Storage Manager storage, but leaving the original file intact on the local file system.

premigration percentage

A space management setting that controls whether the next eligible candidates in a file system are premigrated following threshold or demand migration.

primary storage pool

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files migrated from client nodes. See also *destination* and *copy storage pool*.

privilege class

A level of authority that is granted to an administrator. The privilege class determines which administrative tasks the administrator can perform. See also *node privilege class*, *operator privilege class*, *policy privilege class*, *storage privilege class*, and *system privilege class*.

profile

A named group of configuration information that can be distributed from a configuration manager when a managed server subscribes. Configuration information can include registered administrator IDs, policies, client schedules, client option sets, administrative schedules, storage manager command scripts, server definitions, and server group definitions. See also *configuration manager* and *managed server*.

Q

quota (1) For HSM on AIX, UNIX, or Linux systems, the limit (in megabytes) on the

amount of data that can be migrated and premigrated from a file system to server storage.

(2) For HSM on Windows systems, a user-defined limit to the space that is occupied by recalled files.

R

randomization

The process of distributing schedule start times for different clients within a specified percentage of the schedule's startup window.

raw logical volume

A portion of a physical volume that is comprised of unallocated blocks and has no journaled file system (JFS) definition. A logical volume is read/write accessible only through low-level I/O functions.

read-without-recall recall mode

A mode that causes hierarchical storage management (HSM) to read a migrated file from storage without storing it back on the local file system. The last piece of information read from the file is stored in a buffer in memory on the local file system. Contrast with *normal recall mode* and *migrate-on-close recall mode*.

rebind

To associate all backed-up versions of a file with a new management class name. For example, a file that has an active backup version is rebound when a later version of the file is backed up with a different management class association. See also *bind*.

recall In Tivoli Storage Manager, to copy a migrated file from server storage back to its originating file system using the space management client. See also *transparent recall*, *selective recall*, and *recall mode*.

recall mode

A mode that is assigned to a migrated file with the `dsmatr` command that determines how the file is processed when it is recalled. It determines whether the file is stored on the local file system, is migrated back to Tivoli Storage Manager storage when it is closed, or is read from Tivoli Storage Manager storage without storing it on the local file system.

receiver

A server repository that contains a log of server and client messages as events. For example, a receiver can be a file exit, a user exit, or the Tivoli Storage Manager server console and activity log. See also *event*.

reclamation

The process of consolidating the remaining data from many sequential-access volumes onto fewer, new sequential-access volumes.

reclamation threshold

The percentage of space that a sequential-access media volume must have before the server can reclaim the volume. Space becomes reclaimable when files are expired or are deleted.

reconciliation

The process of synchronizing a file system with the Tivoli Storage Manager server, and then removing old and obsolete objects from the Tivoli Storage Manager server.

recovery log

A log of updates that are about to be written to the database. The log can be used to recover from system and media failures. The recovery log consists of the active log (including the log mirror) and archive logs.

register

To define a client node or administrator ID that can access the server.

registry

A repository that contains access and configuration information for users, systems, and software.

remote

(1) Pertaining to a system, program, or device that is accessed through a communication line.

(2) For HSM products, pertaining to the origin of migrated files that are being moved.

resident file

On a Windows system, a complete file on a local file system that might also be a migrated file because a migrated copy can exist in Tivoli Storage Manager storage. On a UNIX or Linux system, a complete

file on a local file system that has not been migrated or premigrated, or that has been recalled from Tivoli Storage Manager storage and modified. Contrast with *stub file* and *premigrated file*. See *migrated file*.

restore

To copy information from its backup location to the active storage location for use. For example, to copy information from server storage to a client workstation.

retention

The amount of time, in days, that inactive backed-up or archived files are kept in the storage pool before they are deleted. Copy group attributes and default retention grace periods for the domain define retention.

retrieve

To copy archived information from the storage pool to the workstation for use. The retrieve operation does not affect the archive version in the storage pool.

roll back

To remove changes that were made to database files since the last commit point.

root user

A system user who operates without restrictions. A root user has the special rights and privileges needed to perform administrative tasks.

S

SAN See *storage area network*.

schedule

A database record that describes client operations or administrative commands to be processed. See *administrative command schedule* and *client schedule*.

scheduling mode

The type of scheduling operation for the server and client node that supports two scheduling modes: client-polling and server-prompted.

scratch volume

A labeled volume that is either blank or contains no valid data, that is not defined, and that is available for use.

script

A series of commands, combined in a file, that carry out a particular function when the file is run. Scripts are interpreted as

they are run. Contrast with *Tivoli Storage Manager command script*.

Secure Sockets Layer (SSL)

A security protocol that provides communication privacy. With SSL, client/server applications can communicate in a way that is designed to prevent eavesdropping, tampering, and message forgery.

selective backup

The process of backing up certain files or directories from a client domain. The files that are backed up are those that are not excluded in the include-exclude list. The files must meet the requirement for serialization in the backup copy group of the management class that is assigned to each file. Contrast with *incremental backup*.

selective migration

The process of copying user-selected files from a local file system to Tivoli Storage Manager storage and replacing the files with stub files on the local file system. Contrast with *threshold migration* and *demand migration*.

selective recall

The process of copying user-selected files from Tivoli Storage Manager storage to a local file system. Contrast with *transparent recall*.

serialization

The process of handling files that are modified during backup or archive processing. See *dynamic serialization*, *static serialization*, *shared static serialization*, and *shared dynamic serialization*.

server A software program or a computer that provides services to other software programs or other computers.

server options file

A file that contains settings that control various server operations. These settings affect such things as communications, devices, and performance.

server-prompted scheduling mode

A client/server communication technique where the server contacts the client node when tasks must be done. Contrast with *client-polling scheduling mode*.

server storage

The primary, copy, and active-data storage

pools that are used by the server to store user files such as backup versions, archive copies, and files migrated from space manager client nodes (space-managed files). See also *active-data pool*, *primary storage pool*, *copy storage pool*, *storage pool volume*, and *volume*.

session

A logical or virtual connection between two stations, software programs, or devices on a network that allows the two elements to communicate and exchange data.

session resource usage

The amount of wait time, processor time, and space that is used or retrieved during a client session.

shared dynamic serialization

A value for serialization that specifies that a file must not be backed up or archived if it is being modified during the operation. Tivoli Storage Manager retries the backup or archive operation a number of times; if the file is being modified during each attempt, Tivoli Storage Manager will back up or archive the file on its last try. See also *serialization*. Contrast with *dynamic serialization*, *shared static serialization*, and *static serialization*.

shared library

A library device that is used by multiple storage manager servers.

shared static serialization

A copy-group serialization value that specifies that a file must not be modified during a backup or archive operation. Tivoli Storage Manager attempts to retry the operation a number of times. If the file is in use during each attempt, the file is not backed up or archived. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *static serialization*.

snapshot

An image backup type that consists of a point-in-time view of a volume.

space-managed file

A file that is migrated from a client node by the space manager client. The space manager client recalls the file to the client node on demand.

space management

The process of keeping sufficient free storage space available on a local file system for new data by migrating files to server storage. Synonymous with *hierarchical storage management*.

space manager client

A program that runs on a UNIX or Linux system to manage free space on the local file system by migrating files to server storage. The program can recall the files either automatically or selectively. Also called *hierarchical storage management (HSM) client*.

space monitor daemon

A daemon that checks space usage on all file systems for which space management is active, and automatically starts threshold migration when space usage on a file system equals or exceeds its high threshold.

sparse file

A file that is created with a length greater than the data it contains, leaving empty spaces for the future addition of data.

special file

On AIX, UNIX, or Linux systems, a file that defines devices for the system, or temporary files that are created by processes. There are three basic types of special files: first-in, first-out (FIFO); block; and character.

SSL See *Secure Sockets Layer*.

stabilized file space

A file space that exists on the server but not on the client.

stanza A group of lines in a file that together have a common function or define a part of the system. Each stanza is identified by a name that occurs in the first line of the stanza. Depending on the type of file, a stanza is ended by the next occurrence of a stanza name in the file, or by an explicit end-of-stanza marker. A stanza can also be ended by the end of the file.

startup window

A time period during which a schedule must be initiated.

static serialization

A copy-group serialization value that specifies that a file must not be modified

during a backup or archive operation. If the file is in use during the first attempt, the storage manager cannot back up or archive the file. See also *serialization*. Contrast with *dynamic serialization*, *shared dynamic serialization*, and *shared static serialization*.

storage agent

A program that enables the backup and restoration of client data directly to and from storage attached to a storage area network (SAN).

storage area network (SAN)

A dedicated storage network that is tailored to a specific environment, combining servers, systems, storage products, networking products, software, and services.

storage hierarchy

(1) A logical order of primary storage pools, as defined by an administrator. The order is typically based on the speed and capacity of the devices that the storage pools use. The storage hierarchy is defined by identifying the next storage pool in a storage pool definition. See also *storage pool*.

(2) An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include: main storage, such as memory and direct-access storage device (DASD) cache; primary storage (DASD containing user-accessible data); migration level 1 (DASD containing data in a space-saving format); and migration level 2 (tape cartridges containing data in a space-saving format).

storage pool

A named set of storage volumes that are the destination that is used to store client data. A storage pool contains backup versions, archive copies, and files that are migrated from space manager client nodes. A primary storage pool is backed up to a copy storage pool. See also *primary storage pool*, *copy storage pool*, and *active-data pool*.

storage pool volume

A volume that has been assigned to a storage pool. See also *volume*, *active-data pool*, *copy storage pool*, and *primary storage pool*.

storage privilege class

A privilege class that gives an administrator the authority to control how storage resources for the server are allocated and used, such as monitoring the database, the recovery log, and server storage. See also *privilege class*.

stub A shortcut on the Windows file system that is generated by the hierarchical storage management (HSM) client for a migrated file that allows transparent user access. A stub is the sparse file representation of a migrated file, with a reparse point attached.

stub file

A file that replaces the original file on a local file system when the file is migrated to storage. A stub file contains the information that is necessary to recall a migrated file from Tivoli Storage Manager storage. It also contains additional information that can be used to eliminate the need to recall a migrated file.

stub file size

The size of a file that replaces the original file on a local file system when the file is migrated to Tivoli Storage Manager storage. The size that is specified for stub files determines how much leader data can be stored in the stub file. The default for stub file size is the block size defined for a file system minus 1 byte.

subscription

In a Tivoli environment, the process of identifying the subscribers that the profiles are distributed to. For Tivoli Storage Manager, a subscription is the process by which a managed server receives configuration information associated with a particular profile on a configuration manager. See also *managed server*, *configuration manager*, and *profile*.

system privilege class

A privilege class that gives an administrator the authority to issue all server commands. See also *privilege class*.

Systems Network Architecture (SNA)

The description of the logical structure, formats, protocols, and operational sequences for transmitting information through and controlling the configuration and operation of networks.

T**tape library**

A set of equipment and facilities that support an installation's tape environment. The tape library can include tape storage racks, mechanisms for automatic tape mounting, a set of tape drives, and a set of related tape volumes mounted on those drives.

tape volume prefix

The high-level-qualifier of the file name or the data set name in the standard tape label.

target node

A client node for which other client nodes (called agent nodes) have been granted proxy authority. The proxy authority allows the agent nodes to perform operations such as backup and restore on behalf of the target node, which owns the data.

TCA See *trusted communications agent*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

threshold migration

The process of moving files from a local file system to Tivoli Storage Manager storage based on the high and low thresholds that are defined for the file system. Contrast with *demand migration*, *selective migration*, and *migration job*.

throughput

In storage management, the total bytes in the workload, excluding overhead, that are backed up or restored, divided by elapsed time.

timeout

A time interval that is allotted for an event to occur or complete before operation is interrupted.

timestamp control mode

A mode that determines whether commands preserve the access time for a file or set it to the current time.

Tivoli Storage Manager command script

A sequence of Tivoli Storage Manager administrative commands that are stored in the database of the Tivoli Storage Manager server. The script can run from any interface to the server. The script can

include substitution for command parameters and conditional logic.

tombstone object

A small subset of attributes of a deleted object. The tombstone object is retained for a specified period, and at the end of the specified period, the tombstone object is permanently deleted.

Transmission Control Protocol/Internet Protocol (TCP/IP)

An industry-standard, nonproprietary set of communication protocols that provides reliable end-to-end connections between applications over interconnected networks of different types.

transparent recall

The process that is used to automatically recall a file to a workstation or file server when the file is accessed. See also *recall mode*. Contrast with *selective recall*.

trusted communications agent (TCA)

A program that handles the sign-on password protocol when clients use password generation.

U

UCS-2 A 2-byte (16-bit) encoding scheme based on ISO/IEC specification 10646-1. UCS-2 defines three levels of implementation: Level 1-No combining of encoded elements allowed; Level 2-Combining of encoded elements is allowed only for Thai, Indic, Hebrew, and Arabic; Level 3-Any combination of encoded elements are allowed.

UNC See *Universal Naming Convention name*.

Unicode

A character encoding standard that supports the interchange, processing, and display of text that is written in the common languages around the world, plus some classical and historical texts. The Unicode standard has a 16-bit character set defined by ISO 10646.

Unicode-enabled file space

Unicode file space names provide support for multilingual workstations without regard for the current locale.

Unicode transformation format 8

Unicode Transformation Format (UTF), 8-bit encoding form, which is designed for ease of use with existing ASCII-based

systems. The CCSID value for data in UTF-8 format is 1208.

Universal Naming Convention (UNC) name

A name that is used to access a drive or directory containing files shared across a network. The UNC name includes the system name and a SharePoint name that represents the shared drive or directory.

Universally Unique Identifier (UUID)

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

UTF-8 See *Unicode transformation format 8*.

UUID See *Universally Unique Identifier*.

V**validate**

To check a policy set for conditions that can cause problems if that policy set becomes the active policy set. For example, the validation process checks whether the policy set contains a default management class.

version

A backup copy of a file stored in server storage. The most recent backup copy of a file is the active version. Earlier copies of the same file are inactive versions. The number of versions retained by the server is determined by the copy group attributes in the management class.

virtual file space

A representation of a directory on a network-attached storage (NAS) file system as a path to that directory.

virtual volume

An archive file on a target server that represents a sequential media volume to a source server.

volume

A discrete unit of storage on disk, tape or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. See also *scratch volume*, and *storage pool volume*.

volume history file

A file that contains information about volumes that have been used by the server for database backups and for export of administrator, node, policy, or

server data. The file also has information about sequential-access storage pool volumes that have been added, reused, or deleted. The information is a copy of volume information that is recorded in the server database.

Volume Shadow Copy Service

A set of Microsoft application-programming interfaces (APIs) that you can use to create shadow copy backups of volumes, exact copies of files, including all open files, and so on.

VSS See *Volume Shadow Copy Service*.

VSS Backup

A backup operation that uses Microsoft Volume Shadow Copy Service (VSS) technology. The backup operation produces an online snapshot (point-in-time consistent copy) of Microsoft Exchange data. This copy can be stored on local shadow volumes or on Tivoli Storage Manager server storage.

VSS Fast Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on local shadow volumes.

VSS Instant Restore

A volume-level hardware-assisted Microsoft Volume Shadow Copy Service (VSS) function where target volumes that contain the snapshot are copied back to the original source volumes.

VSS offloaded backup

A backup operation that uses a Microsoft Volume Shadow Copy Service (VSS) hardware provider (installed on an alternate system) to move IBM Data Protection for Microsoft Exchange data to the Tivoli Storage Manager server. This type of backup operation shifts the backup load from the production system to another system.

VSS Restore

A function that uses a Microsoft Volume Shadow Copy Service (VSS) software provider to restore VSS Backups (IBM Data Protection for Microsoft Exchange database files and log files) that reside on

Tivoli Storage Manager server storage to their original location.

W**wildcard character**

A special character such as an asterisk (*) or a question mark (?) that can be used to represent one or more characters. Any character or set of characters can replace the wildcard character.

workstation

A terminal or personal computer at which a user can run applications and that is usually connected to a mainframe or a network.

worldwide name

A 64-bit, unsigned name identifier that is unique.

workload partition (WPAR)

A partition within a single operating system instance.

Index

Numerics

128-bit AES encryption support 110

A

absolute mode 236
accessibility features 653
active backup versions
 displaying 94, 162, 581
 restoring 162
Active Directory
 query for communication method and
 server with which to connect 501
active directory objects
 modifying CAD and agent
 services 195
 restore restrictions and
 limitations 193
 restoring 191, 193, 195
 restoring from system state
 backup 191
 restoring using GUI and command
 line 192
adaptive subfile backup 484
 excluding files from 364
 including files for 393
 restore permissions 85
 specifying client cache path 485
 specifying client cache size 486
adlocation option 298
administrative client
 allowing secure sessions within a
 private network 488
archive
 assign description to 205, 208, 334
 associating local snapshot with server
 file space 208, 475
 binding management class to 208,
 299
 case sensitive file name conflicts 311
 command 529
 compressing files 323
 configuring client-node proxy
 support 208
 copy group attributes 233
 copy mode 236
 delete files after 334
 delete individual archives from server
 file space 210, 554
 deleting file spaces 198
 directory tree 206
 display the last modification date and
 last access datecreation date 578
 files only 502
 grace period retention 232
 how managed 231
 improving speed using share
 memory 271
 include files for 393
 information, query 578
 archive (*continued*)
 list of files 208, 377
 more than one file specification 208
 number of attempts to archive open
 files 312
 only files; not directories 208
 overriding management class
 during 239
 packages 205
 primary tasks 205
 process Dfs junction as junction or
 directory 336
 process directories only (not
 files) 340
 query user access 576
 retrieving
 using command line 211
 shared data on multiple clients under
 a single node name 208, 299
 specifying whether to include
 subdirectories for 208
 starting a web client session 91
 summary of options 272
 suppress confirmation prompt before
 deleting 419
 using commands 207
archive copy group 232
archive fastback
 command 532
archive maximum file size 111
archmc option 299
asnodename option 299
ASR
 back up user profiles
 Windows XP, 2003 129
 backup system drive
 Windows 2003 128
 Windows Vista, 7, 2008 126
 client installation package
 Windows 2003 175
 dsm.opt
 Windows Vista, 7, 2008 125
 hardware requirements
 Windows 2003 174
 preparation
 Windows 2003 127
 Windows Vista, 7, 2008 125
 query 581
 recovery diskette
 Windows 2003 127
 restore procedure
 user profiles 177
 Windows 2003 176
 Windows Vista, 7, 2008 172
 software requirements
 Windows 2003 174
 WinPE CD
 Windows Vista, 7, 2008 170
ASR (Automated System Recovery) files
 restoring files using GUI 169
asrmode option 301, 518

attributes

 preserving in tombstone objects 194
auditlogging option 301
auditlogname option 304
authentication
 Tivoli Storage Manager client 86
authorization
 authorizing user to restore or retrieve
 your files 195
 options 287
auto-update 4, 306
autodeploy option 306
autofsrename option 307
Automated System Recovery (ASR)
 ASR concepts
 Windows 2003 259
 back up 124, 535
 computer name for which system
 state recovery is performed 324
 concepts
 Windows 2003 263
 diagnostic tips
 Windows 2003 259, 262
 questions and answers 259
 Windows 2003 259
 reference information
 Windows 2003 267
 restore 614
 automating backup services
 displaying scheduled work 219, 221
 options for 225
 process commands after backup 427
 process commands before
 backup 429
 starting client scheduler 26

B

back up
 configuring support for client node
 proxy backups 119
 include-exclude list 110
 Microsoft Dfs trees and files 157
 network-attached storage (NAS) 545
 new or changed files 112
 NTFS file spaces 155
 number of attempts to back up open
 files 312
 removable media using drive
 label 154
 system services, assigning
 management class 122
 Windows Server 2003 system
 state 122
 Windows Vista system state 122
 with limited bandwidth 109
Backing up cluster groups 56
 Windows 2003 50
backing up volume mount points
 NTFS 569
backmc option 309

- backup 101
 - automating using Client Service Configuration Utility 243
 - case sensitive file name conflicts 311
 - copy mode 236
 - displaying processing status 149
 - excluding domains 103
 - excluding system services components 122
 - excluding system services from 122
 - excluding system state object from 122
 - full VM 145
 - grace period retention 232
 - image 129
 - client domain 346, 348, 350
 - with incremental backup 544
 - image, offline and online 129
 - improving speed using share memory 271
 - incremental
 - associating local snapshot with server file space 572
 - client command line 103
 - command line 103
 - incremental-by-date
 - client command line 103
 - command line 103
 - multi-session, send files contiguously to the server 319
 - Net Appliance CIFS share definitions 149
 - one server session per file specification 319
 - overview 97
 - primary tasks 97
 - process Dfs junction as junction or directory 336
 - process directories only (not files) 340
 - query user access 576
 - selective
 - associating local snapshot with server file space 644
 - back up list of files 103
 - client command line 103
 - command line 103
 - shared data on multiple clients under a single node name 299
 - starting a web client session 91
 - subdirectories 103
 - summary of options 272
 - system state, assigning management class 122
 - VMware Consolidated Backup 146
- backup and restore
 - NAS file servers using CIFS 139
- backup asr command 535
- backup comparison: incremental, journal-based, incremental-by-date 116
- backup copy group 98, 232
 - attributes 233
- backup fastback command 536
- backup files
 - assigning management class 237
- backup group command 538
- backup image command 540

- backup maximum file size 111
- backup nas command 545
- Backup NTFS data on mounted volumes 570
- backup operators group
 - required user security rights for backup and restore 83
- backup planning 97
- backup set
 - enabling GUI for local restore 180
 - restore 161, 180
 - restore systemstate from 615, 620
 - restoring in a SAN environment 620
- backup sets
 - restore considerations 182, 619
- backup systemstate command 547
- backup vm command 549
- backup with client node proxy
 - agent node 119
 - overview 119
 - target node 119
- backup-archive client
 - overview 1
- backup-archive scheduler service
 - install 243
- backupset
 - enabling GUI for local restore of 408
- backupsetname option 310
- batch mode 525
 - starting a session 89
- bottom-up processing
 - include-exclude list 80
- bypass processing of NTFS security information 461

C

- cache files
 - specifying size of 468
- cancel process command 552
- cancel restore command 553
- casesensitiveaware option 311
- central scheduling
 - summary of options 285
- Certificate Authorities
 - root certificates 32
 - Certificate Authorities 32
- changingretries option 312
- class option 313
- classic (standard) restore 166
- client 4
 - client TCP/IP address other than the one for first server contact 490
 - client TCP/IP port number other than the one for first server contact 491
 - registering with server 68
 - setting password 69
 - size for the TCP/IP sliding window for client node 494
- client acceptor daemon
 - manage scheduler, Web client, or both 409
- client acceptor service
 - configuring to manage scheduler 24
- client command options
 - overview 526

- client components
 - Windows client 6
- client node proxy archive
 - overview 208
- client node proxy backup
 - scheduling 121
- client options
 - display current settings 596
 - exclude
 - exclude.archive 71
 - exclude.backup 71
 - exclude.compression 71
 - exclude.dir 71
 - exclude.file 71
 - exclude.file.backup 71
 - exclude.image 71
 - exclude.subfile 71
 - exclude.systemobject 71
 - guitreeviewafterbackup 385
 - order of processing (precedence) 289
 - overriding using command line 289
 - overview 526
 - using with commands 289
- client options file
 - creating and modifying 19
 - generating in shared directory 21
 - required options for 19
 - specifying include-exclude options 69
- client options reference 298
- client scheduler
 - displaying scheduled work 219, 221
 - options for 225
 - starting 26, 639
 - starting automatically 93
- Client scheduler
 - run at startup 217
- Client Service Configuration Utility
 - commands for installing client services 247
 - configure CAD to manage existing scheduler service 245
 - create new scheduler and associate CAD to manage scheduler 245
 - options for installing client services 254
 - using to automate backups 243
- client services
 - considerations 83
- client user-options
 - creating multiple files 21
- client user-options file
 - overriding using commands 289
- client-node proxy support 119, 208
- client-server communication
 - client TCP/IP address other than the one for first server contact 490
 - client TCP/IP port number other than the one for first server contact 491
 - establishing 19
 - identify your workstation to the server 417
 - maximum disk I/O buffer size client uses when reading or writing files 341
 - method 320
 - name of a named pipe 415

- client-server communication *(continued)*
 - query Active Directory for
 - communication method and server with which to connect 501
 - reconnection attempts after failure 321
 - reconnection interval after failure 322
 - size for the TCP/IP sliding window for client node 494
 - size of internal TCP/IP
 - communication buffer 488
 - specifying number of kilobytes client buffers before sending transaction to server 499
 - TCP/IP address for dsmcad 489
 - TCP/IP address of Tivoli Storage Manager server 493
 - TCP/IP port address of Tivoli Storage Manager server 492
 - TCP/IP port address on which to establish shared memory connection 459
 - whether to send small transactions to server without buffering them first 491
- clientview option 313
- closed registration
 - permissions 69
 - using 69
- Cluster configuration wizard 56
- cluster drives
 - enabling management of 317
- Cluster groups
 - backing up 56
- cluster resources
 - permissions 86
- clusterdisksonly command 314
- clusternode option 317
- collocatebyfilespec option 319
- command line
 - assigning description to archive 208
 - display current settings for client options 596
 - displaying
 - Euro characters in prompt 90
 - processing status 149
 - enabling 8.3 short names 352
 - ending a session 95
 - entering commands 526
 - general rules when entering options with commands 290
 - NAS file systems backup 137
 - overriding management class during archive 239
 - overview of parameters 527
 - performing image backup 134
 - performing point-in-time restore 199
 - restoring files and directories 163
 - restrictions for NAS file systems 135
 - return codes for operations 229
 - specifying file specification 527
 - specifying options file during session 421
 - starting a session 88
 - using wildcard characters 529
- command parameters
 - overview 527
- command processing, summary of options 287
- command session
 - ending 524
 - starting 524
- command-line prompt displaying
 - Euro characters 90
- commands
 - archive 529
 - archive fastback 532
 - backup asr 535
 - backup fastback 536
 - backup group 538
 - backup image 540
 - backup nas 545
 - backup systemstate 547
 - backup vm 549
 - batch mode 525
 - cancel process 552
 - cluster restore 553
 - clusterdisksonly 314
 - delete access 553
 - delete archive 554
 - delete backup 556
 - delete filespace 558
 - delete group 559
 - entering 526
 - entering on command line 526
 - expire 561
 - general rules when entering options with 290
 - help 562
 - incremental 564
 - interactive (loop) mode 525
 - loop 572
 - macro 573
 - maximum file specifications permitted 527
 - monitor process 574
 - overview of parameters 527
 - preview archive 574
 - preview backup 575
 - query access 576
 - query adobjects 577
 - query archive 578
 - query ASR 581
 - query backup 581
 - query backupset 585, 587
 - query filespace 588
 - query group 590
 - query image 591
 - query inlxcld 593
 - query mgmtclass 594
 - query node 595
 - query options 596
 - query restore 597
 - query schedule 598
 - query session 598
 - query systeminfo 599
 - query systemstate 601
 - query VM 602
 - restart restore 604
 - restore 605
 - restore adobjects 612
- commands *(continued)*
 - restore asr 614
 - restore backupset 615, 620
 - restore backupset considerations 182, 619
 - restore group 622
 - restore image 624
 - restore NAS 627
 - restore systemstate 629
 - restore vm 631
 - retrieve 635
 - schedule 639
 - scheduled, enabling or disabling 226
 - selective backup 640
 - set access 644
 - set event 646
 - set password 649
 - specifying file specification 527
 - using 521
 - using in executables 229
 - using options with 289
 - using wildcard characters 529
- commmethod option 320
- commrestartduration option 321
- commrestartinterval option 322
- communication methods
 - installable software 7
 - Shared Memory
 - Windows client 7
 - summary 270
 - TCP/IP
 - Windowss client 7
- communications
 - establishing through firewall 27
 - establishing with Secure Sockets Layer (SSL) 29
- compressalways option 322
- compression
 - disabling processing 398
 - enabling processing 398
 - include-exclude statements 398
- compression option 323
- compression processing
 - exclude from backup 398
 - exclude options 398
 - include files for 393
- compression, encryption, and subfile backup processing
 - back up 398
 - exclude from backup 398
 - exclude options 398
- computername option 324
- configure the client for data deduplication 44
- configuring
 - journal engine service 33
 - open file support 66
 - optional tasks 19
 - required tasks 19
 - the client scheduler 24
 - the Web client 22
- Configuring cluster protection 56
- console option 325, 326
- console window displaying
 - Euro characters 90
- containing quotation marks 91

- copy destination attribute 236
- copy frequency attribute 234
- copy group name attribute 233
- copy groups 232
 - archive 232
 - backup 232
- copy mode attribute
 - absolute 236
 - modified 236
- copy serialization attribute 235
- copy type attribute 234
- Createnewbase 326
- customer support
 - contact xix

D

- data deduplication 40
- data deduplication client
 - configuration 44
- data deduplication files
 - exclude 46
- datacenter option 328
- datastore option 329
- date format
 - specifying 329
- dateformat option 329
- dedupcachepath option 331
- dedupcachesize option 332
- deduplication option 332
- default domain
 - excluding domains from backup 103, 343
- default management class 231
- default policy domain 231
- definitions 659
- delete
 - file space 198
 - NAS or client objects 313
- delete access command 553
- delete archive command 554
- delete backup command 556
- delete group command 559
- delete individual backups from server file
 - space 105
- deleted file systems 154
- deletefiles option 334
- deleting
 - authorizations 195
 - individual archives from server file
 - space 210, 554
 - individual backups from server file
 - space 556
- deleting file spaces 198
- description option 334
- detail option 335
- dfsbackupmntpnt option 336
- diagnostics
 - options 289
- diffsnapshot option 337
- directories
 - assigning management class for 339
 - excluding 71
 - excluding from backup
 - processing 364
 - incremental backup processing
 - overview 112

- directories (*continued*)
 - processing during
 - incremental-by-date 116
 - restoring from command line 163
 - restoring from GUI 163
 - specifying on command line 527
- directory
 - archiving 206
- dirmc option 339
- dironly option 340
- disability 653
- disablenqr option 340
- discretionary access control list
 - (permissions)
 - back up 155
- disk space requirements
 - client 6
 - Windows client 6
- diskbuffsize option 341
- diskcachelocation option 342
- displaying
 - archive information 578
 - online help 95
 - restartable restore sessions 597
 - scheduled events 598
 - session information 598
- domain
 - back up using the GUI 102
 - include for image backup 346, 348, 350
 - include for incremental backup 343
 - include for NAS image backup 346
 - specifying drives in the default 102
- domain list
 - using universal naming convention to
 - specify 155
- domain option 343
- domain.image option 346, 348, 350
- domain.nas option 346
- drive label
 - using to backup removable
 - media 154
- DSM_CONFIG environment variable 22
- DSM_DIR environment variable 22
- DSM_LOG environment variable 22
- dsm.opt file
 - creating and modifying 19
 - required options for 19
 - specifying a drive specification using
 - wildcards 76
- dsm.smp file
 - copying to dsm.opt 19
 - location 19
- dsmc command
 - using options 90
- dsmcutil utility
 - commands for installing client
 - services 247
 - options for installing client
 - services 254
 - overview 246
- dsmerror.pru file 359
- dsmerror.log file 359
- dsmsched.log 449, 452
- dsmwebcl.log 449, 452
- duplicate file names
 - avoiding 161

- dynamic and shared serialization 235

E

- education
 - see Tivoli technical training xvi
- enable8dot3namesupport option 352
- enablearchiveretentionprotection
 - option 353
- enablededupcache option 354
- enablelanfree option 355
- encrypting data during archive 110
- encrypting data during backup 110
- encryption
 - multiple clients under a single node
 - name 299
 - of file data 110
 - saving encryption key password 357
- encryption processing
 - determine encryption cipher used in
 - current session 110
 - excluding files from 364
 - include files for 393
 - query systeminfo command 599
- encryptiontype option 357
- encryptkey option
 - encryptkey=generate 357
 - encryptkey=prompt
 - encryptkey=save 357
- enhanced query schedule 219
- enhanced query schedule command 598
- environment prerequisites 6
- environment variables 22
- error log
 - pruning 361
 - specifying path and file name 361
- error processing, summary of
 - options 287
- errorlogmax option 359
- errorlogname option 361
- errorlogretention option 359, 361
- event logging
 - scheduler 221
- event-based policy retention protection
 - archive 241
 - backup 241
- exclude data deduplication files 46
- exclude options 364
 - exclude.archive 71
 - exclude.backup 71
 - exclude.compression 71
 - exclude.dir 71
 - exclude.file 71
 - exclude.file.backup 71
 - exclude.image 71
 - exclude.subfile 71
 - exclude.systemobject 71
 - preview 79
 - processing 80
 - wildcard characters 76, 78
- exclude.image option 71
- excluding files
 - remotely accessed 75
 - system files 75
 - using wildcard characters 78
 - wildcard characters 76
- excluding system objects 122

executable file
 return codes from 229
expire command 561

F

fbbranch option 369
fbclient option 370
fbpolicyname option 372
fbreposlocation option 373
fbserver option 374
fbvolumename option 375
features
 Windows client 7
file
 restoring active and inactive
 versions 611
file names
 avoiding duplicate 161
file space
 delete 198, 558
 determining fsID 335
 excluding 71
 NAS or client objects 313
 performing an image backup 540
file specification
 maximum allowed on
 commands 527
file systems
 deleted 154
 excluding from backup
 processing 364
 GPFS, multinode cluster
 environment 343
 image backup of 129
filelist option 377
filename option 379
files
 archive
 directory tree 206
 archive a list of 208, 377
 archived, overriding management
 class 239
 archiving 529
 archiving more than one file
 specification 208
 assigning management classes 153
 backing up Microsoft Dfs 157
 backing up open 152
 binding management classes to 239
 compressing during archive or
 backup 323
 definition of changed 112
 delete after archive 334
 delete individual archives from server
 file space 210, 554
 delete individual backups from server
 file space 556
 encryption 110
 excluding groups 76, 78
 include-exclude
 creating in Unicode format 393
 including groups 76, 78
 managing growth during
 compression 322
 maximum file size for operations 111
 processing include-exclude 80

files (*continued*)
 query archive information 578
 query backup information 581
 query user access 576
 renaming file spaces that are not
 Unicode to Unicode-enabled 307,
 564, 640
 reset Windows archive attribute after
 backup 441
 restoring files belonging to another
 node 197
 restoring from command line 163
 restoring from GUI 163
 restoring to another workstation 197
 retrieving
 archives using command line 211
 files belonging to another
 node 197
 to another workstation 197
 sorting list of 94
filesonly option 381
firewall
 establishing communications
 through 27, 386, 492
 specifying TCP/IP ports for the Web
 client 519
 using Web client through 519
 whether server or client initiates
 sessions through 457
fixed drives
 backing up 155
fixes, obtaining xviii
format and language
 summary of options 286
fromdate option 381
fromnode option 382
fromtime option 383
frsprimaryrestore option 384
full incremental
 comparing with journal-based,
 incremental-by-date 116
 description 112
 when to use 116
full VM backup 145, 189
 restore using command line 187, 190
fuzzy backup 235

G

getting started
 changing your password 83
 client scheduler 83
 command-line session 83
 displaying online help 83
 ending a session 83
 GUI session 83
 sorting file lists 83
 Web client session 83
glossary 659
GPFS file system
 multinode cluster environment 343
graphical user interface
 changing password 93
 delete individual files or images from
 server file space 556
 displaying active and inactive backup
 versions 94, 162

graphical user interface (*continued*)
 displaying online help 95
 displaying processing status 149
 enabling for local backupset
 restore 408
 enabling local backup set 180
 ending a session 95
 performing image backup 133
 restore files and directories 163
 starting a session 87
 to back up objects 101
group backup
 display active and inactive
 objects 391
 display all members of 460
 overview 118
 specify name of group 385
 specify virtual file space name
 for 504
 specifying full or differential 412
groupname option 385
GUI
 ending a session 95
 overriding management class during
 archive 239
 performing point-in-time restore 199
 starting a session 87
guitreeviewafterbackup option 385

H

hardware requirements
 ASR
 Windows 2003 174
 Windows client 6
help
 displaying online 95
 Internet resources 95
 online forum 95
 service and technical support 95
help command 562
host option 386
httpport option 386
Hyper-V
 backup limitations 147

I

IBM Publications Center xiii
IBM Support Assistant xviii
ieobjtype option 388
ifnewer option 389
image
 restoring 178
 using chkdsk to repair 178
 using chkdsk tool to repair 624
 using fsck to repair 178, 624
image backup
 configuring online image backup 65
 considerations 130
 deleting 556
 excluding files from 364
 file systems or logical volumes 540
 include files for; assign management
 class to 393
 include.dedup 393

- image backup (*continued*)
 - incremental-by-date image
 - backup 133
 - offline and online 129
 - perform 129
 - point-in-time restore 544
 - revoke access 553
 - specifying selective or incremental 412
 - using command line 134
 - using the GUI 133
 - using with file system incremental 133
 - using with incremental-by-date 132
 - with incremental backup 131, 544
- image backup, considerations 130
- imagegapsize option 389
- imagetofile option 391
- inactive backup versions
 - displaying 94, 162, 581
 - restoring 162
- inactive option 391
- includexl option 392
- include option
 - management class 237
 - processing 80
 - wildcard characters 76, 78
- include-exclude list
 - creating 69
 - preview 79
 - query order of processing 593
 - size restriction 80
 - to control processing 110
- include-exclude options file
 - specifying path and file name of 392
 - Unicode-enabled file spaces 392
- include-exclude processing
 - options for 71
 - overview 71
- incrbydate option 399
- incremental and selective commands with snapshotroot option 122
- incremental backup
 - associating local snapshot with server file space 475
 - back up new and changed files with modification date later than last backup 399
 - by date 103
 - client command line 103
 - client domain 343
 - client Java GUI 101
 - command line 103
 - description 112
 - directories, processing overview 112
 - GPFS, multinode cluster environment 343
 - memory-conserving algorithm 411
 - new and changed files 112
 - new and changed files with modification date later than last backup 399
 - of directories
 - processing overview 112
 - process a list of files 377
 - with image backup 131, 544
- incremental command 564

- incremental command (*continued*)
 - journal-based backup 568
- incremental option 400
- incremental-by-date
 - command line 103
 - comparing with incremental, journal-based 116
 - description 116
 - of directories
 - processing overview 116
 - when to use 116
- incremental-by-date backup 116
 - client command line 103
 - client Java GUI 101
 - using with image backup 132
- incrthreshold option 401
- input strings
 - containing blanks 91
- install
 - backup-archive scheduler service 243
- installation requirements 6
 - client 6
- installation steps
 - Tivoli Storage Manager on Windows clients 10
- installing
 - overview 1
 - Windows client 10
- installing Tivoli Storage Manager
 - silent installation 11
- interactive mode 525
- interactive session
 - ending 572
 - starting 89, 572
 - using 572
- Internet, searching for problem resolution xvii, xviii

J

- Java GUI
 - configuration restrictions 88
- journal database files
 - errorlog 35
 - journaldir 35
 - NlsRepos 35
- journal engine service
 - configuring 33
- journal-based backup 114, 568
 - comparing with incremental, incremental-by-date 116
 - excluding directories 73
 - excluding files 73
 - include-exclude options
 - journal-based backup 73
 - performing traditional full incremental, instead of 418, 568
 - specifying how to respond to unsuccessful expire of object 401
 - when to use 116
- journaled file space
 - specifying directories with active files to expire 401
- journalpipe 35
- journalpipe option 402
- JournalSettings stanza 35

K

- keyboard 653
- knowledge bases, searching xvii

L

- LAN-based image backup
 - online and offline image backup 540
- LAN-free data movement 355
 - enabling communications for 106, 402, 404
 - options 107
 - prerequisites 106
 - shared memory port for 404
- lanfreecommmethod option 402
- lanfreshmport option 404
- lanfreetcpport option 404
- lanfreetcpserveraddress option 405
- language option 406
- last access date
 - specifying whether to update during backup or archive 112, 431
- latest option 407
- limitations
 - Hyper-V support 147
- limited bandwidth
 - back up 109
- local backup set
 - enabling GUI for local restore 180
- local snapshot
 - associating local snapshot with server file space 122
- localbackupset option 408
- log
 - See also* schedule log
 - DSM_LOG environment variable 361, 450
 - error log, pruning 359
 - errorlogname option 361
 - errorlogretention option 361
 - schedlogname option 450, 639
 - schedlogretention option 450, 639
 - specifying path and file name 361, 450, 639
 - web client 449
- logical volume
 - image backup of 129
 - restoring 178
- Logical Volume Snapshot Agent (LVSA)
 - configuring for online image backup 65
 - configuring for open file support 66
- logs
 - dsmsched.log 452
 - dsmsched.pru 452
 - dsmwebcl.log 452
 - dsmwebcl.pru 452
- loop command 572
- LVSA
 - bring up an application after LVSA starts 428
 - configuring for online image backup 65
 - configuring for open file support 66
 - quiesce an application before LVSA starts 434

M

- macro command 573
- maintenance 4, 306
- manageservices option 409
- management class
 - assigning 153
- management classes 110
 - assigning to directories 239, 339
 - assigning to files 237
 - binding archive files to 208
 - binding to files 239
 - default 232
 - displaying 233
 - displaying information about 594
 - overriding during archive processing 239
 - overriding the default 237
 - processing 237
 - questions to consider 237
 - selecting for files 237
 - specifying with include option 237
 - using management class, example 237
- maxcmdretries option 410
- memory requirements
 - Windows client 6
- memoryefficientbackup option 411
- messages
 - displaying on screen 503
 - specifying language type 406
 - stop displaying 438
- Microsoft Cluster Server cluster nodes
 - FAQs 63
 - installing scheduler service 50
 - installing Tivoli Storage Manager 50, 63
- Microsoft Dfs junction
 - process as junction or directory 336
 - restore 611
- Microsoft Dfs trees and files
 - back up 157
- migration
 - Web client 3
 - Web client language files 3
- mobile dial-up support 93
- mode option 412
- modes
 - batch 525
 - interactive (loop) 525
- modified mode 236
- monitor option 414
- monitor process command 574

N

- Named Pipe communication method
 - options 271
- namedpipename option 415
- NAS
 - backing up file systems 135
 - deleting file spaces 198, 558
 - query node command 595
 - restore file systems 200, 627
 - restore NAS command 627
- NAS file servers using CIFS
 - backup and restore 139

- NAS file systems backup
 - command line 137
 - Web client
 - GUI 136
- nasnodename option 416
- Net Appliance
 - backing up CIFS share definitions 149
- Network Attached Storage (NAS)
 - assigning management class to file systems 393
 - backup file systems 135
 - specifying full or differential backup 412
- Network Data Management Protocol (NDMP) 8
- Network File System (NFS)
 - backup file systems 149
- network-attached storage (NAS)
 - display nodes for which admin ID has authority 595
- network-attached storage (NAS)
 - backup file systems 545
 - cancel backup and restore processes 552, 574
 - deleting file spaces 198, 558
 - display file spaces on server 588
 - excluding files from backup 364
 - monitoring backup or restore operations 414
 - querying file system images belonging to 581
 - restore file systems 200, 627
 - specifying for query 500
 - specifying node name for operations 416
 - specifying whether to save table of contents for each file system backup 496
- networked file systems
 - include-exclude statements 73
 - networked file systems 73
- new in this release xxiii
- no query restore 166
- node
 - specifying type to query 500
- node name 19
- node name, setting 19
- node option, summary 272
- nodename option 197, 417
- nojournal option 418
- noprompt option 419
- NTFS
 - backing up volume mount points 569
 - Restoring volume mount points 610
- NTFS file spaces
 - back up 155
- NTFS mounted volumes
 - Backing up data on 570
- NTFS security information
 - bypass processing of 461
 - whether to compute CRC for comparison of 462
- numberformat
 - specifying 419
- numberformat option 419

O

- offline image backup 129
- online help
 - displaying 95
 - online forum 95
 - service and technical support 95
- online image
 - specifying the size of the cache files 468
- online image backup 129
 - specifying gap size of striped volumes 389
- open file support
 - for backup operations 99
 - include files for 393
 - installing and configuring 66
 - overview 99
 - snapshot 206
 - specifying the size of the cache files 468
- open registration
 - permissions 69
 - using 69
- operating system re-installation 173
 - Windows Vista, 7, 2008 170
- operating system requirements
 - clients 6
- optfile option 421
- options
 - adlocation 298
 - archive, summary 272
 - archmc 299
 - asnodename 299
 - asrmode 301, 518
 - auditlogging 301
 - auditlogname 304
 - authorization options 287
 - autodeploy 306
 - autofsrename 307
 - backmc 309
 - backup
 - excluding system state 364
 - backup, summary 272
 - backupsetname 310
 - casesensitiveaware 311
 - central scheduling, summary 285
 - changingretries 312
 - class 313
 - clientview 313
 - clusternode 317
 - collocatebyfilespec 319
 - command processing, summary 287
 - commmethod 320
 - commrestartduration 321
 - commrestartinterval 322
 - communication, summary 270
 - compressalways 322
 - compression 323
 - computername 324
 - console 325, 326
 - datacenter 328
 - datastore 329
 - dateformat 329
 - dedupcachepath 331
 - dedupcachesize 332
 - deduplication 332
 - deletefiles 334

- options file
 - ASR recovery
 - Windows Vista, 7, 2008 125
- owner security information (SID)
 - back up 155

P

- parameters
 - yes and no, alternatives 298
- partial incremental
 - incremental by date, running 103
- Passport Advantage xix
- password
 - changing 93, 649
 - number of characters 93
 - setting 422
 - setting for client 69
 - specifying whether to generate
 - automatically or set as user prompt 423
 - using 88
 - valid characters 93
- password option 422
- passwordaccess option 423
- performance
 - improving speed of backups, restores,
 - archives, retrieves 271
 - restore operations 165
 - transaction options 288
 - transaction processing 499
- performing traditional full incremental backup 568
- Persistent Storage Manager 140
 - back up 140
- pick option 424
- pitdate 425
- pittime option 426
- point-in-time restore
 - image backup 544
- policies, storage management 231
- policy domains
 - default policy domain 231
 - standard policy domain 231
- policy sets
 - active policy set 231
- portable media
 - restoring backup sets 180
- postnschedulecmd option 427
- postsnapshotcmd option 428
- preferences editor
 - excluding domains from backup 102
- prenschedulecmd option 429
- preschedulecmd option 429
- preservelastaccessdate option 431
- preservevpath option 432
- Presnapshotcmd option 434
- preview
 - include-exclude list 79
- preview archive command 574
- preview backup command 575
- primary group SID
 - back up 155
- problem determination
 - describing problem for IBM Software Support xix

- problem determination (*continued*)
 - determining business impact for IBM Software Support xix
 - submitting a problem to IBM Software xx
- processing options
 - authorization 287
 - backup and archive 272
 - central scheduling 285
 - communication 270
 - diagnostics 289
 - error processing 287
 - format and language 286
 - node option 271, 272
 - overview 269
 - restore and retrieve 282
 - server and node 271
 - specifying in commands 289
 - transaction processing 288
 - using 269
 - Web client 289
- processing time 101
- Protecting cluster disks 56
- proxied session restrictions 119, 120, 208
- proxy backup of VMware virtual machine data
 - backing up data 141
- publications
 - download xiii
 - order xiii
 - search xiii
 - Tivoli Storage Manager xiv

Q

- query
 - amount of information that displays
 - on screen 455
 - ASR 581
 - backups, establish point-in-time 425, 426
 - based on date and time of backup,
 - archive 381, 383
 - description for 334
 - display active and inactive objects 391
 - files for another node 382
 - group
 - command 590
 - display members of 460
 - include-exclude list 593
 - NAS or client objects 313
 - nodes to which client has proxy
 - authority 208
 - nodes to which client has proxy node
 - authority 119
 - process directories only (not files) 340
 - scrolling preferences after displaying
 - information on screen 456
 - system information 599
 - system state 601
- query access command 576
- query adobjects command 577
- query archive command 578
- query ASR command 581
- query backup command 581

- query backupset command 585, 587
- query filespace command 588
- query group command 590
- query image command 591
- query inclexcl command 593
- query mgmtclass command 594
- query node command 595
- query options command 596
- query restore command 597
- query schedule
 - enhanced 219
- query schedule command 598
- query schedule command, enhanced 598
- query session command 598
- query systeminfo command 599
 - encryption processing 599
- query systemstate command 601
- query VM command 602
- queryschedperiod option 436
- querysummary option 437
- quiet option 438

R

- raw logical volume
 - image backup of 129
 - restoring 178
- reanimate
 - tombstone objects 191
- rebinding files to a different management class 240
- registering
 - client with server 68
 - using closed registration 69
 - using open registration 69
- remote network connection
 - establishing 93
- remotely accessed files
 - excluding 75
 - UNC names 75
- removable media
 - back up 154
- repairing installations 15
- replace option 439
- resetarchiveattribute option 441
- resourceutilization option 442
- restart restore command 604
- restartable restore 166
- restartable restore sessions, display 597
- restore 624
 - active and inactive file versions 611
 - active directory objects 191, 192, 193, 195
 - active version 162
 - ASR (Automated System Recovery) files 169
 - ASR recovery mode 301, 518
 - authorizing another user 195
 - Automated System Recovery (ASR) 614
 - backup set
 - supported tape devices 615, 620
 - backup sets
 - overview 180
 - backups, establish point-in-time 425, 426

- restore (*continued*)
 - based on date and time of
 - backup 381, 383
 - classic (also known as standard) 166
 - create list of backup versions to 424
 - directories 163
 - display active and inactive
 - objects 391
 - estimating processing time 101
 - files 163
 - files and directories 163
 - files belonging to another node 197
 - files for another node 382
 - from file spaces that are not
 - Unicode-enabled 612
 - from portable media
 - overview 180
 - from system state backup 191
 - group
 - command 622
 - GUI, displaying active and inactive
 - versions 94
 - image 178
 - considerations 624
 - enable detection of bad sectors on
 - target volume 504
 - to a file 391
 - using chkdsk tool to repair 178
 - using fsck tool to repair 178
 - image, suppress confirmation
 - prompt 419
 - improving speed using share
 - memory 271
 - inactive version 162
 - large number of files 165
 - list of files 377
 - local backup set using the GUI 180
 - logical volume 178
 - Microsoft Dfs junction 611
 - Microsoft Dfs trees and files 177
 - modify CAD and agent services 195
 - most recent backup version 407
 - NAS file systems 200
 - command line 203
 - Web client 200
 - NAS files and directories using web
 - client 201
 - Net Appliance CIFS shares 186
 - no query 166
 - overview 161
 - primary restoration of sysvol 384
 - primary tasks 161
 - process directories only (not
 - files) 340
 - raw logical volume 178
 - replace existing file with latest
 - backup 389
 - restartable 166
 - restrictions and limitations 193
 - sorting file list 94
 - sparse files 612
 - sparse files to a non-NTFS file
 - system 612
 - sparse files, size restriction for 612
 - standard (also known as classic) 166
 - starting a web client session 91
 - summary of options 282

- restore (*continued*)
 - system state 629
 - to different workstation 505
 - using commands 163
 - using fsck tool to repair 624
 - using GUI and command line 192
 - using the GUI 163
 - using universal naming convention
 - names 162
 - VMware Consolidated Backup 186
 - whether to prompt before overwriting
 - existing files 439
 - Windows Server 2003 system
 - state 167
 - Windows Vista system state 167
 - workstation, to another 197
 - restore adobjects command 612
 - restore asr command 614
 - restore backupset command 615, 620
 - restore backupset command
 - considerations 182, 619
 - restore command 605
 - using multiple 165
 - restore group command 622
 - restore image command 624
 - restore maximum file size 111
 - restore NAS command 627
 - restore procedure
 - ASR 172
 - user profiles 177
 - Windows 2003 176
 - Windows Vista, 7, 2008 172
 - restore systemstate command 629
 - restore vm command 631
 - Restoring data on mounted volumes
 - NTFS 610
 - Restoring data on NTFS mounted
 - volumes 610
 - Restoring NTFS volume mount
 - points 610
 - restoring point-in-time 199
 - using command line 199
 - using GUI 199
 - restoring your system
 - ASR recovery mode
 - Windows 2003 173
 - Windows Vista, 7, 2008 170
 - restrictions
 - asnodename option 299
 - runasservice and encryptkey 446
 - runasservice and passwordaccess 446
 - runasservice and replace 446
 - within a proxied session 119, 120, 208
 - retain extra versions attribute 234
 - retain only versions attribute 235
 - retain versions attribute 237
 - retention grace period
 - archive 232, 240
 - backup 232, 240
 - retrieve
 - archive copies 210
 - archive files by name 211
 - authorizing another user 195
 - based on date and time of
 - archive 381, 383
 - description for 334

- retrieve (*continued*)
 - files belonging to another node 197
 - files for another node 382
 - improving speed using share
 - memory 271
 - list of files 377
 - primary tasks 205
 - process directories only (not
 - files) 340
 - replace existing file with latest archive
 - if existing file is newer 389
 - sorting file list 94
 - starting a web client session 91
 - summary of options 282
 - to different workstation 505
 - whether to prompt before overwriting
 - existing files 439
 - workstation, to another 197
 - retrieve command 635
 - retrieve maximum file size 111
 - retrieving
 - archives using command line 211
 - retryperiod option 445
 - return codes for operations 229
 - revokeremoteaccess option 445
 - runasservice option 446
- S**
- SAN
 - restoring backup sets using 620
- schedcmddisabled option 447, 448
- schedlogmax option 449
- schedlogname option 450
- schedlogretention option 452
- schedmode option 453
- schedrestretretdisabled option 455
- schedule command 639
- schedule log
 - controlling the size 449
 - specifying number of days to keep
 - entries and whether to save pruned
 - entries 452
 - specifying path and file name to store
 - schedule log information 450
- scheduled (automated) backups
 - closing files before backup 152
 - displaying scheduled work 219, 221
 - options for 225
 - process commands after backup 427
 - process commands before
 - backup 429
 - restart applications after backup 152
 - starting 26
- scheduled commands
 - enabling-disabling 226
- scheduled events, displaying 598
- scheduled services
 - disabling scheduled commands 447, 448
 - restrictions for NAS file systems 135
- scheduler
 - configuring 24
 - displaying scheduled work 219, 221
 - event logging 221
 - number of hours between contacts to
 - server for scheduled work 436

- scheduler (*continued*)
 - number of minutes between attempts
 - to process scheduled commands 445
 - options for 225
 - polling mode or prompted mode 453
 - starting 26
 - whether server or client initiates sessions through firewall 457
 - whether to disable execution of restore or retrieve operations 455
- scheduler service
 - installing on Microsoft Cluster Server cluster nodes 50
 - installing on Veritas Cluster Server cluster nodes 50
- Scheduler Service Configuration Utility 24
- scheduler wizard 24
- scheduling
 - client node proxy backup 119, 121
 - scrolllines option 455
 - scrollprompt option 456
- Secure Sockets Layer (SSL)
 - establishing communications with 29
- selective backup 475, 640
 - client command line 103
 - client Java GUI 101
 - command line 103
 - overview 103, 118
- selective command 640
- serialization
 - copy serialization
 - dynamic 235
 - shared static 235
 - static 235
- server
 - communicating with 19
 - establishing communications through firewall 27
 - establishing communications with 19
 - establishing communications with Secure Sockets Layer (SSL) 29
 - query Active Directory for communication method and server with which to connect 501
 - TCP/IP address of Tivoli Storage Manager server 493
 - TCP/IP port address for 492
- server and node options
 - summary 271
- server options
 - Sslfipsmode 481
- service and technical support 95
- Service recovery settings 217
- session information, displaying 598
- sessioninitiation option 457
- set access command 644
 - restore-retrieve authorization 195
- set event command 646
- set password command 649
- setting
 - environment variables
 - DSM_CONFIG 22
 - DSM_DIR 22
 - DSM_LOG 22
 - user privileges 92
- shared dynamic serialization 235, 312
- shared memory communication method
 - options 271
- shared static serialization 235, 312
- shmport option 459
- showmembers option 460
- silent installation 11
- skipmissingsyswfiles option 460
- skipntpermissions option 461
- skipntsecuritycrc option 462
- snaptdiff option 67, 463
- snapshot
 - open file support 206
- snapshot difference 67, 463
- snapshotcachelocation option 467
- snapshotcachesize option 468
- snapshotfsidretries option 469
- snapshotfsidlewait option 470
- snapshotproviderdfs option 473
- snapshotproviderimage option 474
- snapshotroot option 475
- snapshotroot option with incremental and selective commands 122
- software requirements
 - ASR
 - Windows 2003 174
- software support
 - describing problem for IBM Software Support xix
 - determining business impact for IBM Software Support xix
 - submitting a problem xx
- Software Support
 - contact xix
- sparse files
 - restore size restriction 612
 - restoring 612
 - restoring to a non-NTFS file system 612
- specifying whether to update last access date 431
- srvoptsetencryptiondisabled option 477
- srvprepostscheddisabled option 478
- srvprepostsnapdisabled option 479
- SSL (Secure Socket Layer)
 - establishing communications with 29
 - ssl option 480
 - Sslfipsmode option 481
- stagingdirectory option 481
- standard (classic) restore 166
- standard management class
 - copy destination 236
 - copy frequency 234
 - copy group name 233
 - copy mode
 - absolute 236
 - modified 236
 - copy serialization 235
 - copy type 234
 - deduplicate data attribute 237
 - default values 233
 - retain extra versions 234
 - retain only version 235
 - retain versions 237
 - versions data deleted
 - active versions 234
 - inactive versions 234
- standard management class (*continued*)
 - versions data exists 234
- standard policy domain 231
- Start the client scheduler at startup 217
- starting
 - automatically 93
 - overview 1
- starting a session
 - batch mode 89
 - interactive mode 89
- static serialization 235
- storage
 - displaying restartable restore sessions 597
- Storage Agent
 - for LAN-free data movement 106
 - using for LAN-free data movement 355
- storage area network
 - for LAN-free data movement 106
 - restoring backup sets using 355, 620
 - using for LAN-free data movement 355
- storage management policies 231
 - assigning management classes to files 153
 - copy groups 232
 - default management class 231
 - display on backup-archive client or Web client GUI 153
 - include-exclude list 232
 - management classes 232
 - policy domains
 - default 231
 - standard 231
 - policy sets
 - active policy set 231
- subdir option 482
- subdirectories
 - archive 208
 - include in backup 103
- subfilebackup option 484
- subfilecachepath option 485
- subfilecachesize option 486
- support
 - gathering system information
 - for 325, 379, 599
 - support contract xix
 - support information xvi
 - support subscription xix
- swing-enabled browser
 - necessary to run web client 91
- syntax diagram
 - reading xx
 - repeating values xxi
 - required choices xxii
- system access control list (auditing information)
 - back up 155
- system files
 - excluding 75
- system information
 - gathering 325, 379
- system objects
 - primary restoration of sysvol 384
- system recovery
 - Windows 2003 173

- system recovery (*continued*)
 - Windows Vista, 7, 2008 170
- system restore 170
- system services
 - assigning management class 122
 - components, exclude from backup 122
 - exclude from backup processing 364
- system state
 - assigning management class 73, 122, 393
 - back up 122, 547
 - display active and inactive objects 391
 - exclude from backup processing 73, 364
 - query 601
 - restore 629
 - restore from backup set 615, 620
 - restoring 167

T

- tapeprompt option 486
- tasks
 - closed registration 68
 - open registration 68
- TCP/IP communication method
 - options 270
- tcpadminport option 488
- tcpbuffsize option 488
- tcpcadaddress option 489
- tcpclientaddress option 490
- tcpclientport option 491
- tcpnodelay option 491
- tcpserveraddress option 493
- tcpwindowsize option 494
- time format
 - specifying 495
- timeformat option 495
- Tivoli Storage Manager
 - client components
 - Windows client 6
 - communication methods
 - Windows client 7
 - environment prerequisites 6
 - FAQs 63
 - hardware, disk space, memory requirements
 - Windows client 6
 - installation requirements 6
 - installing on Microsoft Cluster Server
 - cluster nodes 50, 63
 - installing on Veritas Cluster Server
 - cluster nodes 50
 - online forum 96
 - password 88
 - upgrading from earlier versions of the product 1
- Tivoli Storage Manager client
 - authentication 86
- Tivoli Storage Manager FastBack
 - configuration 47
- Tivoli Storage Manager FastBack configuration wizard 9, 48
- Tivoli Storage Manager FastBack data
 - backup 148

- Tivoli Storage Manager FastBack data
 - restore 148
- Tivoli Storage Manager FastBack
 - installation requirements 8
- Tivoli Storage Manager on Windows
 - clients
 - installation steps 10
- Tivoli Storage Manager password
 - using 88
- Tivoli technical training xvi
- toc option 496
- todate option 498
- tombstone objects
 - preserving attributes 194
 - reanimate 191
- totime option 498
- traditional full incremental backup 114
- training, Tivoli technical xvi
- transaction processing 499
 - summary of options 288
- txnbytelimit option 499
- type option 500

U

- UNC names
 - excluding files 75
 - remotely accessed files 75
 - restore files 162
- Unicode
 - pre-backup considerations 106, 107
 - renaming file spaces that are not Unicode to Unicode-enabled 307, 564, 640
 - restore from file spaces that are not Unicode-enabled 612
- universal naming convention
 - back up shared files and directories using 156
 - restore 162
 - using to specify domain list 155
- updates_622_client 122, 124, 125, 126, 129, 167, 169, 170, 172, 301, 547, 614, 629
- updates_623_client 412, 549
- updating the client 4
- upgrading
 - online image 3
 - Open File Support (OFS) 3
- upgrading the backup-archive client from earlier versions of the product 1
- usedirectory option 501
- user privileges
 - setting 92
- using multiple sessions 149

V

- v2archive option 502
- verbose option 503
- verifyimage option 504
- Veritas Cluster Server cluster nodes
 - FAQs 63
 - installing scheduler service 50
 - installing Tivoli Storage Manager 50, 63

- versions data
 - deleted attribute 234
 - exists attribute 234
- virtualfsname option 504
- virtualnodename option 505
- VM backup 141
- vmbackdir option 506
- vmbacknodelete option 507
- vmbackuptype option 508, 509
- vmbackvcbtransport 510
- vmchost option 511
- vmcpw option 512
- vmctlmc option
 - options
 - vmctlmc 513
- vmcuser option 513
- vmlist option 514
- vmprocesswithindependent 515
- vmprocesswithprdm 516
- vmvstortransport option 517
- VMware Consolidated Backup
 - backing up data 146
 - restoring data 186
- Volume Shadowcopy Service (VSS)
 - configuring for online image backup 66
 - configuring for open file support 66
- VSS (see Volume Shadowcopy Service) 66
- vStorage backup server
 - backing up data 143
 - off-host backup 143

W

- web client
 - enable to run in a swing-enabled browser 91
 - starting 91
 - supported browsers 91
- Web client
 - configuring in cluster environment 50
 - establishing communications through firewall 27, 386
 - GUI 136
 - NAS
 - file systems backup 136
 - restrict administrator from accessing client running Web client 445
 - restrictions for NAS file systems 135
 - specifying TCP/IP port address for 386
 - summary of options 289
 - unsupported functions 97
 - using through a firewall 519
- webports option 519
- what's new in this release xxiii
- whether to compute CRC for comparison of NTFS security information 462
- wildcard characters
 - guidelines 529
 - include or exclude files 75
 - include or exclude groups of files 76
 - specifying a drive specification in dsm.opt 76

- wildcard characters (*continued*)
 - to include or exclude groups of files 78
- Windows archive attribute
 - reset after backup 441
- Windows client
 - client components 6
 - communication methods 7
 - disk space requirements 6
 - hardware requirements 6
 - installing 10
 - memory requirements 6
 - modify 15
 - reinstall 15
 - upgrade 15
- Windows components
 - installable 6
- Windows supported file systems 8
- WinPE CD
 - Windows Vista, 7, 2008 170



Product Number: 5608-E01, 5608-E02, 5608-E07

Printed in USA

SC23-9792-05

