

Symantec NetBackup™ Administrator's Guide, Volume I

Windows

Release 7.0

Symantec NetBackup™ Administrator's Guide, Volume I

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version 7.0

PN: 20654075

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, and NetBackup are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

www.symantec.com/techsupp/

Contacting Technical Support

Customers with a current maintenance agreement may access Technical Support information at the following URL:

www.symantec.com/techsupp/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and maintenance contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4
Section 1 About NetBackup	31
Chapter 1 Introducing NetBackup interfaces	33
About NetBackup	33
Online documents	35
Using the NetBackup Administration Console	35
Standard and user toolbars	36
Customizing the administration console	37
Configuring NetBackup by using the NetBackup Administration Console	37
NetBackup configuration wizards	37
Activity Monitor utility	39
NetBackup Management utilities	39
Media and Device Management utilities	41
Access Management utility	42
Chapter 2 Administering NetBackup licenses	43
Administering NetBackup licenses	43
Accessing license keys for a NetBackup server	44
Adding new license keys	45
Printing license key lists	45
Deleting license keys	46
Viewing license key properties	46
Exporting license keys	46
Section 2 Configuring hosts	49
Chapter 3 Configuring Host Properties	51
Configuring Host Properties and configuration options	53
About the Host Properties	53
Viewing host properties	54

How to change host properties	55
Exporting host properties	56
About the states of multiple hosts	56
Changing properties on multiple hosts	57
How to interpret the properties of multiple hosts of different operating systems	58
Access Control properties	59
About the Symantec Product Authentication and Authorization tab	60
About the Authentication Domain tab	62
About the Authorization Service tab	65
Active Directory host properties	66
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider	67
Continue with backup if consistency check fails	67
Authorization properties	68
Adding or changing an authorized user	68
Backup Exec Tape Reader properties	69
GRFS advertised name	69
Actual client name	70
Actual path	70
Bandwidth properties	70
From IP address field	72
To IP address field	72
Bandwidth	72
How bandwidth limiting works	72
Busy File Settings properties	73
Working directory	73
Operator's email address	74
Process busy files	74
File action file list	74
Add	74
Add to all	74
Remove	74
Busy file action	75
Retry count	75
How to activate the Busy File Settings host properties	75
Clean-up properties	76
Keep logs	76
Keep vault logs	77
Image cleanup	77
Catalog cleanup wait time	77
Keep true image restoration (TIR) information	77

Move restore job from incomplete state to done state	78
Move backup job from incomplete state to done state	78
Client Name properties	79
Client Attributes properties	80
General tab of the Client Attributes properties	81
Connect Options tab of the Client Attributes properties	83
Windows Open File Backup tab of the Client Attributes properties	85
Client Settings (NetWare) properties	90
Back up migrated files	90
Uncompress files before backing up	90
Keep status of user-directed backups, archives, and restores	90
Client Settings (UNIX) properties	91
Locked file action	91
Keep status of user-directed backups, archives, and restores	92
Reset file access time	92
Megabytes of memory to use for file compression	92
Use VxFS file change log for incremental backups	92
Default cache device path for snapshots	92
Do not compress files ending with list	93
Add button	93
Add to All button	93
Remove button	93
Using the VxFS file change log for incremental backups property	93
Client Settings (Windows) properties	95
General level	96
TCP level	96
Wait time before clearing archive bit	97
Use change journal in incrementals	97
Incrementals based on timestamp	97
Incrementals based on archive bit	97
Time overlap	98
Communications buffer size	98
User directed timeouts	98
Maximum error messages for server	99
Keep status of user-directed backups, archives, and restores	99
Perform default search for restore	99
How to determine if change journal support is useful in your NetBackup environment	99

Guidelines for enabling NetBackup change journal support	100
Credential Access properties	101
NDMP Clients list	101
Disk clients list	101
Data Classification properties	102
Creating a Data Classification	103
Rank column	104
Name column	104
Description column	104
Data Classification ID	104
Default Job Priorities properties	105
Job Type and Job Priority list	106
Job Priority	106
Distributed application restore mapping properties	107
Encryption properties	108
Encryption permissions property	109
Enable encryption property	109
Enterprise Vault properties	110
User Name	111
Password	111
Consistency check before backup	111
Enterprise Vault Hosts properties	112
Exchange properties	112
Exclude Lists properties	114
Use case sensitive exclude list property	115
Exclude list	116
Exceptions to exclude list	116
About the Add to exclude list and Add to exceptions list dialog boxes	117
Syntax rules for exclude lists	117
Traversing excluded directories	119
Fibre Transport properties	120
Preferred	121
Always	121
Never	121
Maximum concurrent FT connections	121
Use defaults from the master server configuration	121
Firewall properties	122
Default connect options	122
Hosts list	123
Attributes for selected hosts	124
Setting up vnetd between a server and a client	125
Setting up vnetd between two servers	126

Enabling logging for vnetd	126
Example setup for using the vnetd port	126
General Server properties	127
Delay on multiplexed restores	128
Check the capacity of disk storage units every	128
Must use local drive	128
Use direct access recovery for NDMP restores	129
Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology	129
Media host override list	129
Forcing restores to use a specific server	130
Disabling the cataloging for duplications of Exchange backups using Granular Recovery Technology (GRT)	131
Global Attributes properties	131
Job retry delay	132
Schedule backup attempts	132
Policy update interval	133
Maximum jobs per client	133
Maximum backup copies	135
Compress catalog interval	135
Maximum vault jobs	135
Administrator email address property	135
Setting up email notifications about backups	136
Configuring the nbmail.cmd script	137
Sending email notifications to the administrator about unsuccessful backups	138
Sending messages to the global administrator about unsuccessful backups	139
Sending messages to the administrator about successful and unsuccessful backups	139
Installing the email utility	140
Logging properties	142
Enable robust logging	144
Global logging level	145
Process specific overrides	145
Debug logging levels for NetBackup services	146
Login Banner Configuration properties	147
Login Banner Heading	148
Text of login banner	148
Show Agree and Disagree buttons on the login banner	148
Removing login banner screen and text	150
Auto log off timeout option	150
Lotus Notes properties	151

Maximum number of logs to restore	151
Transaction log cache path	152
INI file	152
Path	153
Media properties	153
Allow media overwrite property	154
Enable SCSI reserve	155
Allow multiple retentions per media	156
Allow backups to span tape media	156
Allow backups to span disk	157
Enable standalone drive extension	157
Enable job logging	157
Enable unrestricted media sharing for all media servers	158
Media ID prefix (non-robotic)	158
Media unmount delay	158
Media request delay	158
Recommended use for Enable SCSI reserve property	159
NDMP Global Credentials properties	160
NetWare Client properties	161
Network properties	162
NetBackup client service port (BPCD)	162
NetBackup request service port (BPRD)	162
Announce DHCP interval	163
Network Settings Properties	163
Allowed setting	164
Restricted setting	164
Prohibited setting	164
Changing host properties outside of the Administration	
Console	164
REVERSE_NAME_LOOKUP entry	165
Setting the property on UNIX hosts	165
Setting the property on Windows hosts	165
Port Ranges properties	166
Use random port assignments	167
Client port window	167
Client reserved port window	168
Server port window	168
Server reserved port window	169
Restore Failover properties	169
Adding failover servers	171
Retention Periods properties	172
Value	173
Units	173

Retention periods list	173
Schedules list	173
Impact Report button	173
Changing a retention period	173
Suspending volumes	175
Servers properties	175
Master server	176
Additional servers list	176
Media servers list	176
Restricting administrative privileges of media servers	177
Multiple masters that share one Enterprise Media Manager host	178
SharedDisk properties	180
SharePoint properties	180
Symantec Products properties	181
Timeouts properties	182
Client connect timeout	183
Backup start notify timeout	183
File browse timeout	183
Use OS dependent timeouts	184
Media mount timeout	184
Client read timeout	184
Backup end notify timeout	185
Media server connect timeout	185
Universal Settings properties	185
Restore retries	186
Browse timeframe for restores	187
Last full backup	187
Use specified network interface	187
Allow server file writes	187
Accept connections on non reserved ports	187
Enable performance data collection (Windows server only)	188
Client sends mail	188
Server sends mail	188
Client administrator's email	188
Use specified network interface examples	189
UNIX Client properties	190
UNIX Server properties	190
VMware backup hosts properties	191
VSP (Volume Snapshot Provider) properties	192
Windows Client properties	193
Configuration options not found in the Host Properties	194
NBRB_CLEANUP_OBSOLETE_DBINFO	194

	NBRB_ENABLE_OPTIMIZATIONS	194
	NBRB_FORCE_FULL_EVAL	195
	NBRB_REEVAL_PENDING	195
	NBRB_REEVAL_PERIOD	195
	NBRB_RETRY_DELAY_AFTER_EMM_ERR	195
	NBRB_MPX_GROUP_UNLOAD_DELAY	195
	REQUIRED_NETWORK	196
Chapter 4	Configuring server groups	197
	About server groups	197
	Configuring a server group	198
	Server group properties	200
	Deleting a server group	200
Chapter 5	Configuring host credentials	201
	About configuring credentials	201
Chapter 6	Managing media servers	203
	Activating or deactivating a media server	203
	Adding a media server	204
	Decommissioning a media server	205
	Registering a media server	206
	Deleting all devices from a media server	207
	Removing a device host from the EMM database	209
Section 3	Configuring storage	211
Chapter 7	Configuring robots and drives	213
	About optical device support in NetBackup 7.0	214
	About NetBackup robot types	214
	Device configuration prerequisites	215
	About the device mapping file	215
	Updating the device mapping file	216
	About configuring robots and tape drives	216
	Configuring robots and tape drives	217
	About device discovery	217
	Configuring robots and drives by using the wizard	219
	About robot control	219
	Adding a robot	221
	Robot configuration options	222

Adding a tape drive	227
Adding a shared tape drive	229
Tape drive configuration options	229
About drive name rules	233
Configuring drive name rules	233
Adding a tape drive path	235
Correlating tape drives and SCSI addresses on Windows hosts	237
Updating the device configuration by using the wizard	238
Managing robots	239
Changing robot properties	239
Configuring a robot to operate in manual mode	239
Deleting a robot	240
Moving a robot and its media to a new media server	240
Managing tape drives	242
Changing a drive comment	242
About downed drives	242
Changing a drive operating mode	243
Changing a tape drive path	243
Changing a drive path operating mode	244
Changing tape drive properties	244
Changing a tape drive to a shared drive	245
Cleaning a tape drive from the Device Monitor	245
Deleting a drive	246
Resetting a drive	246
Resetting the mount time	247
Setting drive cleaning frequency	248
Viewing drive details	248
Performing device diagnostics	249
About device diagnostic tests	249
Running a robot diagnostic test	249
Running a tape drive diagnostic test	251
Managing a diagnostic test step that requires operator intervention	252
Obtaining detailed information for a diagnostic test step	252
Verifying the device configuration	252
Replacing a device	253
Updating device firmware	255
About the NetBackup Device Manager	256
Stopping and restarting the Device Manager	256

Chapter 8	Configuring tape media	257
	About adding volumes	257
	About adding robotic volumes	258
	About adding stand-alone volumes	258
	About tape volumes	259
	About NetBackup media types	259
	About alternate media types	261
	About WORM media	261
	About WORM media limitations	262
	How to use WORM media in NetBackup	262
	Adding volumes by using the wizard	265
	Adding volumes by using the Actions menu	265
	Volume properties (add volumes)	266
	Managing volumes	269
	Changing the group of a volume	269
	Changing the owner of a volume	270
	Changing the pool of a volume	270
	Changing volume properties	270
	Deassigning a volume	273
	Deleting a volume	274
	Erasing a volume	274
	Exchanging a volume	276
	About frozen media	278
	Freezing or unfreezing a volume	278
	About injecting and ejecting volumes	279
	Injecting volumes	279
	Ejecting volumes	280
	About media ejection timeout periods	282
	About rescanning and updating bar codes	282
	Rescanning and updating bar codes	283
	About labeling NetBackup volumes	284
	Labeling a volume	285
	About moving volumes	286
	Moving volumes by using the robot inventory update option	287
	Moving volumes by using the Actions menu	287
	Recycling a volume	290
	Suspending or unsuspending a volume	291
	About volume pools	292
	About scratch volume pools	292
	Adding a volume pool	293
	Volume pool properties	293

Managing volume pools	295
Changing the properties of a volume pool	295
Deleting a volume pool	296
About volume groups	296
About media sharing	297
Configuring media sharing	298
Configuring unrestricted media sharing	298
Configuring media sharing with a server group	299
 Chapter 9	
Inventorying robots	301
About robot inventory	302
About previewing volume configuration changes	303
When to inventory a robot	303
About showing a robot's contents	307
About inventory results for API robots	308
Showing the media in a robot	309
About comparing a robot's contents with the volume configuration	310
Comparing media in a robot with the volume configuration	311
About updating the volume configuration	312
Determine robot capabilities before you update the volume configuration	313
Updating the volume configuration with a robot's contents	314
Robot inventory options	316
Configuring media settings	317
Media settings - existing media	318
Media settings - new media	320
About bar codes	325
About bar code advantages	325
About bar code best practices	325
About bar code rules	326
About media ID generation rules	329
Configuring bar code rules	329
Barcode rules settings	330
Configuring media ID generation rules	333
Media ID generation options	334
Configuring media type mappings	336
About adding media type mapping entries	337
About the default and allowable media types	338
About the physical inventory utility	343
About the vmphyinv features	343
About vmphyinv requirements and restrictions	343

	When to use vmphyinv	344
	How vmphyinv performs a physical inventory	344
	Example volume configuration updates	349
	Example 1: Removing a volume from a robot	350
	Example 2: Adding existing stand-alone volumes to a robot	351
	Example 3: Moving existing volumes within a robot	353
	Example 4: Adding new volumes to a robot	354
	Example 5: Adding cleaning tapes to a robot	356
	Example 6: Moving existing volumes between robots	357
	Example 7: Adding existing volumes when bar codes are not used	358
Chapter 10	Configuring disk storage	361
	Configuring BasicDisk storage	361
	Configuring NearStore storage	361
	About SharedDisk support in NetBackup 7.0 and later	362
	Configuring disk pool storage	363
Chapter 11	Configuring storage units	365
	About the Storage utility	365
	Using the Storage utility	366
	About storage units	366
	Creating a storage unit using the Device Configuration Wizard	368
	Creating a storage unit using the Actions menu	368
	Creating a storage unit by copying a storage unit	369
	Changing storage unit settings	369
	Deleting storage units	369
	Media Manager storage unit considerations	370
	Disk storage unit considerations	372
	About NDMP storage unit considerations	378
	About storage unit settings	380
	Absolute pathname to directory or volume setting	380
	Density setting	381
	Disk pool setting	381
	Disk type selection	381
	Enable block sharing setting	382
	Enable multiplexing setting	382
	High water mark setting	382
	Low water mark setting	382
	Maximum concurrent write drives setting	383
	Maximum concurrent jobs setting	383

	Maximum streams per drive setting	385
	Media server setting	385
	NDMP host setting	387
	On demand only setting	387
	Only use the following media servers	388
	Properties button	388
	Reduce fragment size setting	390
	Robot number setting	391
	Robot type setting	391
	Staging relocation schedule setting (for basic disk staging only)	391
	Storage device setting	392
	Storage unit name setting	392
	Storage unit type setting	392
	Temporary staging area setting	392
	Transfer throttle setting	392
	Use any available media server setting	393
Chapter 12	Staging backups	395
	About staging backups	395
	About the two staging methods	395
	Basic disk staging	396
	Creating a basic disk staging storage unit	397
	Disk staging storage unit size and capacity	399
	Finding potential free space on a BasicDisk disk staging storage unit	400
	About the Disk Staging Schedule dialog box	402
	Basic disk staging limitations	404
	Initiating a relocation schedule manually	405
Chapter 13	Configuring storage unit groups	407
	About Storage unit groups	407
	Creating a storage unit group	407
	Deleting a storage unit group	409
	Storage unit selection criteria within a group	409
	Prioritized storage unit selection	410
	Failover storage unit selection	410
	Round robin storage unit selection	410
	Media server load balancing storage unit selection	411
	Exception to the storage unit selection criteria	413
	Disk spanning within storage unit groups	414

Chapter 14	Configuring storage lifecycle policies	415
	Storage lifecycle policy overview	415
	Creating a storage lifecycle policy	416
	Storage lifecycle policy name	417
	Data classification option	417
	Duplication job priority setting	419
	Deleting a storage lifecycle policy	419
	Adding a storage destination to a storage lifecycle policy	421
	Use for: Backup, duplication, or Snapshot destination	423
	Storage unit or storage destinations	423
	Volume pool for storage destinations	424
	Media owner for storage destinations	424
	Fixed retention type for storage destinations	424
	Staged capacity managed retention type for storage destinations	425
	Expire after duplication retention type for storage destinations	426
	Alternate read server for storage destinations	426
	Preserve multiplexing for storage destinations	427
	Hierarchical view of storage destinations	427
	Adding a hierarchical duplication destination	429
	Adding a non-hierarchical duplication destination	429
	Modifying the source of a hierarchical duplication destination	430
	Removing a destination from the storage destination list	431
	Hierarchy example	431
	Writing multiple copies using a storage lifecycle policy	433
	Use only one method to create multiple copies	433
	Destination order determines the copy order	433
	Ensuring successful copies using lifecycles	434
	Storage lifecycle policy versions	435
	How to create a new version	435
	Administrator actions that do not create a new version	435
	When do changes to storage lifecycle policies become effective?	437
	Deleting old storage lifecycle policy versions	438
	LIFECYCLE_PARAMETERS file for optional duplication job configuration	438
	CLEANUP_SESSION_INTERVAL_HOURS	438
	DUPLICATION_GROUP_CRITERIA	439
	DUPLICATION_SESSION_INTERVAL_MINUTES	439
	IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS	439

MIN_GB_SIZE_PER_DUPLICATION_JOB	440
MAX_GB_SIZE_PER_DUPLICATION_JOB	440
MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB	
4 4 0	
TAPE_RESOURCE_MULTIPLIER	441
VERSION_CLEANUP_DELAY_HOURS	441
LIFECYCLE_PARAMETERS file example	442
Logic for batch creation	442
Using the nbstlutil command to administrate lifecycle	
operations	443
When to use nbstlutil	444

Section 4 Configuring backups 445

Chapter 15 Creating backup policies 447

Using the Policies utility	448
Planning for policies	449
Group the clients	449
Gather client information	450
Consider storage requirements	450
Backup schedule considerations	451
How to group by general attributes	453
Maximize multiplexed backups	453
Evaluate backup times	453
Creating a policy using the Backup Policy Configuration Wizard	455
Creating a policy without using the Backup Policy Configuration	
Wizard	456
Changing policies	456
Adding or changing schedules in a policy	457
Adding or changing clients in a policy	457
Adding or changing backup selections in a policy	458
Moving policy information from one server to another	458
Cutting, copying, and pasting policy items	459
Changing multiple policies at one time	459
Deleting schedules, backup selections, or clients from a	
policy	461
About the Policy attributes	461
Policy type attribute	462
Data classifications attribute	465
Policy storage attribute	465
Policy volume pool attribute	467
Checkpoint restart for backup jobs	468

Limit jobs per policy attribute	471
Job priority attribute	473
Media owner attribute	473
Go into effect at attribute	474
Backup network drives attribute	474
Follow NFS attribute	476
Cross mount points attribute	478
Compression attribute	480
Encryption attribute	482
Collect disaster recovery information for Intelligent Disaster Recovery attribute	483
Collect disaster recovery information for Bare Metal Restore attribute	483
Collect true image restore information attribute	483
Collect true image restore information with move detection attribute	484
Allow multiple data streams attribute	486
Disable client-side deduplication attribute	488
Enable granular recovery attribute	489
Keyword phrase attribute	489
Snapshot Client Attributes	490
Microsoft Exchange Attributes	490
About the Schedules tab	490
About the Schedule Attributes tab	491
Name attribute	491
Type of backup attribute	492
Synthetic backup attribute	499
Calendar schedule type	500
Frequency schedule type	500
Instant recovery options	502
Multiple copies attribute	503
Override policy storage selection attribute	508
Override policy volume pool attribute	509
Override media owner attribute	509
Retention attribute	510
Media multiplexing attribute	512
Using the Start Windows tab	517
Creating a schedule window	518
Example of schedule duration	519
Creating time windows on successive days	519
Copying a time window	520
Changing a time window	520
Moving a time window	520

Deleting a time window	520
Deleting all time windows	520
Using the Exclude Dates tab	521
Excluding dates from a policy	521
Using the Calendar Schedule tab	521
Scheduling by specific dates	522
Scheduling by recurring week days	523
Scheduling by recurring days of the month	524
Considerations for user schedules	525
How to plan user backup and archive schedules	525
How to create separate policies for user schedules	526
How to use a specific policy and user schedule	526
Backup window considerations	526
How NetBackup determines which schedule to run next	526
Windows that span midnight	529
How open schedules affect the different schedule types	530
Runtime considerations	534
About the Clients tab	535
Adding clients to a policy	535
Browse for Hyper-V virtual machines	536
About the Backup Selections tab	537
Changing backup selections for standard policies	537
Changing backup selections for database policies	539
Changing backup selections for Oracle or DB2 policies	539
Reducing backup time	540
Verifying the backup selections list	541
Path rules for Microsoft Windows file backups	543
Path rules for Windows disk image (Raw) backups	545
Path rules for Windows registry backup	547
Hard links to files (NTFS volumes or UNIX)	547
Pathname rules for UNIX clients	549
About the path rules for NetWare NonTarget clients	555
Path rules for NetWare Target clients	557
Path rules for clients that run extension products	557
Backup selections list directives	558
Backup selections list directives for multiple data streams	563
Excluding files from backups	567
About the Disaster Recovery tab	569
Path	570
Logon	571
Password	571
Send in an email attachment field	571
Critical policies list	572

Creating a Vault policy	573
Performing manual backups	574
Active Directory granular backups and recovery	575
System requirements for Active Directory granular backups and recovery	575
Creating a policy that allows Active Directory granular restores	576
Restoring Active Directory objects	578
Restore issues	580
 Chapter 16	
Synthetic backups	583
About synthetic backups	583
Policy considerations and synthetic backups	584
Schedules that must appear in a policy for synthetic backups	585
Adding clients to a policy for synthetic backups	585
Types of synthetic backups	586
Synthetic full backups	586
Synthetic cumulative incremental backups	587
When to use synthetic backups	589
Synthetic backup jobs create two sets of catalog files	591
Change journal and synthesized backups	592
True image restore and synthesized backups	592
Checkpoint restart and synthesized backups	592
Displaying synthetic backups in the Activity Monitor	593
Logs produced during synthetic backups	593
Synthetic backups and directory and file attributes	594
Using the multiple copy synthetic backups method	594
Configuring multiple copy synthetic backups	596
Configuration variables	596
Simple configuration example	598
Advanced configuration example	598
Optimized synthetic backups using OpenStorage	598
 Chapter 17	
Protecting the NetBackup catalog	599
About NetBackup catalogs	599
Parts of the catalog	600
About the image database	600
About the NetBackup relational database	602
Protecting the catalog	604
About online, hot catalog backups	605
Recovering the catalog	615

	Disaster recovery emails and the disaster recovery file	615
	Archiving the catalog	616
	Creating a catalog archiving policy	618
	Catalog archiving commands	618
	When to catalog archive	620
	Using Vault with the catalog archiving feature	620
	Extracting images from the catalog archives	621
	Estimating catalog space requirements	621
	File size considerations	623
	About the binary catalog format	623
	Moving the image catalog	623
	Indexing the catalog for faster access to backups	625
	Compressing the image catalog	625
	Uncompressing the image catalog	627
Chapter 18	About the NetBackup relational database	629
	Installing the NetBackup relational database (NBDB)	629
	NetBackup master server installation	631
	About the NetBackup configuration entry	637
	Sybase SQL Anywhere server management	637
	Clustered environments	637
	Post-installation tasks	638
	Changing the database password	638
	Moving NBDB database files after installation	639
	Adding a mirrored transaction log	640
	Creating the NBDB database manually	640
	About backup and recovery procedures	642
	Database transaction log	643
	About catalog recovery	643
	Backing up and recovering the relational databases	644
	Unloading the database	645
	Terminating database connections	645
	About the Database Administration tool	646
	Running the Database Administration Tool	647
	Moving the NetBackup database from one host to another	647
	Cluster considerations with the EMM server	650
	Moving the EMM server to a Windows cluster	651
	About moving the EMM server from a Windows cluster	651
Chapter 19	Using the Catalog utility	653
	About the Catalog utility	653
	Searching for backup images	654

Verifying backup images	656
Viewing job results	657
Promoting a copy to a primary copy	657
Duplicating backup images	659
About multiplexed duplication	664
Jobs that appear while making multiple copies	665
Expiring backup images	665
Importing backups	666
Importing backup images, Phase I	667
Importing backup images, Phase II	668
Importing expired images	669
Initiating an import without the Import Wizard	669
Importing Backup Exec media	671
Differences between importing, browsing, and restoring Backup Exec and NetBackup images	673
Section 5	
Monitoring and reporting	677
Chapter 20	
Monitoring NetBackup activity	679
Using the Activity Monitor	679
Activity Monitor topology	681
Filtering topology objects	682
About the Jobs tab	682
Viewing job details	684
Setting job detail selections	684
Monitoring the detailed status of a selected job	684
Running the Troubleshooter from within the Activity Monitor	685
Deleting completed jobs	685
Canceling a job that has not completed	685
Restarting a job	686
Suspending restore or backup jobs	686
Resuming suspended jobs	686
Printing job list information	687
Printing job detail information	687
Copying Activity Monitor text to another document	688
Changing the Job Priority dynamically	688
About the Services tab	688
Types of services	692
Other Symantec services	693
Starting or stopping a service	693
Monitoring NetBackup services	693

	About the Processes tab	694
	Monitoring NetBackup processes	698
	About the Drives tab	698
	Monitoring NetBackup tape drives	699
	Cleaning tape drives from the Activity Monitor	700
	About the jobs database	700
	Retaining job information in the database	700
	About the bpdjobs debug log	703
	About the Device Monitor	704
	About media mount errors	704
	About pending requests and actions	705
	About pending requests for storage units	705
	Managing pending requests and actions	706
	Resolving a pending request	706
	Resolving a pending action	707
	Resubmitting a request	708
	Denying a request	708
Chapter 21	Reporting in NetBackup	709
	About the Reports utility	709
	About the Reports window	710
	About the Reports shortcut menus	711
	Reports settings	712
	Report types	712
	Running a report	715
	Running the Troubleshooter within reports	715
	Copying report text to another document	715
	Saving or exporting a report	716
	Printing a report	716
Section 6	Administering NetBackup	717
Chapter 22	Management topics	719
	NetBackup naming conventions	719
	Wildcards in NetBackup	720
	How to administer devices on other servers	721
	How to access media and devices on other hosts	722
	Example SERVER entries	723
	About the Enterprise Media Manager	723
	Enterprise Media Manager domain requirements	724
	Sharing an EMM server	724

Chapter 23	Accessing a remote server	727
	Accessing remote servers	727
	Adding a NetBackup server to a server list	728
	Adding a server to a remote server list	729
	Choosing a remote server to administer	731
	Using the change server command to administer a remote server	732
	Indicating a remote system upon login	733
	Using the Remote Administration Console	734
	Using the Java Windows Administration Console	735
	Running the Administration Console on a NetBackup client	736
	Troubleshooting remote server administration	736
Chapter 24	Using the NetBackup-Java administration console	739
	Using the NetBackup-Java administration console	739
	Authorizing NetBackup-Java users	742
	Authorization file (auth.conf) characteristics	743
	Authorizing nonroot users for specific applications	745
	Authorizing specific tasks in jbpSA	746
	Authorizing NetBackup-Java users on Windows	747
	Restricting access to NetBackup-Java applications on Windows	748
	Runtime configuration options	748
	BPJAVA_PORT, VNETD_PORT	748
	FIREWALL_IN	749
	FORCE_IPADDR_LOOKUP	750
	INITIAL_MEMORY, MAX_MEMORY	752
	MEM_USE_WARNING	753
	NBJAVA_CLIENT_PORT_WINDOW	753
	NBJAVA_CONNECT_OPTION	754
	NBJAVA_CORBA_DEFAULT_TIMEOUT	754
	NBJAVA_CORBA_LONG_TIMEOUT	754
	How to log the command lines that the NetBackup interfaces use	755
	How to customize jnbSA and jbpSA with bp.conf entries	755
	How to improve NetBackup-Java performance	755
	Running the Java console locally on a UNIX platform	756
	Running the console locally on a Windows platform	756
	How to run a console locally and administer a remote server	757
	How to enhance console performance	757

	Determining better performance when run locally or using remote display back	758
	Scenario 1	759
	Scenario 2	759
	Adjusting time zones in the NetBackup-Java console	760
	Adjusting the time zone	760
	Configuring a custom time zone	761
Chapter 25	Alternate server restores	763
	About alternate server restores	763
	Supported configurations for alternate server restores	764
	Performing alternate server restores	765
	Modifying the NetBackup catalogs	766
	Overriding the original server for restores	767
	Enabling automatic failover to an alternate server	769
	Expiring and importing media for alternate server restores	770
Chapter 26	Managing client restores	771
	Server-directed restores	771
	Client-redirected restores	772
	Restore restrictions	772
	To allow all clients to perform redirected restores	773
	To allow a single client to perform redirected restores	774
	To allow redirected restores of a client's files	774
	Examples of redirected restores	775
	Restoring files and access control lists	779
	Restoring the files that possess ACLs	779
	Restoring files without restoring ACLs	779
	How to improve search times by creating an image list	780
	How to restore System State	781
	Restoring the system state	781
Chapter 27	Powering down and rebooting NetBackup servers	785
	Powering down and rebooting NetBackup servers	785
	Shutting down all NetBackup services	786
	Starting up all NetBackup services	786
	Rebooting a NetBackup server	786
	Rebooting a NetBackup media server	786
	About displaying robotic processes with vmps	787

Chapter 28	About Granular Recovery Technology	789
	About installing and configuring Network File System (NFS) for Active Directory Granular Recovery	789
	About configuring Services for Network File System (NFS) on the Windows 2008 and Windows 2008 R2 NetBackup media server and NetBackup clients	790
	Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2	791
	Disabling the Client for NFS on the media server	794
	Disabling the Server for NFS	795
	About configuring Services for Network File System (NFS) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients	797
	Installing Services for NFS on the Windows 2003 R2 SP2 media server	798
	Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2	801
	Configuring a UNIX or Linux media server and Windows clients for backups and restores that use Granular Recovery Technology	804
	Configuring a different network port for NBFSD	804
	Configuring the log on account for the NetBackup Client Service	805
	Index	807

About NetBackup

- [Chapter 1. Introducing NetBackup interfaces](#)
- [Chapter 2. Administering NetBackup licenses](#)

Introducing NetBackup interfaces

This chapter includes the following topics:

- [About NetBackup](#)
- [Online documents](#)
- [Using the NetBackup Administration Console](#)
- [Configuring NetBackup by using the NetBackup Administration Console](#)
- [NetBackup configuration wizards](#)
- [Activity Monitor utility](#)
- [NetBackup Management utilities](#)
- [Media and Device Management utilities](#)
- [Access Management utility](#)

About NetBackup

NetBackup provides a complete, flexible data protection solution for a variety of platforms. The platforms include Microsoft Windows, UNIX, Linux, and NetWare systems.

NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. An administrator can carefully schedule backups to achieve systematic and complete backups over a period of time, and optimize network traffic during off-peak hours.

The backups can be full or incremental. Full backups back up all client files. Incremental backups back up only the files that have changed since the last backup.


The NetBackup administrator can allow users to back up, restore, or archive the files from their computer. (An archive operation backs up a file, then deletes it from the local disk if the backup is successful.)

NetBackup includes both the server and the client software as follows:

- Server software resides on the computer that manages the storage devices.
- Client software resides on computer(s) that contain data to back up. (Servers also contain client software and can be backed up.)

Figure 1-1 shows an example of a NetBackup storage domain.

Figure 1-1 NetBackup storage domain example



NetBackup accommodates multiple servers that work together under the administrative control of one NetBackup master server in the following ways:

- The master server manages backups, archives, and restores. The master server is responsible for media and device selection for NetBackup. Typically, the master server contains the NetBackup catalog. The catalog contains the internal

databases that contain information about NetBackup backups and configuration.

- Media servers provide additional storage by allowing NetBackup to use the storage devices that are attached to them. Media servers can also increase performance by distributing the network load. Media servers can also be referred to by using the following terms:
 - Device hosts (when tape devices are present)
 - Storage servers (when performing I/O directly to disk)
 - Data movers (when sending data to independent, external disk devices like OpenStorage appliances)

During a backup or archive, the client sends backup data across the network to a NetBackup server. The NetBackup server manages the type of storage that is specified in the backup policy.

During a restore, users can browse, then select the files and directories to recover. NetBackup finds the selected files and directories and restores them to the disk on the client.

Online documents

NetBackup documents are delivered on a documentation CD that is included with the NetBackup media kit. Contact your NetBackup administrator to obtain the location of this CD or to have the files installed on your computer.

These online documents are in Adobe® Portable Document Format (PDF). To view PDF documents, you must use the Adobe Acrobat Reader. You can download the reader from:

<http://www.adobe.com>

Symantec assumes no responsibility for the installation and use of the reader.

For a complete list of NetBackup technical documents, see the Related Documents appendix in the *NetBackup Release Notes*.

Using the NetBackup Administration Console

The NetBackup Administration Console provides a Windows-based interface through which the administrator can manage NetBackup.

Figure 1-2 NetBackup Administration Console

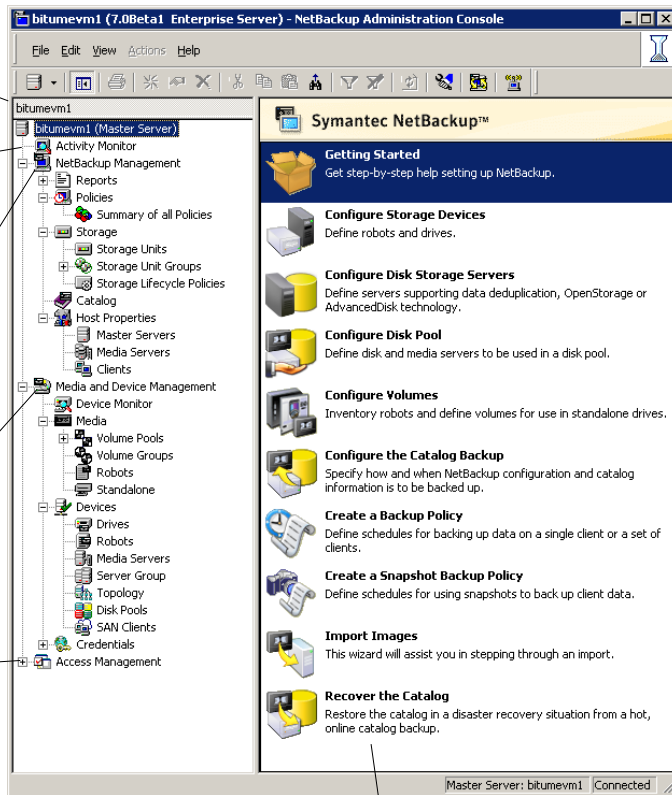
Master server
The information in the NetBackup Administration Console applies to this server only.

Activity Monitor
Displays the NetBackup job information. Provides control over the jobs, services, processes, and drives.

NetBackup Management
Contains the utilities to create and view reports, to configure policies, storage units, catalog backups, and a utility for configuring host properties.

Media and Device Management
Contains the utilities for managing the media and devices that NetBackup uses to store backups.

Access Management
Use to define user groups and grant permissions to these groups. The contents are viewable only by a Security Administrator when NetBackup access control is configured.



Additional licensed utilities
The nodes of other licensed utilities appear under the main NetBackup nodes.

Details pane
Contains the configuration wizards and details specific to the utility that is selected.

Command prompts are used to perform some operations. NetBackup commands are described in *NetBackup Commands*.

The NetBackup Administration Console menus are described in the online Help.

Standard and user toolbars

Upon opening the NetBackup Administration Console, a standard toolbar appears by default.

When certain utilities are selected, a user toolbar appears. The buttons on the toolbar provide shortcuts for menu commands. Slowly drag the pointer over a button to display a button description label.

To display or hide the standard NetBackup toolbar, click **View > Toolbar**.

Customizing the administration console

The **View** menu contains options to customize the NetBackup Administration Console.

For example, the **Options** selection opens a series of tabs that contains various configuration options for the different utilities.

Select the **Administration Console** tab to configure the **Auto log off timeout** option. Use this option of automatically log a user out of the NetBackup Administration Console after a period of inactivity.

Click the **Help** button for more information about the dialog box options.

Configuring NetBackup by using the NetBackup Administration Console

The easiest way to configure NetBackup is to use the configuration wizards. Choose the Getting Started Wizard to configure NetBackup for the first time.

This wizard calls up the following wizards:

- Device Configuration Wizard
- Volume Configuration Wizard
- Catalog Backup Wizard
- Backup Policy Wizard

The wizards help you process configure the basic properties of a NetBackup environment. After you complete these wizards, your NetBackup environment should back up clients and back up your NetBackup catalog.

To configure more advanced properties, you can use the NetBackup Administration Console. You also can use the Administration Console if you prefer not to use the wizards. Alternatively, you can use the wizards to configure the basic properties and then use the Administration Console to configure the more advanced properties.

NetBackup configuration wizards

The easiest way to configure NetBackup is to use the Configuration Wizards. The wizard selection varies in the **Details** pane on the right, depending on what NetBackup utility is selected in the left portion of the screen.

- **Getting Started Wizard**

Use the Getting Started Wizard to configure NetBackup for the first time. The wizard leads the user through the necessary steps to a working NetBackup configuration.

The Getting Started Wizard is comprised of the following wizards, which can also be run separately, outside of the Getting Started Wizard:

- **Device Configuration Wizard**
- **Volume Configuration Wizard**
- **Catalog Recovery Wizard**
- **Backup Policy and Configuration Wizard**

- **Device Configuration Wizard**

Use the Device Configuration Wizard to configure NetBackup to use shared drives or to reconfigure an existing shared drive.

- **Storage Server Configuration Wizard**

Use the Storage Server Configuration Wizard to create the servers that manage disk storage. You can create storage servers for AdvancedDisk, for media server deduplication, for OpenStorage, and for PureDisk storage pools. The wizard appears if an Enterprise Disk Option license or NetBackup Deduplication Option license is installed.

- **Disk Pool Configuration Wizard**

Use the Disk Pool Configuration Wizard to create pools of disk volumes for backup by one or more media servers. For AdvancedDisk, choose disk volumes directly attached to a media server. For OpenStorage, the wizard discovers disk appliances and then you choose disk volumes within the appliance. The wizard appears if an Enterprise Disk Option license or NetBackup Deduplication Option license is installed.

- **Volume Configuration Wizard**

Use the Volume Configuration Wizard to configure removable media to use for backups.

- **Catalog Recovery Wizard**

Use the Catalog Backup Wizard to set up catalog backups. Catalog backups are essential to recover data in the case of a server failure or crash.

- **Backup Policy and Configuration Wizard**

Use the Backup Policy and Configuration Wizard to add a backup policy to the configuration.

- **Import Images Wizard**

Use the Import Images Wizard to import NetBackup images in a two-part process.

- **Catalog Recovery Wizard**

Use the Catalog Recovery Wizard in a disaster recovery situation. Use the Catalog Recovery Wizard only if the NetBackup environment was running the policy-based online, hot catalog backup as the catalog backup type.

Activity Monitor utility

Use the Activity Monitor utility to monitor and control NetBackup jobs, services, processes, and drives.

See [“Using the Activity Monitor”](#) on page 679.

NetBackup Management utilities

The following topics describe the utilities that are found under the **NetBackup Management** node in the NetBackup Administration Console tree:

- **Reports**

Use the **Reports** utility to compile information for to verify, manage, and troubleshoot NetBackup operations.

See [“About the Reports utility”](#) on page 709.

- **Policies**

Use the **Policies** utility to create and specify the backup policies that define the rules for backing up a group of clients.

For example, the backup policy specifies when automatic backups occur for the clients that are specified in the policy. The backup policy also specifies whether users can perform their own backups and when. The administrator can define any number of backup policies, each of which can apply to one or more clients. A NetBackup client must belong to at least one backup policy to be backed up.

See [“Using the Policies utility”](#) on page 448.

- **Storage**

Use the **Storage** utility to display storage unit information and manage NetBackup storage units. A storage unit can be part of a storage unit group as well as part of a storage lifecycle policy, both of which are configured within the **Storage** utility.

Storage units simplify administration because once defined, the NetBackup policy points to a storage unit rather than to the individual devices it contains.

For example, if a storage unit contains two drives and one is busy, NetBackup can use the other drive without administrator intervention.

The media can be one of the following:

- Removable (such as tape in a robot or a stand-alone drive).
The devices in a removable-media storage unit must attach to a NetBackup master or media server and be under control of the NetBackup Media Manager component. The administrator first configures the drives, robots, and media in NetBackup, then defines the storage units. During a backup, NetBackup sends data to the storage unit that the backup policy specifies. During a backup, Media Manager picks a device to which the NetBackup client sends data.

- Disk (such as a file directory within a file system or a collection of disk volumes, either independent file systems or in an appliance).
The administrator specifies the directory, volume, or disk pool during the storage unit setup. For BasicDisk, NetBackup sends the data to that directory during backups. For the Enterprise Disk Options, NetBackup sends the data to the storage server (the host that writes to the storage). Media Manager is not involved.

For disk pool storage, the administrator first defines the storage server and (depending on the disk type) its logon credentials. Depending on disk type, the administrator may have to define logon credentials for the storage itself. The administrator also selects the disk volumes that comprise the disk pool. To create a storage unit, the administrator selects a disk pool and (depending on the disk type) selects the media server(s) to move the data.

Note: Only the storage units that point to shareable disk can specify more than one media server.

See “[About the Storage utility](#)” on page 365.

- **Catalog**

Use the **Catalog** utility to create and configure a catalog backup, which is a special type of backup that NetBackup requires for its own internal databases. These databases, called catalogs, are located on the NetBackup master and media server (default location). The catalogs contain information on every client backup. Catalog backups are tracked separately from other backups to ensure recovery in case of a server crash.

The **Catalog** utility is also used for the following actions:

- To duplicate a backup image

- To promote a backup image from a copy to the primary backup copy
 - To manually expire backup images
 - To import expired backup images or images from another NetBackup server
 - To search for a backup image to verify the contents of the media with what is recorded in the NetBackup catalog
- See [“About the Catalog utility”](#) on page 653.
- **Host Properties**
Use the **Host Properties** utility to customize NetBackup configuration options. In most instances, no changes are necessary. However, **Host Properties** lets the administrator customize NetBackup to meet specific site preferences and requirements for master servers, media servers, and clients.
See [“About the Host Properties”](#) on page 53.

Media and Device Management utilities

The following topics describe the utilities that are found under **Media and Device Management** utilities in the NetBackup Administration Console tree:

- **Device Monitor**
Use the **Device Monitor** utility to manage drives, device paths, and service requests for operators.
- **Media**
Use the **Media** utility to add and manage removable media.
- **Devices**
Use the **Devices** utility to add, configure, and manage storage devices.
- **Credentials**
Use the **Credentials** utility to add, remove, and manage log on credentials for to following:
 - NDMP hosts (requires the NetBackup for NDMP license).
 - Storage servers (requires a NetBackup Deduplication Option or an Enterprise Disk Option license).**Credentials** appears only if one of the previously mentioned license keys is installed.

Access Management utility

NetBackup administrators can protect a NetBackup configuration by defining who may access NetBackup and what functions a user group can perform. This access control is configured by using the **Access Management** utility. **Access Management** is enabled when Symantec Product Authentication and Authorization and NetBackup Access Control (NBAC) is installed and configured.

For installation and configuration information, see Access Management in the *NetBackup Security and Encryption Guide*.

Administering NetBackup licenses

This chapter includes the following topics:

- [Administering NetBackup licenses](#)

Administering NetBackup licenses

License keys are added when the software is installed. Licenses can be added later in the **License Key** dialog box for separately-priced options.

Note: Perform a manual hot catalog backup after updating license keys.

An immediate, manual catalog backup prevents stale keys from being restored in case a catalog restore is necessary before the next scheduled catalog backup.

See [“Backing up catalogs manually”](#) on page 612.

Perform the following tasks from the NetBackup License Keys dialog box:

- Add a new license.
See [“Adding new license keys”](#) on page 45.
- Print a license.
See [“Printing license key lists”](#) on page 45.
- Delete a license.
See [“Deleting license keys”](#) on page 46.
- View the properties of one license.
See [“Viewing license key properties”](#) on page 46.
- Export the license list.

See [“Exporting license keys”](#) on page 46.

Accessing license keys for a NetBackup server

Use the following procedure to access license keys for a NetBackup server.

To access license keys for a NetBackup server

- 1 To view the license keys of the current server:

In the NetBackup Administration Console, click **Help > License Keys**.

To view the license keys of another server:

Click **File > Change Server**, then select another server. Click **Help > License Keys** in the remote server.

- 2 Select the license details to view as follows:

Summary of active licensed features Displays a summary of the active features that are licensed on this server. This view lists each feature and the number of instances of the feature that are licensed.

Summary of active capacity-based licensed features Displays the storage capacity for which the NetBackup environment is licensed and the capacity in use. The summary also notes whether the license is in compliance. The summary does not display the amount of physical storage space.

All capacity values are calculated based on the definition that one terabyte = 1,099,511,627,776 bytes.

The OpenStorage Disk Option, the PureDisk Storage Option, and the Virtual Tape Option do not display all values at this time.

All registered license keys details Displays the details of the license keys that are registered on this server.

The view lists the following:

- Each license key
- The server where the key is registered
- When the key was registered,
- The features that the key provides

- 3 Perform the following tasks from the NetBackup License Keys dialog box:

- Add a new license.
See [“To add new license keys”](#) on page 45.
- Print a license.
See [“To print license key lists”](#) on page 45.

- Delete a license.
See [“To delete license keys”](#) on page 46.
- View the properties of one license.
See [“Viewing license key properties”](#) on page 46.
- Export the license list.
See [“To export license keys”](#) on page 46.

Adding new license keys

Use the following procedure to add new license keys.

To add new license keys

- 1 To add license to the current server:

In the NetBackup Administration Console, click **Help > License Keys**.

To add a license to another server:

Click **File > Change Server**, then select another server. Click **Help > License Keys** in the remote server.

- 2 In the **NetBackup License Keys** dialog box, click the **New** button.
- 3 Enter the license key and click **Add**.
- 4 Perform a manual hot catalog backup after updating license keys.

An immediate, manual catalog backup prevents stale keys from being restored in case a catalog restore is necessary before the next scheduled catalog backup.

See [“Backing up catalogs manually”](#) on page 612.

Printing license key lists

Use the following procedure to print license key lists.

To print license key lists

- 1 In the **NetBackup License Keys** dialog box, select the license key you want to print. If no selection is made, all licenses print.

The printed information includes the following:

- License key
- Name of the host
- Date the key was added
- Name of the product

- Number of instances
 - Name of the feature
 - Whether or not the license is valid
 - Expiration date for the license
- 2 In the **NetBackup License Keys** dialog box, click the **Print** button.
 - 3 Make the print selections and click **OK**.

Deleting license keys

Use the following procedure to delete license keys.

To delete license keys

- 1 Select the license key you want to delete from the license key list. If the key has more than one feature, all the features are listed in the dialog box.
- 2 In the **NetBackup License Keys** dialog box, click the **Delete** button.
- 3 Click **OK** to delete the key and all features that are associated with the key.

If the key appears in the list more than one time, deleting one instance deletes all other instances of the key from the list.

Viewing license key properties

Use the following procedure to view the properties of a license key.

To view the properties of a license key

In the **NetBackup License Keys** dialog box, select one license and click the **Properties** button.

Exporting license keys

Use the following procedure to export license keys.

To export license keys

- 1 In the **NetBackup License Keys** dialog box, click the **Export** button.
- 2 Enter the path and the file name where you want the key properties of all licenses to be exported.
- 3 Click **Save**.

The exported file contains a list of each license key, along with the:

- Name of the host

- Date the license was added
- Name of the product
- Number of instances
- Name of the feature
- Whether or not the license is valid
- Expiration date for the license

Configuring hosts

- [Chapter 3. Configuring Host Properties](#)
- [Chapter 4. Configuring server groups](#)
- [Chapter 5. Configuring host credentials](#)
- [Chapter 6. Managing media servers](#)

Configuring Host Properties

This chapter includes the following topics:

- [Configuring Host Properties and configuration options](#)
- [About the Host Properties](#)
- [Access Control properties](#)
- [Active Directory host properties](#)
- [Authorization properties](#)
- [Backup Exec Tape Reader properties](#)
- [Bandwidth properties](#)
- [Busy File Settings properties](#)
- [Clean-up properties](#)
- [Client Name properties](#)
- [Client Attributes properties](#)
- [Client Settings \(NetWare\) properties](#)
- [Client Settings \(UNIX\) properties](#)
- [Client Settings \(Windows\) properties](#)
- [Credential Access properties](#)
- [Data Classification properties](#)
- [Default Job Priorities properties](#)
- [Distributed application restore mapping properties](#)

- Encryption properties
- Enterprise Vault properties
- Enterprise Vault Hosts properties
- Exchange properties
- Exclude Lists properties
- Fibre Transport properties
- Firewall properties
- General Server properties
- Global Attributes properties
- Logging properties
- Login Banner Configuration properties
- Lotus Notes properties
- Media properties
- NDMP Global Credentials properties
- NetWare Client properties
- Network properties
- Network Settings Properties
- Port Ranges properties
- Restore Failover properties
- Retention Periods properties
- Servers properties
- SharedDisk properties
- SharePoint properties
- Symantec Products properties
- Timeouts properties
- Universal Settings properties
- UNIX Client properties

- [UNIX Server properties](#)
- [VMware backup hosts properties](#)
- [VSP \(Volume Snapshot Provider\) properties](#)
- [Windows Client properties](#)
- [Configuration options not found in the Host Properties](#)

Configuring Host Properties and configuration options

The Host Properties are configuration the options that let an administrator customize NetBackup to meet specific site preferences and requirements. The defaults provide good results in most cases and do not need to be changed.

See [“About the Host Properties”](#) on page 53.

Generally, these options are configured in the NetBackup Administration Console, under **Host Properties**. However, some options cannot be configured by using the NetBackup Administration Console.

See [“Configuration options not found in the Host Properties”](#) on page 194.

See [“About the states of multiple hosts”](#) on page 56.

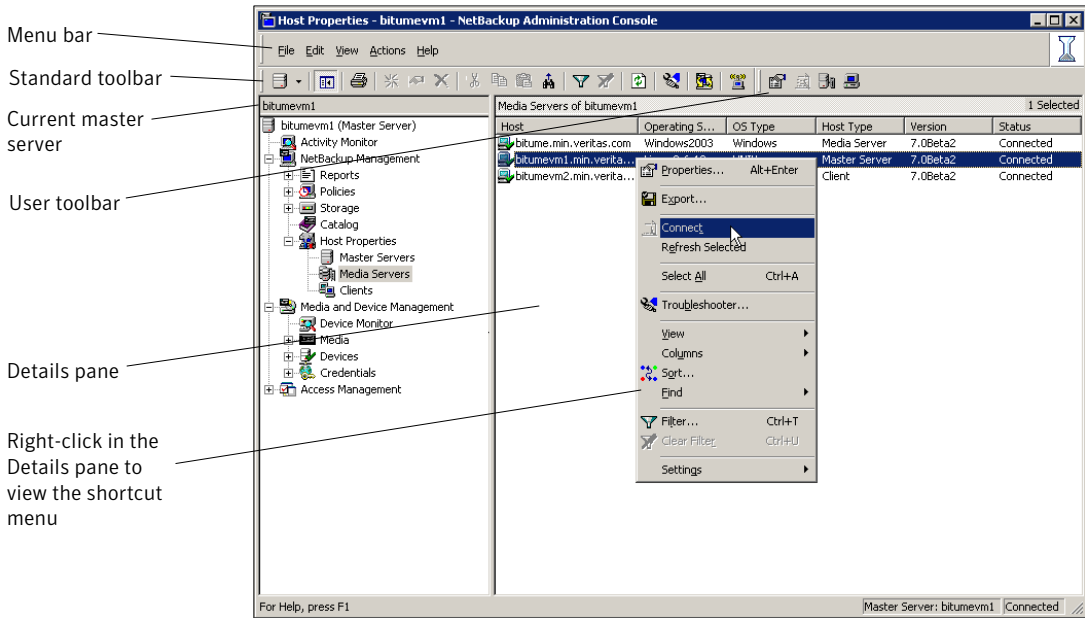
See [“How to change host properties”](#) on page 55.

About the Host Properties

Use the host property dialog boxes in the NetBackup Administration Console to customize NetBackup to meet site preferences. In most instances, however, the NetBackup defaults provide satisfactory results.

[Figure 3-1](#) shows the Host Properties nodes in the **NetBackup Administration Console** tree and the **Details** pane.

Figure 3-1 Host Properties utility



The options on the **Host Properties** menu bar are described in the online Help .

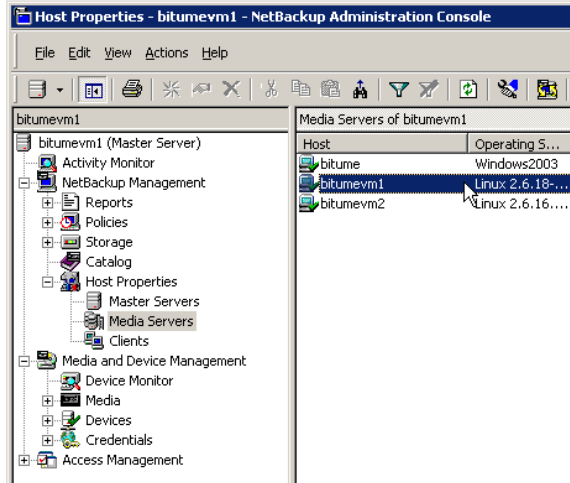
Viewing host properties

The NetBackup Administration Console displays properties for NetBackup master servers, media servers, and clients under Host Properties.

Use the following procedure to view master server, media server, or client properties.

To view master server, media server, or client properties

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties**.



- 2 Select **Master Server, Media Server, or Clients**.
- 3 In the **Details** pane, double-click the server or client to view the properties.

How to change host properties

The NetBackup properties can be changed to customize NetBackup to meet specific site preferences and requirements. In most instances, the NetBackup defaults provide satisfactory results. Host properties can be set for a single host or for multiple hosts all at one time.

One method to change the host properties is to use the NetBackup Administration Console. Another method is to first use the `bpgetconfig` command to obtain a list of configuration entries. Then use `bpsetconfig` to change the entries in the registry.

The commands are described in *NetBackup Commands*.

To change the properties of another client or server, the NetBackup server where you logged on using the NetBackup Administration Console must be in the Servers list on the other system.

See "[Servers properties](#)" on page 175.

For example, if you logged on to server_1 using the NetBackup Administration Console and want to change a setting on client_2, client_2 must include server_1 in its Servers List.

All updates to a destination host (unless it is the same as the host you logged on to using the NetBackup Administration Console) fail if the target host placed a check box in **Allow server file writes** on the Universal Settings properties.

See [“Universal Settings properties”](#) on page 185.

See [“Adding a NetBackup server to a server list”](#) on page 728.

See [“About the states of multiple hosts”](#) on page 56.

See [“Changing properties on multiple hosts”](#) on page 57.

Exporting host properties

Use the following procedure to export the properties of a host.

To export the properties of a host

- 1 Expand **Master Servers**, **Media Servers**, or **Clients**.
- 2 Select one or more hosts.
Click **File > Export**.
- 3 In the **Export Data** dialog box, enter the full path name or browse to the directory and click **Save**.

About the states of multiple hosts

The host properties dialog boxes use the following conventions regarding multiple host selections:

Title of dialog box	If a dialog box contains a Selected Host (or similarly named) box, all controls on the dialog box reflect the values for the host currently selected in the Selected Host box. If a dialog box does not contain a Selected Host (or similarly named) box, settings of all the selected hosts are combined to arrive at a value that is displayed to the user.
Button selection	When multiple hosts are selected, no options appear selected. Selecting any option updates the setting on all selected hosts. To leave each host configured independently, do not select any option while multiple hosts are selected.

Number spinners	When multiple hosts are selected, number spinners appear blank. Selecting any value updates the setting on all selected hosts. To leave each host configured independently, do not select any option while multiple hosts are selected.
Check box states	<p>The host property check boxes may appear in one of the following states:</p> <ul style="list-style-type: none">■ Selected (checked) if the attribute has been set the same for all selected hosts. To set the property on all selected hosts, select the check box.■ Clear (unchecked) if the property has been set the same for all selected hosts. To clear the property on all selected hosts, clear the check box.■ Gray check if the property is set differently on the selected hosts. To leave the property unchanged, set the box to a gray check.
Edit field states	<p>If the property contains a text field for specifying a value, the field may be in one of the following states:</p> <ul style="list-style-type: none">■ The field may contain a value if the property has the same value for all selected hosts.■ The field may be empty or indicate <<Multiple Entries>> if the property was not set the same for all selected hosts. When the cursor is moved to such a field, a small notice appears at the bottom of the dialog box noting that the value is different on the selected hosts.

Note: In a clustered environment, host properties must be made on each node of the cluster separately.

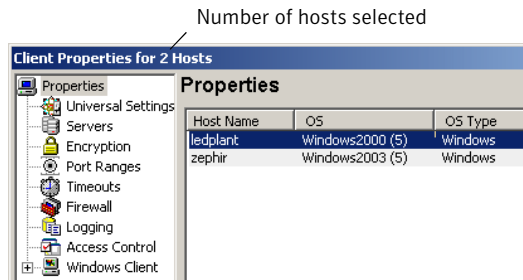
Changing properties on multiple hosts

You can select more than one host and change multiple hosts at one time. Use the following procedure to change properties on multiple hosts at the same time.

To simultaneously change the properties on multiple hosts

- 1 Expand **NetBackup Management > Host Properties > Master Servers, Media Servers, or Clients.**
- 2 Select a host. Hold down the **Shift** key and select another host.
- 3 With multiple hosts still selected, click **Actions > Properties.**

The Properties dialog box displays the names of the selected hosts that are affected by subsequent host property changes.



The following information about each selected host appears:

- Server or client name
- Operating system
- Type of machine in the configuration
- Identifier
- IP address

How to interpret the properties of multiple hosts of different operating systems

If the selected hosts are of various operating systems, none of the operating system-specific information appears.

For example, select a Linux client and a Windows 2008 client. Neither the **Windows Client** properties nor the **UNIX Client** properties appear in the Host Properties tree. If all the selected hosts are of the same operating system, the corresponding properties node appears.

If the property contains a text field for specifying a value, choose from the following options:

- To set the property to the same value for all selected hosts, select the associated check box and type the value in the field.
- To leave the property unchanged, clear the associated check box. The field changes to gray.

At any time you can choose from the following options:

- Click **Default** to set all the fields on the current dialog box to the default values.
- Click **OK** to apply all changes since **Apply** was last clicked. **OK** also closes the dialog box.

- Click **Cancel** to cancel the changes that were made since the last time changes were applied.
- Click **Apply** to save changes to all of the properties for the selected host(s). However, to apply changes click **OK**.
- Click **Help** for information on the properties that appear on the current dialog box.

Access Control properties

Use the Access Control host properties to configure Symantec Authentication and Authorization. The properties apply to currently selected master servers, media servers, and clients.

The following tabs may display:

- See [“About the Symantec Product Authentication and Authorization tab”](#) on page 60.
- See [“About the Authentication Domain tab”](#) on page 62.
- See [“About the Authorization Service tab”](#) on page 65.

The tabs that display depend on whether the host that is selected is a master server, a media server, or a client.

The **Symantec Product Authentication and Authorization** property displays, regardless of which tab is selected. It determines whether the local system uses access control and how the system uses it.

Choose one of the following selections for the **Symantec Product Authentication and Authorization** property:

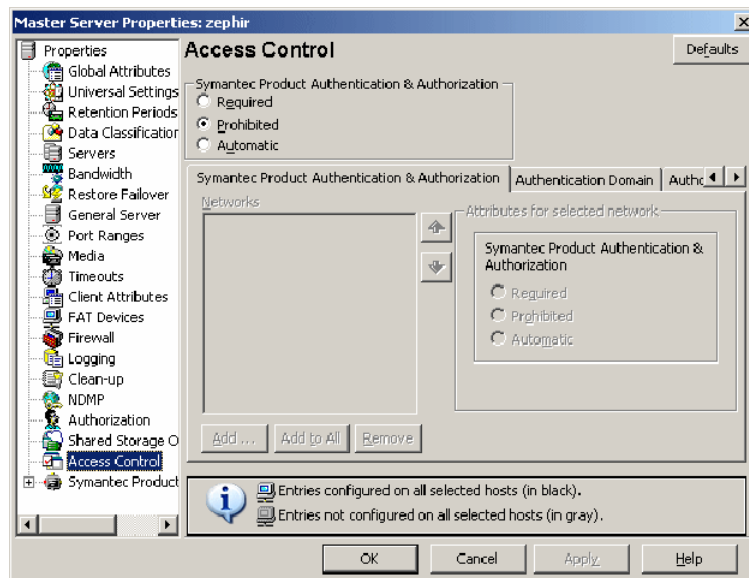
- | | |
|-------------------|--|
| Required | Select if the local system should accept requests only from remote systems that use Symantec authentication and authorization. Connections from remote systems that do not use Symantec authentication and authorization are rejected. Select Required if all systems are at NetBackup 5.0 or later and maximum security is required. |
| Prohibited | Select if the local system should reject connections from any remote system that uses Symantec authentication and authorization. Select Prohibited if the network is closed and maximum performance is required. |
| Automatic | Select if the local system should negotiate with the remote system on whether to use Symantec authentication and authorization. Select Automatic if the network contains mixed versions of NetBackup. |

For more information about controlling access to NetBackup, see the *NetBackup Security and Encryption Guide*.

About the Symantec Product Authentication and Authorization tab

The Symantec Product Authentication and Authorization tab contains a list of networks that are allowed (or not allowed) to use Symantec authentication and authorization with the local system.

Figure 3-2 Access Control dialog box



The Symantec Product Authentication and Authorization tab on the Access Control dialog box contains the following properties:

About the Networks list

This property indicates whether specific networks can or cannot use Symantec authentication and authorization with the local system. The names on the list are relevant only if the **Symantec Product Authentication and Authorization** property in the **Access Control** dialog box is set to **Automatic** or **Required**.

Symantec recommends setting the master server **Symantec Product Authentication and Authorization** property to **Automatic** until the clients are configured for access control. Then, change the **Symantec Product Authentication and Authorization** property on the master server to **Required**.

If a media server or client does not define a Symantec Authentication and Authorization network, it uses the networks of its master server.

To add a network to the **Network** list in the **Access Control** properties, click **Add** on the **Symantec Product Authentication and Authorization** tab to display the **Add Network** dialog box.

The Add Network dialog box contains the following properties:

Host/Domain	Indicate whether the network to be added is a Host name or a Domain name .
Host Details	If the network is a host, enter one of the following items: <ul style="list-style-type: none">■ The host name of the remote system. (host.domain.com)■ The IP address of the remote system. (10.0.0.29)
Domain Details	<ul style="list-style-type: none">■ Domain Name/IP Enter a dot followed by the Internet domain name of the remote systems. (.domain) or the network of the remote system, followed by a dot. (10.0.0.)■ IP Details If the domain is specified by IP, select one of the following items:<ul style="list-style-type: none">■ Select Bit count to indicate that the mask is based on bit count. Select from between 1 and 32. For example: Mask 192.168.10.10/16 has the same meaning as subnet mask 192.168.20.20:255:255:0.0■ Select Subnet mask to enter a subnet mask in the same format as the IP address.
Remove button	To delete a network, select the name, then click Remove .

About the Symantec Product Authentication and Authorization property

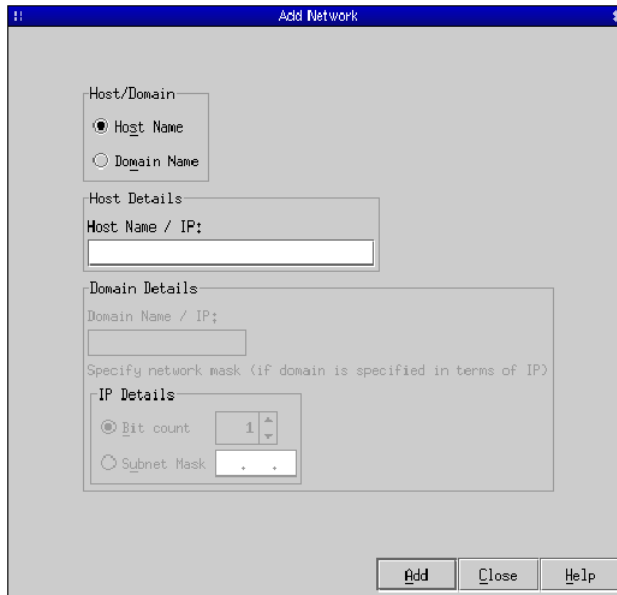
This property determines whether or not the selected network uses Symantec Authentication and Authorization.

Select from one of the following properties:

Required	Select if the local system should accept requests only from remote systems that use Symantec authentication and authorization. Connections from remote systems that do not use Symantec authentication and authorization are rejected. Select Required if all systems are at NetBackup 5.0 or later and maximum security is required.
-----------------	--

- Prohibited** Select if the local system should reject connections from any remote system that uses Symantec authentication and authorization. Select **Prohibited** if the network is closed and maximum performance is required.
- Automatic** Select if the local system should negotiate with the remote system on whether to use Symantec authentication and authorization. Select **Automatic** if the network contains mixed versions of NetBackup.

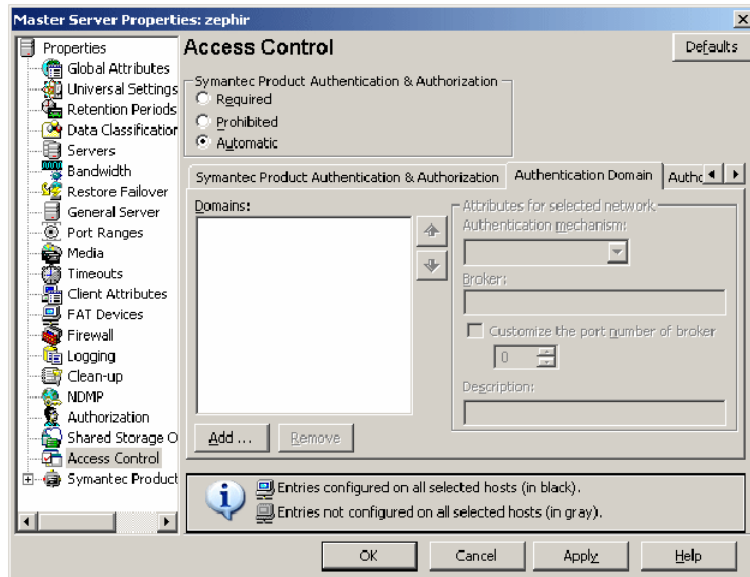
Figure 3-3 Add Network dialog box



About the Authentication Domain tab

The Authentication Domain tab contains the properties that determines which authentication broker a machine uses. A master server that uses Symantec authentication and authorization must have at least one authentication domain entry.

Figure 3-4 Authentication Domain dialog box



The **Authentication Domain** tab on the Access Control dialog box contains the following properties:

Domains list

To add a domain to the **Domains** list in the Access Control properties, click **Add** on the Authentication tab to display the Add Authentication Domain dialog box.

The Add Authentication Domain dialog box contains the following properties:

- **Domain**

Enter an Internet or Windows domain name.

- **Authentication mechanism**

Indicate the authentication mechanism:

- **NIS**: The Network Information Service, version 1.
- **NIS+**: The Network Information Service, version 2.
- **PASSWD**: The local UNIX password file on the specified broker.
- **VXPD**: A VxSS Private Database.
- **WINDOWS**: A Windows Active Directory or Primary Domain Controller.

Note: If the authentication domain is UNIX , enter the fully qualified domain name of the host that performs the authentication.

- **Broker**

The operating system of the broker machine supports the domain type of the authentication service.

Indicate the host name or the IP address of the authentication broker.

- **Customize the port number of service**

Indicate the port number of the authentication broker (optional).

- **Description**

Include a description of the domain (optional).

- **Remove** button

To delete an authorization domain, select the name, then click **Remove**.

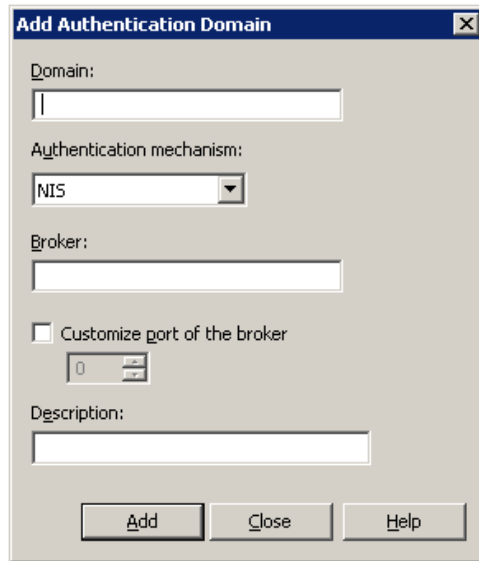
Attributes of the selected network

The attributes of the network selected:

- **Authentication mechanism**
- **Broker**
- **Description**

If a media server or client does not define an authentication domain, it uses the authentication domains of its master server.

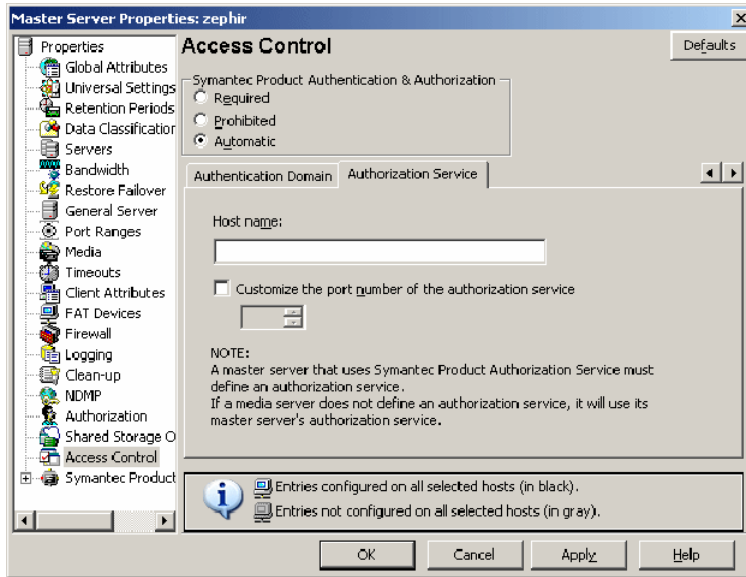
Figure 3-5 Add Authentication Domain dialog box



About the Authorization Service tab

The Authorization Service tab refers to the authorization service that is used by the local NetBackup server. The Authorization Service tab does not appear as a property for clients.

Figure 3-6 Authorization Service dialog box



To configure the Authorization Service for a master server or media server, complete the following options:

- | | |
|---|---|
| Host name | Enter the host name or IP address of the authorization service. |
| Customize the port number of the authorization service | To use a nonstandard port number, select Customize the port number and enter the port number of the authorization service. |

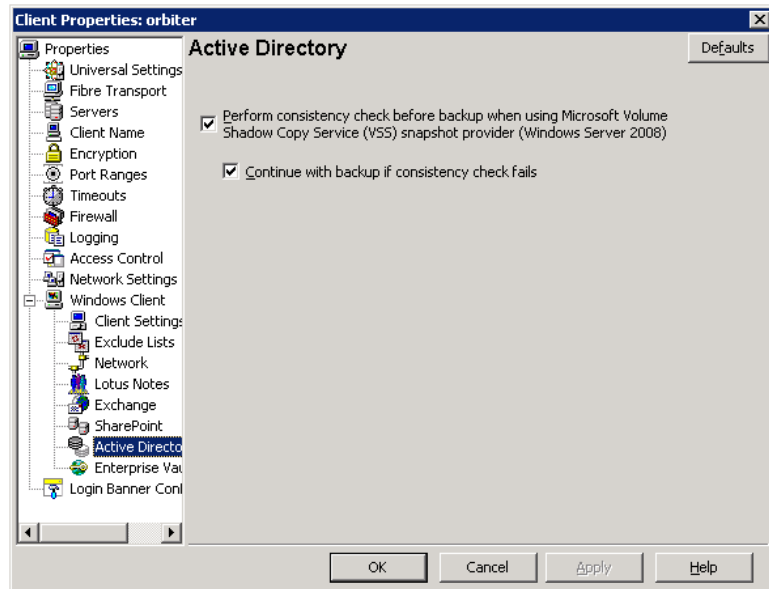
Note: Define a host to perform authorization if you configure this tab for a media server to use access control.

Active Directory host properties

The **Active Directory** properties apply to the backup of currently selected Windows Server 2008 clients. The **Active Directory** properties determine how the backups that allow Active Directory granular restores are performed.

See “[Creating a policy that allows Active Directory granular restores](#)” on page 576.

Figure 3-7 Active Directory dialog box



The following topics describe the Active Directory host properties.

Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider

Select this option to check snapshots for data corruption. This option applies only to snapshots that the Microsoft Volume Shadow Copy Services (VSS) performs.

If corrupt data is found and this option is not selected, the job fails.

See "[Snapshot Provider](#)" on page 87.

Continue with backup if consistency check fails

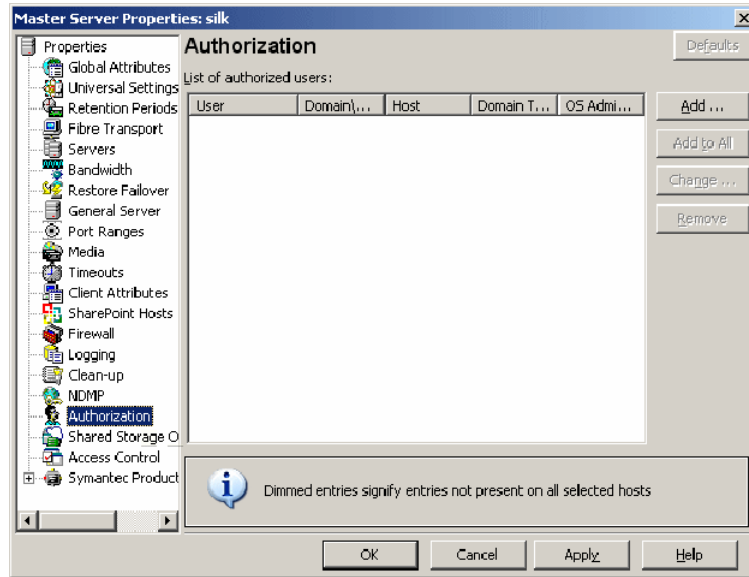
Select this option for the backup job to continue even if the consistency check fails.

It may be preferable for the job to continue, even if the consistency check fails. For example, a backup of the database in its current state may be better than no backup at all. Or, it may be preferable for the backup of a large database to continue if it encounters only a small problem.

Authorization properties

Use the Authorization host properties to configure a list of users that are authorized to use NetBackup. The Authorization host properties apply to currently selected master servers and media servers.

Figure 3-8 Authorization dialog box



Adding or changing an authorized user

Click **Add** to add an authorized user, or select a user in the list and click **Change** to change the configuration of an existing authorized user.

The New User or Change User dialog box contains the following options:

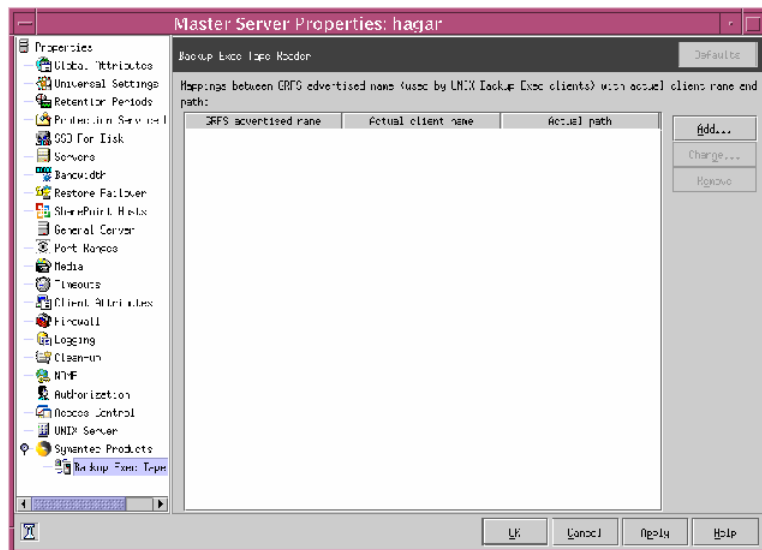
- User** In the **User** text box, type the name that identifies this user to NetBackup. To indicate any user, enter a single asterisk (*).
- Domain\Group** In the **Domain\Group** text box, type the Windows domain or group name. (Use the format **domain\group**.)
Or, type the UNIX local group name or the UNIX netgroup name. To indicate all groups, enter a single asterisk (*).
- Host** In the **Host** text box, type the name of the remote NetBackup Administration Console host from which this user can use NetBackup. To indicate all hosts, enter a single asterisk (*).

- Group/Domain type** Select whether this user is authorized to use NetBackup in a Local group or a Network group.
- User must be an OS administrator** A checkmark in the **User must be an OS administrator** check box indicates that the user must be a system administrator of the host from which they connect.

Backup Exec Tape Reader properties

The Backup Exec Tape Reader properties let NetBackup read the media that Backup Exec writes. Media is read by using a two-phase import process. The Backup Exec Tape Reader properties apply to currently selected master servers.

Figure 3-9 Backup Exec Tape Reader dialog box



In the Backup Exec Tape Reader properties, click **Add** to enter a GRFS mapping.

The **Add a GRFS Mapping** dialog box contains the following options.

See “[Importing backups](#)” on page 666.

GRFS advertised name

The **GRFS advertised name** is the name that the Backup Exec UNIX agent uses to identify itself to the Backup Exec server. The advertised name may not be the same as the real computer name and path.

To set the correct client name and paths in Backup Exec UNIX images .f file paths, map the master server between the GRFS advertised name (generic file system name) and the actual client name and path.

The **GRFS advertised name** uses the following format:

```
ADVERTISED_HOST_NAME/advertised_path
```

where **ADVERTISED_HOST_NAME** is the advertised host name and **advertised_path** is the advertised path. Enter the **ADVERTISED_HOST_NAME** in capital letters.

A Backup Exec service maps the advertised name to the actual computer name and path, and then backs up the advertised name and path. When NetBackup imports Backup Exec UNIX backups, the mapping service is not present; therefore the names and paths must be indicated.

If the host properties do not list any entries, NetBackup assumes that the advertised name is the same as the real computer name. NetBackup assumes that the advertised path is the same as the real path.

To change a selected GRFS entry, select an entry in the Backup Exec Tape Reader properties list and click **Change**.

To remove a GRFS entry, select an entry in the Backup Exec Tape Reader properties list and click **Remove**.

Actual client name

The **Actual client name** maps the advertised name to the real computer name.

If the host properties do not list any entries, NetBackup assumes that the advertised name is the same as the real computer name. NetBackup assumes that the advertised path is the same as the real path.

Actual path

The **Actual path** maps the advertised path to the real path.

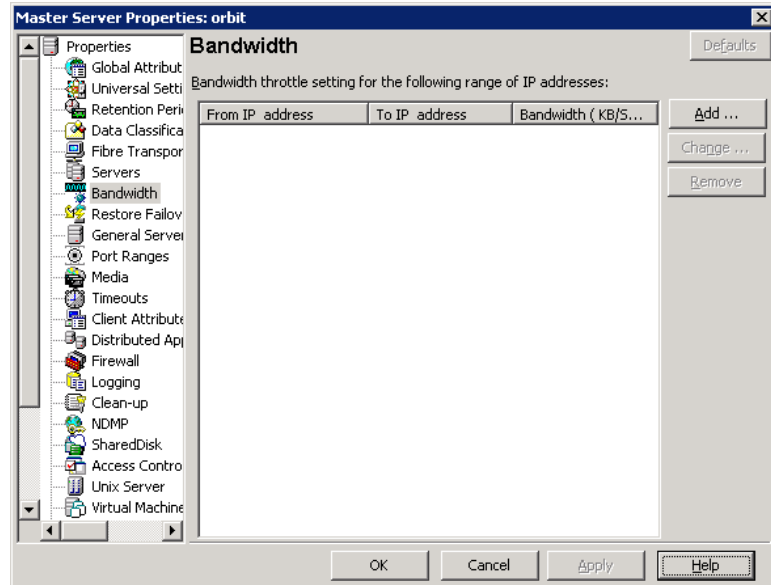
If the host properties do not list any entries, NetBackup assumes that the advertised name is the same as the real computer name. NetBackup assumes that the advertised path is the same as the real path.

Bandwidth properties

Bandwidth properties specify limits for the network bandwidth that one or more NetBackup clients of the selected server use. The actual limiting occurs on the client side of the backup connection. By default, the bandwidth is not limited.

The Bandwidth properties apply to currently selected master servers.

Figure 3-10 Bandwidth dialog box



The bandwidth limits only restrict bandwidth during backups.

Consider the following items about bandwidth limits:

- NetBackup does not currently support bandwidth limits on NetBackup for Microsoft SQL-Server clients
- NetBackup does not currently support bandwidth limits on the following clients:
 - NetBackup for Oracle clients
 - NetBackup for DataTools SQL-BackTrack clients
 - NetBackup for Microsoft SQL-Server clients
- Bandwidth limits have no effect on a local backup (where the server is also a client and data does not go over the network).
- Bandwidth limits restrict maximum network usage and do not imply required bandwidth. For example, if you set the bandwidth limit for a client to 500 kilobytes per second, the client can use up to that limit. It does not mean, however, that the client requires 500 kilobytes per second.

- You cannot use bandwidth limits to distribute the backup workload of active backups by having NetBackup pick the most available network segment. NetBackup does not pick the next client to run based on any configured bandwidth limits.

To add an entry to the bandwidth table, click the **Add** button. An entry is added for each of the selected clients.

Select an entry and click **Remove** to remove a selected entry from the bandwidth table.

The Add or Change Bandwidth Settings dialog box contains the following options.

From IP address field

The **From IP address** field specifies the beginning of the IP address range of the clients and networks to which the entry applies. For example: 10.1.1.2

To IP address field

The **To IP address** field specifies the end of the IP address range of the clients and networks to which the entry applies. For example: 10.1.1.9

Bandwidth

The **Bandwidth** field specifies the bandwidth limitation in kilobytes per second. A value of 0 disables the limits for an individual client or the range of IP addresses covered by the entry.

For example, a value of 200 indicates 200 kilobytes per second.

How bandwidth limiting works

When a backup starts, NetBackup reads the bandwidth limit configuration as configured in the Bandwidth host properties. NetBackup then determines the appropriate bandwidth value and passes it to the client. NetBackup computes the bandwidth limit that is based on the current set of active backups on the subnet and the new backup that starts. Backups that start later are not considered. NetBackup does not include local backups in its calculations.

The NetBackup client software enforces the bandwidth limit. Before a buffer is written to the network, client software calculates the current value for kilobytes per second and adjusts its transfer rate if necessary.

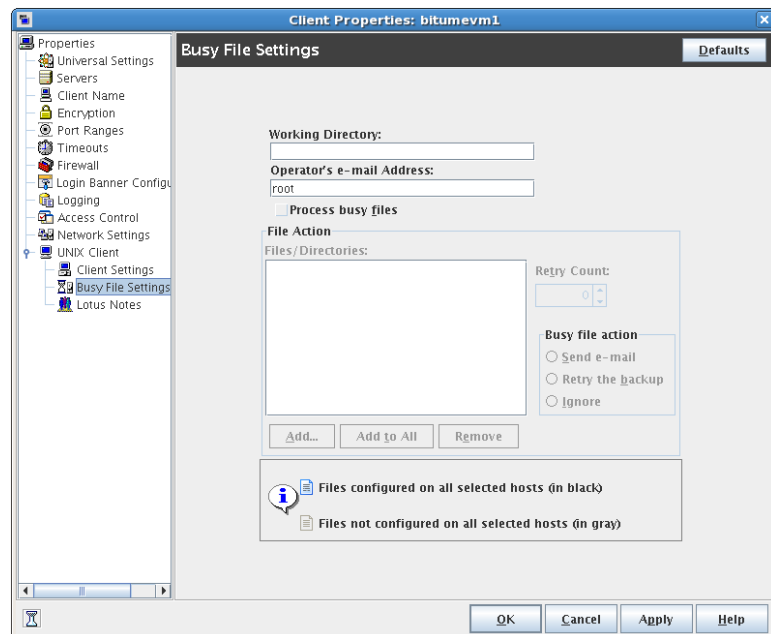
As the number of active backups increase or decrease on a subnet, NetBackup dynamically adjusts the bandwidth limits on that subnet. If additional backups

are started, the NetBackup server instructs the other NetBackup clients that run on that subnet to decrease their bandwidth setting. Similarly, bandwidth per client is increased if the number of clients decreases. Changes to the bandwidth value occur on a periodic basis rather than as backups stop and start. The periodic changes reduce the number of bandwidth value changes that are required.

Busy File Settings properties

The **Busy File Settings** properties apply to currently selected UNIX clients. The Busy File properties define what occurs when NetBackup encounters a busy file during a backup of a UNIX client.

Figure 3-11 Busy File Settings dialog box



The **Busy File Settings** properties contain the following options.

Working directory

The **Working directory** property specifies the path to the busy-files working directory.

On a UNIX client, the value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, NetBackup creates the `busy_files` directory in the `/usr/opensv/netbackup` directory.

Operator's email address

The **Operator's email address** property specifies the recipient of the busy-file notification message when the action is set to **Send email**. By default, the mail recipient is the administrator.

On a UNIX client, the value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BUSY_FILE_NOTIFY_USER` is not in any `bp.conf` file and the mail recipient is `root`.

Process busy files

Enable **Process busy files** for NetBackup to process busy files according to the host property settings. NetBackup follows the Busy File Settings if it determines that a file is changing during a backup. By default, **Process busy files** is not enabled and NetBackup does not process the busy files.

File action file list

The **File action** list specifies the absolute path and file name of the busy file. The metacharacters `*`, `?`, `[]`, `[-]` can be used for pattern matching of file names or parts of file names.

Add

Click **Add** to add a new file entry. Enter the file and path directly, or browse to select a file.

Add to all

Click **Add to all** to add a new file entry for all of the clients currently selected. Enter the file and path directly, or browse to select a file.

Remove

Select the file or directory and click **Remove** to remove the file from the file action list.

Busy file action

The **Busy file action** property directs the action that NetBackup performs on busy files when busy-file processing is enabled. On a UNIX client, the value in the user's \$HOME/bp.conf file takes precedence if it exists.

Select one of the following Busy file options:

1. **Send email**

Directs NetBackup to mail a busy file notification message to the user that is specified in the Operator's email address field in this dialog box.

2. **Retry the backup**

Directs NetBackup to retry the backup on the specified busy file. The Retry count value determines the number of times NetBackup tries a backup.

3. **Ignore**

Directs NetBackup to exclude the busy file from busy file processing. The file is backed up, then a log entry that indicates it was busy appears in the All Log Entries report.

Retry count

The **Retry count** property specifies the number of times to try the backup. The default retry count is 1.

How to activate the Busy File Settings host properties

To activate the settings in the Busy File Settings host properties, complete the following tasks:

- Copy the `bpend_notify_busy` script:

```
/usr/opensv/netbackup/bin/goodies/bpend_notify_busy
```

to the path:

```
/usr/opensv/netbackup/bin/bpend_notify
```

Set the file access permissions to allow group and others to run `bpend_notify`.

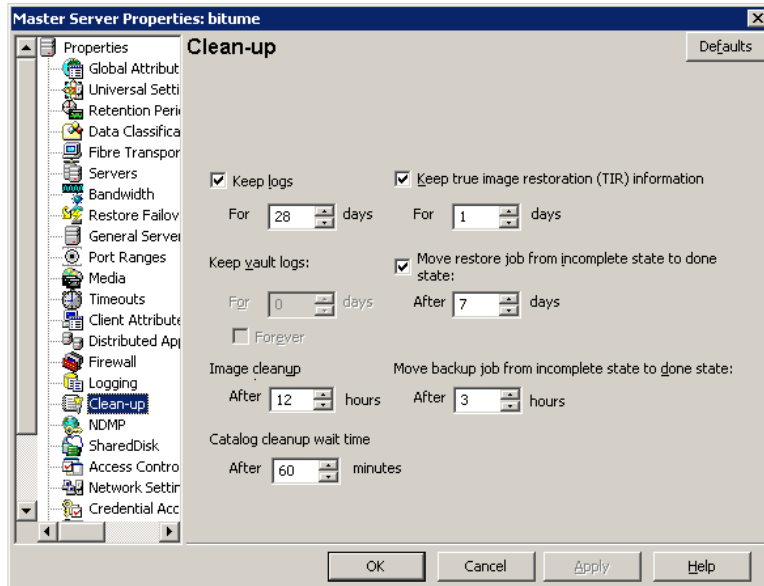
- Configure a policy with a user backup schedule for the busy file backups. This policy services the backup requests that the repeat option in the actions file generates. The policy name is significant. By default, NetBackup alphabetically searches (uppercase characters first) for the first available policy

with a user backup schedule and an open backup window. For example, a policy name of AAA_busy_files is selected ahead of B_policy.

Clean-up properties

The **Clean-up** properties pertain to the retention of various logs and incomplete jobs. The **Clean-up** properties apply to currently selected master servers.

Figure 3-12 Clean-up dialog box



The Clean-Up properties dialog box contains the following options.

Keep logs

The **Keep logs** property specifies the length of time, in days, that the master server keeps its error catalog, job catalog, and debug log information. NetBackup derives the Backup Status, Problems, All Log Entries, and Media Log reports from the error catalog. Also limits the time period that these reports can cover. When this time expires, NetBackup also deletes these logs (that exist) on UNIX media servers and UNIX clients.

Specify how many days you want to keep the logs in case you need the logs to evaluate failures. For example, if you check the backups every day, you can delete the logs sooner than if you check the backups once a month. However, the logs

can consume a large amount of disk space, so do not keep the logs any longer than necessary. The default is 28 days.

Keep vault logs

If Vault is installed, the **Keep vault logs** option is enabled. It specifies the amount of time that the Vault session directories are kept.

Session directories are found in the following location:

```
install_path\netbackup\vault\sessions\vaultname\  
session_x
```

where *x* is the session number. This directory contains vault log files, temporary working files, and report files.

Image cleanup

The **Image cleanup** property specifies the maximum interval that can elapse before an image cleanup is run. Image cleanup is run after every successful backup session (that is, a session in which at least one backup runs successfully). If a backup session exceeds this maximum interval, an image cleanup is initiated.

Catalog cleanup wait time

The **Catalog cleanup wait time** specifies the minimum interval that can elapse before an image cleanup is run. Image cleanup is not run after a successful backup session until this minimum interval has elapsed since the previous image cleanup.

Keep true image restoration (TIR) information

The **Keep true image restoration information** property specifies the number of days to keep true image restore information on disk. After the specified number of days, the images are pruned (removed). The property applies to all policies for which NetBackup collects true image restore information. The default is one day.

When NetBackup performs a true image backup, it stores the following images on the backup media:

- Backed up files
- True image restore information

NetBackup also stores the true image restore information on disk in the *install_path\NetBackup\db\images* directory. NetBackup retains the information for the number of days that this property specifies.

To keep the information on disk speeds up restores. If a user requests a true image restore after the information was deleted from disk, NetBackup retrieves the required information from the media. The only noticeable difference to the user is a slight increase in total restore time. NetBackup deletes the additional information from disk again after one day.

Move restore job from incomplete state to done state

This property indicates the number of days that a failed restore job can remain in an Incomplete state. After that time, the Activity Monitor shows the job as Done.

The default is 7 days. The maximum setting is 365 days.

If Checkpoint Restart for restores is used, the Restore retries property allows a failed restore job to be retried automatically.

See [“Universal Settings properties”](#) on page 185.

See [“Checkpoint restart for restore jobs”](#) on page 470.

Move backup job from incomplete state to done state

This property indicates the maximum number of hours that a failed backup job can remain in an incomplete state. After that time, the Activity Monitor shows the job as Done. The minimum setting is 1 hour. The maximum setting is 72 hours. The default is 3 hours.

When an active job has an error, the job goes into an Incomplete state. In the Incomplete state, the administrator can correct the condition that caused the error. If an Incomplete job does not complete successfully and is moved to the Done state, the job retains the error status.

Note: A resumed job reuses the same job ID, but a restarted job receives a new job ID. The job details indicate that the job was resumed or restarted.

Note: This property does not apply to suspended jobs. Suspended jobs must be resumed manually before the retention period of the job is met and the image expires. If a suspended job is resumed after the retention period is met, the job fails and is moved to the Done state.

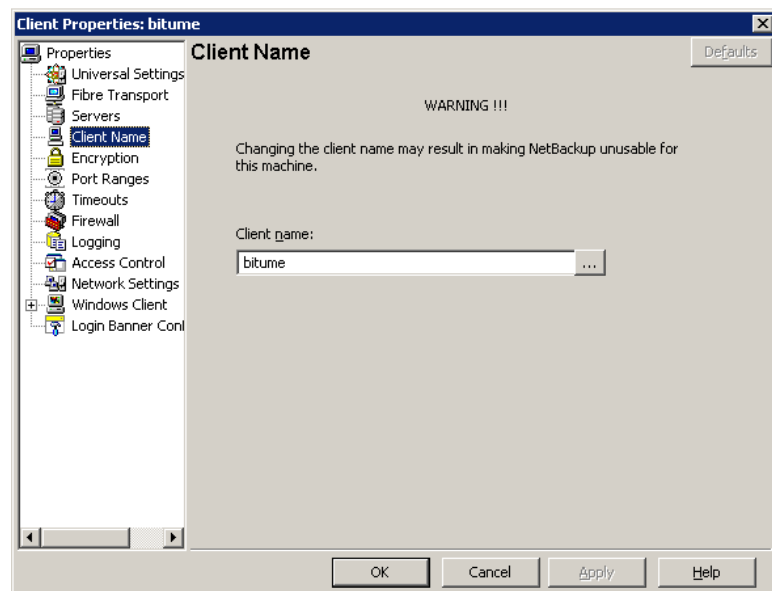
Client Name properties

The host that is specified in the **Client name** properties is the NetBackup client name for the selected client. The **Client name** is the name by which the client is known to NetBackup. The name must match the name the policy uses to back up the client. The only exception is for a redirected restore, where the name must match that of the client whose files are to be restored. The client name is initially set during installation.

The name that is entered here must also match the client name in the **Client Attributes** dialog box for the master server. If it does not, the client cannot browse for its own backups.

See “[Client Attributes properties](#)” on page 80.

Figure 3-13 Client Name dialog box



If the value is not specified, NetBackup uses the name that is set in the following locations:

- For a Windows client
In the Network application from the Control Panel.
- For a UNIX client
The name that is set by using the `hostname` command.

The name can also be added to a `$HOME/bp.conf` file on a UNIX client. However, the name is normally added in this manner only for redirected restores. The value in the `$HOME/bp.conf` file takes precedence if it exists.

Client Attributes properties

Client Attributes properties apply to clients of currently selected master servers.

The **Global client attributes** apply to all clients, unless overridden as described:

Allow client browse	Select this option to allow all clients to browse files for restoring. This attribute is overridden if the Browse and restore ability option on the General tab is set to Deny both for a particular client(s).
Allow client restore	Select this option to allow all clients to restore files. This attribute is overridden if the Browse and restore ability option on the General tab is set to Allow browse only or Deny both .
Clients list	<p>This is a list of clients in the client database on the currently selected master server(s). A client must be in the client database before you can change the client properties in the Client Attributes dialog box.</p> <p>The client database consists of directories and files in the following directory:</p> <p>If a client is not listed in the Clients list, click Add to add clients. To remove a client from the Clients list, select the client, then click Remove.</p> <p>The name that is entered here must match the Client Name property for the specific client. If it does not, the client cannot browse its own backups.</p> <p>See “Client Name properties” on page 79.</p> <p>Use the <code>bpclient</code> command to add clients to the client database if dynamic addressing (DHCP) is in use.</p> <p>Additional information about dynamic host names and IP addressing is available in the <i>NetBackup Administrator’s Guide, Volume II</i>.</p>

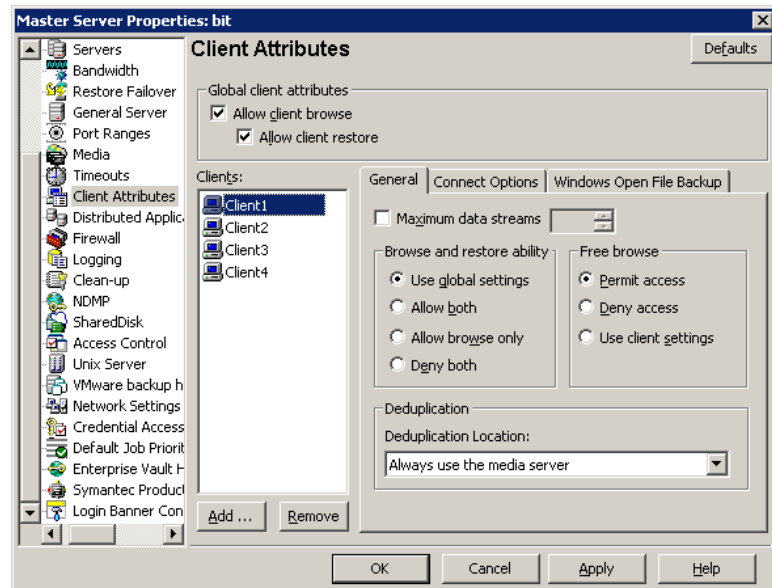
In addition to the **Global client attributes**, the Client Attributes properties contains the following subtabs that can apply to individual clients:

- See “[General tab of the Client Attributes properties](#)” on page 81.
- See “[Connect Options tab of the Client Attributes properties](#)” on page 83.
- See “[Windows Open File Backup tab of the Client Attributes properties](#)” on page 85.

General tab of the Client Attributes properties

The properties on the **General** tab apply to selected Windows master servers. The tab appears on the **Client Attributes** dialog box.

Figure 3-14 General tab of Client Attributes dialog box



The **General** tab contains the following options.

Maximum data streams

This property specifies the maximum number of jobs that are allowed at one time for each selected client. (This value applies to the number of jobs on the client, even if multistreaming is not used.)

To change the setting, select **Maximum data streams**. Then scroll to or enter a value up to 99.

The **Maximum data streams** property interacts with **Maximum jobs per client** and **Limit jobs per policy** as follows:

- If the **Maximum data streams** property is not set, the limit is either the one indicated by the **Maximum jobs per client** property or the **Limit jobs per policy** property, whichever is lower.

- If the **Maximum data streams** property is set, NetBackup ignores the **Maximum jobs per client** property. NetBackup uses either **Maximum data streams** or **Limit jobs per policy**, whichever is lower.
See “[Maximum jobs per client](#)” on page 133.
See “[Limit jobs per policy attribute](#)” on page 471.

Browse and restore ability

This property specifies the client permissions to list and restore backups and archives. Select the client(s) in the **General** tab of the **Client Attributes** dialog box and choose a **Browse and restore ability** property.

To use the **Global client attributes** settings, select **Use global settings**.

- To allow users on the selected clients to both browse and restore, select **Allow both**.
- To allow users on the selected clients to browse but not restore, select **Allow browse only**.
- To prevent users on the selected clients from the ability to browse or restore, select **Deny both**.

Free browse

This property applies to the privileges that are allowed to a non-Windows administrator who is logged into the client. This property also applies to the users that do not have backup and restore privileges.

The **Free browse** property specifies whether the clients can list and restore from scheduled backups. (This setting does not affect user backups and archives.)

Windows administrators can list and restore from scheduled backups as well as user backups regardless of the **Free browse** setting.

Where should deduplication occur

This property specifies the deduplication action for clients if you use one of the following NetBackup deduplication options:

- NetBackup Deduplication Option.
- PureDisk Deduplication Option.

The following are the client direct options:

- To always deduplicate the data on the media server, select **Always use the media server** (the default). Jobs fail if one of the following are true:

- The NetBackup Deduplication Engine on the deduplication storage server is inactive.
- The PureDisk storage pool is inactive.
- To deduplicate data on the client and then send it directly to the storage server, select **Prefer to use client-side deduplication**. NetBackup first determines if the client direct library on the storage server is active. If it is active, the client deduplicates the backup data and sends it directly to the storage server, bypassing media server processing. If it is not active, the client sends the backup data to a deduplication media server. The deduplication media server deduplicates the data.
- To always deduplicate the backup data on the client and then send it directly to the storage server, select **Always use client-side deduplication**. If a job fails, NetBackup does not retry the job.

You can override the **Prefer to use client-side deduplication** or **Always use client-side deduplication** host property in the backup policies.

See “[Disable client-side deduplication attribute](#)” on page 488.


More information about client deduplication is available.

See the *NetBackup Deduplication Guide*.

Connect Options tab of the Client Attributes properties

The properties in the **Connect Options** tab describe how a NetBackup server connects to NetBackup client tabs. The tab appears on the **Client Attributes** dialog box.

Figure 3-15 Connect Options tab of Client Attributes dialog box



The **Connect Options** tab contains the following options.

BPCD connect back

Specify how daemons are to connect back to the NetBackup Client daemon (BPCD):

- **Use default connect options**

Use the value that is defined in the Firewall host properties of the client's NetBackup server.

See [“Default connect options”](#) on page 122.

- **Random port**

NetBackup randomly chooses a free port in the allowed range to perform the legacy connect-back method.

- **VNETD port**

NetBackup uses the `vnetd` port number for the connect-back method.

Ports

Select the method that the selected clients should use to connect to the server:

- **Use default connect options**

Use the value that is defined in the Firewall host properties of the client's NetBackup server.

See “[Default connect options](#)” on page 122.

- **Reserved port**
Use a reserved port number.
- **Non-reserved port**
Use a non-reserved port number.

Daemon connection port

Select the method that the selected clients should use to connect to the server:

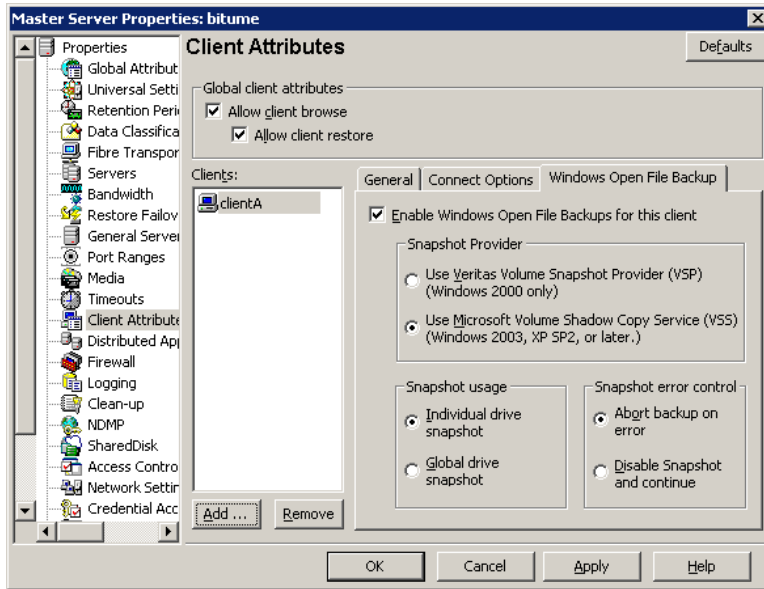
- **Use default connect options**
Use the value that is defined in the Firewall host properties of the client’s NetBackup server.
- **Automatic**
Connect to the daemons on the server using `vnetd` if possible. If the daemons cannot use `vnetd`, the connection is made by using the daemon’s legacy port number.
- **VNETD only**
Connect to the daemons on the server by using only `vnetd`. If the firewall rules prevent a server connection using the legacy port number, check this option. If this option is selected, the **BPCD connect back** setting is not applicable. If this option is selected, the **Ports** setting uses **Non-reserved port**, regardless of the value selected.
- **Daemon port only**
Connect to the daemons on the server by using only the legacy port number.

Windows Open File Backup tab of the Client Attributes properties

The Windows Open File Backup properties specify whether a client uses Windows Open File Backup. The properties also specify whether **Volume Snapshot Provider** or **Volume Shadow Copy Service** is used as the snapshot provider.

Snapshots are a point-in-time view of a source volume. NetBackup uses snapshots to access busy or active files during a backup job. Without a snapshot provider, active files are not accessible for backup.

Figure 3-16 Windows Open File Backup tab of Client Attributes dialog box



The Windows Open File Backup tab contains the following options.

Add button

Click **Add** to add NetBackup clients only if you want to change the Windows Open File Backup defaults.

By default, no clients are listed in the **Client Attributes** dialog box and the server uses the following Windows Open File Backup defaults for all Windows clients:

- Windows Open File Backup is enabled on the client.
- Microsoft Volume Shadow Copy Service (VSS) is used for NetBackup 7.0 clients. See [“Backlevel and upgraded clients that use Windows Open File Backup”](#) on page 87.
- Snapshots are taken of individual drives (**Individual drive snapshot**) as opposed to all drives at once (**Global drive snapshot**).
- Upon error, the snapshot is terminated (**Abort backup on error**).

Remove button

To delete a client from the list, select the client, then click **Delete**.

Enable Windows Open File Backups

This option specifies that Windows Open File Backups be used for the selected clients. Add clients to the list only if you want to change the default property settings.

This option functions independently from the **Perform Snapshot backups** policy option that is available when the Snapshot Client is licensed.

If a client is included in a policy that has the **Perform Snapshot backups** policy option disabled and you do not want snapshots, the **Enable Windows Open File Backups** for this client property must be disabled as well for the client. If both options are not disabled, a snapshot is created, though that may not be the intention of the administrator.

For more information, see the *NetBackup Snapshot Client Administrator's Guide*.

Snapshot Provider

Select the snapshot provider for the selected clients:

■ Use Veritas Volume Snapshot Provider (VSP)

This option specifies that **Veritas VSP** be used as the snapshot provider. VSP is required for Windows 2000 clients and can also be used on 6.x Windows 2003 clients.

■ Use Microsoft Volume Shadow Copy Service (VSS)

This option specifies that **Microsoft VSS** be used to create volume snapshots of volumes and logical drives for the selected clients.

In 7.0, **Microsoft VSS** should be selected for all Windows clients, as VSP is not available. VSS is available for all supported Windows clients, XP SP2 and later. Configure VSS through the Microsoft VSS configuration dialog boxes.

For information about performing Active Directory granular restores when using VSS, see the following topic:

See [“Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider”](#) on page 67.

Backlevel and upgraded clients that use Windows Open File Backup

[Table 3-1](#) shows the expected Open File Backup behavior based on the client version and the Snapshot Provider setting.

Table 3-1 Snapshot Provider behavior for clients in a 7.0 environment

Client version	Snapshot Provider setting	Behavior
6.x	Veritas VSP (6.5 default setting)	Veritas VSP is used for Open File Backup.
6.x	Veritas VSP	Veritas VSP is used for Open File Backup.
6.x	Windows VSS	Windows VSS is used for Open File Backup.
7.0	Windows VSS (7.0 default setting)	Note that using VSS for Open File Backup is a new default behavior in 7.0.
7.0	Veritas VSP	Even if Veritas VSP is indicated, Windows VSS is used for Open File Backup. For upgraded clients: <ul style="list-style-type: none"> ■ For 6.x clients that used VSP and have been upgraded to 7.0: VSP settings are ignored and VSS snapshots are automatically implemented. ■ For 6.x VSS users: You no longer need to create a Client Attribute entry to enable VSS. VSS is the only snapshot provider available to the NetBackup 7.0 Windows client.
7.0	Windows VSS	Windows VSS is used for Open File Backup.

Snapshot usage

Select how snapshots are made for the selected clients:

- **Individual drive snapshot**

Specifies that the snapshot should be of an individual drive (default). When this property is enabled, snapshot creation and file backup are done sequentially on a per volume basis. For example, assume that drives C and D are to be backed up.

If the **Individual drive snapshot** property is selected, NetBackup takes a snapshot of drive C, backs it up, and discards the snapshot. It then takes a snapshot of drive D, backs it up, and discards the snapshot.

Volume snapshots are enabled on only one drive at a time, depending on which drive is to be backed up. This mode is useful when relationships do not have to be maintained between files on the different drives.

Use this configuration if snapshot creation fails when all volumes for the backup are snapshot at once when the **Global drive snapshot** property is enabled. Individual drive snapshot is enabled by default for all non-multistreamed backups by using the Windows Open File Backup option.

- **Global drive snapshot**

The property specifies that the snapshot is of a global drive. All the volumes that require snapshots for the backup job (or stream group for multistreamed backups) are taken at one time.

For example, assume that drives C and D are to be backed up.

In this situation, NetBackup takes a snapshot of C and D. Then NetBackup backs up C and backs up D.

NetBackup then discards the C and D snapshots.

This property maintains file consistency between files in different volumes. The backup uses the same snapshot that is taken at a point in time for all volumes in the backup.

Note: The **Individual drive snapshot** property and the **Global drive snapshot** property only apply to non-multistreamed backups that use Windows Open File Backup. All multistreamed backup jobs share the same volumes snapshots for the volumes in the multistreamed policy. The volume snapshots are taken in a global fashion.

Snapshot error control

Select the processing instructions that NetBackup should follow if it encounters an error during processing:

- **Abort backup on error**

Specifies that a backup aborts if it fails for a snapshot-related issue after the snapshot is created and while the backup is using the snapshot to back up open or active files on the file system.

The most common reason for a problem after the snapshot is created and is in use by a backup, is that the cache storage is full. If the **Abort backup on error** property is checked (default), the backup job aborts with a snapshot error status if the backup detects a snapshot issue.

This property does not apply to successful snapshot creation. The backup job continues regardless of whether a snapshot was successfully created for the backup job. The **Abort backup on error** property applies only to the snapshot

errors that occur after the snapshot is successfully created and is in use by a backup job.

■ **Disable snapshot and continue**

Specifies that if the snapshot becomes invalid during a backup, the volume snapshots for the backup are destroyed. The backup continues with Windows open file backups disabled.

Regarding the file that had a problem during a backup—it may be that the file was not backed up by the backup job. The file may not be able to be restored.

Note: Volume snapshots typically become invalid during the course of a backup because insufficient cache storage was allocated for the volume snapshot. Reconfigure the cache storage configuration of the Windows Open File Backup snapshot provider to a configuration that best suits your client's installation.

Client Settings (NetWare) properties

The Client Settings properties apply to currently selected NetWare clients.

The Client Settings (NetWare) properties dialog box includes the following options.

Back up migrated files

Specifies that the files in secondary storage be moved back to primary storage and backed up. If the property is not selected, only the metadata for the file is backed up and the file is not moved back to primary storage. The metadata is the information still in the primary storage that marks where the file would be. Metadata includes any information that is needed to retrieve the file from secondary storage.

Uncompress files before backing up

The property specifies that compressed files are uncompressed before backing up. Uncompression is useful if the file is restored to a version of NetWare that does not support compression. If the option is not selected (default), the file is backed up in its compressed state.

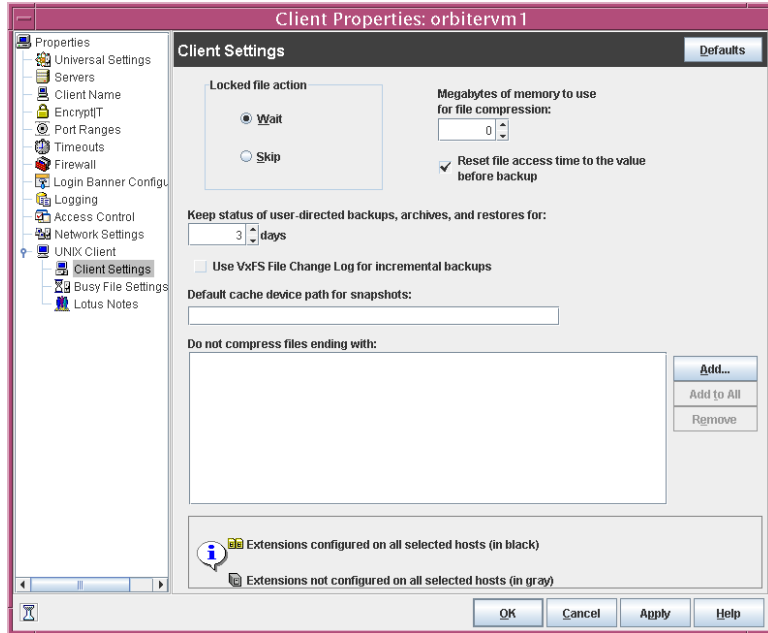
Keep status of user-directed backups, archives, and restores

Specifies how long the system keeps progress reports before it automatically deletes the reports. The default is three days.

Client Settings (UNIX) properties

The UNIX Client Settings apply to currently selected UNIX clients.

Figure 3-17 Client Settings (UNIX) dialog box



The UNIX Client Settings contain the following properties.

Locked file action

Determines what happens when NetBackup tries to back up a file with mandatory file locking enabled in its file mode.

Select one of the following options:

- **Wait**

By default, NetBackup waits for files to become unlocked. If the wait exceeds the **Client read timeout** host property that is configured on the master server, the backup fails with a status 41.

See “[Timeouts properties](#)” on page 182.

- **Skip**

NetBackup skips the files that currently have mandatory locking set by another process. A message is logged if it was necessary to skip a file.

Keep status of user-directed backups, archives, and restores

Specifies the number of days to keep progress reports before the reports are deleted. The default is three days. The minimum is 0. The maximum is 9,999 days.

Logs for user-directed operations are stored on the client system in the following directory:

```
install_path\NetBackup\logs\user_ops\ loginID\logs
```

Reset file access time

Specifies that the access time (`atime`) time for a file displays the backup time. By default, NetBackup preserves the access time by resetting it to the value it had before the backup.

Note: This setting affects the software and the administration scripts that examine a file's access time.

Megabytes of memory to use for file compression

Specifies the amount of memory available on the client when files are compressed during backup. If you select compression, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to compress code, the greater the compression and the greater the percentage of machine resources that are used. If other processes also need memory, use a maximum value of half the actual physical memory on a machine to avoid excessive swapping.

The default is 0. This default is reasonable; change it only if problems are encountered.

Use VxFS file change log for incremental backups

Determines if NetBackup uses the File Change Log on VxFS clients.

The default is off.

See [“Using the VxFS file change log for incremental backups property”](#) on page 93.

Default cache device path for snapshots

For additional information, see the *NetBackup Snapshot Client Administrator's Guide*.

Do not compress files ending with list

The **Do not compress files ending with** list specifies a list of file extensions. During a backup, NetBackup does not compress files with these extensions because the file may already be in a compressed format.

Do not use wildcards to specify these extensions. For example, `.A1` is allowed, but not `.A*` or `.A[1-9]`

Files that are already compressed become slightly larger if compressed again. If compressed files with a unique file extension already exist on a UNIX client, exclude it from compression by adding it to this list.

Add button

Use the **Add** button to add file endings to the list of file endings that you do not want to compress. Click **Add**, then type the file extension in the **File Endings** dialog box. Use commas or spaces to separate file endings if more than one is added. Click **Add** to add the ending to the list, then click **Close** to close the dialog box.

Add to All button

Use the **Add to All** button to add a file extension that you do not want to compress, to the lists of all clients. To add the file extension to the lists of all clients, select it in the list on the Client Settings host property, then click **Add to All**.

Remove button

Click the **Remove** button to remove a file extension from the list. To remove a name, either type it in the box or click the browse button (...) and select a file ending. Use commas or spaces to separate names.

Using the VxFS file change log for incremental backups property

The **Use VxFS file change log** feature is supported on all platforms and versions where VxFS file systems support FCL.

The following VxFS file systems support FCL:

- Solaris SPARC platform running VxFS 4.1 or greater
- AIX running VxFS 5.0 or greater.
- HP 11.23 running VxFS 5.0 or greater.
- Linux running VxFS 4.1 or greater

The File Change Log (FCL) tracks changes to files and directories in a file system. Changes can include files created, links and unlinks, files renamed, data that is appended, data that is overwritten, data that is truncated, extended attribute modifications, holes punched, and file property updates.

NetBackup can use the FCL to determine which files to select for incremental backups, which can potentially save unnecessary file system processing time. The FCL information that is stored on each client includes the backup type, the FCL offset, and the timestamp for each backup.

The advantages of this property depend largely on the number of file system changes relative to the file system size. The performance impact of incremental backups ranges from many times faster or slower, depending on file system size and use patterns.

For example, enable this property for a client on a very large file system that experiences relatively few changes. The incremental backups for the client may complete sooner since the policy needs to read only the FCL to determine what needs to be backed up on the client.

If a file experiences many changes or multiple changes to many files, the time saving benefit may not be as great.

See [“About the Backup Selections tab”](#) on page 537.

The following items must be in place for the **Use VxFS file change log** feature to work:

- Enable the **Use VxFS file change log** property for every client that wants NetBackup to take advantage of the FCL.
- Enable the FCL on the VxFS client.
See the *Veritas File System Administrator's Guide* for information about how to enable the FCL on the VxFS client.
- Enable the **Use VxFS file change log** property on the client(s) in time for the first full backup. Subsequent incremental backups need this full backup to stay synchronized.
- Specify the VxFS mount point in the policy backup selections list in some manner:
 - By specifying ALL_LOCAL_DRIVES.
 - By specifying the actual VxFS mount point.
 - By specifying a directory at a higher level than the VxFS mount point, provided that **Cross mount points** is enabled.
See [“Cross mount points attribute”](#) on page 478.

If the policy has **Collect true image restore information** or **Collect true image restore information with move detection** enabled, it ignores the **Use VxFS file change log** property on the client.

Additional information about the VxFS file change log feature:

- Activity Monitor messages

The Activity Monitor displays any messages that note when the file change log is used during a backup as follows:

```
Using VxFS File Change Log for backup of pathname
```

The Activity Monitor also notes when full and incremental backups are not synchronized.

- Keeping the data files synchronized with the FCL

The data files must be in sync with the FCL for this property to work. To keep the data files synchronized with the FCL, do not turn the FCL on the VxFS client off and on.

Note: If NetBackup encounters any errors as it processes the FCL, it switches to the normal files system scan. If this switch occurs, it appears in the Activity Monitor.

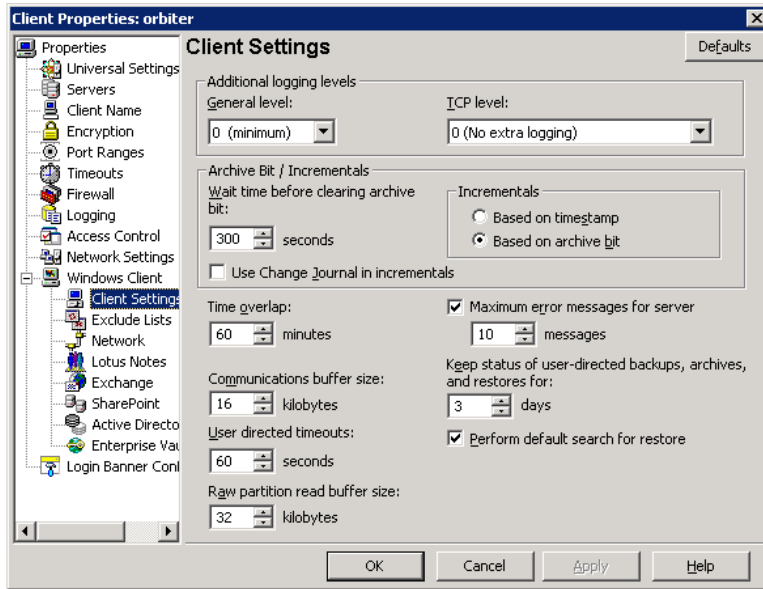
- VxFS administration

Additional VxFS commands are available to administrate the FCL in the *Veritas File System Administrator's Guide*.

Client Settings (Windows) properties

The Windows Client properties apply to currently selected Windows clients.

Figure 3-18 Client Settings (Windows) dialog box



The Client Settings (Windows) contain the following properties.

General level

Enables logs for `bpnetd`, `bbpkar`, `tar`, and `nbwin`. The higher the level, the more information is written. The default is 0.

TCP level

Enables logs for TCP.

Scroll to one of the following available log levels:

- **0 No extra logging (default)**
- **1 Log basic TCP/IP functions**
- **2 Log all TCP/IP functions, including all read and write requests**
- **3 Log contents of each read or write buffer**

Note: Setting the TCP level to 2 or 3 can cause the status reports to be very large. It can also slow a backup or restore operation.

Wait time before clearing archive bit

Specifies how long the client waits before the archive bits for a differential incremental backup are cleared. The minimum allowable value is 300 (default). The client waits for acknowledgment from the server that the backup was successful. If the server does not reply within this time period, the archive bits are not cleared.

This option applies only to differential-incremental backups. Cumulative-incremental backups do not clear the archive bit.

Use change journal in incrementals

NetBackup offers support for the Microsoft change journal to enhance performance of incremental backups on supported Windows OS levels. By enabling the **Use change journal in incrementals** check box, NetBackup can provide faster incremental backups for NTFS 5 (and later) volumes that store large numbers—possibly millions—of files. **Use change journal in incrementals** is available only when a valid tracker database exists on the applicable volumes. The default is not enabled.

When this property is enabled, it automatically enables the **Incrementals are based on timestamp** property.

The Microsoft change journal is a disk file that records and retains the most recent changes to an NTFS volume. By monitoring the change journal, NetBackup can determine which file system objects have changed and when. This information is used to shorten the discovery process that NetBackup performs during an incremental backup by making a file system scan unnecessary.

See [“How to determine if change journal support is useful in your NetBackup environment”](#) on page 99.

See [“Guidelines for enabling NetBackup change journal support”](#) on page 100.

Incrementals based on timestamp

Specifies that files are selected for the backups that are based on the date that the file was last modified. When **Use change journal in incrementals** is selected, **Incrementals based on timestamp** is automatically selected.

Incrementals based on archive bit

Specifies that NetBackup include files in an incremental backup only if the archive bit of the file is set. The system sets this bit whenever a file is changed and it normally remains set until NetBackup clears it.

A full backup always clears the archive bit. A differential-incremental backup clears the archive bit if the file is successfully backed up. The differential-incremental backup must occur within the number of seconds that the **Wait time before clearing archive bit** property indicates. A cumulative-incremental or user backup has no effect on the archive bit.

Disable this property to include a file in an incremental backup only if the datetime stamp for the file has changed since the last backup. For a differential-incremental backup, NetBackup compares the datetime stamp to the last full or incremental backup. For a cumulative-incremental backup, NetBackup compares the timestamp to the last full backup.

If you install or copy files from another computer, the new files retain the date timestamp of the originals. If the original date is before the last backup date on this computer, then the new files are not backed up until the next full backup.

Note: Symantec recommends that you do not combine differential incremental backups and cumulative incremental backups within the same Windows policy when the incremental backups are based on archive bit.

Time overlap

Specifies the number of minutes to add to the date range for incremental backups when using date-based backups. This value compensates for differences in the speed of the clock between the NetBackup client and server. The default is 60 minutes.

This value is used during incremental backups when using the archive bit and when examining the create time on folders. This comparison is done for archive bit-based backups as well as date-based backups.

Communications buffer size

Specifies the size (in kilobytes) of the TCP and IP buffers used to transfer data between the NetBackup server and client. For example, specify 10 for a buffer size of 10 kilobytes. The minimum allowable value is 2, with no maximum allowable value. The default is 16 kilobytes.

User directed timeouts

Specifies the seconds that are allowed between when a user requests a backup or restore and when the operation begins. The operation fails if it does not begin within this time period.

This property has no minimum value or maximum value. The default is 60 seconds.

Maximum error messages for server

Defines how many times a NetBackup client can send the same error message to a NetBackup server. For example, if the archive bits cannot be reset on a file, this property limits how many times the message appears in the server logs. The default is 10.

Keep status of user-directed backups, archives, and restores

Specifies how many days the system keeps progress reports before NetBackup automatically deletes them. The default is 3 days.

Perform default search for restore

Instructs NetBackup to search the default range of backup images automatically. The backed up folders and files within the range appear whenever a restore window is opened.

Clear the **Perform default search for restore** check box to disable the initial search. With the property disabled, the NetBackup Restore window does not display any files or folders upon opening. The default is that the option is enabled.

How to determine if change journal support is useful in your NetBackup environment

Using NetBackup support for the change journal is beneficial only where the volumes are large and relatively static.

Suitable candidates for enabling NetBackup change journal support are as follows:

- If the NTFS volume contains more than 1,000,000 files and folders and the number of changed objects between incremental backups is small (less than 100,000), the volume is a good candidate for enabling NetBackup change journal support.

Unsuitable candidates for enabling NetBackup change journal support are as follows:

- Support for the change journal is intended to reduce scan times for incremental backups by using the information that is gathered from the change journal on a volume. Therefore, to enable NetBackup change journal support is not recommended if the file system on the volume contains relatively few files and folders. (For example, hundreds of thousands of files and folders.) The normal file system scan is suitable under such conditions.

- If the total number of changes on a volume exceeds from 10% to 20% of the total objects, the volume is not a good candidate for enabling NetBackup change journal support.
- Be aware that virus scanning software can interfere with the use of the change journal. Some real-time virus scanners intercept a file open for read, scan for viruses, then reset the access time. This results in the creation of a change journal entry for every scanned file.

Guidelines for enabling NetBackup change journal support

The following items are guidelines to consider for enabling NetBackup change journal support:

- A NetBackup client using change journal support must belong to only one policy. To use one policy avoids the confusion that multiple backup settings causes. Multiple backup settings can cause conflicted update sequence number (USN) information in the permanent record.
- After **Use change journal in incrementals** is selected, restart the NetBackup client service on the target system. A full backup of the target system must be completed under change journal monitoring to enable change journal-based incremental backups.
- Change journal support is not offered for user-directed backups. The USN stamps for full and incremental backups in the permanent record do not change.
- NetBackup support for change journal works with Checkpoint Restart for restores.
See [“Checkpoint restart for restore jobs”](#) on page 470.
- Support for change journal is not offered with several NetBackup options or Symantec products.

If **Use change journal in incrementals** is enabled, it has no effect while using the following options or products:

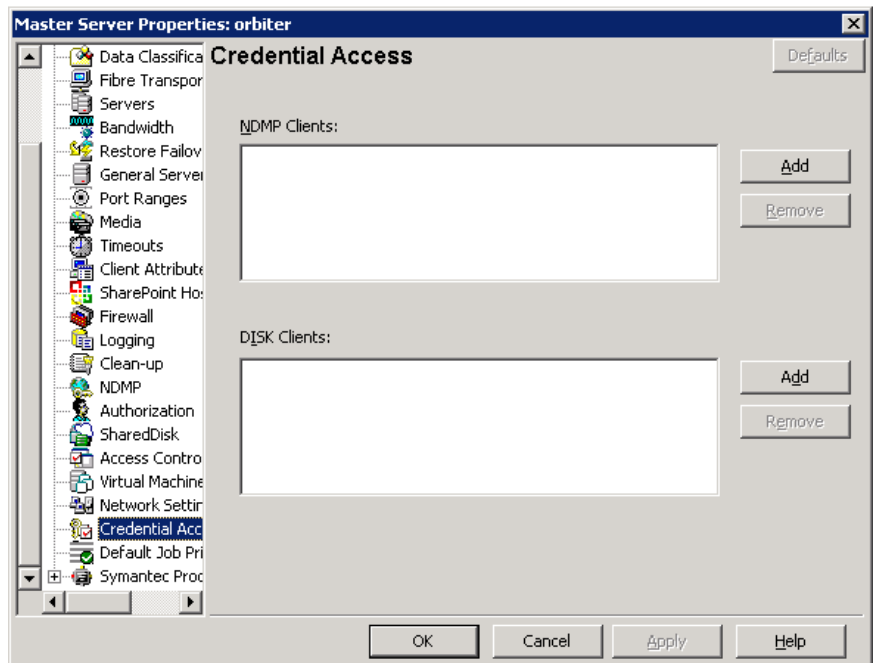
- True image restore (TIR)
See [“Collect true image restore information attribute”](#) on page 483.
- True image restore with Move Detection
See [“Collect true image restore information with move detection attribute”](#) on page 484.
- Synthetic backups
See [“About synthetic backups”](#) on page 583.
- Bare Metal Restore (BMR)

For more information, see the *NetBackup Bare Metal Restore Administrator's Guide*.

Credential Access properties

Certain NetBackup hosts that are not named as clients in a policy must be enabled to access NDMP or disk array credentials. Use the **Credential Access** properties dialog box to enter the names of those NetBackup hosts.

Figure 3-19 Credential Access dialog box



NDMP Clients list

To add an NDMP client to the **NDMP Clients** list, click **Add**. Enter the names of the NDMP hosts that are not named as clients in a policy.

Disk clients list

To add a Disk Client to the **DISK Clients** list, click **Add**. Enter the names of the NetBackup hosts that meet all of the following criteria:

- The host must be designated in a policy as the Off-host backup host in an alternate client backup.
- The host that is designated as the Off-host backup computer must not be named as a client on the **Clients** tab in any NetBackup policy.
- The policy for the off-host backup must be configured to use one of the disk array snapshot methods for the EMC CLARiiON, HP EVA, or IBM disk arrays.

Note: The credentials for the disk array or NDMP host are specified under **Media and Device Management > Credentials**.

Note: Off-host alternate client backup is a feature of NetBackup Snapshot Client, which requires a separate license. The NetBackup for NDMP feature requires the NetBackup for NDMP license.

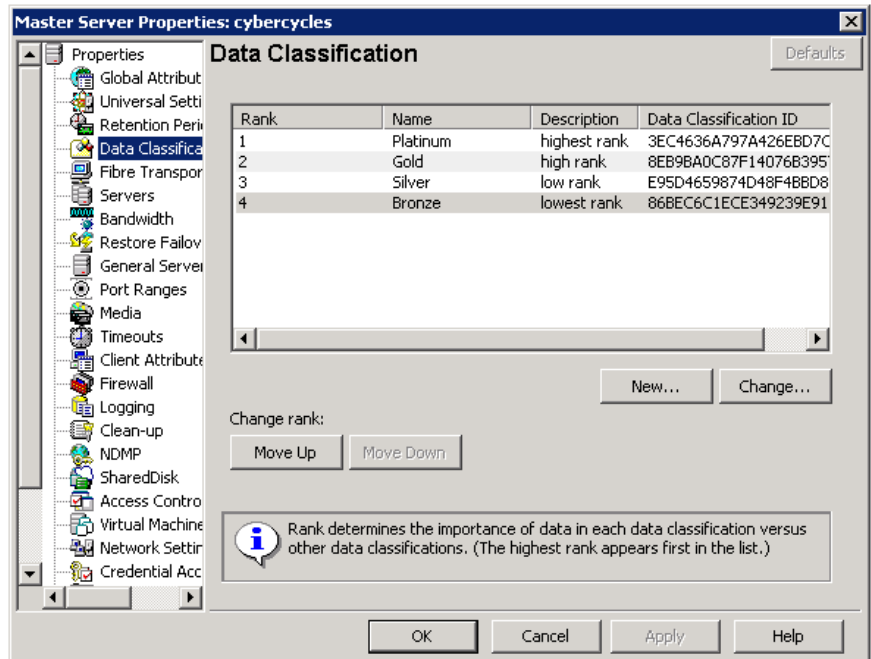
Data Classification properties

The Data Classification properties apply to selected master and media servers.

Data classifications must be configured in the **Data Classification** host properties before storage lifecycle policies can be configured.

See “[Storage lifecycle policy overview](#)” on page 415.

Figure 3-20 Data Classification dialog box



Note: Data classifications cannot be deleted. However, the name, description, and the rank can be changed. The classification ID remains the same.

Creating a Data Classification

Use the following procedures to create or change a data classification.

To create a data classification

- 1 Click **New**.
- 2 Add the name and description in the **New Data Classification** dialog box.
- 3 Click **OK** to save the classification and close the dialog box.

Note: Data classifications cannot be deleted.

- 4 Select a line in the **Data Classification** host properties and use the **Move Up** and **Move Down** buttons to move the classification level up or down in the list.

Rank column

The **Rank** column displays the rank of the data classifications. The order of the data classifications determines the rank of the classification in relationship to the others in the list. The lowest numbered rank has the highest priority.

Use the **Move Up** and **Move Down** buttons to move the classification up or down in the list.

To create a new data classification, click **New**. New data classifications are added to bottom of the list. To increase the rank of a data classification, select a line and click **Move Up**. To decrease the rank of a data classification, select a line and click **Move Down**.

See “[Creating a Data Classification](#)” on page 103.

Name column

The **Name** column displays the data classification name. While data classifications cannot be deleted, the data classification names can be modified.

NetBackup provides the following data classifications by default:

- Platinum (highest rank by default)
- Gold (second highest rank by default)
- Silver (third highest rank by default)
- Bronze (lowest rank by default)

Description column

In the **Description column**, enter a meaningful description for the data classification. Descriptions can be modified.

Data Classification ID

The **Data Classification ID** is the GUID value that identifies the data classification and is generated when a new data classification is added and the host property is saved.

A data classification ID becomes associated with a backup image by setting the Data Classification attribute in the policy dialog box. The ID is written into the image header. The storage lifecycles use the ID to identify the images that are associated with classification.

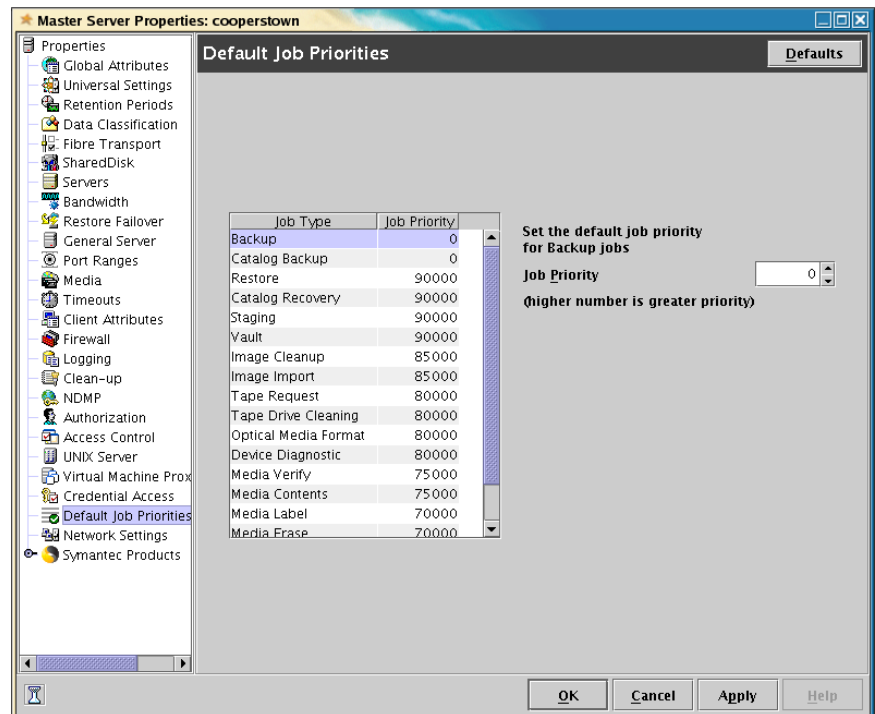
See “[Data classifications attribute](#)” on page 465.

ID values can exist in image headers indefinitely, so data classifications cannot be deleted. The name, description, and rank can change without changing the identity of the data classification.

Default Job Priorities properties

The **Default Job Priorities** host properties let administrators configure the default job priority for different job types. The **Default Job Priorities** host properties list 18 job types and the configurable default priority for each.

Figure 3-21 Default Job Priorities dialog box



The job priority can be set for individual jobs in the following utilities:

- In the **Jobs** tab of the **Activity Monitor** for queued or active jobs.
 See “[Changing the Job Priority dynamically](#)” on page 688.
- In the **Catalog** utility for verify, duplicate, and import jobs.
- In the **Reports** utility for a Media Contents report job.
- In the Backup, Archive, and Restore client interface for restore jobs.

Job Type and Job Priority list

This listing includes 18 job types and the current configurable priority for each.

Job Priority

The **Job Priority** value specifies the priority that a job has as it competes with other jobs for backup resources. The value can range from 0 to 99999. The higher the number, the greater the priority of the job.

A new priority setting affects all the policies that are created after the host property has been changed.

A higher priority does not guarantee that that a job receives resources before a job with a lower priority. NetBackup evaluates jobs with a higher priority before those with a lower priority.

However, the following factors can cause a job with a lower priority to run before a job with a higher priority:

- To maximize drive use, a low priority job may run first if it can use a drive that is currently loaded. A job with a higher priority that requires that the drive be unloaded would wait.
- If a low priority job can join a multiplexed group, it may run first. The job with a higher priority may wait if it is not able to join the multiplexed group.
- If the NetBackup Resource Broker (`nbrb`) receives a job request during an evaluation cycle, it does not consider the job until the next cycle, regardless of the job priority.

See “[Job priority attribute](#)” on page 473.

Understanding the Job Priority setting

NetBackup uses the Job Priority setting as a guide. Requests with a higher priority do not always receive resources before a request with a lower priority.

The NetBackup Resource Broker (NBRB) maintains resource requests for jobs in a queue.

NBRB evaluates the requests sequentially and sorts them based on the following criteria:

- The request's first priority.
- The request's second priority.
- The birth time (when the Resource Broker receives the request).

The first priority is weighted more heavily than the second priority, and the second priority is weighted more heavily than the birth time.

Because a request with a higher priority is listed in the queue before a request with a lower priority, the request with a higher priority is evaluated first. Even though the chances are greater that the higher priority request receives resources first, this is not always definite.

The following scenarios present situations in which a request with a lower priority may receive resources before a request with a higher priority:

- A higher priority job needs to unload the media in a drive because the retention level (or the media pool) of the loaded media is not what the job requires. A lower priority job can use the media that is already loaded in the drive. To maximize drive utilization, the Resource Broker gives the loaded media and drive pair to the job with the lower priority.
- A higher priority job is not eligible to join an existing multiplexing group but a lower priority job is eligible to join the multiplexing group. To continue spinning the drive at the maximum rate, the lower priority job joins the multiplexing group and runs.
- The Resource Broker receives resource requests for jobs and places the requests in a queue before processing. New resource requests are sorted and evaluated every 5 minutes. Some external events (a new resource request or a resource release, for example) can also trigger an evaluation. If the Resource Broker receives a request of any priority while it processes requests in an evaluation cycle, the request is not evaluated until the next evaluation cycle starts.

Distributed application restore mapping properties

Some applications, such as SharePoint and Exchange, distribute and replicate data across multiple hosts. Special configuration is required to allow NetBackup to restore databases to the correct hosts in a SharePoint farm or Exchange Database Availability (DAG) environment. In the **Distributed application restore mapping** properties, add each host in the environment.

The Distributed Application Restore Mapping dialog box contains the following options:

Add Click **Add** to add a host that is authorized to run restores on SharePoint component hosts or Exchange hosts. Provide the name of the **Application host** and the name of the **Component host** in the SharePoint farm or Exchange Database Availability Group (DAG).

Note: For restores to be successful in an Exchange 2010 DAG environment, you must add the CAS server to the list.

Change Click **Change** to change the application host or component host of the currently selected mapping.

Remove Click **Remove** to remove the currently selected mapping..

For more information, see the following:

NetBackup for Microsoft SharePoint Server Administrator's Guide.

NetBackup for Microsoft Exchange Server Administrator's Guide.

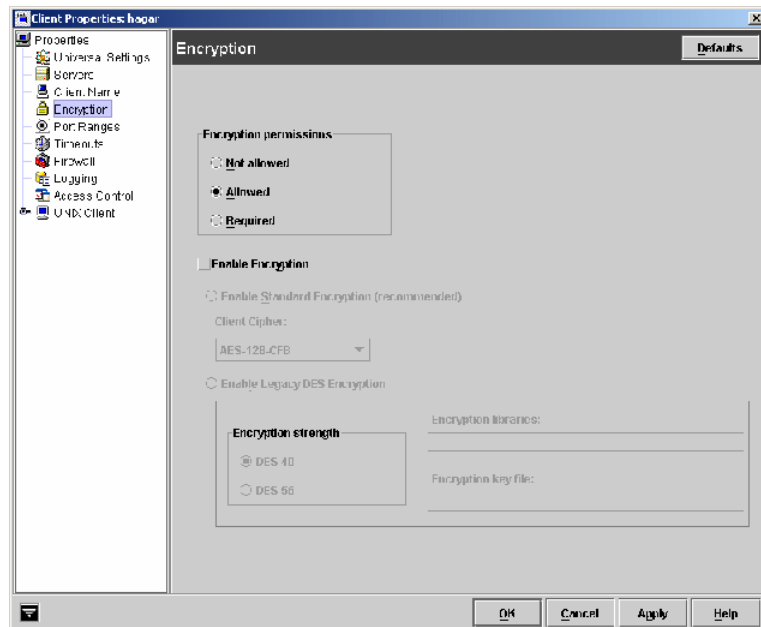
Encryption properties

The **Encryption** properties control encryption on the currently selected client. Multiple clients can be selected and configured at one time only if all selected clients are running the same version of NetBackup. If not, the Encryption properties dialog box is hidden.

The separately-priced NetBackup Encryption option must be installed on the client for these settings (other than Allowed) to take effect.

More information is available in the *NetBackup Security and Encryption Guide*.

Figure 3-22 Encryption dialog box



The **Encryption** dialog box contains the following options.

Encryption permissions property

Indicates the encryption setting on the selected NetBackup client as determined by the master server.

If it is necessary to change this property, select one of the following options:

- **Not allowed**
Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, the backup job ends due to error.
- **Allowed**
Specifies that the client allows either encrypted or unencrypted backups. Allowed is the default setting for a client that has not been configured for encryption.
- **Required**
Specifies that the client requires encrypted backups. If the server requests an unencrypted backup, the backup job ends due to error.

Enable encryption property

Select the **Enable encryption** property if the NetBackup Encryption option is used on the selected client.

After **Enable Encryption** is selected, choose from the following options:

- **Enable standard encryption**
Pertains to the 128-bit and the 256-bit options of NetBackup Encryption. If the selected client does not use Legacy encryption, **Enable standard encryption** is automatically selected.
- **Client Cipher**
The following cipher types are available: BF-CFB, DES-EDE-CFB, AES-256-CFB, and AES-128-CFB. AES-128-CFB is the default. More information on the ciphers file is found in the **NetBackup Security and Encryption Guide**.
- **Enable legacy DES encryption**
Pertains to the 40-bit and the 56-bit data encryption standard (DES) NetBackup encryption packages.
Encryption strength
Defines the encryption strength on the NetBackup client when Legacy encryption is used:
 - DES_40

Specifies the 40-bit DES encryption. DES_40 is the default value for a client that has not been configured for encryption.

- DES_56
Specifies the 56-bit DES encryption.

- **Encryption libraries**

In the **Encryption libraries** field, specify the folder that contains the encryption libraries on NetBackup clients.

The default location is as follows:

- On Windows systems

```
install_path\netbackup\bin\
```

Where *install_path* is the directory where NetBackup is installed and by default is C:\Program Files\VERITAS.

- On UNIX systems

```
/usr/opensv/lib
```

If it is necessary to change the setting, specify the new name.

- **Encryption key file**

In the **Encryption key file** field, specify the file that contains the encryption keys on NetBackup clients.

The default location is as follows:

- On Windows systems

```
install_path\NetBackup\bin\keyfile.dat
```

Where *install_path* is the folder where NetBackup is installed and by default is C:\Program Files\VERITAS.

- On UNIX systems

```
/usr/opensv/netbackup/keyfile
```

If it is necessary to change the setting, specify the new name.

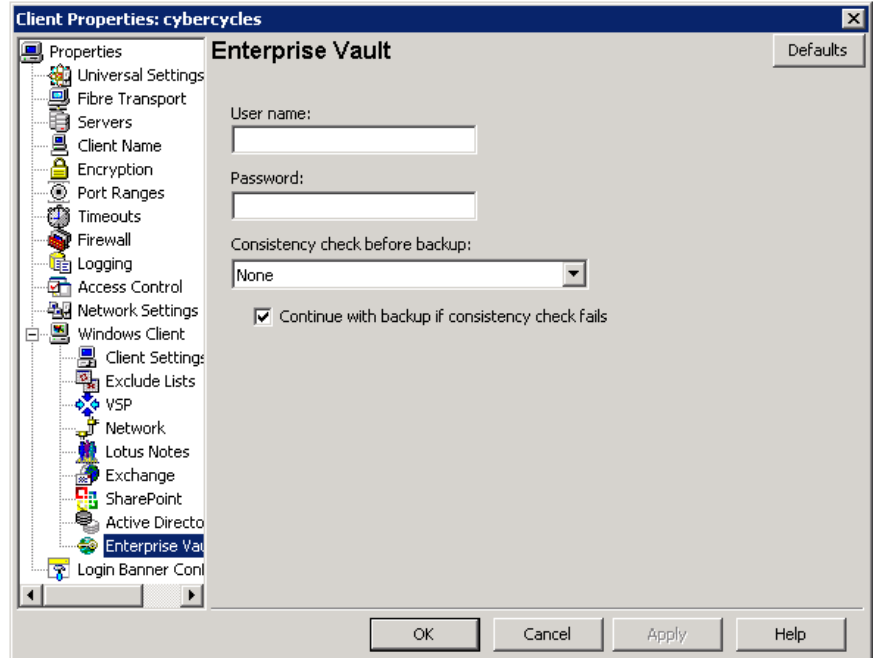
Enterprise Vault properties

The Enterprise Vault properties apply to currently selected clients.

To perform backups and restores, NetBackup must know the user name and password for the account that is used to log on to the Enterprise Vault Server and

to interact with the Enterprise Vault SQL database. The user must set the logon account for every NetBackup client that runs backup and restore operations for Enterprise Vault components.

Figure 3-23 Enterprise Vault client properties



User Name

Specify the user ID for the account that is used to log on to Enterprise Vault (DOMAIN\user name).

Password

Specify the password for the account.

Consistency check before backup

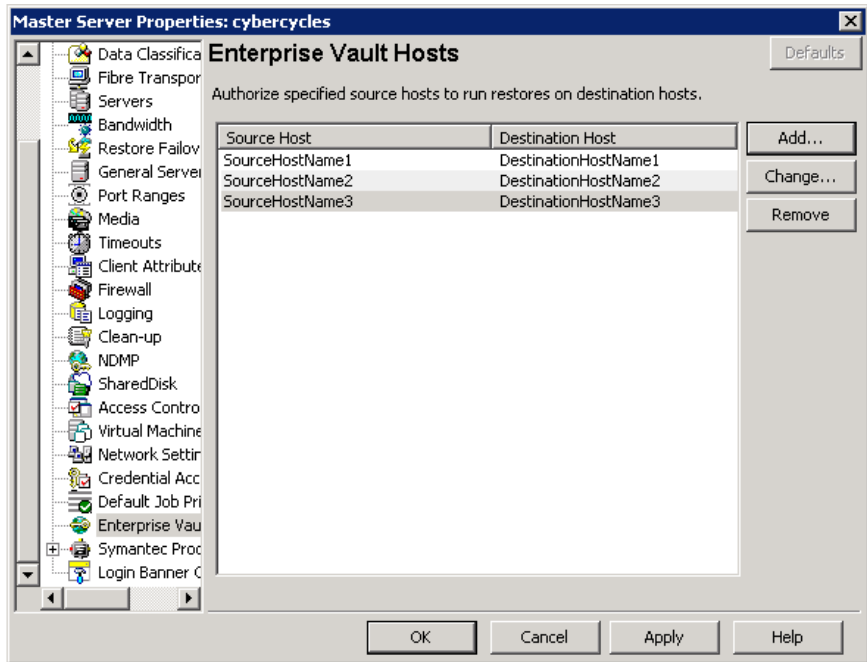
Select what kind of consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation.

Enterprise Vault Hosts properties

The Enterprise Vault Hosts properties apply to currently selected master servers.

Special configuration is required to allow NetBackup to restore SQL databases to the correct hosts in an Enterprise Vault farm. In the Enterprise Vault Hosts master server properties, specify a source and a destination host. By doing so, you specify a source host that can run restores on the destination host.

Figure 3-24 Enterprise Vault Hosts master server properties



Click **Add** to add the source and the destination hosts within the Enterprise Vault configuration. Provide the name of the **Source host** and the name of the **Destination host**.

Click **Change** to change the source host and the destination host, an entry that you select from the Enterprise Vault Hosts field.

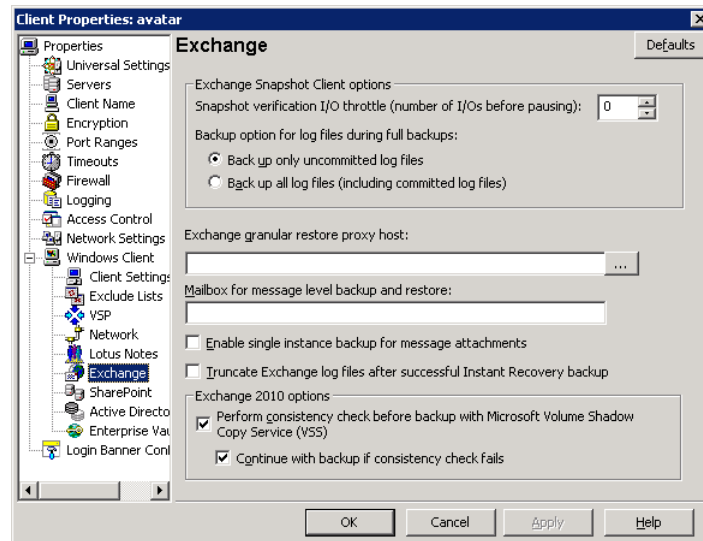
Exchange properties

The Exchange properties apply to the currently selected Windows clients. For clustered or replicated environments, configure the same settings for all nodes.

If you change the attributes for the virtual server name, only the active node is updated.

For complete information on these options, see the *NetBackup for Microsoft Exchange Server Administrator's Guide*.

Figure 3-25 Exchange dialog box



Snapshot verification I/O throttle For snapshot backups, specify the number of I/Os to process for each 1-second pause. This option applies to Exchange 2003 SP2 and to Exchange 2007 if the Exchange Management Console is not installed on the alternate client.

Backup option for log files during full backups Choose which logs to include with snapshot backups:

- **Back up only uncommitted log files**
 Select this option to back up only the log files that are uncommitted. This option is not recommended for Exchange 2010 DAG or Exchange 2007 CCR environments.
- **Back up all log files (including committed log files)**

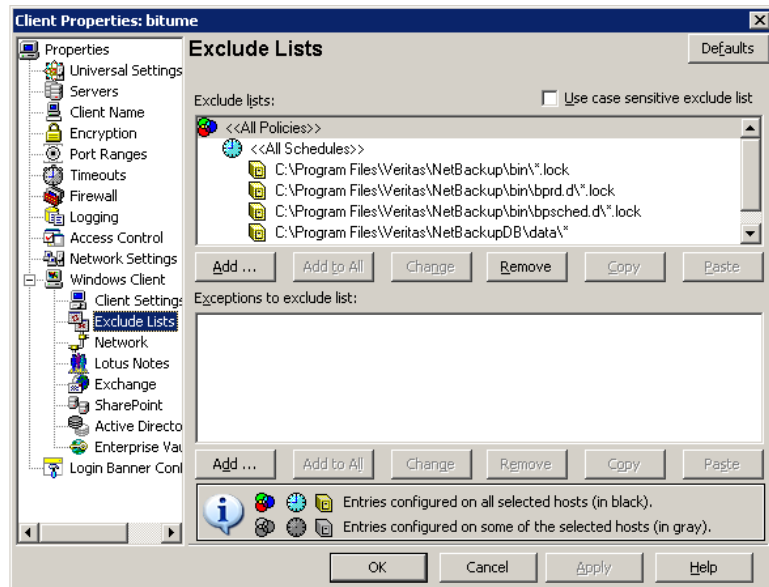
Note: In NetBackup 7.0, the default option is now **Backup all log files (including committed log files)**. If you previously changed this setting for a client, your selection remains the same. For new installations of NetBackup, the default is **Backup all log files (including committed log files)**. For upgrade installations where you did not change this setting for a client, the default is changed to **Backup all log files (including committed log files)**

Truncate log after successful Instant Recovery backup	Enable this option to delete transaction logs after a successful Instant Recovery backup. By default, transaction logs are not deleted for a full Instant Recovery backup that is snapshot only.
Exchange granular restore proxy host	You can specify a different Windows system to act as a proxy for the source client. Use a proxy if you do not want to affect the source client or if it is not available. This situation applies when you duplicate a GRT-enabled backup image from a disk storage unit to a tape storage unit or when you use the <code>bplist</code> command.
Mailbox for message level backup and restore	As of NetBackup 7.0, this setting no longer needs to be configured.
Enable single instance backup for message attachments	Enable this option to back up the data that is stored on a Single Instance Store (SIS) volume. This feature only applies to Exchange Server 2007 and earlier versions.
Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS)	Disable this option if you do not want to perform a consistency check during an Exchange 2010 DAG backup. If you check Continue with backup if consistency check fails , NetBackup continues to perform the backup even if the consistency check fails.

Exclude Lists properties

Use the Exclude Lists properties to create and to modify the exclude lists for Windows clients. An exclude list names policies, schedules, files, and the directories to be excluded from automatic backups.

Figure 3-26 Exclude Lists dialog box



Exclude Lists properties apply only to Windows clients. On NetWare target clients, specify the exclude list when the targets are added. NetWare NonTarget clients do not support exclude lists. For more information, see the NetBackup user's guide for the client.

If more than one exclude or include list exists for a client, NetBackup uses only the most specific one.

For example, assume that a client has the following exclude lists:

- An exclude list for a policy and schedule.
- An exclude list for a policy.
- An exclude list for the entire client. This list does not specify a policy or schedule.

In this example, NetBackup uses the first exclude list (for policy and schedule) because it is the most specific.

The following topics describe the Exclude Lists host properties.

Use case sensitive exclude list property

The **Use case sensitive exclude list** property indicates that the files and directories to exclude are case sensitive.

Exclude list

The **Exclude list** displays the policies that contain schedule, file, and directory exclusions.

Add Click **Add** to exclude a file from being backed up by a policy. The exclusion is configured in the **Add to exclude list** dialog box, then added to the Exclude list.

When the policies on the Exclude list run, the files and directories that are specified on the list are backed up.

See [Figure 3-27](#) on page 116.

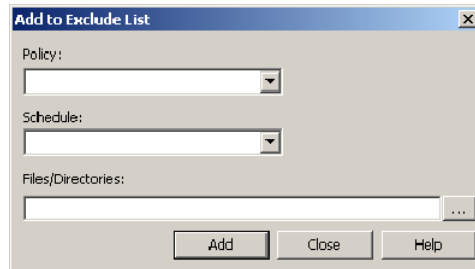
Add to all **Add to all** is enabled only under the following conditions:

- More than one client is selected for configuration and,
- A list item is selected that was not configured on the selected hosts. (Rather, an unavailable list item is selected.)

Click **Add to All** to add the selected list item to all currently selected clients. The item is excluded from the backup list on all selected clients.

Remove Click **Remove** to remove the selected policy, schedule, or file from the Exclude list. The item is included in the backup.

Figure 3-27 Add to Exclude List dialog box



Exceptions to exclude list

The **Exceptions to the exclude list** displays policies, schedules, files, and the directories that are excepted from the Exclude list.

When the policies on the **Exceptions to the exclude list** run, the files and directories on the list are backed up. The list is useful to exclude all files in a directory but one.

Add	<p>Click to create an exception to the Exclude list. The exception is configured in the Add exceptions to exclude list dialog box, then added to the Exceptions to the exclude list.</p> <p>When the policies on the Exceptions to the exclude list run, the items on the exceptions list are backed up. Effectively, you add files back into the backup list of a policy.</p>
Add to all	<p>Click Add to All to add the selected list item to the Exceptions to the exclude list of all currently selected clients. When the policies on the exclude list run, the items on the exceptions list are backed up on all selected clients.</p>
Remove	<p>Click Remove to remove the selected policy, schedule, or file from the Exceptions list. The item is excluded from the backup.</p>

About the Add to exclude list and Add to exceptions list dialog boxes

The **Add to Exclude List** dialog box and the **Add Exceptions to Exclude List** dialog box contain the following fields:

Policy	<p>Enter the policy name that contains the files and the directories to exclude or make exceptions for. You can also select the policy name from the drop-down menu. To exclude or make exceptions for the backup of specific files or directories from all policies, select <All Policies>.</p>
Schedule	<p>Enter the schedule name that is associated with the files and the directories to exclude or make exceptions for. You can also select the schedule name from the drop-down menu. To exclude or make exceptions for the backups of specific files or directories from all schedules, select <All Schedules>.</p>
Files/Directories	<p>Enter the full path to the file(s) and the directories to exclude or make exceptions for.</p>

Syntax rules for exclude lists

Symantec suggests that you always specify automounted directories and CD-ROM file systems in the exclude list. Otherwise, if the directories are not mounted at the time of a backup, NetBackup must wait for a timeout.

The following syntax rules apply to exclude lists:

- Only one pattern per line is allowed.
- NetBackup recognizes standard wildcard use.
See [“Wildcards in NetBackup”](#) on page 720.

See “[NetBackup naming conventions](#)” on page 719.

- Spaces are considered legal characters. Do not include extra spaces unless they are part of the file name.

For example, if you want to exclude a file named

```
C:\testfile (with no extra space character at the end)
```

and your exclude list entry is

```
C:\testfile (with an extra space character at the end)
```

NetBackup cannot find the file until you delete the extra space from the end of the file name.

- End a file path with `\` to exclude only directories with that path name (for example, `C:\users\test\`). If the pattern does not end in `\` (for example, `C:\users\test`), NetBackup excludes both files and directories with that path name.
- To exclude all files with a given name, regardless of their directory path, enter the name. For example:

```
test
```

rather than

```
C:\test
```

This example is equivalent to prefixing the file pattern with

```
\
```

```
\*\
```

```
\*\*\
```

```
\*\*\*\
```

and so on.

The following syntax rules apply only to UNIX clients:

- Do not use patterns with links in the names. For example, assume `/home` is a link to `/usr/home` and `/home/doc` is in the exclude list. The file is still backed up in this case because the actual directory path, `/usr/home/doc`, does not match the exclude list entry, `/home/doc`.
- Blank lines or lines which begin with a pound sign (`#`) are ignored.

Windows client example exclude list

Assume that an exclude list in the Exclude Lists host properties contains the following entries:

```
C:\users\doe\john
```

```
C:\users\doe\abc\
```

```
C:\users\*\test
```

```
C:\*\tempcore
```

Given the exclude list example, the following files and directories are excluded from automatic backups:

- The file or directory named `C:\users\doe\john`.
- The directory `C:\users\doe\abc\` (because the exclude entry ends with `\`).
- All files or directories named `test` that are two levels beneath `users` on drive C.
- All files or directories named `temp` that are two levels beneath the root directory on drive C.
- All files or directories named `core` at any level and on any drive.

Traversing excluded directories

An exclude list can indicate a directory for exclusion, while the client uses an include list to override the exclude list. NetBackup traverses the excluded directories if necessary, to satisfy the client's include list.

Assume the following settings for a Windows client:

- The backup policy backup selection list indicates `ALL_LOCAL_DRIVES`. When a scheduled backup runs, the entire client is backed up.
The entire client is also backed up if the backup selection list consists of only:
`/`
- The exclude list on the client consists of only: `*`
An exclude list of `*` indicates that all files are excluded from the backup.
- However, since the include list on the Windows client includes the following file: `C:\WINNT`, the excluded directories are traversed to back up `C:\WINNT`.
If the include list did not contain any entry, no directories are traversed.

In another example, assume the following settings for a UNIX client:

- The backup selection list for the client consists of the following: `/`
- The exclude list for the UNIX client consists of the following: `/`
- The include list of the UNIX client consists of the following directories:
`/data1`
`/data2`
`/data3`

Because the include list specifies full paths and the exclude list excludes everything, NetBackup replaces the backup selection list with the client's include list.

Fibre Transport properties

The Fibre Transport master server properties apply to the SAN clients whose preferences have not been set explicitly.

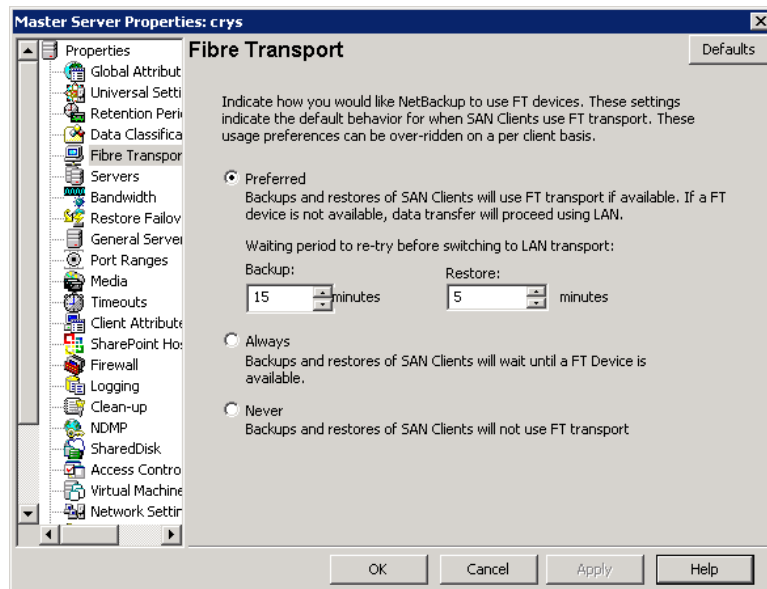
The Fibre Transport media server property applies to the SAN clients for selected media servers.

The Fibre Transport client properties apply to the selected SAN clients. The defaults for clients are the property settings of the master server.

An FT device is the target mode driver on a NetBackup FT media server. An FT pipe is the logical connection that carries backup and restore data between an FT media server and a SAN client.

For more information about NetBackup Fibre Transport, see the *NetBackup Shared Storage Guide*.

Figure 3-28 Master server Fibre Transport host properties



The following topics describe the **Fibre Transport** dialog box.

Preferred

The **Preferred** property specifies to use an FT pipe if an FT device is available within the configured wait period in minutes. If an FT device is not available after the wait period elapses, NetBackup uses a LAN connection for the operation.

If you select this option, also specify the wait period for backups and for restores.

For the global property that is specified on the master server, the default is **Preferred**.

Always

The **Always** property specifies that NetBackup should always use an FT pipe for backups and restores of SAN clients. NetBackup waits until an FT device is available before it begins the operation.

However, an FT device must be active and available. If no FT device exists, NetBackup uses the LAN. An FT device may not exist because none is active, none have been configured, or the SAN Client license expired.

Never

The **Never** property specifies that NetBackup should never use an FT pipe for backups and restores of SAN clients. NetBackup uses a LAN connection for the backups and restores.

If you specify **Never** for the master server, Fibre Transport is disabled in the NetBackup environment. If you select **Never**, you can configure FT usage on a per-client basis.

If you specify **Never** for a media server, Fibre Transport is disabled for the media server.

If you specify **Never** for a SAN client, Fibre Transport is disabled for the client.

Maximum concurrent FT connections

This property applies to the media properties only.

This property specifies the number of FT connections to allow to a media server.

The default is four times the number of HBA target ports (maximum of 16).

Use defaults from the master server configuration

This property applies to the client properties only.

This property specifies that the client follow the properties as they are configured on the master server.

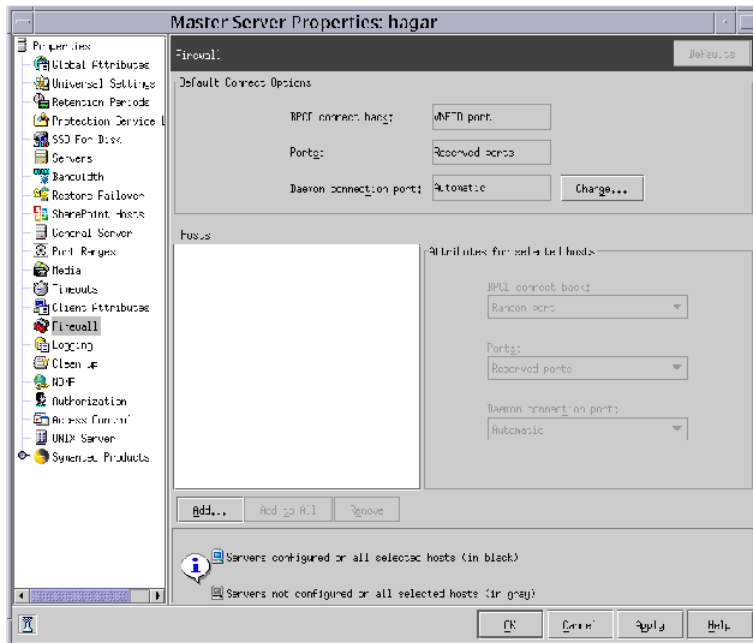
Firewall properties

The Firewall properties describe how the selected master and media servers are connected to by other hosts.

Servers are added to the host list of the Firewall properties. To configure port usage for clients, see the Client Attributes properties.

See “[Client Attributes properties](#)” on page 80.

Figure 3-29 Firewall dialog box



The Firewall host properties contains the following options.

Default connect options

By default, NetBackup selects firewall-friendly connect options under **Default connect options**. However, the default options can be set differently for individual servers under **Attributes for selected Hosts**.

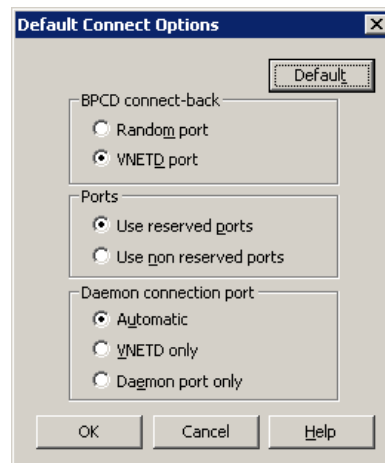
By default, the firewall settings are configured to require the fewest possible ports to be open.

To change the default connect options for the selected server, click **Change**.

Click **Change** to change the **Default connect options**. Change the Firewall properties in the **Default Connect Options** dialog box.

Note: If **VNETD only** is selected as the **Daemon connection port**, the **BPCD connect back** setting is not applicable. If **VNETD only** is selected as the **Daemon connection port**, **Use non-reserved ports** is always used regardless of the value of the Ports setting.

Figure 3-30 Default Connect Options dialog box



Hosts list

To change the default connect options for any server, add the server to the host list. Servers do not automatically appear on the list.

- **Add** button

Click **Add** to add a host entry to the host list. A host must be listed before it can be selected for configuration.

- **Add to all** button

Click **Add to All** to add the listed hosts (along with the specified properties) to all hosts that are selected for host property configuration. (That is, the hosts that are selected upon opening the Host Properties.)

- **Remove** button

Select a host name in the list, then click **Remove** to remove the host from the list.

Attributes for selected hosts

Connect options can be configured for individual servers.

BPCD connect back

This property specifies how daemons are to connect back to the NetBackup Client daemon (`BPCD`) as follows:

- **Use default connect options** (An option for individual hosts)
Use the methods that are specified under **Default connect options**.
- **Random port**
NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method.
- **VNETD port**
This method requires no connect-back. The Veritas Network Daemon (`vnetd`) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. The server initiates all `bpcd` socket connections.
Consider the example in which `bpbxrm` on a media server initially connects with `bpcd` on a client. The situation does not pose a firewall problem because `bpbxrm` uses the well-known `bpcd` port.

Ports

Select whether a reserved or non-reserved port number should be used to connect to the server:

- **Use default connect options** (An option for individual hosts)
Use the methods that are specified under Default attributes.
- **Reserved port**
Connect to the server by a reserved port number.
- **Use non-reserved ports**
Connect to the server by a non-reserved port number. If this property is selected, also enable **Accept connections from non-reserved ports** for the selected server in the Universal Settings properties.
See [“Universal Settings properties”](#) on page 185.

Daemon connection port

Select the **Daemon connection port** method to use to connect to the server:

- **Use default connect options** (An option for individual hosts)
Use the methods that are specified under **Default connect options**.
- **Automatic**
The daemons on the server are connected to by `vnetd` if possible. If it is not possible to use `vnetd`, the daemon's traditional port number makes the connection.
- **VNETD only**
The daemons on the server are connected to by `vnetd` only. Select this property if your firewall rules prevent connections to the server by the traditional port number.
- **Daemon port only**
The daemons on the server are connected to by the traditional port number only.

Note: If **vnetd only** is selected as the Daemon connection port, the **BPCD connect back** setting is not applicable. If **vnetd only** is selected as the Daemon connection port, **Non-reserved port** is always used regardless of the value of the Ports setting.

Setting up vnetd between a server and a client

Use the following procedure to set up `vnetd` between a server and a client.

To set up vnetd between a server and a client

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers** > Double-click on master server > **Client Attributes**.
- 2 In the client list, select the client you want to change.
- 3 Under **BPCD connect back**, select **VNETD port**.
- 4 Click **OK**.

Or, add the client to the client database by running the `bpclient.exe` command, that is located in

```
\Install_path\VERITAS\NetBackup\bin\admincmd.
```

Setting up vnetd between two servers

Use the following procedure to set up `vnetd` between servers.

To set up vnetd between servers

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers** > Double-click on master server > **Firewall**.
- 2 In the host list, select the host you want to change.
- 3 Under **BPCD connect back**, select **VNETD port**.
- 4 Click **OK**.

Enabling logging for vnetd

Use the following procedure to enable logging for `vnetd`.

To enable logging for vnetd

- ◆ Create a `vnetd` directory in the following location:
 - On Windows: `install_path\NetBackup\logs\vnetd`
Or, double-click `mklogdir.bat` in the `install_path\NetBackup\logs\` directory to populate the `logs` directory with log subdirectories, including one for `vnetd`.
 - On UNIX: `/usr/opensv/logs/vnetd`

Example setup for using the vnetd port

The following is a sample configuration to use the `vnetd` port for `bprd`, `bpdbm`, `bpjobjd`, `bpvmd`, and the robotic daemons on master and media servers.

The example uses the **BPCD connect back** connections:

- Change in the configuration file setup
Add the following configuration option to the `vm.conf` file on the machines that can connect to `vmd` or the robotic daemons on *hostname*:

```
CONNECT_OPTIONS = hostname x y z
```

Where:

x is 0 or 1 and is ignored for `vm.conf`.

y is 0 or 1 and is ignored for `vm.conf`.

z is 0 for automatic connections. A `vnetd` connection is tried first when selected. If that fails, a traditional connection is tried.

1 = `vnetd`-only connections.

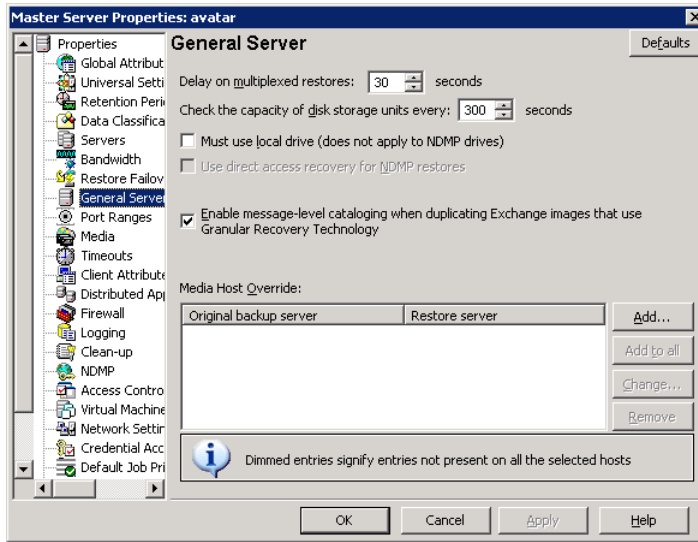
2 = Traditional connections (default)

- Change in the host properties
 - In the Firewall properties for the master server, add an entry in the host list for each remote media server.
(**Host Properties** > **Master Servers** > Selected master server > **Firewall**.)
Under **BPCD connect back**, select **VNETD port**.
Choose **Automatic** for the **Daemon connection port**.
 - In the Firewall properties for each media server, add an entry for each remote server. (**Host Properties** > **Media Servers** > Selected media server > **Firewall**.)
Under **BPCD connect back**, select **VNETD port**.
Choose **Automatic** for the **Daemon connection port**.
 - In the Firewall properties for each Client, add an entry for the Master server. (**Host Properties** > **Clients** > Selected client > **Firewall**.)
Choose **Automatic** for the **Daemon connection port**.
 - In the Client Attributes properties for the Master server, add an entry for each remote client. (**Host Properties** > **Master Servers** > Selected master server > **Client Attributes**.)
Under **BPCD connect back**, select **VNETD port**.

General Server properties

The **General Server** properties apply to selected master and media servers.

Figure 3-31 General Server dialog box



The following topics describe the **General Server** host properties.

Delay on multiplexed restores

This property specifies how long the server waits for additional restore requests of multiplexed images on the same tape. All of the restore requests that are received within the delay period are included in the same restore operation (one pass of the tape). The default is a delay of 30 seconds.

Check the capacity of disk storage units every

This property determines how often NetBackup checks disk storage units for available capacity. If checks occur too frequently, then system resources are wasted. If checks do not occur often enough, too much time elapses and backup jobs are delayed. The default is 300 seconds (5 minutes).

Note: This property applies to the disk storage units of 6.0 media servers only. Subsequent releases use internal methods to monitor disk space more frequently.

Must use local drive

This property appears for master servers only, but applies to all media servers as well. This property does not apply to NDMP drives.

If a client is also a media server or a master server and **Must use local drive** is checked, a local drive is used to back up the client. If all drives are down, another can be used.

This property increases performance because backups are done locally rather than sent across the network. For example, in a SAN environment a storage unit can be created for each SAN media server. Then, the media server clients may be mixed with other clients in a policy that uses ANY AVAILABLE storage unit. When a backup starts for a client that is a SAN media server, the backups go to the SAN connected drives on that server.

Use direct access recovery for NDMP restores

By default, NetBackup for NDMP is configured to use Direct Access Recovery (DAR) during NDMP restores. DAR can reduce the time it takes to restore files by allowing the NDMP host to position the tape to the exact location of the requested file(s). Only the data that is needed for those files is read.

Clear this check box to disable DAR on all NDMP restores. Without DAR, NetBackup reads the entire backup image, even if only a single restore file is needed.

Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology

This option performs message-level cataloging when you duplicate Exchange backup images that use Granular Recovery Technology (GRT) from disk to tape. To perform duplication more quickly, you can disable this option. However, then users are not able to browse for individual items on the image that was duplicated to tape.

See the *NetBackup for Exchange Administrator's Guide*.

Media host override list

Specific servers can be specified in this list as servers to perform restores, regardless of where the files were backed up. (Both servers must be in the same master and media server cluster.) For example, if files were backed up on media server A, a restore request can be forced to use media server B.

The following items describe situations in which the capability to specify servers is useful:

- Two (or more) servers share a robot and each have connected drives. A restore is requested while one of the servers is either temporarily unavailable or is busy doing backups.

- A media server was removed from the NetBackup configuration, and is no longer available.

To add a host to the **Media host override** list, click **Add**.


Click **Add to All** to add a host to the list for all of the hosts currently selected.

To change an entry in the list, select a host name, then click **Change**.

Configure the following options in the Add Media Override settings or Change Media Override settings dialog box:

- **Original backup server**
Type the name of the server where data was backed up originally.
- **Restore server**
Type the name of the server that is to process future restore requests.

Figure 3-32 Add Media Override settings dialog box



Forcing restores to use a specific server

Use the following procedure to force restores to use a specific server.

To force restores to use a specific server

- 1 If necessary, physically move the media to the host to answer the restore requests, then update the Enterprise Media Manager database to reflect the move.
- 2 Modify the NetBackup configuration on the master server. Add the original backup media server and the restore server to the **Media host override** list in the General Server host properties.
- 3 Stop and restart the NetBackup Request Manager service (`bprd`) on the master server.

This process applies to all storage units on the original backup server. Restores for any storage unit on the **Original backup server** go to the server that is listed as the **Restore server**.

To revert to the original configuration for future restores, delete the line from the **Media host override** list.

Disabling the cataloging for duplications of Exchange backups using Granular Recovery Technology (GRT)

Unlike a duplication of a backup that uses Granular Recovery Technology (GRT) from tape to disk, duplication of the same backup from disk to tape takes extra time. NetBackup requires this extra time to catalog the granular Exchange information. You can choose not to catalog the granular information so that the duplication is performed more quickly. However, then users are not able to browse for individual items on the image that was duplicated to tape if the disk copy expires.

During the duplication process, NetBackup writes log entries periodically to show the progress of the job.

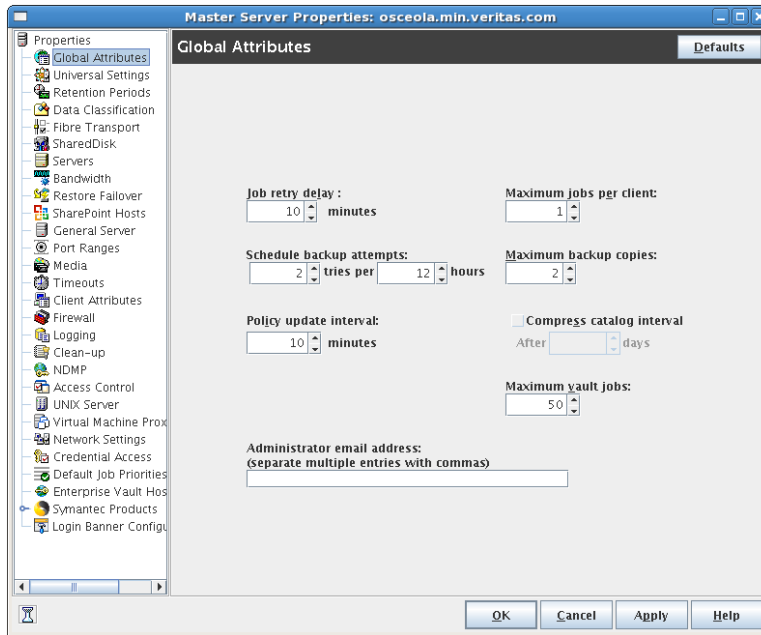
To disable the cataloging of Exchange backups using Granular Recovery Technology

- 1 On the master server, open the NetBackup Administration Console.
- 2 In the left pane, expand **Host Properties**.
- 3 Click **Master Servers**.
- 4 In the right pane, right-click the master server click **Properties**.
- 5 Click **General Server**.
- 6 Uncheck **Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology**.
- 7 Click **OK**.

Global Attributes properties

The **Global Attributes** properties apply to currently selected master servers. The **Global Attributes** properties affect all operations for all policies and clients. The default values are adequate for most installations but can be changed.

Figure 3-33 Global Attributes dialog box



The following topics describe the Global Attributes properties.

Job retry delay

This property specifies how often NetBackup retries a job. The default is 10 minutes. The maximum is 60 minutes; the minimum is 1 minute.

Schedule backup attempts

NetBackup considers the failure history of a policy to determine whether or not to run a scheduled backup job. The **Schedule backup attempts** property sets the timeframe for NetBackup to examine..

This property determines the following characteristics for each policy:

- How many preceding hours NetBackup examines to determine whether to allow another backup attempt (retry). By default, NetBackup examines the past 12 hours.
- How many times a backup can be retried within that timeframe. By default, NetBackup allows two attempts. Attempts include the scheduled backups that start automatically or the scheduled backups that are user-initiated.

Consider the following example scenario using the default setting 2 tries every 12 hours:

- Policy_A runs at 6:00 P.M.; Schedule_1 fails.
- Policy_A is user-initiated at 8:00 P.M.; Schedule_2 fails.
- At 11:00 P.M., NetBackup looks at the previous 12 hours. NetBackup sees one attempt at 6:00 P.M. and one attempt at 8:00 P.M. The **Schedule backup attempts** setting of two has been met so NetBackup does not try again.
- At 6:30 A.M. the next morning, NetBackup looks at the previous 12 hours. NetBackup sees only one attempt at 8:00 P.M. The **Schedule backup attempts** setting of two has not been met so NetBackup tries again. If a schedule window is not open at this time, NetBackup waits until a window is open.

Note: This attribute does not apply to user backups and archives.

Policy update interval

This property specifies how long NetBackup waits to process a policy after a policy is changed. The interval allows the NetBackup administrator time to make multiple changes to the policy. The default is 10 minutes. The maximum is 1440 minutes; the minimum is 1 minute.

Maximum jobs per client

This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. The default is 1 job.

NetBackup can process concurrent backup jobs from different policies on the same client only in the following situations:

- More than one storage unit available
- One of the available storage units can perform more than one backup at a time.

[Figure 3-34](#) shows how the files that are on the same client but in different policies, can be backed up concurrently to different storage devices.

You can specify any number of concurrent jobs within the following constraints:

- Number of storage devices
NetBackup can perform concurrent backups to separate storage units or to drives within a storage unit. For example, a single Media Manager storage unit supports as many concurrent backups as it has drives. A disk storage unit is a directory on disk, so the maximum number of jobs depends on system capabilities.

- Server and client speed

Too many concurrent backups on an individual client interfere with the performance of the client. The best setting depends on the hardware, operating system, and applications that are running.

The **Maximum jobs per client** property applies to all clients in all policies.

To accommodate weaker clients (ones that can handle only a small number of jobs concurrently), consider using one of the following approaches:

- Set the **Maximum data streams** property for those weaker client(s) appropriately. (This property is found under **Host Properties > Master Server > Client Attributes > General** tab.)

See “[Maximum data streams](#)” on page 81.

- Use the **Limit jobs per policy** policy setting in a client-specific policy. (A client-specific policy is one in which all clients share this characteristic).


See “[Limit jobs per policy attribute](#)” on page 471.

- Network loading

The available bandwidth of the network affects how many backups can occur concurrently. Two Exabyte 8500, 8mm tape drives can create up to a 900-kilobyte-per-second network load. Depending on other factors, the load might be too much for a single Ethernet. For loading problems, consider backups over multiple networks or compression.

A special case exists to back up a client that is also a server. Network loading is not a factor because the network is not used. Client and server loading, however, is still a factor.

Figure 3-34 Maximum jobs per client



Note: If online, hot catalog backups are scheduled to occur concurrently with other backups for the master server, set the **Maximum jobs per client** value to greater than two. The higher setting ensures that the catalog backup can proceed while the regular backup activity occurs.

Maximum backup copies

This property specifies the total number of backup copies that can exist in the NetBackup catalog (2 through 10).

NetBackup creates one of the following, whichever is smaller:

- The number of copies that are specified under **Multiple copies**
See “[Multiple copies attribute](#)” on page 503.
- The number of copies that are specified as the **Maximum backup copies** property

Note: To configure multiple copies for a relocation schedule, set the **Maximum backup copies** property to include an additional copy beyond the number of copies to be created in the **Multiple Copies** dialog box. A relocation schedule is created as part of a disk staging storage unit. For example, to create four copies in the **Multiple Copies** dialog box, set the **Maximum Backup Copies** property to five or more.

See “[Criteria for creating multiple copies](#)” on page 503.

Compress catalog interval

This property specifies how long NetBackup waits after a backup before it compresses the image catalog file.

Maximum vault jobs

This property specifies the maximum number of vault jobs that are allowed to be active on the master server. The greater the maximum number of vault jobs, the more system resources are used.

Administrator email address property

This property specifies the address(es) where NetBackup sends notifications of scheduled backups or administrator-directed manual backups.

To send the information to more than one administrator, separate multiple email addresses by using a comma, as follows:

useraccount1@company.com,useraccount2@company.com

Disaster recovery information that is created during online, hot catalog backups is not sent to the addresses indicated here. Disaster recovery information is sent

to the address that is indicated on the **Disaster Recovery** tab in the catalog backup policy.

See [“About the Disaster Recovery tab”](#) on page 569.

Setting up email notifications about backups

Email notifications can be sent to the client's administrator or to the global administrator, specifying that a backup was successful or unsuccessful.

The following represents the contents of a notification email:

```
Backup on client hostname by root was partially successful.  
File list  
-----  
C:\Documents and Settings
```

Before notification emails about backups are sent, the computing environment must be configured correctly.

NetBackup can send notification to specified email addresses about backups on all client or specific clients.

Choose one or both of the following notification methods:

- Send emails about failed backups only.
Send a message to the email address(es) of the NetBackup administrator(s) about any backup that ends in a non-zero status. (**Server sends mail** host property is enabled in Universal Settings.)
- Send emails about successful and failed backups.
Send a message to the local administrator(s) of each client about successful and unsuccessful backups. (**Client sends mail** host property is enabled in Universal Settings.)

Both methods require that the `nbmail.cmd` script be configured.

Both methods require that the host properties be configured with email addresses:

- See [“Sending email notifications to the administrator about unsuccessful backups”](#) on page 138.
- See [“Sending messages to the global administrator about unsuccessful backups”](#) on page 139.
- See [“Sending messages to the administrator about successful and unsuccessful backups”](#) on page 139.

Windows systems require that an application to transfer messages using the Simple Mail Transfer Protocol be installed to accept script parameters. UNIX platforms have an SMTP transfer method built into the system.

See “Installing the email utility” on page 140.

Configuring the nbmail.cmd script

To receive email notifications about backups, the `nbmail.com` script must be configured for Windows.

Locate `install_path\VERITAS\NetBackup\bin\nbmail.cmd` on a NetBackup master server. If configuring the script on the client, copy `nbmail.cmd` from a master server to the client. By default, `nbmail.cmd` does not send email.

The following options are used in the script:

- s The subject line of the email
- t Indicates who receives the email.
- i The originator of the email, though it is not necessarily known to the email server. The default (-i NetBackup) shows that the email is from NetBackup.
- server The name of the SMTP server that is configured to accept and relay emails.
- q Suppresses all output to the screen.

Use the following procedure to configure the `nbmail.cmd` script.

To configure the `nbmail.cmd` script

- 1 Use a text editor to open `nbmail.cmd`. Create a backup copy of `nbmail.cmd` before modifying it.

In some text editors, using the word wrap option can create extra line feeds in the script and render it non-functional.

- 2 Most of the lines are informational in `nbmail.cmd`.

Locate the following lines in the script:

```
@REM @IF "%~4"==" " (  
@REM blat %3 -s %2 -t %1 -i NetBackup -server SERVER_1 -q  
@REM ) ELSE (  
@REM blat %3 -s %2 -t %1 -i NetBackup -server SERVER_1 -q -attach %4  
@REM )
```

- 3 Adjust the five lines as follows:

- Remove `@REM` from each of the five lines to activate the necessary sections for BLAT to run.

- Replace `SERVER_1` with the name of the email server. For example:

```
@IF "%~4"==" " (
blat %3 -s %2 -t %1 -i NetBackup -server emailserver.company.com -q
) ELSE (
blat %3 -s %2 -t %1 -i NetBackup -server emailserver.company.com -q -attach %4
)
```

- 4 Save `nbmail.cmd`.

Sending email notifications to the administrator about unsuccessful backups

Use the following procedure to send email notifications to a client's administrator only if the backups have a non-zero status.

To send email notifications to the administrator for backups with a non-zero status

- 1 Install and configure a mail client on the server.
See [“Installing the email utility”](#) on page 140.
- 2 Edit the `nbmail.cmd` script on the server.
See [“Configuring the nbmail.cmd script”](#) on page 137.
- 3 Open the NetBackup Administration Console on the master server.
- 4 Expand **NetBackup Management > Host Properties > Master Server**.
- 5 Open the properties of the master server.
- 6 Select **Universal Settings**.
- 7 In the **Client administrator’s email** field, enter the email address of the administrator to receive the notification emails. (Separate multiple addresses with commas.)
See [“Universal Settings properties”](#) on page 185.
- 8 Enable the **Server sends mail** option and click **Apply**.

See [“Sending messages to the global administrator about unsuccessful backups”](#) on page 139.

See [“Sending messages to the administrator about successful and unsuccessful backups”](#) on page 139.

Sending messages to the global administrator about unsuccessful backups

Use the following procedure to send messages to the global administrator about backups with a non-zero status.

To send messages to the global administrator about backups with a non-zero status

- 1 Install and configure a mail client on the server.
See [“Installing the email utility”](#) on page 140.
- 2 Edit the `nbmail.cmd` script on the server.
See [“Configuring the nbmail.cmd script”](#) on page 137.
- 3 Open the NetBackup Administration Console on the master server.
- 4 Expand **NetBackup Management > Host Properties > Master Server**.
- 5 Open the host properties of the master server.
- 6 Select **Global Attributes**.
- 7 In the **Administrator’s email address** field, enter the email address of the administrator to receive the notification emails. (Separate multiple addresses with commas.) Click **Apply**.

The global administrator’s email address can also be changed by using the `bpconfig` command on the master server:

```
Install_Path\NetBackup\bin\admincmd\bpconfig -ma email_address
```

For example:

```
C:\Program Files\VERITAS\NetBackup\bin\admincmd\bpconfig  
-ma name@company.com
```

See [“Sending email notifications to the administrator about unsuccessful backups”](#) on page 138.

See [“Sending messages to the administrator about successful and unsuccessful backups”](#) on page 139.

Sending messages to the administrator about successful and unsuccessful backups

An alternative to sending all emails through the master server is to send emails through each client. An email can be sent to each client’s administrator after all backups.

To send email notifications for all backups from a client

- 1 Install and configure a mail client on the client.
See “[Installing the email utility](#)” on page 140.
- 2 Edit the `nbmail.cmd` script on the client.
See “[Configuring the nbmail.cmd script](#)” on page 137.
- 3 Open the NetBackup Administration Console on the master server.
- 4 Expand **NetBackup Management > Host Properties > Clients**.
- 5 Open the host properties for a client. Multiple clients can also be selected.
- 6 Select **Universal Settings**.
- 7 In the **Client administrator’s email** field, enter the email address of the administrator to receive the notification emails. (Separate multiple addresses with commas.)
See “[Universal Settings properties](#)” on page 185.
- 8 Enable the **Client sends mail** option and click **Apply**.
See “[Sending email notifications to the administrator about unsuccessful backups](#)” on page 138.
See “[Sending messages to the global administrator about unsuccessful backups](#)” on page 139.

Installing the email utility

BLAT is the most common application is used for email notification. It is a mail client in the public domain. BLAT is used as an example in the following discussions.

Use the following procedure to install and configure the email utility.

To install and configure the email utility

- 1 Download the `.ZIP` file from the BLAT download page, currently: www.blat.net
- 2 Extract the files to a directory.

3 Copy the `blat.exe` file to the Windows System32 directory.

4 From a command prompt, run the following command:

```
blat -install emailserver.company.com useraccount@company.com
```

Where:

emailserver.company.com is the hostname or IP address of the email server that sends the email notifications.

useraccount@company.com is the primary account to send the emails from the specified server.

Use the following procedure to test the email utility.

To test the email utility

1 Create a test text file that contains a message. For example, create

```
C:\testfile.txt
```

2 From a command prompt, run:

```
blat C:\testfile.txt -s test_subject -to useraccount@company.com
```

A correct setup sends the contents of `testfile.txt` to the email address specified.

3 Use the following list to troubleshoot problems if NetBackup notification does not work correctly:

- Make sure that the BLAT command is not commented out in the `nbmail.cmd` script.
See [“Installing the email utility”](#) on page 140.
- Make sure that the path to `blat.exe` is specified in `nbmail.cmd` if the command is not in the `\system32` directory.
- Make sure that BLAT syntax has not changed in the later versions of BLAT. Check the README for the version of BLAT running on the system.
- The BLAT command may need the `-t n` timeout parameter if the system experiences delays. (*n* represents seconds.)
- The BLAT binary must not be corrupt or incompatible with the email system. Download the latest version.
- Configure the email addresses correctly in the host properties.
- The email account that is specified must be a valid on the email server.

- If the email server requires authentication for SMTP, make sure that the account that is used for the NetBackup client process is authorized. The default account is the local system.

Logging properties

The Logging properties apply to the master servers, media servers, and clients that are currently selected. The available properties differ between master servers, media servers, and clients.

The Logging properties contain the processes that continue to use legacy logging as well as processes that use unified logging.

Unified logging Unified logging creates log file names and messages in a format that is standardized across Symantec products. Some NetBackup processes on the server use unified logging.

Unified logging writes the logs into subdirectories in the following locations:

- UNIX: `/usr/opensv/logs`
- Windows: `install_path\NetBackup\logs`

Note: Do not save logs to a remote file system such as NFS or CIFS. Logs that are stored remotely and then grow large can cause critical performance issues.

Unlike legacy logging, subdirectories for the processes that use unified logging are created automatically.

To control the size and number of unified logs, use the `vxlogcfg` command and the `vxlogmgr` command.

Legacy logging

For those processes that use legacy logging, administrators must first create a log directory for each process to be logged. A logging level selection on the Logging properties page does not enable logging.

Create the NetBackup legacy log directories in the following locations:

- UNIX: `/usr/opensv/netbackup/logs/process_name`
- Windows: `install_path\NetBackup\logs\process_name`

Note: Do not save logs to a remote file system such as NFS or CIFS. Logs that are stored remotely and then grow large can cause critical performance issues.

On a Windows server, you can create all of the NetBackup debug log directories at one time by double-clicking `mklogdir.bat` in the following directory:

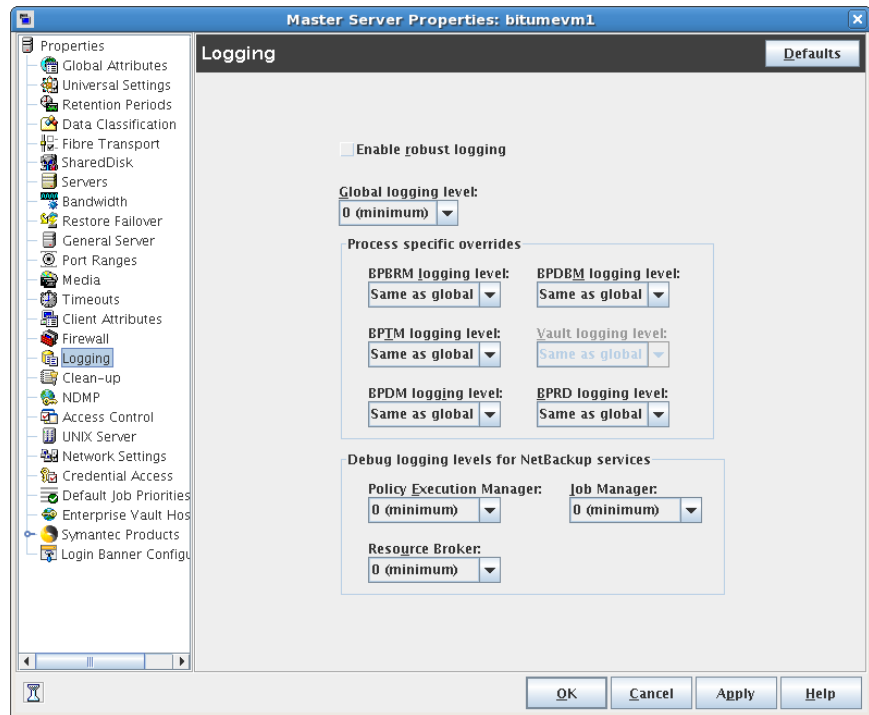
```
install_path\NetBackup\logs\
```

Create the Media Manager legacy log directories in the following locations:

- UNIX: `/usr/opensv/volmgr/debug`
- Windows: `install_path\Volmgr\debug`

For details on both unified and legacy logging, see the *NetBackup Troubleshooting Guide*.

Figure 3-35 Logging dialog box



The Logging properties can contain the following options.

Enable robust logging

A check in the **Enable robust logging** check box indicates that when a log file grows to the maximum size, the log file is closed. When the log file is closed, a new log file is opened. If the new log file causes the maximum number of log files in the directory to be exceeded, the oldest log file is deleted.

See the *NetBackup Troubleshooting Guide* for more information about controlling the log file size.

If this property is enabled, the following processes produce log files:

- bprd
- bpbkar
- bpbzm
- bpcd

- bpdbm
- bptm
- bpdm

The logs are named using the following convention:

MMDDYY_NNNNN.log

where *NNNNN* is an incrementing counter from 00001 - 99999

If the **Enable robust logging** property is disabled, a single log file is produced each day:

MMDDYY.log

Whether **Enable robust logging** is selected or not, the log file is pruned by using `KEEP_LOGS_DAYS` and `DAYS_TO_KEEP_LOGS` settings.

Note: If a NetBackup environment uses scripts depending on the *MMDDYY.log* naming convention, either update the scripts or disable Robust Logging.

Global logging level

This property is used for debugging purposes. The logging levels control the amount of information that the NetBackup server writes to logs. Six levels are supported. Select from between **0** (minimum logging level) through **5** (maximum logging level).

Note: Use the default setting of 0 unless advised otherwise by Symantec Technical Support. Other settings can cause the logs to accumulate large amounts of information.

Some NetBackup processes allow individual control over the amount of information the process writes to logs. For those processes, specify a different logging level other than the **Global logging level**.

Process specific overrides

To override the **Global logging level** property for any of the following services, select a different logging level. Select from between **0** (minimum) through **5** (maximum).

The services within this section use legacy logging:

- **BPBRM logging level**
- **BPTM logging level**
- **BPDM logging level**
- **BPDBM logging level**
- **Vault logging level**
Select a logging level for `bpvault`.
- **BPRD logging level**

These services require that you first create a log directory in the following location:

- **UNIX:** `/usr/opensv/netbackup/logs/process_name`
- **Windows:** `install_path\NetBackup\logs\process_name`

Debug logging levels for NetBackup services

The Logging properties page offers configurable debug levels for the services that use unified logging.

Each service creates a log automatically in the following directories:

- **UNIX:** `/usr/opensv/logs`
- **Windows:** `install_path\NetBackup\logs`

To override the **Global logging level** property for any of the following services, select a different logging level. Select from between **0** (minimum) through **5** (maximum). To retain no logging information, select **No logging** for the service.

You can also use the `vxlogcfg` command to change debug levels.

See the *NetBackup Troubleshooting Guide* for more information.

The following services use unified logging:

- **Policy Execution Manager**
This property appears for EMM servers. The Policy Execution Manager (`NBPEM`) creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, `NBPEM` is notified and the appropriate Policy/Client tasks are updated.
- **Job Manager**
This property appears for EMM servers. The Job Manager (`NBJM`) accepts the jobs that the Policy Execution Manager submits and acquires the necessary resources.
- **Resource Broker**

The Resource Broker (NBRB) makes the allocations for storage units, tape drives, client reservations.

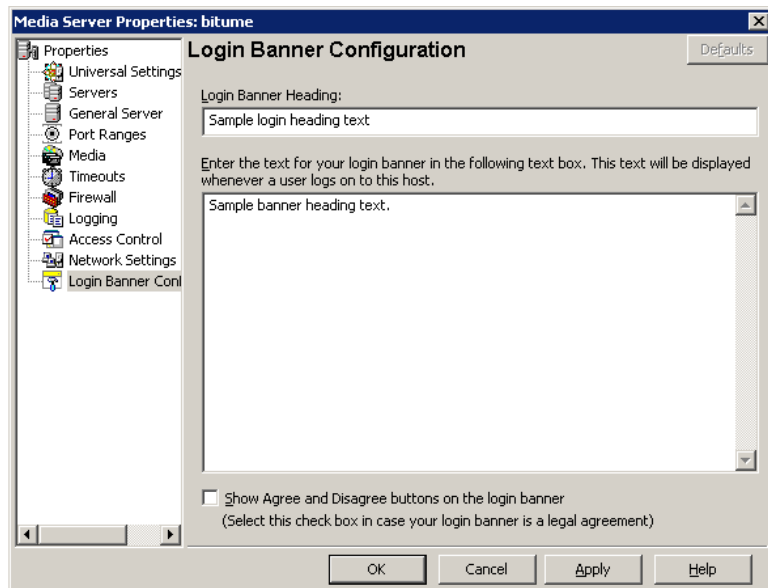
Login Banner Configuration properties

Use the **Login Banner Configuration** properties to configure a banner screen that appears each time a user logs into the NetBackup Administration Console or the Backup, Archive, and Restore client console. The **Login Banner Configuration** properties can be configured to make it mandatory for the user to acknowledge the login banner screen before the user can access the console.

A different login banner can be configured for any master server, media server, or client.

Figure 3-36 shows example banner text for a media server.

Figure 3-36 Login Banner Configuration dialog box



The first time that the NetBackup Administration Console is launched, the **Login Banner Configuration** properties are not configured so no banner appears to the user. The **Login Banner Configuration** host properties must be configured in order for the banner to appear.

The user can change the server once they log into the console. (From the **File** menu, select **Change Server**.) If the banner is configured for the remote server, the banner appears on the remote server as well.

Note: The banner is not available on NetBackup versions earlier than 6.5.4. If a user changes to a host that is at NetBackup version 6.5.3 or earlier, no banner appears.

If a user opens a new console or window from the existing console, the banner does not appear for the new window. (From the **File** menu, select the **New Console** option or the **New Window from Here** option.)

The following topics describe the **Login Banner Configuration** host properties.

Login Banner Heading

Enter the text that is to appear in the banner.

Text of login banner

Enter the text for the banner message. The maximum is 29,000 characters.

Show Agree and Disagree buttons on the login banner

Configure this option when approval is necessary to use the NetBackup Administration Console or the Backup, Archive, and Restore client console. Specific approval may be required due to a legal agreement at the company in which the NetBackup environment resides.

If this option is enabled, users are required to click the **Agree** option and then click **OK** before the console opens. The agreement is meant only for the user that reads and agrees to the message.

If the user chooses the **Disagree** option, the screen is closed.

Figure 3-37 Login Banner with agreement option enabled

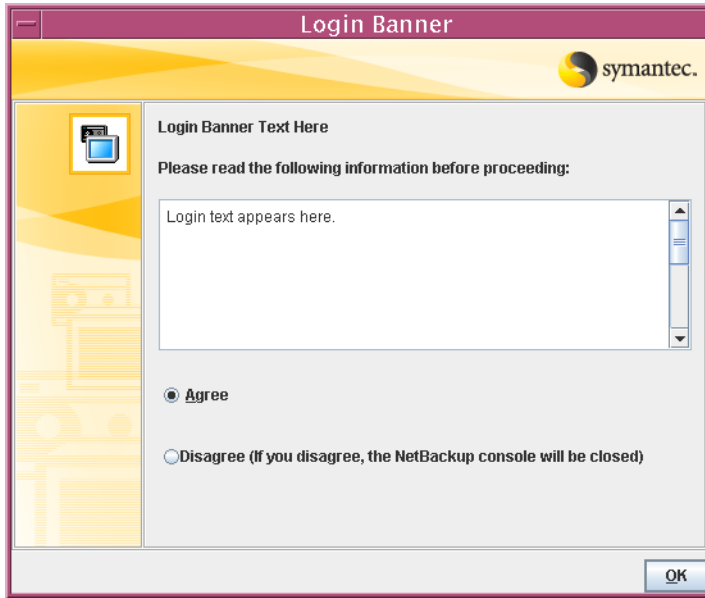
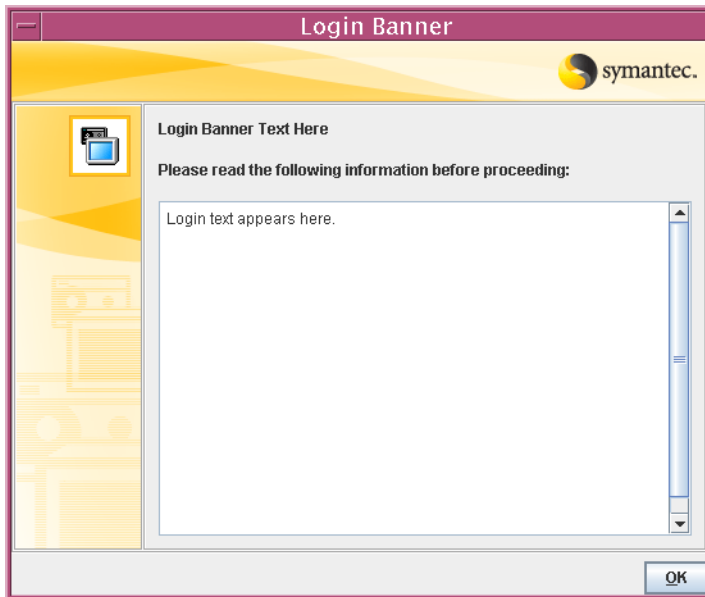


Figure 3-38 Login Banner without agreement option



Removing login banner screen and text

To remove the banner and the text that appears after a user logs into NetBackup, use the following procedure:

To remove the login banner screen and text

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties**.
- 2 Select the host type (**Master Server**, **Media Server**, or **Client**), depending on the host that displays the login banner. Then, select the host name to display the properties.
- 3 Select the **Login Banner Configuration** host properties.
- 4 Clear the **Login Banner Heading** text and the login banner text. Then, click **OK** to save the changes.

Auto log off timeout option

A related option, but one not configured in the **Login Banner Configuration** host properties, is the **Auto log off timeout** option.

The **Auto log off timeout** option allows NetBackup to automatically log a user out of the NetBackup Administration Console after a period of inactivity. The session must be inactive for the configurable number of minutes, hours, or days before the logoff.

To access the **Auto log off timeout** option, select **View > Options**. Then select the **Administration Console** tab.

To log off users automatically, enable the **Auto log off timeout** option. Then, select the duration after which the user is logged off from an inactive session. The minimum logoff duration is 10 minutes and the maximum is two days.

Five minutes before the timeout value is reached, NetBackup warns that the session is to expire in five minutes.

If the logoff warning appears, the user can choose one of the following options:

■ Ignore

If the user selects this option (or does not respond to the warning), a dialog box displays the time that remains before the session ends. Countdown warnings display every minute until the timeout value is reached. When the session ends, the user is logged out of the NetBackup Administration Console or the Backup, Archive, and Restore console.

■ Extend

If the user selects this option, the session continues and the timeout is extended by the logoff timeout value.

If the user begins to work at the console again, the logoff is canceled until the console is left idle again.

■ **Log off**

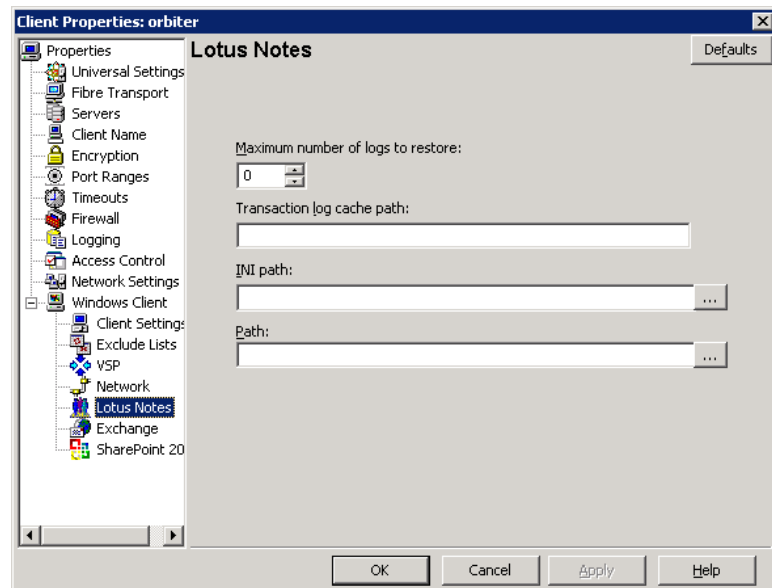
If the user selects this option, the session ends and NetBackup logs off the user immediately.

Lotus Notes properties

The Lotus Notes properties apply to the clients that are currently selected and that run NetBackup for Lotus Notes.

For more information, see the *NetBackup for Lotus Notes Administrator's Guide*.

Figure 3-39 Lotus Notes dialog box



The following topics describe the **Lotus Notes** host properties.

Maximum number of logs to restore

The maximum number of logs that can be prefetched in a single restore job during recovery. Specify a value greater than 1.

Note: If this value is less than or equal to 1, NetBackup does not gather transaction logs during recovery. One transaction log per job is restored to the Domino Server's log directory.

Transaction log cache path

Specify a path where NetBackup can temporarily store the prefetched transaction logs during recovery.

For example:

- UNIX: `/tmp/logcache`
- Windows: `D:\LogCache`

Note: If no path is specified, then NetBackup restores the logs to the Domino server's transaction log directory.

Note the following before specifying the **Transaction log cache path**:

- If the specified path does not exist then it is created during restore.
- The restore job fails with a Status 5 error if the user does not have write permission for the folder.
- Transaction logs are restored to the original location, the Domino transaction log directory, if a path is not specified.
- If the value of **Maximum number of logs to restore** is less than or equal to 1 then this path is ignored. The logs are not prefetched; one transaction log per job is restored to the Domino Server's log directory.
- If there is not sufficient space to restore the specified number of logs, then NetBackup tries to restore only the number of logs that can be accommodated.

INI file

Enter the NOTES.INI file that is associated with the server used to back up and restore the Lotus database. Use this setting to specify the correct .INI file to back up and restore from Domino partitioned servers. Specifying the .INI file for non-partitioned servers is not necessary.

Specify the absolute path to the NOTES.INI file:

- Windows
If the notes.ini file is not located in the default directory, indicate its location in the INI path box. For example:


```
D:\Lotus\Domino\notes.ini
```

- **UNIX**

If the notes.ini is not located in the directory that is specified in the Path, indicate its location here. For example:

```
/db/notesdata/notes.ini
```

Include the directory and the notes.ini file name.

Path

Specify the path where the Lotus Notes program files reside on the client. NetBackup must know where these files are to perform backup and restore operations. The value in this box overrides the Lotus registry key, if both are defined.

Specify the path where the Lotus Notes program files reside on the client:

- **Windows**

Specify the path for Lotus program directory (where nserver.exe resides). For example:

```
D:\Lotus\Domino
```

- **UNIX**

Specify a path that includes the Domino data directory, the Lotus program directory, and the Lotus resource directory. For example:

```
/export/home/notesdata:/opt/lotus/notes/latest
```

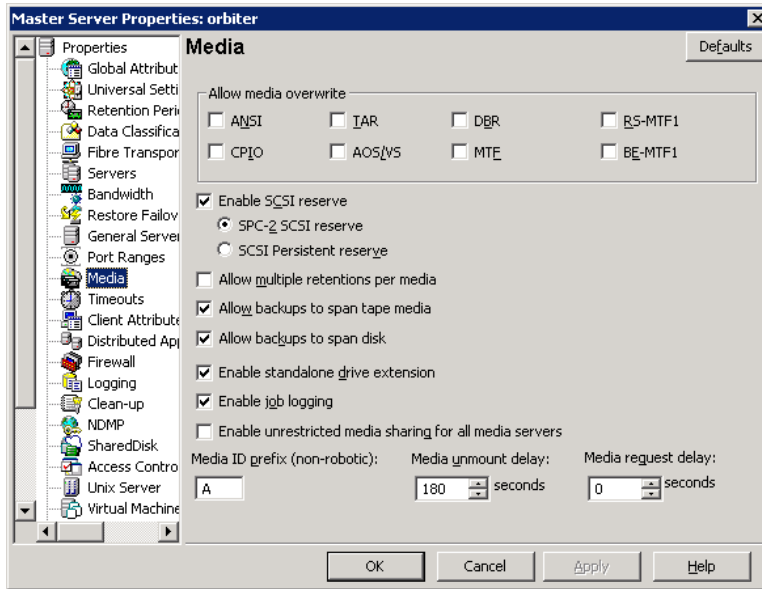
```
/sunspa:/opt/lotus/notes/latest/sunspa/res/C
```

The **Path** value overrides the Lotus registry value, if both are defined.

Media properties

The Media properties apply to the master servers and media servers that are currently selected. Media properties control how NetBackup manages media.

Figure 3-40 Media dialog box



The following topics describe the Media host properties.

Allow media overwrite property

This property overrides NetBackup's overwrite protection for specific media types. Normally, NetBackup does not overwrite certain media types. To disable overwrite protection, place a check in the check box of one or more of the listed media formats.

For example, place a check in the CPIO check box to permit NetBackup to overwrite the cpio format.

By default, NetBackup does not overwrite any of the formats on removable media, and logs an error if an overwrite attempt occurs. This format recognition requires that the first variable length block on a media be less than or equal to 32 kilobytes.

The following media formats on removable media can be selected to be overwritten:

- When ANSI is enabled, ANSI labeled media can be overwritten.
- When AOS/VS is enabled, AOS/VS media can be overwritten. (Data General AOS/VS backup format.)
- When CPIO is enabled, CPIO media can be overwritten.

- When DBR is enabled, DBR media can be overwritten. (The DBR backup format is no longer used.)
- Remote Storage MTF1 media format. When MTF1 is enabled, Remote Storage MTF1 media format can be overwritten.
- When TAR is enabled, TAR media can be overwritten.
- When MTF is enabled, MTF media can be overwritten. With only MTF checked, all other MTF formats can be overwritten. (The exception is Backup Exec MTF (BE-MTF1) and Remote Storage MTF (RS-MTF1) media formats, which are not overwritten.)
- When BE-MTF1 is enabled, Backup Exec MTF media can be overwritten.

If media contains one of the protected formats and media overwrites are not permitted, NetBackup takes the following actions:

- If the volume has not been previously assigned for a backup
 - Sets the volume's state to FROZEN
 - Selects a different volume
 - Logs an error
- If the volume is in the NetBackup media catalog and was previously selected for backups
 - Sets the volume's state to SUSPENDED
 - Aborts the requested backup
 - Logs an error
- If the volume is mounted for a backup of the NetBackup catalog, the backup is aborted and an error is logged. The error indicates the volume cannot be overwritten.
- If the volume is mounted to restore files or list the media contents, NetBackup aborts the request and logs an error. The error indicates that the volume does not have a NetBackup format.

Enable SCSI reserve

This property allows exclusive access protection for tape drives. With access protection, other host bus adaptors cannot issue commands to control the drives during the reservation.

SCSI reservations provide protection for NetBackup Shared Storage Option environments or any other multiple-initiator environment in which drives are shared.

The protection setting configures access protection for all tape drives from the media server on which the option is configured. You can override the media server setting for any drive path from that media server.

See “[Recommended use for Enable SCSI reserve property](#)” on page 159.

See “[Add Path options](#)” on page 235.

The following are the protection options:

- The **SCSI persistent reserve** option provides SCSI persistent reserve protection for SCSI devices. The devices must conform to the SCSI Primary Commands - 3 (SPC-3) standard. SCSI persistent reserve is valid for NetBackup 6.5 and later servers only. If you enable SCSI persistent reserve, NetBackup does not send persistent reserve commands to NetBackup media servers earlier than release 6.5.
- The **SPC-2 SCSI reserve** option (default) provides SPC-2 SCSI reserve protection for SCSI devices. The devices must conform to the reserve and release management method in the SCSI Primary Commands - 2 standard.
- To operate NetBackup without tape drive access protection, clear the **Enable SCSI reserve** property. If unchecked, other HBAs can send the commands that may cause a loss of data to tape drives.

Note: Ensure that all of your hardware processes SCSI persistent reserve commands correctly. All of your hardware includes Fibre Channel bridges. If the hardware does not process SCSI persistent reserve commands correctly and NetBackup is configured to use SCSI persistent reserve, no protection may exist.

Allow multiple retentions per media

This property allows NetBackup to mix retention levels on tape volumes. It applies to media in both robotic drives and nonrobotic drives. The default is that the check box is clear and each volume can contain backups of only a single retention level.

Allow backups to span tape media

This property, when checked, allows backups to span to multiple tape media. This property allows NetBackup to select another volume to begin the next fragment. The resulting backup has data fragments on more than one volume. The default is that **Allow backups to span tape media** is checked and backups are allowed to span media.

If the end of media is encountered and this property is not selected, the media is set to FULL and the operation terminates abnormally. This action applies to both robotic drives and nonrobotic drives.

Allow backups to span disk

This property allows backups to span disk volumes when one disk volume becomes full. The default is that this property is enabled.

The **Allow backups to span disk** property does not apply to AdvancedDisk or OpenStorage storage units. Backups span disk volumes within disk pools automatically.

The following destinations support disk spanning:

- A BasicDisk storage unit spanning to a BasicDisk storage unit. The units must be within a storage unit group.
- An OpenStorage or AdvancedDisk volume spanning to another volume in the disk pool.

For disk spanning to occur, the following conditions must be met:

- The storage units must share the same media server.
- The multiplexing level on spanning storage units should be the same. If there are any differences, the level on the target unit can be higher.
See [“Enable multiplexing setting”](#) on page 382.
- A disk staging storage unit cannot span to another storage unit. Also, a disk staging storage unit is not eligible as a target for disk spanning.
- Disk spanning is not supported on NFS.

Enable standalone drive extension

This property allows NetBackup to use whatever labeled or unlabeled media is found in a nonrobotic drive. The default is that the **Enable standalone drive extension** property is enabled.

Enable job logging

This property allows the logging of the job information. This is the same information that the NetBackup Activity Monitor uses. The default is that job logging occurs.

Enable unrestricted media sharing for all media servers

This property controls media sharing, as follows:

- Enable this property to allow all NetBackup media servers and NDMP hosts in the NetBackup environment to share media for writing. Do not configure server groups for media sharing.
- Clear this property to restrict media sharing to specific server groups. Then configure media server groups and backup policies to use media sharing.
- Clear this property to disable media sharing. Do not configure media server groups.

The default is that media sharing is disabled. (The property is cleared and no server groups are configured.)

See [“About server groups”](#) on page 197.

Media ID prefix (non-robotic)

This property specifies the media ID prefix to use in media IDs when the unlabeled media is in nonrobotic drives. The prefix must be one to three alpha-numeric characters. NetBackup appends numeric characters. By default, NetBackup uses A and assigns media IDs such as A00000, A00001, and so on.

For example, if FEB is specified, NetBackup appends the remaining numeric characters. The assigned media IDs become FEB000, FEB001, and so on. (Note that this numbering does not work with the Configure Volumes wizard).

Media unmount delay

To specify a **Media unmount delay** property indicates that the unloading of media is delayed after the requested operation is complete. Media unmount delay applies only to user operations, to include backups and restores of database agent clients, such as those running NetBackup for Oracle. The delay reduces unnecessary media unmounts and the positioning of media in cases where the media is requested again a short time later.

The delay can range from 0 seconds to 1800 seconds. The default is 180 seconds. If you specify 0, the media unmount occurs immediately upon completion of the requested operation. Values greater than 1800 are set to 1800.

Media request delay

This property specifies how long NetBackup waits for media in nonrobotic drives. A configurable delay is useful if a gravity feed stacker is used on a nonrobotic

drive. A delay often exists between dismounting one media and mounting another. The default is 0 seconds.

During the delay period, NetBackup checks every 60 seconds to see if the drive is ready. If the drive is ready, NetBackup uses it. Otherwise, NetBackup waits another 60 seconds and checks again. If the total delay is not a multiple of 60, the last wait is the remainder. If the delay is less than 60 seconds, NetBackup checks after the end of the delay.

For example, set the delay to 150 seconds. NetBackup waits 60 seconds, checks for ready, waits 60 seconds, checks for ready, waits 30 seconds, and checks for ready the last time. If the delay was 50 seconds (a short delay is not recommended), NetBackup checks after 50 seconds.

Recommended use for Enable SCSI reserve property

All tape drive and bridge vendors support the SPC-2 SCSI reserve and release method. NetBackup has used SPC-2 SCSI reserve since NetBackup 3.4.3, and it is the default tape drive reservation method in NetBackup. SPC-2 SCSI reserve is effective for most NetBackup environments.

The SCSI persistent reserve method may be more effective in the following environments because it provides device status detection and correction:

- NetBackup media servers operate in a cluster environment.
NetBackup can recover and use a reserved drive after a failover (if NetBackup owns the reservation). (With SPC-2 SCSI reserve, the drive must usually be reset because the reservation owner is inoperative.)
- The drive must have high availability.
NetBackup can resolve NetBackup drive reservation conflicts and maintain high drive availability. (SPC-2 SCSI reserve provides no method for drive status detection.)

However, the SCSI persistent reserve method is not supported or not supported correctly by all device vendors. Therefore, thoroughly analyze the environment to ensure that all of the hardware supports SCSI persistent reserve correctly.

Symantec recommends careful consideration of all of the following factors before **Enable SCSI reserve** is used:

- Only a limited number of tape drive vendors support SCSI persistent reserve.
- SCSI persistent reserve is not supported or not supported correctly by all fibre channel bridge vendors. Incorrect support in a bridge means no access protection. Therefore, if the environment uses bridges, do not use SCSI persistent reserve.

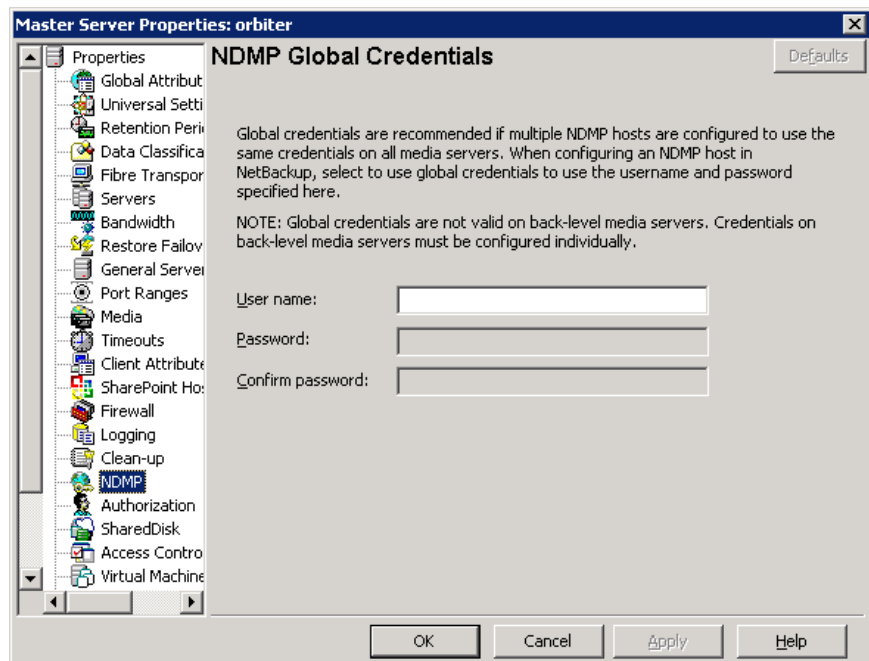
- If parallel SCSI buses are used, carefully consider the use of SCSI persistent reserve. Usually, parallel drives are not shared, so SCSI persistent reserve protection is not required. Also, parallel drives are usually on a bridge, and bridges do not support SCSI persistent reserve correctly. Therefore, if the environment uses parallel SCSI buses, do not use SCSI persistent reserve.
- The operating system tape drivers may require extensive configuration to use SCSI persistent reserve. For example, if the tape drives do not support SPC-3 Compatible Reservation Handling (CRH), ensure that the operating system does not issue SPC-2 reserve and release commands.

If any of the hardware does not support SCSI persistent reserve, Symantec recommends that SCSI persistent reserve is not used.

NDMP Global Credentials properties

The credentials that are entered for **NDMP Global Credentials** can apply to any NDMP host in the configuration. However, the **Use global NDMP credentials for this NDMP host** option must be selected in the **Add NDMP Host** dialog box for the NDMP host.

Figure 3-41 NDMP Global Credentials dialog box



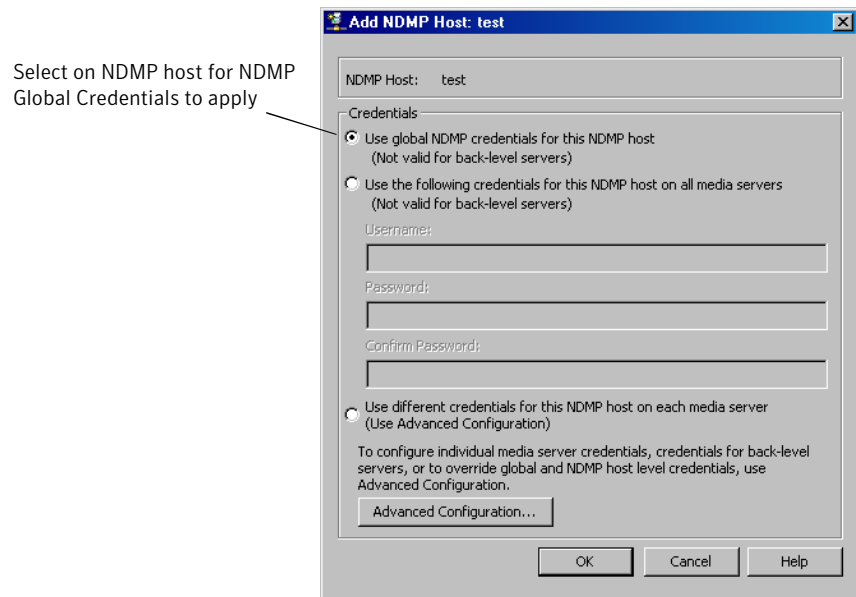
The **NDMP Global Credentials** properties dialog box contains the following options:

- User name** The user name under which NetBackup accesses the NDMP server. This user must have permission to run NDMP commands.
- Password** Enter the password.
- Confirm password** Re-enter the password.

To access the Add NDMP Host dialog box, add an NDMP host under **Media and Device Management > Credentials > NDMP Hosts**.

[Figure 3-42](#) shows the Add NDMP Host dialog box and the option to select on NDMP hosts for NDMP Global Credentials to apply.

Figure 3-42 Add NDMP Host dialog box



NetWare Client properties

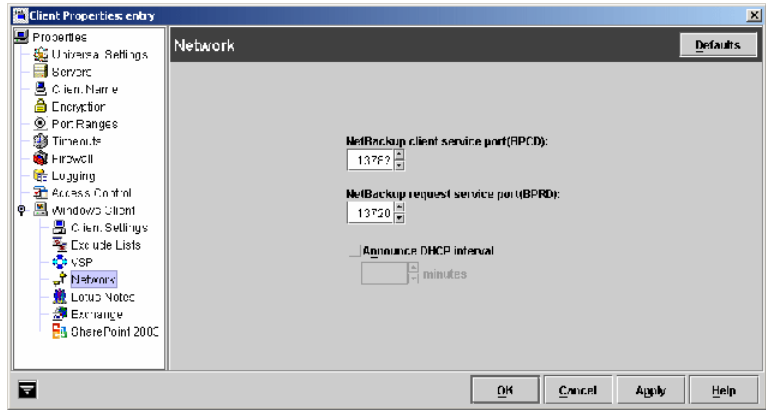
The Netware Client properties define NetBackup properties of NetWare clients. Netware Client properties include the Client Settings for NetWare clients as a subnode:

See [“Client Settings \(NetWare\) properties”](#) on page 90.

Network properties

Use the Network properties to set the properties that define requirements for communications between clients and the master server. The Network properties apply to currently selected Windows clients.

Figure 3-43 Network dialog box



The following topics describe the **Network** host properties.

NetBackup client service port (BPCD)

This property specifies the port that the NetBackup client uses to communicate with the NetBackup server. The default is 13782.

Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.

NetBackup request service port (BPRD)

This property specifies the port for the client to use when it sends requests to the NetBackup request service (`bprd` process) on the NetBackup server. The default is 13720.

Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.

Announce DHCP interval

This property specifies how many minutes the client waits before it announces that a different IP address is to be used. The announcement occurs only if the specified time period has elapsed and the address has changed since the last time the client announced it.

Network Settings Properties

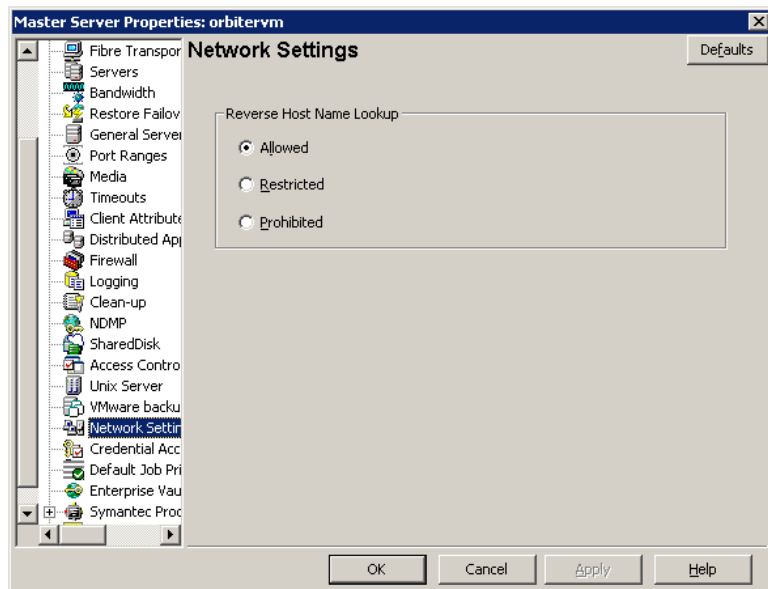
The **Reverse Host Name Lookup** property in the Network Settings host properties applies to master servers, media servers, and clients.

The domain name system (DNS) reverse host name lookup is used to determine what host and domain name a given IP address indicates.

Some administrators cannot or do not want to configure the DNS server for reverse host name lookup. For these environments, NetBackup offers the **Reverse Host Name Lookup** property to allow, restrict, or prohibit reverse host name lookup.

Administrators can configure the **Reverse Host Name Lookup** property for each host.

Figure 3-44 Network Settings dialog box



The following topics describe the **Reverse Host Name Lookup** host properties.

Allowed setting

The **Allowed** setting indicates that the host requires reverse host name lookup to work to determine that the connection comes from a recognizable server.

By default, the host resolves the IP address of the connecting server to a host name by performing a reverse lookup.

If the conversion of the IP address to host name fails, the connection fails.

Otherwise, it compares the host name to the list of known server host names. If the comparison fails, the host rejects the server and the connection fails.

Restricted setting

The **Restricted** setting indicates that the NetBackup host first attempts to perform reverse host name lookup. If the NetBackup host successfully resolves the IP address of the connecting server to a host name (reverse lookup is successful), it compares the host name to the list of known server host names.

If the resolution of the IP address to a host name fails (reverse lookup fails), based on the **Restricted** setting, the host converts the host names of the known server list to IP addresses (using a forward lookup). The host compares the IP address of the connecting server to the list of known server IP addresses.

If the comparison fails, the host rejects the connection from server and the connection fails.

Prohibited setting

The **Prohibited** setting indicates that the NetBackup host does not try reverse host name lookup at all. The host resolves the host names of the known server list to IP addresses using forward lookups.

The NetBackup host then compares the IP address of the connecting server to the list of known server IP addresses.

If the comparison fails, the NetBackup host rejects the connection from the server and the connection fails.

Changing host properties outside of the Administration Console

In some cases, a master server may not be able to view the host properties of a media server or client in the Administration Console. The NetBackup customer's DNS reverse host name lookup configuration may be one possible reason why the host properties may not be visible.

In this case, since changing the NetBackup **Reverse Host Name Lookup** host property involves being able to view the host properties, you'll need to use another method to change the **Reverse Host Name Lookup** host property.

Use the methods that are described in the following sections to add the `REVERSE_NAME_LOOKUP` entry to the `bp.conf` file (UNIX) or to the Windows registry.

REVERSE_NAME_LOOKUP entry

The `REVERSE_NAME_LOOKUP` entry uses the following format:

```
REVERSE_NAME_LOOKUP = ALLOWED | RESTRICTED | PROHIBITED
```

For example:

```
REVERSE_NAME_LOOKUP = PROHIBITED
```

The values of `ALLOWED`, `RESTRICTED`, and `PROHIBITED` represent the same meaning as the values in the **Network Settings** host properties.

Setting the property on UNIX hosts

To set the **Reverse Host Name Lookup** property on a UNIX system outside of the Administration Console, manually add the `REVERSE_NAME_LOOKUP` entry to the `bp.conf` file on the master server, media server, or client.

To add the `REVERSE_NAME_LOOKUP` entry to the `bp.conf` file, use one of the following methods:

- On master and media servers
Use the `bpsetconfig` command to add the entry. The `bpsetconfig` command is described in *NetBackup Commands*.
- On UNIX clients
Edit the `bp.conf` directly to add the entry.

Setting the property on Windows hosts

On master and media servers, the `bpsetconfig` command is available to add the `REVERSE_NAME_LOOKUP` entry to the registry. The `bpsetconfig` command is described in *NetBackup Commands*.

To set the **Reverse Host Name Lookup** property on a Windows client, add the `REVERSE_NAME_LOOKUP` entry to the registry using the following method.

Port Ranges properties

NetBackup communicates between computers by using a combination of registered and dynamically allocated ports. Use the Port Ranges properties to determine how hosts connect to one another. The Port Ranges properties apply to selected master servers, media servers, and clients.

The following items describe the different types of ports:

- Registered ports are registered with the Internet Assigned Numbers Authority (IANA) and are permanently assigned to specific NetBackup services. For example, the port for the NetBackup client service (`bpcd`) is 13782. These ports are specified in a system configuration file:

```
%systemroot%\system32\drivers\etc\services
```

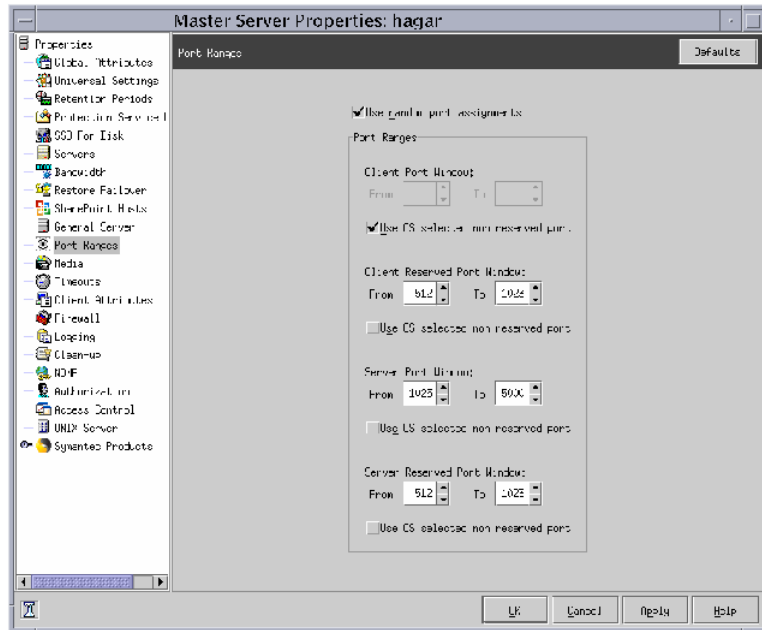
Media Manager services include tape library control daemons, which accept connections from daemons on other servers that share the same library. See the `services` file on the media server to determine the ports that are required for a specific library.

- Dynamically-allocated ports are assigned as needed, from configurable ranges in the Port Ranges host properties for NetBackup servers and clients.

In addition to the range of numbers, you can configure the following for dynamically-allocated ports:

- Whether NetBackup selects a port number at random or starts at the top of the range and uses the first one available.
- Whether connections to `bpcd` on a client use reserved or non-reserved ports.

Figure 3-45 Port Ranges dialog box



The following topics describe the **Port Ranges** host properties.

Use random port assignments

This property specifies how NetBackup communicates with NetBackup on other computers. This property specifies that NetBackup randomly chooses a port from those that are free in the allowed range. For example, if the range is from 1023 through 5000, it chooses randomly from the numbers in this range.

By default, **Use random port assignments** is selected, and ports are randomly chosen.

If this property is not selected, NetBackup chooses numbers sequentially. NetBackup starts with the highest number that is available in the allowed range. For example, if the range is from 1023 through 5000, NetBackup chooses 5000 (if port 5000 is free). If 5000 is in use, port 4999 is chosen.

Client port window

This property specifies the range of non-reserved ports on this computer that are used to connect to NetBackup on other computers. This setting applies to the

selected host that connects to a client that is configured to accept only non-reserved ports.

That means that the **Allow non reserved ports** property is enabled in the Universal Settings dialog box.

See “[Universal Settings properties](#)” on page 185.

To let the operating system determine the non-reserved port to use, enable **Use OS selected non reserved port**.

Client reserved port window

This property specifies the range of reserved ports on this computer that are used for connecting to NetBackup on other computers. This setting applies to the selected host that that is configured to accept only reserved ports.

That means that the **Allow non reserved ports** property is enabled in the Universal Settings dialog box.

See “[Universal Settings properties](#)” on page 185.

The default range is 512 through 1023.

To let the operating system determine the non-reserved port to use, enable **Use OS selected non reserved port**.

Server port window

This property specifies the range of non-reserved ports on which this computer accepts connections from NetBackup on other computers. This setting applies to the selected host that connects to a client that is configured to accept only non-reserved ports.

That means that the **Accept connections on non reserved ports** property is enabled in the Universal Settings dialog box.

See “[Universal Settings properties](#)” on page 185.

Server port window does not appear in a client configuration.

The default range is 1024 through 5000.

To let the operating system determine the non-reserved port to use, enable **Use OS selected non reserved port**.

Server reserved port window

This setting specifies the range of local reserved ports on which this host accepts connections from NetBackup on other hosts. This setting applies to the selected host that connects to a client that is configured to accept only reserved ports.

That means that the **Allow non reserved ports** property is enabled in the Universal Settings dialog box.

See “[Universal Settings properties](#)” on page 185.

The default range is 512 through 1023.

This property does not appear in a client configuration.

To let the operating system determine the non-reserved port to use, enable **Use OS selected non reserved port**.

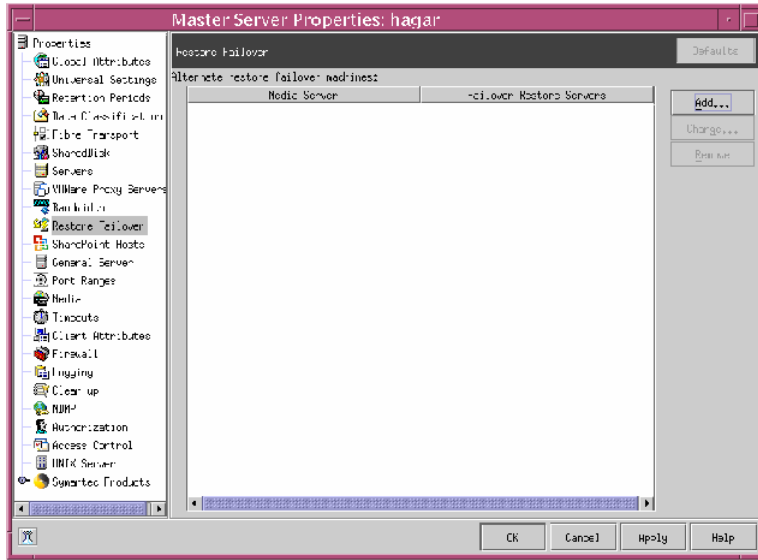
Restore Failover properties

The Restore Failover properties control how NetBackup performs automatic failover to a NetBackup media server. A failover server may be necessary if the regular media server is temporarily inaccessible to perform a restore operation.

The automatic failover does not require administrator intervention. By default, NetBackup does not perform automatic failover.

The **Restore Failover** properties apply to selected master servers.

Figure 3-46 Restore Failover dialog box



The following situations describe examples of when to use the restore failover capability:

- Two or more media servers share a robot and each has connected drives. When a restore is requested, one of the servers is temporarily inaccessible.
- Two or more media servers have stand alone drives of the same type. When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the master server and `bptm` on the media server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The media server is down.
- The media server is up but `bpcd` does not respond. (For example, if the connection is refused or access is denied.)
- The media server is up and `bpcd` is running, but `bptm` has problems. (For example, or `bptm` cannot find the required tape.)

The **Media server** column displays the NetBackup media servers that have failover protection for restores. The Failover restore server column displays the servers that provide the failover protection. NetBackup searches from top to bottom in the **Failover restore server** column until it finds another server that can perform the restore.

A NetBackup media server can appear only once in the **Media server** column but can be a failover server for multiple other media servers. The protected server and the failover server must both be in the same master and media server cluster.


Adding failover servers

To include a NetBackup media server in the **Alternate restore failover machines** list, click **Add**.

Use the following procedure to add or change a media server to the alternate restore failover machine list in the **Restore Failover** host properties.

To add or change a media server to the failover server

- 1 To add an entry, click **Add**. To change an entry, click **Change**.



- 2 In the **Media server** field, specify the media server for failover protection.
- 3 In the **Failover restore servers** field, specify the media server(s) to try if the server that is designated in the **Server** field is unavailable. Separate the names of multiple servers with a single space.
- 4 Click **Add** to add the name to the list. The dialog box remains open for another entry.
Click **OK** if an entry was changed. Click **Close** to close the dialog box.
- 5 Click **Apply** to accept the Restore Failover property changes. Click **OK** to close the host properties dialog box.
- 6 Stop and restart the NetBackup Request daemon on the master server where the configuration was changed.

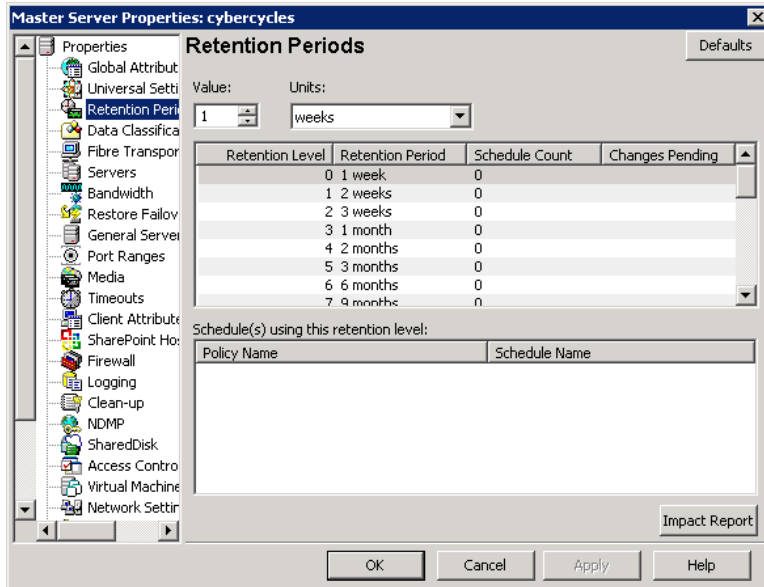
See [“Enabling automatic failover to an alternate server”](#) on page 769.

Retention Periods properties

Use the **Retention Periods** properties to define a duration for each retention level. Retention periods are specified in policy schedules to determine how long NetBackup retains the backups or archives.

The **Retention Periods** properties apply to selected master servers.

Figure 3-47 Retention Periods dialog box



By default, NetBackup stores each backup on a volume that already contains backups at the same retention level. However, NetBackup does not check the retention period that is defined for that level. When the retention period for a level is redefined, it can result in some backups that have different retention periods sharing the same volume.

For example, if the retention level 3 is changed from 1 month to 6 months, NetBackup stores future level 3 backups on the same volumes. That is, the backups are placed on the volumes with the level 3 backups that have a retention period of one month.

No problem exists if the new and the old retention periods are of about the same value. However, before a major change is made to a retention period, suspend the volumes that were previously used for that retention level.

See [“Suspending volumes”](#) on page 175.

The following topics describe the **Retention Periods** host properties.

Value

This **Value** property assigns a number to the retention level setting.

Units

The **Units** property specifies the units of time for the retention period. The list includes hours as the smallest unit of granularity and the special units, **Infinite**, and **Expires immediately**.

Retention periods list

The **Retention periods list** displays a list of the current definitions for the 25 possible levels of retention (0 through 24). By default, levels 9 through 24 are set to infinite. Retention level 9 is the only level that cannot be changed and remains at infinite.

Note that if left at the default, there is no difference between a retention level of 12 and a retention level of 20, for example.

The **Schedule Count** column indicates how many schedules currently use each level. If the retention period is changed for a level, it affects all schedules that use that level.

The **Changes Pending** column uses an asterisk (*) to indicate that the period has been changed and not applied. NetBackup does not change the actual configuration until the administrator clicks **Apply** or **OK**.

Schedules list

The **Schedules list** is a list of the schedules that use the currently selected retention level, and the policy to which each schedule belongs.

Impact Report button

Click the **Impact Report** button to display a summary of how changes affect existing schedules. The list displays all schedules in which the retention period is shorter than the frequency period.

Changing a retention period

Use the following procedure to change a retention period.

To change a retention period

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Server** > Double-click on master server > **Retention Periods**.
- 2 Select the retention level to change.

Note: Level 9 cannot be changed and remains at a setting of infinite.

By default, levels 9 through 24 are set to infinite. If the levels are left at the default, there is no difference between a retention level of 12 and a retention level of 20, for example.

The policy impact list now displays the names of all schedules that use the selected retention level. It also lists the policy to which each schedule belongs.

- 3 Type the new retention period in the **Value** box.
- 4 Select the units of measure (days, weeks, months, years, infinite, or expires immediately).

After **Units** or **Value** is changed, an asterisk (*) appears in the Changes Pending column to indicate that the period was changed. NetBackup does not change the actual configuration until the administrator clicks **Apply** or **OK**.

- 5 Click **Impact Report**.

The policy impact list displays the schedules where the new retention period is less than the frequency period.

To prevent schedules from being listed, redefine the retention period for the schedules or change the retention or frequency for the schedule.

- 6 To discard your changes, click **Cancel**.
- 7 To save your changes and update the configuration, click one of the following:
 - **Apply**
Saves the changes and leaves the dialog box open to make further changes.
 - **OK**
Saves the changes since **Apply** was last clicked. **OK** also closes the dialog box.
- 8 To save the changes, click **OK**.

See “[Retention Periods properties](#)” on page 172.

Suspending volumes

Use the following procedure to suspend volumes.

To suspend volumes

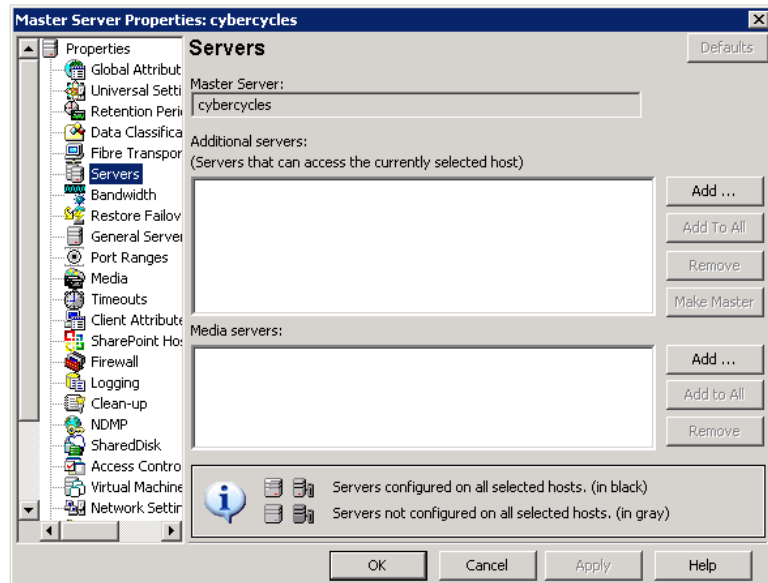
- 1 Use the NetBackup Media List report to determine which volumes are currently at the level to be suspended.
- 2 Use the `bpmedia` command to suspend the volumes.

```
bpmedia -suspend -m media_ID
```

Servers properties

The **Servers** properties display the NetBackup server list on selected master servers, media servers, and clients. The server list displays the NetBackup servers that each host recognizes.

Figure 3-48 Servers dialog box



The following topics describe the **Servers** host properties.

Master server

The **Master Server** property specifies the master server for the selected host. (The name of the selected host appears in the title bar.)

Additional servers list

This list contains the additional servers that can access the server that is specified as **Master server**.

During installation, NetBackup sets the master server to the name of the system where the server software is installed. NetBackup uses the master server value to validate server access to the client. The master server value is also used to determine which server the client must connect to so that files can be listed and restored.

To make changes, do the following:

- To add a server, click **Add** and select a server.
- To delete a server, select a server from the list and click **Remove**.
- To change the master server, select another server from the list, then click **Make Master**.

To configure access to a remote server, add to the server list the name of the host seeking access.

See [“Accessing remote servers”](#) on page 727.

Media servers list

This list specifies that the hosts that are listed are media servers only. Hosts that are listed as media servers can back up and restore clients, but have limited administrative privileges.

Note: If you change the server list on the master server, exit all NetBackup administrator interface programs. Then stop and restart both the NetBackup request service and NetBackup Database Manager service on that server. Restarting the services ensures that the change is recognized.

To add or delete servers, do the following:

- To add a new media server, click **Add** and select a server.
Run `nbemmcmd -addhost` to add a media server to the Enterprise Media Manager (EMM) database of existing master server.
- To delete a media server, select a media server from the list and click **Remove**.

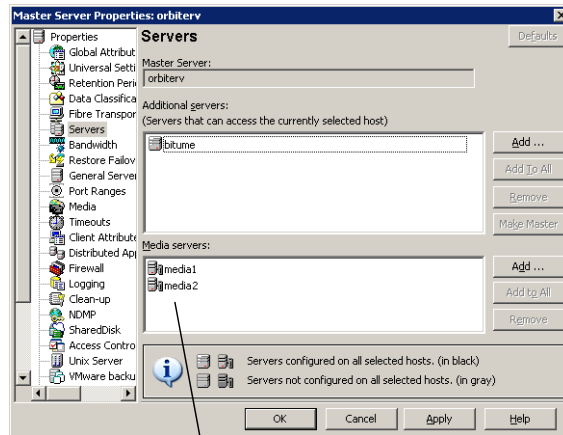
Restricting administrative privileges of media servers

The servers in the **Media servers** list are media servers only. (**Host Properties > Master Server** or **Media Servers > Servers**.)

Computers that are listed as media servers can back up and restore clients, but have limited administrative privileges.

[Figure 3-49](#) shows the **Media server** list in the **Servers** dialog box.

Figure 3-49 Media server list in Servers dialog box



Administrative scope of media servers is limited

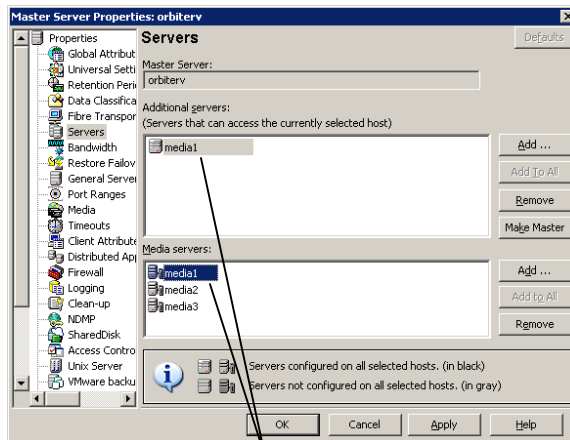
Consider a configuration that consists of master server orbiterv and two media servers—media1 and media2. Set up orbiterv as the master server and media1 and media2 as media servers.

If a computer is defined as both a master server and a media server, the master server entry takes precedence.

A server that is listed as both a master and a media server has consequences. It allows a media server system administrator to be a NetBackup administrator on other master servers.

[Figure 3-50](#) shows a server that is both a media server and a master server.

Figure 3-50 Server that is listed as both media server and master server



A machine that is listed as both an additional server and a media server has full administrative privileges

Multiple masters that share one Enterprise Media Manager host

Multiple master servers can share one EMM database that is located on a single host. The host that contains the EMM database can be either a master server or a media server.

The **Servers** host properties must be set up to allow multiple master servers to access the EMM host.

In the following example, three master servers share one EMM database that is located on one of the servers (meadow).

Table 3-2 describes the `bp.conf` server entries on each master server in the example:

Table 3-2 Server entries example

Meadow	Havarti	Study
SERVER = meadow	SERVER = havarti	SERVER = study
SERVER = havarti	SERVER = meadow	SERVER = meadow
SERVER = study	CLIENT_NAME = havarti	CLIENT_NAME = study
CLIENT_NAME = meadow	EMMSERVER = meadow	EMMSERVER = meadow

Table 3-2 Server entries example (*continued*)

Meadow	Havarti	Study
EMMSERVER = meadow		

SERVER entries are as follows:

- The first SERVER entry must be the name of the master server.
- In order for the NetBackup Administration Console to administer other servers, the servers must be listed. (**File > Change Server.**)
- If the EMM database is on another master server, that server needs to be listed. The table shows how meadow is listed on havarti and study.

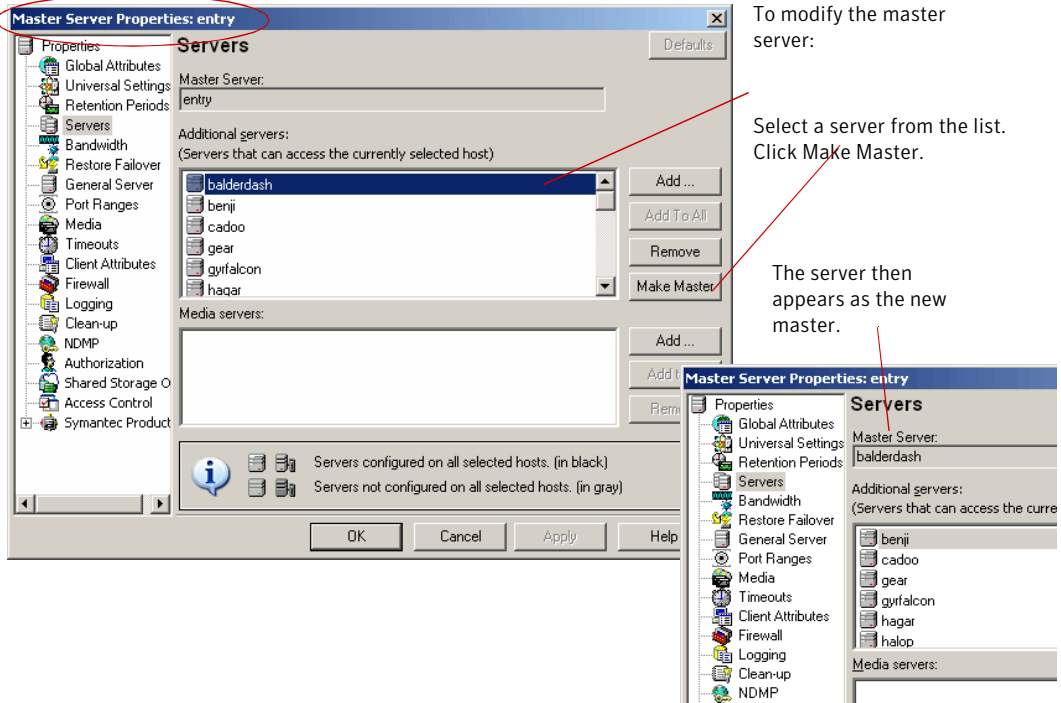
The EMMSERVER entry must be present on all master servers that share the EMM host. The table shows how meadow is listed as the EMMSERVER on havarti, study, as well as on meadow.

[Figure 3-51](#) shows a shared EMM database that is located on a media server.

Figure 3-51 A shared EMM database

If the master server is changed on a media server, the EMM database also needs to be updated.

1



2

To update the EMM database, after changing the master server for a media server, run:

```
install_path \VERITAS\NetBackup\bin\admincmd\nbemcmd -updatehost
```

SharedDisk properties

The SharedDisk master server properties specify the SharedDisk storage option properties for a NetBackup configuration.

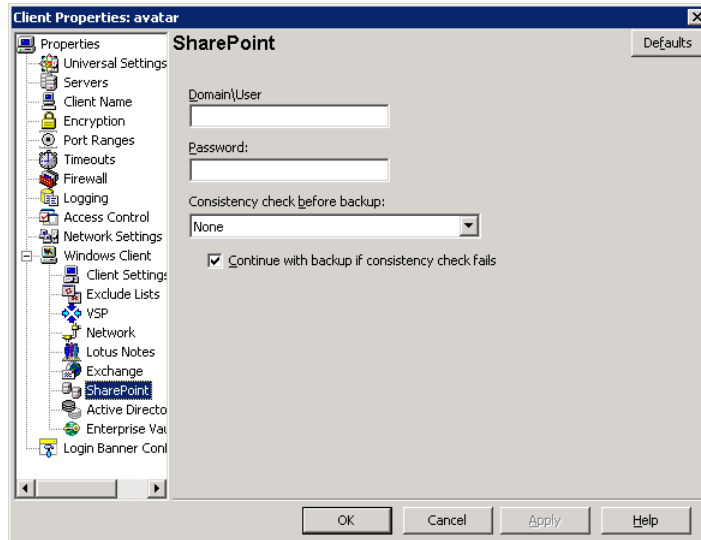
See [“About SharedDisk support in NetBackup 7.0 and later”](#) on page 362.

SharePoint properties

The SharePoint properties apply to currently selected Windows clients to protect SharePoint Server installations.

For complete information on these options, see the *NetBackup for Microsoft SharePoint Server Administrator's Guide* .

Figure 3-52 SharePoint dialog box



- Domain\User** Specify the domain and user name for the account you want to use to log on to SharePoint (DOMAIN\user name).

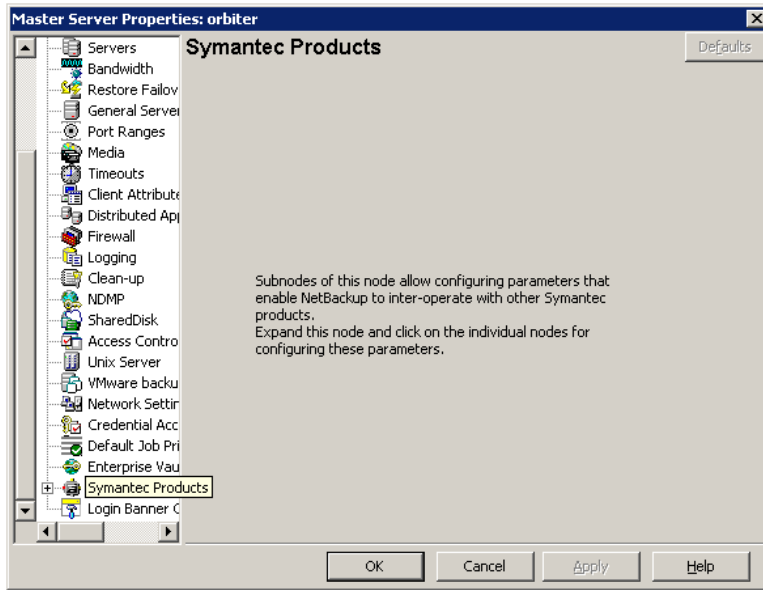
- Password** Specify the password for the account.

- Consistency check before backup** Select the consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation. These checks are performed for both server-directed and user-directed backups.

Symantec Products properties

The Symantec Products properties apply to currently selected master servers.

Figure 3-53 Symantec Products dialog box



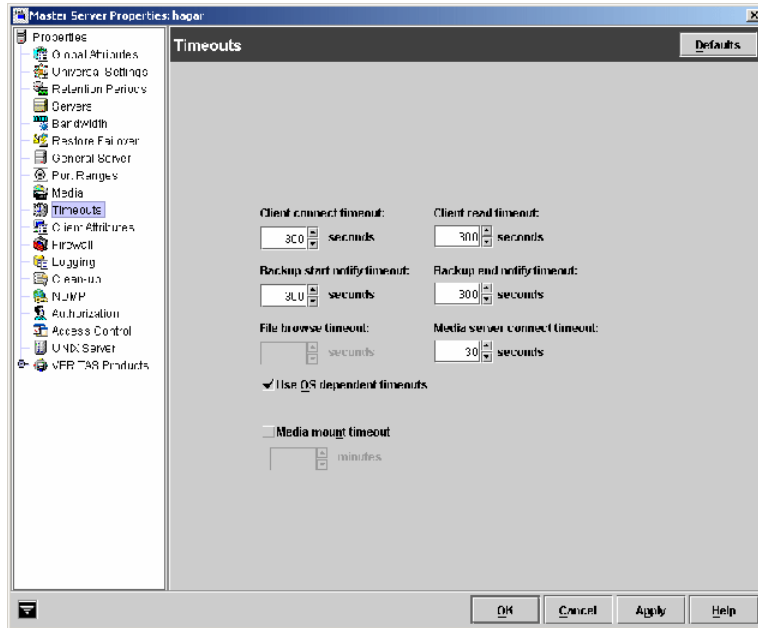
The Symantec Products properties include the subnode, Backup Exec Tape Reader properties.

See [“Backup Exec Tape Reader properties”](#) on page 69.

Timeouts properties

The Timeouts properties apply to selected master servers, media servers, and clients.

Figure 3-54 Timeouts dialog box



The following topics describe the Timeouts properties.

Client connect timeout

This property specifies the number of seconds the server waits before it times out when it connects to a client. The default is 300 seconds.

Backup start notify timeout

This property specifies the number of seconds the server waits for the `bpstart_notify` script on a client to complete. The default is 300 seconds.

Note: If this timeout is changed, verify that **Client read timeout** is set to the same or higher value.

File browse timeout

This property specifies how long the client can wait for a response from the NetBackup master server while it lists files.

Note: If it exists, the value in a UNIX client's `$HOME/bp.conf` file takes precedence to the property here.

If the **File browse timeout** property is exceeded, the user receives a socket read failed error. The timeout can be exceeded even while the server processes the request.

Use OS dependent timeouts

This property specifies that the client waits for the timeout period as determined by the operating system when it lists files, as follows:

- Windows client: 300 seconds
- UNIX client: 1800 seconds

Media mount timeout

This property specifies how long NetBackup waits for the requested media to be mounted, positioned, and ready on backups, restores, and duplications.

The **Media mount timeout** property appears only as a master server property.

Use this timeout to eliminate excessive waiting time during manual media mounts. (For example, when robotic media is out of the robot or is off site.)

Client read timeout

This property specifies the number of seconds to use for the client-read timeout. This timeout can apply to a NetBackup master, remote media server, or database-extension client (such as NetBackup for Oracle). The default is 300 seconds.

The client-read timeout on a database-extension client is a special case. Clients can initially require more time to get ready than other clients. More time is required because database backup utilities frequently start several backup jobs at the same time, slowing the central processing unit.

Note: For database-extension clients, Symantec suggests that the Client read timeout be set to a value greater than 5 minutes. Fifteen minutes are adequate for many installations. For other clients, change **Client read timeout** only if the client encounters problems.

The sequence on a database-extension client is as follows:

- NetBackup on the database-extension client reads the client's client-read timeout to find the initial value. If the option is not set, the standard 5-minute default is used.
- When the database-extension API receives the server's value, it uses it as the client-read timeout.

Backup end notify timeout

This property specifies the number of seconds that the server waits for the `bpend_notify` script on a client to complete. The default is 300 seconds.

Note: If this timeout is changed, verify that **Client read timeout** is set to the same or higher value.

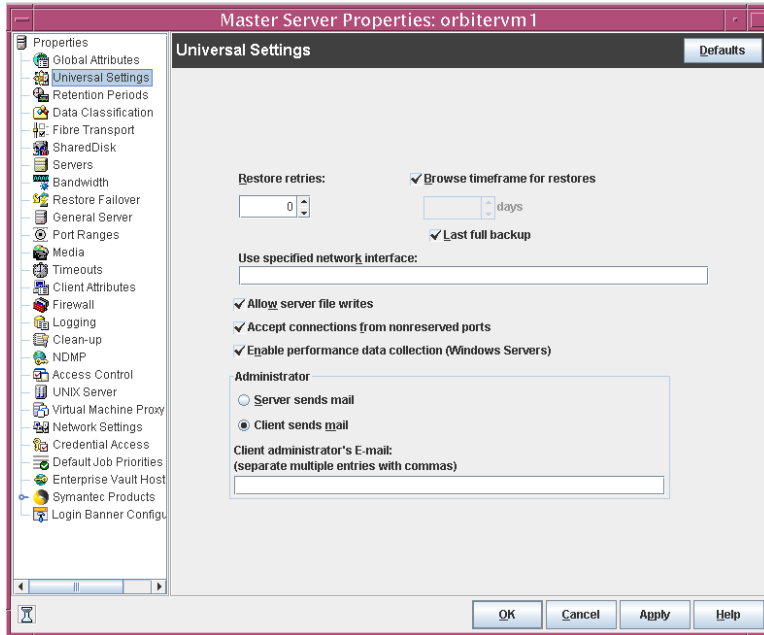
Media server connect timeout

This property specifies the number of seconds that the master server waits before it times out when it connects to a remote media server. The default is 30 seconds.

Universal Settings properties

The Universal Settings properties apply to selected master servers, media servers, and clients.

Figure 3-55 Universal Settings dialog box



The following topics describe the Universal Settings properties.

Restore retries

This property specifies the number of attempts a client has to restore after a failure. (The default is 0; the client does not attempt to retry a restore. The client can try up to three times.) Change **Restore retries** only if problems are encountered.

If a job fails after the maximum number of retries, the job goes into an incomplete state. The job remains in the incomplete state as determined by the **Move restore job from incomplete state to done state** property.

See “[Clean-up properties](#)” on page 76.

A checkpointed job is retried from the start of the last checkpointed file rather than at the beginning of the job.

Checkpoint restart for restores allows a NetBackup administrator to resume a failed restore job from the Activity Monitor.

See “[Checkpoint restart for backup jobs](#)” on page 468.

Browse timeframe for restores

This property specifies how long ago NetBackup searches for files to restore. For example, to limit the browse range to one week before the current date, clear the **Last full backup** check box and specify 7.

This limit is specified on the master server and applies to all NetBackup clients. A limit can be specified on an individual client to reduce the size of the Search window. The client setting cannot make the browse window larger.

By default, NetBackup includes files from the time of the last-full backup through the latest backup for the client. If the client belongs to more than one policy, then the browse starts with the earliest of the set of last-full backups.

Last full backup

This property indicates whether NetBackup includes all backups since the last successful full backup in its browse range. This property must be disabled to enter a value for the **Browse timeframe for restores** property. The default is that this property is enabled.

Use specified network interface

This property specifies the network interface that NetBackup uses to connect to another NetBackup client or server. A NetBackup client or server can have more than one network interface. To force NetBackup connections to be made on a specific network interface, use this entry to specify the network host name of that interface. By default, the operating system determines the one to use.

See [“Use specified network interface examples”](#) on page 189.

Allow server file writes

This property specifies whether a NetBackup server can create or modify files on the NetBackup client. For example, enable this property to prevent server-directed restores and remote changes to the client properties.

After the **Allow server file writes** property is applied, it can be cleared only by modifying the client configuration. The default is that server writes are allowed.

Accept connections on non reserved ports

This property specifies whether the NetBackup client service (`bpcd`) can accept remote connections from non-reserved ports. (Non reserved ports have port numbers of 1024 or greater.) The default is that this property is enabled.

If this property is enabled, the server that connects to the host must also be configured to use non-reserved ports for the client. Select **Accept connections from non reserved ports** on the server properties Client attributes tab.

If the property is disabled (unchecked), `bpcd` requires remote connections to come from privileged ports. (Privileged ports have port numbers that are less than 1024.) **Accept connections on non reserved ports** is useful when NetBackup clients and servers are on opposite sides of a firewall.

When disabled, the source ports for connections to `bpcd` use reserved ports as well.

Enable performance data collection (Windows server only)

This property specifies whether NetBackup updates disk and tape performance object counters. (Applies only to Windows master and media servers. Use the Windows Performance Monitor utility (`perfmon`) to view the NetBackup performance counters. The default is that this property is enabled.

See the *NetBackup Administration Guide, Volume II* for more information about using the System Monitor with NetBackup.

Client sends mail

This property specifies whether the client sends an email to the address that is specified in the Universal Settings properties. If the client cannot send email, select **Server sends mail**. The default is that this property is enabled.

See “[Global Attributes properties](#)” on page 131.

Server sends mail

This property specifies whether the server sends an email to the address that is specified in the **Global Attributes** properties. Enable this property if the client cannot send mail and you want an email notification. The default is that this property is disabled.

See “[Global Attributes properties](#)” on page 131.

Client administrator’s email

This property specifies the email address of the administrator on the client. This address is where NetBackup sends backup status reports for the client. By default, no email is sent. To enter multiple addresses or email aliases, separate entries with commas.

See “[Global Attributes properties](#)” on page 131.

Use specified network interface examples

The following are examples of situations where a network interface is specified:

■ Example 1 - client with multiple network interfaces

Assume a NetBackup client with two network interfaces, as follows:

- One network interface is for the regular network. The host name for the regular interface is *fred*.
- One network interface is for the backup network. The host name for the backup interface is *fred_nb*.

The NetBackup client name setting on both the client and server is *fred_nb*. When client *fred* starts a backup, restore, or list operation, the request goes out on the *fred_nb* interface and over the backup network. The operation assumes that *fred* and the network are set up to do so. If this configuration is not in place, *fred* can send out the request on the *fred* interface and over the regular network. The server receives the request from client *fred_nb* with host name *fred* and refuses it because the host and the client names do not match. One solution is to set Use specified network interface on *fred* to *fred_nb*. All backup, restore, and list requests use the *fred_nb* interface. The server receives requests from client *fred_nb* with host name *fred_nb* and everything works as intended.

Another solution is to set up the master server to allow redirected restores for client *fred*. Redirected restores allow the server to accept the request, but leaves NetBackup traffic on the regular network.

■ Example 2 - server with multiple network interfaces

Assume a NetBackup server with two network interfaces, as follows:

- One network interface is for the regular network. The host name for the regular interface is *barney*.
- One network interface is for the backup network. The host name for the backup interface is *barney_nb*.

The server list on all NetBackup servers and clients have an entry for *barney_nb*.

When *barney* connects to a client for a backup, the request goes out on the *barney_nb* interface and over the backup network. The operation assumes that *barney* and the network are set up to do so. If this configuration is not in place, *barney* can send out the request on the *barney* interface and over the regular network. The client now receives the request from *barney* rather than *barney_nb* and refuses it as coming from an invalid server.

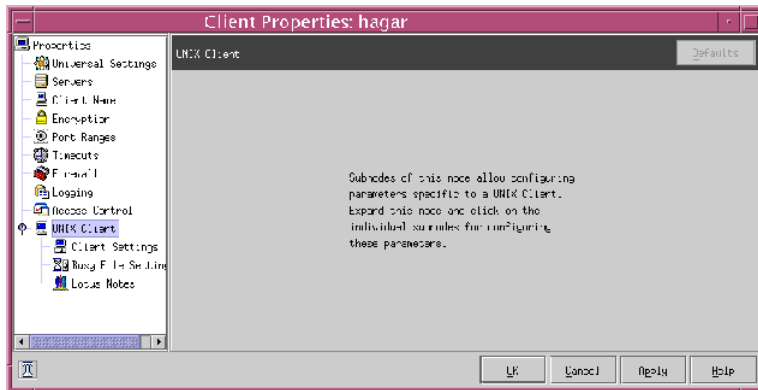
One solution is to set Use specified network interface on barney to barney_nb. Now, when barney connects to a client, the connection is always through the barney_nb interface and everything works as intended.

Another solution is to add an entry for barney to the server list on the client. The client now accepts requests from barney, but NetBackup traffic continues on the regular network.

UNIX Client properties

The UNIX Client properties define NetBackup properties of UNIX clients.

Figure 3-56 UNIX Client dialog box



See “Client Settings (UNIX) properties” on page 91.

See “Busy File Settings properties” on page 73.

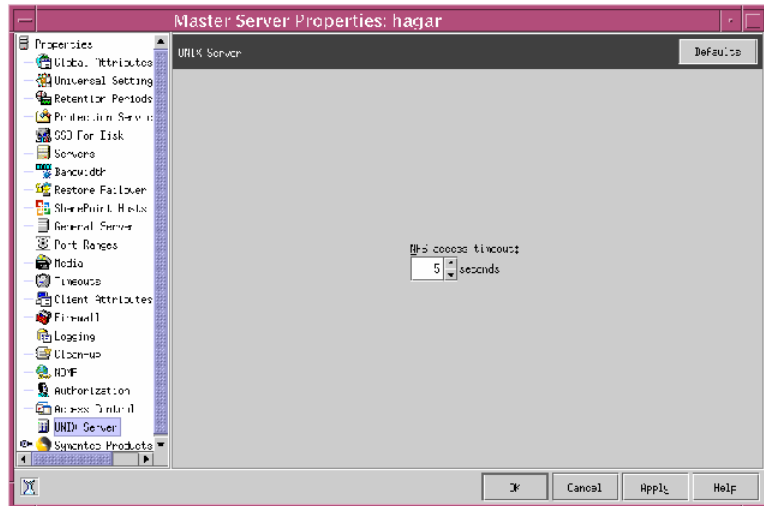
See “Lotus Notes properties” on page 151.

UNIX Server properties

The NFS access timeout property on the UNIX Server properties specifies how long the backup waits to process the mount table before it considers an NFS file system unavailable. The default is 5 seconds.

The UNIX Server properties apply to selected UNIX master servers.

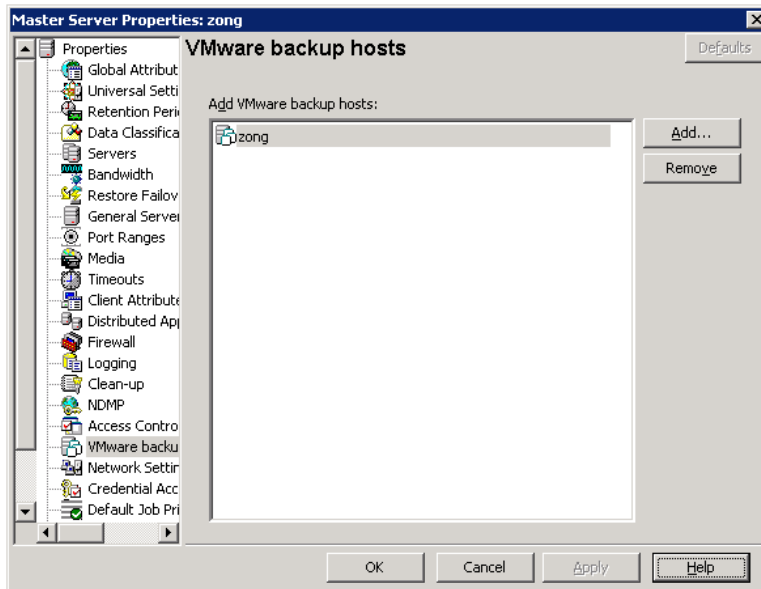
Figure 3-57 UNIX Server dialog box



VMware backup hosts properties

The **VMware backup hosts** properties appear if the NetBackup Enterprise Client license is installed. The properties apply to currently selected master servers.

Figure 3-58 VMware backup hosts dialog box



Click **Add** to add a server to the backup hosts list. To delete a server from the list, select the server and click **Remove**.

A backup host is on the same SAN as a VMware ESX server that can access the snapshot of the virtual machine. A backup host can provide access to the files for third-party backup vendors.

For more information, see the *NetBackup for VMware Administrator's Guide*.

VSP (Volume Snapshot Provider) properties

The Volume Snapshot Provider properties are displayed when the selected client is running NetBackup 6.x. If the client is running NetBackup 7.0, the VSP properties do not appear.

For information about selecting VSP for backlevel and upgraded clients:

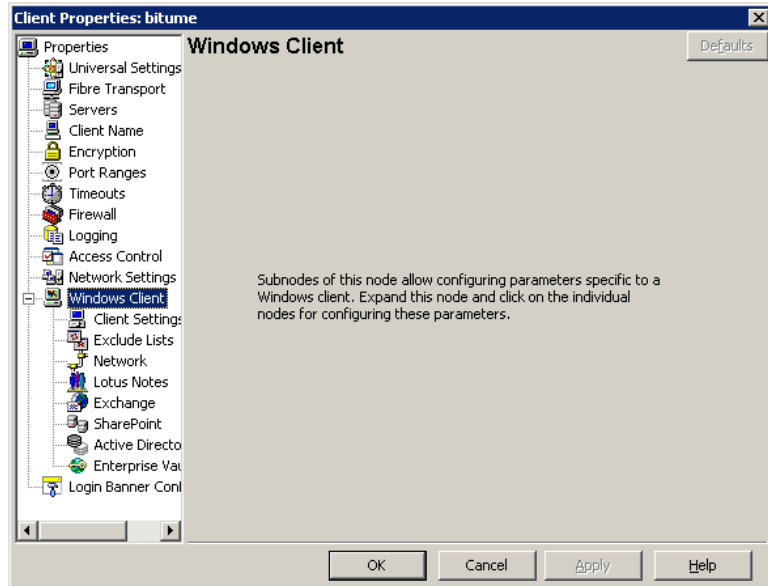
See [“Backlevel and upgraded clients that use Windows Open File Backup”](#) on page 87.

For information about VSP settings, see the *6.5 NetBackup Administrator's Guide, Volume I*.

Windows Client properties

The Windows Client properties define NetBackup properties for Microsoft Windows clients.

Figure 3-59 Windows Client dialog box



Windows Client properties include subnodes for the following host properties:

See [“Client Settings \(Windows\) properties”](#) on page 95.

See [“Exclude Lists properties”](#) on page 114.

See [“Network properties”](#) on page 162.

See [“Lotus Notes properties”](#) on page 151.

See [“Exchange properties”](#) on page 112.

See [“SharePoint properties”](#) on page 180.

See [“Active Directory host properties”](#) on page 66.

See [“Enterprise Vault Hosts properties”](#) on page 112.

Configuration options not found in the Host Properties

To change the default value for an option that is not found in the Host Properties, first use the `bpgetconfig` command to obtain a list of configuration entries. Then use `bpsetconfig` to change the entries as needed.

For information about `bpgetconfig` and `bpsetconfig`, see *NetBackup Commands*.

The following NetBackup administration options cannot be configured by using the NetBackup Administration Console:

- See “[NBRB_CLEANUP_OBSOLETE_DBINFO](#)” on page 194.
- See “[NBRB_ENABLE_OPTIMIZATIONS](#)” on page 194.
- See “[NBRB_FORCE_FULL_EVAL](#)” on page 195.
- See “[NBRB_REEVAL_PENDING](#)” on page 195.
- See “[NBRB_REEVAL_PERIOD](#)” on page 195.
- See “[NBRB_RETRY_DELAY_AFTER_EMM_ERR](#)” on page 195.
- See “[NBRB_MPX_GROUP_UNLOAD_DELAY](#)” on page 195.
- See “[REQUIRED_NETWORK](#)” on page 196.
- The media and device configuration options (`vm.conf` file)

NBRB_CLEANUP_OBSOLETE_DBINFO

The `NBRB_CLEANUP_OBSOLETE_DBINFO` entry serves as a performance tuning option for the Intelligent Resource Manager.

This entry indicates the number of seconds that can elapse between the cleanup of obsolete information in the NetBackup Resource Broker (`nbrb`) database.

The default is 60 seconds.

NBRB_ENABLE_OPTIMIZATIONS

The `NBRB_ENABLE_OPTIMIZATIONS` entry serves as a performance tuning option for the Intelligent Resource Manager. This entry indicates whether the Resource Broker caches states of resource requests. The default is 1 (true).

NBRB_FORCE_FULL_EVAL

The `NBRB_FORCE_FULL_EVAL` entry serves as a performance tuning option for the Intelligent Resource Manager.

This entry indicates the number of seconds that can elapse between full evaluations of all NetBackup Resource Broker (`nbrb`) queues, by using no cached EMM information. For example, full evaluations include matching job resource requests with available resources.

The default is 1800 seconds (30 minutes).

NBRB_REEVAL_PENDING

The `NBRB_REEVAL_PENDING` entry serves as a performance tuning option for the Intelligent Resource Manager. This entry indicates the number of seconds that can elapse between evaluations of the pending request queue. For example, a pending request queue can include, jobs awaiting resources.

The default is 60 seconds.

NBRB_REEVAL_PERIOD

The `NBRB_REEVAL_PERIOD` entry serves as a performance tuning option for the Intelligent Resource Manager and NetBackup Resource Broker (`nbrb`). `NBRB_REEVAL_PERIOD` indicates the time between evaluations if an outstanding request is not satisfied, and if no other requests or resources have been released.

The default is that 5 minutes pass before the initial request is reevaluated.

NBRB_RETRY_DELAY_AFTER_EMM_ERR

The `NBRB_RETRY_DELAY_AFTER_EMM_ERR` entry serves as a performance tuning option for the Intelligent Resource Manager. This entry indicates how long NetBackup waits after an EMM error before it tries again. The error must be one where a retry is possible. For example, if a media server is down.

The default is 60 seconds.

NBRB_MPX_GROUP_UNLOAD_DELAY

The `NBRB_MPX_GROUP_UNLOAD_DELAY` entry serves as a performance tuning option for the Intelligent Resource Manager.

This entry indicates the number of seconds that the NetBackup Resource Broker (`nbrb`) waits for a new job to appear before a tape is unloaded. This setting can help avoid unnecessary reloading of tapes and applies to all backup jobs.

The default is 10 seconds.

During user backups, `nbrb` uses the maximum value of `NBRE_MPX_GROUP_UNLOAD_DELAY` and the **Media mount timeout** host property setting when `nbrb` unmounts the tape.

This host property is found in the NetBackup Administration Console under **NetBackup Management > Host Properties > select master server > Timeouts > Media mount timeout**.

See [“Timeouts properties”](#) on page 182.

During restores, **Media mount timeout** is used, not `NBRE_MPX_GROUP_UNLOAD_DELAY`.

REQUIRED_NETWORK

The `REQUIRED_NETWORK` entry specifies the required route for backup traffic in an environment where the network traffic is segregated.

For example, an environment can contain a production network at `145.21.14.0` and a backup network at `192.132.28.0`.

To indicate that NetBackup use only the backup network, use `bpsetconfig` to add the following entry to the registry:

```
REQUIRED_NETWORK = 192.132.28.0
```

Note: If the variable is set and the network is not available, all connections fail and no backups are performed.

Configuring server groups

This chapter includes the following topics:

- [About server groups](#)
- [Configuring a server group](#)
- [Deleting a server group](#)

About server groups

A server group is a group of NetBackup servers that are used for a common purpose.

A media sharing group is a server group that shares media for write purposes (backups).

A media sharing group can contain the following:

- NetBackup master server
- NetBackup media servers
- NDMP tape servers
- Virtual host names of NetBackup media servers in a cluster

Servers can be in more than one group. All members of a server group must have the same NetBackup master server. Only NetBackup 6.5 and later systems can be in server groups.

See [“Configuring a server group”](#) on page 198.

See [“About media sharing”](#) on page 297.

See [“Configuring media sharing with a server group”](#) on page 299.

Configuring a server group

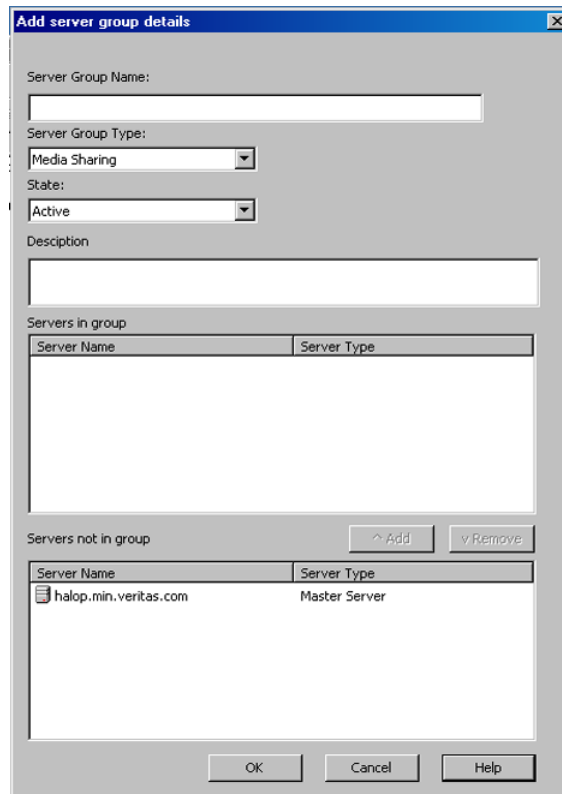
Use the following procedure to configure a server group.

Note: NetBackup allows a server group name to be the same as the name of a media server. However, Symantec recommends that you do not use the same name for a server group and a media server. It may be confusing to use the same name for a media server and a media server group.

See “[About server groups](#)” on page 197.

To configure a server group

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Server Groups**.
- 2 In the **Actions** menu, select **New Server Group**.



- 3 In the **Add Server Group Details** dialog box, enter or select the appropriate information.

See [“Server group properties”](#) on page 200.

To add a server to the group, select it in the **Servers Not in Group** window and click **Add**.

To remove a server from the group, select it in the **Servers in Group** window and click **Remove**.

Server group properties

The server group properties includes the following options:

Table 4-1

Property	Description
Server group name	<p>The name of the server group.</p> <p>You cannot change the name of an existing server group.</p> <p>Symantec recommends that server group names be unique. That is, do not use the same name for a server group that you use for a host such as a media server. If you do, you may not be able to determine easily if a tape is restricted to a specific media server or to a specific media server group.</p>
Server group type	<p>The type of server group.</p> <p>See “About server groups” on page 197.</p> <p>Other server group types (such as Alternate Restore) are reserved for future use.</p>
State	<p>Specify the state of the server group:</p> <ul style="list-style-type: none"> ■ Active. The server group is available for use. ■ Inactive. The server group is not available for use. <p>To change the state, select the new state from the dropdown box.</p>
Description	A description of the media server group.
Servers in group	The servers (and the server type) that belong to the group.
Servers not in group	The servers (and the server type) that do not belong to the group.

Deleting a server group

Use the following procedure to delete a server group.

See [“About server groups”](#) on page 197.

To delete a server group

- 1 In the NetBackup Administration Console, select **Media and Device Management > Devices > Server Groups**.
- 2 Select the group you want to delete.
- 3 Select **Edit > Delete**.
- 4 Click **OK**.

Configuring host credentials

This chapter includes the following topics:

- [About configuring credentials](#)

About configuring credentials

Credentials appears only if a feature that requires external credentials is licensed.

Use **Media and Device Management > Credentials** to manage log on credentials for the following:

- NetBackup Deduplication Engine credentials.
You create the credentials when you configure the storage server.
See the *NetBackup Deduplication Guide*.
- NDMP hosts.
See the *NetBackup for NDMP Administrator's Guide*.
- OpenStorage storage servers.
You configure the credentials when you configure the storage server.
See the *NetBackup Shared Storage Guide*.

Managing media servers

This chapter includes the following topics:

- [Activating or deactivating a media server](#)
- [Adding a media server](#)
- [Decommissioning a media server](#)
- [Registering a media server](#)
- [Deleting all devices from a media server](#)
- [Removing a device host from the EMM database](#)

Activating or deactivating a media server

When you activate a media server, NetBackup can use it for backup and restore jobs. For example, you can deactivate a media server to perform maintenance. When a media server is deactivated, NetBackup does not send job requests to it.

When you deactivate a media server, the following things occur:

- Current jobs are allowed to complete.
- No new jobs are scheduled for the host.
- If the host is part of a shared drive configuration, it does not scan drives.

To activate or deactivate a media server

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media Servers**.
- 2 From the **Media Servers** pane, select the media server to activate or deactivate.
- 3 On the **Actions** menu, select **Activate** or **Actions > Deactivate**.

Adding a media server

This topic applies to NetBackup Enterprise Server only.

Table 6-1 is an overview of how add a media server to an existing NetBackup environment.

Note: The NetBackup Enterprise Media Manager service must be active when a media server is added, devices and volumes are configured, and clients are backed up or restored.

Table 6-1 Adding a media server

Procedure	Section
On the new media server host, attach the devices and install any software that is required to drive the storage devices.	See the vendor’s documentation.
On the new media server host, prepare the host’s operating system.	See the <i>NetBackup Device Configuration Guide</i> .
On the master server, add the new media server to the additional servers list of the master server. Also, add the new media server to the additional servers list of the clients that the new media server backs up. If the EMM server resides on a host other than the master server, add the new media server to the additional servers list on that host. If the new media server is part of a server group, add it to the additional servers list on all media servers in the group. To avoid problems with NetBackup, ensure that the host name used in NetBackup matches the host name in the TCP/IP configuration.	See “ Servers properties ” on page 175.
Restart the NetBackup services on the master server, the EMM server, and the media servers where a new server name was added.	See “ Starting or stopping a service ” on page 693.
On NetWare target clients, add the new media server name by using a <code>server</code> entry in the <code>bp.ini</code> file.	See the <i>NetBackup for Novell NetWare Client System Administrator’s Guide</i> .
Install the NetBackup media server software.	See the <i>NetBackup Installation Guide for Windows</i> .

Table 6-1 Adding a media server (*continued*)

Procedure	Section
On the master server, configure the robots and drives that are attached to the media server.	See “Configuring robots and tape drives” on page 217.
On the master server, configure the volumes.	See “About adding volumes” on page 257.
On the master server, add storage units to the media server. Always specify the media server as the media server for the storage unit. The Device Configuration Wizard can create storage units when you configure robots and drives. Therefore, if you created storage units already, skip this step.	See “Creating a storage unit using the Actions menu” on page 368.
On the master server, configure the NetBackup policies and schedules to use the storage units that are configured on the media server.	See “Using the Policies utility” on page 448.
Test the configuration by performing a user backup or a manual backup that uses a schedule that specifies a storage unit on the media server.	See “Performing manual backups” on page 574.

Decommissioning a media server

This topic applies to NetBackup Enterprise Server.

Use the following procedure to decommission a media server.

If you do not use this procedure and a subsequent restore requires the media that is associated with the decommissioned media server, the restore fails. You then must import the media. Importing media consumes more time than the following procedure.

Note: If you use NetBackup Vault and plan to decommission a media server, contact Symantec Consulting for help with this task.

Table 6-2 Decommission a media server overview

Task	Procedure
If devices attached to the media server contain valid NetBackup media, first move them to a new media server.	See “Moving a robot and its media to a new media server” on page 240.

Table 6-2 Decommission a media server overview (*continued*)

Task	Procedure
Delete all storage units that use the robots that are associated with the media server.	See “Deleting storage units” on page 369.
If the media server has robots, drives, or disk pools, delete them from the media server.	See “Deleting all devices from a media server” on page 207.
Modify any backup policies and storage lifecycle policies that specify the storage units on the media server.	Change the policies to point to any other defined storage units in the NetBackup configuration or to <code>Any Available</code> . See “Changing policies” on page 456.
Remove all references to the media server from the Servers host properties for the master server, all media servers, and all clients.	See “Servers properties” on page 175.
Restart the NetBackup daemons or services on any system that was updated.	See “Starting or stopping a service” on page 693.
Delete the media server.	Run the following command on the master server: <code>install_path\NetBackup\bin\admincmd\nbermcmd -deletehost -machinename server_name -machinetype media</code> Replace <code>server_name</code> with the name of the media server.
Verify that all references to the decommissioned media server have been removed.	Run the following command on the master server: <code>install_path\NetBackup\bin\admincmd\nbermcmd -listhosts</code>

Registering a media server

If the EMM server is not running when you install a media server, the media server is not registered. You cannot discover, configure, and manage the devices of that media server. You must register the media server with the EMM server.

To register a media server

- 1 Start the EMM service on the EMM server.
- 2 On the EMM server host, run the following command (for the *hostname*, use the host name of the media server):

```
nbeemmcmd -addhost -machinename hostname -machinetype media  
-masterserver server_name -operatingsystem  
os_type-netbackupversion level.major_level.minor_level
```

To avoid problems with NetBackup, ensure that the host name that is used in NetBackup matches the host name in the TCP/IP configuration.

Information about `nbeemmcmd` command usage is available.

See the *NetBackup Commands* guide.

Deleting all devices from a media server

You can delete all devices from a media server. The media server can be up, down, or failed and unrecoverable. All devices include robots, drives, and disk pools.

Two procedures exist: one to delete all robots and drives and the other to delete disk pools.

To delete all robots and drives from a media server

- ◆ Enter the following command on the master server:

```
install_path\NetBackup\bin\admincmd\nbeemmcmd -deletealldevices  
-machinename server_name -machinetype media
```

Replace *server_name* with the name of the media server.

To delete disk pools from a media server

- 1 If the media server has disk pools configured, remove the media server from the storage units that use those disk pools. For each storage unit, run the following command on the master server:

```
install_path\NetBackup\bin\admincmd\bpsturep -label  
storage_unit_label -delhost host_name
```

Replace *storage_unit_label* with the name of the storage unit and *host_name* with the name of the media server.

- 2 If the media server is the only storage server for the disk pools, change the state of the disk pools to DOWN. To do so, enter the following command on the master server for each disk pool:

```
install_path\NetBackup\bin\admincmd\nbdev config -changestate  
-stype server_type -dp disk_pool_name -state DOWN
```

Replace *server_type* with the type of storage server: AdvancedDisk, PureDisk, or the vendor string that identifies the OpenStorage server type.

Replace *disk_pool_name* with the name of the disk pool.

- 3 For each disk pool, do the following:
 - Remove the media server from disk pool access by entering the following command on the master server:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -changedp -dp  
-disk_pool_name stype server_type -del_storage_servers  
storage_server
```

Replace *disk_pool_name* with the name of the disk pool.

Replace *server_type* with the type of storage server: AdvancedDisk, PureDisk, or the vendor string that identifies the OpenStorage server type. Replace *storage_server* with the name of the media server.

- If the disk pool is on disk storage available only to the media server and is no longer required, delete the disk pool as follows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -deletedp -dp  
disk_pool_name -stype server_type
```

You cannot delete a disk pool that has unexpired backup images. You must first expire the images and delete the image fragments, as follows:

- Expire the image as follows:

```
install_path\NetBackup\bin\admincmd\bpexpdate -dp  
disk_pool_name -stype server_type -nodelete
```

- Determine the media IDs in the disk pool as follows:


```
install_path\NetBackup\bin\admincmd\bpimmediate -dp  
disk_pool_name -stype server_type -nodelete
```

- Delete each media ID in the disk pool as follows:

```
install_path\NetBackup\bin\nbdelete -dt disk_type -media_id  
name
```

Removing a device host from the EMM database

The following applies only to NetBackup Enterprise Server.

To remove a device host from the EMM database

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices > Media Servers**.
- 2 Select the host.
- 3 On the **Actions** menu, select **Enterprise Media Manager Database > Remove Device Host**.
- 4 Click **Yes** in the confirmation dialog box.

Configuring storage

- [Chapter 7. Configuring robots and drives](#)
- [Chapter 8. Configuring tape media](#)
- [Chapter 9. Inventorying robots](#)
- [Chapter 10. Configuring disk storage](#)
- [Chapter 11. Configuring storage units](#)
- [Chapter 12. Staging backups](#)
- [Chapter 13. Configuring storage unit groups](#)
- [Chapter 14. Configuring storage lifecycle policies](#)

Configuring robots and drives

This chapter includes the following topics:

- [About optical device support in NetBackup 7.0](#)
- [About NetBackup robot types](#)
- [Device configuration prerequisites](#)
- [About the device mapping file](#)
- [Updating the device mapping file](#)
- [About configuring robots and tape drives](#)
- [Configuring robots and tape drives](#)
- [Updating the device configuration by using the wizard](#)
- [Managing robots](#)
- [Managing tape drives](#)
- [Performing device diagnostics](#)
- [Verifying the device configuration](#)
- [Replacing a device](#)
- [Updating device firmware](#)
- [About the NetBackup Device Manager](#)
- [Stopping and restarting the Device Manager](#)

About optical device support in NetBackup 7.0

Beginning with the 7.0 release, NetBackup media servers do not support optical drives or libraries. However, you can use optical devices on NetBackup 6.x media servers.

For information about how to use optical devices, see the documentation for your NetBackup 6.x release.

About NetBackup robot types

A robot is a peripheral device that mounts and unmounts media in tape drives. NetBackup uses robotic control software to communicate with the robot firmware.

NetBackup classifies robots according to one or more of the following characteristics:

- The communication method the robotic control software uses; SCSI and API are the two main methods.
- The physical characteristics of the robot. Library usually refers to a larger robot, in terms of slot capacity or number of drives. Stacker usually refers to a robot with one drive and low media capacity (6 - 12 media slots).
- The media type commonly used by that class of robots. HCART (1/2-inch cartridge tape) and 8 mm are examples of media types.

[Table 7-1](#) lists the NetBackup robot types, with drive and slot limits for each type.

To determine which robot type applies to the model of robot that you use, see the Symantec support web site at the following URL:

<http://entsupport.symantec.com>

Table 7-1 NetBackup robot types

Robot type	Description	Drive limits	Slot limits	Note
ACS	Automated Cartridge System	1680	No limit	API control. Drive limit determine by ACS library software host.
TL4	Tape Library 4mm	2	15	SCSI control.
TL8	Tape Library 8mm	No limit	16000	SCSI control.
TLD	Tape Library DLT	No limit	16000	SCSI control.
TLH	Tape Library Half-inch	256	No limit	API control.

Table 7-1 NetBackup robot types (continued)

Robot type	Description	Drive limits	Slot limits	Note
TLM	Tape Library Multimedia	250	No limit	API control.

Device configuration prerequisites

Before you configure storage devices in NetBackup, ensure that the following prerequisites are accomplished:

- The storage devices must be attached to the computer and recognized by the operating system. The server platforms that NetBackup supports may require operating system configuration changes to allow device discovery.

The *NetBackup Device Configuration Guide* provides information about how to configure device drivers for the systems that NetBackup supports.

- If the host on which you configure devices in NetBackup is not the Enterprise Media Manager server, add it to the NetBackup additional servers list.

See “[Servers properties](#)” on page 175.

NetBackup hosts are added automatically to the list of additional servers if the EMM server is running when the host is installed.

If the EMM server is not running, use the `nbbemcmd -addhost` command to add the host.

See the *NetBackup Commands* guide.

About the device mapping file

NetBackup uses a file to determine which protocols and settings to use to communicate with storage devices. NetBackup also uses the file during device discovery and configuration.

In some cases, you can add support for new or upgraded devices without waiting for a release update from Symantec. To do so, download the current device mapping file from the Symantec support Web site and configure NetBackup to use that file. Refer to the README file that is supplied with the device mapping file for additional instructions.

Note: The contents of this file do not indicate support for any of the devices, only the ability to recognize and automatically configure them.

See “[Updating the device mapping file](#)” on page 216.

Updating the device mapping file

Use the following procedure to download the current device mapping file and update the Enterprise Media Manager database with its information.

To update the current device mapping file

- 1 Find the latest external types file for your devices on the Symantec support Web site. The following is the address for the site:

`http://entsupport.symantec.com`

The files that you download are named similarly to the following files (*x* represents the NetBackup release, such as 7):

`Mappings_x_nnnnnn.ZIP`

- 2 Download the file to the following location on the system that hosts the EMM server:

- `/usr/opensv/var/global` (UNIX)

- `install_path\netbackup\var\global` (Windows)

- 3 Uncompress the file.

It contains a `Readme.txt` and a `device_mappings.txt` file.

- 4 Update the EMM database with the new version of the device mapping file by using the following command (no command-line parameters are required):

`install_path\volmgr\bin\tpext`

- 5 Stop and restart the NetBackup Device Manager.

See [“Stopping and restarting the Device Manager”](#) on page 256.

About configuring robots and tape drives

Symantec recommends that you use the Device Configuration Wizard to add, configure, and update the following types of devices:

- Robots, including those attached to NDMP hosts
- Tape drives, including those attached to NDMP hosts
- Shared drives (for NetBackup Shared Storage Option configurations only)

The wizard discovers the devices attached to the media servers and helps you configure them.

See [“About device discovery”](#) on page 217.

Alternatively, you can add robots and drives manually as follows:

- Use menu options in the NetBackup Administration Console.
See “[Adding a robot](#)” on page 221.
See “[Adding a tape drive](#)” on page 227.
- Use NetBackup commands.
See the *NetBackup Commands* guide.

Manual methods do not use device discovery.

If you do not use the wizard, first add the robot and then add the drives that are in the robot.

Device configuration examples are available.

See the *NetBackup Device Configuration Guide*.

Configuring robots and tape drives

Symantec recommends that you use the **Device Configuration Wizard** to configure robots and drives. However, you can add robots and drives manually.

About device discovery

Device discovery is an exploratory method that determines which peripheral devices a host can detect. Detection depends on physical attachment (SCSI, Fibre Channel, and so on) and device state (on and responding or off and not responding). Detection also depends on host operating system device-layer configuration.

The goal of device discovery is to provide information to enable fully or partially automatic configuration of peripherals for use with NetBackup. Device discovery provides data that correlates the devices that are interconnected across multiple hosts or multiple host bus adapters on the same host.

To discover devices, NetBackup issues SCSI pass-through commands through operating system device files (on UNIX) or APIs (on Windows). The storage devices must be attached to the computer and recognized by the operating system. A pass-through path to a device must exist.

About operating system changes

The operating systems that NetBackup supports may require configuration changes to allow device discovery.

The *NetBackup Device Configuration Guide* provides information about how to configure device drivers for the systems that NetBackup supports.

About device serialization

Device serialization is a firmware feature that allows device identification and configuration. A unique serial number identifies a device.

NetBackup determines device relationships by comparing serial numbers from multiple sources that refer to the same device. If both a robotic library and a drive fully support serialization, NetBackup can determine the drive's position (or address) in the robotic library.

Most robots and drives support device serialization.

If a device supports serialization, the following actions occur when NetBackup queries the devices:

- Each robot and each drive return a unique serial number.
- Each robot also returns the number of drives and the serial number for each of the drives in the robot. NetBackup uses the information to determine the correct drive number for each drive in the robot.

If a device does not support serialization, ask the vendor for a new firmware revision that returns serial numbers. Even with the proper firmware, some devices require the vendor to perform other actions to enable serialization for the device.

If you know that the devices do not support serialization, make sure that you follow the maximum configuration limits that the devices allow. You also must coordinate the drives to their device files or SCSI addresses so you can configure them correctly.

See [“Correlating tape drives and SCSI addresses on Windows hosts”](#) on page 237.

The more devices in the configuration that do not support serialization, the greater the chance of configuration problems by using the **Device Configuration Wizard**.

About the devices that can be discovered

NetBackup can discover the following types of devices:

- SCSI-based robotic libraries (such as changers, autoloaders, and stackers)
- SCSI-based tape drives
- Native parallel SCSI, Fibre Channel Protocol (FCP) and FC-AL (loop) connections
- SCSI over IP (reported)
- API type robots, such as ACS, TLM, and TLH robots
- NDMP devices that run NDMP version 3 or later

About adding devices without discovery

NetBackup supports some devices that cannot be discovered automatically. NetBackup also supports some devices that require user intervention during the discovery process. To add and configure those devices, use **Media and Device Management** in the NetBackup Administration Console or the `tpconfig` command.

For the devices that NetBackup cannot discover or that do not have serial numbers, automated device path correction when the `ltid` device manager starts is limited.

Configuring robots and drives by using the wizard

Symantec recommends that you use the Device Configuration Wizard to configure robots and drives. The wizard configures a robot, its drives, and a storage unit.

To configure robots and drives by using the wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

The properties you can configure depend on the robot type, the host type, and the robot control.

About robot control

When you add a robot to NetBackup manually, you must configure how the robot is controlled. The **New Robot** dialog box includes a section named **Robot control**, in which you configure the control options.

See “[Robot control section](#)” on page 223.

Table 7-2 lists the information required to configure the three robot control types (local, NDMP, and remote). The information required depends on the robot type and the media server type.

Table 7-2 Robot control information

Robot type	Media server type	Robot control	Information required for configuration
ACS	Windows, AIX, Solaris SPARC, HP-UX (except HP IA64), and Linux (except Linux64)	NDMP	NDMP host name and robot device
ACS	All	Remote	ACSLS host

Table 7-2 Robot control information (*continued*)

Robot type	Media server type	Robot control	Information required for configuration
TL4	UNIX	Local	Robotic device file
TL4	Windows	Local	Robot device or SCSI coordinates
TL8	UNIX	Local	Robotic device file
TL8	Windows	Local	Robot device or SCSI coordinates
TL8	Windows, AIX, Solaris SPARC, HP-UX (except HP IA64), and Linux (except Linux64)	NDMP	NDMP host name and robot device
TL8	All	Remote	Robot control host
TLD	UNIX	Local	Robotic device file
TLD	Windows	Local	Robot device or SCSI coordinates
TLD	Windows, AIX, Solaris SPARC, HP-UX (except HP IA64), and Linux (except Linux64)	NDMP	NDMP host name and robot device
TLD	All	Remote	Robot control host
TLH	All (except Solaris Opteron, HP IA64, AIX, Linux, and Linux64)	Local	Library name
TLH	AIX	Local	LMCP device file
TLH	Windows, AIX, Solaris SPARC, HP-UX (except HP IA64), and Linux (except Linux64)	NDMP	NDMP host name and robot device
TLH	All (except Solaris Opteron, Linux64)	Remote	Robot control host
TLM	All (except Linux64 and HP IA64)	Remote	DAS/SDLC server

Library sharing example


The following example applies only to NetBackup Enterprise Server.

[Figure 7-1](#) shows library sharing with two servers using two drives in a TLD robot.

The robotic control for the robot is on the host that is named eel. One drive in the robot is connected to eel and the other is connected to the host shark.

Host eel is the robot control host. To configure this robot on host eel, select **Robot is controlled locally by this device host**. To configure this robot on host shark, select **Robot control is handled by a remote host**. Then, enter eel for the **Robot control host**.

Figure 7-1 Robot control host example



Adding a robot

Symantec recommends that you use the **Device Configuration Wizard** to add, configure, and update tape storage devices.

After you add a robot, you should add the robot's drives.

See [“Adding a tape drive”](#) on page 227.

See [“About NetBackup robot types”](#) on page 214.

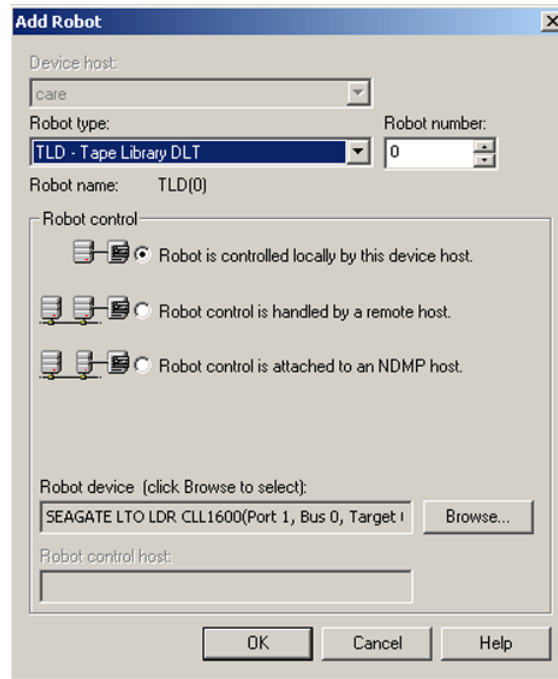
To add a robot using the Actions menu

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **New > New Robot**.

- 3 In the **Add Robot** dialog box, specify the properties for the robot.

The properties you can configure depend on the robot type, the host type, and the robot control.

See “[Robot configuration options](#)” on page 222.



- 4 After you specify properties, click **OK**.
- 5 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box. If you intend to make other changes, click **No**; you can restart the Device Manager after you make the final change.

If you restart the Device manager, any backups, archives, or restores that are in progress also may be stopped.

Robot configuration options

The following topics describe the robot properties you can configure. The properties you can configure depend on robot type, host type, and robot control selections you make in the dialog box.

Device host

The **Device host** option applies only to NetBackup Enterprise Server.

Device host option specifies the host to which the device is attached.

Robot type

Specifies the type of robot. To locate the robot type to use for specific vendors and models, see the Symantec support Web site:

`http://entsupport.symantec.com`

See “[About NetBackup robot types](#)” on page 214.

Robot number

Specifies a unique, logical identification number for the robotic library. This number identifies the robotic library in displays (for example, TLD (21)) and is also used when you add media for the robot.

For NetBackup Enterprise Server environments, do the following:

- Robot numbers must be unique for all robots on all hosts in the configuration, regardless of the robot type or the host that controls them. For example, if you have two robots, use different robot numbers even if different hosts control them.
- If you add a robot that is controlled by a remote device host, use the same robot number for that robot on all device hosts.
- If the robot has its robotic control and drives on different hosts, specify the same robot number in all references to that library. That is, use the same robot number on the hosts with the drives as you do on the host that has the robotic control. A Tape Library DLT robot is one that allows separate robotic control and drive hosts.

Examples are available.

See the *NetBackup Device Configuration Guide*.

Robot control section

The Robot control section defines the type of control for the robot. The options you configure depend on the robot type and the media server type.

Robot is controlled locally by this device host

This option specifies that the host to which the robot is attached controls the robot.

You must configure other options (depending on the robot type and device host type).

See “[Library name](#)” on page 225.

See “[LMCP device file](#)” on page 225.

See “[Robot device](#)” on page 226.

See “[Robotic device file](#)” on page 226.

Robot control is attached to an NDMP host

The following applies only to NetBackup Enterprise Server.

This option specifies that an NDMP host controls the robot.

You must configure other options (depending on the robot type and device host type).

See “[NDMP host name](#)” on page 226.

See “[Robot device path](#)” on page 227.

See “[SCSI coordinates](#)” on page 227.

Robot control is handled by a remote host

The following applies only to NetBackup Enterprise Server.

Specifies that a host other than the device host controls the robot.

You must configure other options (based on the selected robot type and device host platform).

See “[ACSLS host](#)” on page 224.

See “[DAS server](#)” on page 225.

See “[Robot control host](#)” on page 226.

ACSLS host

The **ACSLS host** option applies only to NetBackup Enterprise Server.

The name of the Sun StorageTek ACSLS host; the ACS library software resides ACSLS host. On some UNIX server platforms, this host can also be a media server or EMM server.

The ACS library software component can be any of the following:

- Automated Cartridge System Library Software (ACSLS)
Examples are available.
See the *NetBackup Device Configuration Guide*.
- STK Library Station

- **Storagenet 6000 Storage Domain Manager (SN6000).**
This STK hardware serves as a proxy to another ACS library software component (such as ACSLS).

Note: If the device host that has drives under ACS robotic control is a Windows server, STK LibAttach software must also be installed. Obtain the appropriate LibAttach software from STK. See the Symantec support Web site for the latest compatibility information.

An overview of ACS robots is available.

See the *NetBackup Device Configuration Guide*.

DAS server

The following applies only to NetBackup Enterprise Server.

The name of the ADIC DAS/SDLC server that controls TLM robots.

This server is an OS/2 workstation near or within the robot cabinet or a Windows server near the ADIC Scalar library.

An overview of TLM robots is available.

See the *NetBackup Device Configuration Guide*.

Library name

The following applies only to a TLH robot on NetBackup Enterprise Server only.

For UNIX device hosts (except AIX), the library name that is configured on the UNIX host.

For Windows devices hosts, do the following:

- **Determine the library name by viewing the C:\winnt\ibmat1.conf file.**
For example, in the following example entry in that file, 3494AH is the library name:

```
3494AH 176.123.154.141 ibmpc1
```

- **Enter the library name.**

An overview of TLH robots is available.

See the *NetBackup Device Configuration Guide*.

LMCP device file

Applies to NetBackup Enterprise Server on an AIX device host only.

The name of the Library Manager Control Point device file name for TLH robot types. Use the same name that is configured on the AIX device host.

NDMP host name

The name of the NDMP host to which the robot is attached.

Robot control host

The following applies only to NetBackup Enterprise Server.

This option specifies the host that controls the robot.

The name of the host on which the robot information is defined for TL8, TLD, or TLH robots.

See “[Robot control section](#)” on page 223.

Robot device

The following applies to a Windows device host only.

The name of the robot device.

Click **Browse** and then select a robot from the list that appears in the **Devices** dialog box.

If the discovery operation fails to discover a robot, click **More** in the **Devices** dialog box. Enter either the **Port**, **Bus**, **Target**, and **LUN** numbers or the device name in the next dialog box. If the browse operation fails for any other reason, a dialog box appears that lets you enter the information.

You can find Port, Bus, Target, and LUN numbers by using Windows management tools.

If the browse operation does not find attached robots, an error dialog box appears.

Robotic device file

The following applies to a UNIX device host only.

The device file that is used for SCSI connections. The device files are located in the `/dev` directory tree on the device host.

To specify the robotic device file, click **Browse** and then select a robotic device file from the list that appears in the **Devices** dialog box.

If the browse operation fails to show all of the attached robots, click **More**. Enter the path of the device file in the **robotic device file** field.

If the browse operation fails to show all of the attached robots, click **Other Device**. Enter the path of the device file in the next dialog box.

If the browse operation does not find attached robots, an error dialog box appears.

Information about how to add device files is available.

See the *NetBackup Device Configuration Guide*.

Robot device path

NDMP host only.

The name of the robotic device that is attached to the NDMP host.

SCSI coordinates

Windows systems only.

The device attributes on Windows system cannot change during NetBackup operation.

The Port, Bus, Target, and LUN SCSI coordinates for the robotic device.

Adding a tape drive

Symantec recommends that you use the **Device Configuration Wizard** to add, configure, and update tape storage devices.

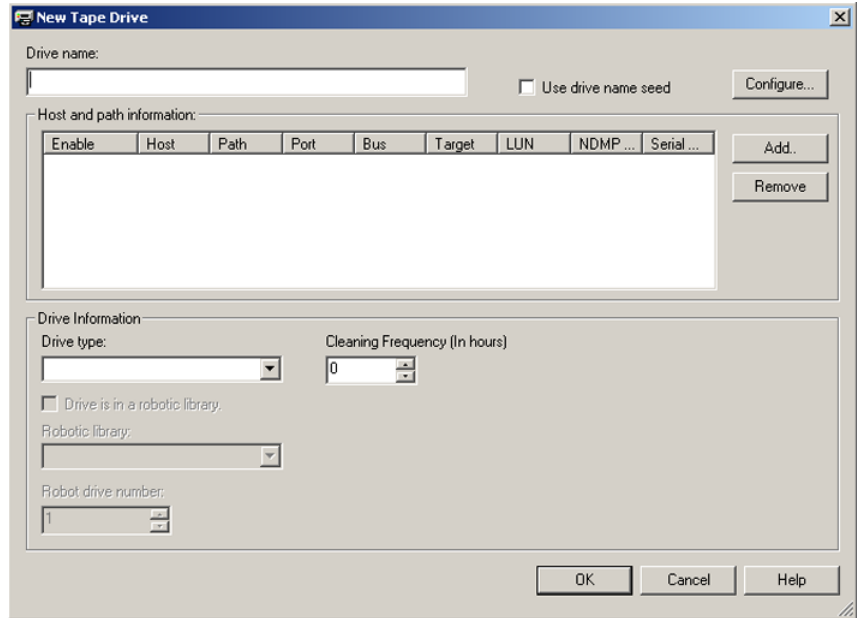
Use the following procedures to add or change a tape drive manually.

When you add a drive, you can configure drive name rules.

See [“About drive name rules”](#) on page 233.

To add a drive using the Actions menu

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **New > New Tape Drive**.



- 3 Specify the properties of the drive.
The properties depend on the drive type and host server type.
See [“Tape drive configuration options”](#) on page 229.
- 4 After you configure all of the properties, click **OK**.
- 5 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box. If you intend to make other changes, click **No**; you can restart the Device Manager after you make the final change.

If you restart the device manager, any backups, archives, or restores that are in progress also may be stopped.

The initial drive status is UP, so the drive is available as soon as you restart the Device Manager. To change the status of the drive, select **Device Monitor**.

Adding a shared tape drive

Procedures for shared drives in a Shared Storage Option configuration are documented in the *NetBackup Shared Storage Guide*.

Symantec recommends that you use the Device Configuration Wizard to add, configure, and update shared drives. The Device Configuration Wizard is the easiest method for adding shared drives in a Shared Storage Option configuration.

Tape drive configuration options

You specify properties when you add a tape drive or change the properties of a drive. The properties you can specify depend on drive type, server platforms, or NetBackup server types.

Drive name

The name of the drive. Each drive name must be unique. Symantec recommends that you use descriptive names. Drive names are limited to 48 characters.

Alternatively, use drive name rules to create a unique drive name.

See [“Use drive name seed”](#) on page 229.

See [“NetBackup naming conventions”](#) on page 719.

Use drive name seed

Adding a drive only.

Select to use drive name rules to assign names to drives automatically.

To configure drive name rules, click **Configure**.

See [“About drive name rules”](#) on page 233.

See [“Configuring drive name rules”](#) on page 233.

Host and path information

For tape drives, use the host and path information list to add or change paths to the drive. You can specify multiple paths to the same physical device. If you specify multiple paths for a drive, it becomes a shared drive.

To add a drive path, click Add.

See [“Add Path options”](#) on page 235.

Drive information

The drive information properties you can specify depend on drive type, server platforms, or NetBackup server types.

Drive type

This property specifies the type of drive.

The following are the valid drive types:

- 4MM (4mm cartridge)
- 8MM (8mm cartridge)
- 8MM2 (8mm cartridge 2)
- 8MM3 (8mm cartridge 3)
- DLT (DLT cartridge)
- DLT2 (DLT cartridge 2)
- DLT3 (DLT cartridge 3)
- DTF (DTF cartridge)
- HCART (1/2-inch cartridge)
- HCART2 (1/2-inch cartridge 2)
- HCART3 (1/2-inch cartridge 3)
- QSCSI (1/4-inch cartridge)

Drive is in a robotic library

This property specifies that the drive is in a robot. If the drive is a stand-alone drive (not in a robot), do not select this option.

If you select this option, configure the **Robotic library** and **Robot drive number** fields.

See [“Robotic library”](#) on page 231.

See [“Robot drive number”](#) on page 231.

Cleaning frequency

The following applies only to tape drives. NetBackup does not support drive cleaning in some robot types.

If you want to set up a frequency-based cleaning schedule for the drive, set the number of mount hours between each drive cleaning. When you add a drive or

reset the mount time to zero, NetBackup records the amount of time that volumes have been mounted in that drive. The default frequency is zero.

When the accumulated mount time exceeds the time you specify for cleaning frequency, drive cleaning occurs if the following are true:

- If the drive is in a robotic library that supports drive cleaning
- If a cleaning cartridge is defined in that robotic library
- If the cleaning cartridge is compatible with the drive that needs to be cleaned
- If the cleaning cartridge has a nonzero number of cleanings that remain

NetBackup resets the mount time when the drive is cleaned.

If you do not specify a cleaning frequency, you can still use automated drive cleaning with the TapeAlert feature if the following are true:

- The drive supports TapeAlert.
- You configured a cleaning volume for the robot.
- The host platform, robot type, and drive support drive cleaning.
- If the cleaning cartridge is compatible with the drive that needs to be cleaned
- If the cleaning cartridge has a nonzero number of cleanings that remain

Drives can also be cleaned from the **Device Monitor**.

Additional information about drive cleaning is available.

See the *NetBackup Administrator's Guide for Windows, Volume II*.

Serial number

This read-only field shows the serial number of the drive.

Robotic library

The robot that controls the drive.

You can select any configured robot that can control the drive.

Robot drive number

Note: Robot drive number does not apply when you add drives to API robots. API robots are ACS, TLH, and TLM type in NetBackup.

See “[ACS, LSM, Panel, Drive](#)” on page 232.

See “[IBM device number](#)” on page 232.

See “[DAS drive name](#)” on page 232.

Specifies the physical location in the robot of the drive. When you add more than one drive to a robot, you can add the physical drives in any order. For example, you can add drive 2 before drive 1.

The correct robot drive number is critical to the proper mounting and utilization of media. You must determine which logical device name (Windows) or the device file (UNIX) identifies which physical drive in the robot. You should correlate the drive serial number with drive serial number information from the robot.

You must determine which physical drive in the robot is identified by the logical device name.

See [“Correlating tape drives and SCSI addresses on Windows hosts”](#) on page 237.

NetBackup does not detect incorrect drive number assignment during configuration; however, an error occurs when NetBackup tries to mount media on the drive.

ACS, LSM, Panel, Drive

ACS robot drive only.

The physical location of the drive within the robot.

During installation, the correlation between the physical drive in the robot and the device file you specified earlier represents. You establish this correlation during installation.

Table 7-3 Drive location properties

ACS Number	The index (in ACS library software terms) that identifies the robot that has this drive.
LSM Number	The Library Storage Module that has this drive.
Panel Number	The robot panel where this drive is located.
Drive Number	The physical number of the drive (in ACS library software terms).

IBM device number

TLH robot drive only.

The IBM device number of the drive within the robot.

DAS drive name

TLM robot drive only.

The DAS/SDLC drive name of the drive within the robot.

About drive name rules

Drive name rules define the rules NetBackup uses to name drives.

The default, global drive name rule creates names in the following format:

vendor ID.product ID.index

If you use the default global rule when you add Quantum DLT8000 drives, the drives are named as follows: The first one that you add is named QUANTUM.DLT8000.000, the second one QUANTUM.DLT8000.001, and so on.

You can change the default, global drive name rule.

You also can create drive name rules for specific device hosts (each device host can have its own rule). Host-specific rules override the global rule for the devices that are attached to the specified host.

Only one global rule can exist; it is used for all connected device hosts. The global rule is used for the drive name unless a host-specific rule or local rule is specified.

Drive names are limited to 48 characters.

Use any of the following drive attributes as part of a drive name rule:

- Host name
- Robot number
- Robot type
- Drive position
Drive position information varies depending on the robot type. Drive position information can be ACS coordinates, TLM or TLH vendor drive name, or the robot drive number.
- Drive type
- Serial number
- Vendor ID
- Product ID
- Index

A **Custom Text** field is also available which accepts any of the allowable drive name characters.

See [“Configuring drive name rules”](#) on page 233.

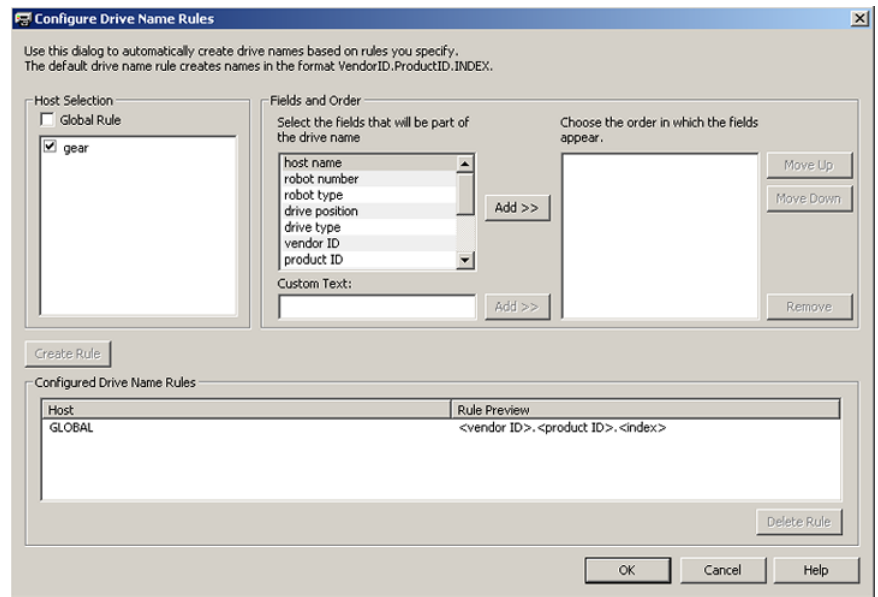
Configuring drive name rules

Use the following procedure to configure drive name rules.

To configure drive name rules

- 1 Open the **New Tape Drive** dialog box.
See “[Adding a tape drive](#)” on page 227.
- 2 In the **New Tape Drive** dialog box, click **Configure**.

Alternatively, if you use the Device Configuration Wizard, click **Configure Drive Name Rules** in the **Device Hosts** screen.



- 3 In the **Configure Drive Name Rules** dialog box, configure the rules for naming drives:
 - To change the global rule, select **Global Rule**.
 - To create a local rule, select the check box for the device host.
 - Select the fields from which to create the drive name from the list of available fields. Click **Add >>** to make a field part of the rule.
 - To add own text to the drive name rule, enter the text in the **Custom Text** field and click the **Add** button.
 - Use the **Move Up** and **Move Down** buttons to change the order of the fields that are defined for the rule.
 - Click **Create Rule** to finalize the rule.

If you use **<host name>** in the rule and the drive is a shared drive, the name of the first host that discovers the drive is used as the host name. The name for a shared drive must be identical on all servers that share the drive.

Adding a tape drive path

Use the following procedure to add a drive path.

To add a drive path

- 1 In the **New Tape Drive** dialog box, click **Add**.
- 2 In the **Add Path** dialog box, configure the properties for the drive path.
 The properties you can specify depend on drive type, server platform, or NetBackup server type.

See [“About SCSI reserve on drive paths”](#) on page 235.

See [“Add Path options”](#) on page 235.

About SCSI reserve on drive paths

NetBackup lets you configure exclusive access protection to tape drives so that other host bus adaptors (HBAs) cannot control the drives during the reservation. The **Enable SCSI Reserve** host property configures the protection for each media server.

See [“Media properties”](#) on page 153.

More information about how NetBackup reserves drives is available.

See the *NetBackup Administrator's Guide for Windows, Volume II*.

Add Path options

[Table 7-4](#) describes the options to add a drive path.

Table 7-4 Add drive path options.

Host name	Applies only to NetBackup Enterprise Server. This attribute specifies the device host for the drive.
Enable host path	Check this attribute to specify that the path is active and that NetBackup can use it for backups and restores.

Table 7-4 Add drive path options. (continued)

NDMP host	<p>This attribute specifies the NDMP host for the device (if an NDMP host is configured in your NetBackup environment).</p> <p>Additional information is available about NDMP drives.</p> <p>See the <i>NetBackup for NDMP Administrator's Guide</i>.</p>
Override SCSI Reserve settings	<p>This attribute specifies the SCSI reserve override setting for the drive path.</p> <ul style="list-style-type: none"> ■ Server Default. Use the SCSI reserve protection setting configured for the media server. If the setting for the media server is no protection, other HBAs can send the commands that can cause a loss of data to the tape drives. ■ SPC-2 SCSI Reserve. This option provides SCSI reserve and release protection for SCSI devices that conform to the reserve and release management method that is defined in the SCSI Primary Commands - 2 (SPC-2) standard. ■ SCSI Persistent Reserve. This option provides SCSI persistent reserve in and persistent reserve out protection for SCSI devices that conform to the SCSI Primary Commands - 3 (SPC-3) standard. <p>Global SCSI reserve properties are configured in the Media host properties.</p> <p>See “Media properties” on page 153.</p>
Port, Bus, Target, and LUN	<p>The device attributes on Windows systems cannot change during NetBackup operation.</p> <p>To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN.</p>
This path is for a Network Attached Storage Device	<p>Specifies that the path is for a network attached storage (NAS) device.</p>

About no rewind device files

UNIX servers only.

Although both rewind and no rewind on close device files are usually available, NetBackup requires only the no rewind device file. A no rewind device remains at its current position on a close operation. On some versions of UNIX, the device file name may be preceded or followed by the letter n.

Device files are in the /dev directory on the UNIX host. If the entries do not exist, create them as explained in the NetBackup Device Configuration Guide.

Correlating tape drives and SCSI addresses on Windows hosts

If your tape drives do not support device serialization, you may have to determine which logical device name or SCSI address matches the physical drive. You also may have to do so if you add the tape drives manually.

To correlate tape drives and SCSI addresses on Windows hosts

- 1 Note the SCSI target of the drive.
- 2 Correlate the SCSI target to the drive address by using the robot's interface panel. Alternatively, examine the indicators on the rear panel of the tape drive.
- 3 Determine the physical drive address (for example, number) by checking labels on the robot.
- 4 Configure the robot in NetBackup and then add the drives.

When you add the drives, ensure that you assign the correct drive address to each set of SCSI coordinates.

Optionally, use the appropriate NetBackup robotic test utility to verify the configuration.

Information about the robotic test utilities is available.

See the *NetBackup Troubleshooting Guide*.

To verify the device correlation Windows

- 1 Stop the NetBackup Device Manager (`ltid`).
- 2 Restart `ltid`, which starts the Automatic Volume Recognition process (`avrd`). Stop and restart `ltid` to ensure that the current device configuration has been activated.

The following point applies only to NetBackup Enterprise Server.

If robotic control is not local to this host, also start the remote robotic control daemon.

- 3 Use the robotic test utility to mount a tape on a drive.
- 4 Use the NetBackup Device Monitor to verify that the tape was mounted on the correct robot drive.

Windows device correlation example

Assume a TLD robot includes three drives at the following SCSI addresses:

Drive 1	5,0,0,0
Drive 2	5,0,1,0
Drive 3	5,0,2,0

Also assume that you requested that the tape be mounted on drive 1.

If the SCSI coordinates for the drive are configured correctly, the Administration Console Device Monitor shows that the tape is mounted on drive 1.

If the Device Monitor shows that the tape is mounted on a different drive, the SCSI coordinates for that drive are not correctly configured. For example, if the Device Monitor shows the tape mounted on drive 2, the SCSI coordinates for drive 1 are incorrect. Replace the drive 1 SCSI coordinates (5,0,0,0) with the correct SCSI coordinates (5,0,1,0) for drive 2. You also know that the SCSI coordinates for drive 2 are incorrect. Possibly, the SCSI coordinates were swapped during configuration.

Use the robotic test utility to unload and unmount the tape from drive 1. Repeat the test for each drive.

The following point applies only to NetBackup Enterprise Server.

If the data path to the drive where the tape was mounted is not on the host with direct robotic control, you may have to unload the drive with a command from another host or from the drive's front panel.

Updating the device configuration by using the wizard

Symantec recommends that you use the Device Configuration Wizard to update the NetBackup device configuration when hardware changes occur.

Update the configuration for all storage device changes. For example, if you add or delete a robot or drive or add a new SCSI adapter in a host, update the configuration.

Do not update the device configuration during backup or restore activity.

To update the device configuration by using the wizard

- 1 In the NetBackup Administration Console, select Media and Device Management > Devices.
- 2 From the list of wizards in the Details pane, click Configure Storage Devices and follow the wizard instructions.

Managing robots

You can perform various tasks to manage your robots.

Changing robot properties

Use the following procedure to change the configuration information for a robot.

To change robot properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices > Robots**.
- 2 In the Robots pane, select the robotic library you want to change.
- 3 Click **Edit > Change**.
- 4 In the Change Robot dialog box, change the properties as necessary.
The properties you can change depend on the robot type, the host type, and the robot control.
See [“Robot configuration options”](#) on page 222.
- 5 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box. If you intend to make other changes, click **No**; you can restart the Device Manager after you make the final change.

If you restart the Device manager, any backups, archives, or restores that are in progress also may be stopped.

Configuring a robot to operate in manual mode

You can configure NetBackup so that storage unit mount requests are displayed in the **Device Monitor** if the robot or drive is down. Pending requests appear in the **Device Monitor**, and you can assign these mount requests to drives manually.

See [“About pending requests for storage units”](#) on page 705.

To configure a robot so that storage unit mount requests appear in the Device Monitor

- ◆ Set the robot to operate in Pend If Robot Down (PIRD) mode by using the following command:

```
installpath\Volmgr\bin\tpconfig -update -robot robot_number -pird  
yes
```

Deleting a robot

Use the following procedure to delete a robot or robots when the media server is up and running.

Any drives that are configured as residing in a robot that you delete are changed to standalone drives.

Any media in the deleted robot is also moved to standalone. If the media is no longer usable or valid, delete it from the NetBackup configuration.

See [“Deleting a volume”](#) on page 274.

If the media server is down or the host has failed and cannot be recovered, you can delete its robots by using a different procedure.

See [“Deleting all devices from a media server”](#) on page 207.

To delete a robot

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 Select **Robots** in the tree pane.
- 3 In the **Robots** pane, select the robot or robots you want to delete.
- 4 On the **Edit** menu, select **Delete**.
- 5 At the prompt, click **Yes**.

Moving a robot and its media to a new media server

Use the following process to move a robot and its media from one server (the *old_server*) to a different media server (the *new_server*).

Table 7-5 Move a robot and media to a new server overview

Task	Procedure
Determine which tapes on the <i>old_server</i> contain NetBackup images that have not expired:	Run the following <code>bpmedialist</code> command: <code>bpmedialist -mlist -l -h old_server</code> The <code>-l</code> option produces one line of output per tape.
Move the tapes in the robot that is attached to the <i>old_server</i> to non-robotic status (standalone).	See “Moving volumes by using the Actions menu” on page 287.

Table 7-5 Move a robot and media to a new server overview (*continued*)

Task	Procedure
<p>Move the media logically from the <i>old_server</i> to the <i>new_server</i>.</p>	<p>If both the <i>old_server</i> and the <i>new_server</i> are at NetBackup 6.0 or later, run the following command:</p> <pre>bpmedia -movedb -allvolumes -oldserver <i>old_server</i> -newserver <i>new_server</i></pre> <p>If either server runs a NetBackup version earlier than 6.0, run the following command for each volume that has active images:</p> <pre>bpmedia -movedb -ev <i>media_ID</i> -oldserver <i>old_server</i> -newserver <i>new_server</i></pre> <p>For the media that has active images, see the <code>bpmedialist</code> command output from the first step of this process.</p>
<p>Configure NetBackup so that restore requests are directed to the <i>new_server</i>.</p>	<p>See “Forcing restores to use a specific server” on page 130.</p>
<p>Shut down both the <i>old_server</i> and the <i>new_server</i>.</p>	<p>See the vendor's documentation.</p>
<p>Disconnect the robot from the <i>old_server</i>.</p>	<p>See the vendor's documentation.</p>
<p>Connect the robot to the <i>new_server</i>. Verify that the operating system on the new media server recognizes the robots.</p>	<p>See the vendor's documentation.</p>
<p>Use the NetBackup Device Configuration Wizard to add the robots and drives to the media servers.</p>	<p>See “Configuring robots and drives by using the wizard” on page 219.</p>
<p>Create the appropriate NetBackup storage units.</p>	<p>See “Creating a storage unit using the Actions menu” on page 368.</p>
<p>Inventory the robots that are attached to the <i>new_server</i>. The inventory updates the location of all tapes in the robot.</p>	<p>See “Updating the volume configuration with a robot's contents” on page 314.</p>

Managing tape drives

You can perform various tasks to manage tape drives.

Changing a drive comment

You can change the comment associated with a drive. Drive comments appear in the **Drive Status** pane.

To change a drive comment

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive or select multiple drives.
- 4 On the **Actions** menu, select **Change Drive Comment**. The dialog box shows the current comment (if any is currently configured).
- 5 (Shared Storage Option.) For a shared drive, select the host and the device path to the selected drive that you want to change. You can change the comment for any or all of the host and the device paths.
- 6 Add a comment or change the current drive comment.
See [“NetBackup naming conventions”](#) on page 719.
- 7 Click **OK**.

About downed drives

NetBackup downs a drive automatically when read or write errors surpass the threshold within the time window. The default drive error threshold is 2. That is, NetBackup downs a drive on the third drive error in the default time window (12 hours).

Common reasons for write failures are dirty write heads or old media. The reason for the action is logged in the NetBackup error catalog (view the Media Logs report or the All Log Entries report). If NetBackup downs a device, it is logged in the system log.

You can use the NetBackup `nbeemmcmd` command with the `--drive_error_threshold` and `-time_window` options to change the default values.

Additional information about `nbeemmcmd` is available.

See the *NetBackup Commands* guide.

To reverse a down action, use the NetBackup Device Monitor to set the device to Up.

See [“Changing a drive operating mode”](#) on page 243.

Changing a drive operating mode

Usually you do not need to change the operating mode of a drive. When you add drive, NetBackup sets the drive state to UP in Automatic Volume Recognition (AVR) mode. Other operating mode settings are used for special purposes.

The drive operating mode is displayed and changed in the **Device Monitor** window.

To change the mode of a drive

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive or select multiple drives.
- 4 From the **Actions** menu, choose the command for the new drive operating mode.

Note that **Up Drive, Operator control** applies only to standalone drives.

- 5 If the drive is configured with multiple device paths or is a shared drive (Shared Storage Option), a dialog box appears that contains a list of all device paths to the drive. Select the path or paths to change.
- 6 Click **OK**.

Changing a tape drive path

Use the following procedure to change a drive path.

See [“Changing a drive path operating mode”](#) on page 244.

To change a drive path

- 1 In the **Change Tape Drive** dialog box, select the drive path.
- 2 Click **Change**.
- 3 In the **Change Path** dialog box, configure the properties for the drive path.

The properties you can change depend on drive type, server platform, or NetBackup server type.

See [“About SCSI reserve on drive paths”](#) on page 235.

See [“Add Path options”](#) on page 235.

Changing a drive path operating mode

A **Drive Paths** pane in the NetBackup Administration Console **Device Monitor** shows path information for drives if one of the following is true:

- Multiple (redundant) paths to a drive are configured
- Any drives are configured as shared drives (Shared Storage Option)

To change a drive path operating mode

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the Drive Paths pane, select a path or select multiple paths.
- 4 On the **Actions** menu, choose a command for the path action, as follows:
 - **Up Path**
 - **Down Path**
 - **Reset Path**

Changing tape drive properties

Use the following procedure to change the configuration information for a drive.

To change drive properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices > Drives**.
- 2 In the details pane, select the drive you want to change.
- 3 Click **Edit > Change**.
- 4 In the Change Tape Drive dialog box, change the properties of the drive.

The properties depend on the drive type and host server type.

See [“Tape drive configuration options”](#) on page 229.

- 5 After you change the properties, click **OK**.
- 6 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box. If you intend to make other changes, click **No**; you can restart the Device Manager after you make the final change.

If you restart the Device Manager, any backups, archives, or restores that are in progress also may be stopped.

The initial drive status is UP, so the drive is available as soon as you restart the Device Manager.

Changing a tape drive to a shared drive

Change a drive to a shared drive by adding paths to a currently configured drive.

To configure and use a shared drive, a Shared Storage Option license is required on each master server and media server.

To change a drive to a shared drive

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 Select **Drives** in the tree pane.
- 3 Select the drive you want to change in the Drives pane.
- 4 Click **Edit > Change**.
- 5 In the **Change Tape Drive** dialog box, click **Add**.
- 6 In the **Add Path** dialog box, configure the properties for the hosts and paths that share the drive.

Cleaning a tape drive from the Device Monitor

When you add a drive to NetBackup, you configure the automatic, frequency-based cleaning interval.

Also, you can perform an operator-initiated cleaning of a drive regardless of the cleaning frequency or accumulated mount time of the drive. However, appropriate cleaning media must be added to NetBackup.

After you clean a drive, reset the mount time.

See [“Resetting the mount time”](#) on page 247.

Additional information about drive cleaning and TapeAlert is available.

See the *NetBackup Administrator’s Guide for Windows, Volume II*.

Drive cleaning functions can also be performed from the **Activity Monitor**.

See [“Cleaning tape drives from the Activity Monitor”](#) on page 700.

To clean a tape drive

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**.
- 2 If a license that activates disk based features is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select the drive to clean.
- 4 On the **Actions** menu, select **Clean Now**. NetBackup initiates drive cleaning regardless of the cleaning frequency or accumulated mount time.

Clean Now resets the mount time to zero, but the cleaning frequency value remains the same. If the drive is a stand-alone drive and it contains a cleaning tape, NetBackup issues a mount request.

- 5 For a shared drive (Shared Storage Option), do the following:

In the list of hosts that share the drive, choose only one host on which the function applies. The **Clean Now** function can take several minutes to complete, so the cleaning information in the **Drive Details** dialog box may not be updated immediately.

Deleting a drive

Use the following procedure to delete a drive or drives when the media server is up and running.

If the media server is down or the host has failed and cannot be recovered, you can delete its drives by using a different procedure.

See [“Deleting all devices from a media server”](#) on page 207.

To delete a drive

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 Select **Drives** in the tree pane.
- 3 Select the drive or drives that you want to delete from the **Drives** pane.
- 4 On the **Edit** menu, select **Delete**.
- 5 At the prompt, click **Yes**.

Resetting a drive

Resetting a drive changes the state of the drive.

Usually you reset a drive when its state is unknown, which occurs if an application other than NetBackup uses the drive. When you reset the drive, it returns to a known state before use with NetBackup. If a SCSI reservation exists on the drive, a reset operation from the host that owns the reservation can help the SCSI reservation.

If the drive is in use by NetBackup, the reset action fails. If the drive is not in use by NetBackup, NetBackup tries to unload the drive and set its run-time attributes to default values.

Note that a drive reset does not perform any SCSI bus or SCSI device resets.

Use the following procedure to reset a drive.

To reset a drive

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive or select multiple drives.
- 4 Select **Actions > Reset Drive**. If the drive is in use by NetBackup and cannot be reset, restart the NetBackup Job Manager to free up the drive.
- 5 Determine which job controls the drive (that is, which job writes to or reads from the drive).
- 6 In the NetBackup Administration Console, expand **NetBackup Management > Activity Monitor** and on the **Jobs** tab, cancel the job.
- 7 In the **Activity Monitor**, restart the NetBackup Job Manager, which cancels all NetBackup jobs in progress.

Resetting the mount time

You can reset the mount time of the drive. Reset the mount time to zero after you perform a manual cleaning.

To reset the mount time

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive.

- 4 Select **Actions > Reset Mount Time**. The mount time for the selected drive is set to zero.
- 5 If you use the Shared drive (Shared Storage Option), do the following:
In the list of hosts that share the drive, choose only one host on which the function applies.

Setting drive cleaning frequency

When you add a drive to NetBackup, you configure the automatic, frequency-based cleaning interval. You can use the **Device Monitor** to change the cleaning frequency that was configured when you added the drive..

To set the cleaning frequency

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive.
- 4 On the **Actions** menu, select **Set Cleaning Frequency**.
- 5 Enter a time (hours) or use the arrow controls to select the number of mount hours between drive cleaning.

Cleaning Frequency is not available for the drives that do not support frequency-based cleaning. This function is not available for shared drives.

The drive cleaning interval appears in the **Drive Details** dialog box (**Actions > Drive Details**).

Viewing drive details

You can obtain detailed information about drives (or shared drives), such as drive cleaning, drive properties, drive status, host, and robotic library information.

Use the following procedure to view the drive details.

To view the drive details

- 1 In the NetBackup Administration Console, select **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the Drive Status pane, select a drive.

4 Select **Actions > Drive Details**.

5 The following applies only to NetBackup Enterprise Server:

If you use the Shared drive for shared drives, you can view the drive control mode and drive index for each host that shares a drive. You also can view a list of hosts that share a drive.

Performing device diagnostics

Diagnostic functions let you run and manage drive and robot diagnostic tests. Diagnostics are executed in an ordered sequence to verify the functionality of hardware devices. These tests can help you to troubleshoot drive or robot problems.

About device diagnostic tests

NetBackup diagnostic functions let you run and manage diagnostic tests. Diagnostics are performed in an ordered sequence to verify the functionality of hardware devices. These tests can help you to troubleshoot and drive problems.

Running a robot diagnostic test

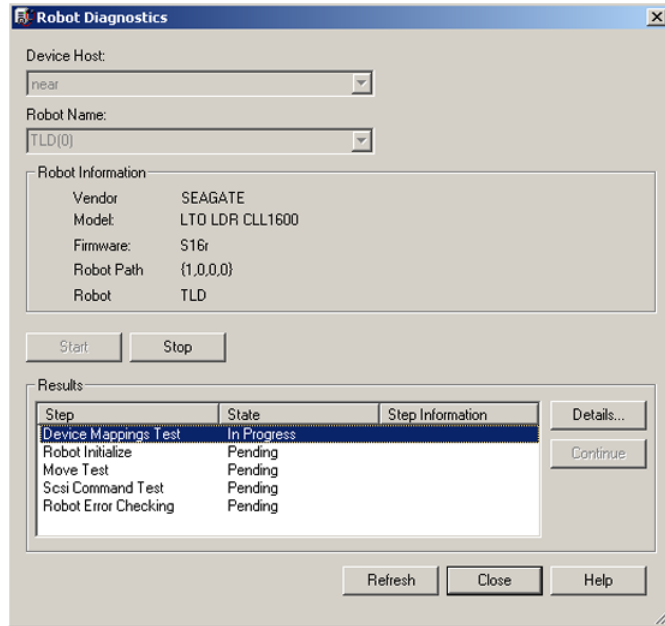
Use this procedure to run diagnostic tests on TLD or TL8 robotic libraries.

Ensure that the library to be tested is properly configured for use with NetBackup. The existing NetBackup robotic control daemons or processes are used for the test.

Note: NetBackup does not support diagnostic tests for API-attached robotic tape libraries and other types of SCSI-attached libraries.

To run a robot diagnostic test

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Robot Diagnostics**.



- 3 In the **Robot Diagnostics** dialog box, select the media server that is the **Device Host** for the robot that you want to test.
- 4 In the **Robot Name** field, select the robot that you want to diagnose.
- 5 Click **Start** to start the diagnostic tests.

The **Results** window shows results of each step in the test.

Operator intervention is required if the **State** column of the **Results** window contains **Waiting**. For example, a test step may prompt you to load a new tape into a drive before the test can continue.

- 6 If operator intervention is required, select the test step in the **Results** window and click **Details** to determine what you must do. Complete the requested operation task and then click **Continue** in the **Test Details** dialog box to resume the test

To stop a test and change the device

- 1 Click **Stop**.

The test ends after it performs any necessary clean-up work and updates the test records to reflect that the test run has been stopped.

- 2 In the **Device Host** and the **Robot Name** boxes, select the host and the robot that you want to test.
- 3 Click **Start** to restart the diagnostic test.

Running a tape drive diagnostic test

NetBackup diagnostic functions let you run and manage diagnostic tests. Diagnostics are performed in an ordered sequence to verify the functionality of hardware devices. These tests can help you to troubleshoot drive problems.

To run a tape drive diagnostic test

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Drive Diagnostics**.
- 3 In the **Drive Diagnostics** dialog box, select the media server that contains the drive that you want to test in the **Device Host** box.
- 4 In the **Drive Name** box, select the drive.
- 5 Click **Start** to start the diagnostic tests.

For robotic drives, the test media is loaded automatically.

For a stand-alone drive, insert the prelabeled test tape that is shown in the **Step Information** column of the **Results** window.

The **Results** window shows results of each step in the test.

- 6 If operator intervention is required, the State column of the Results window displays Waiting. For example, a test step may require that you to load a new tape into a drive before the test can continue.

Complete the intervention and then click **Continue**.

Select the test step in the **Results** window and click **Details** to determine what you must do. Complete the requested operation task and then click **Continue** in the **Test Details** dialog box to resume the test

To stop a test and change the device

- 1 Click **Stop**.

The test ends after it performs any necessary clean-up work and updates the test records to reflect that the test run has been stopped.

- 2 In the **Device Host** and the **Drive** boxes, select the host and the drive that you want to test.
- 3 Click **Start** to restart the diagnostic test.

Managing a diagnostic test step that requires operator intervention

Operator intervention is required if the **Status** column of the **Results** display contains **Waiting**. For example, a test step may prompt for a new tape to be loaded into a drive before the test continues.

To manage a diagnostic step

- 1 Complete the requested operations task.
- 2 Click **Continue** to resume the test.

If you clicked **Details** for a test step that requires operator intervention, you can click **Continue** from the **Test Details** dialog box.

Obtaining detailed information for a diagnostic test step

You can get information for a test step at any time during the test.

To obtain detailed information for a diagnostic test step

- 1 Select a test step in the **Results** display.
- 2 Click **Details**. A dialog box appears that displays information for the step.

The information includes a brief explanation of the checks that are performed by a specific step and the instructions that are associated with any step that requires manual intervention. For example, a step may prompt for a new tape to be loaded into a tape drive before the diagnostic session continues.

- 3 Click **Close** to return to the **Device Diagnostics** dialog box.

Verifying the device configuration

Verify the device configuration by running the Device Configuration Wizard. However, some details of a device configuration cannot be validated without attempting tape mounts. Use the NetBackup `robtest` utility to mount tapes and validate the configuration.

To verify robots and drives by using the wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

Replacing a device

NetBackup can be configured to run an automated form of device discovery during `ltid` startup (which is the default behavior for Windows).

See `AUTO_PATH_CORRECTION` in the *NetBackup Administrator's Guide for Windows, Volume II*.

NetBackup recognizes if you change a device because the serial number of the new device is different than the serial number of the old device. NetBackup updates the device configuration automatically.

NetBackup recognizes device changes as follows:

- When the Device Manager (`ltid`) performs automatic path correction.
- When the Windows Plug-n-Play feature performs serial number checks.

[Table 7-6](#) describes the process to replace a device on a single host.

[Table 7-7](#) describes the process to replace a shared device.

In some circumstances, NetBackup may be unable to determine the correct serial number in a small number of tape drives and robotic libraries. For example, NetBackup may configure serialized devices as unserialized or configure a device with the wrong serial number. If so, a device may be unusable (such as the tape drive may be downed).

To resolve such a problem, do one of the following actions:

- Configure the new device by using the **Device Configuration Wizard**.
See [“Configuring robots and drives by using the wizard”](#) on page 219.
The server operating system must recognize the device before you can configure it in NetBackup. Device configuration can require remapping, rediscovery, and possibly a restart of the operating system.
See the *NetBackup Device Configuration Guide*.
- Disable the automated device discovery by using the `vm.conf` file `AUTO_PATH_CORRECTION` option.

Table 7-6 To replace a device on a single host

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See “Changing a drive operating mode” on page 243.
Replace the device. Specify the same SCSI ID for the new device as the old device.	See the vendor's documentation.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 243.
If either of the following are true, configure the new device by using the Device Configuration Wizard : <ul style="list-style-type: none"> ■ You replaced a drive with a different drive type. ■ You replaced a serialized drive with an unserialized drive. 	See “Configuring robots and drives by using the wizard” on page 219.

Table 7-7 To replace a shared device

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See “Changing a drive operating mode” on page 243.
Replace the device. Specify the same SCSI ID for the new device as the old device.	See the vendor's documentation.
Produce a list of new and missing hardware.	The following command scans for new hardware and produces a report that shows the new and the replaced hardware: <pre>install_path\Veritas\Volmgr\bin\tpautoconf -report_disc</pre>
Ensure that all servers that share the new device are up and that all NetBackup services are active.	See “Starting or stopping a service” on page 693.

Table 7-7 To replace a shared device (*continued*)

Task	Instructions
Read the serial number from the new device and update the EMM database.	<p>If the device is a robot, run the following command:</p> <pre><i>install_path</i>\Veritas\Volmgr\bin\tpautoconf -replace_robot <i>robot_number</i> -path <i>robot_path</i></pre> <p>If the device is a drive, run the following commands:</p> <pre><i>install_path</i>\Veritas\Volmgr\bin\tpautoconf -replace_drive <i>drive_name</i> -path <i>path_name</i></pre>
<p>If the new device is an unserialized drive, run the NetBackup Device Configuration Wizard on all servers that share the drive.</p> <p>If the new device is a robot, run the NetBackup Device Configuration Wizard on the server that is the robot control host.</p>	See “Configuring robots and drives by using the wizard” on page 219.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 243.

Updating device firmware

By default, NetBackup recognizes if you update the firmware of a device.

[Table 7-8](#) is an overview of how to update device firmware.

Table 7-8 How to update device firmware

Task	Instructions
If the device is a drive, change the drive state to DOWN.	See “Changing a drive operating mode” on page 243.
Update the firmware.	See the vendor's documentation.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 243.

About the NetBackup Device Manager

The NetBackup Device Manager processes requests to mount and unmount tapes in robotically controlled devices through the robotic control processes. If you stop and restart the Device Manager (`ltid.exe`), it stops and restarts the Volume Manager (`vmd.exe`), the automatic volume recognition process (`avrd.exe`), and any robotic processes.

Note: If you stop and restart the Device Manager, any backups, archives, or restores that are in progress may fail.

Stopping and restarting the Device Manager

Use the following procedure to stop and restart the NetBackup Device Manager.

When you make device configuration changes, NetBackup asks if you want to restart the Device Manager.

To start or stop the Device Manager Service

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Stop/Restart Device Manager Service**
- 3 Select a device host (NetBackup Enterprise Server only).
- 4 Select the action to perform.
- 5 Click **Apply** or **OK**.

By using **Apply**, you can select device hosts and actions for more than one device host.

- 6 Click **OK** to close the dialog box.

Configuring tape media

This chapter includes the following topics:

- [About adding volumes](#)
- [About tape volumes](#)
- [About NetBackup media types](#)
- [About WORM media](#)
- [Adding volumes by using the wizard](#)
- [Adding volumes by using the Actions menu](#)
- [Managing volumes](#)
- [About volume pools](#)
- [Adding a volume pool](#)
- [Managing volume pools](#)
- [About volume groups](#)
- [About media sharing](#)
- [Configuring media sharing](#)

About adding volumes

Adding volumes is a logical operation that assigns NetBackup attributes to physical media. The media can reside in storage devices already, or you can add them to the storage devices when you add them to NetBackup. How you add volumes depends on the type of volume, robotic or stand-alone.

To import Backup Exec volumes, use a procedure that is documented elsewhere in this guide.

See [“Importing Backup Exec media”](#) on page 671.

About adding robotic volumes

Robotic volumes are the volumes that are located in a robotic tape library.

Table 8-1 Methods to add robotic volumes

Method	Description
The Volume Configuration Wizard	See “Adding volumes by using the wizard” on page 265.
Robot inventory	See “Updating the volume configuration with a robot’s contents” on page 314.
The Actions menu	See “Adding volumes by using the Actions menu” on page 265.
NetBackup commands	See <i>NetBackup Commands</i> .

About adding stand-alone volumes

Stand-alone volumes are the volumes that reside in the drives that are not in a robot or are allocated for stand-alone drives.

Because NetBackup does not label volumes until it uses them, you can add volumes even though they do not reside in a drive. The additional volumes are available for use if the volume in a drive becomes full or unusable. For example, if a volume in a stand-alone drive is full or unusable because of errors, NetBackup ejects (logically) the volume. If you add other stand-alone volumes, NetBackup requests that volume; NetBackup does not generate an `out of media error`.

The easiest way to add stand-alone volumes is to use the Volume Configuration Wizard. Then, when NetBackup requests one of the volumes, insert it into the stand-alone drive and NetBackup labels it.

The `DISABLE_STANDALONE_DRIVE_EXTENSIONS` option of the `nbemmcmd` command can turn off the automatic use of stand-alone volumes.

Table 8-2 Methods to add stand-alone volumes

Method	Description
The Volume Configuration Wizard	See “Adding volumes by using the wizard” on page 265.

Table 8-2 Methods to add stand-alone volumes (*continued*)

Method	Description
The Actions menu	See “Adding volumes by using the Actions menu” on page 265.
NetBackup commands	See <i>NetBackup Commands</i> .

About tape volumes

A tape volume is a data storage tape or a cleaning tape. NetBackup assigns attributes to each volume and uses them to track and manage the volumes. Attributes include the media ID, robot host, robot type, robot number, and slot location.

Volume information is stored in the EMM database.

See [“About the Enterprise Media Manager \(EMM\) database”](#) on page 603.

NetBackup uses two volume types, as follows:

Robotic volumes	Volumes that are located in a robot.
Stand-alone volumes	Volumes that are re in or are allocated for the drives that are not in a robot.

Catalog backup volumes are not a special type in NetBackup. They are the data storage volumes that you assign to the **CatalogBackup** volume pool. To add NetBackup catalog backups, use any of the add volume methods. Ensure that you assign them to the volume pool you use for catalog backups. After you add volumes, use the NetBackup Catalog Backup wizard to configure a catalog backup policy.

See [“About NetBackup catalogs”](#) on page 599.

You can use WORM media with NetBackup.

See [“About WORM media”](#) on page 261.

About NetBackup media types

NetBackup uses media types to differentiate the media that have different physical characteristics. Each media type may represent a specific physical media type; for example, NetBackup media type of 8MM, 8MM2, or 8MM3 can represent Sony AIT media.

The NetBackup media types are also known as Media Manager media types.

Table 8-3 describes the NetBackup media types.

Table 8-3 NetBackup media types

Media type	Description
4MM	4MM cartridge tape
4MM_CLN	4MM cleaning tape
8MM	8MM cartridge tape
8MM_CLN	8MM cleaning tape
8MM2	8MM cartridge tape 2
8MM2_CLN	8MM cleaning tape 2
8MM3	8MM cartridge tape 3
8MM3_CLN	8MM cleaning tape 3
DLT	DLT cartridge tape
DLT_CLN	DLT cleaning tape
DLT2	DLT cartridge tape 2
DLT2_CLN	DLT cleaning tape 2
DLT3	DLT cartridge tape 3
DLT3_CLN	DLT cleaning tape 3
DTF	DTF cartridge tape
DTF_CLN	DTF cleaning tape
HCART	1/2 inch cartridge tape
HCART2	1/2 inch cartridge tape 2
HCART3	1/2 inch cartridge tape 3
HC_CLN	1/2 inch cleaning tape
HC2_CLN	1/2 inch cleaning tape 2
HC3_CLN	1/2 inch cleaning tape 3
QCART	1/4 inch cartridge tape

NetBackup writes media in a format that allows the position to be verified before appending new backups.

See "Media Formats" in the *NetBackup Administrator's Guide for Windows, Volume II*.

About alternate media types

Alternate media types let you define more than one type of tape in the same library. You can use the alternate types to differentiate between different physical cartridges.

The following are examples of alternate media types:

- 8MM, 8MM2, 8MM3
- DLT, DLT2, DLT3
- HCART, HCART2, HCART3

For example, if a robot has DLT4000 and DLT7000 drives, you can specify the following media types:

- DLT media type for the DLT4000 tapes
- DLT2 media type for the DLT7000 tapes

NetBackup then does not load a tape that was written in a DLT4000 drive into a DLT7000 drive and vice versa.

You must use the appropriate default media type when you configure the drives. (When you configure drives in NetBackup, you specify the default media type to use in each drive type.)

In a robot, all of the volumes (of a specific vendor media type) must be the same NetBackup media type. For example, for a TLH robot that contains 3490E media, you can assign either NetBackup HCART, HCART2, or HCART3 media type to that media. You cannot assign HCART to some of the media and HCART2 (or HCART3) to other of the media.

About WORM media

You can use WORM (Write-Once-Read-Many) media to protect key data from unwanted modification or to meet compliance regulations.

NetBackup uses the QIC/WORM tape format for WORM media. This format lets NetBackup append images to WORM tape.

See "Media Formats" in the *NetBackup Administrator's Guide for Windows, Volume II*.

Tape error recovery is disabled for WORM media. NetBackup has job resume logic, which tries to resume a job that has been interrupted (such as an interruption on the fibre channel). However, NetBackup fails a job that uses WORM media and then retries the failed job. Symantec recommends that you use checkpoint and restart for backups.

The `bp1abel` command labels only LTO-3 WORM tapes. All other WORM media cannot be labeled because the label cannot be overwritten when the media is used.

About WORM media limitations

The following are the limitations for WORM tape:

- Third-party copy backups are not supported with WORM media.
- NetBackup does not support resume logic with WORM tape. NetBackup fails a job that uses WORM media and then retries the failed job. Alternatively, if checkpoint and restart are used, NetBackup restarts the job from the last checkpoint. Symantec recommends that you use checkpoint and restart for backups.
- WORM tape is not supported with NetWare media servers.

How to use WORM media in NetBackup

Two methods exist to ensure that data that is intended for WORM media is written on WORM media.

See “[About using volume pools to manage WORM media](#)” on page 263.

See “[About using unique drive and media types to manage WORM media](#)” on page 264.

About supported WORM drives

NetBackup requires an SCSI pass-through driver to use WORM tape drives. NetBackup queries the drive to verify that drive is WORM-capable and that the media in the drive is WORM media. SCSI pass-through paths are provided on the server platforms NetBackup supports. SCSI pass-through paths may require special operating system configuration changes.

See the *NetBackup Device Configuration Guide*.

For information about the drives that NetBackup supports for WORM media, see the NetBackup Hardware Compatibility List on the Symantec support Web site:

<http://entsupport.symantec.com>

All of the vendors except Quantum require the use of special WORM media.

Quantum lets NetBackup convert standard tape media to WORM media. To use Quantum drives for WORM media on Solaris systems, modify the `st.conf` file.

Information is available about how to configure nonstandard tape drives and how to edit the `st.conf` file.

See the *NetBackup Device Configuration Guide*.

About using volume pools to manage WORM media

You can dedicate volume pools for the WORM media. This method lets a WORM-capable tape drive back up and restore standard and WORM media.

Create a new volume pool and specify WORM (uppercase letters) as the first four characters of the pool name.

See [“Adding a volume pool”](#) on page 293.

NetBackup compares the first four characters of the volume pool name to determine if it is a volume pool that contains WORM media. The first four characters must be WORM.

To disable the volume pool name verification, create the following touch file on the media server of the WORM drive:

```
install_path\netbackup\db\config\DISABLE_WORM_POOLCHECK
```

Note the following cases:

- If the drive contains WORM media and the media is in a WORM volume pool, NetBackup writes the media as WORM.
- If the drive contains WORM media and the media is not in a WORM volume pool, NetBackup freezes the media.
- If the drive contains standard media and the media is in a WORM volume pool, NetBackup freezes the media.
- If the drive contains Quantum media that has never been used or all of its NetBackup images have expired, NetBackup uses the media.

About using a WORM scratch pool

For all supported WORM-capable drives (except the Quantum drive), the scratch pool must only contain one type of media. Symantec recommends that you add the most commonly used media to the scratch pool. For example, if most NetBackup jobs use standard media, put standard media in the scratch pool.

If the scratch pool contains standard media, ensure that the WORM volume pool does not run out of media to complete backup jobs.

If the WORM volume pool runs out of media, NetBackup performs the following actions:

- Moves the standard media from the scratch pool into the WORM pool
- Loads the standard media into a WORM-capable drive
- Freezes the media

NetBackup repeats this process until all of the standard media in the scratch pool is frozen.

The opposite also is true. If a standard volume pool runs out of media and the scratch pool contains WORM media, standard backups can fail because appropriate media are unavailable.

About WORM media and the Quantum drive

When you use the Quantum drive, only one kind of media can be used as either standard media or WORM media.

If a WORM volume pool runs out of media, media are moved from the scratch volume pool into the WORM pool. NetBackup determines whether the media are configured as standard or WORM media. For a standard media volume, NetBackup reads the tape label and verifies that the media is unused or that all images are expired. NetBackup also verifies that the media is not currently assigned to a server. After verification, NetBackup configures the media as WORM media and continues with the NetBackup job.

About using unique drive and media types to manage WORM media

You can assign a different drive and media type to all WORM drives and media. For example, configure standard drives and media as HCART and WORM-capable drives and media as HCART2.

This method lets you add both types of media in the scratch pool because NetBackup selects the correct media type for the drive type.

However, because each drive is limited to back ups and restores with a specific type of media, optimal drive usage may not be achieved. For example, the WORM-capable drives cannot be used for backups with standard media even if no WORM backups are in progress.

If you do not use WORM volume pools to manage WORM media, disable the WORM volume pool name verification. To disable the volume pool name verification, create the following touch file on the media server of the WORM drive:

```
install_path\netbackup\db\config\DISABLE_WORM_POOLCHECK
```


Because Quantum drives use only a single media type, this method for managing the WORM media is unnecessary.

Adding volumes by using the wizard

The easiest way to add volumes is to use the Volume Configuration Wizard. NetBackup assigns media IDs and labels the volumes automatically.

To configure volumes by using the wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 From the list of wizards in the Details pane, click **Configure Volumes** and follow the wizard instructions.

Adding volumes by using the Actions menu

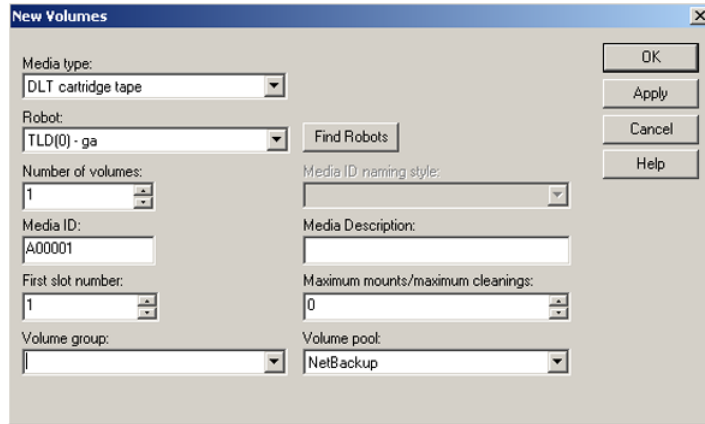
Symantec recommends that you use the Volume Configuration Wizard or the robot inventory option to add volumes.

Be careful when you specify properties. You cannot change some properties later, such as the media ID or type. If you specify them incorrectly, you must delete the volume and add it again.

To add volumes by using the Actions menu

- 1 For new volumes in a robotic library, insert them into the proper slots.
- 2 In the **NetBackup Administration Console**, expand **Media and Device Management > Media**.

3 On the **Actions** menu, select **New > New Volumes**.



4 In the **Add Volumes** dialog box, specify the attributes for the volumes.
See “[Volume properties \(add volumes\)](#)” on page 266.

5 Click **Apply** or **OK**.

If the robot has a bar code reader, NetBackup performs the following actions:

- Adds the volume to the EMM database using the specified media ID.
- Reads the bar code of each new volume.
- Adds the bar codes as attributes in the EMM database.
The **Apply** option adds the volume without closing the dialog box or refreshing the display. You can then add more volumes.

Volume properties (add volumes)

The following topics describe the properties to configure when you add volumes. The topics are arranged alphabetically.

First media ID property

This property appears only if the number of volumes is more than one.

The ID of the first volume in the range of volumes. Media IDs can be from 1 to 6 characters in length. Valid only when you add a range of volumes.

Use the same pattern that you chose in the **Media ID naming style** box. NetBackup uses the pattern to name the remaining volumes by incrementing the digits.

NetBackup allows specific characters in names.

See [“NetBackup naming conventions”](#) on page 719.

First slot number property

The number of the first slot in the robot in which the range of volumes resides. NetBackup assigns the remainder of the slot numbers sequentially.

Note: You cannot enter slot information for volumes in an API robot. The robot vendor tracks the slot locations for API robot types.

Maximum cleanings property

The number of cleanings that are allowed for a cleaning tape. The number must be greater than zero.

This number is decremented with each cleaning and when it is zero, NetBackup stops using the tape. You then must change the cleaning tape or increase the number of cleanings that remain.

Additional information about drive cleaning is available.

See the *NetBackup Administrator's Guide, Volume II*.

Media ID naming style property

This property appears only if the number of volumes is more than one.

The style to use to name the range of volumes. Media IDs can be from 1 to 6 characters in length. Using the pattern, NetBackup names the remaining volumes by incrementing the digits.

NetBackup media IDs for an API robot must match the bar code on the media. For API robots, NetBackup supports bar codes from 1 to 6 characters. Therefore, obtain a list of the bar codes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software.

NetBackup allows specific characters in names.

See [“NetBackup naming conventions”](#) on page 719.

Media ID property

This property appears only if the number of volumes is one.

The ID for the new volume. Media IDs can be from 1 to 6 characters in length.

Media IDs for an API robot must match the bar code on the media (for API robots, NetBackup supports bar codes from 1 to 6 characters). Therefore, obtain a list of

the bar codes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software.

NetBackup allows specific characters in names.

See [“NetBackup naming conventions”](#) on page 719.

Media description property

A description of the media, up to 25 character maximum.

NetBackup allows specific characters in names.

See [“NetBackup naming conventions”](#) on page 719.

Number of volumes property

The number of volumes to add. For a robotic library, enough slots must exist for the volumes.

Robot property

The robotic library to add the volumes to.

To add volumes for a different robot, select a robot from the dropdown list. The list shows robots on the selected host that can contain volumes of the selected media type.

To find a robot that does not appear in the **Robot** box, click **Find Robots** to open the Find Robot dialog box.

To add volumes to a stand-alone drive, select **Standalone**.

Volume group property

If you specified a robot, select from a volume group already configured for that robot. Alternatively, enter the name for a volume group; if it does not exist, NetBackup creates it and adds the volume to it.

If you do not specify a volume group (you leave the volume group blank), the following occurs:

- Stand-alone volumes are not assigned to a volume group.
- NetBackup generates a name for robotic volumes by using the robot number and type. For example, if the robot is a TL8 and has a robot number of 50, the group name is 000_00050_TL8.

See [“About volume groups”](#) on page 296.

Volume pool property

The pool to which the volume or volumes should be assigned.

Select a volume pool you created or one of the following standard NetBackup pools:

- None.
- NetBackup is the default pool name for NetBackup.
- DataStore is the default pool name for DataStore.
- CatalogBackup is the default pool name used for NetBackup hot, online catalog backups of policy type NBU-Catalog.

When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool.

See [“About volume pools”](#) on page 292.

Managing volumes

The following sections describe the procedures to manage volumes.

Changing the group of a volume

If you move a volume physically to a different robot, change the group of the volume to reflect the move.

See [“About volume pools”](#) on page 292.

To change the volume group of a volume

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the Volumes list, select the volumes that you want to change the volume group assignment for.
- 3 On the **Actions** menu, select **Change Volume Group**.
- 4 In the **New volume group name** field, enter the name of the new volume group or select a name from the list of volume groups.
- 5 Click **OK**.

The name change is reflected in the volume list entry for the selected volumes. If you specified a new volume group (which creates a new volume group), the group appears under **Volume Groups** in the tree pane.

Changing the owner of a volume

You can change the media server or server group that owns the volume.

See [“About server groups”](#) on page 197.

See [“About media sharing”](#) on page 297.

To change the owner of a volume

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the Volumes list, select the volume that you want to change.
- 3 On the **Actions** menu, select **Change Media Owner**.
- 4 In the **Media Owner** field, select one of the following:

Any (default)	Allows NetBackup to choose the media owner. NetBackup chooses a media server or a server group (if one is configured).
None	Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Specify a server group. A server group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the drop-down list.

- 5 Click **OK**.

Changing the pool of a volume

Change the **Volume pool** property in the Change Volumes dialog box.

See [“Changing volume properties”](#) on page 270.

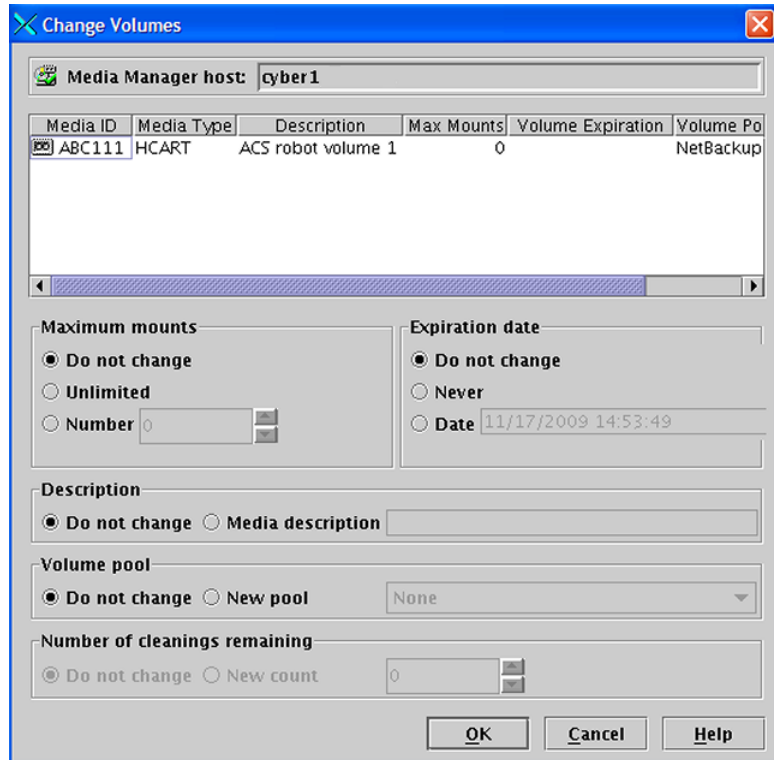
Changing volume properties

You can change some of the properties of a volume, including the volume pool.

To change volume properties

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the **Volumes** pane, select a volume or volumes.

- 3 On the **Edit** menu, select **Change**.



- 4 In the **Change Volumes** dialog box, change the properties for the volume. See [“Volume properties \(change a volume\)”](#) on page 271.

Volume properties (change a volume)

The following topics describe the volume properties that you can change.

Description property

A description of the media, up to 25 character maximum.

NetBackup allows specific characters in names as described in the following topic:

See [“NetBackup naming conventions”](#) on page 719.

Expiration date property

The following does not apply to cleaning tapes.

The date after which the volume is too old to be reliable.

When the expiration date has passed, NetBackup reads data on the volume but does not mount and write to the volume. You should exchange it for a new volume.

See [“Exchanging a volume”](#) on page 276.

When you add a new volume, NetBackup does not set an expiration date.

The expiration date is not the same as the retention period for the backup data on the volume. You specify data retention periods in the backup policies.

Maximum mounts property

The following topic does not apply to cleaning tapes.

The **Maximum mounts** property specifies the number of times that the selected volumes can be mounted.

When the limit is reached, NetBackup reads data on the volume but does not mount and write to the volume.

A value of zero (the default) is the same as **Unlimited**.

To help determine the maximum mount limit, consult the vendor documentation for information on the expected life of the volume.

Number of cleanings remaining property

The number of cleanings that are allowed for a cleaning tape. This number is decremented with each cleaning and when it is zero, NetBackup stops using the tape. You then must change the cleaning tape or increase the number of cleanings that remain.

Additional information about drive cleaning is available.

See the *NetBackup Administrator's Guide, Volume II*.

Volume pool property

The following topic does not apply to cleaning tapes.

The pool to which the volume or volumes should be assigned.

Select a volume pool you created or one of the following standard NetBackup pools:

- None.
- NetBackup is the default pool name for NetBackup.
- DataStore is the default pool name for DataStore.
- CatalogBackup is the default pool name used for NetBackup hot, online catalog backups of policy type NBU-Catalog.

When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool.

See [“About volume pools”](#) on page 292.

Deassigning a volume

NetBackup deassigns volumes as part of normal operations.

See [“About assigning and deassigning volumes”](#) on page 273.

To deassign a volume, you expire the images on the volume. After you expire a volume, NetBackup deassigns it and does not track the backups that are on it. NetBackup can reuse the volume, you can delete it, or you can change its volume pool.

See [“Expiring backup images”](#) on page 665.

You can expire backup images regardless of the volume state (Frozen, Suspended, and so on).

NetBackup does not erase images on expired volumes. You can still use the data on the volume by importing the images into NetBackup (if the volume has not been overwritten).

See [“Importing backups”](#) on page 666.

Note: Symantec recommends that you do not deassign NetBackup volumes. If you do, be certain that the volumes do not contain any important data. If you are uncertain, copy the images to another volume before you deassign the volume.

About assigning and deassigning volumes

An assigned volume is one that is reserved for exclusive use by NetBackup. A volume is set to the assigned state when either application writes data on it for the first time. The time of the assignment appears in the **Time Assigned** column for the volume in the **NetBackup Administration Console Volumes** pane. When a volume is assigned, you cannot delete it or change its volume pool.

A volume remains assigned until NetBackup deassigns it.

NetBackup deassigns a volume only when the data is no longer required, as follows:

- For regular backup volumes, when the retention period has expired for all the backups on the volume.
- For catalog backup volumes, when you stop using it for catalog backups.

To determine which application currently uses a volume, see the **Application** column of the **Volumes** pane.

Deleting a volume

You can delete volumes from the NetBackup configuration. You cannot delete a volume if it is still assigned.

For example, if any of the following situations apply, you may want to delete the volume:

- A volume is no longer used and you want to recycle it by relabeling it with a different media ID.
- A volume is unusable because of repeated media errors.
- A volume is past its expiration date or has too many mounts, and you want to replace it with a new volume.
- A volume is lost and you want to remove it from the EMM database.

After a volume is deleted, you can discard it or add it back under the same or a different media ID.

Before you delete and reuse or discard a volume, ensure that it does not have any important data. You cannot delete NetBackup volumes if they are assigned.

See [“Deassigning a volume”](#) on page 273.

To delete volumes

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media**.
- 2 In the **Volumes** pane, select the volume or volumes that you want to delete. You cannot delete a volume if it is still assigned.
- 3 On the **Edit** menu, select **Delete**.
- 4 In the **Delete Volumes** dialog box, click **OK**.
- 5 Remove the deleted volume or volumes from the storage device.

Erasing a volume

You can erase the data on a volume if the following are true:

- The volume is not assigned.
- The volume contains no valid NetBackup images.

After NetBackup erases the media, NetBackup writes a label on the media.

If you erase media, NetBackup cannot restore or import the data on the media.

If a volume contains valid NetBackup images, deassign the volume so NetBackup can label it.

See “[Deassigning a volume](#)” on page 273.

Table 8-4 Types of erase

Type of erase	Description
SCSI long erase	<p>Rewinds the media and the data is overwritten with a known data pattern. An SCSI long erase is also called a secure erase because it erases the recorded data completely.</p> <p>Note: A long erase is a time-consuming operation and can take as long as 2 hours to 3 hours. For example, it takes about 45 minutes to erase a 4-mm tape on a standalone drive</p>
SCSI quick erase	<p>Rewinds the media and an erase gap is recorded on the media. The format of this gap is drive dependent. It can be an end-of-data (EOD) mark or a recorded pattern that the drive does not recognize as data.</p> <p>Some drives do not support a quick erase (such as QUANTUM DLT7000). For the drives that do not support a quick erase, the new tape header that is written acts as an application-specific quick erase.</p>

Note: NetBackup does not support erase functions on NDMP drives.

To erase a volume

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the volumes pane, select a volume or volumes that you want to erase.
If you select multiple volumes, they must all be in the same robot
- 3 Select either **Actions > Quick Erase** or **Actions > Long Erase**.
- 4 In the erase dialog box, specify the name of the media server to initiate the erase operation.

To overwrite any existing labels on the media, do not select **Verify media label before performing operation**.

5 Click **OK**.

A dialog box warns you that this action is irreversible.

6 Click **OK** if you are certain you want to start the erase action.

A dialog box reminds you to use the **Activity Monitor** to view the progress and status of the action. (For many types of drives, you may not be able to cancel a label or erase media job from the **Activity Monitor**.) Click **OK**.

If you selected **Verify media label before performing operation** and the actual volume label does not match the expected label, the media is not erased.

Exchanging a volume

You should exchange a volume (replace one volume with another volume) if a volume meets any of the following conditions:

- Full (in this case, to exchange a volume means to remove the volume from a robotic tape library).
- Past the maximum number of mounts.
- Old (past the expiration date).
- Unusable (for example, because of repeated media errors).

The exchange volumes procedures in the following subsections depend on whether you want to reuse the old media ID or not.

Exchanging a volume and using a new media ID

Use this procedure when the follow are true:

- The volume contains current and valid NetBackup images.
- You require slots in the robotic library for additional backups, duplications, vault functions, or other purposes.

Table 8-5 Exchange a volume and using a new media ID

Task	Instructions
Move the volume to another location If the volume is in a robotic library, remove it from the robotic library and move it to a stand-alone group.	See “About moving volumes” on page 286.

Table 8-5 Exchange a volume and using a new media ID (*continued*)

Task	Instructions
<p>Add a new volume or move an existing volume in as a replacement for the volume you removed.</p> <p>If you add a new volume, specify a new media ID. Specify the same values for the other attributes as the removed volume (such as robotic residence, volume pool, and the media type).</p>	See “About adding volumes” on page 257.
<p>Physically replace the old volume.</p> <p>Do not delete the old volume in case you need to retrieve the data on the volume.</p>	Beyond the scope of the NetBackup documentation.

Exchanging a volume and using the old media ID

You can exchange a volume and reuse the same media ID, which may be convenient in some instances.

Reuse a media ID only if all data on the old volume is not required and you recycle or discard the volume.

Warning: If you exchange a media ID for a volume that has unexpired backup images, serious operational problems and data loss may occur.

Table 8-6 Exchange a volume and use the old media ID

Task	Instructions
Delete the volume.	See “Deleting a volume” on page 274.
Remove the old volume from the storage device. Physically add the new volume to the storage device.	See “About injecting and ejecting volumes” on page 279.
Add the new volume to the NetBackup volume configuration and specify the same attributes as the old volume, including the old media ID.	See “About adding volumes” on page 257.
Set a new expiration date for the volume.	See “Changing volume properties” on page 270.

Table 8-6 Exchange a volume and use the old media ID (*continued*)

Task	Instructions
Optionally, label the volume. Although you do not have to label the volume, the label process puts the media in a known state. The external media label matches the recorded media label, and the mode is known to be compatible with the drives in the robotic library.	See “Labeling a volume” on page 285.

About frozen media

A frozen volume is unavailable for future backups. A frozen volume never expires, even after the retention period ends for all backups on the media. The media ID is never deleted from the NetBackup media catalog, and it remains assigned to NetBackup. A frozen volume is available for restores. If the backups have expired, you must import the backups first.

See [“Importing backups”](#) on page 666.

NetBackup freezes media automatically when read or write errors surpass the threshold within the time window. The default media error threshold is 2. That is, NetBackup freezes media on the third media error in the default time window (12 hours).

NetBackup also freezes a volume if a write failure makes future attempts at positioning the tape unreliable.

Common reasons for write failures are dirty write heads or old media. The reason for the action is logged in the NetBackup error catalog (view the Media Logs report or the All Log Entries report).

You can use the NetBackup `nbemmcmd` command with the `-media_error_threshold` and `-time_window` options to change the default values.

Additional information about `nbemmcmd` is available.

See the *NetBackup Commands* guide.

To reverse a freeze action, use the `bpmedia` command to unfreeze the volume.

See [“Freezing or unfreezing a volume”](#) on page 278.

Freezing or unfreezing a volume

NetBackup freezes volumes under circumstances.

You can freeze or unfreeze a volume manually.

To freeze or unfreeze media

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the **Volumes** list, select the volume that you want to freeze or unfreeze.
- 3 On the **Actions** menu, select **Freeze** or **Actions > Unfreeze**.
- 4 In the dialog box, click **OK**.

About injecting and ejecting volumes

Media access port (MAP) functionality differs between robotic libraries. For many libraries, NetBackup opens and closes the MAP as needed. However, some libraries have the front-panel inject and the eject functions that conflict with NetBackup's use of the media access port. And for other libraries, NetBackup requires front-panel interaction by an operator when using the media access port.

Read the operator manual for the library to understand the media access port functionality. Some libraries may not be fully compatible with the inject and eject features of NetBackup unless properly handled. Other libraries may not be compatible at all.

Injecting volumes

You can inject volumes into the robots that contain media access ports.

Any volumes to be injected must be in the media access port before the operation begins. If no volumes are in the port, you are not prompted to place volumes in the media access port and the update operation continues.

Each volume in the MAP is moved into the robotic library. If the MAP contains multiple volumes, they are moved to empty slots in the robotic library until the media access port is empty or all the slots are full.

After the volume or volumes are moved, NetBackup updates the volume configuration.

Some robots report only that media access ports are possible. Therefore, **Empty media access port prior to update** may be available for some robots that do not contain media access ports.

Inject volumes into the robots that contain media access ports

- 1 Load the volumes in the MAP.
- 2 Inventory the robot
See [“Updating the volume configuration with a robot's contents”](#) on page 314.
- 3 Select **Empty media access port prior to update** on the Robot Inventory dialog box.

Ejecting volumes

Eject single or multiple volumes.

Volumes that reside in multiple robots can be ejected. Multiple eject dialog boxes appear for each robot type.

Operator intervention is required only if the media access port is too small to contain all of the selected volumes. For these robot types, you are prompted to remove the media from the media access port so the eject can continue with the remaining volumes.

See [“About media ejection timeout periods”](#) on page 282.


To eject volumes

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media**.
- 2 In the **Volumes** pane, select one or more volumes that you want to eject.
- 3 On the **Actions** menu, select **Eject Volumes From Robot**.

Depending on the robot type, one of the following dialog boxes appears:

- **Eject Volume** (singular)
- **Eject Volumes** (plural)

- 4 If the **Eject Volume** (singular) dialog box appears, click **OK** to eject the volumes.

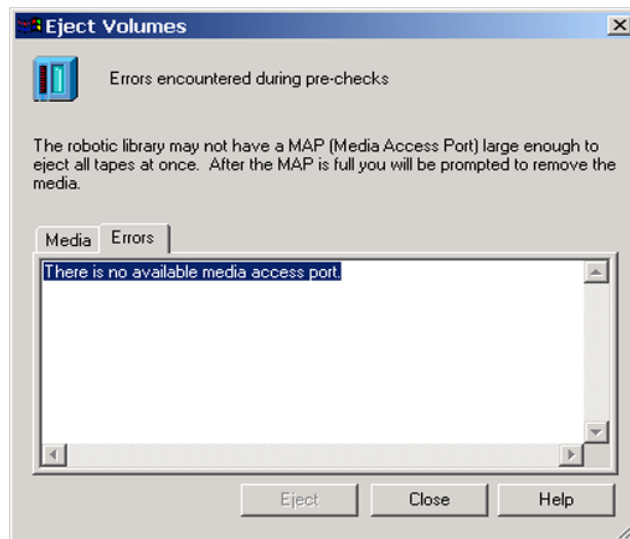


If you select multiple volumes, operator action is required to remove each volume after each eject (a dialog box prompts you to remove volumes).

- 5 If the Eject Volumes (plural) dialog box appears, continue by reading the following:
- After NetBackup completes prechecks for the eject, the **Media** tab of the **Eject Volumes** dialog box shows the volumes that you selected to eject.
 - If no errors occur, the **Errors** tab is empty.
 - If an error occurs or a hardware limitation exists, the eject may not be possible; if so, the **Errors** tab is opened.

The following classes of errors can occur:

- For serious errors, the **Eject** button is not active. Correct the error to eject the media.
- For other errors, the **Errors** tab shows an explanation of the error. Either continue the eject action (**Eject**) or exit (**Close**) depending on the type of error.



- 6 ACS and TLM robots only: In the **Eject Volumes** dialog box, select the media access port to use for the eject.
- 7 In the **Eject Volumes** dialog box, click **Eject** to eject the volumes.

The robotic library may not contain a media access port large enough to eject all of the selected volumes. For most robot types, you are prompted to remove the media from the media access port so the eject can continue with the remaining volumes.

About media ejection timeout periods

The media ejection period (the amount of time before an error condition occurs) varies depending on the capability of each robot.

[Table 8-7](#) shows the ejection timeout periods for robots.

Table 8-7 Media ejection timeout periods

Robot types	Timeout period
Applies only to NetBackup Enterprise Server: Automated Cartridge System (ACS) Tape Library Multimedia (TLM)	One week
Tape Library 8MM (TL8) Tape Library DLT (TLD)	30 minutes.
Applies only to NetBackup Enterprise Server: Tape Library Half-inch (TLH)	None. The robot allows an unlimited amount of time to remove media.

Note: If the media is not removed and a timeout condition occurs, the media is returned to (injected into) the robot. Inventory the robot and eject the media that was returned to the robot.

Some robots do not contain media access ports. For these robots, the operator must remove the volumes from the robot manually.

Note: After you add or remove media manually, use NetBackup to inventory the robot.

About rescanning and updating bar codes

You can rescan the media in a robot and then update NetBackup with the bar codes of that media.

You should rescan and update only in certain circumstances.

Note: Rescan and update bar codes does not apply to volumes in API robot types.

See [“About bar codes”](#) on page 325.

See [“When to rescan and update”](#) on page 283.

See [“When not to rescan and update”](#) on page 283.

When to rescan and update

Rescan and update bar codes only to add the bar codes that are not in the EMM database.

For example: if you add a new volume but do not insert the tape into the robot, NetBackup does not add the bar code to the database. Use this command to add the bar code after you insert the tape into the robotic library.

When not to rescan and update

Do not rescan and update to correct the reports that show a media ID in the wrong slot.

To correct that problem, perform one of the following actions:

- Logically move the volume by selecting a volume and then on the **Actions** menu select **Move**.
- Logically move the volume by updating the volume configuration.
See [“Updating the volume configuration with a robot's contents”](#) on page 314.
- Physically move the volume into the correct slot.

To obtain an inventory of the robot without updating the bar code information in the database, inventory the robot and use the show contents option.

See [“Showing the media in a robot”](#) on page 309.

Rescanning and updating bar codes

Use the following procedure to rescan the media in a robot and update NetBackup with the bar codes.

Note: Rescan and update bar codes does not apply to volumes in API robot types.

See [“About rescanning and updating bar codes”](#) on page 282.

To rescan bar codes and update the EMM database

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media > Robots**.
- 2 Select the robotic library that contains the volumes that you want to scan and update.

- 3 In the Volumes pane, select the volumes.
- 4 On the **Actions** menu, select **Rescan/Update Barcodes**.
The rescan begins immediately.

About labeling NetBackup volumes

When NetBackup labels a volume, it writes a record on the magnetic tape of the volume; the record (or label) includes the NetBackup media ID.

Normally, NetBackup controls the labeling of its volumes. In most cases, NetBackup labels a volume the first time it is used for a backup.

The volume label depends on whether or not the media has a bar code, as follows:

- If the robot supports bar codes and the media has bar codes, NetBackup uses the last six characters of the bar code for the media ID.
To change this default action, specify and select specific characters by using Media ID generation rules.
See [“Configuring media ID generation rules”](#) on page 333.
- For volumes without bar codes, by default NetBackup uses a prefix of the letter A when it assigns a media ID to a volume (for example, A00001).
To change the default prefix, use the `MEDIA_ID_PREFIX` configuration option in the `vm.conf` file.
See the *NetBackup Administrator's Guide for Windows, Volume II*.

Media are not labeled automatically in the following situations:

- They were last used for NetBackup catalog backups.
Do not label catalog backup volumes unless they are no longer used for catalog backups.
- They contain data from a recognized non-NetBackup application and NetBackup is configured to prohibit media overwrite for that media type.

To label these media, the following must be true:

- NetBackup has not assigned them
- They contain no valid NetBackup images

About prelabeling media

It can be beneficial to prelabel media. A successful label operation validates that the media is usable, compatible, and is not write-protected.

In addition, the recorded label can assist with media management if the following occurs:

- The media is misplaced,
- The bar code label or external label is gone or damaged,
- You use the physical inventory utility (vmphyinv) to manage media.

Labeling a volume

If a volume contains valid NetBackup images, deassign the volume so that it can be labeled.

See [“Deassigning a volume”](#) on page 273.

If you want to label media and assign specific media IDs (rather than allow NetBackup to assign IDs), use the `bp1label` command.

Note: If you label a volume, NetBackup cannot restore or import the data that was on the media after you label it.

Note: For many types of drives, you may not be able to cancel a label job from the Activity Monitor.

See [“About labeling NetBackup volumes”](#) on page 284.

To label media

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 In the **Volumes** pane, select a volume or the volumes that you want to label. If you select multiple volumes, they all must be in the same robot.
- 3 On the **Actions** menu, select **Label**.
- 4 In the Label dialog box, specify the following properties for the label operation.

Media server	Enter tname of the media server that controls the drive to write the label.
Verify label before performing operation	Select this option to verify that the media in the drive is the expected media. To overwrite any existing labels on the media, do not select Verify media label before performing operation .

- 5 Click **OK**.
- 6 In the warning dialog box, click **OK**.

If you selected **Verify media label before performing operation** and the actual volume label does not match the expected label, the media is not relabeled.

About moving volumes

When you move volumes in or out of a robotic library or from one robot to another, move the volumes physically and logically, as follows:

- Physically move volumes by inserting or by removing them. For some robot types, use the NetBackup inject and eject options.
- Logically move volumes using NetBackup, which updates the EMM database to show the volume at the new location.

When you move volumes from one robotic library to another robotic library, perform the following actions:

- Move the volumes to stand-alone as an intermediate step
- Move the volumes to the new robotic library

The following types of logical moves are available:

- Move single volumes
- Move multiple volumes
- Combinations of single and multiple volumes
- Move volume groups

You cannot move volumes to an invalid location (for example, move DLT media to an 8-mm robot).

Symantec recommends that you perform moves by selecting and by moving only one type of media at a time to a single destination.

The following are several examples of when to move volumes logically:

- A volume is full in a robotic library and no slots are available for new volumes in the robotic library. Move the full volume to stand-alone, remove it from the robot, then configure a new volume for the empty slot or move an existing volume into that slot. Use the same process to replace a defective volume.
- Moving volumes from a robotic library to an off-site location or from an off-site location into a robotic library. When you move tapes to an off-site location, move them to stand-alone.

- Moving volumes from one robotic library to another (for example, if a library is down).
- Changing the volume group for a volume or volumes.

Moving volumes by using the robot inventory update option

Use this procedure for the following:

- To move volumes within a robot.
The robot must have a bar code reader and the volumes must contain readable bar codes.
- To remove volumes from a robot.
Use this procedure even if the volumes do not contain bar codes or if the robot does not have a reader.

To move volumes by using a robot inventory update

- 1 Physically move the volumes to their new location.
- 2 On the **Actions** menu, select **Inventory Robot**.
- 3 In the Robot Inventory dialog box, select **Update volume configuration**.
- 4 Select other options as appropriate.
See [“About robot inventory”](#) on page 302.

Moving volumes by using the Actions menu

If you move a volume to a robotic library that has a bar code reader, NetBackup updates the EMM database with the correct bar code.

To move volumes by using the Actions menu

- 1 Physically move the volumes to their new location.
- 2 In the **NetBackup Administration Console**, expand **Media and Device Management > Media**.
- 3 In the **Volumes** pane, select the volumes that you want to move.
- 4 On the **Actions** menu, select **Move**.

If you selected volumes of different media types or volume residences, a **Move Volumes** dialog box appears for each residence and each media type.

See [“Multiple Move Volumes dialog boxes may appear”](#) on page 288.

- 5 In the **Move Volumes** dialog box, specify the properties for the move.
See [“Move Volumes properties”](#) on page 288.

Multiple Move Volumes dialog boxes may appear

If you selected volumes of different media types or volume residences, a Move Volumes dialog box appears for each residence and each media type.

For example, you select two full volumes to move out of a robotic library and two stand-alone volumes to move in as replacements. A dialog box appears for the two full volumes and another dialog box for the two replacement volumes. In this example, complete both move dialog boxes to perform the move (complete the move for the volumes that are full first).

Note: These multiple **Move Volumes** dialog boxes may appear on top of each other and need to be repositioned.

Figure 8-1 and Figure 8-2 show examples of moving multiple types or residences.

Figure 8-1 Move volumes to stand-alone

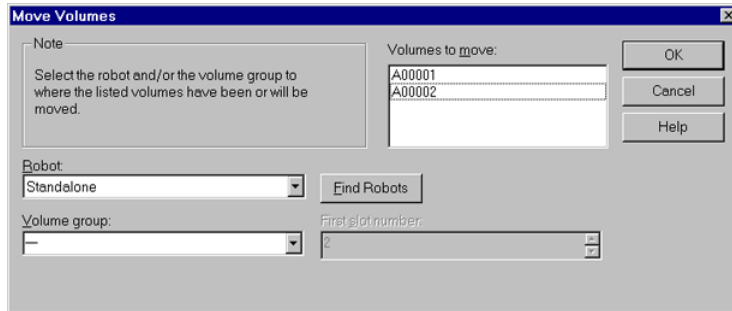
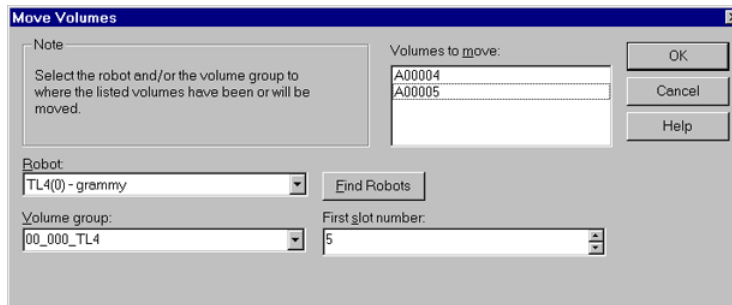


Figure 8-2 Move volumes to the robot



Move Volumes properties

Table 8-8 describes the properties to configure in the **Move Volumes** dialog box.

Table 8-8 Move volumes properties

Property	Description
First slot number	<p>For volumes in a robotic library, specify the first slot number to be used in the destination robotic library. By default, this box shows the slot number where the volume currently resides. NetBackup assigns the remainder of the slot numbers sequentially.</p> <p>Note: You cannot enter slot information for volumes in an API robot. The robot vendor tracks the slot locations for these robot types.</p>
Device host	<p>The Device host specifies the name of the device host where the robot is defined.</p> <p>For single volumes, the current location of the volume appears.</p> <p>NetBackup Enterprise Serve only: To select a robot on another device host, select from the list of device hosts shown.</p>
Find Robots	<p>Use Find Robots to find a robot that does not appear in the Robot box (for example, a new robot).</p>
Robot	<p>Robot specifies the new robotic library for the volumes. You can specify a different robot as the destination or Standalone.</p> <p>The list shows the robot type, number, and control host for any robot that already has at least one volume in the EMM database.</p>
Volume group	<p>Enter or select the volume group to assign to the volumes.</p> <p>If you leave the volume group blank, the following occurs:</p> <ul style="list-style-type: none"> ■ Stand-alone volumes are not assigned a volume group. ■ Robotic volumes are assigned to a new volume group; NetBackup generates the name by using the robot number and type. For example, if the robot is a TL8 and has a robot number of 50, the group name is 000_00050_TL8.
Volume is in a robotic library	<p>To inject a volume into a robotic library, select Volume is in a robotic library.</p> <p>Select a robot and the slot number for the volume.</p> <p>See “Robot property” on page 268.</p> <p>See “First slot number property” on page 267.</p> <p>To eject a volume from a robot, clear Volume is in a robotic library.</p>
Volumes to move	<p>The Volumes to move section of the dialog box shows the media IDs of the volumes that you selected to move.</p>

Recycling a volume

If you recycle a volume, you can use either the existing media ID or a new media ID.

Caution: Recycle a volume only if all NetBackup data on the volume is no longer needed or if the volume is damaged and unusable. Otherwise, you may encounter serious operational problems and a possible loss of data.

Recycling a volume and using the existing media ID

NetBackup recycles a volume and returns it to the volume rotation when the last valid image on the volume expires.

To recycle a volume that contains unexpired backup images, you must deassign the volume.

See [“Deassigning a volume”](#) on page 273.

Recycling a volume and using a new media ID

Recycle a volume if it is a duplicate of another volume with the same media ID. Also recycle a volume if you change how you name volumes and you want to match the barcodes on the volume.

Table 8-9 Recycling a volume and using a new media ID

Task	Instructions
Physically remove the volume from the storage device.	See “Ejecting volumes” on page 280.
If the volume is in a robotic library, move it to stand-alone.	See “About moving volumes” on page 286.
Record the current number of mounts and expiration date for the volume.	See the values in the Media (Media and Device Management > Media in the Administration Console).
Delete the volume entry.	See “Deleting a volume” on page 274.

Table 8-9 Recycling a volume and using a new media ID (*continued*)

Task	Instructions
Add a new volume entry.	See “Adding volumes by using the Actions menu” on page 265. Because NetBackup sets the mount value to zero for new volume entries, you must adjust the value to account for previous mounts. Set the maximum mounts to a value that is equal to or less than the following value: The number of mounts that the manufacturer recommends minus the value that you recorded earlier.
Physically add the volume to the storage device.	See “Injecting volumes” on page 279.
Configure the number of mounts	Set the number of mounts to the value you recorded earlier by using the following command: <pre>install_path\Volmgr\bin\vmchange -m media_id -n number_of_mounts</pre>
Set the expiration date to the value you recorded earlier.	See “Changing volume properties” on page 270.

Suspending or unsuspending a volume

You cannot use a suspended volume for backups until retention periods for all backups on it have expired. At that time, NetBackup deletes the suspended volume from the NetBackup media catalog and unassigns it from NetBackup.

A suspended volume is available for restores. If the backups have expired, import the backups first.

To suspend or unsuspend media

- 1 In the NetBackup Administration Console, select **Media and Device Management > Media**.
- 2 In the Volumes list, select the volume that you want to freeze or unfreeze.
- 3 Select **Actions > Suspend** or **Actions > Unsuspend**.
- 4 In the dialog box, click **OK**.

About volume pools

A volume pool identifies a set of volumes by usage. Volume pools protect volumes from access by unauthorized users, groups, or applications. When you add media to NetBackup, you assign them to a volume pool (or assign them as standalone volumes, without a pool assignment).

By default, NetBackup creates the following volume pools:

NetBackup	The default pool to which all backup images are written (unless you specify otherwise).
DataStore	For DataStore use.
CatalogBackup	For NetBackup catalog backups.
None	For the volumes that are not assigned to a pool.

You can add other volume pools. For example, you can add a volume pool for each storage application you use. Then, as you add volumes to use with an application, you assign them to that application's volume pool. You can also move volumes between pools.

You also can configure a scratch pool from which NetBackup can transfer volumes when a volume pool has no volumes available.

The volume pool concept is relevant only for NetBackup storage units and does not apply to disk storage units.

Examples of volume pool usage are available.

See the *NetBackup Administrator's Guide for Windows, Volume II*.

About scratch volume pools

The scratch pool is an optional pool that contains the media that NetBackup can allocate to other pools as needed. If you configure a scratch pool, NetBackup moves volumes from that scratch pool to other pools that do not have volumes available.

Only one scratch pool is allowed. You cannot add a scratch pool if one exists.

You cannot change the **NetBackup** or **DataStore** pools to be scratch volume pools.

If you create a scratch pool, be aware of the following conditions:

- If the scratch pool contains assigned volumes, these volumes remain in the scratch pool.
NetBackup does not move assigned volumes to other pools as it does with unassigned volumes.

- NetBackup does not assign volumes while they are in a scratch pool. For example if a NetBackup policy or schedule specifies the scratch pool, all requests for those volumes are denied.
- NetBackup returns expired media to the scratch volume pool automatically (media that is returned must have been originally in the same scratch pool).
- To use NetBackup to manage the allocation of volumes to volume pools, do the following:
 - Create volume pools as required, but do not add any volumes to the pools.
 - Define a scratch pool and add all of the volumes to it. NetBackup moves volumes to the other pools as volumes are needed.

Adding a volume pool

Use this procedure to add a new volume pool. After you add a new pool, add volumes to it by adding new volumes to NetBackup or by changing the pool of existing volumes.

To add a volume pool

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media**.
- 2 On the **Actions** menu, select **New > New Volume Pool**.
- 3 In the New Volume Pool dialog box, specify the attributes for the volume pool. See [“Volume pool properties”](#) on page 293.

Volume pool properties

You can specify various properties for a volume pool.

The following are the properties you can configure for volume pools, either when you add a new pool or change an existing one.

Catalog backup pool property

Select this option to use this volume pool for hot, online backups of the NetBackup catalog. This check box creates a dedicated catalog backup pool to be used for **NBU-Catalog** policies. A dedicated catalog volume pool facilitates quicker catalog restore times.

Multiple catalog backup volume pools are allowed.

Description property

The **Description** option is a brief description of the volume pool.

Maximum number of partially full media property

The following option does not apply to the None pool, catalog backup pools, or scratch volume pools.

Specifies the number of partially full media to allow in the volume pool for each of the unique combinations of the following in that pool:

- Robot
- Drive type
- Retention level

The default value is zero, which does not limit the number of full media that are allowed in the pool.

NetBackup writes data only to the number of partially full media with a given combination of attributes in the volume pool. When the number of partially full media is reached for a given combination of attributes, NetBackup queues backup jobs until media becomes available. If a volume becomes full, NetBackup assigns another volume for use if one is available in the pool or in a scratch pool.

When the number of partially full media is reached, NetBackup queues backup jobs until media becomes available. If a media becomes full, NetBackup assigns another media for use if one is available in the pool or in a scratch pool.

The total number of partially full media in the pool may be more than the value you specify. For example, if the value is 10, each of the following can have 10 partially full volumes in the same volume pool:

- Robot TLD(0), 1/2" cartridge (HCART) drive type, 1/2" cartridge tape media type, and two-week retention level.
- Robot TLD(0), 1/2" cartridge (HCART) drive type, 1/2" cartridge tape media type, and four-week retention level.
- Robot TLD(0), 1/2" cartridge (HCART2) drive type, 1/2" cartridge tape 2 media type, and two-week retention level.
- Robot TLD(1), 1/2" cartridge (HCART) drive type, 1/2" cartridge tape media type, and two-week retention level.
- And so on.

A partially full volume that has multiple retention levels can count toward more than one combination of attributes. (That is, if all of the other attributes are the same.) For example, a volume that has both two week and four week retention

levels counts toward the first two examples in the previous list. (Again, if all of the other attributes are the same.)

Frozen, suspended, and imported media do not count against **Maximum number of partially full media**. Therefore, if you unfreeze or unsuspend the media that are partially full, media may exceed the **Maximum number of partially full media** value. Also, if you lower the **Maximum number of partially full media** value, media may exceed the limit. If media exceed the **Maximum number of partially full media** value, NetBackup uses all of the partially full media until the limit is reached.

This option lets you maximize the media usage. NetBackup writes data to the media until they are full, and assigns other media to which it writes data until they are full.

See [“About server groups”](#) on page 197.

Pool name property

The **Pool name** is the name for the new volume pool. Volume pool names are case sensitive and can be up to 20 characters.

See [“NetBackup naming conventions”](#) on page 719.

Scratch pool property

Specifies that the pool should be a scratch pool.

Symantec recommends that you use a descriptive name for the pool and use the term `scratch pool` in the description.

Add sufficient type and quantity of media to the scratch pool to service all scratch media requests that can occur. NetBackup requests scratch media when media in the existing volume pools are allocated for use.

See [“About scratch volume pools”](#) on page 292.

Managing volume pools

The following sections describe the operations you can perform to manage volume pools.

Changing the properties of a volume pool

Use this procedure to change the properties of a volume pool. The properties you can change include the pool type (scratch pool or catalog backup pool).

To change a volume pool

- 1 In the NetBackup Administration Console, select **Media and Device Management > Media > Volume Pools**.
- 2 Select a pool from the pools in the **Volume Pools** list.
- 3 Select **Edit > Change**.
- 4 In the Change Volume Pool dialog box, change the attributes for the volume pool.

See “[Volume pool properties](#)” on page 293.

Deleting a volume pool

You cannot delete any of the following pools:

- A volume pool that contains volumes
- The **NetBackup** volume pool
- The **None** volume pool
- The default **CatalogBackup** volume pool
- The **DataStore** volume pool

To delete a volume pool

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media > Volume Pools**.
- 2 Select a volume pool from the pools in the **Volume Pools** list.
- 3 Ensure that the volume pool is empty. If the pool is not empty, change the pool name for any volumes in the pool. If the volumes are not needed, delete them.
- 4 On the **Edit >** menu, select **Delete**.
- 5 Click **Yes** or **No** in the confirmation dialog box.

About volume groups

A volume group identifies a set of volumes that reside at the same physical location. The location can be either the robot in which the volumes reside, standalone storage, or off-site storage if you use the NetBackup Vault option.

When you add media to NetBackup, NetBackup assigns all volumes in a robot to that robot's volume group. Alternatively, you can assign the media to a different group.

Volume groups are convenient for tracking the location of volumes, such as the case when a volume is moved off site. Volume groups let you perform operations on a set of volumes by specifying the group name rather than each individual media ID of each volume. Operations include moves between a robotic library and a standalone location or deletions from NetBackup.

If you move a volume physically, you also must move it logically. A logical move means to change the volume attributes to show the new location.

The following are the rules for assigning volume groups:

- All volumes in a group must be the same media type.
However, a media type and its corresponding cleaning media type are allowed in the same volume group (such as DLT and DLT_CLN).
- All volumes in a robotic library must belong to a volume group.
You cannot add volumes to a robotic library without specifying a group or having Media Manager generate a name for the group.
- The only way to clear a volume group name is to move the volume to standalone and not specify a volume group.
- More than one volume group can share the same location.
For example, a robotic library can contain volumes from more than one volume group and you can have more than one standalone volume group.
- All volumes in a group must be in the same robotic library or be standalone.
That is, you cannot add a group (or part of a group) to a robotic library if it already exists in another robotic library.

Examples of volume group usage are available.

See the *NetBackup Administrator's Guide for Windows, Volume II*.

About media sharing

Media sharing allows media servers to share media for write purposes (backups).

You can allow all media servers to share media, or you can configure a group of servers to share media.

See [“About server groups”](#) on page 197.

Media sharing provides the following benefits:

- Increases media utilization by reducing the number of partially full media.
- Reduces media-related expenses because fewer media are required and fewer media are vaulted (NetBackup Vault option).

- Reduces administrative overhead because you inject fewer scratch media into the robotic library.
- Increases media life because media are mounted fewer times. Media are not repositioned and unmounted between write operations from different media servers.

Reducing media mounts requires appropriate hardware connectivity between the media servers that share media and the drives that can write to that media. Appropriate hardware connectivity may include Fibre Channel hubs or switches, SCSI multiplexors, or SCSI-to-fibre bridges.

See “[Configuring media sharing](#)” on page 298.

Note: The access control feature of Sun StorageTek ACSLS controlled robots is not compatible with media sharing. Media sharing restricts volume access by the requesting hosts IP address. Use caution when you implement media sharing in an ACSLS environment.

Configuring media sharing

You can configure the following media sharing:

- Unrestricted media sharing
- Media media sharing with server groups

Configuring unrestricted media sharing

Unrestricted media sharing means that all NetBackup media servers and NDMP hosts in your NetBackup environment can share media for writing.

Note: Do not use unrestricted media sharing and media sharing server groups. If you use both, NetBackup behavior is undefined.

To configure unrestricted media sharing

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the **Details** pane, double-click the master server.
- 3 Select **Media**

4 Select **Enable Unrestricted Media Sharing for All Media Servers.**

If you allow unrestricted allow media sharing in your NetBackup environment, you do not need to create media sharing groups.

5 Click **OK.**

Configuring media sharing with a server group

Media sharing with a server group restricts the sharing to members of the group.

See [“About server groups”](#) on page 197.

Table 8-10 Configuring media sharing with a server group process overview

Task	Instructions
Ensure the appropriate connectivity between and among the media servers and robots and drives.	Beyond the scope of the NetBackup documentation.
Configure the media sharing server group.	See “Configuring a server group” on page 198.
Optionally, configure the volume pools for media sharing.	Set the Maximum number of partially full media property for those pools. See “Adding a volume pool” on page 293. See “Changing the properties of a volume pool” on page 295.
Configure backup policies that use the volume pools and media sharing groups.	Set the Policy Volume Pool and Media Owner properties of the backup policies. See “Creating a policy using the Backup Policy Configuration Wizard” on page 455.

Inventorying robots

This chapter includes the following topics:

- [About robot inventory](#)
- [When to inventory a robot](#)
- [About showing a robot's contents](#)
- [Showing the media in a robot](#)
- [About comparing a robot's contents with the volume configuration](#)
- [Comparing media in a robot with the volume configuration](#)
- [About updating the volume configuration](#)
- [Updating the volume configuration with a robot's contents](#)
- [Robot inventory options](#)
- [Configuring media settings](#)
- [About bar codes](#)
- [Configuring bar code rules](#)
- [Configuring media ID generation rules](#)
- [Configuring media type mappings](#)
- [About the physical inventory utility](#)
- [Example volume configuration updates](#)

About robot inventory

[Table 9-1](#) describes the NetBackup Administration Console robot inventory options for the robotic libraries that contain barcode readers and contain barcoded media. Robot inventory is a logical operation that verifies the presence of media, it does not inventory the data on the media.

Table 9-1 Robot inventory options

Inventory option	Description
Show contents	Displays the media in the selected robotic library; does not check or change the EMM database. See “About showing a robot's contents” on page 307.
Compare contents with volume configuration	Compares the contents of a robotic library with the contents of the EMM database but does not change the database. See “About comparing a robot's contents with the volume configuration” on page 310.
Preview changes	Compares the contents of a robotic library with the contents of the EMM database. If differences exist, NetBackup recommends changes to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 303.
Update volume configuration	Updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur. See “About updating the volume configuration” on page 312.

For the robotic libraries without bar code readers or that contain media without bar codes, you can show the contents of a robot. However, more detailed information is required to perform automated media management. Use the `vmphyinv` physical inventory utility to inventory such robots.

See [“About the physical inventory utility”](#) on page 343.

After you physically add, move, or remove volumes in a robot, use a robot inventory to update the NetBackup volume configuration.

See [“When to inventory a robot”](#) on page 303.

See [“Example volume configuration updates”](#) on page 349.

See [“How to access media and devices on other hosts”](#) on page 722.

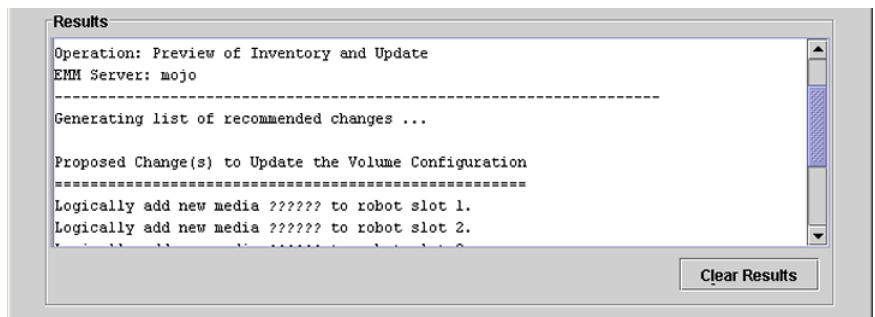
About previewing volume configuration changes

Use this option to preview the changes before you update the EMM database. This option lets ensure that all new media have bar codes before you add them to the EMM database.

If you select **Preview changes** and any recommended changes are found, a dialog box asks if you want to accept the recommended changes. If you click **Yes**, you do not need to perform a separate **Update volume configuration** operation.

Note: If you preview the configuration changes first, then update the EMM database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the EMM database, to the bar code rules, and so on.

Figure 9-1 Preview volume configuration changes (not an API robot)



See [“Updating the volume configuration with a robot's contents”](#) on page 314.

When to inventory a robot

[Table 9-2](#) describes the criteria to use to determine when to inventory a robot and which options to use for the inventory.

Table 9-2 Robot inventory criteria

Action	Inventory option to use
To add new volumes to a robot (a new volume is one that does not have a NetBackup media ID)	<p>For any robot NetBackup supports, use the Update volume configuration option.</p> <p>The update creates media IDs (based on bar codes or a prefix that you specify).</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p>
To determine if volumes were moved physically within a robot	<p>For robots with bar code readers and that contain media with bar codes, use the Compare contents with volume configuration option.</p> <p>See “Comparing media in a robot with the volume configuration” on page 311.</p>
To determine the contents of a robot	<p>Use the Show contents option to determine the media in a robot and possibly their bar code numbers.</p> <p>See “Showing the media in a robot” on page 309.</p>
To determine whether new media have bar codes before you add them to NetBackup	<p>Use the Preview changes option, which compares the contents of the robot with the NetBackup volume configuration information.</p> <p>After you examine the results, use the Update volume configuration option to update the volume configuration if necessary.</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p>

Table 9-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
<p>To insert existing volumes into a robot (an existing volume is one that already has a NetBackup media ID)</p>	<p>If the robot supports bar codes and the volume has a readable bar code, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic location. NetBackup also updates the robot host, robot type, robot number, and slot location. Specify the volume group to which the volume is assigned.</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p> <p>If the robot does not support bar codes or the volumes do not contain readable bar codes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 286.</p> <p>See “About the physical inventory utility” on page 343.</p>
<p>To move existing volumes between robotic and stand-alone (an existing volume is one that already has a NetBackup media ID)</p>	<p>If the robotic library supports bar codes and the volume has a readable bar code, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic or stand-alone location.</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p>

Table 9-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
<p>To move existing volumes from one robot to another (an existing volume is one that already has a NetBackup media ID)</p>	<p>If the robotic library supports bar codes and the volume has a readable bar code, use the Update volume configuration option. NetBackup updates the NetBackup volume configuration information.</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p> <p>If the robots do not support bar codes or the volumes do not contain readable bar codes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 286.</p> <p>See “About the physical inventory utility” on page 343.</p> <p>For either operation, perform the following updates:</p> <ul style="list-style-type: none"> ■ First move the volumes to stand-alone ■ Then move the volumes to the new robot <p>If you do not perform both updates, NetBackup cannot update the entries and writes an "Update failed" error.</p> <p>See “Example 6: Moving existing volumes between robots” on page 357.</p>
<p>To move existing volumes within a robot (an existing volume is one that already has a NetBackup media ID)</p>	<p>If the robot supports bar codes and the volume has a readable bar code, use the Update volume configuration option. NetBackup updates the residence information to show the new slot location.</p> <p>See “Updating the volume configuration with a robot’s contents” on page 314.</p> <p>If the robot does not support bar codes or if the volumes do not contain readable bar codes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 286.</p> <p>See “About the physical inventory utility” on page 343.</p> <p>See “Example 7: Adding existing volumes when bar codes are not used” on page 358.</p>

Table 9-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
To remove existing volumes from a robot (an existing volume is one that already has a NetBackup media ID)	For any robot NetBackup supports, use the Update volume configuration option to update the NetBackup volume configuration information. See “Updating the volume configuration with a robot's contents” on page 314.

About showing a robot's contents

Show contents inventories the selected robotic library and generates a report. This operation does not check or change the EMM database. Use this option to determine the contents of a robot.

The contents that appear depend on the robot type.

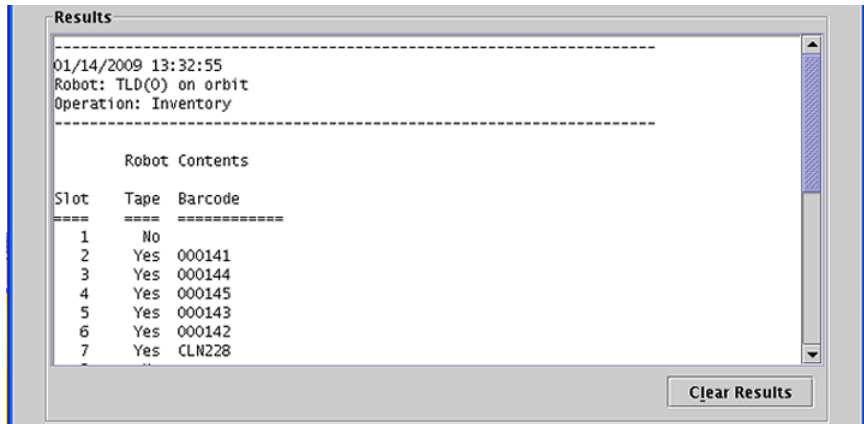
[Table 9-3](#) describes the report contents.

Table 9-3 Show contents description

Robot and media	Report contents
The robot has a barcode reader and the robot contains media with bar codes.	Shows if each slot has media and lists the barcode for the media.
The robot does not have a barcode reader or the robot contains media without bar codes.	Shows if each slot has media.
API robot.	Shows a list of the volumes in the robot. See “About inventory results for API robots” on page 308.

[Figure 9-2](#) is an example of the report.

Figure 9-2 Show contents report



See [“Showing the media in a robot”](#) on page 309.

About inventory results for API robots

[Table 9-4](#) describes the contents of the robot inventory for the API robots.

Table 9-4 API robot report contents

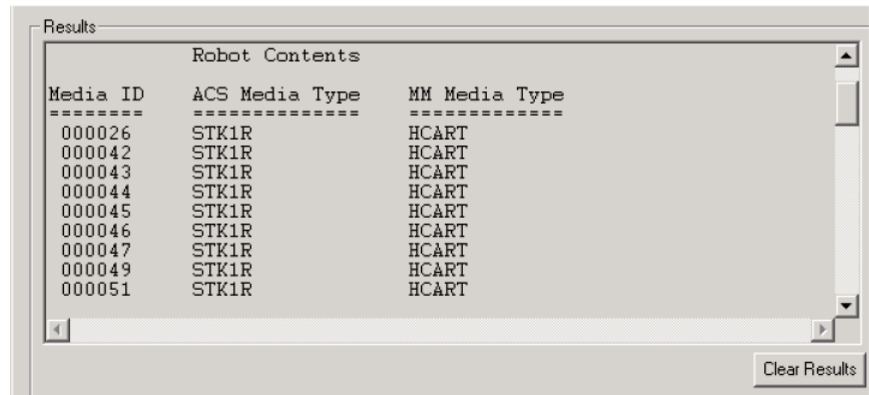
Robot type	Report contents
ACS	<p>The results, received from ACS library software, show the following:</p> <ul style="list-style-type: none"> ■ The ACS library software volume ID. The NetBackup media ID corresponds to the ACS library software volume ID. ■ The ACS media type. ■ The NetBackup Media Manager media type. ■ The mapping between the ACS library software media type and the corresponding NetBackup Media Manager media type (without considering optional bar code rules).
TLH	<p>The results, received from the Automated Tape Library (ATL) library manager, show the following:</p> <ul style="list-style-type: none"> ■ The volume serial number (volser). The Media Manager media ID corresponds to the ATL volser. ■ The ATL media type. ■ The Media Manager media type. ■ The mapping between the ATL media type and the corresponding Media Manager media type (without considering optional bar code rules).

Table 9-4 API robot report contents (*continued*)

Robot type	Report contents
TLM	<p>The results, received from the DAS/SDLC server, show the following:</p> <ul style="list-style-type: none"> ■ The volume serial number (volser). The Media Manager media ID corresponds to the DAS/SDLC volser. ■ The DAS/SDLC media type ■ The Media Manager media type. ■ The mapping between the DAS/SDLC media type and the corresponding Media Manager media type (without considering optional bar code rules).

Figure 9-3 shows the results for an ACS robot; the results for other API robots are similar.

Figure 9-3 Show contents report (API robot)



Showing the media in a robot

Use the following procedure to show the media that is in a robot.

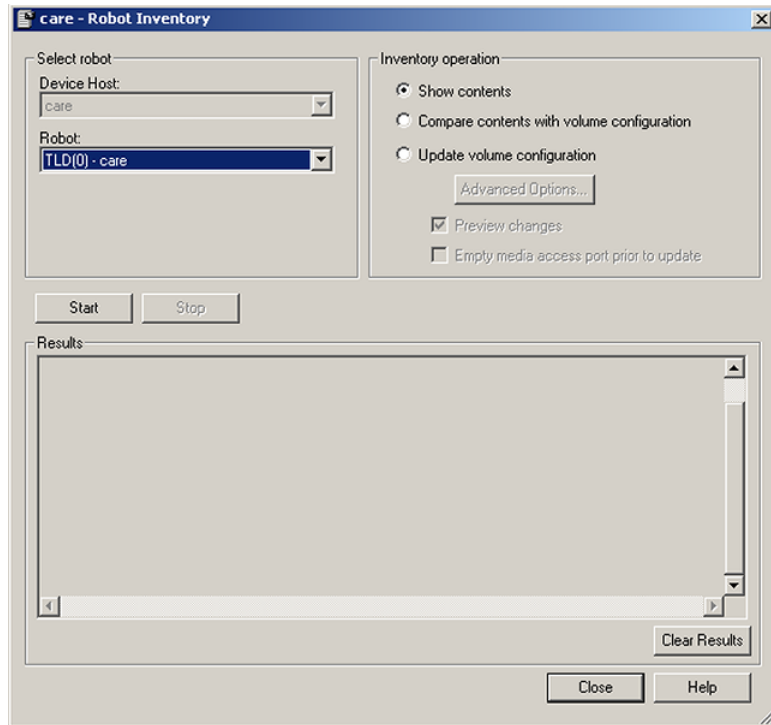
See [“About robot inventory”](#) on page 302.

See [“Robot inventory options”](#) on page 316.

To show robot contents

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Media > Robots**.
- 2 Select the robot you want to inventory.

3 On the **Actions** menu, select **Inventory Robot**.



4 In the **Robot Inventory** dialog box, select **Show contents**.

5 Click **Start** to begin the inventory.

About comparing a robot's contents with the volume configuration

Compare contents with volume configuration compares the contents of a robotic library with the contents of the EMM database. Regardless of the result, the database is not changed.

Table 9-5 Compare contents description

Robot and media	Report contents
The robot can read bar codes	The report shows the differences between the robot and the EMM database

Table 9-5 Compare contents description (*continued*)

Robot and media	Report contents
The robot cannot read bar codes	The report shows only whether a slot contains a volume If the media cave bar codes, this operation is useful for determining if volumes have been physically moved within a robot.
For API robots	The media ID and media type in the EMM database are compared to the information that is received from the vendor’s robotic library software.

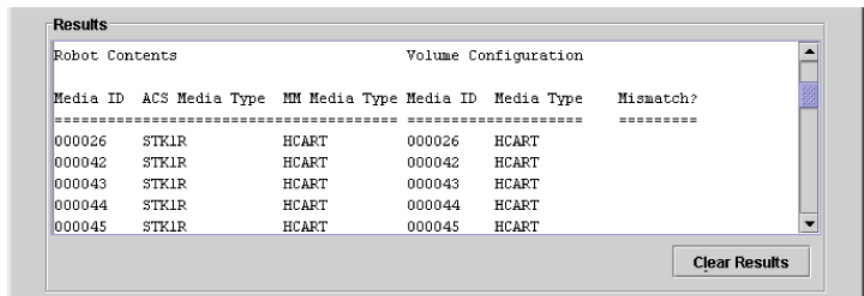
If the results show that the EMM database does not match the contents of the robotic library, perform the following actions:

- Physically move the volume.
- Update the EMM database. Use **Actions > Move** or use the **Update volume configuration** option.

See “[About updating the volume configuration](#)” on page 312.

[Figure 9-4](#) shows a sample compare report.

Figure 9-4 Compare contents report (API robot)



See “[Comparing media in a robot with the volume configuration](#)” on page 311.

Comparing media in a robot with the volume configuration

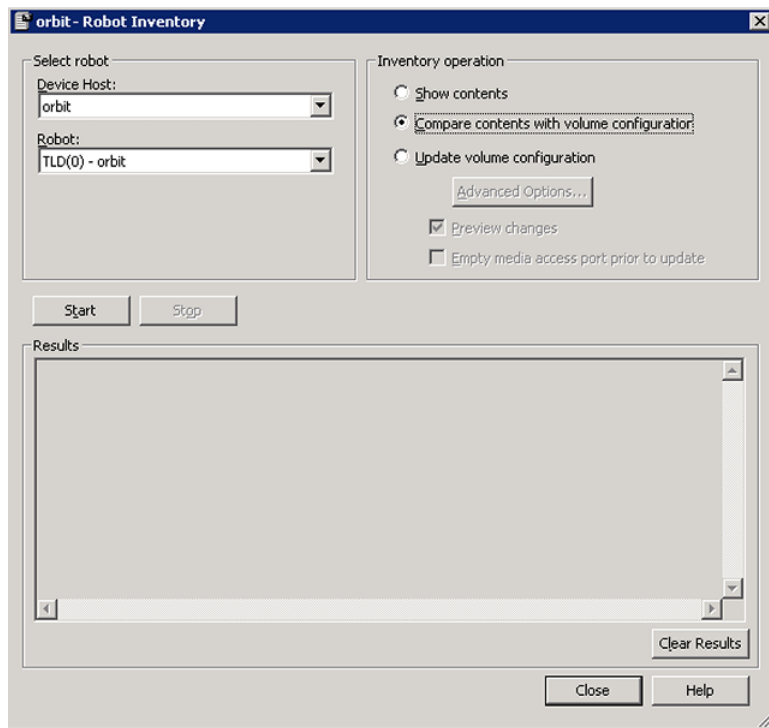
Use the following procedure to compare the media in a robot with the EMM database.

See “[About robot inventory](#)” on page 302.

See “[Robot inventory options](#)” on page 316.

To compare robot media with the database

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots**.
- 2 Select the robot you want to inventory.
- 3 On the **Actions** menu, select **Inventory Robot**.



- 4 In the **Robot Inventory** dialog box, select **Compare contents with volume configuration**.
- 5 Click **Start** to begin the inventory.

About updating the volume configuration

Update volume configuration updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur.

For a new volume (one that does not have a NetBackup media ID), the update creates a media ID. The media ID depends on the rules that are specified on the **Advanced Robot Inventory Options** dialog box.

See [“Robot inventory options”](#) on page 316.

For API robots, the update returns an error if the volume serial number or the media ID contain unsupported characters.


For robots without bar code readers, the new media IDs are based on a media ID prefix that you specify. Similarly, for volumes without readable bar codes, the new media IDs are based on a media ID prefix that you specify

[Figure 9-5](#) is an example for an ACS robot. Results for other API robots are similar.

Robot inventory update returns an error if it encounters unsupported characters in the volume serial number or media identifier from API robots.

See [“Determine robot capabilities before you update the volume configuration”](#) on page 313.

Figure 9-5 Update volume configuration for API robot report



See [“Updating the volume configuration with a robot's contents”](#) on page 314.

Determine robot capabilities before you update the volume configuration

Before you update the volume configuration, determine the following:

- If the robotic library supports bar codes
- If any new volume are inserted into the library has readable bar codes.

Check the bar code capabilities of the robotic library and the volume by comparing the robot contents with the NetBackup volume configuration.

See [“Comparing media in a robot with the volume configuration”](#) on page 311.

If the robotic library does not support bar codes or the volume does not have readable bar codes, save the results of the compare operation. The results can help you determine a media ID prefix if you use the **Media Settings** tab of the **Advanced Options** dialog box to assign a prefix.

Updating the volume configuration with a robot's contents

Use the following procedure to update the EMM database with the contents of a robot.

See [“About robot inventory”](#) on page 302.

You can change the default settings and rules that NetBackup uses to name and assign attributes to new media. For most configurations, the default settings work well. Change the settings only if the configuration has special hardware or usage requirements.

[Table 9-6](#) shows the rules you can configure.

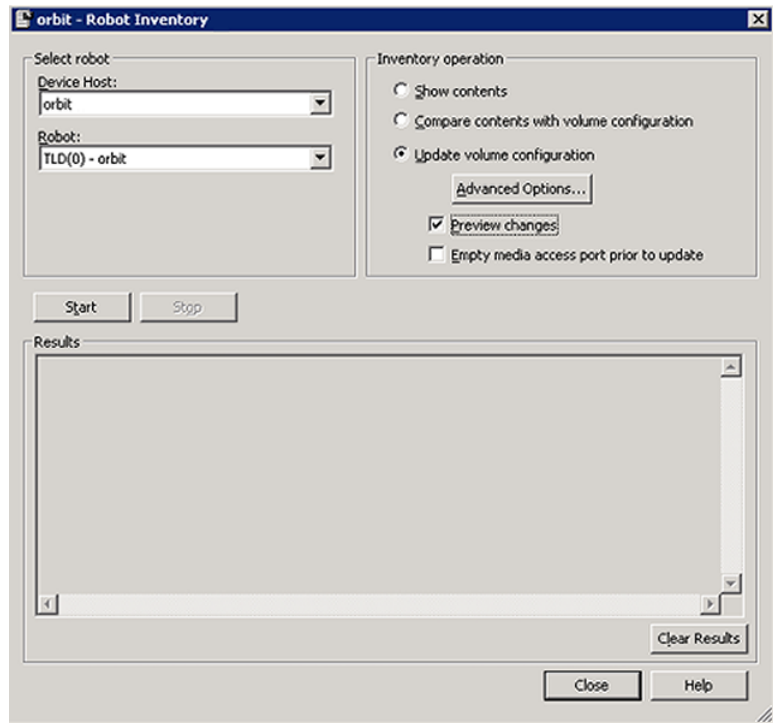
Table 9-6 Attributes for new media

What	Where
Media settings.	See “Configuring media settings” on page 317.
Bar code rules.	See “Configuring bar code rules” on page 329.
Media ID generation rules.	See “Configuring media ID generation rules” on page 333.
Map media for API robots.	See “Configuring media type mappings” on page 336.

To update the volume configuration with a robot's contents

- 1 If necessary, insert new volume(s) into the robotic library.
- 2 In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots**.
- 3 Select the robot you want to inventory.

- 4 On the **Actions** menu, select **Inventory Robot**.



- 5 In the **Robot Inventory** dialog box, select **Update volume configuration**.
- 6 By default, **Preview changes** is selected. To update without previewing changes, clear **Preview changes**.

Note: If you preview the configuration changes first, then update the EMM database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the EMM database, to the bar code rules, and so on.

- 7 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, click **Advanced Options**.

[Table 9-6](#) shows the settings and rules you can configure.

- 8 Click **Start** to begin the inventory.

Robot inventory options

The following are available robot inventory options by using the NetBackup Administration Console:

- Advanced options** The **Advanced Options** option is active if **Update volume configuration** is selected.
- It opens the **Advanced Robot Inventory Options** dialog box, from which you can configure more options.
- See [“Configuring media settings”](#) on page 317.
- See [“Configuring bar code rules”](#) on page 329.
- See [“Configuring media ID generation rules”](#) on page 333.
- See [“Configuring media type mappings”](#) on page 336.
- For most configurations, the default settings work well. Change the settings only if the configuration has special hardware or usage requirements.
- Device host** The **Device host** option is the host that controls the robot. In NetBackup Enterprise Server, specify the device host.
- Empty media access port prior to update** The **Empty media access port prior to update** operation is active only for the robots that support that function.
- To inject volumes in the robot’s media access port into the robot before you begin the update, select **Empty media access port prior to update**.
- The volumes to be injected must be in the media access port before the operation begins. If you select **Empty media access port prior to update** and the MAP is empty, you are not prompted to place volumes in the media access port.
- Note:** If you use NetBackup to eject volumes from the robot, remove the volumes from the media access port before you begin an inject operation. Otherwise, if the inject port and eject port are the same, the ejected volumes may be injected back into the robotic library.
- Robot** Use the **Robot** option to select a robot to inventory.
- If you selected a robot in the Administration Console, that robot appears in this field.
- Show contents** Displays the media in the selected robotic library; does not check or change the EMM database.
- See [“About showing a robot’s contents”](#) on page 307.

Compare contents with volume configuration	Compares the contents of a robotic library with the contents of the EMM database but does not change the database. See “About comparing a robot's contents with the volume configuration” on page 310.
Preview changes	Compares the contents of a robotic library with the contents of the EMM database. If differences exist, NetBackup recommends changes to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 303.
Update volume configuration	Updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur. See “About updating the volume configuration” on page 312.

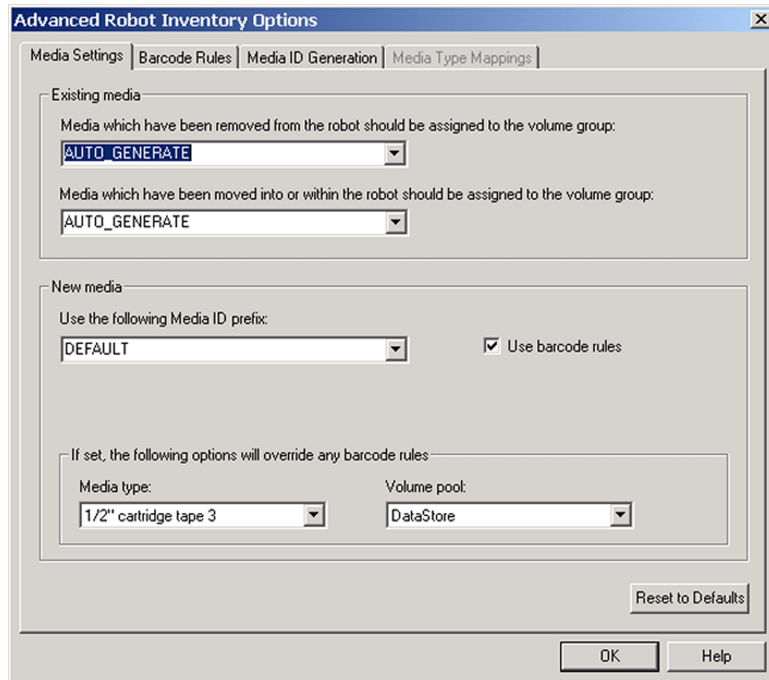
Configuring media settings

Use the **Media Settings** tab of the **Advanced Robot Inventory Options** dialog box to perform the following actions:

- For existing media, specify the volume group
- For new media, specify the media settings

To configure media settings

- 1 In the **Robot Inventory** dialog box, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media Settings** tab.



- 3 Configure the settings.
See [“Media settings - existing media”](#) on page 318.
See [“Media settings - new media”](#) on page 320.

Media settings - existing media

For the media that already exists in your volume configuration, you can specify the volume group for two conditions: if the media are removed from the robot or if the media are moved into or within the robot.

Media which have been removed from the robot

The volume group to assign to the media that are removed from the robot.

The list contains the following selections:

- AUTO GENERATE NetBackup automatically generates a new volume group.
- DEFAULT If there is an existing group with a compatible residence for the volume, the volume is added to that group. If a suitable volume group does not exist, NetBackup generates a new volume group name.
- NO VOLUME GROUP The media are not assigned to a volume group.

Other selections may be available, depending on the setting of the **Media type** field as follows:

- DEFAULT The selection includes the volume groups that are valid for the robot's default media type.
- Other than DEFAULT The selection includes the volume groups that are valid for the specified media type.
 To specify a volume group other than DEFAULT, enter a volume group name or select one from the list.

See ["Media type"](#) on page 321.

Media which have been moved into or within the robot

The volume group to assign to the existing media that you have inserted into the robot (or moved to a new location within the robot).

The list contains the following selections:

- AUTO GENERATE NetBackup automatically generates a new volume group.
- DEFAULT If there is an existing group with a compatible residence for the volume, the volume is added to that group. If a suitable volume group does not exist, NetBackup generates a new volume group name.

Other selections may be available, depending on the setting of the **Media type** field as follows:

- DEFAULT The selection includes the volume groups that are valid for the robot's default media type.
- Other than DEFAULT The selection includes the volume groups that are valid for the specified media type.
 To specify a volume group other than DEFAULT, enter a volume group name or select one from the list.

If the robotic library contains multiple media types, Symantec recommends a `DEFAULT` setting. If you specify a volume group and volumes of different media types were moved into or within the robot, the new update fails. Volumes of different media types cannot have the same volume group.

See “[Media type](#)” on page 321.

Media settings - new media

For new media in the robot to add to your volume configuration, specify the attributes for the new media.

See “[Use the following Media ID prefix](#)” on page 320.

See “[Use barcode rules](#)” on page 321.

See “[Media type](#)” on page 321.

See “[Volume pool](#)” on page 324.

Use the following Media ID prefix

If the robot supports bar codes and the volume has readable bar codes, a prefix is not required because NetBackup creates media IDs automatically.

If either of the following conditions exist, specify a media ID prefix for any new media :

- The robot does not support bar codes.
- The volume that was inserted does not have readable bar codes.

You can either select from a list or enter a prefix.

The list contains the following selections:

DEFAULT	If DEFAULT is selected, NetBackup performs the following actions: <ul style="list-style-type: none">■ Assigns the last <code>MEDIA_ID_PREFIX</code> entry as the default prefix if <code>MEDIA_ID_PREFIX</code> entries are defined in the <code>vm.conf</code> file.■ Uses the letter A if no <code>MEDIA_ID_PREFIX</code> entries are defined.
NOT USED	If NOT USED is selected, the operation succeeds only if the robot supports bar codes and the volume has readable bar codes. NOT USED can be useful if you use bar coded volumes and want updates to fail when unreadable or missing bar codes are encountered.
Other prefixes	If <code>MEDIA_ID_PREFIX</code> entries are defined in the <code>vm.conf</code> file, they appear in the list.

To specify a prefix that is not in the list, enter the new prefix in the list box. NetBackup uses the prefix only for the current operation.

You can specify a prefix of one to five alphanumeric characters. NetBackup assigns the remaining numeric characters to create six characters. For example, if the prefix is NETB, the media IDs are: NETB00, NETB01, and so on.

Information about the `vm.conf` file is available.

See the *NetBackup Administrator's Guide for Windows, Volume II.*

Use barcode rules

Specifies whether or not to use bar code rules to assign attributes for new media.

To enable bar code rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

See “[About bar codes](#)” on page 325.

See “[Configuring bar code rules](#)” on page 329.

Information about the `vm.conf` file is available.

See the *NetBackup Administrator's Guide for Windows, Volume II.*

Media type

Specifies the type for the new media that are added to a robot. The list includes the media types that are valid for the robot.

Note: For API robots, the **Media type** is always set to DEFAULT. To specify a media type for API robots, use the **Media Type Mappings** tab of the dialog box.

See “[Configuring media type mappings](#)” on page 336.

See “[Media type when using bar code rules](#)” on page 321.

See “[Media type when not using bar code rules](#)” on page 323.

See “[About NetBackup media types](#)” on page 259.

Media type when using bar code rules

If you use bar code rules, choose one of the following:

DEFAULT NetBackup uses the bar code rules to determine the media type that is assigned.

Each media type to be added should have a bar code rule. For example, assume you want to add DLT and half-inch cartridges to a TLD robot with a single update operation. First create separate bar code rules for DLT and half-inch cartridges and then select the specific media types when you create the bar code rules. Finally, select DEFAULT on the **Media Settings** tab. The correct media type is assigned to each media.

If you choose DEFAULT on the **Media Settings** tab and DEFAULT in the bar code rule, NetBackup assigns the default media type for the robot.

A specific media type from the list. You can use a single bar code rule to add media of different types, such as DLT and half-inch cartridges (HCART) to a TLD robot. First, select a specific media type on the **Media Settings** tab. Second, select DEFAULT for the bar code rule media type when you create the bar code rule. You can perform one update for DLT and another for half-inch cartridge, and the bar code rule assigns the correct media type.

If you specify a value other than DEFAULT, the bar code rule media type must be the same as the media or be DEFAULT. If not, the bar code rule does not match the media (except for cleaning media).

Table 9-7 shows some combinations of media types on the **Media Settings** tab and bar code rule media types for a TLD (non-API) robot. It also shows the results when the media are added to the volume configuration.

Table 9-7 Example media type and bar code rule combinations

Media type on Media Settings tab	Bar code rule media type	Rule matches?	Media type added to volume configuration
DLT	DEFAULT	Yes	DLT
HCART	DEFAULT	Yes	HCART
DLT	DLT	Yes	DLT
DLT	DLT_CLN	Yes	DLT_CLN
DLT_CLN	DLT	No	DLT_CLN
DLT_CLN	DLT_CLN	Yes	DLT_CLN

Table 9-7 Example media type and bar code rule combinations (*continued*)

Media type on Media Settings tab	Bar code rule media type	Rule matches?	Media type added to volume configuration
DLT_CLN	DEFAULT	Yes	DLT_CLN
DLT	8MM, 4MM, and so on	No	DLT
DEFAULT	DEFAULT	Yes	DLT
DEFAULT	DLT	Yes	DLT
DEFAULT	DLT_CLN	Yes	DLT_CLN
DEFAULT	8 MM, 4 MM, and so on	No	Depends on robot type

The fourth row in the table shows how both cleaning cartridges and regular volumes are added using one update operation.

All the following conditions must be true:

- The media type on the **Media Settings** tab is for regular media (DLT, in this example).
- The bar code matches a bar code tag.
- The media type for the bar code rule is cleaning media (DLT_CLN).

Another example is available:

See [“Example 5: Adding cleaning tapes to a robot”](#) on page 356.

The sixth row and seventh row in the table show how to add only a cleaning tape. In the sixth row, you specify the cleaning media type on the **Media Settings** tab and in the bar code rule. In the seventh, specify the cleaning media on the **Media Settings** tab and specify default when you configure the bar code rule.

See [“Configuring bar code rules”](#) on page 329.

Media type when not using bar code rules

Choose one of the following if bar code rules are not used:

- DEFAULT** NetBackup uses the media type that is configured for the drives if:
- The drives in the robot are configured on the robot control host
 - All drives the same type
 - At least one drive is configured on the robot control host
- If the drives are not the same type, NetBackup uses the default media type for the robot.
- A specific media type** If the robot supports multiple media types and you do not want to use the default media type, select a specific type.
- The following applies only to NetBackup Enterprise Server. Select a specific media type if: the drives are not configured on the robot control host and the drives are not the default media type for the robot.

Table 9-8 shows the default media types for robots when drives are not configured on the robot control host:

Table 9-8 Default media types for non-API robots

Robot type	Default media type
Tape Library 4 MM (TL4)	4 MM cartridge tape.
Tape Library 8 MM (TL8)	8 MM cartridge tape. Also supports 8 MM cartridge tape 2 and 8 MM cartridge tape 3.
Tape Library DLT (TLD)	DLT cartridge tape. Also supports the following: <ul style="list-style-type: none"> ■ DLT cartridge tape 2 and 3, 1/2-inch cartridge tape ■ 1/2-inch cartridge tape 2, 1/2-inch cartridge tape 3 ■ 8 MM cartridge tape, 8 MM cartridge tape 2, 8 MM cartridge tape 3 ■ DTF cartridge tape ■ 1/4-inch cartridge tape

Volume pool

The volume pool for the new media. The actions depend on whether you use barcode rules to assign media attributes, as follows:

DEFAULT	<p>DEFAULT. If you select DEFAULT and:</p> <ul style="list-style-type: none"> ■ Use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned ■ Do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool ■
A specific volume pool.	If you use barcode rules, this volume pool setting always overrides the rule.

About bar codes

When a robotic library has a bar code reader, it scans the media for bar codes and saves the results. The results associate the slot number and the bar code with the media in that slot. NetBackup obtains bar code and slot information from the robotic library.

About bar code advantages

NetBackup functions well whether or not bar codes are used. However, Symantec suggests using media with bar codes in the robots that can read bar codes.

bar codes offer the following advantages:

- Automatic media ID assignment
 When you add new media to a robot, NetBackup is able to assign media IDs according to specified criteria.
- More accurate tracking of volume location
 A robot inventory update can determine which volumes are in a robot.
- Increased performance
 Not using bar codes can adversely affect performance for some robots. A robot that reads bar codes performs a scan each time it moves a tape. The robot stores the correct bar code in memory or verifies a previously saved bar code. However, if a tape does not have a bar code, the robot retries the scan multiple times, degrading performance.

About bar code best practices

Consider the following practices when you select bar codes for volumes:

- bar codes usually appear on the labels that are attached to the outside of tape volumes.

- The maximum bar code length that NetBackup supports depends on the type of robot.
See the *NetBackup Device Configuration Guide*.
- Always follow the robotic library vendor's recommendations when purchasing bar code labels for use with NetBackup.
Ensure that the bar codes have the correct number of characters.
- bar codes can represent any combination of alpha and numeric characters, but different robots support different lengths of bar codes.
See the robot vendor's documentation to determine the requirements for a specific robot type.
- Use bar codes without spaces (at the beginning, at the end, or between any characters).
Otherwise, the robot or NetBackup may not read them correctly.
- Volumes in an API robot have a real or a logical bar code.
This volume identifier is used as the NetBackup media ID. This volume identifier is the volume serial number in ACS, TLH, and TLM robots.
- For API robots, the bar code for a volume must be identical to the NetBackup media ID.
Match bar codes to media IDs by getting custom labels in the same series as the media IDs. For example, to match a set of media IDs from AA0000 to ZZ9999, get bar code labels in that series.
- When a robotic library can contain more than one media type, assign specific characters in the bar code to different media types. Do so by using media ID generation rules.
Also, use bar codes to differentiate between data tapes and cleaning tapes or to differentiate between volume pools.

About bar code rules

A bar code rule specifies criteria for assigning attributes to new robotic volumes. NetBackup assigns these attributes by using the bar code for the volume that the robotic library provides and your bar code rules.

In NetBackup, you choose whether to use bar code rules when you set up the robot inventory update operation. The bar code rules are stored on the EMM server.

Note: NetBackup does not use bar code rules if a volume already uses a bar code.

About NetBackup actions for bar codes

When a robot inventory update operation uses NetBackup bar code rules and a new bar code is detected in the robot, NetBackup does the following:

- Searches the list of rules (from first to last) for a rule that matches the new bar code.
- If the bar code matches a rule, verifies that the media type in the rule is compatible with the media type specified for the update.
- If the media types match, assigns the attributes in the rule to the volume. The attributes include the media type, volume pool, maximum number of mounts (or number of cleanings), and description.

About checking bar codes

In the robots that have bar code readers, NetBackup verifies the bar code to ensure that the robot loads the correct volume.

If the bar code on the volume does not match the bar code in the EMM database, NetBackup does one of the following:

- Assigns the request a pending status (for media-specific jobs such as a restore)
- Use another volume (for backup or duplicate jobs)

If a requested volume is not in a robot, a pending request message appears in the NetBackup Administration Console Device Monitor.

The operator must find the volume and do one of the following:

- Check the Device Monitor to find a suitable drive and mount the requested volume in that drive.
- Move the volume into the robot, update the volume configuration to reflect the correct location for the media, and resubmit the request.

If the volume is labeled, the automatic volume recognition daemon reads the label and the drive is assigned to the request. If the volume is unlabeled and not associated with a robot, the operator manually assigns the drive to the request.

Example bar code rules

[Table 9-9](#) shows some example bar code rules. Rules are sorted first according to the number of characters in the bar code tag and then by the order added. Two exceptions are the <NONE> and <DEFAULT> rules, which are always located at the end of the list.

Table 9-9 Example bar code rules

bar code tag	Media type	Volume pool	Max mounts and cleanings	Description
0080	8MM	b_pool	55	New 008 volumes
DLT	DLT	d_pool	200	DLT backup
CLD	DLT_CLN	None	30	DLT cleaning
CLT	8MM_CLN	None	20	8-mm cleaning
TL8	8MM	t_pool	0	8-mm backup
TL	8MM	None	0	8-mm no pool
<NONE>	DEFAULT	None	0	No bar code
<DEFAULT>	DEFAULT	NetBackup	0	Other bar codes

Assume that you select the following media settings (update options) for the update operation for a new 8-mm volume in a TL8 robot:

Media type = 8MM

Volume group = 00_000_TL8

Use bar code rules = YES

Volume pool = DEFAULT

If a new volume in this robotic library has a bar code of TL800001, NetBackup uses the rule with the bar code tag of TL8. NetBackup assigns the following attributes to the volume:

- Media ID = 800001 (last six characters of bar code)
- Volume group = 00_000_TL8
- Volume pool = t_pool
- Maximum mounts = 0 (no maximum)

If a new volume has a bar code of TL000001, NetBackup uses the rule with the bar code tag of TL. NetBackup assigns the following attributes to the volume:

- Media ID = 000001 (last six characters of bar code)
- Volume group = 00_000_TL8
- Volume pool = None
- Maximum mounts = 0 (no maximum)

About media ID generation rules

Use media ID generation rules to override the default media ID naming method NetBackup uses. The default method uses the last six characters of the bar code the robot provides to generate the media ID.

Note: To use media ID generation rules, the robot must support bar codes and the robot cannot be an API robot. Media ID generation rules are saved in the Media Manager configuration file (`vm.conf`).

For example, two eight-character bar codes are `S00006L1` and `000006L1`. Without any media ID generation rules NetBackup uses the last six characters of the bar code to generate media IDs. In this example, the same media ID for the two bar codes is created (`0006L1`).

Use a rule to control how NetBackup creates media IDs by specifying which characters of a bar code are used in the media ID. Or, specify that alphanumeric characters are to be inserted into the ID.

Define multiple rules to accommodate the robots and the bar code lengths. Define rules to specific robots and for each bar code format that has different numbers or characters in the bar code. Multiple rules allow flexibility for the robots that support multiple media types.

Configuring bar code rules

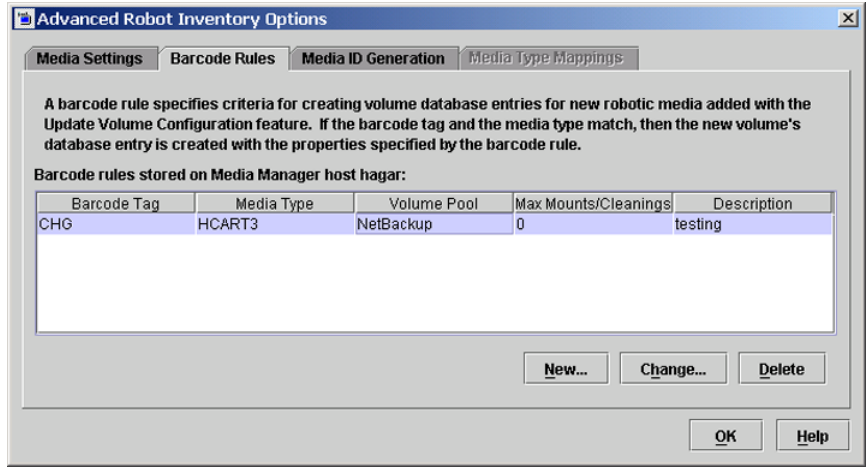
Use the **Barcode Rules** tab of the **Advanced Robot Inventory Options** dialog box to configure rules for assigning attributes to the new volumes that are added to a robot. NetBackup assigns bar codes when you select **Use barcode rules** on the **Media Settings** tab.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

See [“About bar codes”](#) on page 325.

To configure bar code rules

- 1 In the **Robot Inventory** dialog box, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Barcode Rules** tab.



- 3 To add a rule, click **New** and then configure the rule in the dialog box.
See "[Barcode rules settings](#)" on page 330.
- 4 To change a rule, select the rule, click **Change**, and then change the rule in the dialog box.

You can select and change multiple rules with one operation. The Change Barcode Rule dialog box appears for each rule that you selected for change.

You cannot change the barcode tag of a barcode rule. You first must delete the old rule and then add a rule with a new barcode tag.

See "[Barcode rules settings](#)" on page 330.
- 5 To delete a rule, select the rule, click **Delete**, and click **OK** in the confirmation dialog box. You can select and delete multiple rules with one operation.

Barcode rules settings

The following describe the settings you can configure for barcode rules. NetBackup uses these rules to assign bar codes to new media.

Barcode tag

A unique string of characters from the barcode that identifies the type of media. For example, use DLT as the barcode tag for a barcode rule if the following is true:

- You use DLT on the bar codes to identify DLT tapes
- DLT is not used on any other bar codes in the robot

Similarly, if you use CLND for DLT cleaning media, use CLND as the barcode tag for the rule for DLT cleaning media.

The barcode tag can have from 1 to 16 characters but cannot contain spaces.

The following are the special barcode rules that can match special characters in the barcode tags:

NONE	Matches when rules are used and the volume has an unreadable barcode or the robot does not support bar codes.
DEFAULT	For volumes with bar codes, this tag matches when none of the other barcode tags match. However, the following must be compatible: the media type in the DEFAULT rule and the media type on the Media Settings tab.

You cannot change the barcode tag of a barcode rule. Instead, first delete the old rule, then add a rule with a new barcode tag.

Use the **Media Settings** tab to set up the criteria for a robot update.

See [“Configuring media settings”](#) on page 317.

Description

A description of the barcode rule. Enter from 1 to 25 characters.

Maximum mounts

The maximum number of mounts (or cleanings) that are allowed for the volume.

For data volumes, a value of zero means the volume can be mounted an unlimited number of times.

For cleaning tapes, zero means that the cleaning tape is not used. Symantec recommends that you use bar codes for the cleaning media that cannot be confused with bar codes for data media. Doing so can avoid a value of 0 for cleaning tapes.

Media type option

The media type to assign to the media.

The media type that is specified on the **Media Settings** tab always overrides the media type of the barcode rule. If you specify a value other than `DEFAULT` on the **Media Settings** tab, the barcode rule media type must be the same as the media or be `DEFAULT`. If not, the barcode rule does not match the media (except for cleaning media).

See [“Media type when using bar code rules”](#) on page 321.

Note: When a media type is selected, the maximum mounts value may revert to the default value for the specified media type. For example, it may revert to 0 for unlimited when you select a non-cleaning media type.

See [“About NetBackup media types”](#) on page 259.

Barcode rule media type for API robots

NetBackup uses a barcode rule only if the barcode rule media type is compatible with the media type on the **Media Type Mappings** tab.

If you specify `DEFAULT`, NetBackup uses the media type you select on the **Media Type Mappings** tab. If you do not specify a media type on the **Media Type Mappings** tab, NetBackup uses the default media type for the robot.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

Note: You can write a barcode rule that contains the media types that are incompatible with vendor media types. However, the robot inventory update may assign NetBackup media types that are inconsistent with the vendor media types. Avoid this problem by grouping barcode rules by media type.

See [“Configuring media type mappings”](#) on page 336.

Barcode rule media type for non-API robots

NetBackup uses a barcode rule only if the media type in the rule is compatible with the media type you select on the **Media Settings** tab.

[Table 9-10](#) shows the media type to select for non-API robots.

Table 9-10 Media type selection

If you want the media type for the barcode rule to match	Select the following media type for the barcode rule	Media type NetBackup uses
Any media type that you select on the Media Settings tab	DEFAULT	The media type that you select on the Media Settings tab. If you also select DEFAULT on the Media Settings tab, the Media Manager default media type for the robot is used.
Only when you select a specific media type or you select DEFAULT on the Media Settings tab	The same media type	The media type that you select for the barcode rule.

Volume pool

The volume pool for the new media. The actions depend on whether you use barcode rules to assign media attributes.

Select from the following:

- DEFAULT If DEFAULT is selected, NetBackup performs the following actions:
- If you use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned.
 - If you do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool.

A specific volume pool This volume pool setting always overrides any barcode rules.

Configuring media ID generation rules

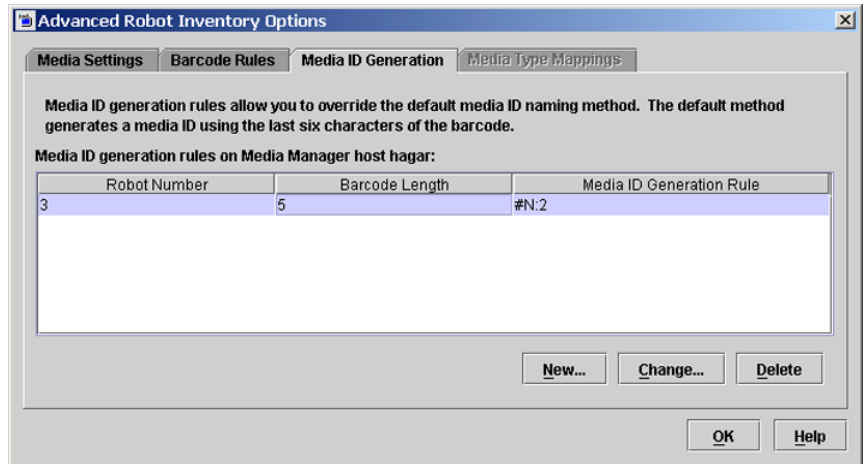
For non-API robots only.

To use media ID generation rules, the robot must support bar codes and the robot cannot be an API robot.

See [“About media ID generation rules”](#) on page 329.

To configure media ID generation rules

- 1 In the **Robot Inventory** dialog box, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media ID Generation** tab.



- 3 To add a rule, click **New** and then configure the rule in the dialog box.
See "[Media ID generation options](#)" on page 334.
- 4 To change a rule, select the rule, click **Change**, and then change the rule in the dialog box.
You cannot change the robot number or barcode length of a rule. To change those properties, first delete the old rule and then add a rule.
You can select and change multiple rules with one operation. A separate change rule dialog box appears for each rule that you selected for change.
- 5 To delete a rule, select the rule, click **Delete**, and click **OK** in the confirmation dialog box. You can select and delete multiple rules with one operation.

Media ID generation options

NetBackup uses rules to generate the IDs for media in robots. The default rule uses the last six characters of the bar code label from the tape.

You can configure media ID generation rules to override the default rule. Control how NetBackup creates media IDs by defining the rules that specify which characters of a bar code label to use for the media ID.

The following subsections describe the media ID generation rule options.

Barcode length

The **Barcode length** is the number of characters in the bar code for tapes in the robot.

You cannot change the barcode length of a rule. Rather, first delete the rule and then add a new rule.

Media ID generation rule

A **Media ID generation rule** consists of a maximum of six colon-separated fields. Numbers define the positions of the characters in the bar code that are to be extracted. For example, the number 2 in a field extracts the second character (from the left) of the bar code. You can specify numbers in any order.

To insert a specific character in a generated media ID, precede the character by a pound sign (#). Any alphanumeric characters that are specified must be valid for a media ID.

Use rules to create media IDs of many formats. However, it may be difficult to manage media if the label on the media and the generated media ID are different.

[Table 9-11](#) shows some examples of rules and the resulting media IDs.

Table 9-11 Example rules and resulting media ID

Bar code on tape	Media ID generation rule	Generated media ID
032945L1	1:2:3:4:5:6	032945
032945L1	3:4:5:6:7	2945L
032945L1	#N:2:3:4:5:6	N32945
543106L1	#9:2:3:4	9431
543106L1	1:2:3:4:#P	5431P

Robot number

The number of the robot to which the rule applies.

You cannot change the robot number of a rule. Rather, first delete the rule and then add a new rule.

Configuring media type mappings

Applies to API robots only.

For API robots, NetBackup contains default mappings from a vendor's media types to NetBackup media types. API robots are ACS, TLH, or TLM robot types.

You can change the default mappings. Changes apply only to the current volume configuration update.

You also can add media type mappings.

See [“About adding media type mapping entries”](#) on page 337.

See [“About the default and allowable media types”](#) on page 338.

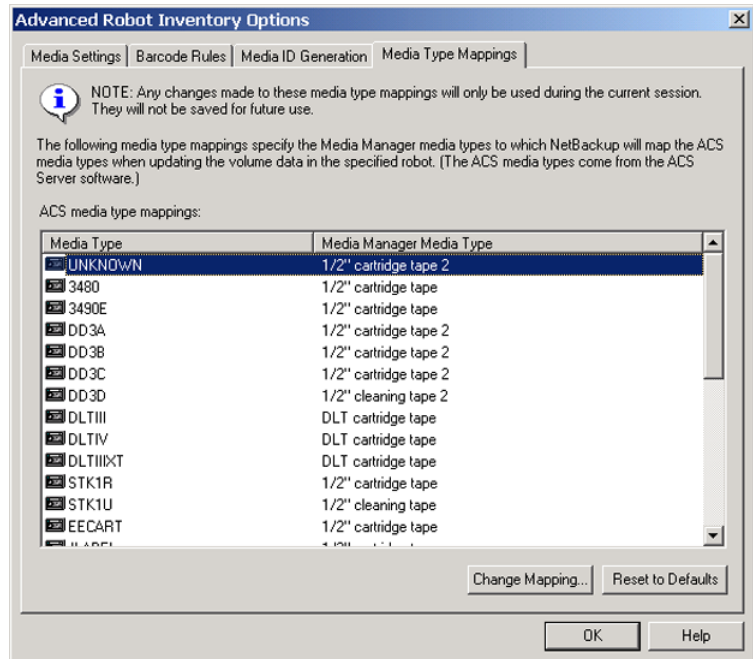
See [“About NetBackup media types”](#) on page 259.

Note: You can write a barcode rule that contains the media types that are incompatible with vendor media types. However, the robot inventory update may assign NetBackup media types that are inconsistent with the vendor media types. Avoid this problem by grouping barcode rules by media type.

Use the following procedure to change media type mappings.

To change media type mappings

- 1 In the **Robot Inventory** dialog box, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media Type Mappings** tab.



The mappings that appear are only for the robot type that was selected for inventory.

The tab shows the default mappings and any mappings you add.

- 3 Select the row that contains the robot-vendor media type mapping that you want to change and click **Change Mapping**.
- 4 In the **Change Media Mapping** dialog box, select a Media Manager media type from the list of allowed selections.
- 5 Click **OK**.

To reset the mappings to the default, click **Reset to Defaults**.

About adding media type mapping entries

Applies to API robots only.

The default media type mappings may not provide the wanted mappings. If not, add robot-specific media mappings to the `vm.conf` file on the host on which you are run the NetBackup Administration Console.

Information about how to do so is available.

See the *NetBackup Administrator's Guide for Windows, Volume II*.

Table 9-12 Examples of robot-specific media mappings

vm.conf entry	Result	Robot default without a vm.conf entry
ACS_3490E = HCART2	Maps the ACS 3490E to the HCART2 media type.	HCART
ACS_DLTIV = DLT2	Maps ACS DLTIV to the DLT2 media type.	DLT for all ACS DLT media types, including DLTIV
TLH_3490E = HCART2	Maps the TLH 3490E to the HCART2 media type.	HCART

About the default and allowable media types

Applies to API robots only.

The default media types on the **Media Type Mappings** tab are the media types provided by each robot vendor.

The following tables contain the default and allowable media types for the API robots as follows:

- The second column of each table shows the default media type.
- The third column shows the media types to which you can map the defaults. To do so, first add the allowable mapping entries to the `vm.conf` file. Some map entries are not allowed. For example, you cannot specify either of the following map entries for ACS robots:

```
ACS_DD3A = DLT
ACS_DD3A = HCART4
```

[Table 9-13](#) shows the default media types and the allowable media types for ACS robots.

[Table 9-14](#) shows the default and allowable media types for TLH robots.

[Table 9-15](#) shows the default and allowable media types for TLM robots.

Table 9-13 Default and allowable media types for ACS robots

ACS media type	Default media type	Allowable media types through mappings
3480	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3490E	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
DD3A	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3B	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3C	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
DD3D	1/2-inch cartridge cleaning tape 2 (HC2_CLN)	HC_CLN, HC2_CLN, HC3_CLN
DLTIII	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
DLTIIIEXT	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
DLTIV	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
EECART	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
JLABEL	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
KLABEL	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_100G	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_10GB	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_200G	1/2-inch cartridge (HCART2)	HCART, HCART2, HCART3
LTO_35GB	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_400G	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_400W	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
LTO_50GB	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
LTO_800G	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3

Table 9-13 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
LTO_800W	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_CLN1	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLN2	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLN3	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
LTO_CLNU	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
SDLT	Digital Linear Tape 3 (DLT3)	DLT, DLT2, DLT3
SDLT_2	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_4	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S1	Digital Linear Tape 2 (DLT2)	DLT, DLT2, DLT3
SDLT_S2	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S3	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
SDLT_S4	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
STK1R	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
STK1U	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
STK1Y	1/2-inch cartridge cleaning tape (HC_CLN)	HC_CLN, HC2_CLN, HC3_CLN
STK2P	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
STK2W	1/2-inch cartridge cleaning tape 2 (HC2_CLN)	HC_CLN, HC2_CLN, HC3_CLN
T10000CT	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3

Table 9-13 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
T1000T1	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T1000TS	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
UNKNOWN (for unknown ACS media types)	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3, HC_CLN, HC2_CLN, HC3_CLN, DLT, DLT2, DLT3, DLT_CLN, DLT2_CLN, DLT3_CLN
VIRTUAL	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3, HC_CLN, HC2_CLN, HC3_CLN, DLT, DLT2, DLT3, DLT_CLN, DLT2_CLN, DLT3_CLN

Table 9-14 Default and allowable media types for TLH robots

TLH media type	Default Media Manager media type	Allowable media types through mappings
3480	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3490E	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3590J	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3590K	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3592JA	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JB	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JX	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JJ	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JR	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3

Table 9-14 Default and allowable media types for TLH robots (*continued*)

TLH media type	Default Media Manager media type	Allowable media types through mappings
3592JW	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
UNKNOWN (for unknown TLH media types)	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3

Table 9-15 Default and allowable media types for TLM robots

TLM media type	Default media type	Allowable media types through mappings
3480	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
OD_THICK	NONE (OD_THICK is translated to media type REWR_OPT for robot contents reports. OD_THICK is ignored for all other robotic inventory operations)	NONE
DECDLT	Digital Linear Tape (DLT)	DLT, DLT2, DLT3
8MM	8mm cartridge (8MM)	8MM, 8MM2, 8MM3
4MM	4mm cartridge (4MM)	4MM
3590	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
DTF	DTF cartridge (DTF)	DTF
SONY_AIT	8mm cartridge (8MM)	8MM, 8MM2, 8MM3
LTO	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
UNKNOWN (for unknown TLM media types)	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3, DLT, DLT2, DLT3, 8MM, 8MM2, 8MM3

Note: The following TLM media types are not supported: OD_THIN, D2, VHS, CD, TRAVAN, BETACAM, AUDIO_TAPE, BETACAMCL, DVCM, and DVCL.

About the physical inventory utility

For the robots without bar code readers or that contain media without bar codes, the **NetBackup Administration Console** inventory reports only the presence of media in a robotic library. More detailed information is required to perform automated media management. Use the `vmphyinv` physical inventory utility for such robots.

The `vmphyinv` physical inventory utility inventories nonbarcoded tape libraries by performing the following actions:

- Mounts each tape
- Reads the tape header
- Identifies the tape in each slot
- Updates the NetBackup volume configuration

`vmphyinv` is a command-line utility. Additional information about the syntax of the `vmphyinv` command is available.

See the *NetBackup Commands* guide.

About the `vmphyinv` features

The `vmphyinv` utility has the following features:

- Can be run from any master server, media server, or SAN media server.
- Can be used with bar coded tape libraries because it verifies the contents of the media.
- Recognizes the NetBackup and the Backup Exec tape formats.
- Supports the remote administration. You do not need to run `vmphyinv` from the host to which the drives are attached.
- Tries to use multiple drives in a robot even if the drives are attached to different hosts.
- Works with shared drives (NetBackup Shared Storage Option).
- Supports all supported SCSI-based robot types.
- Can be used to inventory a single media in a standalone drive. Use the `-u` option or the `-n` option to specify the drive; the drive must contain media and it must be ready.

About `vmphyinv` requirements and restrictions

The `vmphyinv` utility has the following requirements and restrictions:

- It cannot distinguish between the volume records based on the application type.
- When you move the media from robotic drives to standalone drives, you cannot specify a new volume group for the media.

When to use `vmphyinv`

Use the `vmphyinv` utility to update the EMM database for NetBackup and Backup Exec media.

Use `vmphyinv` in the following typical cases:

- You want to inventory a robot that does not have a bar code reader or that contains nonbarcoded media.
- You insert new media into a robotic library and no NetBackup volume records correspond to the media. Use the slot range or list option of `vmphyinv` to perform the inventory operation. You do not need to add volume records to the EMM database.

- You insert some media that have unknown media IDs or globally unique identifiers (GUIDs) into a robot.

For example, you insert 10 media from a different tape library in slots 11 to 20. You do not know the IDs on the tapes. Use the slot range or list option of `vmphyinv` to perform the inventory operation. The `vmphyinv` utility mounts the media, reads the tape header, determines the media ID, and adds media records to the EMM database.

- Some of the media are misplaced and the EMM database does not reflect the correct physical location of these media. Inventory the robot or inventory a subset of media in the robot by using options in `vmphyinv`.

How `vmphyinv` performs a physical inventory

For a physical inventory, the `vmphyinv` utility performs the following sequence of operations:

- Obtains a list of drives to mount the media
See [“About the list of drives”](#) on page 345.
- Obtains a list of media to mount
See [“About the media to be mounted”](#) on page 345.
- Mounts the media and reads the tape headers
See [“About mounts the media and reads the tape header”](#) on page 346.
- Updates the EMM database

See [“About updates the EMM database”](#) on page 347.

About the list of drives

The list of drives the `vmphyinv` utility uses to mount the media is obtained from the EMM database. The drives do not need to be configured locally.

You cannot specify which drives to use. However, you can specify the maximum number of drives to use, which lets you reserve drives for NetBackup backup or restore operations. Specify the number of drives by using the `-drv_cnt` *drive_count* option.

About the media to be mounted

The `vmphyinv` command accepts several options for the media to be mounted, as follows:

- **NetBackup robot number** (`-rn robot_number`).
 The `vmphyinv` utility obtains a list of volume records for that robot and inventories each of the media in the list.
 To use this option, the NetBackup configuration must contain a volume record that corresponds to the robot number in the EMM database for the robot.
- **NetBackup robot number with filter options.**
 If you do not want to inventory all of the media in a robot, specify a subset of the media by using filter options. Some filter options are volume pool, volume group, or slot range. To use these options, NetBackup volume records must exist.

The following are some filter examples.

<code>vmphyinv -rn 4 -pn bear</code>	Mounts media only in robot 4 and in the volume pool bear.
<code>vmphyinv -rn 2 -v moon</code>	Mounts media in robot 2 and in the volume group moon.
<code>vmphyinv -rn 1 -rc1 2 -number 3</code>	Mounts media in robot 1 and slot range 2 to 4.
<code>vmphyinv -rn 5 -pn NetBackup -v mars -rc1 2 -number 6</code>	Mounts media in robot 5, slot range 2 to 7, in volume group mars, and in the NetBackup volume pool.

- **NetBackup robot number and a list of media that belong to a specific robot.**
 For example, if the `-rn robot_number` and `-ml A00001:A00002:A00003` options are specified, only the three specified media are inventoried. If any of these

media do not belong to the specified robot, the media are skipped and are not inventoried. To use this option, NetBackup volume records must exist.

- NetBackup robot number and a slot range or list.
Sometimes, media from a different robot or some other source are moved to a robot and the media ID on the tape is unknown. In these cases, specify a slot range option or list option.
With these options, the NetBackup volume record does not need to exist in the EMM database. However, you must specify the density (using the `-d` option).

Note: For a robot that supports multiple media types, specify the density carefully. If you specify the incorrect density, `vmphyinv` cannot complete the mount and permanent drive failure can occur.

The following are some filter examples.

```
vmphyinv -rn 1 -slot_range 2 10 Mounts media in slot range 2 to 10 in robot  
-d dlt 1.
```

```
vmphyinv -rn 0 -slot_list 3:4:5 Mounts media in slots 3, 4, and 5 in robot  
-d 8mm 0.
```

```
vmphyinv -rn 2 -slot_range 2 4 Mounts media in slots 2, 3, 4, 5, 6, and 7  
-slot_list 5:6:7 -d dlt in robot 2.
```

About mounts the media and reads the tape header

The following sequence of operations explains the mount process:

- The `vmphyinv` utility contacts the NetBackup Volume Manager, `vmd`, on the local host or remote host depending on where the drive is attached.
- The NetBackup Volume Manager starts a process, `opr`.
- The `vmphyinv` utility communicates with `opr` and sends the mount request to `opr`. After `opr` receives the request, it issues a mount request to `ltid`.
- The `vmphyinv` utility reads the tape header to determine the recorded media ID or globally unique identifier (GUID). GUID is an identifier used by Symantec Backup Exec.

Note: The default mount timeout is 15 minutes. Specify a different mount time by using the `-mount_timeout` option.

About media that is not recognized

If the media is not NetBackup media or Backup Exec media, the media is unmounted and the next media is mounted. `vmphyinv` does not generate a new record in the EMM database. To generate volume records for that media, use the `vmupdate` command.

About cleaning media

If the following conditions are all true, `vmphyinv` does not try to mount the media and the next media in the list is mounted:

- You do not specify the `vmphyinv` slot range or list option.
- The robot contains cleaning media.
- The media type is specified as cleaning media in the volume record (such as `4mm_clean` or `dlt_clean`).

If the robot contains cleaning media and any of the following conditions are true, `vmphyinv` tries to determine if the media is cleaning media:

- You use the slot range or list option and the media type of volume record in the EMM database is not a cleaning media type.
- You use the slot range or list option, and the EMM database does not contain a volume record that corresponds to the cleaning media.
- You do not use the slot range or list option, and the EMM database does not contain a volume record that corresponds to the cleaning media.

The `vmphyinv` utility tries to determine if the media is cleaning media. It uses the SCSI parameters (sense keys, tape alert flags, and physical (SCSI) media types) returned by the robot. If `vmphyinv` cannot determine if the media is cleaning media, it tries to mount the media until the mount request times out.

Note: NetBackup may not detect the presence of cleaning media for all drives. Some drives report the presence of cleaning media in a manner NetBackup cannot read.

About updates the EMM database

After all the media are mounted and the tape headers are read, `vmphyinv` displays a list of recommended changes. Accept or reject the changes. If you accept the changes, `vmphyinv` updates the EMM database.

About using the verbose option

Use the `vmphyinv -verbose` option to display more information about the suggested changes. The `-verbose` option shows the number of drives available, the contents of each tape, if the media is a catalog tape. (The media format column of the summary contains NetBackup database for NetBackup catalog tapes.)

This verbose information is written to `stderr`. To save the information, redirect `stderr` to a file.

About update criteria

For valid media types, `vmphyinv` performs the following actions:

- Changes the residence fields and description fields of any NetBackup media record if those fields do not match the media header. The description field is changed only if the media is Symantec Backup Exec media.
- Conditionally changes the media type of an unassigned NetBackup volume record. The media type is changed only if the new media type belongs to the same family of media types as the old media type. For example, the media type DLT can only be changed to DLT2 or DLT3.
- Never changes the volume pool, media type, and ADAMM_GUID of an assigned record.
- Never unassigns an assigned NetBackup volume.

About updating NetBackup media

The `vmphyinv` utility searches the EMM database. It checks if the media ID from the tape is present in the media ID field of any record in the EMM database. If the media ID exists, `vmphyinv` updates the NetBackup volume record that corresponds to the media ID. If the media ID does not exist, `vmphyinv` creates a new NetBackup volume record that corresponds to the NetBackup media.

About updating Backup Exec media

The `vmphyinv` utility searches the EMM database. It checks if the media GUID from the tape is present in the ADAMM_GUID field of any record in the EMM database. If the media GUID exists, `vmphyinv` updates the NetBackup record that contains the GUID. If a media GUID does not exist, `vmphyinv` creates a new NetBackup record that corresponds to the Backup Exec media. `vmphyinv` may use an existing NetBackup volume record if the record does not correspond to any media in the tape library.

For each NetBackup volume record, `vmphyinv` does the following:

- In the NetBackup record, updates the ADAMM_GUID field with the GUID and the Description field with the Backup Exec cartridge label in the tape header.
- Adds the media ID of the NetBackup record to the EMM database (if not already present). Each record is assigned to NetBackup (if not already assigned) and its state is set to Frozen in the EMM database.
- Changes the volume pool of the unassigned NetBackup volume records that are associated with Backup Exec media to the Backup Exec pool. If the Backup Exec pool does not exist, `vmphyinv` creates it.

Note: If a `MEDIA_ID_PREFIX` entry is not specified in the `vm.conf` file, NetBackup uses BE as the default prefix for Backup Exec media.

About error cases

The `vmphyinv` utility may not be able to update the EMM database correctly in the following cases. These cases are reported as errors.

If any of the following cases are encountered, you must intervene to continue:

- Duplicate media IDs are found.
Two or more media in the same robot have the same media ID.
- A NetBackup volume record that belongs to a different robot is found.
It contains the same media ID as the media ID read from the tape header.
- The media type, media GUID, or volume pool of an assigned volume record needs to be changed.
- The bar code of an existing volume record needs to be changed.

Example volume configuration updates

The following examples show only the relevant volume attributes.

See the following for examples of different types of volume configuration updates:

- [Example 1: Removing a volume from a robot](#)
- [Example 2: Adding existing stand-alone volumes to a robot](#)
- [Example 3: Moving existing volumes within a robot](#)
- [Example 4: Adding new volumes to a robot](#)
- [Example 5: Adding cleaning tapes to a robot](#)
- [Example 6: Moving existing volumes between robots](#)

- [Example 7: Adding existing volumes when bar codes are not used](#)

Example 1: Removing a volume from a robot

The following is an example of how to remove a volume from a robotic library. It does not matter whether the robot supports bar codes.

The following are the attributes for media ID 800001:

media ID	800001
media type	8MM cartridge tape
bar code	TL800001
media description	t18 backup volume
volume pool	NetBackup
robot type	TL8 - Tape Library 8MM
volume group	EXB220
max mounts allowed	0 (unlimited)

Assume that you remove the volume from the robotic library, specify the following on the **Media Settings** tab, then run the update:

media type	DEFAULT
volume group	NONROB_8MM
volume pool	DEFAULT

The resulting volume attributes for media ID 800001 are as follows:

media ID	800001
media type	8MM cartridge tape
bar code	TL800001
media description	t18 backup volume
volume pool	NetBackup
robot type	NONE - Not Robotic

```

volume group      NONROB_8MM
max mounts       0 (unlimited)
allowed

```

The new residence information in the EMM database shows a stand-alone location in the volume group. The volume group is specified on the **Media Settings** tab. The media type and volume pool remain unchanged.

The results are the same for a volume that does not have a bar code.

Example 2: Adding existing stand-alone volumes to a robot

The following is an example of how to add a stand-alone volume that has a bar code to a robotic library that supports bar codes (TL8).

When you move volumes from one robot to another robot, perform separate updates.

See [“Example 6: Moving existing volumes between robots”](#) on page 357.

The following are the volume attributes for media ID 800021, which has a readable bar code and already exists as a stand-alone volume:

```

media ID          800021
media type        8MM cartridge tape
bar code          TL800021
media description 8MM stand-alone
volume pool       None
robot type        None (stand-alone)
volume group      NONROB_8MM
max mounts       0 (unlimited)
allowed

```

Assume that you insert the volume into a TL8 robot, specify the following on the **Media Settings** tab, then run the update:

```

media type        DEFAULT
volume group      EXB220
use bar code rules YES (selected)

```

volume pool NetBackup

Assume that the bar code rules in [Table 9-16](#) exist:

Table 9-16 Example bar code rules

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
CLND	DLT_CLN	None	30	dlt cleaning
CLN8	8MM_CLN	None	20	8mm cleaning
TL8	8MM	NetBackup	0	tl8 backup
DLT	DLT	d_pool	200	dlt backup
TS	8MM	None	0	8mm no pool
<NONE>	DEFAULT	None	0	no bar code
<DEFAULT>	DEFAULT	NetBackup	0	other bar codes

NetBackup recognizes that the media ID exists and changes the EMM database to reflect the new robotic location. NetBackup does not create a new media ID.

The volume attributes for media ID 800021 are as follows:

media ID 800021
media type 8MM cartridge tape
bar code TL800021
media description 8MM stand-alone
volume pool NONE
robot type TL8 - Tape Library 8MM
robot number 0
robot slot 1
robot host shark
volume group EXB220
max mounts
allowed 0 (unlimited)

The bar code matches the bar code of an existing stand-alone volume in the configuration. Therefore, NetBackup updates the residence information in the EMM database to reflect the new robotic location. Because the volume is not new, bar code rules are ignored.

The only setting used on the **Media Settings** tab is the volume group for added or moved volumes. The media type setting was not used because this example was for a single existing volume that already had a media type.

Example 3: Moving existing volumes within a robot

The following is an example of how to move a volume from one slot to another slot within the same robot. The robot supports bar codes and the volume has a readable bar code.

Note: To move volumes within a robotic library, use **Update volume configuration** only if the robotic library supports bar codes and the volumes have readable bar codes. Otherwise, NetBackup cannot properly recognize the move.

The following are the attributes for media ID 800002, which currently resides in slot 1 of the robotic library:

media ID	800002
media type	8MM cartridge tape
bar code	TL800002
media description	tl8 backup
volume pool	NetBackup
robot type	TL8 - Tape Library 8MM
robot number	0
robot slot	1
robot host	shark
volume group	EXB220
max mounts allowed	0 (unlimited)

Assume that you move the volume to empty slot 10, specify the following on the **Media Settings** tab, then run the update.

Example volume configuration updates

media type	DEFAULT
volume group	EXB220
use bar code rules	NO (not selected)
volume pool	DEFAULT

The resulting volume attributes are the following:

media ID	800002
media type	8MM cartridge tape
bar code	TL800002
media description	tl8 backup
volume pool	NetBackup
robot type	TL8 - Tape Library 8MM
robot number	0
robot slot	10
robot host	shark
volume group	EXB220
max mounts allowed	0 (unlimited)

The updated volume attributes show the new slot number, but all other information is unchanged.

Example 4: Adding new volumes to a robot

The following is an example of how to add new volumes with bar codes to a robot that supports bar codes.

Assume the following:

- The new volume is an 8MM tape with a readable bar code of TL800002.
- No media generation rules are defined.
- The drives in the robot all have a drive type of 8MM or no drives are configured on the robot control host.

Specify the following on the **Media Settings** tab and run the update:

Media type DEFAULT
 Volume group EXB2220
 Use bar code rules YES (selected)
 Volume pool DEFAULT

[Table 9-17](#) contains the example bar code rules.

Table 9-17 Example bar code rules

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
CLND	DLT_CLN	None	30	dlt cleaning
CLN8	8MM_CLN	None	20	8mm cleaning
TL8	8MM	NetBackup	0	tl8 backup
DLT	DLT	d_pool	200	dlt backup
TS	8MM	None	0	8mm no pool
<NONE>	DEFAULT	None	0	no bar code

The bar code on the media matches the bar code rule named TL8 and the resulting volume attributes for the new volume are as follows:

Media ID 800002
 Media type 8MM cartridge tape
 Bar code TL800002
 Media description tl8 backup
 Volume pool NetBackup
 Robot type TL8 - Tape Library 8MM
 Robot number 0
 Robot slot 1
 Robot host shark
 Volume group EXB2220

Maximum mounts 0 (unlimited)
 allowed

No media ID generation rules exist. Therefore, the media ID is from the last six characters of the bar code. The new residence information in the EMM database shows the robot host, robot type, robot number, slot, and host. The volume group is from the **Media Settings** tab. The volume pool and max mounts allowed are from the bar code rule.

If bar code rules (or bar codes) are not used, the media description, volume pool, and max mounts allowed would be set to the following defaults:

Media description Added by NetBackup
 Volume pool NetBackup for data tapes or None for cleaning tapes
 Max mounts 0 (unlimited)

Note: If the robot does not support bar codes or the bar code is unreadable, specify a Media ID prefix on the **Media Settings** tab. Alternatively, specify DEFAULT for the media ID. If you do not, NetBackup does not add new media IDs.

Example 5: Adding cleaning tapes to a robot

A special case exists when you add cleaning tapes. For example, assume you update a TLD robot.

The tapes you inserted include regular tapes with bar codes that range from DLT00000 to DLT00010 and a cleaning tape with a bar code of CLN001.

[Table 9-18](#) contains the example bar code rules:

Table 9-18 Example bar code rules

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
CLN	DLT_CLN	None	30	dlt cleaning
DL	DLT	d_pool	200	dlt backup
<NONE>	DEFAULT	None	0	no bar code

Specify the following on the **Media Settings** tab, then run the update.

media type DLT

volume group STK7430
use bar code rules YES (selected)

The bar codes on the regular tapes match the DL bar code rule. The media type of the DL bar code rule matches the Media type on the **Media Settings** tab. The tapes are added as DLT.

The cleaning tape matches the CLN bar code rule. NetBackup recognizes that DLT_CLN is the cleaning tape for DLT. NetBackup adds the cleaning tape CLN001 as DLT_CLN type media along with the regular volumes.

This example shows NetBackup's ability to add cleaning cartridges along with regular volumes when you use Update volume configuration.

If the volumes you insert include a cleaning tape, NetBackup adds the volumes correctly if the following are true:

- The Media type on the **Media Settings** tab is the regular media (DLT in this example).
- The bar code on the volume matches a bar code tag (CLN in this example).
- The media type for the bar code rule is the correct cleaning media (DLT_CLN in this example).

To add only cleaning media, specify the cleaning media type on the **Media Settings** tab and in the bar code rule (DLT_CLN in this example).

Example 6: Moving existing volumes between robots

When you move volumes from one robot to another and the volumes in both robots are in the same EMM database, perform two separate updates.

These updates move the volumes to stand-alone, as an intermediate step, and then to the new robot. Otherwise, NetBackup is unable to update the entries and you receive an "Update request failed" error.

This example assumes that robot 2 is able to read bar codes and the volume has readable bar codes. If not, NetBackup cannot manage the volumes..

See [“Example 7: Adding existing volumes when bar codes are not used”](#) on page 358.

To move existing volumes between robots, use the following process:

- Remove the volume from robot 1 and insert the volume in robot 2.
- Perform an Update volume configuration on robot 1. This updates the volume attributes to show the volume as stand-alone.

- Perform an Update volume configuration on robot 2. This updates the configuration to show the volume in robot 2.

Example 7: Adding existing volumes when bar codes are not used

This example is not recommended and is included only to illustrate the undesirable results.

The following is an example of how to add an existing stand-alone volume to a TL4 robot. A TL4 robot supports media inventory (detects media presence), but not bar codes.

The following are the attributes for media ID 400021, which already exists as a stand-alone volume:

media ID	400021
media type	4MM cartridge tape
bar code	-----
media description	4MM stand-alone
volume pool	None
robot type	NONE - Not Robotic
volume group	NONROB_4MM
max mounts allowed	0 (unlimited)

Assume that you insert the volume into the robot, specify the following on the **Media Settings** tab, and run the update:

media type	DEFAULT
volume group	00_000_TL4
media ID prefix	C4
volume pool	DEFAULT

The resulting volume attributes are as follows:

media ID	C40000
media type	4MM cartridge tape

```

bar code          -----
media description Added by NetBackup
volume pool       NetBackup
robot type        TL4 - Tape Library 4MM
robot number      0
robot slot        1
robot host        shark
volume group      00_000_TL4
max mounts        0 (unlimited)
allowed
  
```

Note that NetBackup assigned a new media ID to the volume (C40000). This undesired result occurs if you use **Update volume configuration** and the volumes do not contain readable bar codes or the robot does not support bar codes. Without a bar code, NetBackup cannot identify the volume and assumes it is new. The media ID C40000 is generated from the media ID prefix specified on the **Media Settings** tab.

The old media ID (400021) remains in the configuration. The information for the new media ID (C40000) shows the robotic location, which includes the robot host, robot type, number, slot, and host. The volume group and volume pool are configured according to the **Media Settings** tab selections. The maximum mounts allowed is set to the default (0).

For this situation, use the physical inventory utility.

See [“About the physical inventory utility”](#) on page 343.

Configuring disk storage

This chapter includes the following topics:

- [Configuring BasicDisk storage](#)
- [Configuring NearStore storage](#)
- [About SharedDisk support in NetBackup 7.0 and later](#)
- [Configuring disk pool storage](#)

Configuring BasicDisk storage

A BasicDisk type storage unit consists of a directory on locally-attached disk or network-attached disk that is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory.

No special configuration is required for BasicDisk storage. You specify the directory when you configure a storage unit.

See [“About storage units”](#) on page 366.

Configuring NearStore storage

A NearStore disk type storage unit is used to store images on Network Attached Storage (NAS) from NetApp. The NearStore disk storage unit features are available on all supported media server platforms.

Information about configuring NearStore storage units is described in the *NetBackup Administrator's Guide, Volume II*.

About SharedDisk support in NetBackup 7.0 and later

The SharedDisk option is not supported beginning with the NetBackup 7.0 release.

You can use a NetBackup 7.0 master server to configure, manage, and operate SharedDisk on NetBackup 6.5 media servers.

For information about using SharedDisk, see the documentation for your NetBackup 6.5 release.

With these changes, the following behavior is to be expected in NetBackup 7.0:

- All configuration attempts to a SharedDisk storage server on a 7.0 or later media server fail with a `storage server not found error`.
- All read or write requests to a SharedDisk disk pool use 6.5 media servers only. If no 6.5 media servers are available, the requests fail.
- If you upgrade a 6.5 SharedDisk media server to 7.0, NetBackup marks the storage servers as DOWN. It no longer functions as a SharedDisk storage server.

To ensure that the media server is not considered for SharedDisk jobs, do one of the following: Restart the Enterprise Media Manager service after the upgrade or remove the storage server from all disk pools and then delete it.

- You can delete the SharedDisk disk pools and the SharedDisk storage servers that reside on 7.0 media servers. However, all delete operations on images fail. To delete images, do the following:

- Expire the images and delete them from the catalog by using one of the following `bpexpdate` commands:

```
bpexpdate -backupid backupid -d 0 -nodelete
```

With this command, NetBackup does not run an image cleanup job. You can use **NetBackup Management > Catalog** to determine the *backupid*.

```
bpexpdate -backupid backupid -d 0 -force
```

With this command, NetBackup attempts an image cleanup job. It fails with error 174; you can ignore the error. You can use **NetBackup Management > Catalog** to determine the *backupid*.

```
bpexpdate -stype SharedDisk
```

With this command, NetBackup attempts an image cleanup job. It fails with error 174; you can ignore the error.

- Delete the fragments of the expired images by using the following command:
- ```
nbdelete -allvolumes -force
```

---

**Note:** Symantec recommends that you use solutions other than SharedDisk. The AdvancedDisk storage option is another solution.

---

## Configuring disk pool storage

You can configure disk pools if you license a NetBackup feature that uses disk pools.

For more information, see the NetBackup online Help or the following guides:

- *The NetBackup Deduplication Guide.*
- *The NetBackup Shared Storage Guide.*



# Configuring storage units

This chapter includes the following topics:

- [About the Storage utility](#)
- [About storage units](#)
- [About storage unit settings](#)

## About the Storage utility

The data that is generated from a backup job or another type of job is recorded in storage. A storage destination can be a single tape or disk volume, a named group of storage units, or a storage lifecycle policy.

A NetBackup administrator must define storage destinations with the **Storage** utility before a backup job or another type of job can be run.

The **Storage** utility contains subnodes to define three different storage configurations:

- **Storage Units**

The primary storage destination is a storage unit, since storage units can be included as part of a storage unit group or a storage lifecycle policy.

A storage unit is a label that NetBackup associates with physical storage. The label can identify a robot, a path to a volume, or a disk pool.

See [“About storage units”](#) on page 366.

- **Storage Unit Groups**

Storage unit groups let you identify multiple storage units as a group. How the storage units are selected within the group is determined when the group is created.

See [“About Storage unit groups”](#) on page 407.

- **Storage Lifecycle Policies**

Storage lifecycle policies let you apply the same behavior to all the backup images in the lifecycle.

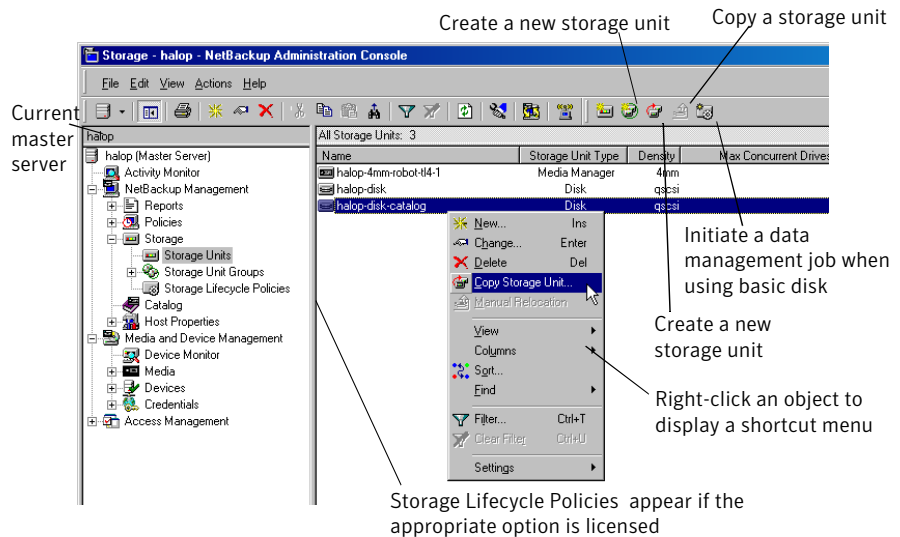
See “Storage lifecycle policy overview” on page 415.

## Using the Storage utility

Expand **Storage**, and then **Storage Units**, **Storage Unit Groups**, or **Storage Lifecycle Policies** to display the storage destinations that were created for the selected server. The storage configuration can be displayed for other master servers.

See “Accessing remote servers” on page 727.

**Figure 11-1** Storage Unit node of the Storage utility



## About storage units

A storage unit is a label that NetBackup associates with physical storage. The label can identify a robot, a path to a volume, or a disk pool.

The creation of any storage unit type consists of the following general steps:

- Name the storage unit. A configured storage unit indicates to NetBackup the existence of physical storage.
- Choose the storage unit type: Media Manager, disk, or NDMP.  
See Figure 11-2 on page 368.

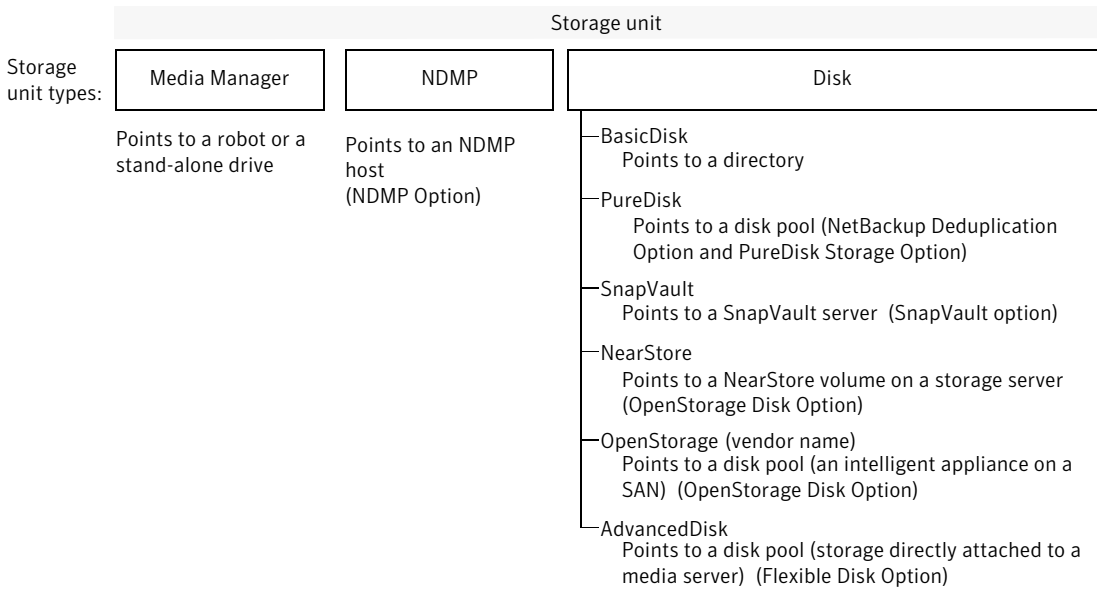
- Select a media server. The selection indicates that the media server(s) have permission to write to the storage unit.
- Indicate the destination where the data is written.
- For Media Manager storage units: Data is written to tape robots and stand-alone tape drives.

For disk storage: NetBackup permits an unlimited number of disk storage units. Disk storage may be one of the following types:

- AdvancedDisk storage units: The destination is a disk pool.
- BasicDisk storage units: The destination is a path to a volume on a host.
- NearStore storage units: The destination is a NearStore volume on a storage server.
- OpenStorage storage units: The destination is a disk pool.
- PureDisk storage unit: The destination is a disk pool.
- SharedDisk storage units: The destination is a disk pool.  
See [“About SharedDisk support in NetBackup 7.0 and later”](#) on page 362.
- SnapVault storage: The destination is a SnapVault server.
- For NDMP storage: The NDMP protocol is used to perform backups and recoveries. The destination is an NDMP host.

[Figure 11-2](#) shows the different storage unit types and the option that needs to be installed, if necessary.

**Figure 11-2** Storage unit types



## Creating a storage unit using the Device Configuration Wizard

The following procedure describes how to create a storage unit by using the Device Configuration Wizard.

### To create a storage unit with the Device Configuration Wizard

- 1 In the NetBackup Administration Console tree, select the **Master Server** or **Media and Device Management**.
- 2 From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

For help while running the wizard, click the **Help** button in the wizard screen.

## Creating a storage unit using the Actions menu

The following procedure describes how to create a storage unit from the Actions menu.

### To create a storage unit from the Actions menu

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage**.
- 2 Click **Actions > New > New Storage Unit**.



- 3 Complete the fields on the New Storage Unit dialog box.  
See [“About storage unit settings”](#) on page 380.
- 4 Click **OK** to add the storage unit to the configuration.

## Creating a storage unit by copying a storage unit

The following procedure describes how to create a storage unit by copying a storage unit.

### To create a storage unit by copying an existing storage unit

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage**.
- 2 Select a storage unit in the Details pane.
- 3 Click **Actions > Copy Storage Unit**.
- 4 Type a unique name for the new storage unit. For example, describe the type of storage. Use this name to specify a storage unit for policies and schedules.  
See [“NetBackup naming conventions”](#) on page 719.
- 5 Complete the fields in the Copy Storage Unit dialog box.  
See [“About storage unit settings”](#) on page 380.

## Changing storage unit settings

Symantec suggests that changes be made only during periods when no backup activity is expected for the policies that use the affected storage units.

### To change storage unit settings

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage**.
- 2 Double-click the storage unit you want to change from those listed in the Details pane.  
  
Hold down the Control or Shift key to select multiple storage units.
- 3 Complete the fields on the Change Storage Unit dialog box.  
See [“About storage unit settings”](#) on page 380.

## Deleting storage units

To delete a storage unit from a NetBackup configuration means to delete the label that NetBackup associates with the physical storage.

To delete a storage unit does not prevent files from being restored that were written to that storage unit, provided that the storage was not physically removed and the backup image has not expired.

#### To delete a BasicDisk or Media Manager storage unit

- 1 Use the **Catalog** utility to expire any images that exist on the storage unit. This removes the image from the NetBackup catalog.

See [“Expiring backup images”](#) on page 665.

- Do not manually remove images from the BasicDisk or Media Manager storage unit.
- Once the images are expired, they cannot be restored unless the images are imported.

See [“Importing backups”](#) on page 666.

NetBackup automatically deletes any image fragments from a disk storage unit or a disk pool. This generally occurs within seconds of expiring an image. However, to make sure that all of the fragments are deleted, check the directory on the storage unit to make sure that it is empty.

- 2 Select the **Storage** utility, then **Storage Units**.
- 3 In the Details pane, select the storage unit you want to delete. Hold down the Control or Shift key to select multiple storage units.
- 4 Select **Edit > Delete**.
- 5 In the confirmation dialog box, select the storage units to delete.
- 6 Click **OK**.
- 7 Modify any policy that uses a deleted storage unit to use another storage unit.

If a storage unit points to disk pool, the storage unit can be deleted without impacting the disk pool.

See [“Expiring backup images”](#) on page 665.

## Media Manager storage unit considerations

To create a storage unit of a tape robot or a stand-alone tape drive, select Media Manager as the **Storage unit type**.

See [“About storage unit settings”](#) on page 380.

Figure 11-3 Media Manager storage unit settings

**Change Storage Unit**

Storage unit name: orbiter-hcart2-robot-tld-0

Storage unit type: Media Manager  On demand only

Disk type:

Properties

Storage device: tld(0) - hcart2

|               |                               |
|---------------|-------------------------------|
| Robot type:   | TLD - Tape Library DLT        |
| Density:      | hcart2 - 1/2 Inch Cartridge 2 |
| Robot number: | 0                             |

Media server: orbiter

Maximum concurrent write drives: 2  Reduce fragment size to: 1048576 Megabytes

Enable Multiplexing  
Maximum streams per drive: 1

OK Cancel Help

When NetBackup sends a job to a Media Manager storage unit, it requests resources from the Enterprise Media Manager (EMM). Then NetBackup requests that Media Manager mount the volume in a drive.

If a stand-alone drive does not contain media or if a required volume is not available to a robot, a mount request appears in the **Pending Requests** pane of the Device Monitor. An operator can then find the volume, mount it manually, and assign it to the drive.

## About adding a Media Manager storage unit

Take the following items into consideration when adding a Media Manager storage unit:

- If using NetBackup Enterprise Server:  
Add the storage unit to the master server. Specify the media server where the drives attach.
- If using NetBackup Server:  
Add the storage unit to the master server where the drives attach. The robotic control must also attach to that server.
- The number of storage units that you must create for a robot depends on the robot's drive configuration as follows:
  - Drives with identical densities must share the same storage unit on the same media server. If a robot contains two drives of the same density on the same media server, add only a single storage unit for the robot. Set the **Maximum concurrent write drives** setting to 2.  
See [“Maximum concurrent write drives setting”](#) on page 383.
  - Drives with different densities must be in separate storage units. Consider an STK SL500 library that is configured as a Tape Library DLT (TLD). It can have both half-inch cartridge and DLT drives. Here, you must define a separate storage unit for each density.
  - Applies only to NetBackup Enterprise Server. If a robot's drives and robotic control attach to different NetBackup servers, specify the server where the drives attach as the media server. Always specify the same robot number for the drives as is used for the robotic control.
- Stand-alone drives with the same density must be in the same storage unit.  
For example, if a server has two 1/4-inch qscsi drives, add a storage unit with **Maximum concurrent write drives** set to 2. Media and device selection logic chooses the drive to use when NetBackup sends a backup to this storage unit. The logic is part of the Enterprise Media Management (`nibemm`) service.
- Stand-alone drives with different densities must be in different storage units.
- A robot and a stand-alone drive cannot be in the same storage unit.

## Disk storage unit considerations

NetBackup permits the creation of an unlimited number of disk storage units.

Disk media can be one of the following disk types:

See [“AdvancedDisk storage units”](#) on page 374.

See [“BasicDisk storage units”](#) on page 374.

See [“NearStore storage units”](#) on page 374.

See [“OpenStorage storage units”](#) on page 374.

See “[PureDisk storage units](#)” on page 375.

See “[SnapVault storage units](#)” on page 376.

Not all settings are available on each disk storage unit type.

See “[About storage unit settings](#)” on page 380.

---

**Note:** Symantec recommends that quotas are not imposed on any file systems that NetBackup uses for disk storage units. Some NetBackup features may not work properly when file systems have quotas in place. (For example, the capacity-managed retention selection in lifecycles and staging to storage units.)

---

## About the disk storage model

The NetBackup model for disk storage accommodates all Enterprise Disk Options. That is, it is the model for all disk types except for the Basic Disk type.

The following items describe components of the disk storage model:

- Data mover

An entity that moves data between the primary storage (the NetBackup client) and the storage server. NetBackup media servers function as data movers.

- Depending on the Enterprise Disk Option, a NetBackup media server also may function as a storage server.

- Storage server

An entity that writes data to and reads data from the disk storage. A storage server is the entity that has a mount on the file system on the storage.

Depending on the NetBackup option, the storage server is one of the following:

- A host that is part of a storage appliance or filer

- A NetBackup media server

- Disk pool

A collection of disk volumes that are administered as an entity. NetBackup aggregates the disk volumes into pools of storage (a disk pool) you can use for backups.

A disk pool is a storage type in NetBackup. When you create a storage unit, you select the disk type and then you select a specific disk pool.

The following topics describe the disk storage unit types.

## BasicDisk storage units

A BasicDisk type storage unit consists of a directory on a locally-attached disk or a network-attached disk that is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory.

Notes about the BasicDisk type storage unit:

- Do not include the same volume or file system in multiple BasicDisk storage units.
- BasicDisk storage units cannot be used in a storage lifecycle policy.

## AdvancedDisk storage units

An AdvancedDisk disk type storage unit is used for a dedicated disk that is directly attached to a NetBackup media server. An AdvancedDisk selection is available only when the Flexible Disk Option is licensed.

NetBackup assumes exclusive ownership of the disk resources that comprise an AdvancedDisk disk pool. If the resources are shared with other users, NetBackup cannot manage disk pool capacity or storage lifecycle policies correctly.

For AdvancedDisk, the NetBackup media servers function as both data movers and storage servers.

See the *NetBackup Shared Storage Guide*.

## NearStore storage units

A NearStore disk type storage unit is used to store images on Network Attached Storage (NAS) from NetApp. NearStore appears as a selection only when the OpenStorage Disk Option is licensed.

For NearStore, the NetBackup media servers function as the data movers. The NearStore host is the storage server.

---

**Note:** NearStore storage units cannot be used as part of a storage unit group or used in a storage lifecycle policy.

---

Information about configuring NearStore storage units is described in the *NetBackup Administrator's Guide, Volume II*.

## OpenStorage storage units

An OpenStorage disk type storage unit is used for disk storage on an intelligent disk appliance. The actual name of the disk type depends on the vendor. An

OpenStorage selection is available only when the OpenStorage Disk Option is licensed.

The disk appliance is integrated into NetBackup through an API. The storage vendor partners with Symantec to integrate the appliance into NetBackup. The disk appliance is the storage server

For OpenStorage, the NetBackup media servers function as the data movers. The storage vendor's plug-in must be installed on each media server that functions as a data mover. The logon credentials to the storage server must be configured on each media server.

See the *NetBackup Shared Storage Guide*.

## PureDisk storage units

A PureDisk disk type storage unit is used to store deduplicated data for the following NetBackup options:

- Media server deduplication pool.  
NetBackup deduplication must be configured.  
See the *NetBackup Deduplication Guide*.
- PureDisk deduplication pool (PureDisk 6.6 and later).  
NetBackup deduplication must be configured.  
See the *NetBackup Deduplication Guide*.
- PureDisk Deduplication Option (PDDO) storage pool (PureDisk 6.5 and later).  
PureDisk Deduplication Option (PDDO) must be configured.  
See the *NetBackup PureDisk Deduplication Option Guide*.

---

**Note:** PDDO storage units cannot be used as part of a storage unit group.

---

PureDisk appears as a selection when the NetBackup Deduplication Option or the PureDisk Storage Option is licensed.

## SharedDisk storage units

NetBackup SharedDisk is not support on NetBackup 7.0 and later media servers.

See “[About SharedDisk support in NetBackup 7.0 and later](#)” on page 362.

## SnapVault storage units

A SnapVault storage unit is used to store images on Network Attached Storage (NAS). The SnapVault selection is available only when the NetBackup Snapshot Client option is licensed.

SnapVault storage units cannot be used in a storage unit group or as part of a staging operation.

For SnapVault, the NetBackup media servers function as the data movers. The SnapVault host is the storage server.

## Disk storage units in storage lifecycle policies

This topic discusses how storage units can interact with storage lifecycle policies.

[Figure 11-4](#) is an example of how policies can interact with volumes in a disk pool that a storage unit references.

Two policies are created as follows:

- Policy\_gold has a gold classification. It is configured to use Lifecycle\_Gold, which has a gold data classification.
- Policy\_silver has a silver classification. It is configured to use Any Available storage unit. That means it can use any available storage unit or any lifecycle that has a silver classification.

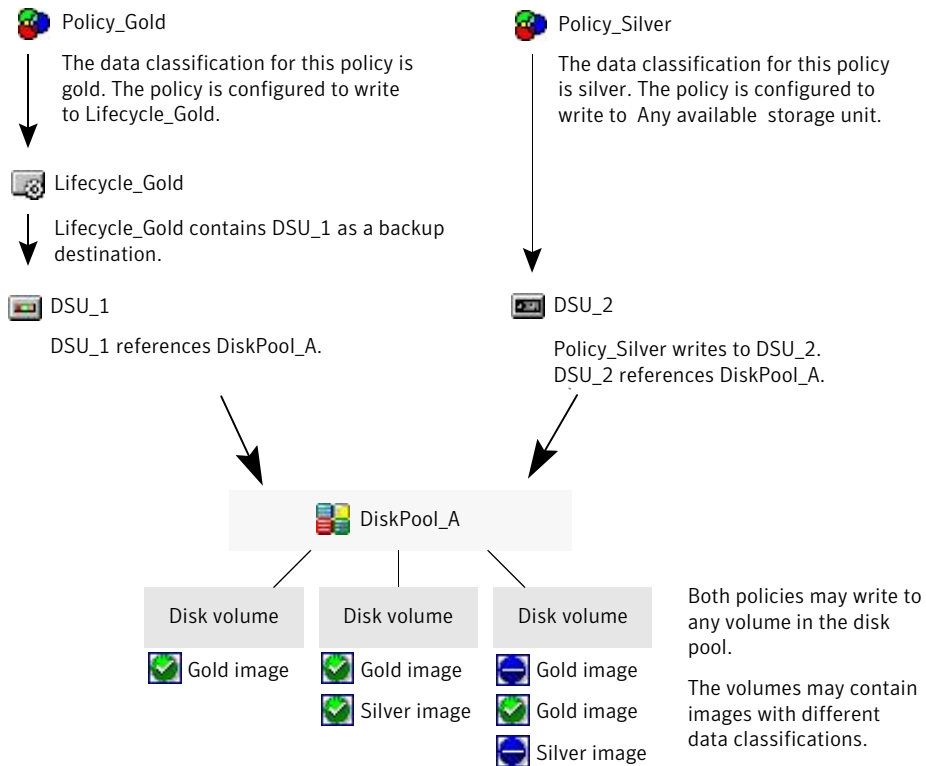
Two storage units are available as follows:

- DSU\_1 is a destination in Lifecycle\_Gold and references DiskPool\_A.
- DSU\_2 is not in a lifecycle. It references DiskPool\_A.

DiskPool\_A contains three disk volumes. Both the gold and the silver images can be written to any disk volume in the pool.



**Figure 11-4** Storage lifecycle policies and disk storage units referencing disk pools



See “[Storage lifecycle policy overview](#)” on page 415.

## Maintaining available space on disk storage units

Disk storage units can be managed so that they do not become entirely full and cause backups to fail.

The following list describes how space can be created for more images on a disk storage unit:

- Add new disk space.
- Set the **High water mark** to a value that best works with the size of backup images in the environment.

See “[High water mark setting](#)” on page 382.

To maintain space on basic disk staging storage units, consider the following:

- Increase the frequency of the relocation schedule. Or, add resources so that all images can be copied to a final destination storage unit in a timely manner.
- Upon NetBackup installation or upgrade, the `nb_updatedssu` script runs. The script deletes the `.ds` files that were used in previous releases as pointers to relocated data. Relocated data is tracked differently in the current release and the `.ds` files are no longer necessary. Under some circumstances, a `.ds` file cannot be deleted upon installation or upgrade. In that case, run the script again:

```
install_path\netbackup\bin\goodies\nb_updatedssu
```

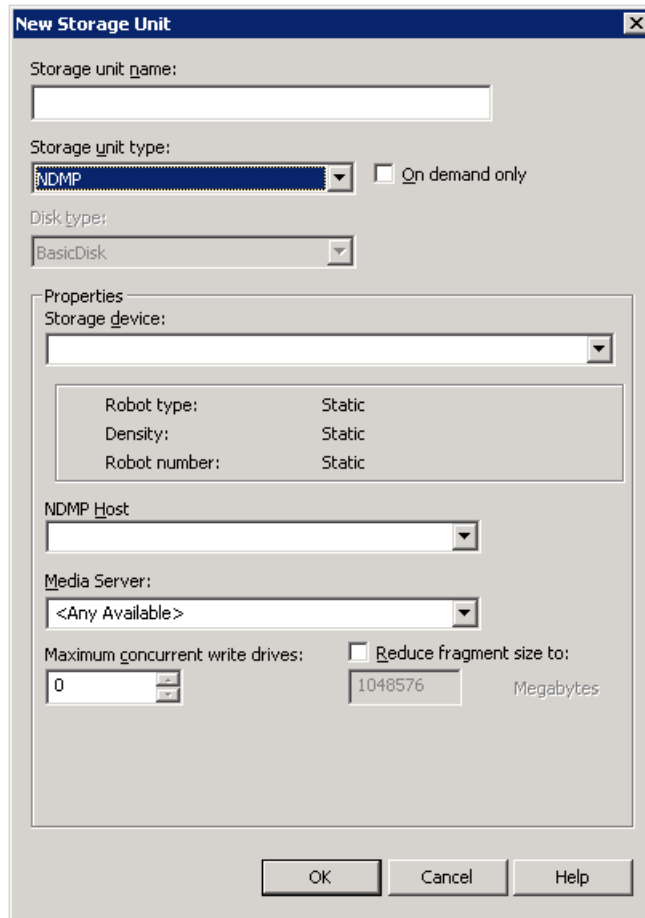
- The potential free space can be determined on a disk staging storage unit. See [“Finding potential free space on a BasicDisk disk staging storage unit”](#) on page 400.
- The General Server host property **Check the capacity of disk storage units** property determines how often NetBackup checks 6.0 disk storage units for available capacity. Subsequent releases use internal methods to monitor disk space more frequently. See [“Check the capacity of disk storage units every”](#) on page 128.

## About NDMP storage unit considerations

The NetBackup for NDMP license must be installed on the media server to use the hosts as storage units. Media Manager controls NDMP storage units but the units attach to NDMP hosts.

See [“About storage unit settings”](#) on page 380.

**Figure 11-5** NDMP storage unit settings



Create NDMP storage units for drives directly attached to NAS filers. Any drive that is attached to a NetBackup media server is considered a Media Manager storage unit, even if used for NDMP backups.

---

**Note:** Remote NDMP storage units may already be configured on a media server from a previous release. Upon upgrade of the media server, those storage units are automatically converted to Media Manager storage units.

---

See the *NetBackup for NDMP Administrator's Guide* for more information.

## About storage unit settings

The following topics describe the settings that appear for all types of storage units. The settings are listed alphabetically. Each setting does not appear for each storage unit type.

### Absolute pathname to directory or volume setting

**Absolute pathname to directory** or **Absolute pathname to volume** is available for any storage units that are not based on disk pools. For example, BasicDisk storage units.

The setting specifies the absolute path to a file system or a volume available for disk backups. Enter the path directly in the field, then click **Add**. Use any location on the disk, providing that sufficient space is available.

The **Properties** button displays properties about the directory or volume.

See [“Properties button”](#) on page 388.

Do not configure multiple BasicDisk storage units to use the same volume or file system. Not only do the storage units compete for space, but different **Low water marks** can cause unexpected behaviors.

If the BasicDisk storage unit is used as a disk staging storage unit, Symantec recommends dedicating a disk partition or file system to it. Dedicating space allows the disk staging space management logic to operate successfully. Or, consider defining AdvancedDisk storage units, which use disk pools that are composed of the disk volumes that are dedicated file systems for disk backup.

Use platform-specific file path separators (/ and \) and colon (:) within a drive specification.

See [“NetBackup naming conventions”](#) on page 719.

### Directory can exist on the root file system or system disk

Enable this BasicDisk storage unit setting to allow the directory that is created in the **Absolute pathname to directory** field to exist on the following:

- The root file system (UNIX)
- A system drive (Windows)

When this setting is enabled, the directory is created automatically. If a storage unit is configured on C: drive and this option is not enabled, backups fail with error code 12.

---

**Note:** With this setting enabled, the system drive can fill up.

---

A job fails under the following conditions:

## Density setting

The **Storage device** selection determines the media **Density**. This setting appears for Media Manager and NDMP storage units only.

## Disk pool setting

A **Disk pool** is selected for any disk storage unit with an Enterprise Disk Option licensed. Select the disk pool that contains the storage for this storage unit.

The following items describe which disk pools appear in the drop-down list:

- For AdvancedDisk, all NetBackup disk pools appear in the **Disk pool** list.
- For OpenStorage, only the disk pools for that OpenStorage vendor's appliance appear in the list.
- For PureDisk, the media server deduplication pools, the PureDisk deduplication pools, and the PureDisk Storage Pool Authority (SPA) appear in the list.

## Disk type selection

A Disk type is selected for a disk storage unit.

A disk storage unit can be one of the following types:

- AdvancedDisk (NetBackup Flexible Disk Option needed)
- BasicDisk
- NearStore (OpenStorage Disk Option needed)
- OpenStorage (vendor name) (NetBackup OpenStorage Disk Option needed)
- PureDisk (NetBackup Deduplication Option or PureDisk Storage Option needed)
- SharedDisk (NetBackup Flexible Disk Option needed)  
See "[About SharedDisk support in NetBackup 7.0 and later](#)" on page 362.

SnapVault (NetBackup Snapshot Client option needed).

For information on SnapVault storage units, see the *NetBackup Snapshot Client Administrator's Guide*.

## Enable block sharing setting

The **Enable block sharing** setting is available to some disk storage types. The **Enable block sharing** setting lets the data blocks that have not changed from one backup to the next be shared. To share data blocks can significantly save disk space in the storage unit.

## Enable multiplexing setting

The **Enable multiplexing** setting allows multiple backups to multiplex onto a single drive in a storage unit.

## High water mark setting

The **High water mark** setting is a threshold that, when reached, signals to NetBackup that the disk storage unit should be considered full. Default: 98%.

As the disk storage capacity grows closer to the **High water mark**, NetBackup reduces the number of jobs sent to the storage unit. NetBackup does not assign new jobs to a storage unit that is considered full. Once the capacity of the storage unit is less than the **High water mark**, NetBackup assigns jobs to the storage unit.

Try to prevent the situation where multiple jobs write to a storage unit at one time and fill it to capacity. Once the storage unit is full, none of the jobs can complete and all the jobs fail due to a disk full condition. To reduce the number of jobs that are allowed to write to the storage unit, decrease the **Maximum concurrent jobs** setting.

If the **High water mark** is set for a disk staging storage unit, nearing the high water mark triggers a cleanup of the disk staging storage unit until the **Low water mark** is met.

See [“Maximum concurrent jobs setting”](#) on page 383.

## Low water mark setting

Once the **High Water Mark** is reached, space is created on the disk storage unit until the **Low Water Mark** is met. Default: 80%.

The **Low water mark** setting has no effect unless backups are written through a storage lifecycle policy, using the capacity managed retention type. NetBackup copies expired images to a final destination storage unit to create space.

See [“Staged capacity managed retention type for storage destinations”](#) on page 425.

The **Low water mark** setting cannot be greater than the **High water mark** setting.

---

**Note:** Basic disk staging storage units may already be configured on a media server of a previous release. Upon upgrade, the disk storage units are set with the **Low water mark** at 100%. To make the best use of upgraded storage units, adjust the level.

---

For the disk storage units that reference disk pools, the low water mark applies to the disk pool.

See the *NetBackup Deduplication Guide*.

See the *NetBackup Shared Storage Guide*.

## Maximum concurrent write drives setting

The **Maximum concurrent write drives** setting specifies the number of tape drives that NetBackup can use at one time for jobs to this storage unit. The number of tape drives available is limited to the maximum number of tape drives in the storage device. If a job contains multiple copies, each copy applies toward the **Maximum concurrent write drives** count.

Select one number for the following:

- Storage unit that contains only stand-alone tape drives.  
Specify a number that is less than or equal to the number of tape drives that are in the storage unit.
- Robot.  
Specify a number that is less than or equal to the number of tape drives that attach to the NetBackup media server for the storage unit.

Assume that you have two stand-alone drives of the same density and specify 1. Both tape drives are available to NetBackup but only one drive can be used for backups. The other tape drive is available for restores and other non-backup operations. (For example, to import, to verify, and to duplicate backups.)

---

**Note:** To specify a **Maximum concurrent write drives** setting of 0 disables the storage unit.

---

## Maximum concurrent jobs setting

The **Maximum concurrent jobs** setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. (Default: one job. The job count can range from 0 to 256.) This setting corresponds to the **Maximum concurrent write drives** setting for a Media Manager storage unit.

If three backup jobs are ready to be sent to the storage unit and **Maximum concurrent jobs** is set to two, the first two jobs start and the third job waits. If a job contains multiple copies, each copy applies toward the **Maximum concurrent jobs** count.

---

**Note:** To specify a **Maximum concurrent jobs** setting of 0 disables the storage unit.

---

The **Maximum concurrent jobs** setting can be used to balance the load between disk storage units. A higher value (more concurrent jobs) means that the disk may be busier than if the value was set for fewer jobs.

The media server load balancing logic considers all storage units and all activity. A storage unit can indicate three media servers. If **Maximum concurrent jobs** is set to three and two of the media servers are busy or down, the third media server is assigned all three jobs.

The **Maximum concurrent jobs** setting depends on the available disk space and the server's ability to run multiple backup processes. Where disk pools are used, the setting also depends on the number of media servers in the storage unit.

If multiple storage units reference the same disk pool, the number of concurrent jobs that can access the pool is the sum of the **Maximum concurrent jobs** settings on all of the disk storage units. The setting applies to the storage unit and not to the disk pool. Therefore, the job load is automatically spread across the media servers that the storage unit configuration indicates.

---

**Note:** Increase the **Maximum concurrent jobs** setting if the storage unit is used for catalog backups as well as non-catalog backups. Increase the setting to ensure that the catalog backup can proceed while regular backup activity occurs. Where disk pools are used, increase the setting if more than one server is in the storage unit.

---

## Impact when two disk storage units reference one disk pool

Figure 11-6 shows how the **Maximum concurrent jobs** settings are combined when two disk storage units share one disk pool.

For example, DSU\_1 is configured as follows:

- To use MediaServer\_A
- To have a **Maximum concurrent jobs** setting of two
- To reference Disk\_pool1

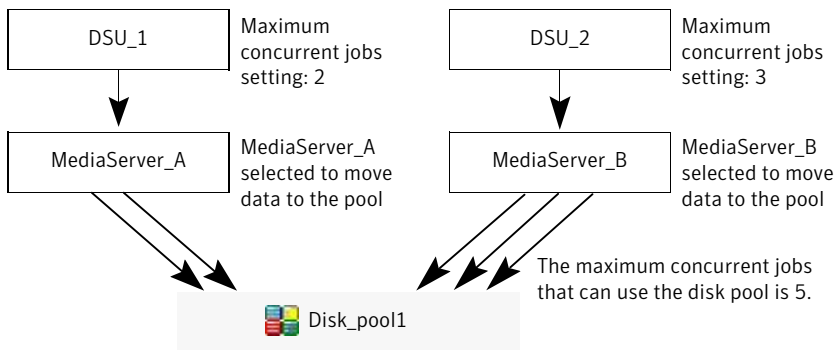
DSU\_2 is configured as follows:



- To use MediaServer\_B
- To have a **Maximum concurrent jobs** setting of three
- To reference Disk\_pool1

Both storage units reference the same disk pool. Combined, the storage units have a **Maximum concurrent jobs** setting of five. However, only two jobs can run concurrently on MediaServer\_A; three on MediaServer\_B.

**Figure 11-6** Impact when disk storage units use one disk pool but different media servers



If the storage units were configured to use both media servers, the media servers could run five concurrent jobs: two from DSU\_1 and three from DSU\_2.

## Maximum streams per drive setting

The **Maximum streams per drive** setting determines the maximum number of concurrent, multiple client backups that NetBackup can multiplex onto a single drive. The range is from 2 to 32.

See “[About multiplexing](#)” on page 513.

See “[Media multiplexing attribute](#)” on page 512.

## Media server setting

The following setting applies only to NetBackup Enterprise Server.

The Media server setting specifies one of the following:

- The NetBackup media server where the drives in the storage unit attach.
- The NetBackup media server that controls the disk storage unit.
- The NetBackup media servers that can move data to and from the disk pool.

- The NetBackup media servers that function as deduplication servers.

To make this storage unit available to any media server (default), select **Any Available**. NetBackup selects the media server dynamically at the time the policy is run.

Consider the following items, depending on the type of storage:

- BasicDisk storage

To configure a disk storage unit, select a single media server.

- AdvancedDisk storage units

The **Media server** setting specifies the NetBackup media servers that can move data to and from the disk pool.

The media servers that are configured as storage servers appear in the media servers list. The disk storage must be directly attached to the media server that is configured as the storage server.

NetBackup selects a media server when the policy runs.

- NDMP storage

To configure an NDMP storage unit, the **Media server** setting specifies the name of the media server that is to back up the NDMP host. Only those media servers that can talk to the specified NDMP storage device appear in the drop-down menu.

An NDMP host can be authenticated on multiple media servers. Select **Any Available** to have NetBackup select the media server and storage unit at the time the policy is run.

- OpenStorage storage units

The **Media server** setting specifies the NetBackup media servers that can move data to the storage server.

To allow any media server in the media server list to access the storage (default), select **Use Any Available Media Server**.

To restrict the media servers that can access the storage, select **Only Use The Following Media Servers**. Then select the media servers that are allowed to access the storage.

The following is required on each media server that accesses the storage:

- The vendor's software plug-in is installed.
- The login credentials to the storage server are configured.

Only the media servers on which storage server credentials are configured appear in the media servers list. If a server does not appear, verify that the software plug-in is installed and that login credentials are configured for that media server.

---

**Note:** Run the `tpconfig` command line utility directly on the media server to configure and verify credentials.

---

NetBackup selects a media server when the policy runs.

- SharedDisk storage units  
See “[About SharedDisk support in NetBackup 7.0 and later](#)” on page 362.
- PureDisk storage units (media server deduplication pool and PureDisk deduplication pool)  
To allow any media server in the list to deduplicate data, select **Use Any Available Media Server**.  
To restrict the media servers that can deduplicate data, select **Only Use The Following Media Servers**. Then select the media servers that are allowed to deduplicate the data.  
Each media server must be configured as a deduplication media server.  
See the *NetBackup Deduplication Guide*.
- PureDisk storage units (PureDisk Deduplication Option storage pool)  
To allow any media server in the list to access the storage (default), select **Use Any Available Media Server**.  
To restrict the media servers that can access the storage, select **Only Use The Following Media Servers**. Then select the media servers that are allowed to access the storage.  
NetBackup selects a media server when the policy runs.  
The following is required on each media server that accesses the storage:
  - The PureDisk agent is installed.
  - The logon credentials to the PureDisk server are configured on the media server.See the *NetBackup PureDisk Remote Office Edition Administrator’s Guide* for the media server requirements.

## NDMP host setting

The **NDMP host** setting specifies the NDMP tape server that is used to write data to tape. Select the host name from the drop-down menu or click **Add** to add a host.

## On demand only setting

The **On demand only** setting specifies whether the storage unit is available exclusively on demand. That is, only when a policy or schedule is explicitly

configured to use this storage unit. Clear the **On demand only** check box to make the storage unit available to any policy or schedule.

For SnapVault and NearStore storage units, **On demand only** is selected by default and cannot be changed.

---

**Note:** If **On demand only** is selected for all storage units, be sure to designate a specific storage unit for each policy or schedule. Otherwise, NetBackup is unable to find a storage unit to use.

---

## Only use the following media servers

To restrict the media servers for the storage, select **Only Use The Following Media Servers**. Then select the media servers that you want to use with the storage.

The following items describe the media server functionality for each type of storage:

- For AdvancedDisk storage, the media servers are both storage servers and data movers. The media servers that are configured as the storage servers and data movers appear in the media servers list.
- For OpenStorage, the media servers that are configured as data movers for the OpenStorage implementation appear in the media server list. (For OpenStorage, NetBackup media servers function as data movers.) If a media server does not appear in the list, verify that the software plug-in is installed and that logon credentials are created.

The following is required on each media server that accesses the storage:

- The vendor's software plug-in is installed.
- The logon credentials to the storage server are configured.
- For PureDisk (media server deduplication pool and PureDisk deduplication pool), the media servers function as deduplication servers. NetBackup deduplication must be configured.
- For PureDisk (PureDisk Deduplication Option storage pool), the NetBackup media servers function as the data movers. The PureDisk Linux servers function as the storage servers. PureDisk Deduplication Option (PDDO) must be configured.

## Properties button

Click **Properties** to display information about the volume or the disk pool, as follows:

- **Available storage or Available**

This value reflects the space that remains for storage on a disk storage unit. The following equation determines the available space:

Available space = free space + potential free space - committed space

The `df` command may report a value for the available space that is slightly different from the actual free space value that appears as a result of the `nbdevquery` command:

```
nbdevquery -listdv -stype server_type -dp disk_pool
```

The available space that the `df` command lists does not include the space that the operating system reserves. Since NetBackup runs as `root`, the `nbdevquery` command includes the reserved space in the available space equation.

- **Capacity**

The **Capacity** value reflects the total amount of space that the disk storage unit or pool contains, both used and unused.

- **Disk pool comments**

Comment that are associated with the disk pool.

- **High water mark**

The high water mark for the disk pool applies to both the individual disk volumes in the pool and the disk pool:

- Individual volumes

When a disk volume reaches the high water mark, new jobs are not assigned to the volume. This is true for all disk types except BasicDisk staging storage units. The high water mark event triggers the deletion of images that have been relocated, to attempt to bring the used capacity of the disk volume down to the low water mark.

- Disk pool

When all volumes are at the high water mark, the disk pool is full. When a disk pool approaches the high water mark, NetBackup reduces the number of jobs that are allowed to write to the pool.

NetBackup does not assign new jobs to a storage unit in which the disk pool is full. (Default: 99%.)

- **Low water mark**

The low water mark for the disk pool. Once a disk volume fills to its high water mark, NetBackup attempts to delete enough relocated images to reduce the used capacity of the disk volume down to the low water mark. The low water mark setting cannot be greater than the high water mark setting.

---

**Note:** The **Low water mark** setting has no effect unless backups are written through a storage lifecycle policy, using the capacity-managed retention type.

---

- **Name**  
The name of the disk pool.
- **Number of volumes**  
The number of disk volumes in the disk pool.
- **% full**  
The percentage of storage that is currently in use on the volume.  
The `df` command may report a percentage used (Use%) value that is different from the **% full** value. (See the preceding **Available Storage** topic for a description of why the values appear differently.)
- **Raw size**  
The raw, unformatted size of the storage in the disk pool.
- **Usable size**  
The amount of usable storage in the disk pools.

## Reduce fragment size setting

The **Reduce fragment size** setting specifies the largest fragment size that NetBackup can create to store backups.

- For Media Manager storage units as follows:  
The default maximum fragment size for a Media Manager storage unit is 1000 terabytes. To specify a maximum fragment size other than the default, place a check in the **Reduce fragment size** check box. Then enter a value from 50 megabytes to 1,048,575 megabytes.  
Fragmenting multiplexed tape backups can expedite restores. Fragments allow NetBackup to skip to the specific fragment before searching for a file. Generally, NetBackup starts at the beginning of the multiplexed backup and reads tar headers until it finds the file.
- For disk storage units  
The default maximum fragment size for a disk storage unit is 524,288 megabytes. To specify a maximum fragment size other than the default, enter a value from 20 megabytes to 524,288 megabytes.  
Backups to disk are usually fragmented to ensure that the backup does not exceed the maximum size that the file system allows.  
The **Reduce fragment size** setting is intended primarily for storing large backup images on a disk type storage unit.

For media server deduplication pools and PureDisk deduplication pools, Symantec recommends against specifying the largest fragment size allowed (512 GB). For best results, the default fragment size is set to 50 GB.

---

**Note:** OpenStorage vendors may have special requirements for the maximum fragment size. Consult the vendor's documentation for guidance.

---

If an error occurs in a backup, the entire backup is discarded. The backup restarts from the beginning, not from the fragment where the error occurred. (An exception is for backups where checkpoint restart is enabled. In that case, fragments before and including the last checkpoint are retained; the fragments after the last checkpoint are discarded.)

---

**Note:** Basic disk staging units with different maximum fragment sizes may already be configured on a media server from a previous release. Upon upgrade, the disk storage units are not automatically increased to the new default of 524,288 megabytes. To make the best use of upgraded storage units, increase the fragment size on the upgraded storage units.

---

## Robot number setting

The **Storage device** selection determines the **Robot number**. It is the same robot number used in the Media Manager configuration.

## Robot type setting

The **Robot type** and indicates the type of robot (if any) that the storage unit contains. The **Storage device** setting determines the **Robot type**.

For the specific vendor types and models that correspond to each robot type, see the Supported Peripherals section of the NetBackup Release Notes.

## Staging relocation schedule setting (for basic disk staging only)

Click the **Staging schedule** button to configure the relocation schedule for this storage unit. A schedule is what makes the disk storage unit into a basic disk staging storage unit. During the relocation schedule, the backup image is duplicated from the temporary staging area to the final destination storage unit.

See “[Temporary staging area setting](#)” on page 392.

See “[Basic disk staging](#)” on page 396.

See “[About staging backups](#)” on page 395.

## Storage device setting

The **Storage device** list contains all possible storage devices available. Storage units can be created for the listed devices only.

## Storage unit name setting

Type a unique **Storage unit name** for the new storage unit. The name could describe the type of storage. The **Storage unit name** is the name used to specify a storage unit for policies and schedules.

The storage unit name cannot be changed after creation. The **Storage unit name** is inaccessible in a Change Storage Unit operation.

See “[NetBackup naming conventions](#)” on page 719.

## Storage unit type setting

The **Storage unit** type setting specifies the type of storage that this storage unit uses, as follow:

- Disk  
See “[Disk storage unit considerations](#)” on page 372.
- Media Manager  
See “[Media Manager storage unit considerations](#)” on page 370.
- NDMP  
See “[About NDMP storage unit considerations](#)” on page 378.

## Temporary staging area setting

If the storage unit is to be used as a temporary staging area, click the temporary staging checkbox. Then, configure the staging schedule.

See “[Staging relocation schedule setting \(for basic disk staging only\)](#)” on page 391.

The Staging column in the **Storage units** details pane indicates whether or not the unit is used as a temporary staging area for basic disk staging.

See “[Basic disk staging](#)” on page 396.

See “[Staging relocation schedule setting \(for basic disk staging only\)](#)” on page 391.

## Transfer throttle setting

The **Transfer throttle** setting appears for SnapVault storage units only.



This setting makes it possible to limit the amount of network bandwidth that is used for the SnapVault transfer. (In case bandwidth needs to be reserved for other applications.) Zero (default) means no network bandwidth limit for the SnapVault transfer; SnapVault uses all available bandwidth. (Range: 0 to 9999999.)

A value greater than 0 indicates a transfer speed for SnapVault in kilobytes per second. For example, a value of 1 sets a transfer speed limit for SnapVault of 1 kilobyte per second, which is a very slow transfer rate.

## Use any available media server setting

To allow any media server in the media server list to access the storage (default), select **Use Any Available Media Server**.

The following items describe the media server functionality for each type of storage:

- For AdvancedDisk storage, the media servers are both storage servers and data movers. The media servers that are configured as the storage servers and data movers appear in the media servers list.
- For OpenStorage, the media servers that are configured as data movers for the OpenStorage implementation appear in the media server list. (For OpenStorage, NetBackup media servers function as data movers.) If a media server does not appear in the list, verify that the software plug-in is installed and that logon credentials are created.

The following is required on each media server that accesses the storage:

- The vendor's software plug-in is installed.
- The login credentials to the storage server are configured.
- For PureDisk (media server deduplication pool and PureDisk deduplication pool), the media servers function as deduplication servers. NetBackup deduplication must be configured.
- For PureDisk (PureDisk Deduplication Option storage pool), the NetBackup media servers function as the data movers. The PureDisk Linux servers function as the storage servers. PureDisk Deduplication Option (PDDO) must be configured.



# Staging backups

This chapter includes the following topics:

- [About staging backups](#)

## About staging backups

Staging backups is the process in which a backup is written to a storage unit, then duplicated to a second storage unit. Eligible backups are deleted on the initial storage unit when space is needed for more backups.

This two-stage process allows a NetBackup environment to leverage the advantages of disk-based backups for recovery in the short term.

Staging also meets the following objectives:

- Staging allows for faster restores from disk.
- Staging allows the backups to run when tape drives are scarce.
- Staging allows data to be streamed to tape without image multiplexing.

## About the two staging methods

NetBackup offers the following methods to stage backups:

- **Basic disk staging**  
Basic disk staging consists of two stages. First, data is stored on the initial storage unit (disk staging storage unit). Then, per a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.  
See [“Basic disk staging”](#) on page 396.  
The following storage unit types are available for basic disk staging: BasicDisk, NearStore, and tape.

- Staging using the **Storage Lifecycle Policies** utility  
Staged backups that are configured within the **Storage Lifecycle Policies** utility also consist of two stages: Data on the staging storage unit is copied to a final destination, however, the data is not copied per a specific schedule. Instead, the administrator can configure the data to remain on the storage unit until either a fixed retention period is met, or until the disk needs additional space, or until the data is duplicated to the final location. No BasicDisk, SnapVault, or disk staging storage units can be used as destinations in a lifecycle.  
See [“Storage lifecycle policy overview”](#) on page 415.

The staging method that is used is determined in the policy Attributes tab. The **Policy storage unit/lifecycle** selection determines whether the backup goes to a storage unit or a lifecycle.

---

**Note:** Symantec recommends that a disk partition or file system be dedicated to any disk storage unit that is used for staging. Dedicated space allows the disk staging space management logic to operate successfully.

---


## Basic disk staging

The basic disk staging method is conducted in the following stages:

- Stage I  
Clients are backed up by a policy. The **Policy storage** selection in the policy indicates a storage unit that has a relocation schedule configured. The schedule is configured in the **New** or **Change Storage unit** dialog box by clicking **Staging Schedule**.
- Stage II  
Images are copied from the Stage I disk staging storage unit to the Stage II storage unit. The relocation schedule on the disk staging storage unit determines when the images are copied to the final destination. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

The image continues to exist on both the disk staging storage unit and the final destination storage units until the image expires or until space is needed on the disk staging storage unit.

[Figure 12-1](#) shows the stages is basic disk staging.

**Figure 12-1** Stage I and II of basic disk staging

When the relocation schedule runs, NetBackup creates a data management job. The job looks for any data that can be copied from the disk staging storage unit to the final destination. The Job Details in the Activity Monitor identify the job as one associated with basic disk staging. The Job Details list displays Disk Staging in the job's Data Movement field.

When NetBackup detects a disk staging storage unit that is full, it pauses the backup. Then, NetBackup finds the oldest images on the storage unit that successfully copied onto the final destination. NetBackup expires the images on the disk staging storage unit to create space.

## Creating a basic disk staging storage unit

Use the following process to create a basic disk staging storage unit.

### To create a basic disk staging storage unit

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage > Storage Units**.
- 2 Click **Actions > New > New Storage Unit**.
- 3 In the **New Storage Unit** dialog box do the following tasks:
  - Name the storage unit.  
See "[Storage unit name setting](#)" on page 392.
  - Select Disk as the **Storage unit type**.  
See "[Storage unit type setting](#)" on page 392.
  - Select the **Disk type** of disk storage unit that is to be a disk staging storage unit: BasicDisk or NearStore.
  - Select a media server.  
See "[Media server setting](#)" on page 385.
  - Enter the absolute path to the directory to be used for storage.

- See “ [Absolute pathname to directory or volume setting](#)” on page 380.
  - Select whether this directory can reside on the root file system or system disk.  
See “[Directory can exist on the root file system or system disk](#)” on page 380.
  - Enter the maximum concurrent jobs that are allowed to write to this storage unit at one time.  
See “[Maximum concurrent jobs setting](#)” on page 383.
  - The **Low water mark** setting has no effect unless backups are written through a storage lifecycle policy, using the capacity-managed retention type.
  - Enter a **High water mark** value.  
The high water mark works differently for the BasicDisk disk type. NetBackup assigns new jobs to a BasicDisk disk staging storage unit, even if it is over the indicated high water mark. For BasicDisk, the high water mark is used to trigger the deletion of images that have been relocated. NetBackup continues to delete images until the disk reaches the low water mark.
  - Check the **Enable temporary staging area** option. Once the option is enabled, the **Staging Schedule** button is enabled.
  - Click the Staging Schedule button.
- 4 The Disk Staging Schedule is similar to the scheduling dialog box used to configure policies. The differences appear on the **Attributes** tab.
- In the **Disk Staging Schedule** dialog box, perform the following tasks:
- The schedule name defaults to the storage unit name.
  - Select the priority that the relocation jobs that are started from this schedule have compared to other types of jobs.  
See “[Priority of relocation jobs started from this schedule](#)” on page 402.
  - Select whether to create Multiple Copies. With the **Multiple copies** attribute enabled, NetBackup can create up to four copies of a backup simultaneously.  
See “ [Multiple copies attribute](#)” on page 503.  
For disk staging storage units, the **Maximum backup copies** Global host property must include an additional copy beyond the number of copies that are indicated in the Copies field.  
See “[Maximum backup copies](#)” on page 135.
  - Select a storage unit to contain the images from this storage unit upon relocation.

- See [“Final destination storage unit”](#) on page 403.
- Select a volume pool to contain the images from this storage unit upon relocation.  
See [“Final destination volume pool”](#) on page 403.
  - Select a media owner to own the images from this storage unit upon relocation.  
See [“Final destination media owner”](#) on page 403.
  - Select whether to use an alternate server for the images from this storage unit upon relocation.  
See [“Use alternate read server”](#) on page 404.
- 5 Click **OK** to accept the disk staging schedule.
  - 6 Click **OK** to add the storage unit.

## Disk staging storage unit size and capacity

To take advantage of basic disk staging requires that the NetBackup administrator understand the life expectancy of the image on the Stage I storage unit.

The size and use of the file system of the Stage I storage unit directly affects the life expectancy of the image before it is copied to the Stage II storage unit. Symantec recommends a dedicated file system for each disk staging storage unit.

Consider the following example: A NetBackup administrator wants incremental backups to be available on disk for one week.

Incremental backups are done Monday through Saturday, with full backups done on Sunday. The full backups are sent directly to tape and do not use basic disk staging.

Each night's total incremental backups are sent to a disk staging storage unit and average from 300 MB to 500 MB. Occasionally a backup is 700 MB. Each following day the relocation schedule runs on the disk staging storage unit and copies the previous night's incremental backups to the final destination, a Media Manager (tape) storage unit.

### Minimum disk size for a basic disk staging storage unit

The minimum disk size is the smallest size that is required for the successful operation of the disk staging logic.

The minimum size must be greater than or equal to the largest combined size of the backups that are placed on the storage unit between runs of the disk staging schedule. (In our example, the disk images remain on the disk for one week.)

In this example, the relocation schedule runs nightly, and the largest nightly backup is 700 MB. Symantec recommends that you double this value to allow for any problems that may occur when the relocation schedule runs. To double the value gives the administrator an extra schedule cycle (one day) to correct any problems.

To determine the minimum size for the storage unit in this example, use the following formula:

Minimum size = Max data per cycle × (1 cycle + 1 cycle for safety)

For example: 1.4 GB = 700 MB × (1+1)

### **Average disk size for a basic disk staging storage unit**

The average disk size represents a good compromise between the minimum and the maximum sizes.

In this example, the average nightly backup is 400 MB and the NetBackup administrator wants to keep the images for one week.

To determine the average size for the storage unit in this example, use the following formula:

Average size = Average data per cycle × (number of cycles to keep data + 1 cycle for safety)

2.8 GB = 400 MB × (6 + 1)

### **Maximum disk size for a basic disk staging storage unit**

The maximum disk size is the recommended size needed to accommodate a certain level of service. In this example, the level of service is that disk images remain on disk for one week.

To determine the maximum size for the storage unit in this example, use the following formula:

Maximum size = Max data per cycle × (# of cycles to keep data + 1 cycle for safety)

For example: 4.9 GB = 700 MB × (6 + 1)

## **Finding potential free space on a BasicDisk disk staging storage unit**

Potential free space is the amount of space on a disk staging storage unit that NetBackup could free if extra space on the volume is needed. The space is the total size of the images that are eligible for expiration plus the images ready to be deleted on the volume.



To find the potential free space on a BasicDisk storage unit, use the `bpstulist` and the `nbdevquery` commands. Note that the name of the storage unit and disk pools are case sensitive.

Run `bpstulist -label` to find the disk pool name. In the case of BasicDisk storage units, the name of the disk pool is the same as the name of the BasicDisk storage unit. In the following example, the name of the storage unit is NameBasic:

```
bpstulist -label basic
NameBasic 0 server1 0 -1 -1 1 0 "C:\\" 1 1 524288 *NULL* 0 1 0 98 80 0 NameBasic server1
```

Run the `nbdevquery` command to display the status for the disk pool, including the potential free space. Use the following options, where:

`-stype server_type`

Specifies the vendor-specific string that identifies the storage server type. For a BasicDisk storage unit, enter `BasicDisk`.

`-dp`

Specifies the disk pool name. For a basic disk type, the disk pool name is the name of the BasicDisk storage unit.

```
nbdevquery -listdv -stype BasicDisk -dp NameBasic -D
```

The value is listed as `potential_free_space`.

```
Disk Volume Dump
name : <Internal_16>
id : <C:\>
diskpool : <NameBasic::server1::BasicDisk>
disk_media_id : <@aaaaf>
total_capacity : 0
free_space : 0
potential_free_space: 0
committed_space : 0
precommitted_space : 0
nbu_state : 2
sts_state : 0
flags : 0x6
num_read_mounts : 0
max_read_mounts : 0
num_write_mounts : 1
max_write_mounts : 1
system_tag : <Generic disk volume>
```

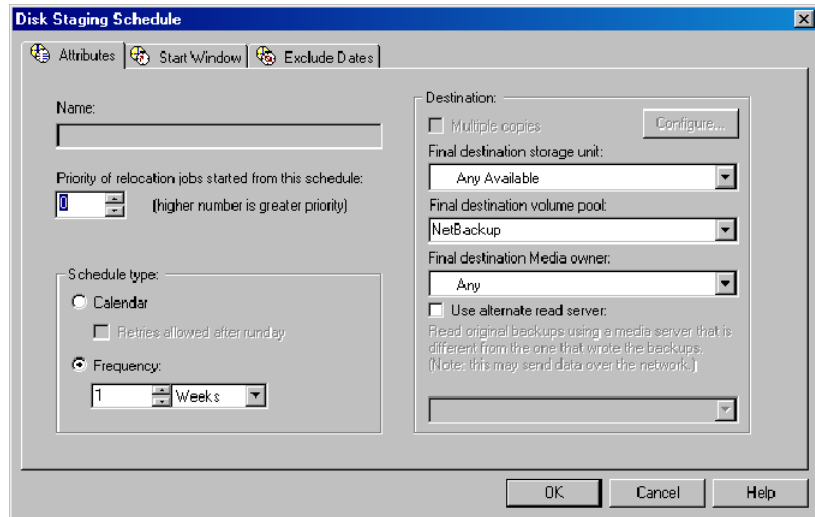
## About the Disk Staging Schedule dialog box

Click the **Staging Schedule** button to display the Disk Staging Schedule dialog box. The dialog box is similar to the scheduling dialog box that appears when a policy is configured.

The schedule that is created for the disk staging storage unit is not listed under **Schedules** in the NetBackup Administration Console when the **Policies** utility is selected.

**Figure 12-2** shows the disk staging schedule for a basic disk staging storage unit.

**Figure 12-2** Disk Staging Schedule for a basic disk staging storage unit



The **Attributes** tab on the **Disk Staging Schedule** dialog box differs from the **Attributes** tab of a regular policy. The differences are described in the following topics.

### Name field

The **Name** on the **Disk Staging Schedule** dialog box automatically defaults to the name of the storage unit.

### Priority of relocation jobs started from this schedule

The **Priority of relocation jobs started from this schedule** field indicates the priority that NetBackup assigns to relocation jobs for this policy. Range: 0 (default) to 99999 (highest priority).

## Final destination storage unit

If the schedule is a relocation schedule, a **Final destination storage unit** must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A **Final destination storage unit** is the name of the storage unit where the images reside after a relocation job copies them.

To copy images to tape, NetBackup uses all of the drives available in the **Final destination storage unit**. However, the **Maximum concurrent write drives** setting for that storage unit must be set to reflect the number of drives. The setting determines how many duplication jobs can be launched to handle the relocation job.

NetBackup continues to free space until the **Low water mark** is reached.

See “[Low water mark setting](#)” on page 382.

See “[Maximum concurrent write drives setting](#)” on page 383.

See “[About staging backups](#)” on page 395.

## Final destination volume pool

If the schedule is a relocation schedule, a **Final destination volume pool** must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A **Final destination volume pool** is the volume pool where images are swept from the volume pool on the basic disk staging storage unit.

See “[About staging backups](#)” on page 395.

---

**Note:** The relocation schedule that was created for the basic disk staging storage unit is not listed under **Schedules** in the NetBackup Administration Console when the **Policies** utility is selected.

---

## Final destination media owner

If the schedule is a relocation schedule, a **Final destination media owner** must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A **Final destination media owner** is the media owner where the images reside after a relocation job copies them.

Specify one of the following:

- **Any** lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured).
- **None** specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

- A server group. A server group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the **Final destination media owner** drop-down list.

## Use alternate read server

The **Alternate Read Server** attribute applies to NetBackup Enterprise Server only.

An alternate read server is a server allowed to read a backup image originally written by a different media server.

The path to the disk or directory must be identical for each media server that is to access the disk.

If the backup image is on tape, the media servers must share the same tape library or the operator must find the media.

If the backup image is on a robot that is not shared or a stand-alone drive, the media must be moved to the new location. An administrator must move the media, inventory the media in the new robot, and execute `bpmedia -oldserver -newserver` or assign a failover media server.

To avoid sending data over the network during duplication, specify an alternate read server that meets the following conditions:

- Connected to the storage device that contains the original backups (the source volumes).
- Connected to the storage device that contains the final destination storage units.

If the final destination storage unit is not connected to the alternate read server, data is sent over the network.

## Basic disk staging limitations

The basic disk staging method does not support the backup images that span disk storage units.

To avoid spanning storage units: do not use Checkpoint restart on a backup policy that writes to a storage unit group that contains multiple disk staging storage units.

See [“Checkpoint restart for backup jobs”](#) on page 468.

## Initiating a relocation schedule manually

A relocation schedule may be started manually to copy images to the final destination before the schedule is due to run.

### To initiate a relocation schedule

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage > Storage Units**.
- 2 Select a basic disk staging storage unit in the **Details** pane.
- 3 Select **Actions > Manual Relocation** to initiate the schedule.

If the relocation schedule finds data that can be copied, NetBackup creates a job to copy the data to the final destination storage unit.

The image then exists both storage units until the disk staging (Stage I) storage unit becomes full and the oldest images are deleted.

See [“Maintaining available space on disk storage units”](#) on page 377.



# Configuring storage unit groups

This chapter includes the following topics:

- [About Storage unit groups](#)
- [Creating a storage unit group](#)
- [Deleting a storage unit group](#)
- [Storage unit selection criteria within a group](#)
- [Disk spanning within storage unit groups](#)

## About Storage unit groups

Storage unit groups let you identify specific storage units as a group. You can specify a storage unit group name in a policy in the same way that you specify individual storage units. When a storage unit group is used in a policy, only the storage units that are specified in the group are candidates for the backup.

## Creating a storage unit group

The following procedure describes how to create a storage unit group.

**To create a storage unit group**

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Storage**.
- 2 Right-click **Storage Unit Groups** and select **New Storage Unit Group**.

- 3 Enter a storage unit group name for the new storage unit group.  
See “[NetBackup naming conventions](#)” on page 719.



The storage unit group name is case sensitive.

- 4 Add to or remove storage units from the group:
  - To add storage units to the group, select the storage units from the **Storage units not in the group** list. Click **Add**.
  - To remove storage units from the group, select the storage units from the **Storage units in group** list. Click **Remove**.
  - Storage units are listed in order of priority: The units at the top of the list have the highest priority in the group. To change the priority of a storage unit, select the storage unit and click **Move Up** or **Move Down**.



SnapVault, NearStore, and PureDisk storage units cannot be included in storage unit groups.

- 5 Choose how storage units are to be selected within the group:
  - **Prioritized.** Choose the first storage unit in the list that is not busy, down, or out of media.  
See [“Prioritized storage unit selection”](#) on page 410.
  - **Failover.** Choose the first storage unit in the list that is not down or out of media.  
See [“Failover storage unit selection”](#) on page 410.
  - **Round Robin.** Choose the least recently selected storage unit in the list.  
See [“Round robin storage unit selection”](#) on page 410.
  - **Media server load balancing.**  
Symantec recommends the **Media server load Balancing** criteria for disk staging storage units within a storage unit group.  
See [“Media server load balancing storage unit selection”](#) on page 411.

One exception to the selection criteria is in the case of a client that is also a media server with locally connected storage units.  
See [“Exception to the storage unit selection criteria”](#) on page 413.
- 6 Click **OK**.

## Deleting a storage unit group

The following procedure describes how to delete a storage unit group.

### To delete a storage unit group

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage > Storage Unit Groups**.
- 2 Select the storage unit group you want to delete from those listed in the **Details** pane. Hold down the Control or Shift key to select multiple storage units.
- 3 Select **Edit > Delete**.
- 4 Click **OK**.

## Storage unit selection criteria within a group

The storage unit selection criteria determines the order in which storage units are selected within a storage unit group.

Choose from one of the following selection criteria.

- [Prioritized storage unit selection](#)
- [Failover storage unit selection](#)
- [Round robin storage unit selection](#)
- [Media server load balancing storage unit selection](#)

The only difference between the selection criteria options is the order in which the storage units are selected.

A queue can form for a storage unit if the storage unit is unavailable.

The following are some reasons why a storage unit can be considered unavailable:

- The storage unit is busy.
- The storage unit is down.
- The storage unit is out of media.
- The storage unit has no available space.
- The storage unit has reached the **Maximum concurrent jobs** setting.  
See [“Maximum concurrent jobs setting”](#) on page 383.

## Prioritized storage unit selection

The **Prioritized** option indicates that NetBackup choose the next available storage unit in the list. (Default.)

If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.

## Failover storage unit selection

The **Failover** option indicates that if a job must queue for a storage unit, the job will queue rather than try another storage unit in the group.

## Round robin storage unit selection

The **Round robin** option indicates that NetBackup choose the least recently selected storage unit in the list as each new job is started.

If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.

## Media server load balancing storage unit selection

The **Media server load balancing** option indicates that NetBackup select a storage unit based on a capacity-managed approach. In this way, NetBackup avoids sending jobs to busy media servers.

If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.

The selection is based on the following factors:

- The rank of the media server.  
NetBackup considers the number of processes that are running on each CPU along with the memory thresholds on each server to determine the rank of a media server. If the free memory drops below a determined threshold, or if the number of running processes per CPU rises over a determined threshold, then the overall rank of the media server drops.
- The number of jobs on the media server. NetBackup considers the number of scheduled jobs on each media server.
- Whether the media server has enough disk space to accommodate the estimated size of the image. (Physical and virtual tapes ignore this requirement.)  
NetBackup estimates the size of any of the new or any current jobs on each media server. It then determines whether the jobs fit on a given volume.  
NetBackup estimates the amount of space that the job may require, based on previous backup history. If no history is available, the high water mark for the storage unit serves as a guide.

**Media server load balancing** cannot be selected for a storage unit group that includes a BasicDisk storage unit. Also, a BasicDisk storage unit cannot be included in an existing storage unit group with **Media server load balancing** enabled.

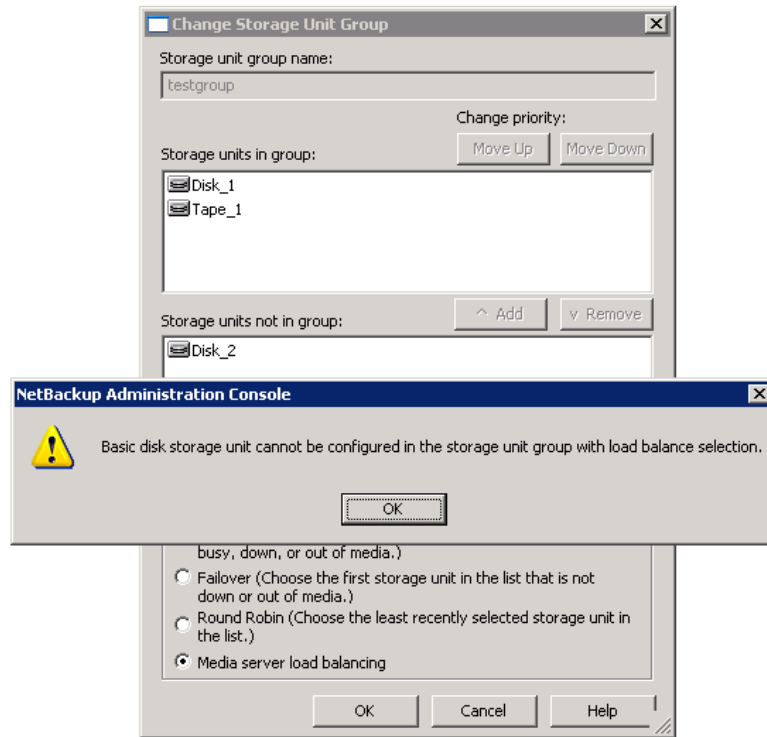
[Figure 13-1](#) shows the message that displays when this option is selected for a storage group that contains a BasicDisk storage unit.

---

**Note:** Symantec recommends that you select **Media server load balancing** for disk staging storage units within a storage unit group.

---

Figure 13-1 Message for a prohibited option in a storage unit group



## Other load balancing methods

The **Media server load balancing** option is one load balancing option and requires a license.

The following methods to distribute the backup workload do not require additional licenses:

- Adjust the backup load on a media server.
  - Change the **Limit jobs per policy** policy attribute for one or more of the policies that are sent to a media server. To decrease **Limit jobs per policy** reduces the workload on a media server on a specific network segment. See [“Limit jobs per policy attribute”](#) on page 471.
  - Reconfigure policies or schedules to use storage units on other media servers.
  - Consider changing the Bandwidth host properties on one or more clients. See [“Storage unit selection criteria within a group”](#) on page 409.

- Distribute the backup load on media servers during peak periods.  
Reconfigure policy schedules so that they write to storage units on the media servers that can handle the load. (Assuming that master servers and media servers are on separate hosts).
- Adjust the backup load on client.  
Change the **Maximum jobs per client** global attribute. For example, to increase **Maximum jobs per client** increases the number of concurrent jobs that any one client can process and therefore increases the load.  
See “[Storage unit selection criteria within a group](#)” on page 409.
- Reduce the time to back up clients.  
Increase the number of jobs that clients can perform concurrently, or use multiplexing. Another possibility is to increase the number of jobs that the media server can perform concurrently for the policies that back up the clients.
- Give preference to a policy.  
Increase the **Limit jobs per policy** attribute for the preferred policy relative to other policies. Or, increase the priority for the policy.  
See “[Limit jobs per policy attribute](#)” on page 471.
- Adjust the load between fast and slow networks.  
Increase the **Limit jobs per policy** and **Maximum jobs per client** for policies and clients in a faster network. Decrease these numbers for slower networks. Another solution is to use NetBackup’s capability to limit bandwidth.  
See “[Limit jobs per policy attribute](#)” on page 471.  
See “[Storage unit selection criteria within a group](#)” on page 409.
- Maximize the use of devices.  
Use multiplexing. Allow as many concurrent jobs per storage unit, policy, and client as possible without causing server, client, or network performance problems.
- Prevent backups from monopolizing tape devices.
  - Place some drives in a down state or limit the number that are used concurrently in a specific storage unit. For example, if there are four drives in a robot, allow only two to be used concurrently.
  - Do not place all devices under Media Manager control.

## Exception to the storage unit selection criteria

The only exception to the storage unit selection criteria order is in the case of a client that is also a media server with locally connected storage units. The locally available storage units take precedence over the defined sequence of storage units in the group.

You may have set up a storage unit to be **On demand only**. If the unit is in a storage unit group that a policy requires, the **On demand only** option is satisfied and the device is used.

See [“On demand only setting”](#) on page 387.

## Disk spanning within storage unit groups

A backup may span storage units if a disk full condition is detected. Backups can span from one BasicDisk storage unit to another BasicDisk storage unit if the storage units are in the same storage unit group. The storage units must also share the same media server.

See [“Storage unit selection criteria within a group”](#) on page 409.

# Configuring storage lifecycle policies

This chapter includes the following topics:

- [Storage lifecycle policy overview](#)
- [Creating a storage lifecycle policy](#)
- [Deleting a storage lifecycle policy](#)
- [Adding a storage destination to a storage lifecycle policy](#)
- [Hierarchical view of storage destinations](#)
- [Writing multiple copies using a storage lifecycle policy](#)
- [Storage lifecycle policy versions](#)
- [LIFECYCLE\\_PARAMETERS file for optional duplication job configuration](#)
- [Using the `nbstlutil` command to administrate lifecycle operations](#)

## Storage lifecycle policy overview

A storage lifecycle policy is a storage plan for a set of backups. A lifecycle policy is configured within the **Storage Lifecycle Policies** utility.

Essentially, a lifecycle is a list of destinations where copies of the backup images are stored, along with the prescribed retention period for each copy. After a lifecycle is configured, the lifecycle process works to create copies of the images on each destination. NetBackup retries the copies as necessary to ensure that all copies are created.

Lifecycles offer the opportunity for users to assign a classification to the data at the policy level. A data classification represents a set of backup requirements, which makes it easier to configure backups for data with different requirements. For example, email data and financial data.

Storage lifecycle policies can be set up to provide staging behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the storage lifecycle. This process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term. It also preserves the advantages of tape-based backups for long-term storage.

A storage lifecycle operation consists of the following steps:

- A backup is written to all destinations in the lifecycle.  
This process can occur if the NetBackup administrator has set up a lifecycle policy that contains at least one backup destination. The policy that writes the data must indicate that the backup data is to go to a lifecycle policy.
- NetBackup automatically copies the image to all duplication destinations in the lifecycle. The backup is retained on the backup destination until the retention period is met. Duplication destinations are optional and can provide another method for disk staging.
- The retention type that is selected for the destinations determines how long the backup resides on the destination. Eventually, NetBackup deletes the backup from the destinations to create more disk space.

## Creating a storage lifecycle policy

A storage lifecycle can be selected within a backup policy similarly to how a storage unit is selected in a policy. If a storage lifecycle is selected, the images that the policy creates are written to all the destinations that are defined in the storage lifecycle.

### To create a storage lifecycle policy

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage > Storage Lifecycle Policies**.
- 2 Click **Actions > New > New Storage Lifecycle Policy**.
- 3 In the **New Storage Lifecycle Policy** dialog box, enter a **Storage lifecycle policy name**.
- 4 Select a **Data classification**. (Optional.)  
See [“Data classification option”](#) on page 417.



- 5 Select the **Duplication job priority**. This is the priority that duplication jobs have in relationship to all other jobs. In duplication jobs, NetBackup duplicates data from a backup destination to a duplication destination within a lifecycle. See [“Duplication job priority setting”](#) on page 419.
- 6 Click **Add** to add storage destinations to the lifecycle. See [“Adding a storage destination to a storage lifecycle policy”](#) on page 421. See [“Adding a hierarchical duplication destination”](#) on page 429.
- 7 Click **OK** to create the storage lifecycle. After they are created, data classifications cannot be deleted.

## Storage lifecycle policy name

The **Storage lifecycle policy name** describes the storage lifecycle. The name can be modified.

## Data classification option

The **Data classification** allows the NetBackup administrator to classify data based on relative importance. A classification represents a set of backup requirements. When data must meet different backup requirements, consider assigning different classifications.

For example, email backup data can be assigned to the silver data classification and financial backup data backup may be assigned to the platinum classification. The financial data is assigned to the higher classification of platinum because backups of the financial data are consider more important.

The **Data classification** defines the level of data that the storage lifecycle is allowed to process. The **Data classification** drop-down menu contains all of the defined classifications. To select a classification is optional.

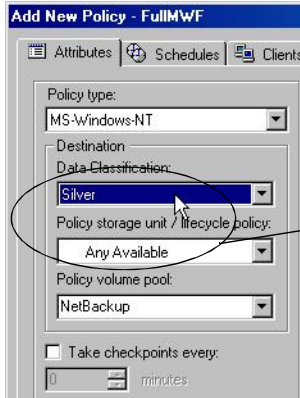
One data classification can be assigned to each storage lifecycle policy and applies to all destinations in the lifecycle. A storage lifecycle is not required to have a data classification.

If a data classification is selected, the storage lifecycle stores only those images from the policies that are set up for that classification. If no classification is indicated, the storage lifecycle accepts images of any classification or no classification.

## How to associate backup data with a data classification

A data classification is assigned in the backup policy to associate backup data with a data classification. Data from the policy can be stored only in a storage lifecycle policy with the same data classification assigned to it.

**Figure 14-1** Data classification assignment in the policy



If a classification is indicated, the Policy storage selection must contain a lifecycle with that assigned classification

After data is backed up to a storage lifecycle policy with an assigned classification, the data is managed according to the storage lifecycle configuration. The storage lifecycle defines what happens to the data from the initial backup until the last copy of the image expires.

If a classification is selected, all the images that the policy creates are tagged with that classification ID.

## Creating or changing a data classification

NetBackup contains four default data classifications. The name, the description, and the rank of each can be changed in the Data Classification host properties.

(In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers > Double-click on master server > Data Classifications.**)

New data classifications can also be created. However, data classifications cannot be deleted.

See [“Data Classification properties”](#) on page 102.

## Duplication job priority setting

The **Duplication job priority** setting is the priority that duplication jobs have in relationship to all other jobs. In duplication jobs, NetBackup duplicates data from a backup destination to a duplication destination within a lifecycle. Range: 0 (default) to 99999 (highest priority).

For example, the **Duplication job priority** for a policy with a gold data classification may be set higher than for a policy with a silver data classification.

The priority of the backup job is set in the backup policy on the **Attributes** tab.

See “[Job priority attribute](#)” on page 473.

## Deleting a storage lifecycle policy

To delete a storage lifecycle policy, use the following procedure:

### To delete a storage lifecycle policy

- 1 This step prevents new backup jobs from writing to the storage lifecycle policy.

Remove the storage lifecycle policy from all backup policies.

- 2 This step addresses in-process backup jobs writing to the storage lifecycle policy.

Wait for all in-process backup jobs to the storage lifecycle policy to complete or cancel the jobs using the Activity Monitor or command line.

- 3 This step prevents new duplication jobs from writing to the storage lifecycle policy.

To prevent any duplication jobs from writing to the storage lifecycle policy, use the following command:

```
nbstlutil cancel -lifecycle name
```

This command prevents the jobs by canceling any duplication jobs that were submitted to the storage lifecycle policy.

- 4 This step addresses in-process duplication jobs writing to the storage lifecycle policy.

Use the Activity Monitor to cancel in-process duplication jobs. Since one duplication job can contain images from multiple storage lifecycle policies, it can be difficult to determine which duplication job is associated with which storage lifecycle policy.

- 5 Once all of the operations are complete, delete the storage lifecycle policy. To delete a storage lifecycle policy deletes all versions of the definition.

---

**Note:** If orphaned images are detected due to a system error, NetBackup logs the fact that the images exist and alerts the administrator to address the situation.

---

If the administrator tries to delete a storage lifecycle policy with active images, a 1519 error appears. Wait several minutes and try to delete the storage lifecycle policy definition again until the error no longer appears.

- To use the Administration Console to delete a storage lifecycle policy:
  - Expand **Storage > Storage Lifecycle Policies**.
  - Select the storage lifecycle policy name.
  - Select **Edit > Delete**.
  - In the **Delete Storage Lifecycle Policies** dialog box, select the storage lifecycle policy name and click **OK**.

If images are still active for the storage lifecycle policy, a dialog box displays the following message:

```
The storage lifecycle policy, storage_lifecycle_name, could not be deleted. Status 1519.
```

- To use the `nbstl` command to delete a storage lifecycle policy, enter the following:

```
nbstl storage_lifecycle_name -delete
```

If images are still active, the following error appears:

```
C:\>nbstl storage_lifecycle_name -delete
Exit error: images are in process
EXIT status = 1519
```

# Adding a storage destination to a storage lifecycle policy

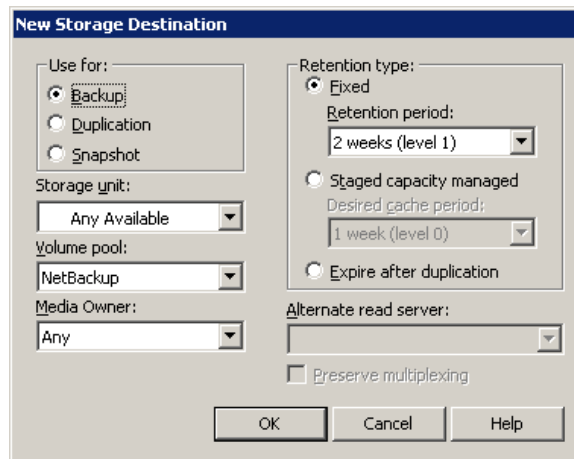
Use the following procedure to add a storage destination to a storage lifecycle policy:

## To add a storage destination to a lifecycle policy

- 1 In the NetBackup Administration Console, select **NetBackup Management > Storage > Storage Lifecycle Policies**.
- 2 Click **Actions > New > New Storage Lifecycle Policy**.
- 3 In the **New Storage Lifecycle Policy** dialog box, click **Add**.

To create a hierarchical duplication destination, select a destination to become the source of the destination to be added, then click **Add**.

See [“Adding a hierarchical duplication destination”](#) on page 429.



- 4 In the **New Destination** dialog box, select the purpose for which images are to be written to the new destination:
  - **Backup** images to be written to the destination as part of a backup operation.
  - **Duplication** images to be written to the destination as part of a duplication operation.
  - **Snapshot** images to be written to the destination as part of the snapshot operation.

---

**Note:** Snapshot destinations cannot be used as the source for any duplication.

---

See [“Use for: Backup, duplication, or Snapshot destination”](#) on page 423.

See [“Storage destination list requirements”](#) on page 423.

- 5 Indicate the **Storage unit** or storage unit group where the backups are to be written. No BasicDisk, SnapVault, or disk staging storage units can be used as destinations in a lifecycle.

See [“Storage unit or storage destinations”](#) on page 423.

- 6 Indicate the **Volume pool** where the backups (or copies) are to be written.

See [“Volume pool for storage destinations”](#) on page 424.

- 7 Indicate the **Media owner** if the storage unit is a Media Manager type and server groups are configured.

Specify a Media owner to allow only those media servers in the group to write to the media on which backup images for this policy are written.

See [“Media owner for storage destinations”](#) on page 424.

- 8 Select the retention type for the destination:

- **Fixed.**

See [“Fixed retention type for storage destinations”](#) on page 424.

- **Staged capacity managed.**

See [“Staged capacity managed retention type for storage destinations”](#) on page 425.

- **Expire after duplication.**

See [“Expire after duplication retention type for storage destinations”](#) on page 426.

If a policy is configured to back up to a lifecycle, the retention that is indicated in the lifecycle is followed. The **Retention** attribute in the schedule is not followed.

See [“Retention attribute”](#) on page 510.

- 9 Indicate an **Alternate read server** that is allowed to read a backup image originally written by a different media server.

See [“Alternate read server for storage destinations”](#) on page 426.

- 10 Select whether to **Preserve multiplexing**. This option is available for duplication destinations that use tape media.  
See [“Preserve multiplexing for storage destinations”](#) on page 427.
- 11 Click **OK** to create the storage destination.

## Use for: Backup, duplication, or Snapshot destination

Select whether images are to be written to the destination as part of the backup operation or as part of the duplication operation. (During the duplication operation, backups are duplicated to secondary storage.)

If a storage lifecycle contains multiple backup destinations, a multiple copies operation is implied.

See [“Writing multiple copies using a storage lifecycle policy”](#) on page 433.

### Storage destination list requirements

The storage destinations for a storage lifecycle policy must meet the following requirements:

- At least one destination must be a backup destination.  
(Limit: four backup storage destinations per storage lifecycle policy.)
- All backup destinations must be on the same media server.
- One of the destinations must be of a fixed retention type.

If these requirements are not met, an error dialog appears upon saving the storage lifecycle policy.

---

**Note:** A storage lifecycle policy does not need to contain a duplication destination if staging behavior is not the objective. The storage lifecycle policy can be used to create multiple copies.

---

## Storage unit or storage destinations

Indicate the storage unit where the backups are to be written.

You can select the following destinations:

- Any available
- Media Manager storage units (tape)
- Disk storage units (no BasicDisk, SnapVault, or disk staging storage units)

- Storage unit groups (may contain no BasicDisk, SnapVault, or disk staging storage units). A storage lifecycle policy can point to a storage unit group that contains a BasicDisk storage unit. However, NetBackup does not select BasicDisk storage units from a storage group for a lifecycle policy.

---

**Note:** The storage destination list cannot contain other storage lifecycles.

---

Storage units or storage unit groups may appear in more than one lifecycle. Storage units or storage unit groups may be used in a storage lifecycle while also being used as stand-alone units.

## Volume pool for storage destinations

The **Volume pool** option is enabled for tape storage units.

## Media owner for storage destinations

A **Media owner** can be selected. A **Media owner** is a group of NetBackup servers that are used for a common purpose.

## Fixed retention type for storage destinations

A **Fixed** retention type means that the backup data is retained for a specific length of time before the backups are expired. When the retention period is reached, NetBackup deletes information about the expired backup. The files in the backup then become unavailable for restore.

---

**Note:** Every lifecycle must contain at least one destination with a fixed retention period. The destination can be a backup or a duplication destination.

---

The images are not deleted if all copies have not completed. For example, the administrator selects a fixed retention for a tape device to keep images on tape for two days. If the images have not been duplicated to all of the destinations in the lifecycle after two days, the images are not expired. The image remains on the tape device until all copies have been created. (Or, until the administrator uses `nbsstlutil` utility to intervene.)

The **Retention period** specifies exactly how long NetBackup retains the backups before the backups are expired. To set the retention period, select a time period (or level) from the drop-down list.



## Staged capacity managed retention type for storage destinations

A **Staged capacity managed** storage destination means that NetBackup automatically manages the space on the (disk) destination. (This option is not available to tape storage units since tape capacity is considered to be infinite.)

The **High water mark** and **Low water mark** settings on the disk storage unit or disk pool determine how the space is managed.

See [“High water mark setting”](#) on page 382.

See [“Low water mark setting”](#) on page 382.

If space is needed for new images, expired backup copies are removed from the storage destination when the storage unit reaches the high water mark. NetBackup removes backup images on the storage destination until the low water mark is reached.

It searches for images to remove in the following order:

- Any Backup images that have passed the **Desired cache period** setting.
- Data classifications: Images that belong to a data classification with a lower rank are deleted before those that belong to a data classification with a higher rank.

If more space is needed, any images that are not past the **Desired cache period** can be deleted. However, an image is never deleted if it has not been duplicated to all destinations in the lifecycle, even if the image is past its retention period.

See [“Writing multiple copies using a storage lifecycle policy”](#) on page 433.

To see exactly when the storage destination reaches the low water mark value is difficult. A backup can occur at the same time as the expiration process occurs. After the backup is complete, the low water mark may be slightly greater than its lowest possible value.

The retention period for a staged capacity managed storage destination is not assured as it is for a fixed retention period. The **Desired cache period** becomes a target that NetBackup tries to maintain. If the space is not required, the backup data could remain on the storage destination longer than the **Desired cache period** indicates.

Symantec does not recommend allowing capacity-managed images and fixed-retention images to be written to the same volume in a disk storage unit. The volume may fill with fixed-retention images and not allow the space management logic to operate as expected.

Keep the following points in mind when configuring lifecycle destinations or selecting the storage location for a policy:

- All lifecycles that write to a volume in a disk storage unit should write images of the same retention type: fixed or capacity-managed.
- Do not write images both to a volume in a disk storage unit within a lifecycle and to the same volume (by the storage unit) directly from a policy.
- Mark all disk storage units that are used with lifecycles as **On demand only**.
- Check any storage unit groups to make sure that fixed and capacity-managed images cannot be written to the same volume in a disk storage unit.

**Staged capacity managed** is selectable for any disk storage unit that is allowed in a lifecycle. However, for the disk types that support single-instance store (SIS), **Staged capacity managed** functions to various degrees. In order for **Staged capacity managed** to operate, NetBackup must know how much space a backup image uses. With SIS enabled on the storage unit, NetBackup cannot know exactly how much space a particular backup image occupies.

The following storage unit configurations use SIS:

- PureDisk storage units
- NearStore storage units that have either the **Enable file system export** option enabled or the **Enable block sharing** option enabled.
- Some OpenStorage storage units, depending on the vendor characteristics.

## Expire after duplication retention type for storage destinations

The **Expire after duplication** retention can be applied to backup as well as duplication destinations.

When data is staged to a tape device, the following retention options are available:

- Staged capacity managed
- Expire after duplication

NetBackup does not allow the last destination in the lifecycle to use the **Expire after duplication** retention type because no subsequent copy is configured. NetBackup requests that another retention type is chosen for the destination.

## Alternate read server for storage destinations

An **Alternate read server** is available to duplication destinations only. It specifies the name of the server that is allowed to read a backup image originally written by a different media server.

An **Alternate read server** can be specified for both backup and duplication destinations. However, to use an **Alternate Read Server** as part of a duplication

operation, the name of the alternate server must be specified in the backup destination configuration.

## Preserve multiplexing for storage destinations

The **Preserve Multiplexing** option is available for the duplication destinations that use tape media. If the backup to be duplicated is multiplexed and you want the backups to remain multiplexed, check **Preserve Multiplexing**.

To preserve multiplexing significantly improves performance of duplication jobs because it eliminates the need to request the write-side duplication media for every image.

## Hierarchical view of storage destinations

The storage destination list contains all the destinations (storage units and storage unit groups) that the storage lifecycle can use. The list includes the storage that is used for the original backups as well as storage that is used for duplication at a later time.

Figure 14-2 shows how after the first copy is created, all subsequent copies can be made locally from that source, without tying up network resources.

**Figure 14-2** Hierarchical destinations

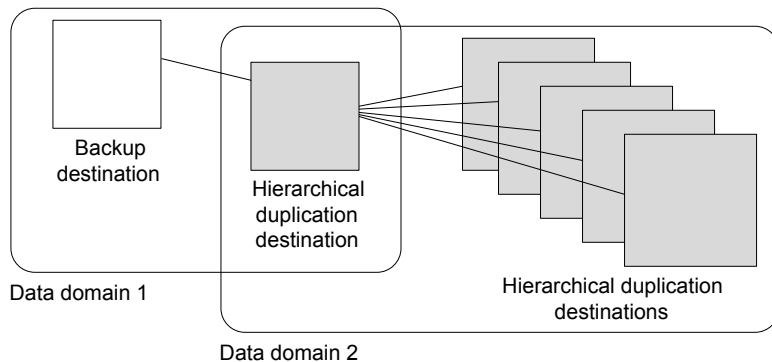
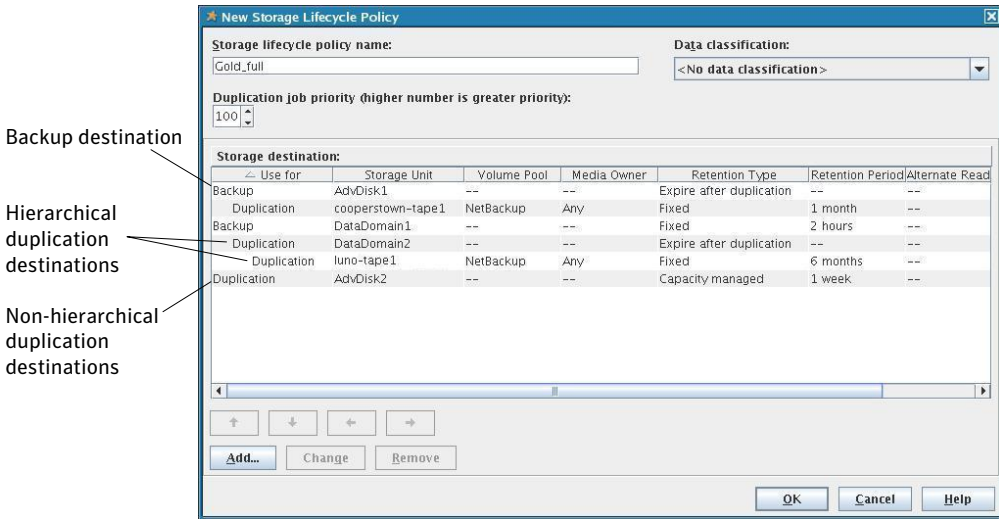


Figure 14-3 shows how the hierarchy of destinations is represented graphically in the **Storage Lifecycle Policy** dialog box. The indentation (or hierarchy) indicates the backup source for each copy. One copy can be the source for many other copies.

**Figure 14-3** Hierarchical storage destinations in a lifecycle policy



Changing the location of a destination in the hierarchy changes the storage unit that serves as the source for the subsequent copies. Changing the location cannot change the destination type. (For example, make a backup destination into a duplication destination.)

The **Maximum backup copies** host property setting in the Global Attributes properties limits the number of destinations that can be added to a lifecycle.

The destination list in the **Storage Lifecycle Policy** dialog box can contain the following destinations at various hierarchical levels:

- Backup**                      A backup destination is never indented in the list or is never a child of another destination.

The first backup destination in the list is generally the primary copy, from which duplication copies are created if the destination is non-hierarchical.
- Snapshot**                      A snapshot destination is never indented and is never a child of another destination.

|                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Duplication (hierarchical)     | <p>A hierarchical duplication destination is a duplication destination that uses a specific source for duplication. It is always indented under a backup destination or another duplication destination. A hierarchical duplication destination can have siblings.</p> <p>The source (or parent) for a hierarchical destination is the destination that appears above the destination in the hierarchy. The source can be a backup or a duplication destination.</p> <p>If a hierarchical duplication destination has children, it serves as the source for the children.</p> |
| Duplication (non-hierarchical) | <p>A non-hierarchical duplication destination is a duplication destination that does not use a specific source for duplication. It is never indented or is never a child of another destination.</p> <p>It can serve as the source for hierarchical duplication destinations.</p> <p>Any backup that is marked as the primary copy can provide the source for a non-hierarchical duplication destination.</p>                                                                                                                                                                 |

## Adding a hierarchical duplication destination

A hierarchical duplication destination means that the destination uses a specific source.

### To add a hierarchical duplication destination

- 1 Select a backup or a duplication storage destination to become the source of the destination to be added.
- 2 Click **Add**.
- 3 Select the **Duplication** type in the **New Storage Destination** dialog box. Complete the remaining fields.
- 4 Click **OK** to add the duplication destination. The hierarchical duplication destination is indented under the selected backup or duplication destination.

## Adding a non-hierarchical duplication destination

A non-hierarchical duplication destination means that the destination does not have a specific backup source. It uses either the primary copy or the best copy.

### To add a non-hierarchical duplication destination

- 1 Make sure that no destination is selected in the **Change Storage Lifecycle policy** dialog.
- 2 Click **Add**.

- 3 Select the **Duplication** type in the **New Storage Destination** dialog box. Complete the remaining fields.
- 4 Click **OK** to add the duplication destination. The duplication destination is added at the end of the destination list without any indentation.

## Modifying the source of a hierarchical duplication destination

Modifying the source of a hierarchical destination does not modify the children of the hierarchical destination.

### To modify the source of a hierarchical duplication destination

- 1 Select the hierarchical duplication destination.
- 2 Click the arrows to move the destination into the new position.

The function of the arrow keys varies depending on the location of the selected destination in the hierarchy:

- Up arrow

Swaps the position of the selected destination with the sibling above it, if one exists.

Using the up arrow does not change the source of the selected destination. The up arrow also moves the children of a destination and preserves their relationship with the selected destination.

The up arrow is disabled if no sibling appears above the selected destination.

- Down arrow

Swaps the position of the selected destination with the sibling below it, if one exists.

Using the down arrow does not change the source of the selected destination. The down arrow also moves the children of a destination and preserves their relationship with the selected destination.

The down is disabled if no sibling appears below the selected destination.

- Right arrow

Moves the destination right in the hierarchy, making the sibling above the destination the source for the duplication destinations.

If no sibling exists above the destination in the hierarchy, the right arrow is disabled. It is always disabled for **Backup** and **Snapshot** destinations.

Moving the destination to the right does not change the position number of the destination in the list.

The right arrow also moves the children of the destination and preserves their relationship with the selected destination.

- Left arrow

Moves the destination to the left in the hierarchy, turning the parent into a sibling.

The left arrow is enabled for duplication destinations. For the left arrow to be enabled the selected duplication destination must be either the first or last in a list of siblings.

If the destination is the first sibling of a parent, click the left arrow to make it into a sibling of its parent.

Note that the left arrow also moves the children along with the selected destination to preserve the relationship with the destination.

The left arrow is disabled for **Backup** and **Snapshot** destinations.

- 3 Click **OK** to save the hierarchy change.

---

**Note:** The order of the destinations at the time the lifecycle is saved may differ from the next time the lifecycle is opened. NetBackup reorders the destinations while it stores them in the catalog configuration file. How the hierarchy works is not changed, however, and the parent-child relationships are preserved.

---

## Removing a destination from the storage destination list

Removing a destination from the storage destination lists can affect the hierarchy. If a destination is removed, and that destination serves as a source for other destinations, those destinations have no source. Without a source, the destinations use the primary backup and the benefits of creating hierarchical destinations are lost.

### To remove a destination from the storage destination list

- 1 Select the destination.
- 2 Click **Remove**. The destination is removed from the list of destinations. The children shift left in the hierarchy.

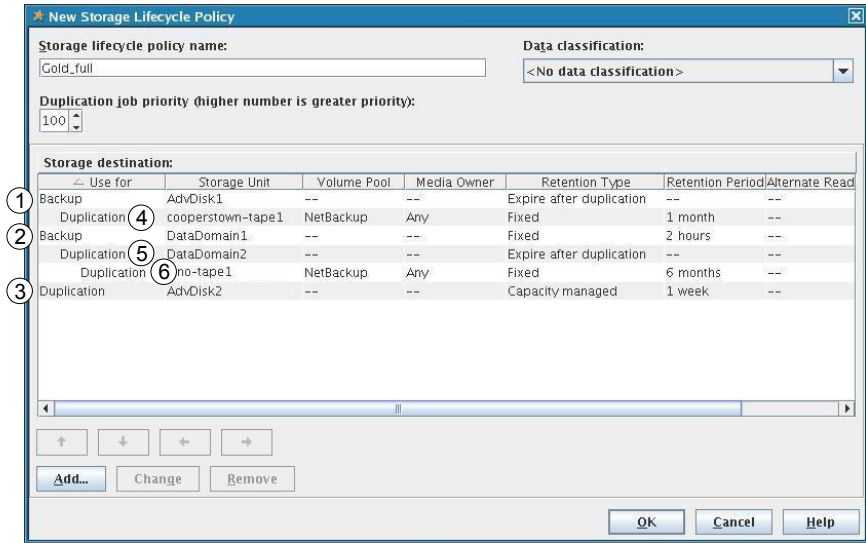
## Hierarchy example

Figure 14-4 shows a configuration with hierarchical and non-hierarchical storage destinations.

Let's examine the following aspects of this example:

- The backup source for each copy.
- The order in which the copies are created.

**Figure 14-4** Copy creation order and hierarchy example



The example shows a storage lifecycle policy that contains six storage destinations. The numbers indicate the order in which the copies are created.

**Table 14-1** Copy creation order and hierarchy example

| Order of creation | Used for                 | Hierarchy                                                                                                                                                                                                                                                                                                                                               |
|-------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1                 | Backup                   | No indentation. The NetBackup Policy Execution Manager ( <code>nbpem</code> ) runs backups to the backup destinations first.                                                                                                                                                                                                                            |
| 2                 | Backup                   | No indentation.                                                                                                                                                                                                                                                                                                                                         |
| 3                 | Duplication              | No indentation; the copy uses either backup destination 1 or 2 as the source, depending on which is marked as the primary copy. Duplications to the duplication destinations that run after the source backup is created. The Duplication Manager ( <code>nbstserv</code> ) runs every 5 minutes (default) and looks for backups eligible to duplicate. |
| 4                 | Hierarchical duplication | Indented under 1; the copy uses backup destination 1 as the source.                                                                                                                                                                                                                                                                                     |
| 5                 | Hierarchical duplication | Indented under 2; the copy uses backup destination 2 as the source.                                                                                                                                                                                                                                                                                     |



**Table 14-1** Copy creation order and hierarchy example (*continued*)

| Order of creation | Used for                 | Hierarchy                                                                |
|-------------------|--------------------------|--------------------------------------------------------------------------|
| 6                 | Hierarchical duplication | Indented under 5; the copy uses duplication destination 5 as the source. |

## Writing multiple copies using a storage lifecycle policy

NetBackup writes backups to all of the destinations in the storage destination list. Therefore, if a storage lifecycle policy contains multiple destinations, a multiple copies operation is implied.

The following topics are considerations when using storage lifecycle policies to create multiple copies.

### Use only one method to create multiple copies

NetBackup permits only one method to create multiple copies to be in use at one time.

To create multiple copies, use one of the following methods:

- Enable the **Multiple copies** option in a policy configuration. If a policy has the **Multiple copies** option enabled, the policy cannot select a storage lifecycle policy as a storage destination.  
 See “[Multiple copies attribute](#)” on page 503.
- Add more than one destination to the storage destination list of a storage lifecycle policy.  
 See “[Adding a storage destination to a storage lifecycle policy](#)” on page 421.

The same criteria for creating copies applies to both methods.

### Destination order determines the copy order

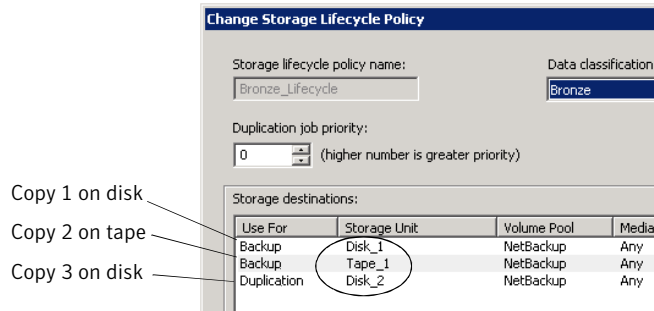
The order in which the destinations appear in a lifecycle determines the copy number of the backup.

For example, in [Figure 14-5](#) a lifecycle is configured to create three copies:

- Two copies to two different backup destinations
- One copy to a duplication destination

To make sure that copy 1 is written to disk, enter the disk type destination before the tape destination.

Figure 14-5 Destination order determines copy order



## Ensuring successful copies using lifecycles

The process to create copies as part of a storage lifecycle policy differs from the process to create copies as set up in a policy. The policy's **Configure Multiple Copies** dialog box includes the option to **Fail all copies**. That option means that if one copy fails, the remaining copies can be set to either continue, or to fail.

In a storage lifecycle policy, all copies must be completed. A lifecycle initially tries three times to create a copy. If no copy is created, NetBackup continues to try, but less frequently.

The successful completion of copies is important because a lifecycle does not allow a copy to be expired before all copies are completed to each destination in the lifecycle. Expiration is necessary to free up space on the storage unit for new backups. NetBackup changes the retention period of an image to Infinite until all copies are created. After all copies are complete, the retention returns to the level as set in the policy that writes to the storage destination.

To complete successful backups in the lifecycle, a backup destination may duplicate a backup onto another backup destination.

Consider the following example: A lifecycle contains two backup destinations (BU\_1, BU\_2) and three duplication destinations. The backup to BU\_1 is successful, but the backup to BU\_2 is not successful. To fulfill the backup on BU\_2, a duplication job is created from BU\_1 to BU\_2. The duplication job is in addition to the jobs that are run for the three duplication destinations.

Duplication jobs can be controlled by using the `nbstlutil` command.

See [“Using the nbstlutil command to administrate lifecycle operations”](#) on page 443.

## Storage lifecycle policy versions

Once a storage lifecycle policy is configured, it runs according to a single configuration or definition. The definition affects both the backups once they begin to run as well as the duplication jobs once the image is in process.

The ability to create storage lifecycle policy versions lets administrators safely modify a definition without waiting until all images that are associated with the storage lifecycle policy have been fully processed.

Each image that a storage lifecycle policy manages is tagged with the storage lifecycle policy name and the storage lifecycle policy version number. These two attributes are written into the image header, in the NetBackup image catalog.

Whenever an administrator creates or changes a storage lifecycle policy, NetBackup creates a new version (between 0 and  $n$ ). New backup jobs use the most recent version.

When a new backup job is submitted to the Activity Monitor, the backup is tagged with the most recent storage lifecycle policy version number. The processing of an image that is associated with a version remains fixed according to that version of the storage lifecycle policy definition. It is fixed at backup time and does not change, unless the administrator uses the `nbstl` command to modify an existing version. Whenever the storage lifecycle policy is modified using the NetBackup Administration Console or `bpadm`, a new version is created.

A storage lifecycle policy version remains as long as there are any incomplete images that refer to the version.

### How to create a new version

Any change that an administrator makes to a storage lifecycle policy using the NetBackup Administration Console or `bpadm` creates a new storage lifecycle policy version. The new version is created when the changes to the storage lifecycle policy are committed or saved. The NetBackup Administration Console and `bpadm` always display the most recent version.

### Administrator actions that do not create a new version

If an administrator uses `nbstl` to make a change to a storage lifecycle policy, `nbstl` creates a new version by default.

However, the `nbstl` command contains options to view different versions and to modify the definitions of existing storage lifecycle policy versions without creating a new version:

- `-all_versions`

Use to display all versions of a storage lifecycle policy definition. Without specifying this option, only the most recent version is displayed by default.

- `-version number`

Use to display a specific version.

- `-modify_current`

Use with most `nbstl` configuration options to make changes to the current storage lifecycle policy version without creating a new version. Knowing the current version number is not necessary if this option is used.

- `-modify_version -version number`

Use with most `nbstl` configuration options to make changes to a specific version without creating a new version.

Use `-modify_current` or `-modify_version` to change any of the following configuration options:

- The duplication priority (`-dp`)
- The storage unit to be used for each destination (`-residence`)
- The volume pool for each destination (`-pool`)
- The server group for each destination (`-server_group`)
- The retention level for each destination (`-rl`)
- The alternate read server for each destination (`-as`)
- The preserve multiplexing option for duplication copies (`-mpx`)

Some fields require values for all of the destinations in the storage lifecycle policy. Make sure that the number of values that are specified for the fields matches the existing destination count.

For example, in a storage lifecycle policy that contains three destinations, to change the value of one, a value must be given for all three destinations. Note that the values for all three destinations are replaced. To change the value for the second destination, provide the existing values for the first and the third destinations.

Some configuration options cannot be changed using `-modify_current` or `-modify_version`. To change any of the following options, you must create an entirely new storage lifecycle policy version:

- The type of the destination: Backup, duplication, snapshot (`-uf`)
- The retention type for the destination: Fixed, capacity managed, expire after duplication (`-managed`)

- The source of a destination, used primarily in hierarchical storage lifecycle policy configurations (`-source`)
- The data classification of an existing version (`-dc`)
- The number of destinations. You cannot add a destination or remove a destination from the storage lifecycle policy definitions.

You cannot instruct a lifecycle to follow the configuration of a previous version that has been superseded. To revert to the behavior of a previous version, change the definition to match the earlier definition. The change creates a version with the same content as the previous version, but with a new version number.

## When do changes to storage lifecycle policies become effective?

For the changes to become effective for a backlog of jobs, it may be necessary to cancel the applicable backup or duplication jobs.

When the `nbstl` command is used to alter an existing version, those changes may not become effective immediately. The images that are managed by the storage lifecycle policy version that was altered may already belong to a duplication job that is Active or Queued, as seen in the Activity Monitor. Once a duplication job is queued, the characteristics (storage lifecycle policy attributes) are fixed for that job and subsequent changes to the definition have no effect. To make changes effective for a backlog of images, cancel the duplication jobs. The storage lifecycle policy manager creates and submits new duplication jobs for those images, using the changes to the configuration.

The following are conditions under which changes to an existing version are not immediately effective:

- Changes to a backup destination have no effect because the backup job is already underway or completed.
- Changes to a duplication destination do not affect the image copies that previous duplication jobs created.
- Changes to a duplication destination do not affect the image copies that have already been submitted and are currently represented by a duplication job in the Activity Monitor, whether it be Active or Queued. If you want your changes to apply to those active duplication jobs, you need to cancel the applicable duplication jobs. Once canceled, `nbstserv` re-forms and re-submits new duplication jobs for these image copies, using the changes to the appropriate version of the storage lifecycle policy.
- Changes to a duplication destination affect the image copies that have not yet been created and have not yet been submitted. That is, they are not yet represented by a duplication job in the Activity Monitor. Your changes become

effective for the next duplication session. Whenever `nbstserv` begins a new session, it re-reads the definitions for processing instructions.

- If a duplication job does not complete successfully, unfinished images in the job are submitted as part of a new job. Changes to the version affect the resubmitted job.

## Deleting old storage lifecycle policy versions

When a version of a storage lifecycle policy is no longer the active (or most recent) version, the version is subject to deletion. NetBackup automatically deletes the inactive storage lifecycle policy version after all the images that refer to it have finished processing. When the images have finished processing, they are considered storage lifecycle policy-complete.

By default, NetBackup deletes an inactive version after 14 days.

The following LIFECYCLE\_PARAMETER entries apply to version deletion:

- See “[CLEANUP\\_SESSION\\_INTERVAL\\_HOURS](#)” on page 438.
- See “[VERSION\\_CLEANUP\\_DELAY\\_HOURS](#)” on page 441.

# LIFECYCLE\_PARAMETERS file for optional duplication job configuration

The NetBackup administrator can customize how the NetBackup Storage Lifecycle Manager (`nbstserv`) runs duplication jobs. The `nbstserv` default values work well for most environments.

To change the values, the administrator must create a file named LIFECYCLE\_PARAMETERS and save it in the following location:

```
install_path\NetBackup\db\config
```

One or all of the parameters can appear in the LIFECYCLE\_PARAMETERS file in any order. If the file does not exist, NetBackup uses the defaults as indicated.

## CLEANUP\_SESSION\_INTERVAL\_HOURS

This parameter concerns the deletion of storage lifecycle policy versions where a more recent version exists.

This parameter controls how often `nbstserv` looks for the versions that have been deleted.

**Syntax:** CLEANUP\_SESSION\_INTERVAL\_HOURS *nn\_hours*

Default: 24 hours.

See [“Deleting old storage lifecycle policy versions”](#) on page 438.

## DUPLICATION\_GROUP\_CRITERIA

This parameter indicates how batches are created. The entry applies to the use of tape and disk.

Syntax: `DUPLICATION_GROUP_CRITERIA 0 | 1`

0 = Select 0 to indicate that batches be created based on the storage lifecycle policy name.

1 = Select 1 to indicate that batches be created based on the duplication job priority from the storage lifecycle policy definition. This setting allows multiple storage lifecycle policies of the same priority together in a job.

Default: 1 (Use the storage lifecycle policy name).

See [“Duplication job priority setting”](#) on page 419.

## DUPLICATION\_SESSION\_INTERVAL\_MINUTES

This parameter indicates how frequently `nbstserv` starts a duplication session. During a duplication session, NetBackup looks for completed backups on backup storage destinations and decides whether or not it is time to start a new duplication job.

Default: 5 minutes. Minimum: 1 minute.

Syntax: `DUPLICATION_SESSION_INTERVAL_MINUTES 5`

## IMAGE\_EXTENDED\_RETRY\_PERIOD\_IN\_HOURS

All copies must be completed in a lifecycle. If necessary, NetBackup initially tries three times to duplicate an image to a duplication destination. The limit prevents NetBackup from retrying too frequently. If, after three tries, the copy is still unsuccessful, this parameter indicates how long NetBackup waits before an image copy is added to the next duplication job. (The `DUPLICATION_SESSION_INTERVAL_MINUTES` parameter determines the frequency.)

The NetBackup administrator may need more than two hours (the default) to solve the problem. Alternatively, the administrator can temporarily de-activate a lifecycle using `nbstlutil`.

Default: 2 hours. Minimum: 1 hour.

Syntax: IMAGE\_EXTENDED\_RETRY\_PERIOD\_IN\_HOURS 2

## MIN\_GB\_SIZE\_PER\_DUPLICATION\_JOB

This parameter indicates the size that the batch of images should reach before one duplication job is run for the entire batch.

The lifecycle does not request a duplication job until either:

- The aggregate size of the images in a batch reaches the minimum size as indicated by MIN\_GB\_SIZE\_PER\_DUPLICATION\_JOB
- The MAX\_MINUTES\_TIL\_FORCE\_SMALL\_DUPLICATION\_JOB time passes. This parameter determines the maximum time between batch requests.

Default: 7 gigabytes.

Syntax: MIN\_GB\_SIZE\_PER\_DUPLICATION\_JOB *GB\_value*

## MAX\_GB\_SIZE\_PER\_DUPLICATION\_JOB

This parameter determines how large the batch of images is allowed to grow.

When the size reaches the size as indicated by this parameter, no additional images are added to the batch.

Default: 25 gigabytes.

Syntax: MAX\_GB\_SIZE\_PER\_DUPLICATION\_JOB *GB\_value*

## MAX\_MINUTES\_TIL\_FORCE\_SMALL\_DUPLICATION\_JOB

This parameter indicates how old any image in the group can become before the batch is submitted as a duplication job. It applies to both disk and tape.

The MAX\_MINUTES\_TIL\_FORCE\_SMALL\_DUPLICATION\_JOB entry working differently in this release than it did in previous releases.

A very small batch is not submitted to `nbstserv` until one source job in the batch has finished at least 30 minutes ago.

---

**Note:** The timer does not come into effect if the total size of all the images in the batch exceeds the parameter value. Or, if all of the source media that the duplication job requires is full.

---

This parameter helps to ensure a balance between the following conditions:

- Submitting many small duplication jobs too soon or too frequently.



On the one hand, `nbstserv` doesn't want to submit a small job if there's additional work available to make the job bigger and more efficient.

- Waiting too long before submitting a small job.

On the other hand, `nbstserv` should not wait too long to submit a small job.

Default: 30 minutes.

Syntax: `MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB 30`

## TAPE\_RESOURCE\_MULTIPLIER

This parameter indicates a value which serves as the multiplier for the number of concurrently active duplication jobs that can access a single storage unit. This parameter applies to tape media.

Storage unit configuration includes limiting the number of jobs that can access the resource at one time. (The **Maximum concurrent write drives** value.) This value specifies the optimal number of jobs that the Resource Broker can consider running on that resource.

This parameter helps administrators ensure a balance in the following situation:

- To overload the Resource Broker with jobs that it can't run is not prudent.
- Make sure that there's enough work that is queued so that the devices won't become idle. The `TAPE_RESOURCE_MULTIPLIER` entry lets administrators tune the amount of work that the Resource Broker can evaluate for a particular storage unit.

For example, a particular storage unit contains three write drives. If the `TAPE_RESOURCE_MULTIPLIER` parameter is set to two, then the limit on concurrently active jobs is six. Other duplication jobs requiring the storage unit remain queued.

Default: 2

Syntax: `TAPE_RESOURCE_MULTIPLIER n`

## VERSION\_CLEANUP\_DELAY\_HOURS

This parameter concerns the deletion of storage lifecycle policy versions where a more recent version exists.

This parameter controls how much time must pass since an inactive version was the active version. If the version has been inactive for at least as long as the `VERSION_CLEANUP_DELAY_HOURS` value, NetBackup considers it for deletion.

Syntax: `VERSION_CLEANUP_DELAY_HOURS nn_hours`

Default: 336 hours (14 days).

See [“Deleting old storage lifecycle policy versions”](#) on page 438.

## LIFECYCLE\_PARAMETERS file example

The following is an example of the contents and syntax for a LIFECYCLE\_PARAMETERS file using the default values:

```
DUPLICATION_SESSION_INTERVAL_MINUTES 5
IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS 2
MIN_GB_SIZE_PER_DUPLICATION_JOB 7
MAX_GB_SIZE_PER_DUPLICATION_JOB 25
MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB 30
```

## Logic for batch creation

The Storage Lifecycle Manager service (`nbstserv`) is in charge of creating duplication jobs for storage lifecycle policies. Part of duplication job creation includes grouping the backup (or source) jobs into batches.

One objective of the batching logic is to prevent media contention for tape operations (including VTL).

Batching logic applies to both disk and tape. (Though the method to prevent media contention for disk is to use disk pools and then to limit I/O streams to disk pools.)

The batching logic requires that for each evaluation cycle, `nbstserv` consider all completed source jobs when determining which duplication job to run next. By default, `nbstserv` performs the evaluation once every 5 minutes.

`nbstserv` avoids overloading the Resource Broker (`nbrb`) queue with jobs. Too many jobs in the queue make the role of the Resource Broker harder and slows down system performance.

By default, `nbstserv` now creates groups based on the **Duplication job priority** setting of each storage lifecycle policy. Multiple storage lifecycle policies with the same priority can be batched together. Even if a NetBackup environment does not use the **Duplication job priority** setting, it benefits from allowing multiple storage lifecycle policies in one duplication job.

See [“Duplication job priority setting”](#) on page 419.

This batching logic change affects how duplication jobs appear in the Activity Monitor. Storage lifecycle policies that have been combined into one job appear under a single policy name: `SLP_MultipleLifecycles`. If a storage lifecycle policy has not been combined with another, the name appears in the Activity Monitor under the name of the SLP: `SLP_name`.

Users may see some duplication jobs that, although in the running state, are not duplicating data because they have no resources to read or write. These jobs continue to run until they receive resources to complete the job.

To turn off grouping by duplication job priority, change the `DUPLICATION_GROUP_CRITERIA` entry, a `LIFECYCLE_PARAMETER`.

See “[DUPLICATION\\_GROUP\\_CRITERIA](#)” on page 439.

## Using the `nbstlutil` command to administrate lifecycle operations

The NetBackup storage lifecycle utility command (`nbstlutil`) gives administrators the ability to intervene between pending storage lifecycle operations. Specifically, the `nbstlutil` command can be used to cancel, inactivate, or activate the processing of existing lifecycle-managed images.

`nbstlutil` cannot affect the jobs that are currently running or queued. Use the Activity Monitor to intervene in the jobs that are running or queued.

The command is found in the following location:

```
install_path\NetBackup\bin\admincmd\nbstlutil
```

Use `nbstlutil` to perform the following administrative actions:

- List the status of lifecycle-managed images. The EMM table that tracks the status of lifecycle-processed images can be printed. Support may request this information to troubleshoot a lifecycle problem.
- Cancel pending duplication operations on the selected images or image copies. When a duplication is canceled, NetBackup considers the image or image copy to be lifecycle complete and does not attempt to create any more copies of the backup image.
- Inactivate (suspend) pending and future lifecycle operations on selected images or image copies. NetBackup retains the image information so that processing can be resumed by the administrator at a later time.
- Activate (resume) suspended lifecycle operations on selected images or image copies.

See *NetBackup Commands for Windows* for a description of all the options available for `nbstlutil`.

## When to use `nbstlutil`

NetBackup starts a duplication session every five minutes to copy data from a backup destination to a duplication destination. (Five minutes, or the frequency as designated by the `DUPLICATION_SESSION_INTERVAL_MINUTES` parameter.)

If the copy fails, the next three duplication sessions retry the copy. If the copy fails all three times, the copy is retried every two hours until it succeeds. (Two hours, or the frequency as designated by the `IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS` parameter.)

Use the `nbstlutil` command in the case of a hardware problem that may require more than 15 minutes to resolve. That is, the problem may take longer to resolve than three duplication sessions five minutes apart.

For example, a duplication job fails because the library has a hard failure. It may take longer than two hours to repair the library. The administrator may not want duplication jobs to begin every two hours. Use the `nbstlutil` command to inactivate the lifecycle while the library is repaired. When ready, the lifecycle can be activated and duplication jobs can begin.

---

**Note:** Once reactivated, the administrator may want to temporarily change the `IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS` parameter to one hour to begin duplication jobs sooner.

---

# Configuring backups

- [Chapter 15. Creating backup policies](#)
- [Chapter 16. Synthetic backups](#)
- [Chapter 17. Protecting the NetBackup catalog](#)
- [Chapter 18. About the NetBackup relational database](#)
- [Chapter 19. Using the Catalog utility](#)



# Creating backup policies

This chapter includes the following topics:

- [Using the Policies utility](#)
- [Planning for policies](#)
- [Creating a policy using the Backup Policy Configuration Wizard](#)
- [Creating a policy without using the Backup Policy Configuration Wizard](#)
- [Changing policies](#)
- [About the Policy attributes](#)
- [About the Schedules tab](#)
- [About the Schedule Attributes tab](#)
- [Using the Start Windows tab](#)
- [Using the Exclude Dates tab](#)
- [Using the Calendar Schedule tab](#)
- [Considerations for user schedules](#)
- [Backup window considerations](#)
- [About the Clients tab](#)
- [About the Backup Selections tab](#)
- [About the Disaster Recovery tab](#)
- [Creating a Vault policy](#)
- [Performing manual backups](#)

- [Active Directory granular backups and recovery](#)

## Using the Policies utility

Policies define the rules that NetBackup follows when clients are backed up. A backup policy can apply to one or more clients. Every client must be in at least one backup policy. The best approach to configure backup policies is to divide clients into groups according to the backup requirements and archive requirements. Then, create a policy for each group.

The left pane contains a hierarchical view of the policies on the master server currently under management. The Details pane is a detailed view that displays the information that pertains to the policy that is selected.

To display information about all policies on the current master server, click **Summary of All Policies**. A summary of all policies appears in the **Details** pane, subdivided into panes that display **Policies**, **Schedules**, **Clients**, and **Selections**.

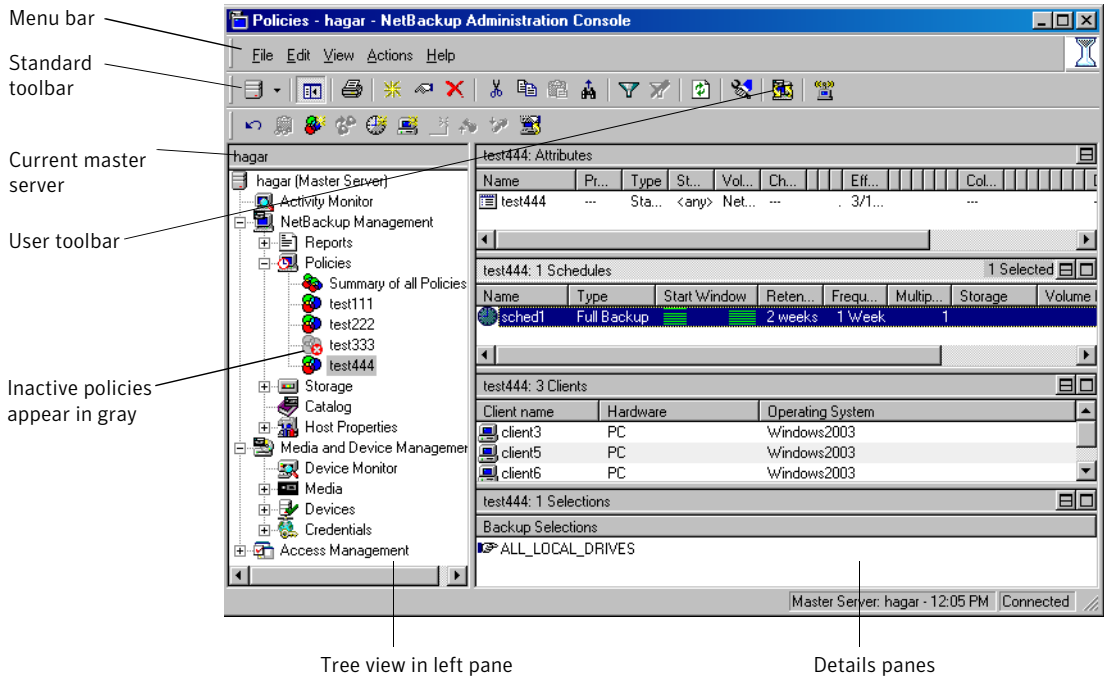
The options on the **Policies** menu bar are described in the online Help .

To display the general attributes for a specific policy, select that policy in the left pane. The Details pane shows the general attributes for that policy only.

Double-click on a policy to display the attributes in tabs, available for editing.



**Figure 15-1** NetBackup Policy utility



## Planning for policies

Policies can meet the needs of a wide variety of clients in a single NetBackup configuration. To take advantage of the capabilities of policies, plan before you start to configure the policies. The following topic provides guidelines for planning.

### Group the clients

Divide clients into groups according to the types of work the clients perform.

Clients that are used for similar tasks generally have much in common regarding the backup requirements. For example, most clients in an engineering department create the same types of files at similar levels of importance.

In some instances, you can create a single policy for each group of clients. In other cases, you need to subdivide the clients and include them in the separate policies that are based on their backup requirements.

In this example, assume that the clients are in the same work group. Initially, try to cover all of the clients in the same backup policy.

## Gather client information

Gather information about each client that is relevant to the backups. For example, the names of the clients and the number of files and typical file sizes on each client.

In this example, client1 is a file server that contains a large amount of data. To avoid long backup times, include client1 in one policy and the workstations in another policy. Later, we may find that we need more than one policy for client1.

## Consider storage requirements

Create backup policies to accommodate special storage requirements.

The storage unit and volume pool settings apply to all files that are backed up by the policy. If files have special storage requirements, create separate policies for the files, even if other factors are the same, such as schedules.

In [Table 15-1](#), separate policies are used for

D:\h002\DevExp

and

D:\h002\DesDoc

on client1. Those files are sent to DLT tape. Other files on client1 go on 8mm tape. If it is necessary to keep backups for some files on separate media, create a policy that specifies a unique volume pool for those backups. Then, add the media for that volume pool.

**Table 15-1** Separate policies for a single client

| Policy | Clients            | Files                                         | Storage |
|--------|--------------------|-----------------------------------------------|---------|
| S1     | client1            | C:\<br>D:\User<br>D:\h001<br>E:\h002\Projects | 8mm     |
| S2     | client1<br>client1 | E:\h002\DevExp<br>E:\h002\DesDoc              | DLT     |

## Backup schedule considerations

Create additional backup policies if one set of schedules does not accommodate all clients and files.

Consider the following factors to create schedules in a policy:

- **Best times for backups to occur.**  
 To back up different clients on different schedules, create more policies. For example, create different policies for night-shift and day-shift clients. In our example, we can back up all clients during the same hours so additional policies are not necessary.
- **How frequently the files change.**  
 If some files change infrequently compared to other files, back up the files on a different schedule. To use a different schedule, create another policy that has an appropriate schedule and then include the files and clients in that policy.  
 In [Table 15-2](#), `C:\` on client1 is in a different policy (S3). The `C:\` drive on the workstations is also in a separate policy (WS2).
- **How long backups need to be kept.**  
 Each schedule has a retention setting that determines how long NetBackup keeps the files that are backed up by the schedule. Because the schedule backs up all the files in the backup selection list, all files should have similar retention requirements. Do not include the files whose full backups must be retained forever, together in a policy where full backups are retained for only four weeks.  
 In [Table 15-2](#), `E:\h002\DesDoc` on client1 is in a different policy (S4). `E:\h002\DesDoc` requires full backups every 12 weeks and those backups must be retained much longer than the other files on client1.

**Table 15-2** Example policies

| Policy | Clients | File selections                        | Frequency of change | Storage | Frequency of automatic backups                                                                                                                     |
|--------|---------|----------------------------------------|---------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| S1     | client1 | D:\User<br>D:\h001<br>E:\h002\Projects | high                | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every 4 weeks</li> </ul> |

**Table 15-2** Example policies (*continued*)

| Policy | Clients                    | File selections                   | Frequency of change | Storage | Frequency of automatic backups                                                                                                                                                            |
|--------|----------------------------|-----------------------------------|---------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S2     | client1                    | E:\h002\DevExp                    | high                | DLT     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every 4 weeks</li> </ul>                                        |
| S3     | client1                    | C:\                               | low                 | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Full backups every 4 weeks</li> </ul>                                                                       |
| S4     | client1                    | E:\h002\DesDoc                    | high                | DLT     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every 4 weeks</li> <li>■ Full backups every 12 weeks</li> </ul> |
| WS1    | mars                       | D:\User<br>D:\Programs            | high                | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every 4 weeks</li> </ul>                                        |
|        | jupiter                    | D:\User<br>D:\Programs            |                     |         |                                                                                                                                                                                           |
|        | neptune                    | D:\User<br>D:\Programs<br>D:\Util |                     |         |                                                                                                                                                                                           |
| WS2    | mars<br>jupiter<br>neptune | C:\<br>C:\<br>C:\                 | low                 | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Full backups every 4 weeks</li> </ul>                                                                       |

## How to group by general attributes

Create separate policies for the clients that require different general attribute settings than other clients.

General attributes include the following:

|                              |                                                                                                                                                                                                                                                                                                    |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Policy Type</b>           | You must use the correct policy type for each client.<br>See <a href="#">“Policy type attribute”</a> on page 462.                                                                                                                                                                                  |
| <b>Follow NFS</b>            | Select this attribute if a UNIX client has NFS mounted files to be backed up. Consider placing these clients in a separate policy so problems with NFS do not affect the other clients.<br>See <a href="#">“Follow NFS attribute”</a> on page 476.                                                 |
| <b>Backup Network Drives</b> | Select this attribute to back up the files that the client stores on network drives (applies only to the MS-Windows policy type).<br>See <a href="#">“Backup network drives attribute”</a> on page 474.                                                                                            |
| <b>Compression</b>           | Set this attribute if you want a client to compress its backups before sending them to the server. Note that the time to compress can increase backup time and make it unsuitable to use for all clients.<br>See <a href="#">“Compression attribute”</a> on page 480.                              |
| <b>Job Priority</b>          | Use this attribute to control the order in which NetBackup starts the backups. The clients with the higher priority are backed up first.<br>See <a href="#">“Job priority attribute”</a> on page 473.<br><br>In the example, no extra policies are required because of general attribute settings. |

## Maximize multiplexed backups

Create separate policies as necessary to maximize the benefits of multiplexed backups.

To maximize drive use, multiplex the slower clients that produce small backups. The higher-performance clients that produce long backups are likely to use drives fully and not benefit from multiplexing.

## Evaluate backup times

Evaluate total backup times for each schedule and further subdivide policies to reduce backup times to an acceptable level.

For example, if the backup of

D:\User, D:\h001, and E:\h002\Projects

on client1 takes too much time, create a new policy for E:\h002\Projects.

Policy S5 has the same requirements as S1. Backup E:\h002\Projects separately to reduce backup time.

In addition to reducing the backup time for each policy, separate policies can reduce the total backup time for the server client1. NetBackup processes files within a backup selection list serially, in the order they appear in the backup selection list. However, separate policies are processed in parallel if enough drives are available and the maximum jobs attributes are set to allow it.

The **Multiplexing** and **Allow Multiple Data Streams** options also allow backup policies to be processed in parallel.

See [“Allow multiple data streams attribute”](#) on page 486.

**Table 15-3** Example policies

| Policy | Clients | File selections    | Frequency of change | Storage | Frequency of automatic Backups                                                                                                                                   |
|--------|---------|--------------------|---------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S1     | client1 | D:\User<br>D:\h001 | high                | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Cumulative incremental backups</li> <li>■ Full backups every four weeks</li> </ul> |
| S2     | client1 | E:\h002\DevExp     | high                | DLT     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Cumulative incremental backups</li> <li>■ Full backups every four weeks</li> </ul> |
| S3     | client1 | C:\                | low                 | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Cumulative incremental backups</li> </ul>                                          |

**Table 15-3** Example policies (*continued*)

| Policy | Clients                    | File selections                   | Frequency of change | Storage | Frequency of automatic Backups                                                                                                                                                               |
|--------|----------------------------|-----------------------------------|---------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S4     | client1                    | E:\h002\DesDoc                    | high                | DLT     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every four weeks</li> <li>■ Full backups every 12 weeks</li> </ul> |
| S5     | client1                    | D:\h002\Projects                  | high                | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every four weeks</li> </ul>                                        |
| WS1    | mars                       | D:\User<br>D:\Programs            | high                | 8mm     | <ul style="list-style-type: none"> <li>■ Daily incremental backups</li> <li>■ Weekly full backups</li> <li>■ Full backups every four weeks</li> </ul>                                        |
|        | jupiter                    | D:\User<br>D:\Programs            |                     |         |                                                                                                                                                                                              |
|        | neptune                    | D:\User<br>D:\Programs<br>D:\Util |                     |         |                                                                                                                                                                                              |
| WS2    | mars<br>jupiter<br>neptune | C:\<br>C:\<br>C:\                 | low                 | 8mm     | <ul style="list-style-type: none"> <li>■ Daily Incr</li> <li>■ 4 Weeks Full</li> </ul>                                                                                                       |

## Creating a policy using the Backup Policy Configuration Wizard

The easiest method to set up a backup policy is to use the Backup Policy Configuration Wizard. This wizard guides you through the setup process, which simplifies the process as it automatically chooses the best values for most configurations.

The Backup Policy Configuration Wizard cannot be used to configure a calendar-based schedule. You can change the schedule to a calendar-based schedule after running the wizard.

See [“Using the Calendar Schedule tab”](#) on page 521.

Use the following procedure to create a policy using the Backup Policy Configuration Wizard.

#### To create a policy with the Backup Policy Configuration Wizard

- 1 In the NetBackup Administration Console, select **Master Server** or **NetBackup Management**.
- 2 From the list of wizards in the **Details** pane, click **Create a Backup Policy**.  
Click **Help** on any wizard screen for assistance while running the wizard.

## Creating a policy without using the Backup Policy Configuration Wizard

Use the following procedure to create a policy without using the Backup Policy Configuration Wizard.

#### To create a policy without the Backup Policy Configuration Wizard

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select **Actions > New > New Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.  
See [“NetBackup naming conventions”](#) on page 719.
- 4 Click **OK**.

## Changing policies

Change policies only when no backup activity is expected for the affected policies and clients. Make adjustments before backups begin to ensure an orderly transition from one configuration to another.

Changing a policy causes NetBackup to recalculate when the policy is due.

See [“Events that cause the schedules to be recalculated”](#) on page 527.

See [“Moving policy information from one server to another”](#) on page 458.

See [“Changing multiple policies at one time”](#) on page 459.



See [“Adding or changing schedules in a policy”](#) on page 457.

See [“Adding or changing clients in a policy”](#) on page 457.

See [“Adding or changing backup selections in a policy”](#) on page 458.

See [“Cutting, copying, and pasting policy items”](#) on page 459.

See [“Using the composite change dialog box”](#) on page 460.

See [“Deleting schedules, backup selections, or clients from a policy”](#) on page 461.

## Adding or changing schedules in a policy

### To add or change schedules in a policy

- 1 Expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 Perform one of the following actions:
  - To add a schedule, select **Actions > New > New Schedule**.
  - To change an existing schedule, double-click the schedule name in the **Details** pane.
- 4 Complete the entries in the **Attributes** tab, **Start Window** tab, **Exclude Dates** tab, and **Calendar Schedule** tab (if it appears).  
See [“About the Schedules tab”](#) on page 490.
- 5 Click **OK**.
- 6 If this schedule is the last schedule, click **OK**. To add more schedules, click **Add** and repeat the previous step.

## Adding or changing clients in a policy

Use the following procedure to add or change clients in a policy.

### To add or change clients in a policy

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 Perform one of the following actions:
  - To add a new client, select **Actions > New > New Client**.

- To change an existing client, double-click the client name in the **Details** pane.
- 4 Complete the entries in the **Add New Client or Change Client** dialog box.  
See [“To add a client to a policy”](#) on page 535.

## Adding or changing backup selections in a policy

Use the following procedure to add or change backup selections in a policy.

### To add or change backup selections in a policy

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 Perform one of the following actions:
  - To add a new backup selection, select **Actions > New > New Backup Selection**.
  - To change an existing backup selection, double-click the backup selection in the **Details** pane.
- 4 Complete the entries in the **New Backup Selections** or **Change Backup Selections** dialog box.
- 5 Add the new backup selection and click **Add**, or make changes to an existing selection.
- 6 Click **OK**.

## Moving policy information from one server to another

Use the following procedure to move policy information from one server to another.

### To move policy information from one server to another

- 1 Copy the policy information to the clipboard. Copy the entire policy or part of a policy.
- 2 Change to the destination server.
- 3 Paste the policy information in the **Policies** node. For the configuration to work, you must complete the rest of the configuration on the destination server. (For example, select a storage unit and volume pool at the destination server.)

## Cutting, copying, and pasting policy items

You can copy or cut and paste the following items:

- Copy and paste (not cut) attributes
- Copy and paste (not cut) entire policies
- Copy, cut, and paste schedules, backup selections, and clients

Use the following procedure to cut, copy, and paste various policy items.

### To cut, copy, and paste items (general procedure)

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**. In the left pane, select the policy from which you want to copy or move items.
- 2 In the **Details** pane, select the item (backup selections, clients, or schedules).
- 3 Select **Edit > Copy** or **Edit > Cut**. Both the **Copy** and **Cut** commands copy the selected items to the clipboard.
- 4 In the left pane, select the policy where you want to paste or move the items.
- 5 In the **Details** pane, click the pane where you want to paste the contents of the clipboard: attributes, clients, schedules, or backup selections.

To view the contents of the clipboard, select **Edit > Clipboard**.

Any items with the same name are replaced with the contents of the clipboard after pasting. If there are any schedules that do not match the policy type, the schedules are deleted or renamed. The action is indicated in a dialog box.

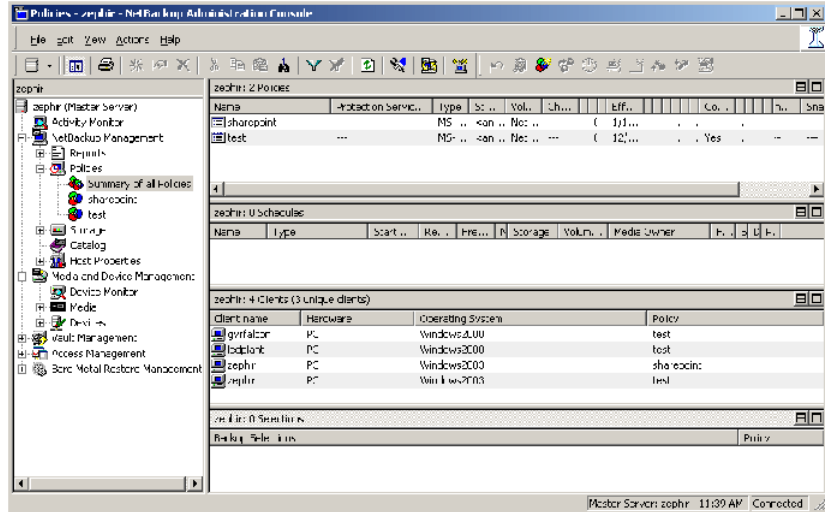
- 6 On the **Edit** menu, click **Paste**.

## Changing multiple policies at one time

Use the following procedure to change more than one policy at one time.

### To change multiple policies

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies > Summary of All Policies**.



- 2 In the **Details** pane, hold down the Shift key and select the items you want to change.
- 3 Select **Edit > Change**.

See “Using the composite change dialog box” on page 460.

### Using the composite change dialog box

When multiple policies are selected, be sure that you know how to interpret the settings before making changes to them at the same time:

- If the property is a text field in which to specify a value:
  - The text field displays a value if the property has the same value for all selected machines.
  - The text field is clear if the property does not have the same value for all selected machines.

Any change is applied to the field for every selected policy.
- If the property is a check box and not a value, the check box displays in one of the following states:
  - Checked, if the attribute is specified for all selected machines.

- Clear, if the property is clear on all selected machines.
- Gray check, if the property is set differently on the selected machines.

To change the selection, do one of the following:

- Select the check box to set the property on all selected machines.
- Clear the check box to clear the property on all selected machines.
- Set (or leave) the box to a gray check to leave the property unchanged.

At any time, click **Cancel** to cancel changes or click **OK** to apply all changes and close the dialog box.

## Deleting schedules, backup selections, or clients from a policy

Use the following procedure to delete, backup selections, or clients from a policy.

### To delete schedules, backup selections, or clients from a policy

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 In the **Details** pane, select the item you want to delete.
- 4 Select **Edit > Delete**.
- 5 Click **Yes**.

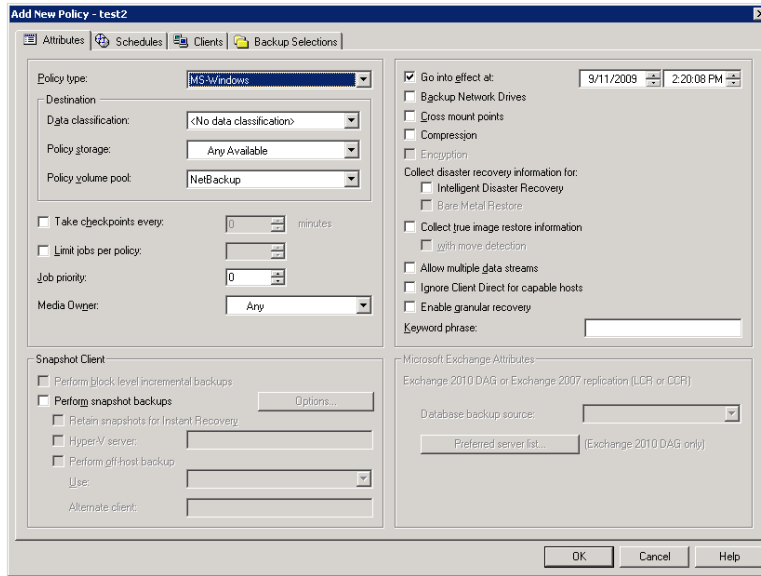
When a client is deleted from the client list, the NetBackup client software is not deleted or uninstalled from the client. Backups for the client can be recovered until the backups expire. Also, when a file is deleted from a backup selection list, the actual file is not deleted from the client.

## About the Policy attributes

The settings on the **Attributes** tab determine the characteristics of all the backups that NetBackup performs according to the selected policy.

[Figure 15-2](#) shows the Attributes tab for a policy.

Figure 15-2 Policy Attributes tab



Policy attributes are configurable depending on the type of policy and the options that are installed.

The following topics describe the settings on the **Attributes** tab.

## Policy type attribute

The **Policy type** attribute determines the purpose of the policy. Usually, the **Policy type** attribute determines the type of clients that can be backed up by this policy. Not all policy types serve to back up clients. (NBU-Catalog, for example.) Select the type of policy from the drop-down list.

You can change the policy type of an existing policy. However, the existing schedules may become invalid for the new policy type. If the schedules become invalid, NetBackup displays an alert notice. NetBackup then either deletes the invalid schedules or changes the schedules to an equivalent type.

[Table 15-4](#) lists and describes that different Policy type attributes.

**Table 15-4** Policy types

| Policy type                                   | Description                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>AFS</b><br>(UNIX only)                     | Use for the policies that back up only AFS file systems on clients.<br><br>For information on setting up these policies, see "Using NetBackup with AFS," in the <i>NetBackup Administrator's Guide, Volume II</i> .                                                                                                                           |
| <b>DataTools-SQL-BackTrack</b><br>(UNIX only) | Use for the policies that contain only clients with the NetBackup <b>SQL-BackTrack</b> extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                              |
| <b>DataStore</b>                              | A policy type that is reserved for use by Symantec or its partners to provide agents for new applications or databases.                                                                                                                                                                                                                       |
| <b>DB2</b>                                    | Use for the policies that contain only clients with the NetBackup DB2 extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                                               |
| <b>FlashBackup</b><br>(UNIX only)             | Applies only to NetBackup Enterprise Server:<br><br>Use for the policies that contain only NetBackup FlashBackup clients on UNIX. This policy is available only when the <b>Snapshot Client</b> is installed.<br><br>For information on setting up this policy type, see the <i>Snapshot Client Administrator's Guide</i> .                   |
| <b>FlashBackup- Windows</b><br>(UNIX only)    | Applies only to NetBackup Enterprise Server:<br><br>Use for the policies that contain only <b>FlashBackup-Windows</b> clients on Windows. This policy is available only when the <b>Snapshot Client</b> is installed.<br><br>For information on setting up this policy type, see the <i>NetBackup Snapshot Client Administrator's Guide</i> . |
| <b>Generic</b><br>(Windows only)              | A policy type that does not require the same pre-processing steps that backup policies require. For example, a LiveUpdate policy is considered a generic policy.                                                                                                                                                                              |
| <b>Informix-On-BAR</b><br>(UNIX only)         | Use for the policies that contain only clients with the NetBackup Informix extension . For information on setting up this policy type, see the guide for this option.                                                                                                                                                                         |
| <b>Lotus-Notes</b>                            | Use for the policies that contain only clients with the NetBackup Lotus Notes extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                                       |
| <b>MS-Exchange-Server</b>                     | Use for the policies that contain only clients with the NetBackup MS Exchange extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                                       |
| <b>MS-SharePoint</b><br>(Windows only)        | Use to configure a policy for NetBackup for SharePoint Portal Server.                                                                                                                                                                                                                                                                         |
| <b>MS-SQL-Server</b>                          | Use for the policies that contain only clients with the NetBackup MS SQL Server extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                                     |

**Table 15-4** Policy types (*continued*)

| Policy type            | Description                                                                                                                                                                                                                                                                                                              |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>MS-Windows*</b>     | Use for the policies that contain only supported Windows OS levels.                                                                                                                                                                                                                                                      |
| <b>NBU-Catalog</b>     | Use for hot catalog backup jobs. Allows for a catalog backup while other jobs are running.                                                                                                                                                                                                                               |
| <b>NCR-Teradata</b>    | Use for the policies that contain only clients with the NetBackup for Teradata option. For information on setting up this policy type, see the guide for this option.                                                                                                                                                    |
| <b>NDMP</b>            | Use for the policies that contain only clients with the NetBackup NDMP option. This policy type is available only when the NetBackup NDMP is installed and licensed. For information on setting up this policy type, see the guide for this option.                                                                      |
| <b>NetWare</b>         | Use for the policies that contain only NonTarget NetBackup Novell NetWare clients (this version uses a Microsoft Windows interface).                                                                                                                                                                                     |
| <b>Oracle</b>          | Use for the policies that contain only clients with the NetBackup Oracle extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                       |
| <b>OS/2</b>            | Use for the policies that contain only OS/2 clients.                                                                                                                                                                                                                                                                     |
| <b>PureDisk-Export</b> | Use for the policies that export data from PureDisk to NetBackup.                                                                                                                                                                                                                                                        |
| <b>SAP</b>             | Use for the policies that contain only clients with the NetBackup SAP extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                          |
| <b>Standard*</b>       | Use for the policies that contain any combination of the following: <ul style="list-style-type: none"> <li>■ UNIX clients (including supported Mac clients), except those covered by specific such as Oracle.</li> <li>■ NetBackup Novell NetWare clients that have the target version of NetBackup software.</li> </ul> |
| <b>Sybase</b>          | Use for the policies that contain only clients with the NetBackup Sybase extension. For information on setting up this policy type, see the guide for this option.                                                                                                                                                       |
| <b>Vault</b>           | Use as a policy type to schedule and run a Vault job. Available only when Vault is licensed.                                                                                                                                                                                                                             |

---

**Note:** Use the **Standard** or **MS-Windows** policy type to implement the following options: **CheckPoint Restart** for backups, **Checkpoint Restart** for restores, synthetic backups, or the **Collect disaster recovery information for Bare Metal Restore** option.

---

For more details on off-host backups, refer to the *NetBackup Snapshot Client Administrator's Guide*.



## Data classifications attribute

Select a **Data Classification** if you want the backup to go to a storage unit that stores backups of a particular classification. For example, a gold backup must go to a storage unit with a gold data classification.

A **Data Classification** selection is optional. If no classification is indicated, the policy uses the storage units and groups that the **Policy storage** attribute indicates.

If a data classification is selected, all the images that the policy creates are tagged with the classification ID.

See “[Data classification option](#)” on page 417.

## Policy storage attribute

The **Policy storage** attribute specifies the storage destination for the policy’s data. Select the storage destination from the drop-down list. The drop-down list may contain storage units, storage lifecycle policies, and storage unit groups.

If the selection is configured to do so, the storage unit or lifecycle policy determines which type of disk staging is used for this policy.

See “[About staging backups](#)” on page 395.

The list includes only those lifecycles that are of the same data classification as the policy. For example, gold backup images cannot be sent to a silver storage lifecycle.

Images that belong to a specific data classification cannot be sent to a storage lifecycle that does not have a classification. To select a data classification is optional.

If the **Any Available** option is selected, NetBackup tries to store data on locally-attached storage units first. To force NetBackup to use only a locally-attached drive, select **Must use local drive** in the **General Server** properties.

See “[General Server properties](#)” on page 127.

If a local device is not found or if the **Must use local drive** attribute is not selected, NetBackup tries to find an available storage unit alphabetically.

If the **Any Available** option is selected, NetBackup uses the first storage unit that meets the following requirements:

- The storage unit must not be designated as **On Demand Only**.
- The storage unit must have available drives.
- The storage unit must have media available in the required volume pool.

An exception is the case in which a client is also a media server with locally-attached storage units. The local storage units take precedence over the sequence that is based on alphabetical order.

The storage unit that is selected on the **Schedule** tab, overrides the **Policy storage** attribute.

See [“Override policy storage selection attribute”](#) on page 508.

See [“About selecting a storage destination”](#) on page 466.

## About selecting a storage destination

Consider the following points before selecting a storage destination:

- If the site contains only one storage unit or if there is no storage unit preference
  - Specify **Any Available** for the **Policy storage** attribute.
  - Do not specify a storage unit at the schedule level.  
See [“Override policy storage selection attribute”](#) on page 508.
  - In this situation, do not configure all storage units to be **On Demand Only**. NetBackup may be unable to find an available storage unit for the backups.  
See [“On demand only setting”](#) on page 387.
- If a specific storage unit is designated and the unit is unavailable, backups cannot run for those policies and the schedules that require the unit.
- If **Any Available** is selected, any Basic disk storage unit that is not assigned to a storage group is considered available for disk spanning.  
See [“Media properties”](#) on page 153.
- To limit the storage units available to a policy, select a storage unit group that contains only the units you want the policy to use.

Another method to limit the storage units available to a policy is the following:

- Create a volume pool that contains the volumes that are available only to the specific storage units.  
See [“Adding a volume pool”](#) on page 293.  
If the **Scratch pool** option is enabled for the volume pool, any storage unit has access to the volumes in the volume pool. For this method to work, do not enable **Scratch pool** for the volume pool.  
See [“About scratch volume pools”](#) on page 292.
- In the policy, set **Policy volume pool** to the volume pool that is defined in the previous step.
- For all policies, set **Policy storage** attribute to **Any Available**.

- A policy may specify a storage unit group. Make sure that one of the storage units within the group is set to **On Demand Only** to satisfy the policy requirement.  
 See “[On demand only setting](#)” on page 387.

## Policy volume pool attribute

The **Policy volume pool** attribute specifies the default volume pool where the backups for the policy are stored. A volume pool is a group of media that is grouped together for use by a single application. The volume pool is protected from access by other applications and users.

Select a volume pool name from the list of volume pools. Whenever a new volume is required, it is allocated from the volume pool indicated.

The volume pool that is selected on the **Schedule** tab overrides the **Policy volume pool** setting.

See “[Override policy volume pool attribute](#)” on page 509.

[Table 15-5](#) describes the default NetBackup volume pools.

**Table 15-5** Default volume pools

| Volume pool name | Volume pool description                                                                                                                                                                                                                                             |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| None             | The default pool for applications, other than NetBackup.                                                                                                                                                                                                            |
| DataStore        | The default pool for DataStore.                                                                                                                                                                                                                                     |
| NetBackup        | Unless otherwise specified in the policy, all backups use media from the NetBackup pool. One exception is the NBU-Catalog policy type. The NBU-Catalog policy type is used for online, hot catalog backups, which selects the CatalogBackup volume pool by default. |
| CatalogBackup    | This pool is selected by default for the NBU-Catalog policy type. It is used exclusively for online, hot catalog backups. Online, hot catalogs are directed to a single, dedicated pool to facilitate faster catalog restores.                                      |

The following are some additional volume pools that can be useful to create:

- A Scratch pool from which NetBackup can automatically transfer volumes when another volume pool does not have media available.
- An Auto volume pool, for use by automatic backups.
- A User volume pool, for use by user backups.

Media is assigned to the volume pools for Media Manager storage devices. Disk-type storage devices are not allocated to a volume pool.

See [“About volume pools”](#) on page 292.

See [“Adding a volume pool”](#) on page 293.

## Volume pool override example

Assume that you want all schedules but one to use the **Backups** pool. The exception is a user-archive schedule that requires the **Archive** pool.

In the policy, set **Policy volume pool** to **Backups**.

When you set up the schedules for the policy, set **Override policy volume pool** as follows:

- For the schedules that use the **Backups** volume pool, clear **Override policy volume pool**.
- For the schedule that requires the **Archive** volume pool, select **Override policy volume pool** and specify **Archive** for the pool name.

## Checkpoint restart for backup jobs

The **Take checkpoints every** check box specifies whether NetBackup takes checkpoints during a backup job. Indicate how often the policy should take checkpoints.

Checkpoints during a backup are beneficial if a backup fails. Without **Take checkpoints every** enabled, a failed backup restarts from the beginning of the job. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the beginning of the last checkpoint rather than restart the entire job. Note that checkpoints cannot occur while a file is backed up. Checkpoints are saved at file boundaries and point to the next file in the list.

The **Schedule backup attempts** Global Attributes host property indicates the number of times that NetBackup tries a failed backup.

See [“Global Attributes properties”](#) on page 131.

Policy types MS-Windows (for Windows clients) and Standard (for UNIX clients) support this policy attribute. To see if this feature is supported for a specific agent or option, refer to the manual for that agent or option.

---

**Note:** Checkpoints are not taken for a user archive schedule. If the user archive is resumed, it restarts from the beginning.

---

## Checkpoint frequency

The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. (Default: 15 minutes.) The administrator determines checkpoint frequency on a policy-by-policy basis. Balance more frequent checkpoints with the likelihood of time that is lost when a backup is resumed. If the frequency of checkpoints affects performance, increase the time between checkpoints.

## Decision to start a new job or resume an incomplete job

NetBackup decides when a new job should be started instead of resuming an incomplete job.

NetBackup starts a new job in the following situations:

- If a new job is due to run.
- If the time since the last incomplete backup was longer than the shortest frequency in any schedule for the policy.
- If the time indicated by the Clean-up property, **Move backup job from incomplete state to done state**, has passed.
- For calendar scheduling, if another run day has arrived.

## Checkpoint restart support

The following topics pertain to the use of checkpoint restart with different products and options.

### Windows clients and checkpoint restart

The following items pertain to the use of checkpoint restart with Windows clients:

- Checkpoint restart is not supported for backup selections indicated by a UNC path.
- No checkpoints are taken during a System State backup.
- No checkpoints are taken during a Windows Disk Image (raw) backup.
- No checkpoints are taken for the remainder of the backup after NetBackup encounters Single-instance Store (SIS).
- When an incremental backup resumes and completes successfully, the archive bits are cleared for the files that were backed up after the job resumed. However, the archive bits are not cleared for the files that were backed up before the resume. Since the archive bits remain, the files that were backed up before the resume are backed up again during the next incremental backup.

### Multiple copies and checkpoint restart

Checkpoint restart is supported for the policies that are configured to create multiple backup copies.

See “ [Multiple copies attribute](#)” on page 503.

The last failed copy that contains a checkpoint can be resumed if all of the following items are true:

- A copy is configured to allow other copies to continue the job if the copy fails and subsequent checkpoints occur.
- The **Take checkpoints every** option is selected for this policy.

### Synthetic backups and checkpoint restart

Checkpoint restart is not supported for use with synthetic backups in the current NetBackup release.

### Snapshot Client and checkpoint restart

Checkpoint restart is supported for use with local or alternate client backups. However, the following methods are not supported: Block Level Incremental Backups, Media Server Copy, Third-Party Copy Device, and Instant Recovery backups.

### Basic disk staging and checkpoint restart

Checkpoint restart is supported for use in Stage I of basic disk staging, during which data is backed up to disk.

Checkpoint restart is unavailable in the Stage II storage unit policy of basic disk staging, during which data is relocated to another storage unit.

See “ [About staging backups](#)” on page 395.

### NearStore storage units and checkpoint restart

NearStore storage units do not support checkpoint restart.

### NetWare clients and checkpoint restart

Although NetWare clients can use the Standard policy type, checkpoint restart for backups is not supported on NetWare clients.

### Checkpoint restart for restore jobs

**Checkpoint Restart** for restore jobs saves time by providing the mechanism for NetBackup to resume a failed restore job. The job resumes automatically from the start of the file last checkpointed rather than from the beginning of the entire

restore job. NetBackup automatically takes checkpoints once every minute during a restore job.

The following two host properties affect Checkpoint Restart for restore jobs:

- Master server host property **Clean-up > Move Restore Job from Incomplete State to Done State**  
See “[Clean-up properties](#)” on page 76.
- Master server host property **Universal > Restore Retries**  
See “[Universal Settings properties](#)” on page 185.

Limitations to the Checkpoint Restart option for restore jobs include the following:

- The restore restarts at the beginning of the last checkpointed file only, not within the file.
- Checkpoint Restart for restore jobs works only on the files that are backed up by using Standard or MS-Windows policy types.
- Third Party Copy and the Media Server Copy images that use Standard policy types are supported. However, they cannot be suspended or resumed if the backup image has changed blocks. Flashbackup is not supported.

### Suspending and resuming a restore job

A NetBackup administrator can choose to suspend a checkpointed restore job and resume the job at a later time.

For example, while an administrator runs a restore job for several hours, the administrator receives a request for a second restore. The request is of a higher priority and requires the resources in use by the first job. The administrator can suspend the first job, start the second restore job and let it complete. The administrator can then resume the first job from the Activity Monitor and let the job complete.

---

**Note:** If a checkpointed restore that has no end date is suspended, then resumed, and a new backup occurs before the resume is initiated, the files from the new backup are included in the restore. For example, a user makes a restore request of a directory. Then that restore is suspended. The request is resumed the next day after another backup of the directory has been performed. The files that are restored are from the latest backup.

---


## Limit jobs per policy attribute

The **Limit jobs per policy** attribute limits the number of jobs that NetBackup performs concurrently when the policy is run. By default, the check box is clear,

and NetBackup performs an unlimited number of backup jobs concurrently. Other resource settings can limit the number of jobs.

A configuration can contain enough devices so that the number of concurrent backups affects performance. To specify a lower limit, select the check box and specify a value from 1 to 999.

**Figure 15-3** Limit jobs per policy attribute



Client A and Client B backups can occur concurrently and back up to different devices

The number of concurrent backup jobs that NetBackup can perform depends on the following:

- **Limit jobs per policy** does not prevent concurrent jobs if the jobs are from different policies.  
For example, if three policies limit concurrent jobs to two, NetBackup can start two jobs from each policy. A total of six policies can be running at one time in this situation.
- Parent jobs do not count toward the limit.  
Only the children jobs count toward the **Limit jobs per policy** setting. The following jobs produce a parent job and children jobs: multistreamed jobs, catalog backups, Snapshot Client snapshots, or Bare Metal Restore jobs. See [“About the Jobs tab”](#) on page 682.
- The number of storage devices available and multiplexing limits.  
To process more than one backup job at a time, the configuration must include one of the following:
  - Multiple storage units.
  - A storage unit with enough drives to perform more than one backup at a time.
  - Storage units that are configured to multiplex.With removable media devices such as tape drives, the number of concurrent jobs depends on the total number of drives in the storage units. With disk



storage, the storage device is defined as a file path and the available disk space determines how many paths are possible.

- **The server speed.**  
 Too many concurrent backups interfere with the performance of the server. The best number depends on the hardware, operating system, and applications that are running.
- **The network load.**  
 The available bandwidth of the network determines how many backups can occur concurrently. If you encounter loading problems, consider multiple networks for backups. Or, configure the backup policy to use Compression. See [“Compression attribute”](#) on page 480.  
 When the client that is backed up is also a server, it is a special case. In this instance, the network load is not a factor because the network is not used. However, the load on the client and server is still a factor.
- **Multiplexing.** If multiplexing is used, set **Limit jobs per policy** high enough to support the specified level of multiplexing.  
 Lower values can limit multiplexing within a policy if jobs from different schedules exist within the policy. For example, **Limit jobs per policy** is set to two and an incremental backup schedule is due to run for four clients. Only two clients are backed up at one time, regardless of the multiplexing settings.

## Job priority attribute

The **Job priority** attribute specifies the priority that a policy has as it competes with other policies for resources. A higher priority means that NetBackup assigns the first available resource to the policy with the highest priority.

To set the priority, enter a number in the **Job priority** field. The value can range from 0 to 99999. The higher the number, the greater the priority of the job.

To set a job priority default for a job type, see the **Default Job Priorities** host properties.

See [“Default Job Priorities properties”](#) on page 105.

## Media owner attribute

Active only for Media Manager type storage units or if the **Policy storage** attribute is **Any Available**.

The **Media Owner** property specifies which media server or server group should own the media on which backup images for this policy are written.

You can specify the following for **Media Owner**:

|                |                                                                                                                                                                                                                             |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any (default)  | Allows NetBackup to choose the media owner. NetBackup chooses a media server or a server group (if one is configured).                                                                                                      |
| None           | Specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.                                                   |
| A server group | Specify a server group. A server group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured appear in the drop-down list. |

See [“About media sharing”](#) on page 297.

See [“Configuring a server group”](#) on page 198.

## Go into effect at attribute

To activate the policy, select the **Go into effect at** check box. The policy must be active for NetBackup to use the policy.

The **Go into effect at** attribute specifies when the policy can begin to schedule backups. For example, if today is Monday and you enter Wednesday at 12:00 AM, the policy does not run until that time or later. Use the **Go into effect at** attribute to configure a series of policies in advance of when the policies need to become active.

To deactivate a policy, clear the check box. Inactive policies appear are unavailable in the Administration Console. To resume backups, recheck the **Go into effect at** check box. Make sure that the **Go into effect at** date and time is set to the time that you want to resume backups.

If the schedule is to be used for a catalog archive, the policy must not be active. The **Go into effect at** check box must be clear. For more information about how to configure a policy to archive the catalog,

See [“Creating a catalog archiving policy”](#) on page 618.

## Backup network drives attribute

The **Backup network drives** attribute is for use on single user systems, Win95, Win98, and ME. These operating systems are not supported with this version of NetBackup. The preferred method for backing up data from a computer that is not a NetBackup client is to use UNC paths. UNC paths are more precise and indicate exactly what should be backed up.

When **Backup network drives** or UNC paths are used, the network drives must be available to the service account that the NetBackup Client service logs into at

startup. By default, the startup account is set to System. You must change this account on each Windows client that is backed up that contains data that is shared from another computer.

**Backup network drives** must be checked when policies back up to CD ROM drives. For scheduled backups, the file list must indicate at least the first level of folders to be backed up. For example, `D:\Folder1` instead of only `D:\`

---

**Note:** Mapped drive letters cannot be backed up. Drive letters do not appear in the Backup, Archive, and Restore console when backups are browsed.

---

## Setup example with UNC paths

In the following example, assume the following:

- `master1` is the NetBackup master server.
- `win_client` is a Windows NetBackup client.
- `win_PC` is a Windows computer (not necessarily a NetBackup client) and contains a shared folder named `TestData`.

The following procedure backs up the folder `TestData` on `win_PC` through `win_client`.

### To back up the example folder

- 1 On the NetBackup master server, `master1`, create a policy for `win_client`.
- 2 Add `\\win_PC\TestData` to the file list of the policy. This step is not necessary if the policy is only used for user-directed backups.
- 3 On `win_client`, the NetBackup client:
  - Change the NetBackup Client Service on `win_client` to **Start Up or Log On** with the same account as the user that performs the backup. This user account must have read permissions for the share that is to be backed up. The account must have write permission to perform restores.
  - Stop and start the **NetBackup Client Service** so the new account takes effect.
- 4 To run a user backup, expand the **Network** node in the Backup, Archive, and Restore client interface to `win_PC`. Select `TestData`.

Backups run as scheduled or when a manual backup is performed.

## Example using the Backup network drives attribute

In the following example, assume the following:

- `master1` is the NetBackup master server.
- `win_client` is a Windows NetBackup client.
- `win_PC` is a Windows computer (not necessarily a NetBackup client) and has a shared folder that is named `share`.

The following procedure backs up the folder `share` on `win_PC` through `win_client`.

The following procedure backs up the folder `TestData` on `win_PC` through `win_client`.

#### To back up the example folder

- 1 On the NetBackup master server, `master1`, select **Backup network drives** in the policy to be used for the backup.
- 2 On `win_client`, the NetBackup client:
  - Change the **NetBackup Client Service** on `win_client` to **Start Up** or **Log On** with the same account as the user that performs the backup. This user account must have read permissions for the share that is to be backed up. The account must have write permission to perform restores.
  - Stop and start the **NetBackup Client Service** so the new account takes effect.
  - Create a `bpstart_notify.bat` file that maps a drive on `win_client` to `\\win_PC\share`.  
Enter the command:  

```
net use X: \\win_PC\share
```

  
Where `x:` is the mapped drive letter.
- 3 To run a user backup, expand the **Network** node in the Backup, Archive, and Restore client interface to `win_PC`. Select **TestData**.
- 4 Scheduled backups run as scheduled or when a manual backup is performed.

## Follow NFS attribute

The **Follow NFS** attribute specifies whether NetBackup is to back up or archive any NFS-mounted files that are named in the backup selection list. Or, by the user in the case of a user backup or archive. Clear the check box to prevent the backup or archive of NFS mounted files.

---

**Note:** The **Follow NFS** attribute applies only to UNIX clients in certain policy types. NetBackup allows it to be selected in those instances only.

---

The following are notes on the **Follow NFS** attribute:

The **Follow NFS** setting eliminates the need to locate and log on to the systems where the files reside. If the files are mounted on the NetBackup client, you can back up, archive, and restore them by working from the NetBackup client. You must have the necessary permissions on the NFS mount. Use this capability to back up the systems that the NetBackup client software does not support.

Generally, do not back up NetBackup clients over NFS. Back up and archive files on the NFS server where the files physically reside. NFS backups have lower performance and sometimes encounter problems.

If **Follow NFS** is selected, you may want to use the policy only for the files and clients that are backed up or archived over NFS.

---

**Note:** If **Follow NFS** is not selected, the backup process reads the client's mount table and evaluates each item in the table. NetBackup resolves any links to the true path. NetBackup must resolve the links so it can accurately avoid backing up any files that reside on NFS-mounted file systems.

---

If NetBackup cannot access a Network File System when it evaluates the mount table, it assumes that the file system is unavailable. (The default time to access the file system is five seconds.)

To change the default, change the UNIX master server host property, **NFS access timeout**.

See "[UNIX Server properties](#)" on page 190.

- The behavior of the **Follow NFS** attribute depends on the **Cross mount points** setting.  
See "[Cross mount points attribute](#)" on page 478.
- The **Follow NFS** option has no effect on raw partitions. The Network File Systems that are mounted in a raw partition are not backed up. Nor can you back up raw partitions from other machines that use NFS mounts to access the raw partitions. The devices are not accessible on other machines through NFS.

---

**Note:** NetBackup does not support raw partition backups on unformatted partitions.

---

- The **Follow NFS** option causes files in automounted file systems to be backed up. Automounted directories can be excluded to allow the backup of other NFS

mounts. To do so, add an entry for the automounter's mount directory to the exclude list on the client.

## Cross mount points attribute

The **Cross Mount Points** attribute controls whether NetBackup crosses file system boundaries during a Windows 2003 or later or UNIX backup or archive.

Consider the following items when setting the **Cross mount points** attribute:

- Enable **Cross mount points** to back up all files and directories in the selected path, regardless of the file system.  
For example, if root (/) is specified as the file path on a UNIX system, NetBackup backs up root (/) and all files and directories under root in the tree.  
NetBackup specifically excludes mapped directories even if **Follow NFS** and **Cross mount points** are enabled. To back up mapped directories, include the directories in the file list.
- The following entries have the same effect on Windows systems:  
/  
:\n  
\*:\n  
ALL\_LOCAL\_DRIVES
- Disable **Cross mount points** to back up only the files that are in the same file system as the selected file path. Prohibit NetBackup from crossing mount points to back up root (/) without backing up all the file systems that are mounted on root. (For example, /usr and /home.)
- **Cross Mount Points** has no effect on UNIX raw partitions. If a raw partition is the root partition and contains mount points for other file systems, the other file systems are not backed up, even if you select **Cross Mount Points**.
- On UNIX systems only, do not use **Cross Mount Points** in policies where you use the ALL\_LOCAL\_DRIVES directive in the backup selection list.
- Do not cross mount points to back up a media server that uses mount points to any disk storage that contains backup images. If the policy crosses mount points, the NetBackup backup images that reside on that disk storage are backed up. The NetBackup **BasicDisk** disk type and the Enterprise Disk Option disk types use mount points for disk storage.

## Creating separate policies with cross mount points disabled

In some cases, consider creating separate policies for the backups that cross mount points and the backups that do not cross mount points.

For example, in one policy, **Cross mount points** is not enabled. The backup selection list contains only `root (/)`. Only the root file system is backed up, and not the file systems that are mounted on it.

In another policy, enable **Cross mount points**. In the backup selection list, include `root (/)` to back up all the data on a client.

## Cross mount points and the Follow NFS attributes

To back up NFS-mounted files, select **Follow NFS**.

[Table 15-6](#) summarizes the behavior of **Cross mount points** and **Follow NFS**:

**Table 15-6** Cross mount point behavior


| Cross mount points | Follow NFS | Result                                                                                                                                           |
|--------------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Disabled           | Disabled   | No crossing of mount points (default).                                                                                                           |
| Disabled           | Enabled    | Back up NFS files if the file path is (or is part of) an NFS mount.                                                                              |
| Enabled            | Disabled   | Cross local mount points but not NFS mounts.                                                                                                     |
| Enabled            | Enabled    | Follow the specified path across mount points to back up files and directories (including NFS), regardless of the file system where they reside. |

**Note:** NetBackup specifically excludes mapped directories even if **Follow NFS** and **Cross mount points** are enabled. To back up mapped directories, include the directories in the file list.

## Cross mount point examples

The following examples assume that the client disks are partitioned as shown in [Figure 15-4](#).

**Figure 15-4** Cross mount point example



In **Figure 15-4**, the client contains `/`, `/usr`, and `/home` in separate partitions on disk `d1`. Another file system that is named `/home/njr` exists on disk `d2` and is mounted on `/home`. In addition, disk `d3` contains a directory named `/net/freddie/home` that is NFS-mounted on `/net/freddie`.

■ **Example 1**

Assume that **Cross mount points** and **Follow NFS** are not selected. Assume that the backup selection list contains the following entries:

```
//usr/home
```

NetBackup considers only the directories and files that are in the same file system as the backup selection list entry it is processing. It does not back up `/home/njr` **OR** `/net/freddie/home`.

■ **Example 2**

Assume that **Cross mount points** and **Follow NFS** are selected. Assume that the backup selection list contains only `/`.

In this case, NetBackup backs up all the files and directories in the tree, including those under `/home/njr` and `/net/freddie/home`.

To back up only `/usr` and individual files under `/`, leave `/` out of the list and separately list the files and directories you want to include. For example:

```
/usr
/individual_files_under_root
```

## Compression attribute

The **Compression** attribute specifies that the backups use the software compression that is based on the policy. Select the check box to enable compression. (Default: no compression.)



The degree to which a file can be compressed depends on the data type. A backup usually involves more than one type of data. Examples include stripped and unstripped binaries, ASCII, and the non-unique strings that repeat. Some data types are more favorable to compression.

---

**Note:** When compression is not used, the server may receive more data than what exists on the client. The discrepancy is due to client disk fragmentation and the file headers that the client adds. (To tell how much space a file occupies, run the `du` command. To tell how much free disk space is available, run the `df` command.)

---

|                                       |                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Data types that compress well:        | Programs, ASCII files, and unstripped binaries (typically 40% of the original size).                                                                                                                                                                                                                                                                                                                                          |
| Best-case compression:                | Files that are composed of the strings that repeat can sometimes be compressed to 1% of their original size.                                                                                                                                                                                                                                                                                                                  |
| Data types that do not compress well: | Stripped binaries (usually 60% of original size).                                                                                                                                                                                                                                                                                                                                                                             |
| Worst-case compression:               | Files that are already compressed become slightly larger if compressed again. On UNIX clients, if a compressed file has a unique file extension, exclude it from compression by adding it under the Client Settings (UNIX) properties.                                                                                                                                                                                        |
| Effect of file size:                  | File size has no effect on the amount of compression. However, it takes longer to compress many small files than a single large one.                                                                                                                                                                                                                                                                                          |
| Client resources that are required:   | Compression requires client computer processing unit time and as much memory as the administrator configures.                                                                                                                                                                                                                                                                                                                 |
| Effect on client speed:               | Compression uses as much of the computer processing unit as available and affects other applications that require the computer processing unit. For fast CPUs, however, I/O rather than CPU speed is the limiting factor.                                                                                                                                                                                                     |
| Files that are not compressed:        | <p>NetBackup does not compress the following files:</p> <ul style="list-style-type: none"> <li>■ Files that are equal to or less than 512 bytes, because that is the tar block size.</li> <li>■ On UNIX clients, the files that end with suffixes specified with the <code>COMPRESS_SUFFIX =.suffix</code> option in the <code>bp.conf</code> file.</li> <li>■ On UNIX clients, files with the following suffixes:</li> </ul> |

|                   |                       |                       |                       |
|-------------------|-----------------------|-----------------------|-----------------------|
| <code>.arc</code> | <code>.gz</code>      | <code>.iff</code>     | <code>.sit.bin</code> |
| <code>.arj</code> | <code>.hqx</code>     | <code>.pit</code>     | <code>.tiff</code>    |
| <code>.au</code>  | <code>.hqx.bin</code> | <code>.pit.bin</code> | <code>.Y</code>       |

|          |       |          |      |
|----------|-------|----------|------|
| .cpt     | .jpeg | .scf     | .zip |
| .cpt.bin | .jpg  | .sea     | .zom |
| .F       | .lha  | .sea.bin | .zoo |
| .F3B     | .lzh  | .sit     | .z   |
| .gif     | .pak  |          |      |

Compression increases the overhead computing on the client and increases backup time due to the time required to compress the files. The lower transfer rate that is associated with compression on the client reduces the ability of some tape devices (notably 8mm) to stream data. The effect of the lower transfer rate causes additional wear on those devices.

The savings in media and network resources continue to make compression desirable unless total backup time or client computing resources become a problem. If total backup time is a problem, consider multiplexing. The NetBackup multiplexing feature backs up clients in parallel, reducing the total time to back them up.

If compressed data is written to a storage unit that has single-instance store (SIS) capabilities, the storage unit may not be able to use data deduplication on the compressed or the encrypted data. In data deduplication, only one instance of the file is stored. Subsequent instances of the file reference the single file.

Compression reduces the size of a backup by reducing the size of files in the backup. In turn, the smaller backup size decreases the amount of media that is required for storage. Compression also decreases the amount of data that travels over the network as well as the network load.

## Encryption attribute

The **Encryption** attribute determines whether the backup should be encrypted. When the server initiates the backup, it passes on the **Encryption** policy attribute to the client in the backup request.

The client compares the **Encryption** policy attribute to the **Encryption** host properties for the client. If the encryption permissions for the client are set to REQUIRED or ALLOWED, the policy can encrypt the backups for that client.

See “[Encryption properties](#)” on page 108.

For additional encryption configuration information, see the *NetBackup Security and Encryption Guide*.

---

**Note:** If encrypted data is written to a storage unit that has single-instance store capabilities, the storage unit may not be able to use data deduplication on the compressed or the encrypted data.

---

## Collect disaster recovery information for Intelligent Disaster Recovery attribute

The **Collect disaster recovery information for Intelligent Disaster Recovery** attribute specifies whether NetBackup collects the information that IDR requires to recover Windows clients.

For more information, see "Configuring NetBackup Policies for IDR" in the *NetBackup Security and Encryption Guide*.

## Collect disaster recovery information for Bare Metal Restore attribute

The **Collect disaster recovery information for Bare Metal Restore** attribute specifies whether the BMR client agent runs on each client. If the attribute is enabled, the BMR client agent runs before each backup to save the configuration information of the client. The Activity Monitor displays the activity as a job separate from the backup.

Bare Metal Restore is a separately-priced option.

For more information, see the *Bare Metal Restore Administrator's Guide for UNIX, Windows, and Linux*.

Only policy types **MS-Windows** (for Windows clients) and **Standard** (for UNIX clients) support this policy attribute. This attribute is enabled by default when an **MS-Windows** or **Standard** policy is created on the master servers that are licensed for BMR.

## Collect true image restore information attribute

The **Collect true image restore information** attribute specifies whether the policy collects the information necessary to perform a true image restore. That is, to restore the directory contents to reflect what the directories had contained at the time of an incremental or a full backup. Files that were deleted before the backup are not restored.

With the attribute enabled, a restore based on an incremental backup includes all files that were backed up since the last full backup. The restore also includes those files that were deleted at any time during that period.

NetBackup starts to collect the true image restore information with the next full or incremental backup for the policy. The true image restore information is collected for each client regardless of whether any files were changed.

NetBackup does not provide true image restores based on the time of a user backup or archive. However, NetBackup uses a user backup for a true image restore if the backup is more recent than the latest automatic full or incremental backup.

To include the files that were moved, renamed, or newly installed in the directories, enable **With move detection** for the true image incremental backups.

The following options require that **Collect true image restore information with move detection** be enabled:

- It must be enabled to create synthetic backups.  
See “[Synthetic backup attribute](#)” on page 499.
- It must be enabled to back up data to the NearStore disk storage units that use the File System Export option.  
For more information about configuring NearStore storage units, see the *NetBackup Administrator's Guide, Volume II*.

## Collect true image restore information with move detection attribute

The **Collect true image restore information with move detection** attribute specifies what true image incremental backups should include. Enable this attribute to include the files that were moved, renamed, or newly installed from a tar or a zip archive. (Depending on how the files were packaged and how they were installed, some newly installed files are not backed up by non-TIR incremental backups.)

Without move detection, NetBackup skips the files and directories that were moved, renamed, or newly installed because their modification times are unchanged. With move detection, NetBackup compares path names and inode numbers with those from the previous full or incremental backup. If a name or an inode number is new or changed, the file or directory is backed up.

---

**Note:** This attribute must be selected to create synthetic backups.

---

The following examples show how move detection backs up the files that otherwise would not be backed up:

- A file that is named `C:\pub\doc` is moved to or installed in `C:\spec\doc`. Here, the archive bit is unchanged but `C:\spec\doc` is new in the `C:\spec\` directory and is backed up.
- A directory that is named `C:\security\dev\` is renamed as `C:\security\devices\`. Here, the archive bit is unchanged but `C:\security\devices\` is a new directory and is backed up.

NetBackup begins to collect the information that is required for move detection with the next full or incremental backup for the policy. This first backup after the attribute is set, always backs up all files, even if it is an incremental backup.

Move detection consumes space on the client and the backup can fail if there is not enough disk space available.

### About true image restores

Table 15-7 lists the files that are backed up in the C:\home\abc\doc directory during a series of backups between 12/01/2009 and 12/04/2009. Assume that **Collect true image restore information** was turned on for the policy that performed the backups.

**Table 15-7** True image restore backup example

| Day        | Type of backup | Files that are backed up in C:\user\doc            |
|------------|----------------|----------------------------------------------------|
| 12/01/2009 | Full           | file1 file2 dirA/fileA dirB/fileB file3            |
| 12/02/2009 | Incremental    | file1 file2 dirA/fileA ----- ----                  |
| 12/03/2009 | Incremental    | file1 file2 dirA/fileA ----- ----                  |
| 12/04/2009 | User backup    | file1 file2 dirA/fileA ----- ---- dirC/fileC file4 |
| 12/04/2009 | Incremental    | file1 file2 ----- ----- ---- ----- file4           |

---

**Note:** Dashes ( ----- ) indicate that the file was deleted before this backup.

---

Assume that the 12/04/2009 version of the C:\user\doc directory needs to be restored.

- After a regular restore, the restored directory contains all files and directories that ever existed in C:\user\doc from 12/01/2009 (last full backup) through 12/04/2009:

```
file1
file2
dirA\fileA
dirB\fileB
file3
dirC\fileC
file4
```

- A true image restore of the 12/04/2009 backup creates a directory that contains only the files and directories that existed at the time of the incremental backup:

```
file1
file2
file4
```

NetBackup does not restore any of the files that were deleted before the 12/04/2009 incremental backup.

The restored directory does not include the subdirectories `dirA` and `dirC`, even though they were backed up on 12/04/2009 with a user backup.

NetBackup did not restore these directories because they did not exist at the time of the incremental backup. The incremental backup was the reference for the true image restore.

## Allow multiple data streams attribute

The **Allow multiple data streams** attribute specifies that NetBackup can divide automatic backups for each client into multiple jobs. The directives, scripts, or templates in the backup selection list specify whether each job can back up only a part of the backup selection list. Since the jobs are in separate data streams, they can occur concurrently.

The directives, scripts, or templates in the backup selection list determine the number of streams (backup jobs) that start for each client. The list also determines how the backup selection list is divided into separate streams.

The following settings determine the number of streams that can run concurrently:

- Number of available storage units
- Multiplexing settings
- Maximum jobs parameters

Multistreamed jobs consist of a parent job to perform stream discovery and children jobs for each stream. In the Activity Monitor, the children jobs display the Job ID of the parent job. Parent jobs display a dash (-) in the Schedule column.

---

**Note:** If this attribute is enabled, and a file system is in a client's exclude list, a NetBackup job appears in the Activity Monitor for the excluded file system. However, no files in the excluded file system are backed up by the job.

---

See [“Backup selections list directives for multiple data streams”](#) on page 563.

See [“When to use multiple data streams”](#) on page 487.

See “[When to use multiple data streams](#)” on page 487.

## When to use multiple data streams

The following items describe the reasons to use multiple data streams:

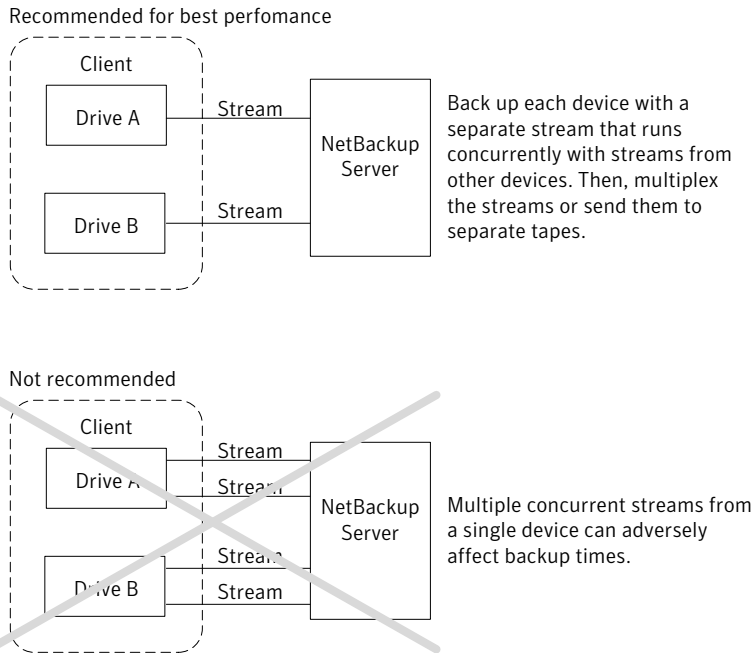
- To reduce backup time

Multiple data streams can reduce the backup time for large backups by splitting the backup into multiple streams. Use multiplexing, multiple drives, or a combination of the two to process the streams concurrently. Configure the backup so each device on the client is backed up by a separate data stream that runs concurrently with streams from other devices. For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The heads must move back and forth between the tracks that contain files for the respective streams. [Figure 15-5](#) shows that multiple concurrent streams from a single device are not recommended.
- To reduce retry time for backup failures

Because the backup streams run independently, the use of multiple data streams can shorten the retry time in the event of a backup failure. A single failure only terminates a single stream. NetBackup can restart the failed stream without restarting the others. For example, assume the backup for a 10-gigabyte partition is split into five streams, each containing 2 gigabytes. If the last stream fails after it writes 1.9 gigabytes (a total of 9.9 gigabytes is backed up), NetBackup retries only the last gigabyte stream. If the 10-gigabyte partition is backed up without multiple data streams and a failure occurs, the entire 10-gigabyte backup must be retried. The **Schedule backup attempts** property in the **Global Attributes** properties, applies to each stream. For example, if the **Schedule backup attempts** property is set to 3, NetBackup retries each stream a maximum of three times. The Activity Monitor displays each stream as a separate job. Use the job details view to determine the files that are backed up by each of these jobs. See “[Global Attributes properties](#)” on page 131.
- To reduce administration by running more backups with fewer policies

Use multiple data streams in a configuration that contains large file servers with many file systems and volumes. Multiple data streams provide more backups with fewer policies than are otherwise required.

**Figure 15-5** Multiple stream recommendations



## Disable client-side deduplication attribute

This attribute appears only if the NetBackup Deduplication Option license key is active.

The **Disable client-side deduplication** attribute affects the behavior of clients that are configured for client direct backup, as follows:

- If you select the attribute, the clients do not deduplicate their own data and do not send their backup data directly to the storage server. The NetBackup clients that are configured for client direct backup send their data to a deduplication media server. That server deduplicates the data and then sends it to the storage server.
- If you do not select the attribute, the clients that are configured for client direct backups deduplicate their data. They also send it directly to the storage server. Media server deduplication and data transport is bypassed.

The **Deduplication Location** property on the master server **Client Attributes** host properties tab configures clients for client direct deduplication. This policy attribute overrides that **Deduplication Location** property.



See “[General tab of the Client Attributes properties](#)” on page 81.

More information about client deduplication is available.

See the *NetBackup Deduplication Guide*.

## Enable granular recovery attribute

The **Enable granular recovery** attribute is selectable for the following policy types:

- MS-Exchange-Server
- MS-SharePoint
- MS-Windows (for Active Directory)

With this option enabled, users can restore the individual objects that reside within a database backup image, such as:

- A user account from an Active Directory database backup
- Email messages or folders from an Exchange database backup
- A document from a SharePoint database backup

Granular-level restores can be performed only if the backup was written to a disk storage unit.

For more information on how to configure NetBackup to perform granular-level backups with a specific agent, see the following:

*NetBackup for Microsoft SharePoint Server Administrator's Guide*

*NetBackup for Microsoft Exchange Server Administrator's Guide*

For more information on how to configure NetBackup to perform granular-level backups with Active Directory, see the following:

See “[Active Directory granular backups and recovery](#)” on page 575.

## Keyword phrase attribute

The **Keyword phrase** attribute is a phrase that NetBackup associates with all backups or archives based on the policy. Only the Windows and UNIX client interfaces support keyword phrases.

Clients can use the same keyword phrase for more than one policy. The same phrase for multiple policies makes it possible to link backups from related policies. For example, use the keyword phrase legal department documents, for backups of multiple clients that require separate policies, but contain similar types of data.

The phrase can be a maximum of 128 characters in length. All printable characters are permitted including spaces and periods. (Default: no keyword phrase.)

Clients can also specify a keyword phrase for a user backup or archive. A user keyword phrase overrides the policy phrase.

## Snapshot Client Attributes

Snapshot Client attributes are available when the NetBackup Enterprise Client license is installed.

A snapshot is a point-in-time, read-only, disk-based copy of a client volume.

For more information about the Snapshot Client attributes, see the following guides:

- *NetBackup Snapshot Client Administrator's Guide.*
- *NetBackup for VMware Administrator's Guide.*
- *NetBackup for Hyper-V Administrator's Guide.*

## Microsoft Exchange Attributes

The Microsoft Exchange Attributes let you indicate what database backup source you want to use for an Exchange 2010 Database Availability Group or for an Exchange 2007 replication backup.

See the *NetBackup for Microsoft Exchange Server Administrator's Guide.*

## About the Schedules tab

The schedules that are defined on the **Schedules** tab determine when backups occur for the policy that is selected. Each schedule also includes various criteria, such as how long to retain the backups.

From the policy Schedules tab, perform the following tasks:

- To create a new schedule click **New**.
- To edit a schedule, select the schedule and click **Properties**.
- To delete a schedule, select the schedule and click **Delete**.

Schedule attributes appear on the following tabs:

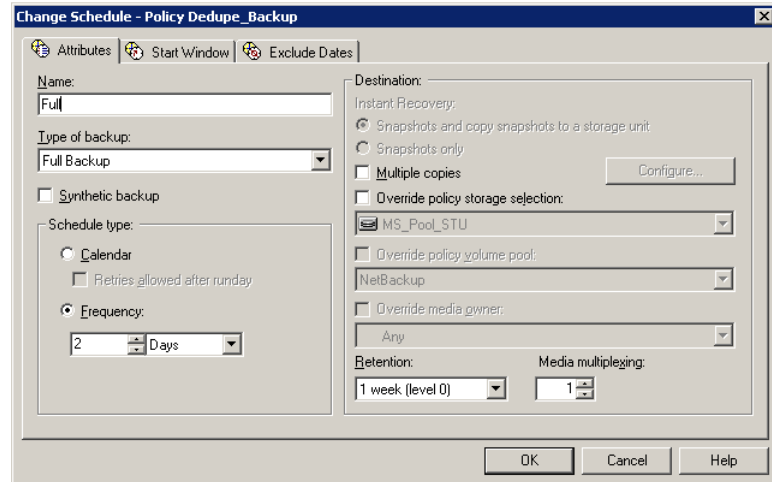
|                       |                                                                                              |
|-----------------------|----------------------------------------------------------------------------------------------|
| <b>Attributes</b> tab | Schedule the time and frequency at which a task runs, along with other scheduled attributes. |
|-----------------------|----------------------------------------------------------------------------------------------|

|                              |                                                                                                                                                                                                                                                               |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Start Window tab</b>      | Schedule the time of each day that a task runs.<br>See <a href="#">“Using the Start Windows tab”</a> on page 517.                                                                                                                                             |
| <b>Exclude Dates tab</b>     | Indicate the dates that a task should not run.<br>See <a href="#">“Using the Exclude Dates tab”</a> on page 521.                                                                                                                                              |
| <b>Calendar Schedule tab</b> | Schedule the run days for a task by indicating specific dates, recurring weekdays, recurring days of the month. (This tab appears only when Calendar is selected as the Schedule type.)<br>See <a href="#">“Using the Calendar Schedule tab”</a> on page 521. |

## About the Schedule Attributes tab

The following sections describe the settings on the **Attributes** tab for schedules. Schedule attributes include the backup type (different from the Policy Type), when the backup can occur, and how long the backup image is retained.

**Figure 15-6** Schedule Attributes tab



### Name attribute

Specify a name for the schedule by typing it in the **Name** attribute.

See [“NetBackup naming conventions”](#) on page 719.

The schedule name appears on screens and messages about the schedule.

If the schedule is a relocation schedule, created as part of a basic disk staging storage unit, the schedule name cannot be changed. The name defaults to the name of the storage unit.

See “[About staging backups](#)” on page 395.

## Type of backup attribute

The **Type of backup** attribute specifies the type of backup that the schedule controls. Select a backup type from the drop-down list. The list displays only the backup types that apply to the current policy.

If the schedule is a relocation schedule created as part of a basic disk staging storage unit, no backup type selection is needed.

See the following backup type descriptions:

- |                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full backups                   | A full backup backs up all of the files that are specified in the backup selections list for the policy. The files are backed up, regardless of when the files were last modified or backed up. Full backups occur automatically according to schedule criteria. If you run incremental backups, you must also schedule a full backup to perform a complete restore. If you configure a policy for a raw partition backup (formatted partitions only), select <b>Full Backup</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Cumulative incremental backups | <p>A cumulative incremental backup backs up the files that are specified in the backup selections list that changed since the last full backup. All files are backed up if no previous backup was done. Cumulative incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup and the last cumulative incremental backup.</p> <p>Do not combine differential incremental backups and cumulative incremental backups within the same Windows policy when the incremental backups are based on archive bit (default).</p> <p>By default, if the time between file creation and a full or a differential incremental backup is less than 5 minutes, the differential or cumulative incremental backup may yield unexpected results. The backups are successful, but the additional files are backed up.</p> <p>See “<a href="#">About incremental backups</a>” on page 494.</p> |

|                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Differential incremental backups         | <p>A differential incremental backup backs up the files that changed since the last successful incremental (differential or cumulative) or full backup. All files are backed up if no previous backup was done. Differential incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup, the last cumulative incremental, and all differential incremental backups that occurred since the last full backup.</p> <p>By default, if the time between file creation and a full or a differential incremental backup is less than 5 minutes, the differential or cumulative incremental backup may yield unexpected results. The backups are successful, but the additional files are backed up.</p> <p>See <a href="#">“About incremental backups”</a> on page 494.</p> |
| User backups                             | <p>A user initiates a user backup through the Backup, Archive, and Restore client interface. A user backup backs up all files that the user specifies. Users can start backups only during the times that are allowed on the schedule <b>Start Window</b> tab.</p> <p>If the schedule is to be used for a catalog archive, select <b>User Backup</b> for the backup type.</p> <p>See <a href="#">“Creating a catalog archiving policy”</a> on page 618.</p>                                                                                                                                                                                                                                                                                                                                                                         |
| User archive backups                     | <p>A user initiates a user archive through the Backup, Archive, and Restore client interface. A user archive backup first backs up the files that the user indicates. Then it deletes the files from the local disk if the backup is successful. Archive backups free local disk space while retaining a copy for future use. The copy is kept until the retention period expires. Users can start archives only during the times that are specified in the schedule <b>Start Window</b> tab.</p> <p><b>Note:</b> The NetBackup administrator should make sure that a full backup of the client exists before a user archives files from the client.</p>                                                                                                                                                                            |
| Application backups                      | <p>An application backup is a backup type that applies to all database agent clients.</p> <p>For more information on how to configure schedules for this type of backup, see the NetBackup guide that came with the product.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Automatic backups                        | <p>An automatic backup is a backup type for all database agent clients, except NetBackup for Informix and Oracle.</p> <p>For more information on how to configure schedules for this type of backup, see the NetBackup guide for the database product.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Automatic incremental backups            | <p>An automatic incremental backup applies only to NetBackup for Informix clients.</p> <p>For more information on how to configure schedules for this type of backup, see the <i>NetBackup for Informix Administrator’s Guide</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Automatic cumulative incremental backups | <p>An automatic cumulative incremental backup applies only to NetBackup for Oracle clients.</p> <p>For more information on how to configure schedules for this type of backup, see the <i>NetBackup for Oracle Administrator’s Guide</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

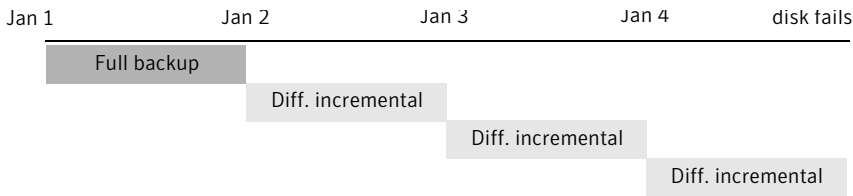
|                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Automatic differential incremental backups | <p>An automatic differential incremental backup applies only to NetBackup for Oracle clients.</p> <p>For more information on how to configure schedules for this type of backup, see the <i>NetBackup for Oracle Administrator's Guide</i>.</p>                                                                                                                                                                                                                                                                                                                                                        |
| Automatic full backups                     | <p>An automatic full backup applies only to NetBackup for Informix and for Oracle clients.</p> <p>For more information on how to configure schedules for this type of backup, see the <i>NetBackup for Informix Administrator's Guide</i> or the <i>NetBackup for Oracle Administrator's Guide</i>.</p>                                                                                                                                                                                                                                                                                                |
| Automatic Vault sessions                   | <p>An automatic Vault session applies only to Vault policies. The option does not run a backup, but instead runs the command that is specified in the Vault policy's backup selections list. In this way it starts an automatic, scheduled vault session or vault eject operation. Available only when Vault is licensed.</p> <p>See <a href="#">“Creating a Vault policy”</a> on page 573.</p>                                                                                                                                                                                                        |
| Vault catalog backups                      | <p>Use when the schedule is for a catalog backup policy to be used by Vault. Available only when Vault is licensed.</p> <p>If the schedule is a Vault Catalog Backup type, You must configure one of the two schedule attribute combinations or the schedule cannot be saved:</p> <ul style="list-style-type: none"><li>■ Check and configure <b>Multiple copies</b>, or</li><li>■ Check <b>Override policy storage selection</b>, <b>Override policy volume pool</b> and specify the <b>Retention</b>.</li></ul> <p><b>Note:</b> The selected storage unit selection should not be Any Available.</p> |

## About incremental backups

The following examples show how data is included in a series of full and incremental backups.

A differential incremental backup backs up the data that changed since the last full or differential incremental backup. [Figure 15-7](#) shows how data is included in a series of full and differential incremental backups between January 1 and January 4.

**Figure 15-7** Full and differential incremental example




Recovery = Jan 1 (full) + Jan 2 (incr) + Jan 3 (incr) + Jan 4 (incr)

The January 1 full backup includes all files and directories in the policy backup selections list. The subsequent differential incremental backups include only the data that changed since the last full or differential incremental backup. If the disk fails sometime on January 4 (after the backup), the full backup and all three of the incremental backups are required for the recovery.

A cumulative incremental backup backs up the data that changed since the last full backup. [Figure 15-8](#) shows how data is included in a series of full and cumulative incremental backups between January 1 and January 4. The January 1 full backup includes all files and directories in the policy backup selections list. Each of the cumulative incremental backups include the data that changed since the last full backup. If the disk fails sometime on January 4 (after the backup), the full backup and the last cumulative incremental backup are required for the recovery.

**Figure 15-8** Full and cumulative incremental example



Recovery = Jan 1 (full) + Jan 4 (cumulative incremental)

**Table 15-8** Retention requirements for incremental backups

| Type         | Retention requirement | Comments                                                                                                                                                                                                                   |
|--------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Differential | Longer                | To restore all files requires the last full backup and all the differential incremental backups that occurred since the last full backup. Therefore, all the differentials must be kept until the next full backup occurs. |
| Cumulative   | Shorter               | Each cumulative incremental backup contains all the changes that occurred since the last full backup. Therefore, a complete restore requires only the most recent cumulative incremental in addition to the full backup.   |

**Table 15-9** Relative backup and restore times for incremental backups

| Type         | Backup time | Restore time | Comments                                                                                                                                                         |
|--------------|-------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Differential | Shorter     | Longer       | Less data in each backup, but all differential incremental backups are required since the last full backup for a restore. This results in a longer restore time. |
| Cumulative   | Longer      | Shorter      | More data in each backup, but only the last cumulative incremental is required for a complete restore (in addition to the full).                                 |

You can use a combination of cumulative and differential incremental backups together to get the advantages of both methods. For example, assume a set of schedules with the following backup frequencies and retention periods (notice that the differential incremental backups occur more often.)

**Table 15-10** Example frequencies and retention periods

| Backup type              | Frequency | Retention period |
|--------------------------|-----------|------------------|
| Full                     | six days  | two weeks        |
| Cumulative incremental   | two days  | four days        |
| Differential incremental | one day   | two days         |



The following set of schedules result in the following series of backups:

|       |              |            |              |            |              |       |              |
|-------|--------------|------------|--------------|------------|--------------|-------|--------------|
| Day 1 | Day 2        | Day 3      | Day 4        | Day 5      | Day 6        | Day 7 | Day 8        |
| Full  | Differential | Cumulative | Differential | Cumulative | Differential | Full  | Differential |

Notes about example:

- Every other day a differential incremental backup occurs, which usually has a minimum backup time.
- On alternate days, a cumulative incremental backup occurs, which requires more time than the differential backup, but not as much time as a full backup. The differential backup can now be expired.
- To recover all files may require (at most), two incremental backups in addition to the most recent full backup. The combination of backups usually means less restore time than if all differential incremental backups were used. The full backups can be done less often if the amount of data being backed up by the incremental backups is small.

## How NetBackup determines that files are due for backup

On Windows clients, NetBackup performs the incremental backups of files that are based on the **Perform incrementals based on archive bit** setting. This setting is found in the Backup, Archive, and Restore client interface, under **File > NetBackup Client Properties**, on the **General** tab.

If **Perform incrementals based on archive bit** is enabled, incremental backups for this client are based on the state of the archive bit of each file. The operating system sets the bit whenever a file changes and it remains set until cleared by NetBackup. The conditions under which NetBackup clears the bit depend on the type of backup being performed.

- For a full backup, NetBackup backs up files regardless of the state of their archive bit. After a full backup, the archive bit is always cleared.
- For a differential incremental backup, NetBackup backs up the files that have the archive bit set and have therefore changed. When the client receives a response from the server that indicates that the backup was successful (or partially successful) the archive bits are cleared. The clear archive bit lets the next differential incremental back up only the files that changed since the previous full or differential incremental backup.
- For a cumulative incremental backup, NetBackup backs up the files that have the archive bit set. However, NetBackup does not clear the archive bits after the backup. Without a clear archive bit, the next cumulative incremental

backup backs up changed files and the files that were in the cumulative incremental backup.

If **Perform incrementals based on archive bit** is disabled, NetBackup includes a file in an incremental backup only if the datetime stamp of the file has changed since the last backup. The datetime stamp indicates when the file was last backed up.

Depending on the timestamp, NetBackup behaves in the following manner:

- For a full backup, NetBackup backs up files regardless of the datetime stamp.
- For a differential incremental backup, NetBackup compares the datetime stamp of the file against the last full or incremental backup.
- For a cumulative incremental backup, NetBackup compares the datetime stamp of the file against the last full backup.

If files are installed or copied from another computer, the new files retain the datetime stamp of the originals. If the original date is before the last backup date, the new files are not backed up until the next full backup.

Incremental backups on UNIX clients look at all files and directories to determine if a backup is due based on a reference date. (That is, back up all the files that have changed since `date_x`).

UNIX files and directories have the following three times that are associated with them:

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>mtime</code> | The file modification time. The file system updates the <code>mtime</code> for a file or directory each time the file is modified. An application can save the <code>mtime</code> of the file before modifying it. The application then resets it with the <code>utime(2)</code> system call.                                                                                                                                                                                            |
| <code>atime</code> | The file access time. The file system updates the <code>atime</code> for a file or directory each time the file is accessed (read or write). An application can save the <code>atime</code> of the file before accessing it. The application then resets it with the <code>utime(2)</code> system call.                                                                                                                                                                                  |
| <code>ctime</code> | The inode change time. The <code>ctime</code> for a file or directory is updated each time the file or directory's inode changes. (For example, changes due to permissions, ownership, and link-counts changes.) The <code>ctime</code> for a file or directory cannot be saved before and reset after a change. The <code>ctime</code> of a file or directory changes when the <code>mtime</code> and <code>atime</code> (changes with the <code>utime(2)</code> system call) is reset. |

When NetBackup reads the data for a file that is included in a backup, it does not affect the file modification time. It does affect the access time of the file. For this reason, NetBackup saves the `atime` and `mtime` of the file before it reads the file.

Then NetBackup resets the `atime` and `mtime` with the `utime(2)` system call. NetBackup does not cause problems for storage migration products or the administrator scripts that use file access times (`atime`) as criteria for their operations. While this benefit is obvious, a side effect is that it does update the `ctime` of the file.

Customers can configure NetBackup so that it does not reset the access time of the file after it reads a file. Customers can choose to have NetBackup use the `ctime` and the `mtime` of the file to determine what files to include in an incremental backup. Normally, these two options are used together, but there may be some sites that want to use one without the other. By default, NetBackup uses only the `mtime` of the file to determine what files and directories to back up.

When a file is moved from one location to another, the `ctime` of the file changes, but the `mtime` remains unchanged. If NetBackup uses only the `mtime` to determine the files that are due during an incremental backup, it does not detect these moved files. For sites where using the `mtime` might create a problem, use the `ctime` to determine files due to be included in an incremental backup. The `ctime` is used if the `bp.conf` file contains the `USE_CTIME_FOR_INCREMENTALS` and `DO_NOT_RESET_FILE_ACCESS_TIME` entries.

When a directory is moved from one location to another, the `ctime` of the directory changes, but the `mtime` remains unchanged. Neither the `mtime` nor the `ctime` are changed for the files or directories within the moved directory. No reliable method using file timestamps can determine that files within a moved directory need to be included in an incremental backup.

In either case, these moved files and directories are included in subsequent full backups.

## Synthetic backup attribute

A synthetic full or synthetic cumulative incremental backup is a backup assembled from previous backups. The backups include one previous, traditional full backup, and subsequent differential backups, and a cumulative incremental backup. (A traditional full backup means a non-synthesized, full backup.) A client can then use the synthesized backup to restore files and directories in the same way that a client restores from a traditional backup.

Synthetic backups can be written to tape or to disk storage units, or a combination of both.

See [“About synthetic backups”](#) on page 583.

## Calendar schedule type

Calendar-based schedules allow administrators to select specific days to run a policy. Select the **Calendar** schedule attribute for the **Calendar Schedule** tab to appear in the **Change Schedule** dialog box.

See [“Using the Calendar Schedule tab”](#) on page 521.

A calendar-based relocation schedule determines the days that images are swept from the disk staging storage unit to the final destination storage unit. (A relocation schedule is created as part of a basic disk staging storage unit configuration.)

### Retries allowed after runday attribute

Enable **Retries allowed after runday** to have NetBackup attempt to complete this schedule until the backup is successful. With this attribute enabled, the schedule attempts to do run, even after a specified run day.

## Frequency schedule type

Use the **Frequency** attribute to specify how much time must elapse between the successful completion of a scheduled task and the next attempt.

For example, assume that a schedule is set up for a full backup with a frequency of one week. If NetBackup successfully completes a full backup for all clients on Monday, it does not attempt another backup for this schedule until the following Monday.

NetBackup recognizes the intervals that suggest schedules based on days, even if the job does not run daily. For example, if the frequency is 48 hours, NetBackup attempts to run the job about the same time on those days that it is to run. (NetBackup checks if the frequency is divisible by 24 hours.) If the interval is 50 hours, NetBackup tries to run the job 50 hours after the last successful backup.

A frequency-based relocation schedule determines how often images are swept from the basic disk staging storage unit to the final destination storage unit. (A relocation schedule is created as part of a basic disk staging storage unit configuration.)

To set the frequency, select a frequency value from the drop-down list. The frequency can be hours, days, or weeks.

---

**Note:** **Frequency** does not apply to user schedules because the user can perform a backup or archive whenever the time window is open.

---

---

**Note:** A policy can contain more than one schedule. However, Symantec recommends that calendar-based and frequency-based schedule types are not mixed within the same policy. Under some conditions, schedule types that are combined in one policy can cause unexpected results.

See [“Backup window considerations”](#) on page 526.

---

## How to set up backup frequency

To determine backup frequency, consider how often data changes. For example, determine if files change several times a day, daily, weekly, or monthly.

Typically, sites perform daily backups to preserve daily work. Daily backups ensure that only one day’s work is lost in case of a disk failure. More frequent backups are necessary when data changes many times during the day and these changes are important and difficult to reconstruct.

Daily backups are usually incremental backups that record the changes since the last incremental or full backup. Incremental backups conserve resources because they use less storage and take less time to perform than full backups.

Full backups usually occur less frequently than incremental backups but should occur often enough to avoid accumulating consecutive incremental backups. A large number of incremental backups between full backups increases the time it takes to restore a file. The time increases because of the effort that is required to merge the incremental backups when files and directories upon restore.

Consider the following when setting the frequency for full backups:

- Extend the time between full backups for the files that seldom change. A longer frequency uses fewer system resources. It also does not significantly increase recovery time because the incremental backups between full backups are smaller.
- Shorter the time between full backups for the files that change frequently. A shorter frequency decreases restore time. A shorter time between full backups can also use fewer resources. It reduces the cumulative effect of the longer incremental backups that are necessary to keep up with frequent changes in the files.

To achieve the most efficient use of resources, ensure that most of the files in a given policy change at about the same rate. For example, assume that half of the files in a policy selection list change frequently enough to require a full backup every week. However, the remaining files seldom change and require monthly full backups only. If all the files are in the same policy, full backups are performed weekly on all the files. This wastes system resources because half the files need full backups only once a month. A better approach is to divide the backups into

two policies, each with the appropriate backup schedule, or to use synthetic backups.

## How backup frequency determines schedule priority

If more than one automatic schedule is due for a client within a policy, the backup frequency determines the schedule that NetBackup uses, as follows:

- Jobs from the schedule with the lower frequency (longer period between backups) always have higher priority. For example, a schedule with a backup frequency of one year has priority over a schedule with a backup frequency of one month.
- If NetBackup encounters a backup policy with two schedules that are each due to run, the schedule that is first alphabetically runs first.

Alphabetical priority occurs if both of the following are true:

- Each schedule is within the defined time window.
- Each schedule is configured with the same frequency value.

For example, NetBackup prioritizes the following three schedules in the following order:

- `monthly_full` (frequency is one month)
- `weekly_full` (frequency is two weeks)
- `daily_incremental` (frequency is one week)

## Instant recovery options

The **Instant recovery** options are available under the following conditions:

- The **Snapshot Client** option is licensed and installed.  
Refer to the *NetBackup Snapshot Client Administrator's Guide*.
- **Perform snapshot backups** is selected.
- **Retain snapshots for Instant Recovery** is selected.

Snapshots and copy snapshots to a storage unit

When the **Snapshots and copy snapshots to a storage unit** attribute is enabled, the snapshot persists on the client volume with a backup copy made to the storage unit on the media server.

### Snapshots only

When the **Snapshots only** attribute is enabled, the snapshot is not backed up to tape or to other storage. NetBackup creates a snapshot on disk only. This option is required for the **NAS\_Snapshot** method.

The snapshot is created on the same device as the one that contains the original data if it uses **VxFS\_Checkpoint** method or is vxvm space optimized. In this case, another policy can be used to back up the data to a separate device.

With this attribute enabled, transaction logs are not be truncated at the end of the backup.

## Multiple copies attribute

When the **Multiple copies** attribute is enabled, NetBackup can create up to four copies of a backup simultaneously. The storage units must be on the same media server with sufficient resources available for each copy. For example, to create four copies simultaneously in a Media Manager storage unit, the unit needs four tape drives. (This option is sometimes referred to as Inline Copy.)

The **Maximum backup copies** property specifies the total number of backup copies that may exist in the NetBackup catalog (2 through 10). NetBackup creates the number of copies that is specified under **Multiple copies**, or the number that the **Maximum backup copies** property specifies, whichever is fewer.

See [“Maximum backup copies”](#) on page 135.

To create more than four copies, additional copies can be created at a later time using duplication.

If multiple original images are created simultaneously, the backup time that is required may be longer than for one copy. Also, if both Media Manager and disk storage units are specified, the duration of disk write operations match that of slower removable media write operations.

### Criteria for creating multiple copies

To create multiple copies, the following criteria must be met:

- The backup destinations must share the same media server with sufficient resources available for each copy.
- The storage units that are used for multiple copies must be configured to allow a sufficient number of concurrent jobs to support the concurrent copies. The pertinent storage unit settings are **Maximum concurrent jobs** and **Maximum concurrent write drives**.

See [“Maximum concurrent jobs setting”](#) on page 383.

See [“Maximum concurrent write drives setting”](#) on page 383.

- When using a storage lifecycle policy to create multiple copies, the number of destinations in the lifecycle cannot exceed the **Maximum Backup Copies** setting in the **Global** host properties. The lifecycle cannot be saved until the destinations are decreased, or until the **Maximum Backup Copies** setting is increased.

See [“Maximum backup copies”](#) on page 135.

Multiple copy operations do not support the following:

- Third-party copies.
- NDMP storage units.
- Storage units that use a QIC (quarter-inch cartridge) drive type.
- The option to create multiple copies is not allowed for synthetic backups.

## Multiple copies and disk staging storage units

Multiple copies can also be configured for a relocation schedule, created as part of basic disk staging storage unit configuration. The **Maximum backup copies** Global host property must be set to include an additional copy beyond the number of copies to be created in the **Multiple Copies** dialog box. For example, to create four copies in the **Multiple Copies** dialog box, the **Maximum backup copies** property must be set to five or more.

Since NetBackup eventually relocates a backup from the initial, temporary staging storage unit to a final destination, NetBackup considers this to be one copy. NetBackup automatically counts this copy against the **Maximum backup copies** value.

## Multiple copies and storage lifecycle policies

If a schedule is configured to create multiple copies, none of the copies can be sent to a storage lifecycle policy.

That is, in the **Configure Multiple Copies** dialog box, the **Storage unit** selection cannot indicate a storage lifecycle policy.

Storage lifecycle policies offer their own method to create multiple copies.

See [“Writing multiple copies using a storage lifecycle policy”](#) on page 433.

## Configuring multiple copies in a policy schedule

To configure a policy schedule to create multiple copies, use the following procedure.



**To configure a schedule to create multiple copies**

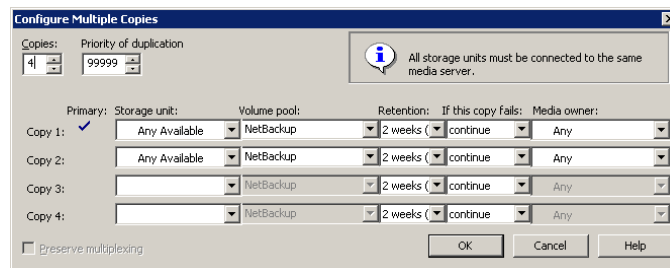
- 1 Expand **NetBackup Management > Policies**.
- 2 Double-click an existing policy or select **Actions > New > New Policy** to create a new policy.
- 3 Select the **Schedules** tab.
- 4 Double-click an existing schedule or click **New** to create a new schedule.
- 5 In the **Attributes** tab, select **Multiple copies**, then click **Configure**.

The **Multiple Copies** option is disabled if the destination for this policy is a storage lifecycle policy. (The **Policy storage** selection is on the policy **Attributes** tab.) NetBackup does not allow the two methods for creating multiple copies to be enabled at the same time.

See [“Use only one method to create multiple copies”](#) on page 433.

- 6 In the **Copies** field, specify the number of copies to be created simultaneously. The maximum is four, or the number of copies that the **Maximum backup copies** setting specifies, whichever is fewer.

See [“Maximum backup copies”](#) on page 135.



Copy 1 is the primary copy. If Copy 1 fails, the first successful copy is the primary copy.

To configure multiple copies as part of a relocation schedule for a basic disk staging storage unit, set the **Maximum backup copies** Global host property to include an additional copy beyond the number of copies to be created in the **Multiple Copies** dialog box. For example, to create four copies in the **Multiple Copies** dialog box, the **Maximum backup copies** property must be set to five or more.

See [“Multiple copies and disk staging storage units”](#) on page 504.

- 7 Specify the priority of the duplication job compared to other jobs in the queue (0 to 99999).

- 8 Specify the storage unit where each copy is stored. Select Any Available to allow NetBackup to select the storage unit at runtime.

If a Media Manager storage unit contains multiple drives, the storage unit can be used for both the original image and the copies.

- 9 Specify the volume pool where each copy is stored.

- 10 Select the retention level for each copy.

See “[Retention attribute](#)” on page 510.

- 11 Select what should happen to the copy in the event that the copy does not complete. Select whether the entire job should fail (**fail all copies**), or whether the remaining copies should continue.

If a copy is configured to allow other copies to continue, and if **Take checkpoints every** is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.

- 12 For tape media, specify who should own the media onto which NetBackup writes the images:

|                |                                                                                                                                                                                                                                                              |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any            | NetBackup chooses the media owner, either a media server or server group.                                                                                                                                                                                    |
| None           | Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.                                                                                              |
| A server group | Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All the media server groups configured in the NetBackup environment appear in the drop-down list. |

These settings do not affect images residing on disk. Images that reside on shared disks are not owned by any one media server. Any media server with access to the shared pool of disk can access the images.

- 13 Click **OK**, until the policy is saved.

#### To configure a basic disk staging relocation schedule to create multiple copies

- 1 Expand **NetBackup Management > Storage Units**.
- 2 Double-click an existing basic disk staging storage unit or Select **Actions > New > New Storage Unit** to create a new basic disk staging storage unit. To create a new basic disk staging storage unit, select the **Temporary staging area** checkbox and configure the other storage unit selections.

See “[About staging backups](#)” on page 395.

- 3 Click **Staging Schedule**.
- 4 In the **Attributes** tab, specify the priority that NetBackup should assign to the duplication jobs compared to other jobs in the queue. Range: 0 (default) to 99999 (highest priority).
- 5 Select a schedule type and schedule when the policy should run.
- 6 Select whether to use an alternate read server by checking **Use alternate read server**. The alternate server that is indicated is allowed to read a backup image originally written by a different media server.
- 7 Select **Multiple copies** and click **Configure**.
- 8 In the **Copies** field, specify the number of copies to create simultaneously. Copy 1 is the primary copy. If Copy 1 fails, the first successful copy is the primary copy.

The **Maximum backup copies** Global host property must include an additional copy beyond the number of copies that are indicated in the Copies field. For example, to create four copies in the **Multiple Copies** dialog box, set the **Maximum backup copies** property to five or more.

See [“Global Attributes properties”](#) on page 131.

See [“Multiple copies and disk staging storage units”](#) on page 504.

- 9 Specify the storage unit where each copy is stored. If a Media Manager storage unit has multiple drives, it can be used for both the source and the destination.
- 10 Specify the volume pool where each copy is stored.
- 11 Select the retention level for each copy.  
 See [“Retention attribute”](#) on page 510.
- 12 Select what should happen in the event that the copy does not complete. Select whether the entire job should fail, or whether the remaining copies should continue.

**13** Click **OK**. If a copy is configured to allow other copies to continue, and **Take checkpoints every** is selected for this policy, then only the last failed copy that contains a checkpoint can be resumed.

**14** For tape media, specify who should own the media onto which NetBackup writes the images:

|                |                                                                                                                                                                                                                                                          |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any            | NetBackup chooses the media owner, either a media server or server group.                                                                                                                                                                                |
| None           | Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.                                                                                          |
| A server group | Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups configured in the NetBackup environment appear in the drop-down list. |

These settings do not affect images residing on disk. Images on shared disk are not owned by any one media server. Any media server with access to the shared pool of disk can access the images.

## How to restore from a specific backup copy

Usually NetBackup restores from the primary copy of an image. However, it is possible to restore from a specific backup copy other than the primary copy. To do so, use the `bprestore` command.

See the Backup, Archive, and Restore online Help or *NetBackup Commands*.

## Override policy storage selection attribute

The **Override policy storage selection** attribute specifies the following:

- Whether to use the policy storage unit or the storage lifecycle policy as specified in the policy's Attributes tab.
- Whether to use a different, specified storage unit or storage lifecycle for this schedule.

Click the check box to override the **Policy storage** selection that is indicated in the **Attributes** tab. Choose the storage unit or lifecycle from the drop-down list of previously configured storage units and lifecycle policies. If the list is empty, no storage units or lifecycles have been configured.

If a data classification is indicated for the policy, only those storage lifecycles with the same data classification appear.

See [“Data classifications attribute”](#) on page 465.

To use only the policy storage selection that is indicated by the **Policy storage** setting in the **Attributes** tab, do not enable the check box.

See [“Policy storage attribute”](#) on page 465.

---

**Note:** Storage lifecycle policies cannot be selected within the multiple copies configuration dialog box.

---

## Override policy volume pool attribute

The **Override policy volume pool** attribute specifies whether to use the policy volume pool or another volume pool for this schedule as follows:

- To override the volume pool that the **Policy Volume Pool** General Attribute specifies, select the check box. Choose the volume pool from the list of previously configured volume pools.
- To use the policy volume pool, do not select the check box. NetBackup uses the volume pool that is specified in the **Policy volume pool** General Attribute. If no policy volume pool is specified, NetBackup uses NetBackup as the default. If the policy is for a NetBackup catalog, NBU-Catalog policies use CatalogBackup.

See [“Policy volume pool attribute”](#) on page 467.

## Override media owner attribute

The **Override media owner** attribute applies to tape media only and specifies whether to use the policy media owner or another owner for this schedule. The rules for shared disk media are more flexible, so override settings are not needed.

Specify one of the following:

|                |                                                                                                                                                                                                                                                          |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any            | NetBackup chooses the media owner, either a media server or server group.                                                                                                                                                                                |
| None           | Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.                                                                                          |
| A server group | Specifies that a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written. All media server groups configured in the NetBackup environment appear in the drop-down list. |

To override the media owner that the **Media Owner** General Attribute specifies, select the check box. Choose the media owner from the drop-down list

To use the policy media owner, do not select the check box. NetBackup uses the media owner that is specified in the **Media Owner** General Attribute.

See [“Media owner attribute”](#) on page 473.

## Retention attribute

The **Retention** attribute specifies how long NetBackup retains the backups. To set the retention period, select a time period (or level) from the drop-down list. When the retention period expires, NetBackup deletes information about the expired backup. Once the backup expires, the files in the backup are unavailable for restores. For example, if the retention is two weeks, data can be restored from a backup that this schedule performs for only two weeks after the backup.

If a policy is configured to back up to a lifecycle, the **Retention** attribute in the schedule is not followed. The retention that the lifecycle indicates is followed instead.

See [“Adding a storage destination to a storage lifecycle policy”](#) on page 421.

### How to assign retention periods

The retention period for data depends on how likely the need is to restore information from media after a certain period of time. Some data (financial records, for example) have legal requirements that determine the retention level. Other data (preliminary documents, for example) can probably be expired when the final version is complete.

A backup’s retention also depends on what needs to be recovered from the backup. For example, if day-to-day changes are critical, keep all the incremental backups in addition to the full backups for as long as the data is needed. If incremental backups only track work in progress toward monthly reports, expire the incremental backups sooner. Rely on the full backups for long-term recovery.

Establish some guidelines that apply to most of the data to determine retention periods. Note the files or the directories that have retention requirements outside of these guidelines. Plan to create separate policies for the data that falls outside of the retention requirement guidelines. For example, place the files and directories with longer retention requirements in a separate policy. Schedule longer retention times for the separate policies without keeping all policies for the longer retention period.

Another consideration for data retention is off-site storage of the backup media. Off-site storage protects against the disasters that occur at the primary site.

Set the retention period to infinite for the backups that must be kept for more than one year as follows:

- One method to implement off-site disaster recovery is to use the duplication feature to make a second copy for off-site storage.
- Another approach is to send monthly or weekly full backups to an off-site storage facility. To restore the data, request the media from the facility. (Note that a total directory or disk restore with incremental backups requires the last full backup plus all incremental backups.)
- Consider configuring an extra set of schedules to create the backups to use as duplicates for off-site storage.

Ensure that adequate retention periods are configured, regardless of the method that is used for off-site storage. Use the NetBackup import feature to retrieve expired backups.

## Precautions for assigning retention periods

For full backups, specify a time period that is longer than the frequency setting for the schedule. (The frequency is how often the backup runs). For example, if the frequency for a full backup is one week, specify a retention period of two to four weeks. Two to four weeks provides enough of a margin to ensure that the current full backup does not expire before the next full backup occurs.

For cumulative incremental backups, specify a time period that is longer than the frequency setting for the schedule. For example, if the frequency setting is one day, specify a retention period of one week. One week provides enough of a margin to ensure that the current cumulative-incremental backup does not expire before the next successful one occurs. A complete restore requires the previous full backup plus the most recent cumulative-incremental backup.

For differential incremental backups, specify a time period that is longer than the period between full backups. For example, if full backups occur weekly, save the incremental backups for two weeks.

A complete restore requires the previous full backup plus all subsequent incremental backups.

- NetBackup does not track backups after the retention period expires. Assign an adequate retention period as recovering files after the retention period expires is difficult or impossible.
- Within a policy, assign a longer retention period to full backups than to incremental backups. It may not be possible to restore all the files if the full backup expires before the incremental backups.
- Archive schedules normally use an infinite retention period.

## Changing retention periods

Set the default retention periods by selecting **NetBackup Management > Host Properties > Master Server > Double-click on master server > Servers > Retention Periods**.

The retention periods are indexed to different levels. For example, the default retention period for level 0 is one week. NetBackup also uses the level to determine the volume to use to store a backup.

See [“Mixing retention levels on tape volumes”](#) on page 512.

See [“Retention Periods properties”](#) on page 172.

## Mixing retention levels on tape volumes

By default, NetBackup stores each backup on a tape volume that contains existing backups at the same retention level. If a backup has a retention level of 2, NetBackup stores it on a tape volume with other backups at retention level 2. When NetBackup encounters a backup with a different retention level, it switches to an appropriate volume. Because tape volumes remain assigned to NetBackup until all the backups on the tape expire, this approach results in more efficient use of media. One small backup with an infinite retention prevents a volume from being reused, even if all other backups on the volume expired.

To mix retention levels on volumes, select **Allow multiple retentions per media** on the **Media** host properties.

If you keep only one retention level on each volume, do not use any more retention levels than necessary. Multiple retention levels increase the number of required volumes.

See [“Media properties”](#) on page 153.

---

**Note:** Retention levels can be mixed on disk volumes with no restrictions.

---

## Media multiplexing attribute

The **Media multiplexing** attribute specifies the maximum number of jobs from the schedule that NetBackup can multiplex onto any one drive. Multiplexing sends concurrent backup jobs from one or several clients to a single drive and multiplexes the backups onto the media.

Specify a number from 1 through 32, where 1 specifies no multiplexing. Any changes take effect the next time a schedule runs.



---

**Note:** Some policy or some schedule types do not support media multiplexing. The option cannot be selected in those instances.

---

To configure multiplexed backups, multiplexing must be indicated in both the storage unit (**Maximum Streams Per Drive** setting) and the schedule (**Media Multiplexing** setting) configuration. Regardless of the **Media multiplexing** setting, the maximum jobs that NetBackup starts never exceeds the **Maximum Streams Per Drive** value for the storage unit.


See “[Maximum streams per drive setting](#)” on page 385.

## About multiplexing

NetBackup multiplexing sends concurrent backups from one or several clients to a single storage device. NetBackup multiplexes the backups sequentially onto the media. Multiplexed and unmultiplexed backups can reside on the same volume. Separate volume pools or media IDs are not necessary.

[Figure 15-9](#) shows the multiplexed flow of client data to a server.

**Figure 15-9** Multiplexed backups



Multiplexing is generally used to reduce the amount of time that is required to complete backups. The performance in the following situations is improved by using multiplexing:

Slow clients

Instances in which NetBackup uses software compression, which normally reduces client performance, are also improved.

Multiple slow networks

The parallel data streams take advantage of whatever network capacity is available.

Many short backups (for example, incremental backups)

In addition to providing parallel data streams, multiplexing reduces the time each job waits for a device to become available. Therefore, the storage device transfer rate is maximized.

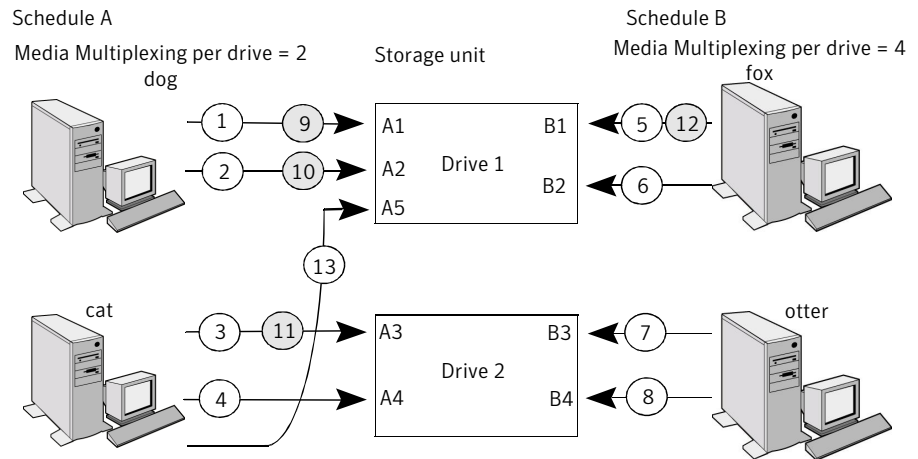
No special action is required to restore a multiplexed backup. NetBackup finds the media and restores the requested backup. Multiplexing reduces performance on restores because it uses extra time to read the images.

To reduce the affect of multiplexing on restore times, set the storage unit maximum fragment size to a value smaller than the largest allowed value.

When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches either of the following:

- This schedule's **Media Multiplexing** setting  
If the limit is reached for a drive, NetBackup sends jobs to other drives.  
In the [Figure 15-10](#), when the Schedule A limit is reached on Drive 1, NetBackup adds Schedule A jobs to Drive 2.
- The storage unit's **Maximum streams per drive** setting  
NetBackup can add jobs from more than one schedule to a drive.  
See "[Maximum streams per drive setting](#)" on page 385.  
In [Figure 15-10](#), unshaded numbers denote a job starting. Shaded numbers denote job completion. For example, 1 denotes the start of job A1 on Drive 1. Nine denotes the completion of job A1 on Drive 1.

**Figure 15-10** Multiplexing and schedules



Assume schedule A begins first (note that the schedules can be in the same or in different policies). Also, assume that Allow Multiple Data Streams is enabled, so a

- ①② Jobs A1 and A2 from client dog start on drive 1. Schedule A Media Multiplexing limit of 2 is reached for this drive.
- ③④ Jobs A3 and A4 from client cat start on drive 2. Schedule A Media Multiplexing limit of 2 is reached for this drive.
- ⑤⑥ Jobs B1 and B2 for client fox start on drive 1. Storage unit max mpx is reached for this drive.
- ⑦⑧ Jobs B3 and B4 from client otter start on drive 2. All jobs are now running for schedule B. Storage Unit Max mpx is reached for drive 2.
- ⑨⑩ Jobs A1 and A2 from client dog finish on drive 1. However, jobs B1 and B2 for client fox continue to run. Schedule A Media Multiplexing limit of 2
- ⑪⑫ Job A3 from client cat finishes on drive 2 and job B1 from client fox finishes on drive 1. Job B2 is the only job currently running on drive 1.
- ⑬ Job A5 from client cat starts on drive 1. Job A5 is the last job for schedule A. Schedule A Media Multiplexing limit of 2 prevents job A5 from starting on Drive 2. Therefore, job A5 starts on Drive 1. NetBackup attempts to add multiplexed jobs to drives that already use multiplexing. If multiplexed jobs are confined to specific drives, other drives are available for non-multiplexed

NetBackup attempts to add multiplexed jobs to drives that are already use multiplexing. If multiplexed jobs are confined to specific drives, other drives are available for non-multiplexed jobs.

If the backup window closes before NetBackup can start all the jobs in a multiplexing set, NetBackup completes only the jobs that have started.

For example, [Figure 15-10](#) assumes that the Activity Monitor shows A1 through A5 as queued and active.

If only A1 and A2 start before the window closes, NetBackup does not perform the other jobs that are in the set. If the window closes before any jobs start, then only the first queued and active job starts and completes. (A1 in this example.)

Consider the following configuration settings when using multiplexing:

**Limit jobs per policy**

Set **Limit jobs per policy** high enough to support the specified level of multiplexing.

**Maximum jobs per client**

The **Maximum Jobs Per Client** property limits the number of backup jobs that can run concurrently on any NetBackup client. **Maximum Jobs Per Client** appears on the **Global** properties dialog box.

Usually, the client setting does not affect multiplexing. However, consider a case where jobs from different schedules on the same client go to the same storage unit. In this case, the maximum number of jobs that are permitted on the client is reached before the multiplexing limit is reached for the storage unit. When the maximum number of jobs on the client is reached, it prevents NetBackup from fully using the storage unit's multiplexing capabilities.

Choose a value that is based on the ability of the central processing unit to handle parallel jobs. Because extra buffers are required, memory is also important. If the server cannot perform other tasks or runs out of memory or processes, reduce the **Maximum Streams Per Drive** setting for the storage unit.

Consider the following items to estimate the potential load that multiplexing can place on the central processing unit:

- The maximum concurrent jobs that NetBackup can attempt equals the sum of the concurrent backup jobs that can run on all storage units.
- The maximum concurrent jobs that can run on a storage unit equals the value of **Maximum Streams Per Drive**, multiplied by the number of drives.

**Maximum jobs this client**

You can set the maximum number of jobs that are allowed on a specific client without affecting other clients.

**MPX restore delay**

The **Delay On Multiplexed Restores** property applies to multiplexed restores. The property specifies how long the server waits for additional restore requests of files and raw partitions in a set of multiplexed images on the same tape. **Delay On Multiplexed Restores** appears on the **General Server** properties dialog box.

See [“Maximum streams per drive setting”](#) on page 385.

See [“Media multiplexing attribute”](#) on page 512.

## About demultiplexing

Demultiplexing speeds up future restores and is useful for creating a copy for off-site storage. Use the duplication process in the Catalog utility to demultiplex a backup.

Duplication allows one multiplexed backup at one time to be copied from the source media to the target media. When duplication is complete, the target contains a single demultiplexed copy of each duplicated backup. (The target can also contain other backups.) The duplicate copy can be made into the primary copy. Do not select **Preserve Multiplexing** in the **Setup Duplication Variables** dialog box when backups are duplicated.

---

**Note:** If you use the `bpduplicate` command instead of the NetBackup Administration Console, do not include the `-mpx` option on that command.

---

See [“Duplicating backup images”](#) on page 659.

## Using the Start Windows tab

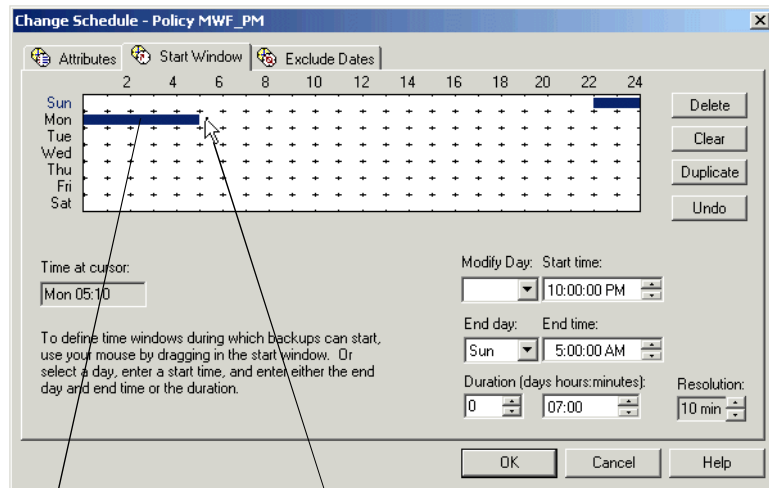
The **Start Window** tab provides controls for setting time periods during which NetBackup can start backups, archives, or basic disk staging relocation when using this schedule. Time periods are referred to as time windows. Configure time windows so that they satisfy the requirements necessary to complete a task or job. For example, create a different window for the backups that open each day for a specific amount of time, or keep the window open all week.

## Creating a schedule window

### To create a schedule window

- 1 Position the cursor over the day and time that the time window is to open. The **Time at cursor** field reflects the day and time that corresponds to the current position of the cursor.

Move the cursor and the value changes. The time is based on a 24-hour clock. For example, 1:00 A.M. is **01:00** and 11:00 P.M. is **23:00**. 12:00 P.M. is **00:00**.



Time window

Current cursor position

To set the resolution for start or end time settings, change the value in the **Resolution** field. For example, a resolution of 10 minutes allows time window adjustments by 10-minute increments.

- 2 When the cursor is over the chosen start time, press the left mouse button. Drag the cursor to the day and time when the window is to close.

Release the button and the fields display the following information:

- The **Modify Day** and **Start Time** fields display the day and time the time window opens.
- The **End Day** and **End Time** fields display the time that the time window closes.
- The **Duration** indicates the scope of the time window that is based on the start time and the end date and times.

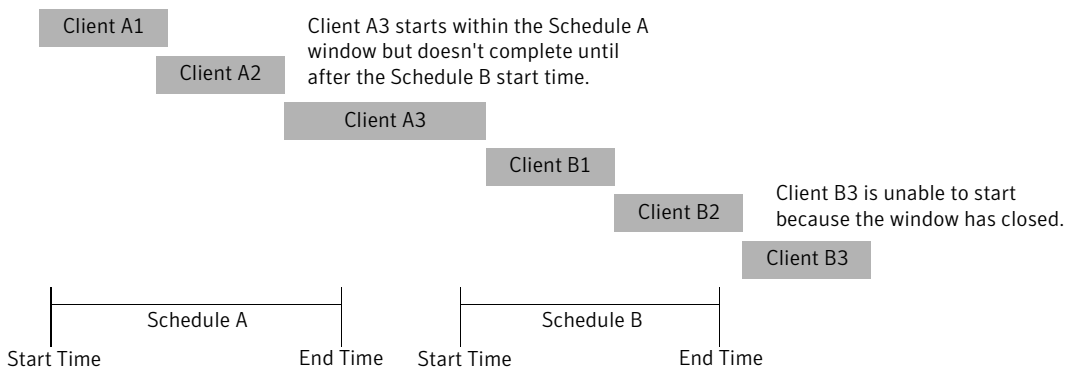
Specify enough time to allow all clients in the policy to complete a backup.

Consider allowing extra time in the schedule in case the schedule starts late due to factors outside of NetBackup. (Delays due to unavailable devices, for example.) Otherwise, all backups may not have a chance to start.

## Example of schedule duration

Figure 15-11 represents the effect of schedule duration on two full backup schedules. The start time for schedule B begins shortly after the end time for previous schedule A. Both schedules have three clients with backups due.

**Figure 15-11** Duration example



The backup for client A3 in Schedule A does not finish until after the Schedule B window opens. Schedule A does not leave enough time for the Schedule B backups. Client B3 must wait until the next time NetBackup runs Schedule B.

Client A3 illustrates that if a backup starts, it runs to completion even if the window closes while the backup is running.

## Creating time windows on successive days

### To create time windows on successive days

- 1 Move the cursor over the chosen start time. Press the Shift key and the left mouse button.
- 2 Drag the cursor to the time that the first time window is to close.
- 3 Keep the button and key pressed, and drag the cursor down to duplicate the window on successive days.

## Copying a time window

### To copy a time window to another day

- 1 Press the Ctrl key.
- 2 Click and drag an existing time window to another day.

If blank days follow the selected window, click **Duplicate** to duplicate the selected time window on those days. Duplication stops when it reaches a day that already contains a defined schedule.

## Changing a time window

### To change a time window

- 1 Select the time window.
- 2 Position the cursor over the end of the window that is to change. The cursor changes to a two-headed arrow.
- 3 Press the left mouse button and drag the time window to the chosen length.

## Moving a time window

### To move a time window

- 1 Select the time window.
- 2 Position the cursor over the center of the selected window so the cursor changes to crossed arrows.
- 3 Press the left mouse button and drag the time window to the chosen location within the schedule area.

## Deleting a time window

### To delete a time window

- 1 Select the time window.
- 2 Click **Delete**.

## Deleting all time windows

### To delete all time windows

- ◆ Click **Clear**.



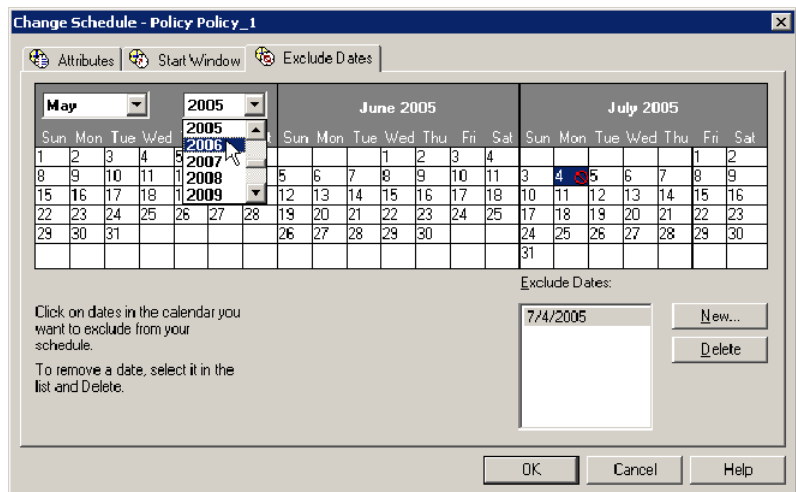
## Using the Exclude Dates tab

Use the **Exclude Dates** tab to exclude specific dates from a schedule. The **Exclude Dates** tab displays a 3-month calendar. Use the controls at the top of the calendar to change the month or year.

### Excluding dates from a policy

To exclude a date from the policy schedule

- 1 Select the **Exclude Dates** tab.



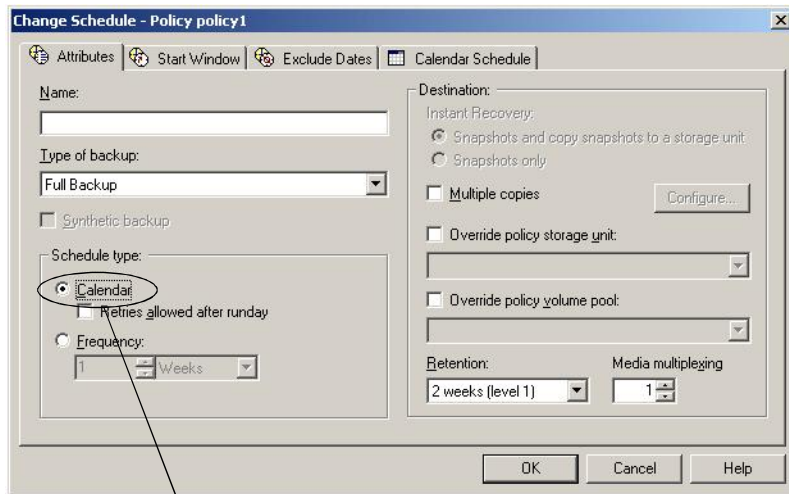
- 2 Use one of the following methods to indicate a date:
  - Click the date on the calendar to exclude.
  - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.

The date appears in the **Exclude Dates** list.
- 3 After the dates are selected, select another tab to make changes or click **OK** to close the dialog box.

## Using the Calendar Schedule tab

The **Calendar Schedule** tab appears when the **Calendar** option is selected as the Schedule type on the **Attributes** tab of the Schedule dialog box. Calendar-based schedules provide several run day options for determining when a task runs.

Figure 15-12 Calendar selection in the Policy Attributes tab



Select Calendar on the Attributes tab to enable the Calendar Schedule tab

The Calendar Schedule tab displays a 3-month calendar. Use the controls at the top of the calendar to change the month or year.

## Scheduling by specific dates

A task can run on specific dates rather than follow a recurring schedule, and specific dates can be added to a recurring schedule. Use the **Specific dates** run day option to schedule specific dates for a task to run.

### To schedule a task on specific dates

- 1 In the **Calendar Schedule** tab, select **Specific Dates**.
- 2 Use one of the following methods to indicate a date:
  - Click the date in the calendar.
  - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.

The date appears in the **Specific Dates** list.

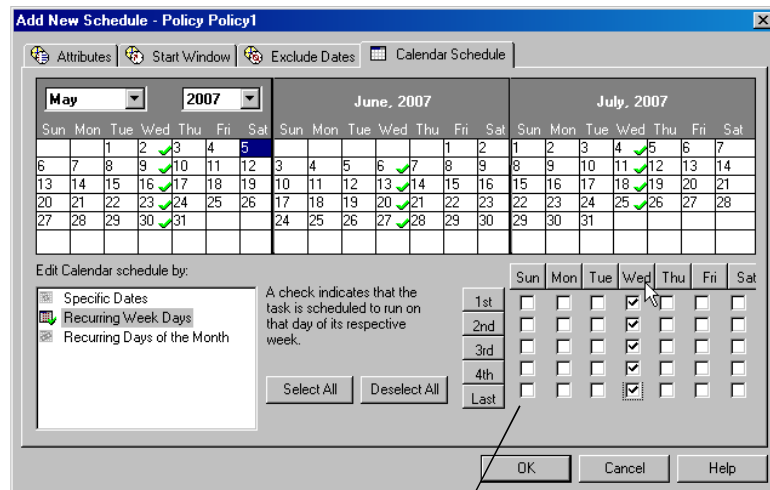
- 3 To remove a date, select it in the calendar schedule list and click **Delete**.
- 4 After the dates are selected, select another tab to make changes or click **OK** to save and close the dialog box.

## Scheduling by recurring week days

The **Recurring Week Days** option presents a matrix of days and weeks to schedule a task. The matrix is not a calendar. A check mark on a day indicates that the task is scheduled to run on that day of that week each month.

For example, schedule a task to run on the first and the third Thursday of every month. Or, schedule a task that runs the last week in every month.

**Figure 15-13** Recurring week days setting on the **Calendar Schedule** tab



Matrix

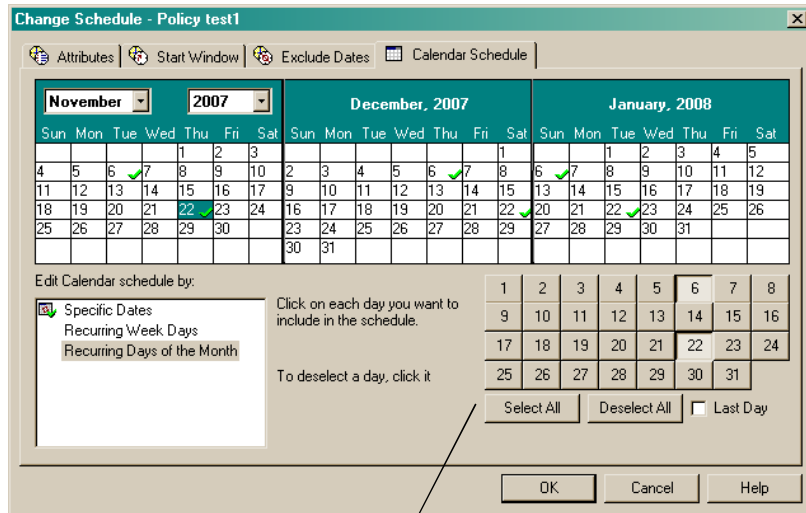
### To schedule a recurring weekly task

- 1 In the **Calendar Schedule** tab, select **Recurring Week Days**.
- 2 If necessary, click **Deselect All** to remove existing selections from the matrix.
- 3 Click a check box in the matrix to select the day. Or, click a check box to clear it.
- 4 Click the name of the day column header to select or clear the corresponding day for each week of the month.
- 5 Click a row number to select or clear the entire week.
- 6 Click the check box for the appropriate day in the **Last** row to schedule a task for the last week of each month. The task is scheduled, regardless of the number of weeks in the month.
- 7 After the dates are selected, select another tab to make changes or click **OK** to save and close the dialog box.

## Scheduling by recurring days of the month

The **Recurring Days of the Month** option presents a matrix to schedule a task for certain days of the month. A task can be scheduled to occur on the last day of the month, regardless of the actual date.

**Figure 15-14** Recurring days of the month setting on the **Calendar Schedule** tab



### To schedule a recurring monthly task

- 1 In the **Calendar Schedule** tab, select **Recurring Days of the Month**.
- 2 If necessary, click **Deselect All** to remove existing selections from the matrix.
- 3 To select all of the days in every month, click **Select All**.
- 4 Select the button for each day to be included in the run schedule. Click the button again to deselect the day.
- 5 Select the **Last Day** check box to run the schedule on the last day of the month, regardless of the date.
- 6 After the dates are selected, select another tab to make changes or click **OK** to save and close the dialog box.

## Considerations for user schedules

In order for users to perform backups and archives, you must create a schedule that allows user backups. A user backup schedule can be included in a policy that contains automatic backup schedules.

Restores can be performed at any time and are not scheduled.

---

**Note:** An archive is different from a backup. During an archive, NetBackup first backs up the selected files, then deletes the files from the local disk if the backup is successful. In this topic, references to backups also apply to the backup portion of archive operations unless otherwise noted.

---

### How to plan user backup and archive schedules

To plan schedules for user backups and archives, consider the following questions:

- What are the most convenient times for users to perform backups?  
If possible, do not permit user backups and archives when automatic backups are running. If an automatic backup is running when a user submits a backup or archive, NetBackup usually queues the user job.  
The job is not queued if there is a limiting setting. (For example, the **Limit Jobs per Policy** policy attribute or the **Maximum Jobs per Client** Global Attributes host property.)  
If the automatic backup continues to run, the user job misses the backup window. User jobs delay automatic backups and can cause backups to miss the backup window.
- Which storage unit should be used for user backups?  
Use a different storage unit to eliminate conflicts with automatic backups.
- Which volume pool should be used for user backups?  
Use a different volume pool to manage the media separate from the automatic backup media.

---

**Note:** If the retention period expires for a backup, it can be difficult or impossible to restore the archives or backups.

---

- How long should an archive be kept?  
Consider setting the retention period for archives to infinite, since the disk copy of the files is deleted.

## How to create separate policies for user schedules

User backup and archive schedules do not need to be in a policy separate from automatic backup schedules. If you create separate policies for user backups or archives, the considerations are similar to those for automatic backups. In user backup schedules, however, no backup selection list is necessary because users select the objects before they start the backup or archive.

## How to use a specific policy and user schedule

To use a specific policy or schedule for user backups or archives, perform the following on the client:

|                              |                                                                                                                                                                                                             |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| On Microsoft Windows clients | Start the Backup, Archive, and Restore client interface. Click <b>File &gt; NetBackup Client Properties</b> and select the <b>Backups</b> tab. Specify the backup policy and backup schedule.               |
| On NetWare target clients    | Specify the policy and schedule with <code>backup_policy</code> and <code>backup_sched</code> entries in the <code>bp.ini</code> file. (See the NetBackup NetWare user's guide).                            |
| On UNIX clients              | Specify the policy and schedule with <code>BPARCHIVE_POLICY</code> , <code>BPARCHIVE_SCHED</code> , <code>BPBACKUP_POLICY</code> , or <code>BPBACKUP_SCHED</code> options in the <code>bp.conf</code> file. |

## Backup window considerations

The following topics describe the details and the issues regarding how backup windows work.

### How NetBackup determines which schedule to run next

When a policy contains one schedule, the schedule that is selected to run next is straightforward. But when a policy contains multiple schedules, choosing the schedule to run next can become more complicated. The following topics describe how NetBackup determines which schedule to run next if a policy contains multiple schedules.

Essentially, NetBackup performs two tasks to determine which schedule to run next.

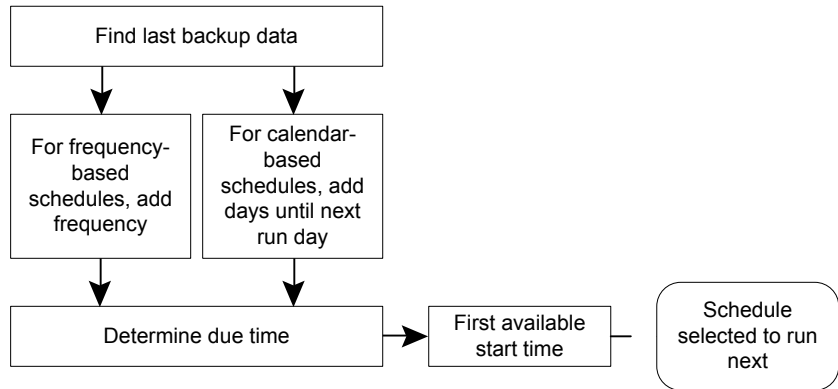
First, NetBackup determines the due time for each schedule. The due time depends on the following:

- The last backup data for each schedule based on comparable schedules.

- The frequency that is added to each schedule to determine which schedule is due next.

Second, NetBackup checks the start time for each schedule. The schedule with the soonest start time runs next. That is, the schedule with the next open window.

**Figure 15-15** Schedule selection overview



## Events that cause the schedules to be recalculated

Any one of the following actions causes NetBackup to recalculate which schedule to run next in a policy:

- When a backup job finishes.
- When a client backup image expires.
- When the administrator changes the policy.  
 NetBackup looks for updated policies every ten minutes. (The length of time NetBackup waits can be configured by changing the **Policy Update Interval** in the **Global Attributes** host properties.) If the policy has just been updated, NetBackup waits an additional minute to be sure that changes are not currently underway.  
 See “[Policy update interval](#)” on page 133.
- When `nbpem` starts.

## Determining the due time for each schedule

The due time is based on the last backup data for the schedule, plus the schedule’s frequency:

$$\text{Last backup data} + \text{frequency} = \text{Due time}$$

The term "last backup data" refers to the schedule that ran most recently among comparable schedules. NetBackup uses the date and time of that schedule to determine the due time for all the schedules that use that schedule as the last backup data.

In some cases, the last backup data for a schedule names the schedule itself. In other cases, the last backup data for a schedule is another comparable schedule.

Comparable schedules are those schedules that fit the following rules:

- Full schedules are compared to other full schedules of the same or longer frequency.
- Cumulative incremental schedules are compared to:
  - Full schedules of the same or longer frequency.
  - Other cumulative incremental schedules of the same or longer frequency.
- Differential incremental schedules are compared to:
  - Full schedules of the same or longer frequency.
  - Cumulative incremental schedules of the same or longer frequency.
  - Other differential incremental schedules of the same or longer frequency.

---

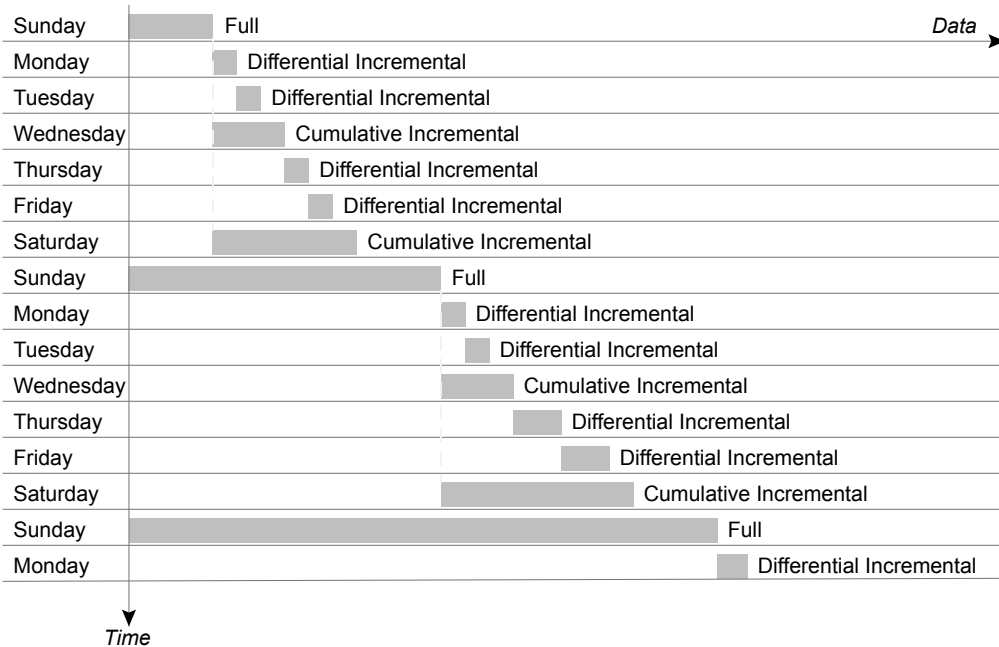
**Note:** To have a longer frequency means that the schedule is configured to run less often.

---

The comparison rules ensure that no schedule is overlooked for consideration, potentially causing a gap in backup coverage.



**Figure 15-16** Schedule coverage



### Scheduling complexities

The following jobs create additional complexities in scheduling:

- **Multistreaming jobs**  
 Each stream is scheduled independently. The data may change in the time between the streamed backups. Two restores that are based on the same backup may not be identical if created from different streams.
- **Synthetic backup jobs**  
 In the case of synthetic backup jobs, NetBackup uses the previous synthetic job as the basis for determining when the next synthetic job should run.

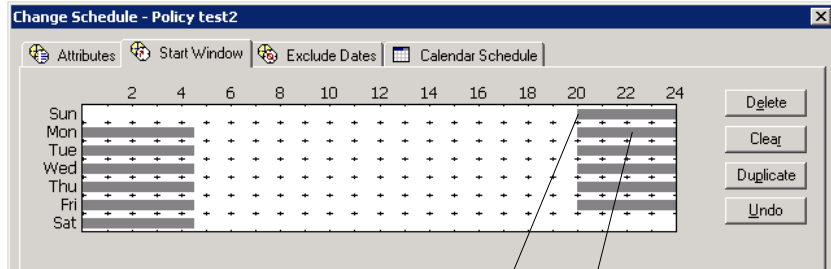
### Windows that span midnight

A backup window may begin in one day and end in another. If a policy is scheduled to run each day, NetBackup does not run the job again immediately after midnight. Instead, even though the window spans into another day, NetBackup considers it to be one window. NetBackup recognizes that the administrator's intention is usually not to have a job run again so soon after the previous backup.

Figure 15-17 shows a window that spans midnight.

If a policy is scheduled to run each day, NetBackup looks to see if another window opens later in the day. If another window is set up to open later, NetBackup waits and runs the job then.

**Figure 15-17** Schedule that spans midnight



The first job begins Sunday.

The job is due Monday as well.  
Instead of running the job again immediately after midnight  
NetBackup looks for a window later in the day and runs it

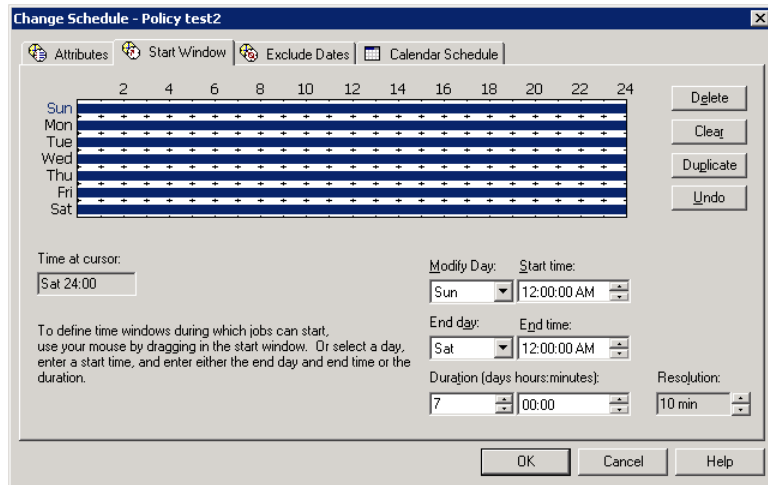
If no other window is scheduled to open later in the day, NetBackup does not wait. If the job has a daily frequency, the job runs again after midnight to meet the daily backup frequency requirement.

## How open schedules affect the different schedule types

A single window can include the entire week. Such a schedule is considered an open schedule because a job may run at any time of day or night.

Figure 15-18 shows an open schedule.

**Figure 15-18** Open schedule



The following topics consider what an open schedule means to calendar-based and frequency-based schedules:

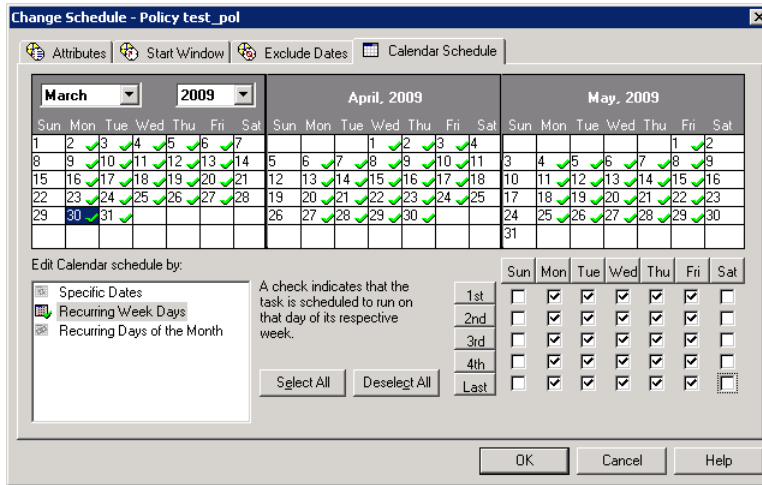
- See [“Open schedules and calendar-based scheduling”](#) on page 531.
- See [“Open schedules and frequency-based scheduling”](#) on page 532.

### Open schedules and calendar-based scheduling

A schedule that is open all day and night allows a job to run whenever the calendar schedule indicates.

Given the calendar schedule in [Figure 15-19](#) and an open schedule, backups should run Monday through Friday.

Figure 15-19 Calendar scheduling and an open schedule



NetBackup determines that a job is due to run by considering when the job last ran successfully and the frequency of the job.

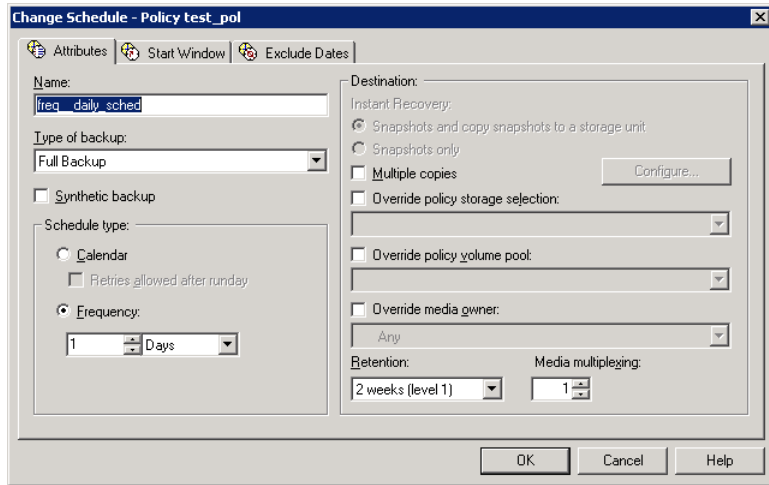
A frequency is not configurable for a calendar-based schedule. NetBackup assumes that an environment requires one backup on each day that is selected on the calendar schedule. Given an open schedule, backups run as soon after midnight as possible to satisfy the daily backup requirement.

## Open schedules and frequency-based scheduling

In a frequency-based schedule, a schedule that is open all day and night allows a job to run as the frequency setting dictates.

Given the frequency-based schedule in [Figure 15-20](#) and an open schedule, backups should run every day of the week, including Saturday and Sunday.

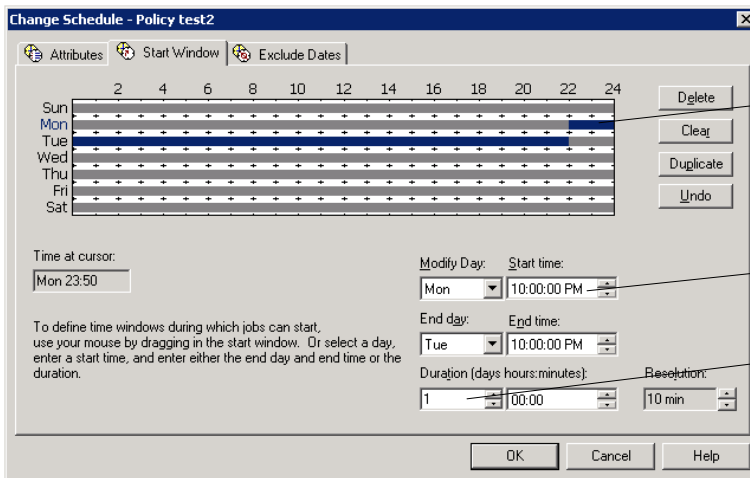
**Figure 15-20** Frequency-based scheduling and an open schedule



With a frequency of one day, NetBackup runs backups at 24-hour intervals based on the start time.

In this example, backups run at 10:00 P.M. nightly because the start time is 10:00 P.M.

**Figure 15-21** Frequency-based schedule with open schedule example



Click on a window to see the and end time of each day.

The start time indicates when can run.

The window has a duration c day is duplicated for each da open schedule.

The following steps describe one method to create an open schedule that runs at 10:00 P.M.:

- Select Sunday as the **Modify Day** and **10:00:00 PM** as the **Start time**.
- Select Monday as the **End Day** and **10:00:00 PM** as the **End time**. The **Duration** is then automatically set to one day.
- Click **Duplicate** to copy this window to each day of the week.

## Runtime considerations

The following topics describe factors that may cause a job to run more frequently than expected, or may prevent a job from meeting its backup frequency requirement.

### Changing a policy causes the policy to run

If the administrator changes or activates a policy, the change prompts NetBackup to run the job as soon as possible. It does not matter if the schedule is calendar-based or frequency-based.

### Window availability

Whether the schedule is calendar-based or frequency-based, a job cannot run if windows are not open on the configured runday.

- For calendar-based schedules, windows must be open on the specific dates, recurring weekdays, or recurring days of the month that the calendar schedule indicates.

---

**Note:** A frequency is not configurable for a calendar-based schedule. For this schedule type, NetBackup assumes a daily backup frequency.

---

- For frequency-based schedules, a daily frequency requires that a window is open each day.

### Backup attempt limit

A **Global Attribute** host property setting determines how many times a failed job can attempt to run. The **Schedule backup attempts** property includes the number of attempts and the time period in which the attempts can take place.

By default, a failed job tries to run two times every 12 hours if an open window is available. Note that this setting supersedes any other frequency requirement and can cause a schedule to skip an open window.

For example, if a job meets the maximum number of job attempts, NetBackup does not try to run the job again during the retry period indicated. It does not attempt, even in an open window and a daily backup frequency has not been met that day.

See “[Schedule backup attempts](#)” on page 132.

## About the Clients tab

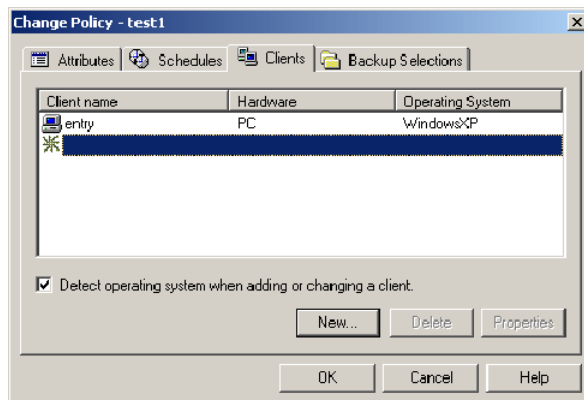
The **Clients** tab contains a list of clients to be backed up (or acted upon) by the selected policy. A client must be included in the list of at least one backup policy to be backed up. Placing a client in more than one backup policy can be useful. For example, place the client name in two policies to back up different sets of files on the client according to different policy rules.

The **Clients** tab does not appear for Vault or Catalog policy types.

## Adding clients to a policy

### To add a client to a policy

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.



- 2 In the right pane, double-click the policy you want to change.
- 3 Select the Clients tab and click **New**.
- 4 In the highlighted field, type the name of the client or browse to find and select the client.
- 5 Press **Enter**.
- 6 Observe the following rules for assigning client names:

- Use a name by which the server knows the client (one that you can use on the server to `ping` or `telnet` to the client).
  - If the client is in multiple policies, use the same name in each policy.
  - If the network configuration has multiple domains, use a more qualified name. For example, use `client1.null.com` or `client1.null` rather than only `client1`.
  - Add only clients with the hardware and the operating systems that this policy supports.
- 7 If the **Detect operating system** check box is not selected, you are prompted to choose the hardware and the operating system.
- Add only clients with the hardware and the operating systems that the policy supports. For example, do not add a Novell NetWare client to an MS-Windows policy.
- 8 Click **OK** to close the Change Policy dialog box or select another tab.
- 9 To add another client, click **New**.

## Browse for Hyper-V virtual machines

- **Enter the VM hostname**

Enter the host name, display name, or GUID of the virtual machine. The format of the host name or display name depends on your system. It may be the fully qualified name or another name, depending on your network configuration and how the name is defined in the guest OS. If NetBackup cannot find the name or GUID you enter, policy validation fails.

If it is checked, uncheck the **Browse and select Virtual Machines** option.
- **Browse and select Virtual Machine**

Click this option to discover Hyper-V servers or cluster nodes (shown in the left pane). You can select virtual machines from a list (in the right pane).

The virtual machine names that are listed may be derived from a cache file. Use of the cache file is faster than rediscovering the machines on the network if your site has a large number of virtual machines. If the virtual machine is turned off but was turned on when the cache file was last created, its name appears in the list.

If the display name of the virtual machine was recently changed in the Hyper-V Manager, note: The virtual machine name that was used for the backup does not change.

If NetBackup cannot obtain the IP address of the virtual machine, the IP address is displayed as NONE.



- **Last Update**

To update the cache file and re-display virtual machines, click the refresh icon to the right of the **Last Update** field. This field shows the date and time of the most recent cache file that contains the names of virtual machines.

## About the Backup Selections tab

The **Backup Selections** tab lists the files, directories, directives, scripts, and the templates that are backed up with this policy. NetBackup uses the same backup selection list for all of the clients that are backed up according to the policy.

Every file on the list does not need to exist on all of the clients. NetBackup backs up the files that it finds that are on the backup selections list. However, each client must contain at least one of the files in the backup selections list, or the client backup fails with a status 71. The policy backup selections list does not apply to user backups or archives. For user backups and archives, users select the objects to back up before they start the operation.

A backup selection list may contain different information based on the policy type, as follows:


- Standard, Exchange, and Lotus Notes policy types list paths and directives. See [“Changing backup selections for standard policies”](#) on page 537.
- Depending on the database type, the backup selection list for database policies contains different types of objects. See [“Changing backup selections for database policies”](#) on page 539.
  - For Exchange and Lotus Notes, the list contains paths and directives.
  - For MS-SQL-Server, Informix-On-BAR, SAP, and Sybase, the list contains the scripts that define and control the database backup, including how the client uses multiple streams.
  - For Oracle and DB2, the list contains scripts and templates.
- Vault policy types list only Vault commands in the backup selections list.

### Changing backup selections for standard policies

Standard, Exchange, and Lotus Notes policy types list paths and directives in the backup selection list.

To add or change backup selections in standard, Exchange, or Lotus Notes policies

- 1 In the NetBackup Administration window, expand **NetBackup Management > Policies**.



- 2 Double-click the policy where you want to change the backup selections list.
- 3 Click the **Backup Selections** tab.
- 4 To add an entry, at the end of the list, click under the last entry in the backup selections list, then click **New**.
- 5 Select a path or directive as follows:
  - Click the folder icon to browse to a remote folder to select a path.
  - Click the directives icon to browse to a directive.  
Click the arrow next to the **Directive** field and select a directive. Click **OK** to include the directive to the backup selections list on the **Selections** tab.  
See [“Backup selections list directives”](#) on page 558.  
See [“Backup selections list directives for multiple data streams”](#) on page 563.  
Paths may contain up to 1023 characters.
- 6 Press **Return** to exit the edit box.
- 7 To rearrange the selections in the backup selection list, do the following:
  - To move an entry, select the entry, then use the **Up** and **Down** buttons, mouse, or keyboard to move an entry.
  - To delete an entry, select the entry and click **Delete**.
  - To rename an entry, select it and click **Rename**. An edit box opens around the entry to modify it.
- 8 Verify that the entries on the backup selections list are accurate  
See [“Verifying the backup selections list”](#) on page 541.

## Changing backup selections for database policies

The type of database determines whether the selections list contains paths, directives, or scripts.

### To change backup selections in database policies

- 1 In the NetBackup Administration window, expand **NetBackup Management > Policies**.
- 2 Double-click the database policy in the Console tree where you want to change the backup selections lists.
- 3 Click the **Backup Selections** tab. To add an entry, at the end of the list, click under the last entry in the backup selections list, then click **New**. An edit box appears.
- 4 Add the backup selections:
  - Enter a script into the text box. Scripts require that you specify the full path. Be sure that the scripts that are listed are installed on each of the clients that are specified on the Client tab.
  - Click the folder icon to browse to a remote folder to specify script paths for a client.
- 5 Click **OK** to add the items to the **Backup Selections** list.

## Changing backup selections for Oracle or DB2 policies

An Oracle backup or XML export policy, or a DB2 backup policy, lists templates and scripts in the backup selection list. The listed templates and scripts are run during manual and automatic backups in the order in which they appear in the backup selection list.

### To change backup selections for Oracle or DB2 policies

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Double-click the policy where you want to add or change templates or scripts.
- 3 Click the **Backup Selections** tab.
- 4 To add an entry, click **New**. An edit box appears.
- 5 Add the backup selections as follows:

- |           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Templates | <ul style="list-style-type: none"><li>■ Type only the template file name, for example, <code>weekly_full_backup.tpl</code>, or,</li><li>■ Click the <b>Template</b> button. For Oracle policies, choose Oracle_RMAN or Oracle_XML_Export from the <b>Template set</b> list.<br/>Then, choose a template name in that set from the <b>Template</b> list. You can also add all available templates by clicking <b>Add all templates to the backup selections list</b>.</li></ul> |
| Scripts   | <ul style="list-style-type: none"><li>■ Use the full path of the client and the file name for the script you want to include.</li><li>■ Click <b>Browse</b> to locate the script in the client Browse window. Click <b>OK</b> to add the selection. Be sure that the shell scripts that are listed are installed on each of the clients that are specified in the <b>Clients</b> tab.</li></ul>                                                                                |

6 To change the order of the backup selections, select one and click **Up** or **Down**.

7 Click **OK** to add the selection to the selection list.

See “[Changing backup selections for standard policies](#)” on page 537.

See “[Changing backup selections for database policies](#)” on page 539.

See “[Changing backup selections for Oracle or DB2 policies](#)” on page 539.

See “[Verifying the backup selections list](#)” on page 541.

## Reducing backup time

Selection list entries are processed serially for each client and in the order that they appear in the backup selections. A client can be added to multiple policies, to divide the client’s files among the different backup selections lists. Multiple policies can reduce the backup time for that client because the files can be backed up in parallel.

Multiple clients can be backed up in parallel in the following situations:

- Multiple storage devices are available (or if the policies are multiplexed).
- The **Maximum Jobs per Client** Global host property, and the **Limit Jobs per Policy** policy attributes are set to allow it.

---

**Note:** Understand disk and controller input and output limitations before configuring including a client in multiple policies. For example, if two file systems overload the client when backed up in parallel, place both file systems in the same policy. Schedule the file systems at different times or set **MaximumJobs per Client** to 1.

---

Another method to reduce backup time is to select **Allow Multiple Data Streams** for a policy. Then, add `NEW_STREAMS` directives to the backup selections list.

For example:

```
NEW_STREAM
file_a
file_b
file_c
NEW_STREAM
file_d
file_e
file_f
```

The example produces two concurrent data streams. The first data string contains `file_a`, `file_b`, and `file_c`. The second data stream contains `file_d`, `file_e`, and `file_f`.

See [“Allow multiple data streams attribute”](#) on page 486.

---

**Note:** For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can cause longer backup times. The tape heads must move back and forth between the tracks that contain files for the respective streams.

---

A directive instructs NetBackup to perform specific actions to process the files in the backup selections list.

## Verifying the backup selections list

Verify a backup selections list to make sure that the file paths are correct for the clients in the policy.

### To verify a backup selections list

- 1 Check all entries against the file path rules for the clients in the policy. If the list includes directives, verify that the syntax for the directives is correct.
- 2 Run a set of backups. Then, check the Problems report or the All Log Entries report for warning messages. The backup status code does not always indicate errors on the backup selection list. NetBackup does not require all paths in the backup selections list to be present on all clients, so the error is not reflected.

See “[Report types](#)” on page 712.

- 3 Run the `check_coverage` script to create a File System Backup Coverage Report. The script is located in `install_path\NetBackup\bin\goodies`.

The script can reveal mistakes in the selections list that make it impossible for NetBackup to find the files. This results in files being skipped in the backup.

If a path is not found, NetBackup logs a trivial (TRV) or warning (WRN) message. However, the same job can end with a backup status code of 0 (successful). Usually, to report files missing from the backup selections list is not helpful, since not all files are expected to be present on every client. However, check the logs or use the `check_coverage` script to ensure that files are not missed due to bad or missing backup selections list entries.

### Example log messages

The following examples show the log message that appear when files on a client are not found. For information on `check_coverage`, see the comments in the script.

Assume that the backup selections list contains the path `c:\worklist` that is not present on all clients. NetBackup backs up `C:\worklist` on the clients where it exists.

For other clients, the Problems report or the All Log Entries report shows a message similar to the following:

```
9/1/09 8:28:17 AM carrot freddie Info from client freddie: TRV
- object not found for file system backup: C:\worklist
```

This message occurs if `c:\worklist` is not the correct path name. For example, the directory name is `c:\worklists`, but `c:\worklist` was typed.

---

**Note:** If the paths seem correct and the message appears, ensure that no trailing spaces appear in the paths.

---

## Path rules for Microsoft Windows file backups

Microsoft Windows path conventions, UNIX path conventions, or a combination of the two can be used in the backup selections list.

The following conventions can be used in the backup selections list:

Microsoft Windows  
conventions

- Enter one path per line.
- Begin all paths with the drive letter followed by a colon (:) and a backslash (\). The drive letter is case-insensitive, however, the path is case sensitive. For example, `c:\Worklists\Admin\`  
To specify an entire volume, append a backslash (\) to the entry to ensure that all data is protected on that volume:  
Correct entry:

```
c:\
```

Incorrect entry:

```
c:
```

- Precede each component in the path with a backslash.  
If the last component in the path is a directory, follow it with a backslash (\) as well. The trailing backslash is not required but serves as a reminder that the path is to a directory instead of a file: `c:\users\net1\`  
If the last component is a file, include the file extension and omit the backslash from the end of the name: `c:\special\list.txt`
- Allowable wildcard characters are the same as those allowed in Windows paths: `* ?`  
See [“Wildcards in NetBackup”](#) on page 720.
- To back up all local drives except for those that use removable media, specify:

```
:\
```

Or

```
*:\ or ALL_LOCAL_DRIVES
```

The following drives are not backed up: Floppy disks, CD-ROMs, and any drives that are located on remote systems but mounted on a system through the network.

- By default, NetBackup does not back up some files.  
See [“Files that are excluded from backups by default”](#) on page 567.
- Exclude specific files from backups by creating an exclusion list on the client.  
See [“Excluding files from automatic backups”](#) on page 568.  
The following backup selection list uses Microsoft Windows conventions:

```
c:\
d:\workfiles\
e:\Special\status
c:\tests*.exe
```



UNIX conventions that are permitted on Windows

UNIX conventions are similar to those for Microsoft Windows, with the following exceptions:

- Begin each line with a forward slash (/).
- Omit the colon (:) after the drive letter.
- Specify / to back up all local drives except for those that are removable.

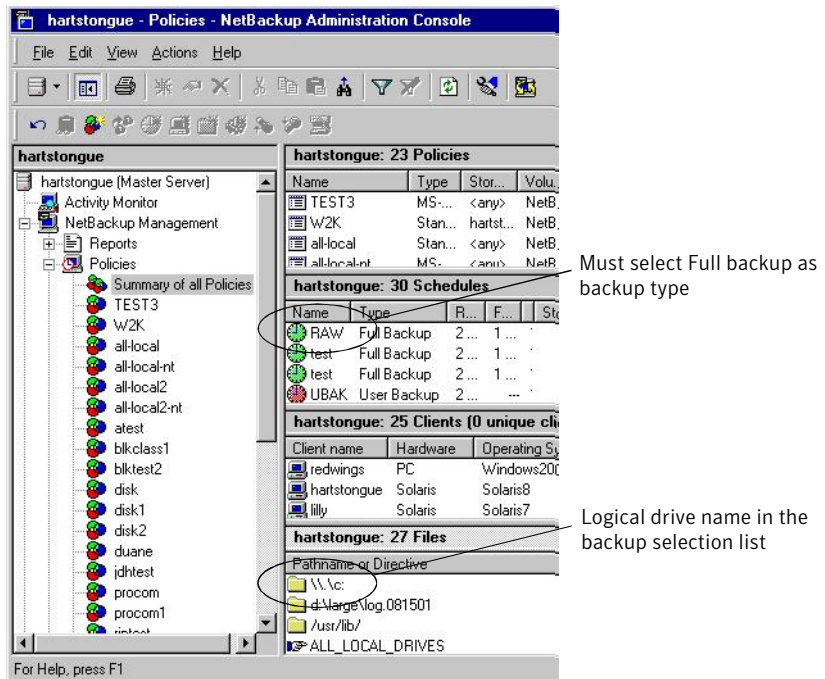
The following example uses UNIX conventions:

```
/c/
/d/workfiles/
/e/Special/status
/c/tests/*.exe
```

## Path rules for Windows disk image (Raw) backups

On Windows clients, you can back up a logical disk drive as a disk image. That is, NetBackup backs up the entire logical drive on a bit-by-bit basis rather than by directories and files.

**Figure 15-22** Disk image backups



Select **Full backup** as the backup type to perform a disk image backup.

To specify a disk image backup, add the logical name for the drive to the policy backup selection list. The format in the following example backs up drive C.

```
\\.\c:
```

Disk images can be included in the same backup selection list with other backups:

```
\\.\c:
```

```
d:\workfiles\
```

```
e:\Special\status
```

```
HKEY_LOCAL_MACHINE:\
```

To restore the backup, the user first chooses **Select for restore > Restore from Normal backup**.

When the backups are listed, the disk image appears as a file with the same name that was specified in the backup selection list. In this example:

```
\\.\c:
```

Select the disk image source, then enter the destination in the following format:

```
\\.\drive:
```

Where drive is the location where the partition is to be restored.

Notes on disk image backups:

- NetBackup first attempts to use Windows Open File Backup methods. If that fails, NetBackup locks the logical drive, which ensures that no changes occur during the backup. If there are open files on the logical drive, a disk image backup is not performed.
- Before a disk image is backed up or restored, all applications that have a file opened on the partition should be shut down. If the applications are not shut down, the operation fails. Examples of such applications are Windows Explorer or Norton AntiVirus.  
Ensure that no active COW (Copy On Write) snapshots are in progress. If there is an active COW snapshot, the snapshot process itself has a handle open to the volume.
- NetBackup does not support raw partition backups on unformatted partitions.
- If the volume is configured to contain a paging file (`pagefile.sys`), a raw partition backup of that volume may fail. In order for a raw partition backup of that volume to succeed, the volume may need to be reconfigured so as not to contain a paging file. The raw partition backup of the volume may work without reconfiguration if a snapshot can successfully be taken of that volume.

## Path rules for Windows registry backup

Consider the following items when configuring a Windows registry backup:

- Back up for disaster recovery

To ensure a successful recovery in case of a disk failure, always back up the entire registry. That is, back up the directory that contains the entire registry. On most Windows systems, this directory is located at:

```
%systemroot%\system32\config
```

Where `%systemroot%` is the directory where Windows is installed.

---

**Note:** To recover the registry, do not include individual registry files or HKEY entries in the selection list that's used to back up the entire registry. If you use a NetBackup exclude list for a client, do not exclude any registry files from your backups.

---

To restore the registry in the case of a disk failure, see the Disaster Recovery chapter in the *NetBackup Troubleshooting Guide*.

- Back up individual HKEYs (do not use for disaster recovery)

Do not include HKEY entries in the same policy backup selection list that is used to back up the entire registry. However, to restore individual keys within the registry, create a separate policy, then specify the specific HKEYs in the backup selection list for that policy.

The following is an example HKEY entry for a policy backup selection list:

```
HKEY_LOCAL_MACHINE:\
```

Remember, you cannot perform a disaster recovery by restoring HKEYs. In addition, backups and restores are slower than if the entire registry was backed up.

## Hard links to files (NTFS volumes or UNIX)

A hard link is a directory entry for a file. Every file can be considered to have at least one hard link. On NTFS volumes or on UNIX systems, each file can have multiple hard links. Therefore, a single file can appear in many directories (or even in the same directory with different names). A Volume serial number (VSN) and a File Index indicates the actual file, unique on the volume. Collectively, the VSN and File Index are referred to as the file ID.

During a backup, if the backup selection list includes hard-linked files, the data is backed up only once. NetBackup uses the first file name reference that is found in the directory structure. If a subsequent file name reference is found, the reference is backed up as a link to the name of the first file. To back up subsequent


references means that only one backup copy of the data is created, regardless of the number of multiple hard links.

If all hard-link references are restored, the hard-linked files continue to point to the same ID as the other files to which they are linked. However, if all the hard links are not restored, you can encounter anomalies as shown in the following examples.

**Example 1**


Assume that three hard links point to the same data. During a backup of Link2 and Link3, Link2 is encountered first and backed up. Then Link3 is backed up as a link to Link2. The three files are all hard-linked to the same data.

**Figure 15-23** Example of hard links to the same data




The original copies of Link2 and Link3 are backed up to tape, then deleted. Only Link1 is left on the disk.

**Figure 15-24** Example of hard links backed up to tape and disk



During a subsequent restore, Link2 and Link3 are restored. The restored files, however, do not point to the same file ID as Link1. Instead, they are assigned a new file ID or inode number and the data is written to a new place on the disk. The data in the new location is an exact copy of what is in Link1. The duplication occurs because the backup does not associate Link2 and L3 with Link1.

**Figure 15-25** Example of restored hard links



### Example 2

Assume that you attempt to restore only Link3. Here, NetBackup cannot link Link3 to Link2 because Link2 does not exist. The restore can complete only if it can link to Link2. A secondary restore request to the NetBackup server automatically restores Link2, which contains the data. Link2 can now be successfully restored.

## Pathname rules for UNIX clients

The following items describe the conventions that are used to specify backups for UNIX clients:

- Enter one pathname per line. NetBackup supports a maximum path length of 1023 characters for UNIX clients.
- Begin all pathnames with a forward slash (/).
- The following wildcard characters are allowed:

```
*
?
[]
{ }
```

See [“Wildcards in NetBackup”](#) on page 720.

- If a backup selection list entry contains trailing spaces and a matching entry is not found, NetBackup deletes the spaces and checks again. If a match is not found, NetBackup skips the entry and logs a message in the Problems report or the All Log Entries report:

```
TRV - cannot process path pathname: No such file or directory.
Skipping TRV - Found no matching file system for pathname
```

See [“Report types”](#) on page 712.

See [“Symbolic links and directory junctions”](#) on page 550.

See [“About hard links to directories”](#) on page 551.

See [“About hard links to files”](#) on page 551.

See [“UNIX raw partitions”](#) on page 551.

See [“Backing up and restoring extended attribute files and named data streams”](#) on page 554.

## Considerations for UNIX pathnames

Consider the following items when indicating UNIX pathnames to specify backups for UNIX clients:

- Pathnames that cross mount points or that the client mounts through NFS can affect the backup configuration. Read about the **Follow NFS** and **Cross mount points** attributes before you create a backup selection list.  
See [“Follow NFS attribute”](#) on page 476.  
See [“Cross mount points attribute”](#) on page 478.
- NetBackup can back up operating system, kernel, and boot files. However, NetBackup cannot, create bootable tapes. Consult your system documentation to create a bootable tape.
- By default, NetBackup does not back up all files.  
See [“Files that are excluded from backups by default”](#) on page 567.
- Exclude specific files from backups by creating an exclusion list on the client.  
See [“Excluding files from automatic backups”](#) on page 568.
- On Solaris, HP-UX, AIX, Linux Red Hat 4 (and later), Linux SUSE SLE 9 (and later), and supported Mac platforms, NetBackup backs up Access Control Lists (ACLs).
- NetBackup can back up and restore Sun PC NetLink files.
- By default, NetBackup backs up and restores Solaris 9 and 10 extended attribute files. The FlashBackup single file restore program (`sfr`) does not restore extended attribute files.
- By default, NetBackup backs up and restores named data streams for VxFS 4.0 (Solaris SPARC) and VxFS 5.0 (Solaris, HP, Linux, and AIX). The FlashBackup single file restore program (`sfr`) does not restore extended attribute files.  
See [“Backing up and restoring extended attribute files and named data streams”](#) on page 554.
- On Hewlett-Packard and Solaris SPARC platforms, NetBackup backs up VxFS extent attributes.

## Symbolic links and directory junctions

Keep in mind the following points when including symbolic links and directory junctions in a backup selections list:

- **Symbolic links**  
NetBackup backs up the symbolic link object and does not attempt to follow the link to back up what it may point to. To achieve a backup of the target of the symbolic link, include that target in the file list.  
Restoring the symbolic link object restores only the object and not the data to which it may point. To restore the target data, select it from the backup image.
- **Directory junctions**

NetBackup backs up the directory junction object and does not attempt to traverse into the directory to which it may point. To achieve a backup of the target of the directory junction, include that target in the file list.

Restoring the directory junction link object restores only the object and not the data to which it may point. To restore the target data, select it from the backup image.

## About hard links to directories

On most UNIX systems, only the root user can create a hard link to a directory. Some systems do not permit hard links, and many vendors recommend that these links be avoided. NetBackup does not back up and restore hard-linked directories in the same manner as files.

The differences are as follows:

- During a backup, if NetBackup encounters hard-linked directories, the directories are backed up once for each hard link.
- During a restore, NetBackup restores multiple copies of the hard-linked directory contents if the directories do not already exist on the disk. If the directories exist on disk, NetBackup restores the contents multiple times to the same disk location.

## About hard links to files

A hard link differs from a symbolic link in that a hard link is not a pointer to another file. A hard link is two directory entries that point to the same inode number.

If the backup selection list includes hard-linked files, the data is backed up only once during a backup. NetBackup uses the first file name reference that is found in the directory structure. If a subsequent file name reference is found, it is backed up as a link to the name of the first file. Backup up only the link means that only one backup copy of the data is created, regardless of the number of hard links. Any hard link to the data works.

See [“Hard links to files \(NTFS volumes or UNIX\)”](#) on page 547.

## UNIX raw partitions

Save a copy of the partition table before performing raw partition backups so that you retain the copy for reference before to a restore. To restore the raw partition, a device file must exist and the partition must be the same size as when it was backed up. Otherwise, the results of the restore are unpredictable.

Consider the following items when creating UNIX raw partition backups:

- Use raw partition backups only if you can ensure that the files have not changed in any way during the backup. Or, in the case of a database, if you can restore the database to a consistent state by using transaction log files.
- Do not perform backup archives of raw partitions on any client. An archive backs up the raw partition, then deletes the device file that is associated with the raw partition. The file system does not recover the space that the raw partition uses.
- Before backing up file systems as raw partitions, unmount the file system. Unmounting the file system allows buffered changes to be written to the disk. Also, it prevents the possibility of any changes in the file system during the backup. Use the `bpstart_notify` and the `bpend_notify` scripts to unmount and remount the backed-up file systems.
- The **Cross mount points** policy attribute has no effect on raw partitions. If the root partition is backed up as a raw partition and contains mount points to other systems, the file systems are not backed up. The other file systems are not backed up, even with **Cross mount points** selected.  
See “[Cross mount points attribute](#)” on page 478.  
The same is true for the **Follow NFS** policy attribute. NFS file systems that are mounted in a raw partition are not backed up. Nor can you back up raw partitions from other machines by using NFS mounts to access the raw partitions. The devices are not accessible on other machines through NFS.  
See “[Follow NFS attribute](#)” on page 476.
- Specify the logical partition names for any disks that disk volume managers manage. (For example, Veritas Volume Manager (VxVM).)
- For clients in a FlashBackup policy, refer to the *NetBackup Snapshot Client Administrator’s Guide* for the differences between Standard and FlashBackup policies.

If there are no file systems to back up and the disks are used in raw mode, back up the disk partitions as raw partitions. For example, databases are sometimes used in raw mode. Use `bpstart_notify` and `bpend_notify` scripts to provide the necessary pre-processing and post-processing of databases when they are backed up as raw partitions.

You can also perform a raw partition backup of a disk partition that is used for file systems. A disadvantage of this method is that you must restore the entire partition to recover a single file (unless FlashBackup is in use). To avoid overwriting the entire partition, use the redirected restore feature to restore the raw partition to another raw partition of the same size. Then, copy individual files to the original file system.



Raw partition backups are also useful for backing up entire disks. Since the file system overhead is bypassed, a raw partition backup is usually faster. The size of the raw partition backup is the size of the entire disk, regardless of whether the entire disk is used.

To specify a UNIX raw partition in the policy backup selection list, enter the full path name of the device file.

For example, on a Solaris system enter:

```
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

---

**Note:** Do not specify wildcards (such as `/dev/rsd*`) in pathnames for raw partition backups. Doing so can prevent the successful restore of entire devices if there is overlap between the memory partitions for different device files.

---

You can include raw partitions in the same backup selection list as other backups. For example:

```
/home
/usr
/etc
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

---

**Note:** NetBackup does not distinguish between full and incremental backups when it backs up a raw partition. The entire partition is backed up in both cases.

---

Raw partition backups occur only if the absolute pathname in the backup selection list is a block or character special device file. You can specify either block or character special device files. Character special device files are often faster because character devices avoid the use of the buffer cache for accessed disk data. Test both a block and character special device file to ensure the optimum backup speed for your platform.

Ensure that you specify the actual block-device or character-device files. Sometimes these are links to the actual device files. If a link is specified, only the link is backed up. If the device files are reached while backing up `/dev`, NetBackup backs up only the inode files for the device, not the device itself.

To perform a raw partition backup, select `Full backup` for the **Type of Backup** from the **Schedules** tab. Any other backup type does not work for backing up raw partitions.

See “[Type of backup attribute](#)” on page 492.

## Backing up and restoring extended attribute files and named data streams

NetBackup can back up and restore the following file attributes:

- Extended attribute files of the Solaris UNIX file system (UFS) and temporary file system (TMPFS)
- Named data streams of the VxFS file system

NetBackup backs up extended attribute files and named data streams as part of normal file system backups.

Extended attribute files and named data streams are normal files contained in a hidden attribute directory that relate to a particular base file. The hidden directory is stored within the file system, but can be accessed only by the base file to which it is related. To view which files have extended attributes on Solaris 9 (or greater) systems, enter: `ls -@`

Neither extended attribute files nor named data streams can be backed up or restored individually. Rather, the files are backed up and restored all at once along with the base file.

The presence of a large number of extended attribute files or named data streams can cause some degradation in backup and restore speed. The speed is affected since the base file and all associated files are backed up.

The speed is especially affected in the case of incremental backups, during which NetBackup checks the `mtime` or `ctime` of each file individually.

To back up or restore named data streams and extended attributes, the client, media server, and master server must run the following versions:

- NetBackup clients
  - HP 11.23 running VxFS 4.1 or greater.

---

**Note:** Access Control Lists (ACLs) are not backed up unless running VxFS 5.0 or greater.

---

- AIX running VxFS 4.0 or greater.

---

**Note:** ACLs are not backed up unless running VxFS 5.0 or greater.

---

- Solaris 10 running VxFS 5.0 or greater
- Solaris SPARC 9, 10 running VxFS 4.0 or greater

- Linux running VxFS 5.0 or greater.
- A NetBackup master server
  - A NetBackup master server of any version can back up and restore named data streams and Solaris extended attributes.

Restored attribute files and named data streams can replace existing files if **Overwrite existing files** is selected in the **Backup, Archive, and Restore** client interface.

If an attempt is made to restore the following items, an error message appears in the **Restore Monitor** to inform the user that the extended attributes or named data streams are not restored.

- The extended attribute files to any non-Solaris 9 client (or greater), or
- Named data streams to any non-VxFS 4.0 client,

NetBackup then continues with the restore job.

To disable the restore of extended attribute files and named data streams, add an empty file to the client. Name the file `IGNORE_XATTR` and place it in the following directory:

```
/usr/obj/RS6000/netbackup/
```

The addition affects only Solaris 9 or VxFS 4.0 clients.

File `IGNORE_XATTR` was formerly known as `IGNORE_XATTR_SOLARIS`.

---

**Note:** Extended attributes and named data streams cannot be compressed.

---

## About the path rules for NetWare NonTarget clients

For NetWare systems that are running the NonTarget version of NetBackup client software, specify the paths in the following form:

```
/SMDR/TSA/TS/resources/directory/file
```

Where:

- **SMDR (Storage Management Data Requestor)** is the name of the NetWare file server that is running the SMDR.NLM that is used for backups. (NLM means NetWare-loadable module.)
- **TSA (Target Service Agent)** is a NetWare software module that prepares the data for backup or restore by the SMDR. The type of TSA that is used depends on the data. For example, NetWare file systems and DOS workstations each have TSAs.

- **TS** is the Target Service, which is the NetWare entity that contains the data that the selected TSA handles. For example, in the case of the DOS TSA (`tsasms.com`) it is a DOS workstation. In the case of a NetWare file system TSA, it is the system with the NetWare file systems to be backed up.
- **resources** are the specific resources on the target service. For example, it can be NetWare file systems such as BINDERY, SYS, and USER.
- *directory/file* is the directory and file that are in the resource (if it is a path to a specific file).

Observe the following rules for paths:

- Give the server access to each path or the scheduled backup fails. To provide this access, use the **Allowed scheduled access** command on the **Backup** menu in the NetBackup interface on the NetWare client.

For more information, see the *NetBackup for Novell NetWare Client Administrator's Guide*.

- Enter one path per line.
- Begin all paths with a forward slash (/).
- Precede each component in the path with a forward slash.  
If the last component in the path is a directory, follow it with a forward slash (/). The trailing slash is not required but is a reminder that the path points to a directory instead of a file.

```
/client1/client1.NetWare File System/client1/SYS/DOC/
```

If the last component is a file, include the file extension and omit the slash from the end of the name.

```
/client1/client1.NetWare File System/client1/SYS/DOC/TEST.TXT
```

- All components in a path must show uppercase and lowercase letters as they appear in the actual path on the client.
- Wildcard usage is the same as for Windows clients.  
See “[Wildcards in NetBackup](#)” on page 720.
- To back up all NetBackup for NetWare clients that are in the policy, enter only one forward slash (/) on a line:

```
/
```

- To back up an entire NetBackup for NetWare client, enter a forward slash (/) followed by the client name and another forward slash:

```
/client1/
```

The following example backs up SYS, BINDERY, and USER file systems under the file system TSA on a client that is named client1:

```
/client1/client1.NetWare File System/client1/SYS/
/client1/client1.NetWare File System/client1/BINDERY/
/client1/client1.NetWare File System/client1/USER/
```

Note that the **Allowed scheduled access** command on the NetBackup NetWare client **Backup** menu must also specify access to these paths.

See the *NetBackup for Novell NetWare Client Administrator's Guide*.

## Path rules for NetWare Target clients

For NetWare clients that are running the Target version of NetBackup client software, use the following format for the paths:

```
/target/
```

Where *target* is the name of a target that is defined on the NetBackup for NetWare client.

For more information, see the *NetBackup Administrator's Guide for Novell NetWare Clients*.

- Enter one target per line.
- Begin all target names with a forward slash (/).
- All target names must be in uppercase.
- Wildcard usage is the same as for Windows clients.  
See “[Wildcards in NetBackup](#)” on page 720.

The following example backs up the targets: NETWARE, SYSTEM, and BINDERY:

```
/NETWARE/
/SYSTEM/
/BINDERY/
```

## Path rules for clients that run extension products

Path rules for the NetBackup clients that are running separately-priced options are covered in the NetBackup guide for the product. (For example, Snapshot Client or NetBackup for MS-Exchange.)

## Backup selections list directives

The backup selections list can contain the directives that signal NetBackup to perform specific actions when it processes the files in the selections list.

The available directives depend on the policy type and whether the **Allow multiple data streams** attribute is enabled for the policy. The following example is a backup selections list that contains the `NEW_STREAM` directive. The example is from an MS-Windows policy with **Allow multiple data streams** enabled.

```
NEW_STREAM
D:\Program Files
NEW_STREAM
C:\Winnt
```

### ALL\_LOCAL\_DRIVES directive

Use the `ALL_LOCAL_DRIVES` directive to back up all local drives except for those drives that use removable media.

The `ALL_LOCAL_DRIVES` directive applies to the following policy types:

- Standard (except for NetWare target clients)
- MS-Windows
- NetWare (NonTarget clients only)

However, `ALL_LOCAL_DRIVES` is not allowed for NetWare policy types if **Allow multiple data streams** is also used.

See [“ALL\\_LOCAL\\_DRIVES directive and multiple data streams”](#) on page 561.

See [“Files that are excluded from backups by default”](#) on page 567.

### System\_State directive

The `System_State:\` directive is needed for the operating system versions which do not support Shadow Copy Components, such as the 32-bit version of Windows 2003 XP.

Windows 2003 Server computers recognize the `System_State:\` directive and behave as if following the `Shadow Copy Components:\` directive. A message informs the user that this directive translation occurred.

The `System_State:\` directive creates a backup for critical system-related components. The exact set of system components that are backed up depends on the operating system version and system configuration.

The list of items that are backed up can include the following:

- Active Directory
- COM+ Class Database
- Cluster Database
- IIS Database
- Registry
- Boot Files and protected files
- SYSVOL
- Certificate Server

The files that comprise the registry can be found in the following location:

```
%SystemRoot%\SYSTEM32\Config
```

At a minimum, the following files are backed up as part of the registry:

- DEFAULT
- SAM
- SOFTWARE
- SECURITY
- SYSTEM

## Shadow Copy Components:\ directive

The `Shadow Copy Components:\ directive` specifies that all of the Volume Shadow Copy component writers get backed up.

The `Shadow Copy Components:\ directive` affects the backups of the following clients:

- Windows 2003 Server computers that use the Volume Shadow Copy components.
- IA64 systems with EFI System partitions.  
 In the policies that back up clients on IA64 platforms, use the `Shadow Copy components:\ directive` instead of the `System_State:\ directive`. The `Shadow Copy components:\ directive` includes System State components and the EFI System partition automatically in the backup.

Since the Shadow Copy Components contain System State information, the Shadow Copy Components need to be backed up by a full backup.

The Volume Shadow Copy components include the following:

- System State writers, which can include:
  - System files
  - COM+ Class Registration Database
  - SYSVOL
  - Active Directory
  - Cluster quorum
  - Certificate Services
  - Registry
  - Internet Information Services
- System Service writers, which can include:
  - Removable Storage Manager
  - Event logs
  - Windows Internet Name Service
  - Windows Management Instrumentation
  - Remote Storage
  - Dynamic Host Configuration Protocol
  - Terminal Server Licensing
  - Background Intelligent Transfer Service
- **User Data** writers, which include any items that the computer does not require to operate. For example, Active Directory Application Mode.
- **Other Data** writers, a category that is intended for future NetBackup releases.

## Directives for multiple data streams

If the **Allow multiple data streams** general attribute is set for a policy, the following directives can be used in the backup selections list:

- NEW\_STREAM
- ALL\_LOCAL\_DRIVES
- UNSET
- UNSET\_ALL

See [“Backup selections list directives for multiple data streams”](#) on page 563.

See [“Allow multiple data streams attribute”](#) on page 486.



See “[Backup selections list directives for multiple data streams](#)” on page 563.

## Directives for specific policy types

Some directives apply only to specific policy types and can appear only in backup selections lists for those policies. NetBackup passes policy-specific directives to the clients along with the backup selections list. The clients then perform the appropriate action according to the directive.

---

**Note:** Include policy-specific directives only in backup selections lists for the policies that support the directives or errors can occur.

---

The following policy types have their own backup selections list directives:

- AFS
- FlashBackup
- NDMP
- Lotus-Notes
- MS-Exchange-Server

For example, the following directives can appear only in the backup selections list of an AFS policy:

```
CREATE_BACKUP_VOLUMES
SKIP_SMALL_VOLUMES
```

Except for AFS, these policy types can be used when their associated separately-priced option is installed.

For information about AFS directives, see the *NetBackup Administrator's Guide, Volume II*.

For information on other policy types and associated backup selections list directives, see the NetBackup guide for the option.

## ALL\_LOCAL\_DRIVES directive and multiple data streams

The `ALL_LOCAL_DRIVES` directive applies only to Standard (except for NetWare target clients), MS-Windows, and NetWare policies. If this directive is used, this directive must be the only entry in the backup selections list for the policy. That is, no other files or directives can be listed.

The action that `ALL_LOCAL_DRIVES` causes depends on whether **Allow multiple data streams** is enabled for the policy.

See [“Allow multiple data streams attribute”](#) on page 486.

- If the **Allow multiple data streams** option is enabled, the `ALL_LOCAL_DRIVES` directive applies only to Standard (except for NetWare clients) or MS-Windows policy types. NetBackup backs up the entire client, then splits the data from each drive (Windows) or file system (UNIX) into its own backup stream. NetBackup periodically preprocesses the client to make necessary changes to the streams.

See [“Setting the preprocess interval for auto-discovery”](#) on page 565.

- If the **Allow multiple data streams** option is not enabled, NetBackup backs up the entire client and includes all drives and file systems in the same stream.

---

**Caution:** Do not select **Cross mount points** for policies where you use the `ALL_LOCAL_DRIVES` directive.

---

## ALL\_LOCAL\_DRIVES example 1

Assume that **Allow multiple data streams** is enabled in the auto-discovery mode. Assume that the client is a Windows system with two drive volumes, C:\ and D:\. The backup selections list contains:

```
ALL_LOCAL_DRIVES
```

For this backup selections list, NetBackup generates the following:

- One stream for C:\
- One stream for D:\

For a UNIX client, NetBackup generates a stream for each file system.

`SYSTEM_STATE` is also backed up because `SYSTEM_STATE` is included in the `ALL_LOCAL_DRIVES` directive.

See [“Allow multiple data streams attribute”](#) on page 486.

## ALL\_LOCAL\_DRIVES example 2

Assume that **Allow multiple data streams** is not enabled. Assume that the client is a Windows system with two drive volumes, C:\ and D:\. The backup selections list contains:

```
ALL_LOCAL_DRIVES
```

Here, NetBackup backs up the entire client in one data stream that contains the data from both C:\ and D:\.

SYSTEM\_STATE is also backed up because SYSTEM\_STATE is included in the ALL\_LOCAL\_DRIVES directive.

See [“Allow multiple data streams attribute”](#) on page 486.

## Backup selections list directives for multiple data streams

If the **Allow multiple data streams** attribute is enabled for the policy, the following directives can be used to control how NetBackup creates backup streams:

- NEW\_STREAM
- ALL\_LOCAL\_DRIVES
- UNSET and UNSET\_ALL

---

**Note:** For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The heads must move back and forth between the tracks that contain files for the respective streams.

---

### About the NEW\_STREAM directive and multiple data streams

The NEW\_STREAM directive is recognized only if **Allow multiple data streams** is set for the policy. NEW\_STREAM directives are ignored if **Allow multiple data streams** is not set.

If this directive is used in a backup selections list, the first instance of it must be on the first line. If it appears on the first line, it can also appear elsewhere in the list.

The presence of NEW\_STREAM on the first line of the backup selections list determines how the backup is performed in the following modes: in administrator-defined streaming or in the auto-discovery streaming.

### About the Administrator-defined streaming mode

If NEW\_STREAM is the first line of the backup selections list, the backup is performed in the administrator-defined streaming mode.

The following actions occur:

- The backup splits into a separate stream at each point in the backup selections list where the **NEW\_STREAM** directive occurs.
- All file paths between **NEW\_STREAM** directives belong to the same stream.

- The start of a new stream (a **NEW\_STREAM** directive) defines the end of the previous stream.
- The last stream in the backup selections list is terminated by the end of the backup selections list.

In the following examples, assume that each stream is from a separate physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times. The backup time is longer if the heads must move back and forth between the tracks that contain files for the respective streams.

For example, consider the following backup selections list:

```
NEW_STREAM
D:\Program Files
C:\Winnt
NEW_STREAM
C:\users
D:\DataFiles
```

This backup selections list contains two data streams:

- The **NEW\_STREAM** directive at the top of the list starts administrator-defined streaming and the first data stream. This stream backs up `D:\Program Files` and `C:\Winnt`.
- The second **NEW\_STREAM** starts a second data stream that backs up `C:\users` and `D:\DataFiles`.

If a backup selections list entry is added to a stream, the entry is not backed up until the schedule is due for the policy. If the next backup due is an incremental, only the files that changed are backed up. To ensure that a new entry gets a full backup the first time, add it to a new stream. NetBackup performs a full backup of new streams that are added to the backup selections list.

In the previous example, assume that you add the following:

```
D:\Utilities

after

D:\Datafiles
```

If an incremental backup is due that night, only changed files in `D:\Utilities` are backed up. Add a **NEW\_STREAM** directive before `D:\Utilities`, to perform a full backup of all files in `D:\Utilities`, regardless of when the files were last changed.

## About the auto-discovery streaming mode

The auto-discovery streaming mode is initiated if the `NEW_STREAM` directive is not the first line of the backup selections list. The list must contain either the `ALL_LOCAL_DRIVES` directive or wildcards.

In this mode, the backup selections list is sent to the client, which preprocesses the list and splits the backup into streams as follows:

- If the backup selections list contains the `ALL_LOCAL_DRIVES` directive, NetBackup backs up the entire client. However, NetBackup splits each drive volume (Windows) or file system (UNIX) into its own backup stream. See [“ALL\\_LOCAL\\_DRIVES directive and multiple data streams”](#) on page 561.
- If wildcards are used, the expansion of the wildcards results in one stream per wildcard expansion. Wildcard usage is the same as for Windows clients. See [“Wildcards in NetBackup”](#) on page 720.

If the backup selections list contains neither the `ALL_LOCAL_DRIVES` directive nor wildcards, the auto-discovery mode is not used. The server preprocesses rather than the client. Each file path in the backup selections list becomes a separate stream.

The auto-discovery streaming mode applies to Standard and MS-Windows policy types, except for NetWare clients.

Before the backup begins, the client uses auto-discovery to preprocess the backup selections list to determine how many streams are required. The first backup that a policy performs preprocesses the backup selections list. However, preprocessing does not necessarily occur before every backup. Whether or not it occurs depends on the preprocess interval.

## Setting the preprocess interval for auto-discovery

The preprocess interval applies only to auto-discovery mode and specifies how often preprocessing occurs. When a schedule is due and NetBackup uses auto-discovery, NetBackup checks whether the previous preprocessing session occurred within the preprocess interval.

NetBackup performs one of the following actions:

- If the preprocessing session occurs within the preprocess interval, NetBackup does not run preprocessing on the client.
- If the preprocessing session did not occur within the preprocess interval, NetBackup preprocesses the client and makes required changes to the streams.

If necessary, you can change the interval by using the `bpconfig` command. The default is four hours and is a good value for most of the sites that run daily backups.

If the interval is too long or too short, the following can occur:

- An interval that is too long can cause missed backups because new streams are not added early enough. For example, assume that the preprocess interval is set to four hours and a schedule has a frequency of less than four hours. A new stream can be omitted from the next backup because the preprocessing interval has not expired when the backup is due.
- An interval that is too short can cause preprocessing to occur often enough to increase scheduling time to an unacceptable level. A short interval is most likely to be a problem when the server must contact a large number of clients for preprocessing.

The form of the `bpconfig` command to use for changing the interval is as follows:

```
install_path\NetBackup\bin\admincmd\bpconfig [-prep hours]
```

For more information on the `bpconfig` command, see *NetBackup Commands*.

## UNSET, UNSET\_ALL directives, and multiple data streams

All policy-specific directives that are passed to a client in a stream are passed in all subsequent streams. The `UNSET` and `UNSET_ALL` directives change this behavior. These directives are recognized only if the **Allow multiple data streams** option is set for the policy.

See [“Directives for specific policy types”](#) on page 561.

See [“Allow multiple data streams attribute”](#) on page 486.

**UNSET**                      The `UNSET` directive interrupts a policy-specific directive so it is not passed with any additional streams. The directive that was unset can be defined again later in the backup selections list to be included in the current and the later streams.

In the following backup selections list, the `set` command is a client-specific directive that is passed to the first and all subsequent streams.

```
NEW_STREAM
set destpath=/etc/home
/tmp
/use
NEW_STREAM
/export
NEW_STREAM
/var
```

For the `set` command to be passed to the first two streams only, use `UNSET` or `UNSET_ALL` at the beginning of the third stream. At this location, it prevents `SET` from being passed to the last stream.

```
NEW_STREAM
set destpath=/etc/home
/tmp
/use
NEW_STREAM
/export
NEW_STREAM
UNSET_ALL [or UNSET set destpath=/etc/home]
/var
```

**UNSET\_ALL**                      `UNSET_ALL` has the same effect as `UNSET` but unsets all policy-specific directives in the backup selections list that have been defined up to this point.

## Excluding files from backups

By default, a number of files and file states are not backed up by NetBackup. You can also exclude specific files from automatic backups by specifying the files or directories in an exclude list on the client.

### Files that are excluded from backups by default

By default, NetBackup does not back up the following files:

- NFS files or directories. To back up NFS files, enable **Follow NFS**.

- Files or directories in a different file system. To back up files in a different file system, enable **Cross mount points**.
- Files or directories with path lengths longer than 1023 characters.
- Files or directories in which the operating system does not return inode information (the `lstat` system call fails).
- Directories that NetBackup cannot access (the `cd` command cannot access).
- Socket special files. (Named pipes are backed up, however.)
- Locked files when locked by an application that currently has the file open.
- Busy files. If a file is open, NetBackup backs up the last saved version of the file.

NetBackup automatically excludes the following file system types on most platforms:

- `cdrom` (all UNIX platforms)
- `cachefs` (AIX, Solaris, UnixWare)
- `devpts` (Linux)
- `mntfs` (Solaris)
- `proc` (UNIX platforms; does not exclude automatically for AIX, so `/proc` must be added manually to the exclude list. If `/proc` is not added manually, partially successful backups may result with the `ALL_LOCAL_DRIVES` directive on AIX)
- `tmpfs` (Linux)
- `usbdevfs` (Linux)

## Excluding files from automatic backups

On most NetBackup clients, you can exclude specific files from automatic backups by specifying the files in an exclude list on the client.

You can also create an include list to add a file(s) specifically that is excluded. The include list is useful to exclude an entire directory except for one file, for example.

---

**Note:** Exclude and include lists do not apply to user backups and archives.

---

The method for specifying files in the exclude and include lists depends on the type of client as follows:



|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| On Microsoft Windows clients | <p>Specify exclude and include lists in the Backup, Archive, and Restore client interface: Start Backup, Archive, and Restore and click <b>File &gt; NetBackup Client Properties</b>. Go to the <b>Exclude</b> list or <b>Include</b> list tab. For further instructions, see the NetBackup user's guide for the client.</p> <p>The <b>Exclude</b> list or the <b>Include</b> list can also be specified through the NetBackup Administration Console on the master server.</p> <p>See <a href="#">"Exclude Lists properties"</a> on page 114.</p> |
| On NetWare target clients    | <p>The exclude and include lists are specified when the targets are added. See the NetBackup user's guide for the client.</p>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| For UNIX clients             | <p>Create the exclude and include lists in the following files on the client:</p> <ul style="list-style-type: none"> <li>■ /usr/openv/netbackup/include_list</li> <li>■ /usr/openv/netbackup/exclude_list</li> </ul>                                                                                                                                                                                                                                                                                                                               |

## Windows excluded files

Windows maintains a list of files and folders that are excluded when Microsoft Windows Backup is used to back up files. This list is known as the **FilesNotToBackup** list. NetBackup excludes those files and directories from automatic backups even if they are not in the NetBackup exclude list for the client. Those items also are excluded from user-directed backups (unlike items in a NetBackup exclude list, which can be backed up by a user-directed operation).

Windows also maintains a list of registry keys that are not to be restored. NetBackup does not restore the registry keys that are listed in the **Windows KeysNotToRestore** list.

## About the Disaster Recovery tab

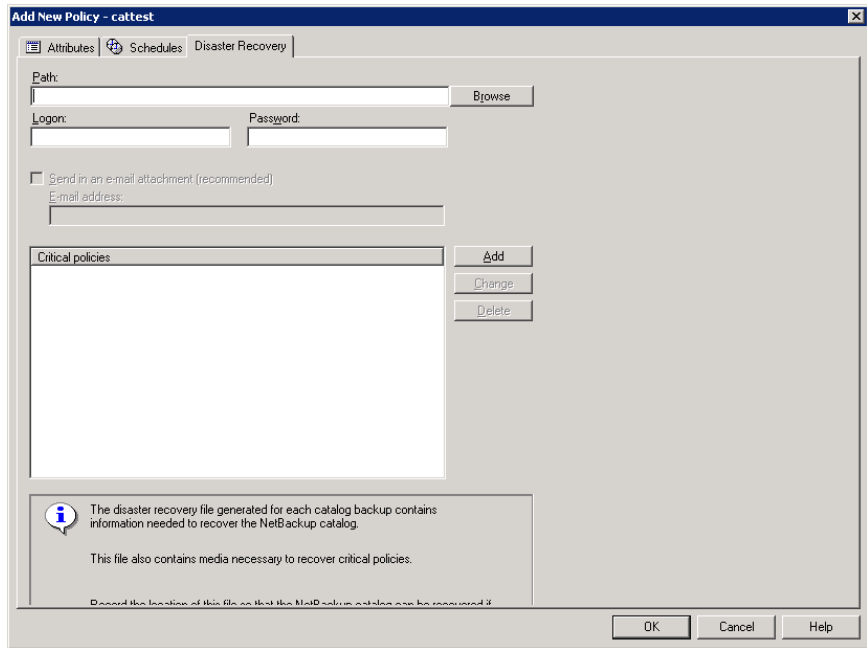
The **Disaster Recovery** tab appears for those policies that are based on the NBU-Catalog policy type for catalog backups. The **Disaster Recovery** tab contains options for configuring disaster recovery protection methods for the catalog data.

---

**Note:** Do not save the disaster recovery information to the local machine. Symantec recommends that the image file be saved to a network share or a removable device.

---

Figure 15-26 Disaster Recovery tab



---

**Note:** Vault protects the disaster recovery data by sending the data to the Vault site as an email attachment of the Vault report email.

---

The following topics describe the options on the **Disaster Recovery** tab.

## Path

The path indicates the directory where the disaster recovery information is to be saved. Symantec recommends that you save the image file to a network share or a removable device. Do not save the disaster recovery information to the local machine.

The share must be established and available before the hot catalog backup runs. Specify an NFS share, or a UNC path (CIFS Windows share).

When indicating a UNC path, note the following:

- A Windows master server can indicate a UNC path to a Windows machine.
- A UNIX master server cannot indicate a UNC path to a Windows machine.

- A UNIX master server cannot indicate a UNC path to a UNIX machine. To do so, first mount that UNC location on the master server, and then provide the UNC path to the UNIX machine.

## Logon

The logon and password information that is required to access an established Windows or NFS share.

If the logon information is not valid, NetBackup returns a message. The message requests that the user either reenter the logon and password information or clear the alternate location option to continue.

## Password

The password that is required to log on to the share.

## Send in an email attachment field

Symantec recommends that the disaster recovery report be sent to at least one email address. To send the information to more than one address, separate email addresses with a comma as follows:

```
email1,email2
```

See [“Setting up email notifications about backups”](#) on page 136.

The `nbmail.cmd` or `mail_dr_info.cmd` script must be configured (`Install_path\NetBackup\bin\`). In addition specify the email address(es) in the **Disaster Recovery** tab.

NetBackup performs the notification by passing the email addresses, subject, and message to `nbmail.cmd` or `mail_dr_info.cmd`. The scripts use the mail program that is specified in the script to send email to the user. See the comments in the script for configuration instructions.

The following points describe how `mail_dr_info.cmd` and `nbmail.cmd` interact:

- If `Install_path\NetBackup\bin\mail_dr_info.cmd` is configured, the disaster recovery report is sent to the administrator(s) that are indicated in the **Disaster Recovery** tab. NetBackup administrators can set up the script to send the disaster recovery information to alternate locations.
- If `mail_dr_info.cmd` is not configured, and `Install_path\NetBackup\bin\nbmail.cmd` is not configured, the disaster recovery report is sent to the administrator(s) that are indicated in the **Disaster Recovery** tab by `nbmail.cmd`.

- If neither file is configured, NetBackup attempts to use Microsoft internal IMAPI services.

---

**Note:** By default, neither `nbmail.cmd` nor `mail_dr_info.cmd` is configured to send email.

---

## Critical policies list

A policy that is listed on the **Critical Policies** list is considered crucial to the recovery of a site in the event of a disaster. The NetBackup **Disaster Recovery** report lists all of the media that is used for backups of critical policies, including the most recent full backup. The NetBackup **Disaster Recovery** wizard warns you if any media for critical policies are not available.

---

**Note:** The **Disaster Recovery** report lists the media for only incremental and full backup schedules so critical policies should use only incremental or full backup schedules. Certain database backups schedules, such as Oracle and Microsoft SQL Server, only use schedule types of Application Backup and Automatic Backup. Because of the schedule types, media listings for these backups do not appear on the **Disaster Recovery** report.

---

## Adding policies to the Critical Policies list of a catalog backup policy

Use the following procedure to add policies to the **Critical Policies** list of a catalog backup policy.

### To add a policy to the critical policies list


- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.

Double-click on a configured catalog backup policy. Or, to create a catalog backup policy, see the following procedure:

See [“Configuring a catalog backup manually”](#) on page 610.

- 2 Select the **Disaster Recovery** tab.

- Near the **Critical Policies** list, click **Add**. An active field appears in the list.



- Click to the right of the active field to display a list of configured policies. Select a policy to add it to the **Critical Policies** list.
- Click **Add** to add another policy to the list.  
 To change a policy, select the policy and click **Change**.  
 To delete a policy from the list, select the policy and click **Delete**.
- Click **OK** to save the policy.

## Creating a Vault policy

A Vault policy differs from other policies in the following respects:

- You must specify Vault as the policy type.
- You do not specify clients for Vault policies; therefore the Clients tab does not appear.
- Specify a Vault command in the backup selections list instead of a file(s).

### To create a Vault policy

- In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- Select **Actions > New > New Policy**.
- Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- On the **Attributes** tab, select **Vault** as the policy type.
- On the **Schedules** tab, click **New** to create a new schedule. The type of backup defaults to **Automatic**. Complete the schedule.  
 The Clients tab does not appear for Vault policy types.
- On the **Backup Selections** tab, enter one of two Vault commands:

- Use `vltrun` to specify the robot, vault name, and profile for the job. The `vltrun` command accomplishes all the steps necessary to select, copy, and eject media. If the vault profile name is unique, use the following format:

```
vltrun profile_name
```

If the vault profile name is not unique, use the following format:

```
vltrun robot_number/vault_name/profile_name
```

- Use the `vlteject` command to eject media or to generate reports for completed Vault sessions. For example:

```
vlteject -eject -report [-vault vault_name
[-sessionid id]] [-auto y|n] [-eject_delay seconds]
```

Both commands are located in the following directory:

```
install_path\netbackup\bin
```

For more information on Vault names, profile names, and command usage, see the *Vault Administrator's Guide*.

- 7 Click **OK**.

## Performing manual backups

A manual backup is user initiated and based on a policy.

A manual backup is useful in the following situations:

- To test a configuration
- To back up a client that missed the regular backup
- To back up a client before installing new software to preserve the old configuration
- To preserve records before a special event such as a company split or merger
- To back up quarterly or yearly financial information

In some cases, it may be useful to create a policy and schedule that you use only for manual backups. Create a policy for manual backups only by creating a policy with a single schedule that has no backup window. Without a backup window, the policy can never run automatically.

### To perform a manual backup

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 Select **Actions > Manual Backup**. (To perform a manual backup, you must enable the **Active. Go into effect at** option.)  
 See [“Go into effect at attribute”](#) on page 474.  
 If the **Go into effect** property is set for a future date and time, the backup does not run.
- 4 In the Manual Backup dialog box, select the schedule and the clients that you want to back up.  
 If you do not select any schedules, NetBackup uses the schedule with the highest retention level. If you do not select any clients, NetBackup backs up all clients.  
 User schedules do not appear in the schedules list. A user schedule cannot be manually backed up because they do not have a backup selection list (the user selects the files).
- 5 Click **OK** to start the backup.

## Active Directory granular backups and recovery

Administrators can now use NetBackup can restore individual objects and attributes in the Active Directory instead of restoring the entire Active Directory. It has also become easier for administrators to restore deleted objects (tombstone objects) from the Active Directory.

The following topics describes how to configure a policy to perform recovery of an Active Directory object:

- System requirements necessary to perform Active Directory granular backups and restores.
- How to configure a policy for an Active Directory backup that allows granular restores.
- How to restore individual objects and attributes in the Active Directory.

## System requirements for Active Directory granular backups and recovery

Active Directory granular restores are supported on the following systems:

- Windows 2003 R2 SP2
- Windows 2008
- Windows 2008 R2

To perform Active Directory granular backups and restores, ensure that you meet the following requirements:

- The master server, the media server, and clients must all have NetBackup 6.5.4 or later installed. And, all must be at the same level.
- The Network File System (NFS) must be installed on the media server and all Active Directory domain controllers or ADAM/LDS hosts.  
See [“About installing and configuring Network File System \(NFS\) for Active Directory Granular Recovery”](#) on page 789.  
See [“About configuring Services for Network File System \(NFS\) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients”](#) on page 797.  
See [“About configuring Services for Network File System \(NFS\) on the Windows 2008 and Windows 2008 R2 NetBackup media server and NetBackup clients”](#) on page 790.
- The NetBackup Client Service must be configured to log on as an account with domain privileges.  
To perform granular backups and restores of the Active Directory, the NetBackup Client Service (`bpineta`) must run under the domain administrator account on the Active Directory domain controller or ADAM server. By default, `bpineta` runs under the Local System account.  
See [“Configuring the log on account for the NetBackup Client Service”](#) on page 805.

For information on the media server platforms that support Granular Recovery Technology, see the following:

*NetBackup Enterprise Server and Server 7.x OS Software Compatibility List*

## Creating a policy that allows Active Directory granular restores

A policy that backs up the Active Directory can be configured to allow the restore of the objects and attributes in the Active Directory. The objects and attributes can be restored locally or remotely without the interruption of restarting the domain controllers where the restore is performed.

The **Active Directory** host properties offer additional configuration options for the backup of Windows Server 2008 computers. Specifically, whether or not NetBackup performs a consistency check if Microsoft Volume Shadow Copy Service (VSS) is used as the snapshot provider.



See [“Active Directory host properties”](#) on page 66.

**To create a policy to allow Active Directory restores**

- 1 Check that the NetBackup Client Service (`bpineta`) is running under the domain administrator account on the Active Directory domain controller. In this case, the Active Directory domain controller is the NetBackup client.  
 See [“Configuring the log on account for the NetBackup Client Service”](#) on page 805.
- 2 In the **Policy** dialog box, on the **Attributes** tab, select **MS-Windows** as the policy type. Specify the other policy attributes as needed.
- 3 Enable the **Enable granular recovery** option. If this option is not enabled, the backup still runs, but the backup cannot produce granular restores.
- 4 In the **Schedules** tab, create schedules as needed.  
 Other items in the policy may use a differential or cumulative incremental backup type, but the Active Directory items are always fully backed up.  
 See [“Active Directory backups are full backups”](#) on page 577.
- 5 In the **Backup Selections** tab, open the **Select Directive** dialog.
- 6 For the **Directive set**, select **Windows 2003** or **Windows 2008**.
- 7 To back up the Active Directory, select any one of the following directives:
  - See [“System\\_State directive”](#) on page 558.
  - See [“Shadow Copy Components:\ directive”](#) on page 559.
  - See [“ALL\\_LOCAL\\_DRIVES directive”](#) on page 558.

---

**Note: Active Directory Application Mode (ADAM)** is a lightweight directory service that runs as a user service. This directive can be used to back up ADAM data on computers where it is installed. However, it does not back up the Active Directory itself.

---

- 8 In the **Clients** tab, select the clients as needed.
- 9 Save the policy.

**Active Directory backups are full backups**

Any Active Directory backup is always a full backup, whether it is a granular backup or not.

Whenever Active Directory is in a policy's **Backup Selections** list, the Active Directory portion is always fully backed up, even when the backup type is incremental, differential or cumulative. Any other items in the **Backup Selections** list may use a differential or cumulative incremental backup type as indicated. Even though a full backup is forced for an Active Directory backup, normal incremental rules are applied to the non-Active Directory items in the policy file list.

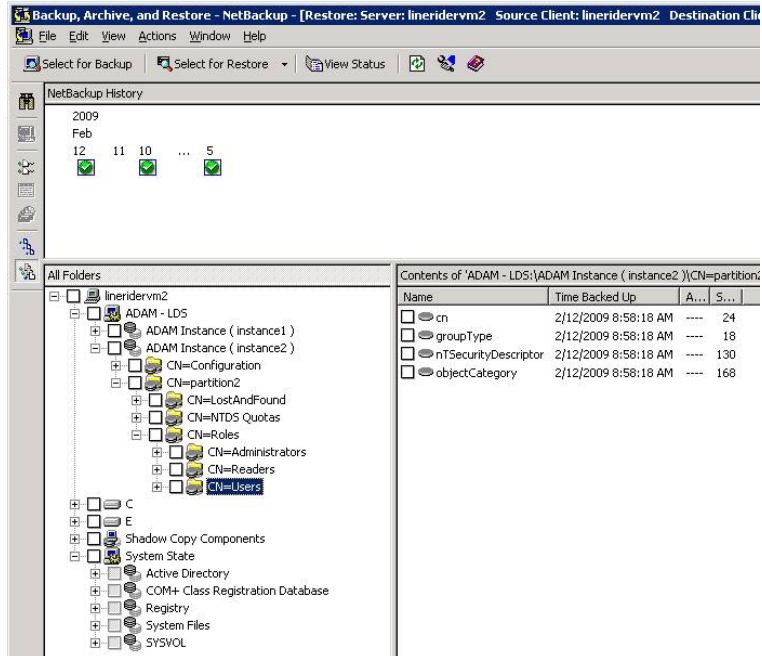
## Restoring Active Directory objects

The following procedure describes how to restore objects from an Active Directory backup in a non-disaster recovery situation:

### To restore individual objects from an Active Directory backup

- 1 Open the NetBackup Backup, Archive, and Restore client interface.
- 2 Select **File > Select Files and Folders to Restore**.
- 3 Expand and browse the **Active Directory** node.

- 4 Select the objects to be restored. Do not select both granular and non-granular objects. When a user explores and expands selections, a delay can occur during communication with the NetBackup server. The delay is a result of dynamically determining the contents from the image on the media server. The approach prevents the NetBackup catalog from unanticipated growth due to numerous granular entries.



- 5 Select **Action > Restore**.
- 6 If an Active Directory object is selected, the **Restore Marked Files** dialog box contains two tabs:
  - **General tab**  
When an Active Directory object is selected, the **Restore Destination Choices** are disabled in the **General** tab. Configure the other restore options as needed.
  - **Active Directory tab**  
The **Active Directory** tab contains an option to recreate the objects that have been deleted: **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.

The **Active Directory** tab contains an option that lets administrators recreate the objects whose tombstone lifetimes have passed. The objects have also been purged from the Active Directory Deleted Objects container. To allow this capability, enable the option labeled **Recreate deleted objects that cannot be restored from the Active Directory Deleted Objects container**.

- 7 Click **Start Restore** in the **Restore Marked Files** dialog box.

Some restore situations require additional steps, depending on what is restored.

See [“Restore issues”](#) on page 580.

## Restore issues

The following sections describe additional information about granular restores. Some situations require additional steps to fully restore the objects. In some situations, a granular restore of some part of the Active Directory is not possible.

### Restores that are disabled

At times, when user and computer accounts are restored from a granular Active Directory restore, they are disabled. The following topics describe possible reasons why the accounts can be disabled.

#### Deleted objects

When objects in Active Directory are deleted, they are removed from their current Active Directory or ADAM/AD LDS container. They are converted into tombstones and placed in the Active Directory Deleted Objects container where their tombstone lifetime is monitored. By default, NetBackup restores deleted objects from this container if the tombstone lifetime has not passed.

After the tombstone lifetime passes, the tombstones are purged from the Active Directory Deleted Objects container. Purging the tombstones has the effect of permanently deleting the objects from the Active Directory and ADAM/AD LDS databases.

#### User objects

When restoring user objects, you must reset the object's user password and enable the object's user account:

- For Active Directory user objects, use the Microsoft Active Directory Users and Computers application.
- For ADAM/AD LDS user objects, use ADSI Edit.

In Active Directory, computer objects are derived from user objects. Some attributes that are associated with a computer object cannot be restored when you restore a deleted computer object. They can only be restored if the attributes were saved through schema changes when the computer object was originally deleted.

### Computer objects

Computer object credentials change every 30 days and the credentials from the backup may not match the credentials that are stored on the actual computer. When a computer object is restored it is disabled if the **userAccountControl** property was not preserved in the deleted object.

Use the Microsoft Active Directory Users and Computers application to reset a computer object.

#### To reset a computer object's account

- 1 Remove the computer from the domain.
- 2 Re-join the computer to the domain. The security identifiers (SID) for the computer remains the same since it is preserved when a computer object is deleted. However, if the tombstone expired and a new computer object was recreated, the SID is different.

### Group and member objects

To restore Active Directory group membership links may require that the restore job be run twice.

For example, consider the case where a group and its member objects are deleted.

If a restore job contains both group objects and member objects, the job restores the objects in alphabetical order. However, the group that is restored has a link dependency on a member that does not exist yet. When the group is restored, the link cannot be restored.

Run the restore again to restore all forward and backward links.

### Group policy objects

NetBackup does not support granular restores of Group Policy Objects.



# Synthetic backups

This chapter includes the following topics:

- [About synthetic backups](#)
- [Policy considerations and synthetic backups](#)
- [Types of synthetic backups](#)
- [When to use synthetic backups](#)
- [Synthetic backup jobs create two sets of catalog files](#)
- [Change journal and synthesized backups](#)
- [True image restore and synthesized backups](#)
- [Checkpoint restart and synthesized backups](#)
- [Displaying synthetic backups in the Activity Monitor](#)
- [Logs produced during synthetic backups](#)
- [Synthetic backups and directory and file attributes](#)
- [Using the multiple copy synthetic backups method](#)
- [Optimized synthetic backups using OpenStorage](#)

## About synthetic backups

Synthetic backups can be written to tape storage units or disk storage units, or a combination of the two.

The following sections describe how synthetic backups work in a NetBackup configuration:

- Processing takes place on master and media server(s) instead of client  
During a traditional full backup, all files are copied from the client to a master server or a media server. The files are copied even though those files may not have changed since the last incremental backup.  
When NetBackup creates a synthetic full backup, NetBackup detects whether new or changed files have been copied to the media server during the last incremental backup. The client does not need to be running to combine the full backups and the incremental backups on the media server to form a new, full backup. The new, full synthetic backup is an accurate representation of the clients' file system at the time of the most recent full backup.
- Reduces the network traffic  
Network traffic is reduced because files are transferred over the network only once. After the backup images are combined into a synthetic backup, the tapes or disk that contain the component images can be recycled or reclaimed. Synthetic backups can reduce the number of tapes or disk space in use.
- Supports disk environments  
Synthetic backups can be created in the environments that are comprised exclusively of disk storage units.
- Uses drives more effectively  
Synthetic backups can be written to tape storage units or disk storage units, or a combination of both. If the backups use tape, the backups can be synthesized when drives are not generally in use. For example, if backups occur primarily at night, the drives can synthesize full backups during the day.

## Policy considerations and synthetic backups

The **Synthetic Backup** option is available under the following conditions:

- The policy type must be either Standard or MS-Windows.
- The **Collect True Image Restore Information With Move Detection** option must be selected on the **Policy Attributes** tab.  
See [“Collect true image restore information with move detection attribute”](#) on page 484.
- The schedule that is created for a synthetic backup must have **Synthetic Backup** selected.  
See [“Synthetic backup attribute”](#) on page 499.
- One of the following must be available:
  - Disk storage unit(s) with adequate space available.



- Tape library(s) with multiple drives to read and write. See [“When to use synthetic backups”](#) on page 589.
- A combination of disk storage unit(s) and tape library(s).

## Schedules that must appear in a policy for synthetic backups

A policy for synthetic backups must contain one of the following types of schedules:

- At least one traditional, full backup must be run successfully to create a full image. The synthetic backup job fails if there is not at least one previous full image.
- Schedule(s) for incremental backups.  
Incremental backups are necessary to capture the changes in the file system since the last full or incremental backup. The synthetic backup job receives a status code of 1 for a policy that contains full or incremental synthetic backup schedules, but no incremental backup schedules.  
The synthetic backup synthesizes all of the incremental backups to create a new full or cumulative backup image. Therefore, the synthetic backup is only as current as the last incremental backup.

---

**Note:** To configure a synthetic cumulative backup for any clients that are archive bit-based (default), use only differential incremental backups for the traditional, non-synthesized backups.

---

- One full and one cumulative backup schedule with the **Synthetic Backup** option selected. See [“Synthetic backup attribute”](#) on page 499.

## Adding clients to a policy for synthetic backups

After clients are added to a synthetic backup policy, run a traditional, full backup of the policy. A traditional backup is necessary before a synthetic backup can be created.

Since **Collect True Image Restore Information With Move Detection** is required for synthetic backups, all of the clients in the policy must support TIR.

See [“Collect true image restore information with move detection attribute”](#) on page 484.

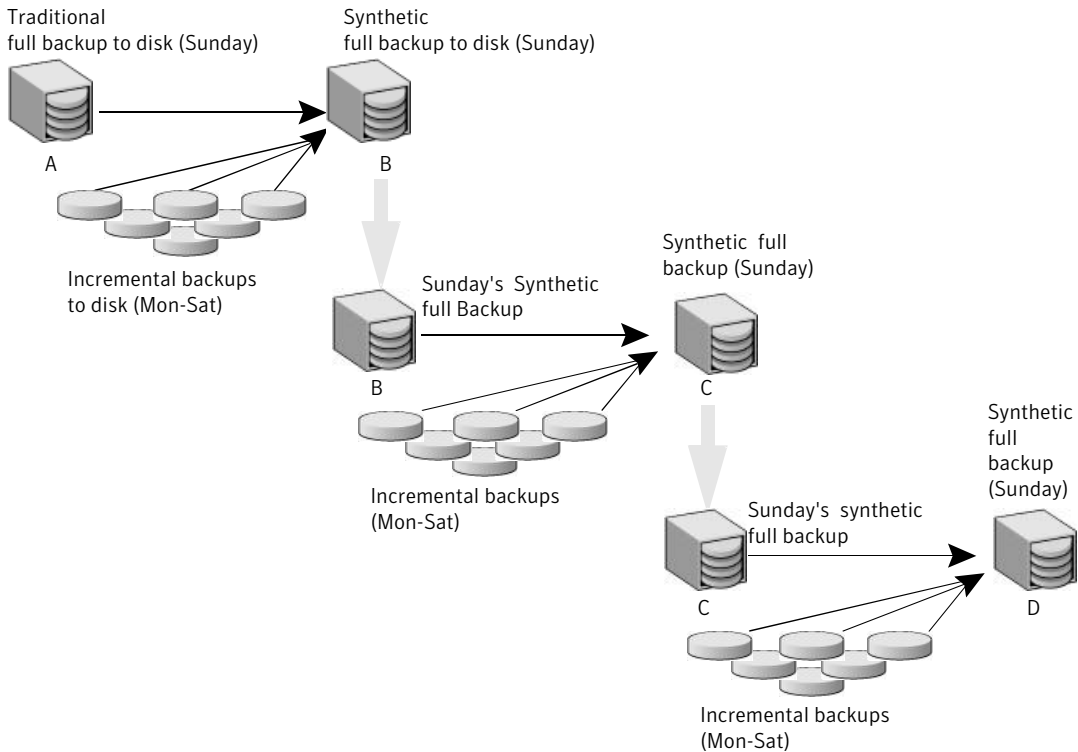
# Types of synthetic backups

Two types of synthetic backup images can be created: synthetic full and cumulative synthetic. The images that are used to create the synthetic image are known as component images. For instance, the component images in a synthetic full are the previous full image and the subsequent incremental images.

## Synthetic full backups

Figure 16-1 illustrates the creation of synthetic full backups (B, C, D) from an existing full backup (A) and shows the incremental backups between full backups.

Figure 16-1 Creation of synthetic full backups



The traditional full backup (A) and the incremental backups are created in the traditional manner: data is scanned, then copied from the client’s file system to the backup media. The synthetic backups do not interact with the client system at all, but are instead synthesized on the media server.

See “[Synthetic cumulative incremental backups](#)” on page 587.

The following is an example of a synthetic full backup:

- Create a Standard or MS-Windows policy for the clients (5.0 or later) you want to back up. Include the following schedules:
  - A schedule for one full, traditional backup to run at least once.
  - A schedule for daily (Monday through Saturday) differential incremental backups.
  - A schedule for weekly full, synthetic backups.
- Make sure that the traditional full backup runs. If the backup does not complete, run the backup manually.
- Per schedule, run daily, differential incremental backups for the clients throughout the week. The last incremental backup for the week runs on Saturday.
- Per schedule, run synthetic full backups for the clients on subsequent Sundays.

---

**Note:** The synthetic full backups in the scenario are only as current as the Saturday incremental backup.

---

## Synthetic cumulative incremental backups

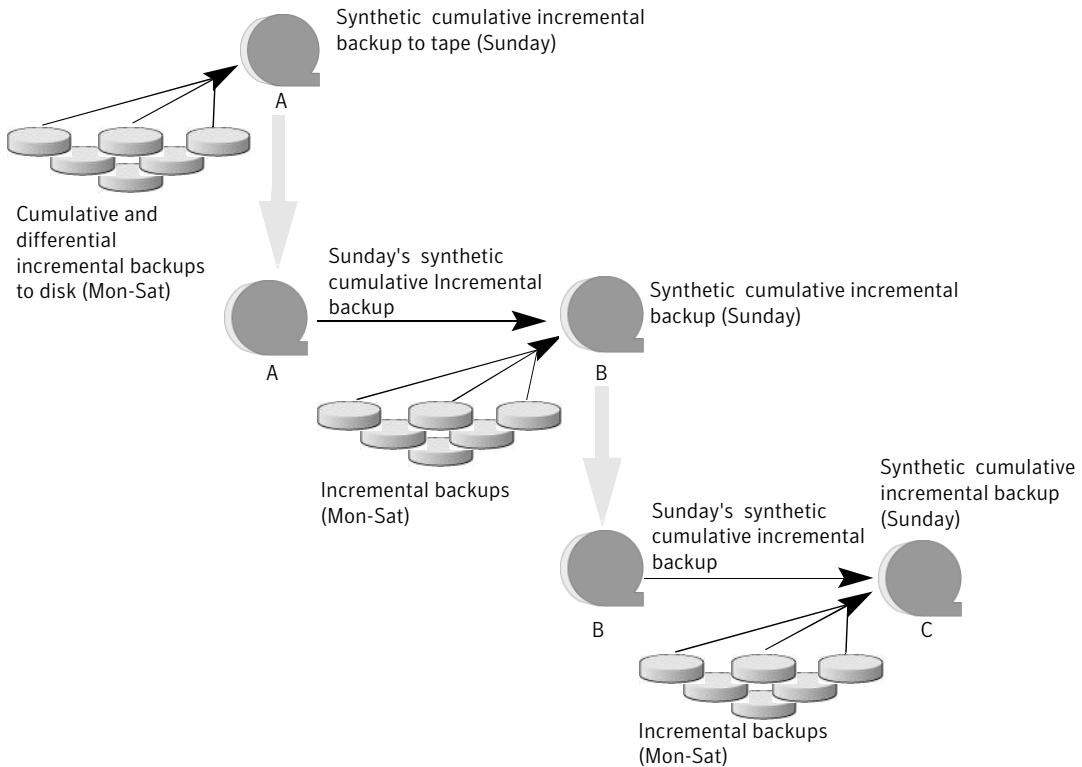
The scenario to create a synthetic, cumulative incremental backup is similar to the scenario to create a synthetic full backup. Remember, a cumulative incremental backup includes all changes since the last full backup.

If a cumulative incremental backup exists that is newer than the last full backup, a synthetic cumulative backup image is produced by consolidating the following component backup images:

- All differential incremental backups that were taken since the last cumulative backup.
- The last cumulative incremental backup. If no cumulative incremental backup is available, only the differential incremental backups are used for the synthetic image.

**Figure 16-2** illustrates the creation of synthetic cumulative incremental backups (A, B, C) from the latest cumulative incremental backup and shows the subsequent differential incremental backups.

**Figure 16-2** Creation of synthetic cumulative backups



The following is an example of a synthetic cumulative backup:

- Create a Standard or MS-Windows policy for the clients (5.0 or later) you want to back up. Include the following schedules:
  - A schedule for one full, traditional backup to run at least once.
  - A schedule for daily (Monday through Saturday) differential incremental backups.
  - A schedule for weekly cumulative incremental synthetic backups.
- Make certain that the traditional full backup runs. If the backup does not complete, run the backup manually.
- Per schedule, run daily differential incremental backups for the clients throughout the week. The last incremental for the week runs on Saturday.
- Per schedule, run synthetic cumulative incremental backups for the clients on subsequent Sundays.

---

**Note:** The synthetic cumulative backups in the scenario are only as current as the Saturday incremental backup.

---

## When to use synthetic backups

The following items concern scenarios and notes about when to use synthetic backups:

- The synthetic full backup is a scalable solution for backing up remote offices with manageable data volumes and low levels of daily change.  
If the clients experience a high rate of change daily, the incremental backups are too large. In this case, a synthetic backup is no more helpful than a traditional full backup.
- Refrain from multiplexing backups that will be synthesized. Do not synthesize multiplexed backups because it is inefficient. To synthesize multiplexed client images requires multiple passes over the source media—one per client.
- Performance issues occur if multiple streams are selected for synthesized backups. The issues are similar to those encountered while multiplexing synthesized backups problems. Back up to disk whenever possible to improve multiple stream performance issues.
- Reduce the gap between the last incremental backup and the synthesized backup. Since a synthetic backup does not involve direct contact with the client, a synthetic backup is only as current as the last incremental backup. If there is a concern to reduce a potential gap in backup coverage, run an incremental backup before the synthetic backup.  
Only frequency-based schedules allow an incremental backup and a synthetic backup to run on the same day. Calendar-based schedules are dependent on one another. If a daily incremental schedule runs earlier in the day, the synthetic cumulative backup does not run later that same day (00:00:00–23:59:59).
- Disk-based images are more efficient for synthesizing. For example, NetBackup processes the newest component images first in a synthesized backup, followed by sequentially older images. When two or more component images are written to the same tape, the tape movement can be somewhat inefficient compared to disk-based images.
- If using tape storage units, consider the following:
  - For tape backups, separate tapes are required. The tape for the synthetic image must be different from the tape where the component images reside.

- The maximum drive usage applies only to the drive that is needed for writing the synthetic backup. If any of the component images reside on tape, an additional drive is needed for reading.
- If a single tape drive device is used to generate synthetic images, place component images in a hard drive location first. In that way, a synthetic image can be generated with the single tape drive device.
- Testing on synthetic backups has found the following about:
  - The time it takes to run a synthetic full backup does not increase significantly over time.
  - Synthetic full backups are generated more quickly when built from disk-based incremental backups. If the synthetic full backup is also generated on disk, the run time is even faster. The disk copy then can be duplicated to tape.
- Testing on restores from synthetic backups has found the following:
  - The time that is required to perform a restore from a synthetic backup does not increase significantly over time.
  - The restore times for both a complete synthetic backup and for a single file is the same. It is the same whether the restore is from a traditional backup or from a synthetic backup.
  - The restore time of a single directory may increase over time when sourced from synthetic backups. The restore time depends on the pattern of file changes within the directory.

Contrast a traditional full backup, which stores the files in file system order with a synthetic full backup, which stores the files in last-file-accessed order. The synthetic full contains the newest files at the front of the media and the unchanged files at the end. Over time, the processing order introduces the potential for fragmentation of a single directory across the synthetic full image.

Note that the scenario is limited to single directory restores. Single file restores and full image restores from synthetic fulls are equal or better than from traditional full backups, as noted in previous bullets.
- Synthetic backups are supported on all media server platforms and tier one master server platforms.
- The option to create multiple copies is not allowed for synthetic backups.
- Synthetic backups are not supported if any of the component images are encrypted.

- A user-generated backup cannot be used to generate a synthetic image. A backup that is generated from a User Backup schedule or a User Archive schedule cannot be used as one of the components of a synthetic backup.

## Synthetic backup jobs create two sets of catalog files

When a synthetic backup job is run, two sets of catalog files are created: an image file and one or more .f files.

Each set uses the following timestamps:

- The catalog for a synthetic image usually has a timestamp one second later than the most recent incremental component image. The timestamp may be more than one second later if there were possible image name conflicts.
- The second set is named with the current timestamp. The set is used to mark the time the synthetic backup job was run. It does not contain any file data.

Do not manually remove any of these catalog files. The catalog files automatically expire after the retention period as specified in the schedule for the policy. The two sets of catalogs have the same expiration times.

For example, these are the catalog files after the incremental backup jobs run:

```
XDisk_1064417510_INCR
XDisk_1064417510_INCR.f

XDisk_1064420508_INCR
XDisk_1064420508_INCR.f

XDisk_1064421708_INCR
XDisk_1064421708_INCR.f
```

The following files are the first set of catalog files after a synthetic full backup job runs. The timestamp displays the most incremental timestamp plus one:

```
XDisk_1064421709_FULLL
XDisk_1064421709_FULLL.f
```

The following file is from the second set of catalog files, showing the current timestamp:

```
XDisk_1064424108_FULLL
```

## Change journal and synthesized backups

If this Windows client host property is enabled, the property has no effect when the client is backed up using the synthetic backup schedule.

See [“Use change journal in incrementals”](#) on page 97.

## True image restore and synthesized backups

Since the **Collect true Image restore information with move detection** policy property must be enabled for synthetic backups, all clients that are included in the policy must support TIR.

See [“Collect true image restore information with move detection attribute”](#) on page 484.

The **Keep true image restoration (TIR) information** property indicates how long TIR information in the image catalog is kept before it is pruned (removed). The property is located in the master server **Clean-Up** host properties.

See [“Keep true image restoration \(TIR\) information”](#) on page 77.

However, if a synthetic full and synthetic cumulative schedule was defined in the policy, the TIR information is pruned from the component images until a subsequent traditional or synthetic full or cumulative backup image has generated successfully.

Consider a situation where **Keep true image restoration (TIR) information** host specifies that TIR information is pruned from the catalog after two days. On the third day the TIR information is pruned only if a traditional or synthetic full backup image has been generated.

If the TIR information was pruned from a component image and you accidentally expire the most recent synthetic image, rerun the synthetic backup job to restore automatically the TIR information to the catalog. In case the TIR information cannot be restored due to bad, missing, or vaulted media, the synthetic backup job fails with error code 136 (TIR info was pruned from the image file). If the problem is correctable, run the synthetic backup again.

## Checkpoint restart and synthesized backups

If Checkpoint Restart is indicated for the policy, the backups that are produced with the synthetic backup schedule are not checkpointed. The option is enabled if **Take checkpoints** on the policy Attributes tab is enabled. If the **Take checkpoints** option is enabled for a synthetic backup, the property has no effect.



# Displaying synthetic backups in the Activity Monitor

A synthetic job is distinguished from a traditional full backup by the notation that is indicated in the Data Movement field of the Activity Monitor. Synthetic jobs display Synthetic as the Data Movement type while traditional backups display Standard.

## Logs produced during synthetic backups

When a synthetic backup is scheduled, NetBackup starts the `bpsynth` program to manage the synthetic backup process. `bpsynth` plans how the synthetic backup is built from the previous backup images.

If it is needed, `bpsynth` then schedules the tape drive resources that are needed for the synthetic backup. If the required resources are not available, the job fails with a status code that indicates that a resource is needed.

If the resources can be obtained eventually but not immediately, the synthetic job waits until the resources become available. A synthetic job may wait while a backup, restore, or another synthetic backup job uses a drive.

`bpsynth` passes the information to programs `bptm` and `bpdm` so that tape and disk images can be read or written. Catalog information is managed using `bpdbm`. Each of these programs has a debug log file in the logs directory.

If problems occur with synthetic backups, the following debug logs are required to diagnose the problem:

- On the master server: `bpsynth`, `bpdbm`, and the log files located in `install_path:\Program Files\VERITAS\NetBackup\logs` as described in the *NetBackup Troubleshooting Guide*.
- On the media server(s): `bptm` (if any tape images), `bpdm` (if any disk images), `bpcd`  
Note that several media servers can be involved if the component images are on different nodes.

However, `bpsynth` is used for each stream or client. To use `bpsynth` can be inefficient with tape images since `bpsynth` needs a tape drive to write the new image. Also, `bpsynth` may use the same component image volumes. One may need to finish before the next can proceed.

## Synthetic backups and directory and file attributes

For a synthetic backup to include directory and the file attribute changes, the change must first be picked up by a component incremental backup. (For example, changes like Access Control Lists (ACLs).)

On UNIX, changing an object's ACL changes the `ctime` (inode change time) for the object but not the `mtime` (data modification time). Since `mtime` triggers incremental backups, the ACL change is not reflected in an incremental backup, and therefore not in a synthetic full backup.

To include ACL changes in backups, enter `USE_CTIME_FOR_INCREMENTALS` in the `bp.conf` file on each UNIX client.

For each Windows client, enable **Incrementals: Based on Archive Bit**. The property is found under **NetBackup Management > Host Properties > Clients > selected client(s) > Windows Client**.

See [“Incrementals based on archive bit”](#) on page 97.

## Using the multiple copy synthetic backups method

The multiple copy synthetic backups method introduces the capability to produce a second copy of a synthetic backup at a remote site as part of a normal synthetic backup job.

This method provides the following benefits:

- It eliminates the bandwidth cost of copying synthetic full backups to another site.  
Instead of duplicating a local synthetic full backup to a remote site to produce a second copy, it is more efficient to produce the second copy by using data movements only at the remote site.
- It provides an efficient method to establish a dual-copy disaster recovery scheme for NetBackup backup images.

[Table 16-1](#) emphasizes how the synthetic full backup produced at the remote site is a clone, or a second copy, of the first copy produced at the local site.

**Table 16-1** Comparing synthetic copy process with and without method enabled

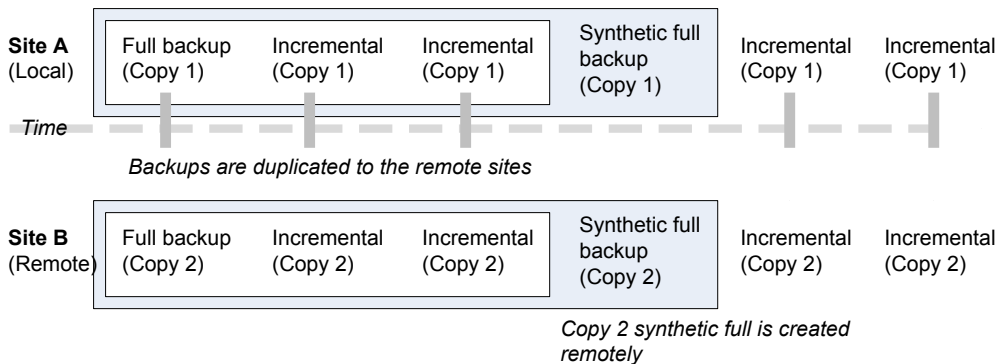
| Step | Without using the multiple copy synthetic backups method: | Using the multiple copy synthetic backups method: |
|------|-----------------------------------------------------------|---------------------------------------------------|
| 1    | A full backup is performed at the local site (Site A).    | Step 1 remains the same.                          |

**Table 16-1** Comparing synthetic copy process with and without method enabled (continued)

| Step | Without using the multiple copy synthetic backups method:          | Using the multiple copy synthetic backups method:                                                                                                                                    |
|------|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2    | The full backup is duplicated to the remote site (Site B).         | Step 2 remains the same.                                                                                                                                                             |
| 3    | An incremental backup is performed at Site A.                      | Step 3 remains the same.                                                                                                                                                             |
| 4    | The incremental backup is duplicated to Site B.                    | Step 4 remains the same.                                                                                                                                                             |
| 5    | Steps 3 and 4 are repeated each time an incremental schedule runs. | Step 5 remains the same.                                                                                                                                                             |
| 6    | A full synthetic backup is produced at Site A.                     | Step 6 remains the same.                                                                                                                                                             |
| 7    | The full backup is duplicated to Site B.                           | A full synthetic backup is produced at Site B from images at Site B.<br><br>The full synthetic backup at the remote site is a second copy of the synthetic backup at the local site. |
| 8    | Steps 2 through 7 repeat per backup scheduling needs.              | Step 8 remains the same.                                                                                                                                                             |

Figure 16-3 shows how no extra bandwidth is used to copy the synthetic full backup from Site A to Site B.

**Figure 16-3** Remote creation of synthetic full backup



## Configuring multiple copy synthetic backups

To configure a multiple copy synthetic backup, create a configuration file on the master server for each synthetic backup policy for which a second copy is to be produced.

The configuration file is a text file that is named after the policy and schedule:

```
multi_synth.policy.schedule
```

Create the file in the following location:

```
install_path\VERITAS\NetBackup\db\config\multi_synth.policy.
schedule
```

## Configuration variables

The file format uses a traditional name-pair scheme for setting configuration preferences. Each preference uses a key name that is separated from the preference value by an equal sign with each name-value pair residing on a single line.

For example:

```
NAME=VALUE
```

Enter all values as integers.

[Table 16-2](#) describes the configuration entries that can be included in the configuration file.

**Table 16-2** Configuration entries

| Entry       | Purpose                                                                                                                                                                                          |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SRC_COPY    | Specifies the copy number of each source component for the second synthetic backup. Every source backup must have a copy by this number unless SRC_COPY_FALLBACK is specified. The default is 2. |
| TARGET_COPY | Specifies the copy number for the second synthetic backup produced. This must be different from the copy number of the first synthetic backup (which is 1). Default is 2.                        |
| COPY        | COPY is an alternate specification for SRC_COPY and TARGET_COPY.<br><br>If COPY is specified and either SRC_COPY and TARGET_COPY is not specified, the value for COPY is used.                   |

**Table 16-2** Configuration entries (*continued*)

| Entry             | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TARGET_STU        | <p>Specifies the storage unit name or storage unit group name where the second copy synthetic backup is to be written. Use the special identifier <code>__ANY__</code> to indicate that Any Available storage unit can be used that is not configured to be on demand only. Note that there are two underscores before and after <code>ANY</code>:</p> <p>TARGET_STU= <code>__ANY__</code></p>                                                                                                                                                                                                                                                                                |
| FAIL_MODE         | <p>The second synthetic backup is produced immediately following the first copy synthetic backup if no errors occur during production of the first copy. If an error occurs during the second copy, the <code>FAIL_MODE</code> value specifies the fate of the first copy job and image.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"> <li>■ <code>FAIL_MODE=ALL</code><br/> <code>ALL</code> means that if the second copy fails, the first copy and its job also fail. (Default.)</li> <li>■ <code>FAIL_MODE=ONE</code><br/> <code>ONE</code> means that if the second copy fails, the failure does not affect the first copy job.</li> </ul> |
| ENABLED           | <p>Specifies whether production of the second copy is enabled or disabled. This entry turns on the feature.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"> <li>■ <code>ENABLED=YES</code><br/> Production of the second copy is enabled. (Default.)</li> <li>■ <code>ENABLED=NO</code><br/> Production of the second copy is disabled.</li> </ul>                                                                                                                                                                                                                                                                                                |
| SRC_COPY_FALLBACK | <p>Specifies that if a copy by the number given in <code>SRC_COPY</code> or <code>COPY</code> does not exist, the synthetic backup should use the primary backup.</p> <p>The only valid value is the following:</p> <p><code>SRC_COPY_FALLBACK=PRIMARY</code></p>                                                                                                                                                                                                                                                                                                                                                                                                             |
| VOLUME_POOL       | <p>Specifies the volume pool for tape media, if one is used. If no volume pool is specified, NetBackup uses the volume pool that is specified in the policy. If a volume pool is entered for disk, the entry is ignored.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## Simple configuration example

The following multiple copy synthetic configuration example takes advantage of default values to produce the second synthetic copy.

```
TARGET_STU=disk_stu
```

The default source of copy 2 and the default destination copy 2.

## Advanced configuration example

In this example, the second copy targets a tape library (`tape_stu`). The configuration specifies a volume pool (`Synthetics`) for the target copy.

The copy number for the multiple copy synthetic backup is copy 3. If copy 3 is unavailable, `SOURCE_COPY_FALLBACK` indicates that copy 3 can be produced using the primary copy.

If copy 3 fails, only copy 3 fails and not the job of the primary copy.

```
TARGET_STU=tape_stu
VOLUME_POOL=Synthetics
SOURCE_COPY_FALLBACK=PRIMARY
COPY=3
ENABLED=YES
FAIL_MODE=ONE
```

# Optimized synthetic backups using OpenStorage

NetBackup environments that use the Enterprise Disk license key environment can benefit from the OpenStorage optimized synthetic backup method.

This method constructs the synthetic image by using calls from the media server to the storage server. The media server tells the storage server which full and incremental images to use to create the synthetic backup. Then, the storage server constructs (or synthesizes) the synthetic image directly on the storage server, reducing network traffic.

See the *NetBackup Shared Storage Guide* for more information.

# Protecting the NetBackup catalog

This chapter includes the following topics:

- [About NetBackup catalogs](#)
- [Parts of the catalog](#)
- [Protecting the catalog](#)
- [Recovering the catalog](#)
- [Disaster recovery emails and the disaster recovery file](#)
- [Archiving the catalog](#)
- [Estimating catalog space requirements](#)

## About NetBackup catalogs

NetBackup catalogs are the internal databases that contain information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and the storage devices.

Since NetBackup needs the catalog information so that it can restore client backups, configure a catalog backup before using NetBackup for regular client backups. Schedule the catalog backups to occur on a regular basis. Without regular catalog backups, you risk losing regular backups if there is a problem with the disk that contains the catalogs.

For information on how to configure catalog backups in cluster environments, see the *NetBackup High Availability Guide*.

## Parts of the catalog

The NetBackup catalog resides on the NetBackup master server.

The catalog consists of the following parts:

- The image database.  
The image database contains information about the data that has been backed up. It is the largest part of the catalog.  
See [“About the image database”](#) on page 600.
- NetBackup data that is stored in relational database files.  
The data includes media and volume data describing media usage and volume information, which is used during the backups.  
See [“About the NetBackup relational database”](#) on page 602.
- NetBackup configuration files.  
The configuration files (`databases.conf` and `server.conf`) are flat files that contain instructions for the SQL Anywhere daemon.  
See [“About the server.conf file”](#) on page 632. and [About the databases.conf file](#).

## About the image database

The image database contains subdirectories for each client that is backed up by NetBackup, including the master server and any media servers.

The image database is located at `Program Files\VERITAS\Netbackup\db\images` and contains the following files:

- Image files (files that store only backup set summary information)
- Image `.f` files (files that store the detailed information of each file backup)

The image database is the largest part of the NetBackup catalog. It consumes about 99% of the total space that is required for the NetBackup catalog. While most of the subdirectories are relatively small in the NetBackup catalogs, `\images` can grow to hundreds of gigabytes. The image database on the master server can grow too large to fit on a single tape. Image database growth depends on the number of clients, policy schedules, and the amount of data that is backed up.

See [“Estimating catalog space requirements”](#) on page 621.

If the image catalog becomes too large for the current location, consider moving it to a file system or disk partition that contains more space.

See [“Moving the image catalog”](#) on page 623.

The image database component of the NetBackup catalog uses the `.f` files in binary format for Windows, Solaris, HP\_UX, AIX, and Linux platforms.



The catalog conversion utility (`cat_convert`) can be used to upgrade an image database to the binary format.

Information about the `cat_convert` command is available in the *Commands Guide*.

See “[Moving the image catalog](#)” on page 623.

## About image files

Each image file is an ASCII file, generally less than 1 kilobyte in size. An image file contains only backup set summary information. For example, the backup ID, the backup type, the expiration date, fragment information, and disaster recovery information.

## About image .f files

The binary catalog can contain one or more image .f files. This type of file is also referred to as a files-file. The image .f file may be large because it contains the detailed backup selection list for each file backup. Generally, image files range in size from 1 kilobyte to 10 gigabytes.

The file layout determines whether the catalog contains one .f file or many .f files. NetBackup configures the file layout automatically, based on the size of the binary catalog. NetBackup uses one of two layouts: single file layout or multiple file layout.

### ■ Image .f file single file layout

NetBackup stores file information in a single image .f file if the information for the catalog is less than 4 megabytes.

When the backup file of one catalog backup is less than 4 megabytes, NetBackup stores the information in a single image .f file. The image .f file is always greater than or equal to 72 bytes, but less than 4 megabytes.

### ■ Image .f file multiple file layout

When the file information for one catalog backup is greater than 4 megabytes, the information is stored in multiple .f files: one main image .f file plus nine additional .f files.

Separating the additional .f files from the image .f file and storing the files in the `catstore` directory improves performance while writing to the catalog. The main image .f file is always exactly 72 bytes.

```
-rw- 1 root other 72 Aug 30 00:40 test_1030680524_INCR.f
-rw- 1 root other 804 Aug 30 00:08 catstore/test_1030680524_INCR.f-list
-rw- 1 root other 1489728 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgDir0
-rw- 1 root other 0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgExtraObj0
```

```
-rw- 1 root other 1280176 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgFile0
-rw- 1 root other 192 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgHeader0
-rw- 1 root other 0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgNDMP0
-rw- 1 root other 9112680 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgRecord0
-rw- 1 root other 2111864 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgStrings0
-rw- 1 root other 11 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgUserGroupNames0
```

## About the NetBackup relational database

NetBackup installs Sybase SQL Anywhere during the master server installation as a private, non-shared server for the NetBackup database. Also known as the Enterprise Media Manager (EMM) database, the NetBackup database (NBDB) contains information about volumes and the robots and drives that are in NetBackup storage units.

The same installation of Sybase SQL Anywhere is used for the optionally-licensed product, Bare Metal Restore (BMR) database. The BMRDB database contains the information that the NetBackup Bare Metal Restore option manages. The BMR database is created during the BMR installation process.

As part of the catalog backup, the database and the configuration files for the NBDB and the BMRDB databases are protected as follows:

### ■ Database files

- *Install\_path\VERITAS\NetBackupDB\data\BMRDB.db* (if BMR is installed)
- *Install\_path\VERITAS\NetBackupDB\data\BMRDB.log* (if BMR is installed)
- *Install\_path\VERITAS\NetBackupDB\data\BMR\_DATA.db* (if BMR is installed)
- *Install\_path\VERITAS\NetBackupDB\data\BMR\_INDEX.db* (if BMR is installed)
- *Install\_path\VERITAS\NetBackupDB\data\DARS\_DATA.db*
- *Install\_path\VERITAS\NetBackupDB\data\DARS\_INDEX.db*
- *Install\_path\VERITAS\NetBackupDB\data\DBM\_DATA.db*
- *Install\_path\VERITAS\NetBackupDB\data\DBM\_INDEX.db*
- *Install\_path\VERITAS\NetBackupDB\data\EMM\_DATA.db*
- *Install\_path\VERITAS\NetBackupDB\data\EMM\_INDEX.db*
- *Install\_path\VERITAS\NetBackupDB\data\NBDB.db*
- *Install\_path\VERITAS\NetBackupDB\data\NBDB.log*

---

**Note:** NetBackup does not support saving the NetBackup relational database (NBDB, including the EMM data) or the configuration files to a remote file system such as NFS or CIFS.

---

■ Configuration files

- *Install\_path\VERITAS\NetBackupDB\data\vxdbms.conf*
- *Install\_path\VERITAS\NetBackupDB\conf\server.conf*
- *Install\_path\VERITAS\NetBackupDB\conf\databases.conf*

---

**Note:** The catalog backup process copies this data to *Install\_path\VERITAS\NetBackupDB\staging* and backs up the copy.

---

## About the Enterprise Media Manager (EMM) database

The Enterprise Media Manager (EMM) database contains information about media and the robots and drives that are in NetBackup storage units. The NetBackup Resource Broker queries the EMM database to allocate storage units, drives (including drive paths), and media. The host on which the EMM database resides is called the EMM server.

The EMM database contains the following information:

- Device attributes
- Robotic library and stand-alone drive residence attributes
- NDMP attributes
- Barcode rule attributes
- Volume pool attributes
- Tape attributes
- Media attributes
- Storage unit attributes
- Storage unit group attributes
- Hosts with assigned tape drives
- Media and device errors
- Disk pool and disk volume attributes
- Storage server attributes

- Logon credentials for storage servers, disk arrays, and NDMP hosts
- Fibre Transport attributes

The EMM database ensures consistency between drives, robotic libraries, storage units, media, and volume pools across multiple servers. The EMM database contains information for all media servers that share devices in a multiple server configuration.

The NetBackup scheduling components use the EMM database information to select the server, drive path, and media for jobs. When the device manager `ltid` starts up, it reads device information from the EMM database into a shared memory segment. Components on the same host communicate by using shared memory IPC or socket protocols. Socket protocols are used between components across multiple hosts. Command line interfaces are available to obtain run-time (shared memory) information and static device configuration information.

See “[About the Enterprise Media Manager](#)” on page 723.

## Protecting the catalog

In order for NetBackup to restore any file, NetBackup needs information from the catalog to determine where the backup for the file is located. Without a catalog, NetBackup cannot restore data.

Because the catalog plays an integral part in a NetBackup environment, a special type of backup protects the catalog. A catalog backup backs up catalog-specific data as well as produces disaster recovery information.

A catalog backup is configured separately from regular client backups by using the Catalog Backup Wizard. The catalog can be stored on a variety of media.

Configure a catalog backup before you run any regular backups.

---

**Note:** If portions of the catalog are relocated, note the changes so that subsequent catalog backups are aware of the locations of all the catalog components. In the event that a catalog recovery is needed, the same alterations must be implemented before the recovery of the catalog.

---

As additional protection for the catalog, consider archiving the catalog.

See “[Archiving the catalog](#)” on page 616.

The *NetBackup Troubleshooting Guide* provides helpful setup information to aid in disaster recovery. Since the catalog plays a critical role in the NetBackup environment, much of the information concentrates on catalog considerations.

## About online, hot catalog backups

In NetBackup 7.0, all catalog backups are online, hot catalog backups. The ability to configure offline, cold catalog backups has been removed in this release.

The online, hot catalog backup is designed for active environments in which continual backup activity occurs. It is considered an online, hot method because it can be performed while regular backup activity occurs.

The online, hot catalog backup is policy-based so it has all of the scheduling flexibility of a regular backup policy. Because the policy allows for incremental backups, catalog backup times for large catalogs can be significantly reduced.

Online, hot catalog backups use media from the **CatalogBackup** volume pool only.

The online, hot catalog backup performs the following tasks:

- Backs up the catalog while continual client backups are in progress
- Spans multiple tapes for a catalog backup
- Allows for a flexible pool of catalog tapes
- Performs a full or an incremental catalog backup
- Restores the catalog to a different location
- Runs scheduled catalog backups
- Appends to existing data on tape

You can configure an online catalog backup by using one of the following methods:

- By using wizards:
  - The Catalog Backup Wizard.  
See [“Using the Catalog Backup Wizard”](#) on page 605.
  - The Backup Policy Configuration Wizard.  
See [“Using the Backup Policy Wizard to configure a catalog backup”](#) on page 608.
- By creating a backup policy manually and indicating the **NBU-Catalog** policy type.  
See [“Using the Backup Policy Wizard to configure a catalog backup”](#) on page 608.

### Using the Catalog Backup Wizard

Catalog backups write only to media in the **CatalogBackup** volume pool. This procedure assumes that a storage device is configured and media is available in the **CatalogBackup** volume pool.

See [“About adding volumes”](#) on page 257.

### To use the Catalog Backup Wizard to configure a catalog backup

- 1 Click **Configure the Catalog Backup** in the right pane to launch the **NetBackup Catalog Backup Wizard**. The wizard is visible when either the **Master Server** or the **NetBackup Management** node is selected in the left pane.

Click Help within any wizard screen for more information on the wizard settings.

- 2 Click **Next** on the Welcome screen.
- 3 On the **NetBackup Catalog Backup Policy** screen, select a policy from the list of existing catalog backup policies.
- 4 Or, to create a new catalog backup policy, select **Create a new catalog backup policy**.
- 5 Click **Next** to launch the **Policy Name and Type** screen of the **Backup Policy Configuration Wizard**.
- 6 In the **Policy Name and Type** wizard screen, enter the policy name. Notice that **NBU-Catalog** is automatically selected as the policy type.

Type a unique name for the new policy in the **Add a New Policy** dialog box.

See “[NetBackup naming conventions](#)” on page 719.

Click **Next**.

- 7 On the **Backup Type** wizard screen, select the backup type. The **User Backup** does not apply for NBU-Catalog policies. Click **Next**.
- 8 On the **Rotation** wizard screen, select the rotation schedule. By default, a frequency-based schedule is selected. A frequency-based schedule ensures that the catalog backup has an opportunity to run in busy environments where backup jobs are running.

The selection **After each backup session** refers to a period when no regular backup policy is running.

Catalog backups can be scheduled to run concurrently with other backup types on the master server.

See “[Running online, hot catalog backups concurrently with other backups](#)” on page 613.

Click **Next**.

- 9 In the **Start Window** wizard screen, define a window of time during which the catalog backup can start and click **Next**. The scheduled windows (**Off hours, Working hours, All day, Custom**) are preset in the wizard. To change these settings, first complete the wizard. Then, select the policy in the **Policies** utility.

User Window selections are disabled, as regular users (those who are not NetBackup administrators) cannot start catalog backups.

- 10 On the **Catalog Disaster Recovery File** wizard screen, enter the path where each disaster recovery image file can be saved on disk. The image file contains the disaster recovery information. Enter the logon and password information, if necessary.

Symantec recommends that you save the image file to a network share or a removable device. Do not save the disaster recovery information to the local machine.

See “[Path](#)” on page 570.

See “[Logon](#)” on page 571.

See “[Password](#)” on page 571.

Click **Next**.

- 11 Symantec recommends that you configure the NetBackup environment to send the disaster recovery information to a NetBackup administrator. This backup-specific information is sent after every catalog backup.

On the **E-mail Disaster Recovery Information** wizard screen, enter one or more addresses. To send the information to more than one administrator, separate multiple email addresses using a comma as follows:

*email1@domain.com, email2@domain.com*

Make sure that email notification is enabled in your environment.

See “[Disaster recovery emails and the disaster recovery file](#)” on page 615.

See “[Send in an email attachment field](#)” on page 571.

---

**Note:** The disaster recovery email is not sent to the address that is specified in the **Global Attributes** properties. The **Administrator’s email Address** in the **Global Attributes** properties specifies the addresses where NetBackup sends notifications of scheduled backups or administrator-directed manual backups.

---

- 12 The last screen of the **Policy Wizard** describes that once the policy is created, you can make changes in **NetBackup Management > Policies**. Click **Finish** to create the policy.
- 13 The Catalog Backup Wizard resumes, with the new catalog backup policy listed.
- 14 Click **Next** to finish the **Catalog Backup Wizard**.
- 15 The final Catalog Backup Wizard screen displays the total number of catalog backup policies for this master server. Click **Finish** to complete the wizard.
- 16 You may want to add critical policies to the **Critical Policies** list. Specify some policies as critical policies after the **Catalog Backup Wizard** is complete. A policy that is listed on the **Critical Policies** list is considered crucial to the recovery of a site in the event of a disaster.

See [“Adding policies to the Critical Policies list of a catalog backup policy”](#) on page 572.

The NetBackup **Disaster Recovery** report lists the media that is used for backups of critical policies. The reports lists the media for only incremental and full backup schedules, so critical policies should use only incremental or full backup schedules.

## Using the Backup Policy Wizard to configure a catalog backup

Catalog backups write only to media in the **CatalogBackup** volume pool. This procedure assumes that a storage device is configured and media is available in the **CatalogBackup** volume pool.

See [“About adding volumes”](#) on page 257.

### To use the Backup Policy Wizard to configure a catalog backup

- 1 Click **Create a Backup Policy** in the right pane to launch the **Backup Policy Configuration Wizard**. The wizard is visible when either the **Master Server** or the **NetBackup Management** node is selected in the left pane.  
  
Click **Help** within any wizard screen for more information on the wizard settings.
- 2 Click **Next** on the Welcome screen.
- 3 In the **Policy Name and Type** wizard screen, enter the policy name. Select **NBU-Catalog** as the policy type.  
  
Click **OK**.

See [“NetBackup naming conventions”](#) on page 719.



- 4 On the **Backup Type** wizard screen, select the backup type. The **User Backup** does not apply for NBU-Catalog policies. Click **Next**.
- 5 On the **Rotation** wizard screen, select the rotation schedule. By default, a frequency-based schedule is selected. A frequency-based schedule ensures that the catalog backup has an opportunity to run in busy environments where backup jobs are running.

The selection **After each backup session** refers to a period when no regular backup policy is running.

Catalog backups can be scheduled to run concurrently with other backup types on the master server.

See [“Running online, hot catalog backups concurrently with other backups”](#) on page 613.

Click **Next**.

- 6 In the **Start Window** wizard screen, define a window of time during which the catalog backup can start and click **Next**. The scheduled windows (**Off hours, Working hours, All day, Custom**) are preset in the wizard. To change these settings, first complete the wizard. Then, select the policy in the **Policies** utility and customize the settings.

User Window selections are disabled, as regular users (those who are not NetBackup administrators) cannot start catalog backups.

- 7 On the **Catalog Disaster Recovery File** wizard screen, enter the path where each disaster recovery image file can be saved on disk. The image file contains the disaster recovery information. Enter the logon and password information, if necessary.

Symantec recommends that you save the image file to a network share or a removable device. Do not save the disaster recovery information to the local machine.

See [“Path”](#) on page 570.

See [“Logon”](#) on page 571.

See [“Password”](#) on page 571.

Click **Next**.

- 8 Symantec recommends that you configure the NetBackup environment to send the disaster recovery information to a NetBackup administrator. This backup-specific information is sent after every catalog backup.

To send the information to more than one administrator, separate multiple email addresses using a comma as follows:

```
email1@domain.com, email2@domain.com
```

Make sure that email notification is enabled in your environment.

See [“Disaster recovery emails and the disaster recovery file”](#) on page 615.

See [“Send in an email attachment field”](#) on page 571.

---

**Note:** The disaster recovery email is not sent to the address that is specified in the **Global Attributes** properties. The **Administrator’s email Address** in the **Global Attributes** properties specifies the addresses where NetBackup sends notifications of scheduled backups or administrator-directed manual backups.

---

- 9 Click **Finish** to complete the wizard.
- 10 You may want to add critical policies to the **Critical Policies** list. Specify some policies as critical policies after the **Backup Policy Wizard** is complete. A policy that is listed on the **Critical Policies** list is considered crucial to the recovery of a site in the event of a disaster.

See [“Adding policies to the Critical Policies list of a catalog backup policy”](#) on page 572.

The NetBackup **Disaster Recovery** report lists all of the media that is used for backups of critical policies, including the most recent full backup. The reports lists the media for only incremental and full backup schedules, so critical policies should use only incremental or full backup schedules.

## Configuring a catalog backup manually

You can configure a catalog backup manually by using the **Policy** utility. This procedure assumes that a storage device is configured and media is available in the **CatalogBackup** volume pool.

### To configure an online, hot catalog backup manually

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select **Actions > New > Policy**.

- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.  
See [“NetBackup naming conventions”](#) on page 719.  
Click **OK**.
- 4 On the **Attributes** tab, complete the following entries:
  - **Policy Type**  
Select **NBU-Catalog** as the policy type.
  - **Policy storage**  
For disk storage units, increase the **Maximum Concurrent Jobs** storage unit setting to ensure that the catalog backup can proceed during regular backup activity.  
See [“Maximum concurrent jobs setting”](#) on page 383.

---

**Note:** The media server that is used for catalog backups must be at the same NetBackup version as the master server. If your installation contains media servers of various levels, do not select **Any Available** for the destination **Policy Storage Unit**. If media servers are at various version, a media server at a level other than the master server could be selected.

---

- **Policy volume pool**  
NetBackup automatically creates a **CatalogBackup** volume pool that is selected by default only for **NBU-Catalog** policy types.
  - For other policy attribute descriptions, see the following topic:  
See [“About the Policy attributes”](#) on page 461.
- 5 Select the **Schedules** tab to set up a schedule for an online catalog backup.  
See [“Running online, hot catalog backups concurrently with other backups”](#) on page 613.  
See [“About catalog policy schedules”](#) on page 613.  
See [“About the Schedules tab”](#) on page 490.

---

**Note:** The Clients tab does not apply to the **NBU-Catalog** policy and does not appear.

---

- 6 The **Disaster Recovery** tab appears for **NBU-Catalog** policies only.  
The tab contains information regarding the location of data crucial to disaster recovery:

- See “[Path](#)” on page 570.  
Enter the path where each disaster recovery image file can be saved on disk. The image file contains the disaster recovery information. Enter the logon and password information, if necessary.  
Symantec recommends that you save the image file to a network share or a removable device. Do not save the disaster recovery information to the local machine.
  - See “[Logon](#)” on page 571.
  - See “[Password](#)” on page 571.
  - See “[Send in an email attachment field](#)” on page 571.
- 7 You may want to add critical policies to the **Critical Policies** list. The **Critical Policies** list contains the names of policies that back up critical data. Media that contains critical policy backups is listed on the **NetBackup Disaster Recovery Report** that is generated when the online catalog backup is run. The reports lists the media for only incremental and full backup schedules, so critical policies should use only incremental or full backup schedules.
- See “[Adding policies to the Critical Policies list of a catalog backup policy](#)” on page 572.
- Click **OK** to save the policy.

## Backing up catalogs manually

Catalog backups typically run automatically per the NBU-Catalog policy. However, a catalog backup can be started manually.

A manual catalog backup is useful in the following situations:

- To perform an emergency backup. For example, if the system is schedule to be moved and you cannot wait for the next scheduled catalog backup.
- If there is only one stand-alone drive and the stand-alone drive is used for catalog backups. In this situation, automatic backups are not convenient. The catalog backup tape must be inserted before each catalog backup and removed when the backup is done. (The tape swap is necessary because NetBackup does not mix catalog and regular backups on the same tape.)

### To perform a manual online, hot catalog backup

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 Select the catalog backup policy you want to run.
- 3 Select **Actions > Manual Backup**.

See [“Performing manual backups”](#) on page 574.

You can also run the `bpbbackup` command from the command line to perform an online, hot catalog backup.

More information is available in *NetBackup Commands*.

### Running online, hot catalog backups concurrently with other backups

You can schedule online, hot catalog to run concurrently with other backup types for the master server.

Make the following adjustments to ensure that the catalog backup can proceed while regular backup activity occurs:

- Set the **Maximum jobs per client** value to greater than one. The property is found in the Global Attributes host properties for the master server. See [“Global Attributes properties”](#) on page 131.
- Increase the **Maximum concurrent jobs** setting on the storage unit where the backups are sent. See [“Maximum concurrent jobs setting”](#) on page 383.

### About catalog policy schedules

The following items are notes about catalog policy schedules:

- The following schedules are supported in the online, hot catalog backup policy type:
  - Full
  - Differential incremental (depends on a full schedule)
  - Cumulative incremental
  - Session-based differential incremental
  - Session-based cumulative incremental
- Symantec recommends that only one catalog backup policy be configured.

- The media server that is used for catalog backups must be at the same NetBackup version as the master server.
- The incremental schedule depends on a full schedule.
- The least frequent schedule runs if many schedules are due at the same time.
- One catalog backup policy can contain multiple session-based incremental schedules:
  - If one is cumulative and the others are differential, the cumulative runs when the backup session ends.
  - If all are cumulative or all are differential, the first schedule that is found runs when the backup session ends.
- The queued scheduled catalog backup is skipped if a catalog backup job from the same policy is running.
- Session end means that no jobs are running. (This calculation does not include catalog backup jobs.)
- The Vault catalog backup is run whenever triggered from Vault, regardless of whether a catalog backup job is running from the same policy.
- When an online catalog backup is run, it generates three jobs: a parent job, a child job for NetBackup relational database tables, and a child job for catalog images and configuration data. The child jobs contain the actual backed up data. Consider both child jobs to duplicate, verify, or expire the backup.

---

**Note:** Additional child catalog jobs are created for the BMR database if a remote EMM server is configured.

---

## How do I know if a catalog backup succeeded?

The All Log Entries, Problems, and Media Log reports, available from the Reports utility, provide information on NetBackup catalog backups. In addition, you can use email.

An email message is sent to the address that is indicated in the **Disaster Recovery** settings for an online catalog backup.

Configure this email with the `mail_dr_info.cmd` script.

See [“Disaster recovery emails and the disaster recovery file”](#) on page 615.

See the *Administrator’s Guide, Volume II* for more information on setting up this script.

## Strategies that ensure successful catalog backups

Use the following strategies to ensure successful catalog backups:

- Use only the methods that are described in this chapter to back up the catalogs. The methods that are described here are the only operations that can track all relevant NetBackup activities and ensure consistency between the catalog files.
- Back up the catalogs often. If catalog backup files are lost, the changes that were made between the last catalog backup and the time of the disk crash are lost.
- Do not use methods other than NTFS compression to compress the catalogs or NetBackup may not be able to read them.
- Never manually compress the catalogs or NetBackup may be unable to restore the catalogs using `bprecover`.
- If you back up your catalogs to disk (not recommended), always back up to a different disk than where the catalog files reside. If you back up the catalog to the disk where the actual catalog resides, both catalog backups are lost if the backup disk fails. Recovering the catalog is much more difficult. Also, ensure that the disk has enough space for the catalogs. Backups to a full disk fail.
- The NetBackup binary image catalog is more sensitive to the location of the catalog. Catalog backups that are stored on a remote file system may have critical performance issues. NetBackup does not support saving catalogs to a remote file system such as NFS or CIFS.

---

**Note:** The catalog backup tape must be removed when the backup is finished or regular backups cannot proceed. NetBackup does not mix catalog and regular backups on the same tape.

---

## Recovering the catalog

Catalog recovery is discussed in the *NetBackup Troubleshooting Guide*.

## Disaster recovery emails and the disaster recovery file

The **Catalog Backup Wizard** and the **Backup Policy Wizard** prompt you to send the disaster recovery information to an email address. If the catalog backup is

configured manually using the Policy utility, this information appears on the **Disaster Recovery** tab.

The disaster recovery email and the accompanying attachment that is sent contain the following important items for a successful catalog recovery:

- A list of the media that contains the catalog backup
- A list of critical policies.
- Instructions for recovering the catalog
- The image file as an attachment.

If a catalog backup policy included both full backups and incremental backups, the attached image file can be a full or an incremental catalog backup.

Recovering from an incremental catalog backup completely recovers the entire catalog if the **Automatically recover the entire NetBackup catalog** option is selected on the wizard screen. The entire catalog is recovered because the incremental catalog backup references information from the last full backup. You do not need to recover the last full catalog backup before you recover the subsequent incremental backups.

You can tailor the disaster recovery email process by providing the `mail_dr_info.cmd` script in the `Install_path\VERITAS\NetBackup\bin` directory. This script is similar to the `nbmail.cmd` script. See the comments in the `nbmail.cmd` script for use instructions.

See [“About the Disaster Recovery tab”](#) on page 569.

## Archiving the catalog

The catalog archiving feature helps users tackle the problems that large amounts of catalog data can pose: large catalogs require a greater amount of disk space and can be time-consuming to back up. Catalog archiving reduces the size of online catalog data by relocating the large catalog `.f` files to secondary storage. NetBackup administration continues to require regularly scheduled catalog backups, but the backups are faster without the large amount of online catalog data.

Catalog archiving is available on both UNIX and Windows platforms.

---

**Note:** When you consider whether to archive the `.f` files, note that additional time is required to mount the tape and perform the restore.

---

Catalog archiving operations must be performed when NetBackup is in an inactive state (no jobs are running).



## To archive the catalog

- 1 Create a policy named **catarc** to reflect that the purpose of the schedule is for catalog archiving.

See [“Creating a catalog archiving policy”](#) on page 618.

- 2 Run `bpcatlist` to display images available for archiving.

Running `bpcatlist` alone does not modify any catalog images. Only when the `bpcatlist` output is piped to `bpcatarc` and `bpcatrm` are the images modified and the image `.f` files removed.

- 3 Determine the images that were previously archived by running:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist -online
```

The command returns the following message if catalog archiving was not performed previously: No entity was found.

- 4 Once the `bpcatlist` output correctly lists all the images to be archived, pipe the output through `bpcatarc` and `bpcatrm`. For example:

```
bpcatlist -client all -before Jan 1 2009 | bpcatarc | bpcatrm
```

The command waits until the backup completes successfully before the command returns the prompt. An error is reported if the catalog archive fails.

The Activity Monitor displays a Job ID for the job. The File List for the job (double-click the job in the Activity Monitor) displays a list of image files that were processed. When the job completes with a status 0, `bpcatrm` removes the corresponding `.f` files. If the job fails, no catalog `.f` files are removed.

- 5 Restore the catalog archive by doing the following:

- Use `bpcatlist` to list the files that need to be restored.
- After the `bpcatlist` command displays the proper files to restore, run `bpcatres` to restore the actual files.

To restore all the archived files from step 2, run the following command:

```
bpcatlist -client all -before Jan 1 2009 | bpcatres
```

This command restores all the catalog archive files before Jan 1, 2009.

See [“Catalog archiving commands”](#) on page 618.

## Creating a catalog archiving policy

The catalog archiving feature requires the presence of a policy named **catarc** before the catalog archiving commands can run properly. The policy can be reused for catalog archiving.

### To create a catalog archiving policy

- 1 Create a new policy and name it **catarc**. The **catarc** policy waits until `bpcatarc` can activate it. Users do not run this policy. Instead, `bpcatarc` activates this special policy to perform a catalog backup job, then deactivates the policy after the job is done.
- 2 Set the backup type on the **Attributes** tab. The type of backup that is indicated for the catalog archive policy must be **User Backup**.
- 3 Deactivate the catalog archive policy by clearing the **Go into effect at** field on the **Attributes** tab of the Policy dialog.
- 4 Set the retention level of the catalog archive for a time at least as long as the longest retention period of the backups being archived. Data can be lost if the retention level of the catalog archive is not long enough.

You may find it useful to set up, then designate a special retention level for catalog archive images.

- 5 Set a schedule for **catarc**. The schedule for **catarc** must include in its window the time `bpcatarc` command is run. If the `bpcatarc` command is run outside of the schedule that is indicated in `catarc`, the operation fails.
- 6 On the **Files** tab, browse to the directory where catalog backup images are placed:

```
Install_path\NetBackup\db\images
```

- 7 On the **Clients** tab, enter the name of the master server.
- 8 Save the policy.

## Catalog archiving commands

The catalog archiving option relies on three commands to designate a list of catalog `.f` files, then archive the files. A fourth command, `bpcatres`, is used to restore the files if necessary.

### **bpcatlist** command

The `bpcatlist` command queries the catalog data. Then, `bpcatlist` lists the portions of the catalog that are based on selected parameters. For example, date,

client, policy, schedule name, backup ID, the age of the backup image, or the date range of the backup image. `bpcatlist` outputs the formatted image summary information of matched images to standard output.

The other catalog archiving commands, `bpcatarc`, `bpcatrm`, and `bpcatres`, all depend on input from `bpcatlist` by a piped command.

For example, to archive (backup and delete) all of the `.f` files that were created before January 1, 2010, the following would be entered:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist -client all
-before Jan 1 2009 | bpcatarc | bpcatrm
```

`bpcatlist` is also used to provide status information.

For each catalog, it lists the following information:

- Backup ID (**Backupid**)
- Backup date (**Backup Date**)
- Catalog archive ID (**catarcid**). After one `.f` file is successfully backed up, a catalog archive ID is entered into the **catarcid** field in the image file. This field is zero if the image was never archived.
- Archived status (S), indicating if the catalog was not archived (1) or was archived (2)
- Compressed status (C), indicating if the catalog is not compressed (0) or compressed (1)
- Catalog file name (Files file)

The following is an example of the `bpcatlist` output, showing all of the backups for client alpha since October 23:

```
bpcatlist -client alpha -since Oct 23
Backupid Backup Date ...Catarcid S C Files file
alpha_0972380832 Oct 24 10:47:12 2009 ... 973187218 1 0 alpha_0972380832_UBAK.f
alpha_0972336776 Oct 23 22:32:56 2009 ... 973187218 1 0 alpha_0972336776_FULLL.f
alpha_0972327197 Oct 23 19:53:17 2009 ... 973187218 1 0 alpha_0972327197_UBAK.f
```

More information is available in *NetBackup Commands*.

## bpcatarc command

The `bpcatarc` command reads the output from `bpcatlist` and backs up the selected list of `.f` files. After one `.f` file is successfully backed up, a catalog archive ID is entered into the **catarcid** field in the image file. For archiving of the `.f` files to proceed, a policy by the name of **catarc** is required. The policy is based on a **User**

**Backup** type schedule. The schedule for **catarc** must include in its window the time `bpcatarc` command is run.

See “[Creating a catalog archiving policy](#)” on page 618.

## bpcatrm command

The `bpcatrm` command reads the output from `bpcatlist` or `bpcatarc`. If the image file has valid **catarcid** entries, `bpcatrm` deletes selected image `.f` files from the online catalog.

`bpcatrm` does not remove one `.f` file unless the file has been previously backed up using the **catarc** policy.

## bpcatres command

Use the `bpcatres` command to restore the catalog. The `bpcatres` command reads the output from `bpcatlist` and restores selected archived `.f` files to the catalog. For example:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist
-client all -before Jan 1 2007 | bpcatres
```

## When to catalog archive

Consider the following items before catalog archiving:

- Perform catalog archiving operations when NetBackup is in an inactive state (no jobs are running).
- To ensure that catalog backup images are not on the same tapes as user backups, create a separate media pool for catalog archives.
- You may find it useful to set up and then designate, a special retention level for catalog archive images.

To specify retention levels, go to **Host Properties > Master Server > Retention Periods**.

See “[Retention Periods properties](#)” on page 172.

## Using Vault with the catalog archiving feature

Since the catalog archiving feature uses a regular **User Backup** schedule in the **catarc** policy, the files are duplicated and vaulted similarly to other backups.

## Extracting images from the catalog archives

The situation may arise in which a storage provider needs to extract all of a specific client's records. The storage provider can extract the customer images from the catalog archive by creating the archives that are based on client name.

### To extract images from the catalog archives based on a specific client

- 1 Create a volume pool for the client.
- 2 Create a catalog archiving policy. Indicate the volume pool for that client in the **Attributes** tab.
- 3 Run `bpcatlist` so only the `.f` files from that client are listed. For example:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist
-client clientname | bpcatarc | bpcatrm
```

- 4 If you do not want to write more images to the client's volume pool, change the volume pool before you run another archiving catalog.

## Estimating catalog space requirements

NetBackup requires disk space to store its error logs and information about the files it backs up.

The disk space that NetBackup needs varies according to the following factors:

- Number of files to be backed up
- Frequency of full and incremental backups
- Number of user backups and archives
- Retention period of backups
- Average length of full path of files
- File information (such as owner permissions)
- Average amount of error log information existing at any given time

### To estimate the disk space that is required for a catalog backup

- 1 Estimate the maximum number of files that each schedule for each policy backs up during a single backup of all its clients.
- 2 Determine the frequency and the retention period of the full and the incremental backups for each policy.

- 3 Use the information from steps 1 and 2 to calculate the maximum number of files that exist at any given time.

For example:

Assume that you schedule full backups to occur every seven days. The full backups have a retention period of four weeks. Differential incremental backups are scheduled to run daily and have a retention period of one week.

The number of file paths you must allow space for is four times the number of files in a full backup. Add to that number one week's worth of incremental backups.

The following formula expresses the maximum number of files that can exist for each type of backup (daily or weekly, for example):

Files per Backup × Backups per Retention Period = Max Files

For example:

A daily differential incremental schedule backs up 1200 files and the retention period for the backup is seven days. Given this information, the maximum number of files that can exist at one time are the following:

$$1200 \times 7 \text{ days} = 8400$$

A weekly full backup schedule backs up 3000 files and the retention period is four weeks. The maximum number of files that can exist at one time are the following:

$$3000 \times 4 \text{ weeks} = 12,000$$

Obtain the total for a server by adding the maximum files for all the schedules together. Add the separate totals to get the maximum number of files that can exist at one time. For example, 20,400.

For the policies that collect true image restore information, an incremental backup collects catalog information on all files (as if it were a full backup). This changes the calculation in the example: the incremental changes from  $1200 \times 7 = 8400$  to  $3000 \times 7 = 21,000$ . After 12,000 is added for the full backups, the total for the two schedules is 33,000 rather than 20,400.

- 4 Obtain the number of bytes by multiplying the number of files by the average length of the file's full paths and file information.

If you are unsure of the average length of a file's full path, use 100. Using the results from the examples in step 3 yields:

$$(8400 \times 150) + (12,000 \times 150) = 3060000 \text{ bytes (or about 2988 kilobytes)}$$

- 5 Add between 10 megabytes to 15 megabytes to the total sum that was calculated in step 4. The additional megabytes account for the average space that is required for the error logs. Increase the value if you anticipate problems.
- 6 Allocate space so all the data remains in a single partition.

## File size considerations

File system limitations include the following:

- Some UNIX systems have a large file support flag. Turn on the flag to enable large file support. For example, AIX disables large file support by default, so the file size limit is 2 GB.
- For UNIX systems, set the file size limit for the root user account to unlimited to support large file support.

## About the binary catalog format

The catalog in a binary file format has several advantages over the catalog in a text format:

- The catalog is more compact. The binary representations of numbers, dates, and other information, takes up less disk space than the text representations.
- The catalog is much faster to browse and search, especially for large file sizes.
- The catalog supports alternate backup methods without the need to post-process images, which improves catalog performance for alternate backup methods.

The following points describe size the limitations that are associated with the binary catalog:

- The maximum number of files that can be backed up per image:  
 $(2^{31}) - 1$  files = 2,147,483,647 files = 7FFFFFFF files
- The maximum number of different user IDs and group IDs (combined):  
 $(2^{31}) - 1$  IDs = 2,147,483,647 IDs = 7FFFFFFF IDs

## Moving the image catalog

An image catalog may become too large for its current location. Consider moving the image catalog to a file system or disk partition that contains more available space.

---

**Note:** NetBackup does not support saving the catalog to a remote file system. Therefore, Symantec advises against moving the image catalog to a remote file system such as NFS or CIFS.

---

---

**Note:** NetBackup only supports moving the image catalog to a different file system or disk partition. It does not support moving the other subdirectories that make up the entire NetBackup catalog. For example, do not use the `ALTPATH` mechanism to move `install_path\NetBackup\db\error`.

---

### To move the image catalog

- 1 Back up the NetBackup catalogs manually.

A backup of the catalogs ensures that you can recover image information in case something is accidentally lost during the move.

- 2 Check the **Jobs** tab in the Activity Monitor and ensure that no backups or restores are running for the client.

If jobs are running, either wait for them to end or stop them by using the **Jobs** tab in the Activity Monitor.

- 3 Use the **Services** tab in the Activity Monitor to stop the Request Manager and the Database Manager services. These services are stopped to prevent jobs from starting. Do not modify the database while this procedure is performed.

- 4 Create a file named `ALTPATH` in the image catalog directory.

For example, if NetBackup is installed in the default location and the client name is `mars`, the path to the image catalog is:

```
C:\Program Files\VERITAS\NetBackup\db\images\mars\ALTPATH
```

- 5 Create the directory to which you intend to move the image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

- 6 On the first line of the `ALTPATH` file, specify the path to the directory where you intend to move the client's image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

The path is the only entry in the `ALTPATH` file.



- 7 Move all files and directories (except the `ALTPATH` file) that are in the current client directory to the new directory.

For example, if the images are currently in

```
C:\Program Files\VERITAS\NetBackup\db\images\mars
```

and the `ALTPATH` file specifies

```
E:\NetBackup\alternate_db\images\mars
```

then move all files and directories (except the `ALTPATH` file) to

```
E:\NetBackup\alternate_db\images\mars
```

- 8 Start the NetBackup Request Manager and NetBackup Database Manager services by using the **Services** tab in the Activity Monitor.

Backups and restores can now resume for the client.

## Indexing the catalog for faster access to backups

If the NetBackup environment contains a large number of backups, consider indexing the catalogs to reduce the time that is required to restore files.

To index the catalog means to create indexes of the files that are recorded in the NetBackup image catalog. NetBackup uses the indexes to go directly to the catalog entry for a file. Without indexing, NetBackup must start searching for a file at the beginning of the catalog entries.

## Compressing the image catalog

The image catalog contains information about all client backups. It is accessed any time a user lists or restores files. NetBackup lets you compress all portions of the catalog or only older portions of the catalog. No method selectively compresses image-catalog files other than by age.

Control image-catalog compression by setting the Global Attributes property, **Compress Catalog Interval**. Use this property to specify how old the backup information must be before it is compressed. Specify the number of days to defer compression information, thus users who restore files from recent backups are unaffected. By default, **Compress Catalog Interval** is set to 0 and image compression is not enabled.

See [“Compress catalog interval”](#) on page 135.

---

**Note:** Symantec discourages manually compressing or decompressing catalog backups using `bpimage -[de]compress` or any other method. Manually compressing or decompressing a catalog backup while any backup (regular or catalog) is running results in inconsistent image-catalog entries. When users list and restore files, the results can be incorrect.

---

The time to perform compression depends on the server speed and the number and size of the files being compressed. Files are compressed serially, and temporary working space is required in the same partition.

The catalog must be in an NTFS partition for compression to occur. If you choose to compress the image catalog, NetBackup uses NTFS compression on the server to perform compression after each backup session. It does not make a difference to NetBackup if the backup session was successful. The operation occurs while NetBackup expires backups and before it runs the `session_notify` script and the backup of the NetBackup catalogs.

When numerous compressed image-catalog files must be processed, the backup session is extended until compression is complete. The additional backup time is especially noticeable the first time you perform the compression. To minimize the effect of the initial sessions, consider compressing the files in stages. For example, begin by compressing the records for the backups older than 120 days. Continue to reduce the number of days over a period of time until you reach a comfortable setting.

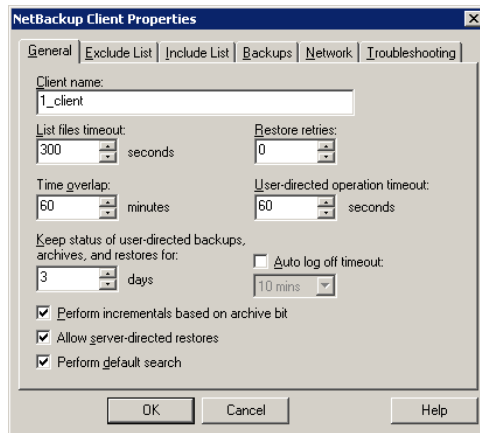
Compressing the image catalog accomplishes the following objectives:

- Reduces greatly the disk space that is consumed.
- Reduces the media that is required to back up the catalog.

The amount of space that is reclaimed varies with the types of backups you perform. Full backups result in a larger percentage of catalog compression than incremental backups. Normally, more data is duplicated in a catalog file for a full backup. Using catalog compression, a reduction of 80% is possible.

This reduction in disk space and media requirements is achieved at the expense of performance when a user lists or restores files. Since the information is uncompressed at each reference, performance degradation is in direct proportion to the number and size of compressed files that are referenced. If the restore requires numerous catalog files to be uncompressed, increase the timeout value that is associated with list requests.

Change the timeout value by changing the **List Files Timeout** General property setting on the client.

**Figure 17-1** List Files Timeout General property on the client

## Uncompressing the image catalog

You may find it necessary to uncompress all records temporarily that are associated with an individual client. Uncompress the records if you anticipate large or numerous restore requests, for example.

### Uncompressing the NetBackup catalog

Use the following procedure to uncompress the NetBackup catalog.

#### To uncompress the NetBackup catalog

- 1 Verify that the partition where the image catalog resides contains enough space to accommodate the uncompressed catalog.  
See “[Estimating catalog space requirements](#)” on page 621.
- 2 Stop the NetBackup Request Manager service, `bp_rmd`. Use the Activity Monitor or the Services application in the Windows Control Panel.
- 3 Verify that the NetBackup Database Manager, `bp_dbm`, is running.
- 4 In the NetBackup Administration Console, expand **NetBackup Management** > **Host Properties** > **Master Server**. Double-click the host to be uncompressed.
- 5 Select the **Global Attributes** properties.  
See “[Global Attributes properties](#)” on page 131.
- 6 Clear the **Compress Catalog Interval** check box and click **OK** to save the host property change.

- 7 Open a command prompt. Change to the following directory:

```
install_path\veritas\netbackup\bin\admincmd
```

Run one of the followings commands.

To decompress the records for a specific client, enter:

```
bpimage -decompress -client_name
```

To decompress the records for all clients, enter:

```
bpimage -decompress -allclients
```

- 8 Restart the NetBackup Request Manager service, bprd.

# About the NetBackup relational database

This chapter includes the following topics:

- [Installing the NetBackup relational database \(NBDB\)](#)
- [Post-installation tasks](#)
- [About backup and recovery procedures](#)
- [Unloading the database](#)
- [Terminating database connections](#)
- [About the Database Administration tool](#)
- [Moving the NetBackup database from one host to another](#)

## Installing the NetBackup relational database (NBDB)

This topic contains the information that describes the proper installation and operation of the Sybase SQL Anywhere relational database management system.

Generally, the implementation of Sybase SQL Anywhere in the NetBackup catalog is transparent. NetBackup installs Sybase SQL Anywhere during the master server installation as a private, non-shared server for the NetBackup database (NBDB). The NetBackup database (NBDB), contains the Enterprise Media Manager (EMM) data as well as other NetBackup data that NetBackup services use.

The same installation of Sybase SQL Anywhere is used for the optionally-licensed product, Bare Metal Restore (BMR) and its associated database (BMRDB). The BMR database is created during the BMR installation process.

By default, the NetBackup relational database (NBDB) is installed on the master server. The master server is also the default location for the Enterprise Media Manager (EMM) server. Since EMM is the primary user of NBDB, the NetBackup database always resides on the same machine as the Enterprise Media Manager.

See [“Enterprise Media Manager domain requirements”](#) on page 724.

For performance reasons, the EMM server and the relational database can be moved to another server.

See [“Moving NBDB database files after installation”](#) on page 639.

---

**Note:** NetBackup does not support saving the NetBackup relational database (NBDB, including the EMM data) to a remote file system such as NFS or CIFS.

---

The following procedure is performed automatically during installation in the order presented. The procedure can be performed manually to install the database independently.

#### NetBackup database installation

- 1 As part of the NetBackup master server installation, the SQL Anywhere server is created. The server parameters are set in the `server.conf` file in the following location:

```
Install_path\VERITAS\NetBackupDB\conf\server.conf
```

See [“About the server.conf file”](#) on page 632.

- 2 The following entry is added to the registry to set the database location:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config\VXDBMS_NB_DATA
```

See [“About the NetBackup configuration entry”](#) on page 637.

- 3 The VxDBMS configuration file for NetBackup is created. This file requires the read and write permissions of a Windows administrator:

```
Install_path\VERITAS\NetBackupDB\data\vxdbms.conf
```

- 4 The NetBackup database is created:

```
Install_path\VERITAS\NetBackupDB\data\NBDB.db
```

- 5 DBA password is set for the NetBackup database in `vxdbms.conf`:

```
VXDBMS_NB_PASSWORD = encrypted_password
```

- 6 A minimum of four additional database files are created with contiguous space pre-allocated:

- The NetBackup system database file that is mentioned in the following step:

*Install\_path\VERITAS\NetBackupDB\data\NBDB.db*

- The EMM database files:

*Install\_path\VERITAS\NetBackupDB\data\EMM\_DATA.db*

*Install\_path\VERITAS\NetBackupDB\data\EMM\_INDEX.db*

- The NetBackup transaction log, necessary for recovering the database:

*Install\_path\VERITAS\NetBackupDB\data\NBDB.log*

- 7 The SQL Anywhere accounts and schema are created for each of the NetBackup components that make use of the NetBackup database. (For example, EMM\_MAIN.)

- 8 The following command initializes the EMM data:

*Install\_path\VERITAS\Volmgr\bin\tpext.exe*

## NetBackup master server installation

SQL Anywhere is installed in the following directories:

- *Install\_path\VERITAS\NetBackupDB*

The files in *Install\_path\VERITAS\NetBackupDB\conf* can be shared within a cluster.

- *Install\_path\VERITAS\NetBackup\bin*

The contents of each directory are examined in the following topics.

### Relocating the NetBackup database

The NetBackup database, NBDB, and its associated files, is created on the master server by default. For performance reasons, NBDB can be moved to another host. Symantec recommends that NBDB be on the same host as the EMM server.

See [“Sharing an EMM server”](#) on page 724.

The NBDB database files can be moved from their default location in

*Install\_path\VERITAS\NetBackupDB\data*.

See [“Moving NBDB database files after installation”](#) on page 639.

---

**Note:** NetBackup does not support saving the NetBackup relational database (NBDB, including the EMM data) to a remote file system such as NFS or CIFS.

---

---

**Note:** If Bare Metal Restore is installed, BMRDB must be located on the master server.

---

## About the server.conf file

Symantec recommends that this file not be edited without assistance from technical support. NetBackup may not start if the `server.conf` file is edited.

`Install_path\VERITAS\NetBackupDB\conf\server.conf` is read when the SQL Anywhere service is started. The SQL Anywhere service gets all configuration information from this file:

```
-n NB_server_name -x tcpip(LocalOnly=YES;ServerPort=13785) -gd DBA
-gk DBA -gl DBA -gp 4096 -ti 0 -c 25M -ch 500M -cl 25M -zl -os 1M -o
"C:\Program Files\Veritas\NetBackupDB\log\server.log"
```

Where `server_name` indicates the name of the SQL Anywhere server. Each Sybase server has a unique name. Use the same name that was used during installation. If a fully qualified name was used at that time, use a fully qualified name here.

---

**Note:** If this name is changed, the Enterprise Media Manager cannot connect to the database.

---

|                                                       |                                                                                                                                                          |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-x tcpip(LocalOnly=YES;ServerPort=13785)</code> | Indicates what kind of connections are allowed in addition to shared memory. For example, local TCP/IP connections using port 13785.                     |
| <code>-gp 4096</code>                                 | Indicates the maximum page size (in bytes) for the database. This parameter is given during database creation.                                           |
| <code>-ct+</code>                                     | Indicates that character set translation is used. UTF8 encoding is used.                                                                                 |
| <code>-gd DBA -gk DBA -gl DBA</code>                  | Indicates that the DBA user is the account used to start, stop, load, and unload data.                                                                   |
| <code>-ti 0</code>                                    | Indicates the client idle time that is allowed before shut down. By default, no idle time is allowed, which prevents the database from shutting down.    |
| <code>-c 25M</code>                                   | Indicates the initial memory that is reserved for caching database pages and other server information. The value may be changed for performance reasons. |



|                                                                                        |                                                                                                                                                                                                    |
|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-ch 500M</code>                                                                  | Indicates the maximum cache size, as a limit to automatic cache growth. The value may be changed for performance reasons.                                                                          |
| <code>-cl 25M</code>                                                                   | Indicates the minimum cache size, as a limit to automatic cache resizing. The value may be changed for performance reasons.                                                                        |
| <code>-o</code><br><code><i>Install_path</i>\VERITAS\NetBackupDB\log\server.log</code> | Indicates the location of server output messages. The messages include start and stop events, checkpoints, error conditions, and cache changing size. This log is not managed, but growth is slow. |
| <code>-ec SIMPLE</code>                                                                | Indicates the encryption method. Default: SIMPLE.<br><br>NONE SIMPLE TLS (TLS_TYPE=cipher; [FIPS={Y N}])<br>CERTIFICATE=server-identity-filename;<br>CERTIFICATE=PASSWORD=password)                |

## About the databases.conf file

The `Install_path\VERITAS\NetBackupDB\conf\databases.conf` configuration file contains the locations of the main database files and the database names for automatic startup when the SQL Anywhere service is started. For example, if NBDB and BMRDB are both located on the master server in the default locations, `databases.conf` contains:

```
"C:\Program Files\VERITAS\NetBackupDB\data\NBDB.db" -n NBDB
"C:\Program Files\VERITAS\NetBackupDB\data\BMRDB.db" -n BMRDB
```

## About the registration.dat file

This file is created for use with Symantec OpsCenter.

It is created in the following location:

```
Install_path\VERITAS\NetBackupDB\conf\registration.dat
```

## About the bin directory

`NetBackup\bin` contains NetBackup-specific binaries and commands for administrating NBDB and BMRDB:

- `NbDbAdmin.exe`

This file launches the Database Administration tool, which provides administrators with a way to more easily perform the tasks based on the `nldb` commands.

See “[About the Database Administration tool](#)” on page 646.

- `create_nldb.exe`  
Used during installation and upgrades to create and upgrade the NetBackup database, NBDB.
- `nldb_admin.exe`  
Among other things, use `nldb_admin.exe` to change the DBA and NetBackup account passwords, or to start and stop individual databases.
- `nldb_backup.exe`  
Use to make an online backup of the SQL Anywhere database files to a file system directory.

---

**Note:** Using this command (or the Database Administration tool) to restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use this command (or the Database Administration tool) to restore the NetBackup catalog only as a precautionary measure.

---

- `nldb_move.exe`  
Use to change the location of the SQL Anywhere database files from the default location.
- `nldb_ping.exe`  
Displays the status of the SQL Anywhere database.
- `nldb_restore.exe`  
Use to recover from an online backup in a file system directory that was created using `nldb_backup`.
- `nldb_unload.exe`  
Use to create a dump of all or part of the NBDB database or the BMRDB database schema and data.
- `nldbms_start_server.exe`  
Use to start and stop the SQL Anywhere service.
- `nldb_upgrade.exe`  
Used internally to upgrade the NetBackup and BMR databases.

---

**Note:** Due to performance issues, NetBackup supports database files only on locally attached drives.

---

The commands are described in *NetBackup Commands* and the online Help.

## About the charsets directory

`Install_path\VERITAS\NetBackupDB\charsets` contains SQL Anywhere-specific information.

## About the data directory

`Install_path\VERITAS\NetBackupDB\data` is the default location of the database, NBDB, and contains the following files:

- `NBDB.db`  
Main NetBackup database file; considered a **dbspace**.
- `EMM_DATA.db`  
An additional **dbspace** that contains EMM data.
- `EMM_INDEX.db`  
Enhances the EMM database performance.
- `NBDB.log`  
The transaction log for the NetBackup database, necessary for recovery. `NBDB.log` is automatically truncated after a successful full or incremental online, hot catalog backup of the SQL Anywhere database.
- `vxdbms.conf`  
Contains the configuration information specific to the Sybase SQL Anywhere installation:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_NB_DATA = C:\Program Files\VERITAS\NetBackupDB\data
VXDBMS_NB_INDEX = C:\Program Files\VERITAS\NetBackupDB\data
VXDBMS_NB_TLOG = C:\Program Files\VERITAS\NetBackupDB\data
VXDBMS_NB_PASSWORD = encrypted_password
VXDBMS_ODBC_DRIVER = NB SQL Anywhere
```

The encrypted password that is used to log into both the DBA accounts for NBDB and BMRDB, and other data accounts is stored in `vxdbms.conf`.

The password is set to a default upon installation (`nbusql`). Symantec recommends that the password is changed after installation.

See “[Changing the database password](#)” on page 638.

If the encryption method was changed from the default (SIMPLE) in the `server.conf` file, change this file to reflect the corresponding encryption method.

- If BMR is installed, the directory also contains: `BMRDB.db`, `BMRDB.log` (transaction log for BMR), `BMR_DATA.db`, `BMR_INDEX.db`

## About the log directory

`Install_path\VERITAS\NetBackupDB\log` contains the SQL Anywhere server log file `server.log` that contains only Sybase logs.

## About the scripts directory

`Install_path\VERITAS\NetBackupDB\scripts` contains the SQL Anywhere scripts that are used to create the database. The directory also contains NetBackup SQL scripts that are used to create the EMM and other schemas.

---

**Note:** Do not edit the scripts that are located in this directory.

---

## About the staging directory

`Install_path\VERITAS\NetBackupDB\staging` is used as a temporary staging area during online, hot catalog backup and recovery.

## About the WIN32 directory

`Install_path\VERITAS\NetBackupDB\WIN32` contains SQL Anywhere commands and `.dll` files.

## About the java directory

`Install_path\VERITAS\NetBackupDB\java` is a directory used by Symantec OpsCenter.

## About the shared directory

`Install_path\VERITAS\NetBackupDB\shared` is a directory used by Symantec OpsCenter.

## About the NetBackup configuration entry

The `VXDBMS_NB_DATA` registry entry is a required entry and is created upon installation. The entry indicates the path to the directory where `NBDB.db`, `BMRDB.db`, and the `vxdbms.conf` file are located.

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\

 Config\VXDBMS_NB_DATA
```

## Sybase SQL Anywhere server management

Upon startup, the Sybase SQL Anywhere server uses the SQL Anywhere service to set the server parameters in the `server.conf` file. Then, the service starts the databases that are indicated in the `databases.conf` file.

To start and stop the Sybase SQL Anywhere service, use one of the following methods:

- In the NetBackup Administration Console, select **NetBackup Relational Database Manager** (SQLANYs\_VERITAS\_NB) in the Activity Monitor Services tab.
- From the Windows Service Manager, select **NetBackup Relational Database Manager** (SQLANYs\_VERITAS\_NB).
- `Install_path\VERITAS\NetBackup\bin\bpdown -e SQLANYs_VERITAS_NB`
- `Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB`

Individual databases can be started or stopped, while leaving the SQL Anywhere service to continue. To do so, use the Database administration tool or the following commands:

- `nbdb_admin [-start | -stop]`

Starts or stops NBDB without shutting down the SQL Anywhere server.

To see whether the database is up, enter `nbdb_ping`.

- `nbdb_admin [-start | -stop BMRDB]`

Starts or stops BMRDB without shutting down the SQL Anywhere server.

To see whether the BMRDB database is up, enter `nbdb_ping -dbn BMRDB`.

See “[About the Database Administration tool](#)” on page 646.

## Clustered environments

Sybase SQL Anywhere is supported in a clustered environment. Sybase SQL Anywhere failover is included with the NetBackup server failover solution. The

software is installed on all machines in the cluster, but the database files are created on a shared disk.

To facilitate the shared files, database and configuration files are installed on a shared drive.

Configuration files are stored in *Shared\_drive\VERITAS\NetBackupDB\conf*.

## Post-installation tasks

The tasks described in the following topics are optional and can be performed after the initial installation:

- Change the database password.  
See [“Changing the database password”](#) on page 638.
- Move NBDB and BMRDB database files (possibly to tune performance).  
See [“Moving NBDB database files after installation”](#) on page 639.
- Add a mirrored transaction log.  
See [“Adding a mirrored transaction log”](#) on page 640.
- Recreate NBDB.  
See [“Creating the NBDB database manually”](#) on page 640.

## Changing the database password

You can change the DBA and application password at any time. The password is encrypted using AES-128-CFB and stored in the `vxdbsms.conf` file. The permissions for the `vxdbsms.conf` file allow only a Windows administrator to read or write to it.

---

**Note:** Symantec recommends changing the password after installation.

---

The default password that is set during installation is `nbusql`. This password is used for NBDB and BMRDB and for all DBA and application accounts. (For example, `EMM_MAIN`.)

### To change the database password

- 1 Log on to the server as a Windows Administrator.
- 2 Use one of the following methods to change the database password:
  - Use the Database Administration tool.  
See [“About the Database Administration tool”](#) on page 646.

- Run the following command to update the `vxdbsms.conf` file with the new, encrypted string:

```
Install_path\NetBackup\bin\nbdb_admin -dba new_password
```

## Moving NBDB database files after installation

In the case of large databases, consider changing the location of the database files or splitting the database files into multiple directories to improve performance.

---

**Note:** Due to performance issues, NetBackup supports database files only on locally attached drives.

---

**Note:** Run a catalog backup to back up NBDB and BMRDB both before and after moving the database files.

---

### To move the NBDB and the BMRDB database files

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 3 Start the SQL Anywhere service by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYS_VERITAS_NB
```

- 4 Use one of the following methods to move the existing data, index, and transaction log files:

- Use the Database Administration tool.  
See “[About the Database Administration tool](#)” on page 646.

- Type the following command:

```
Install_path\VERITAS\NetBackup\bin\nbdb_move.exe
-data data_directory
-index index_directory -tlog log_directory
```

You can run the `nbdb_move` command at any time because it does not drop the database and recreate it. Thus, all data is preserved.

If a mirrored transaction log is in use, type the following command:

```
Install_path\VERITAS\NetBackup\bin\nbdb_move.exe -data
data_directory
-index index_directory -tlog log_directory
-mlog log_mirror_directory
```

- 5 Start all services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

- 6 Perform a catalog backup.

## Adding a mirrored transaction log

The transaction logs `NBDB.log` and `BMRDB.log` are critical files used to recover the SQL Anywhere databases.

For extra protection, you can use a mirrored transaction log. Create this mirrored log in a different directory from the original log.

### To create a mirrored transaction log

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 3 Start the SQL Anywhere service by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB
```

- 4 Use one of the following methods to create the mirrored transaction log:
  - Use the Database Administration tool.  
See [“About the Database Administration tool”](#) on page 646.

- Type the following command:

```
Install_path\NetBackup\bin\nbdb_move.exe
-mloglog_mirror_directory
```

To move the existing data, index, transaction log files, and create the mirrored transaction log, type the following command:

```
Install_path\NetBackup\bin\nbdb_move.exe
-datadata_directory-index index_directory -tlog
log_directory-mlog log_mirror_directory
```

- 5 Start up all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

- 6 Perform a catalog backup.

## Creating the NBDB database manually

The NBDB database is created automatically during NetBackup installation. However, it may be necessary during certain catalog recovery situations to create it manually by using the `create_nbdb` command.



---

**Note:** To recreate the database manually is not recommended in most situations.

---

**Note:** If the `NBDB.db` database already exists, the `create_nbdb` command does not overwrite it. If you want to move the database, move it by using the `nbdb_move` command.

---

### To create the NBDB database manually

- 1 Shut down all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 2 Start the SQL Anywhere service by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB
```

- 3 Run the following command:

```
Install_path\NetBackup\bin\create_nbdb.exe
```

- 4 Start up all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

- 5 The new `NBDB` database is empty and does not contain the `EMM` data that is loaded during a normal installation.

Make sure that you have the most current support for new devices before the data is repopulated. New devices are added approximately every two months.

- 6 Repopulate the `EMM` data by running the `tpext` utility. `tpext` updates the `EMM` database with new versions of device mappings and external attribute files.

```
Install_path\VERITAS\Volmgr\bin\tpext.exe
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads `EMM` data into the database.

### Additional create\_nbdb options

In addition to using the `create_nbdb` command to create the `NBDB` database, you also can use it to perform the actions described in the following topics. In each command, `NB_server_name` matches the name in `server.conf`.

See [“About the server.conf file”](#) on page 632.

- Drop the existing NBDB database and recreate it in the default location by typing the following command:

```
create_nbdb -drop[current_data_directory]
```

The `-drop` option instructs NetBackup to drop the existing NBDB database.

Provide the location of the current NBDB data directory, `current_data_directory`, if the default location is not used.

- Drop the existing NBDB database and do not recreate by typing the following command:

```
create_nbdb -db_server NB_server_name
-drop_only[current_data_directory]
```

Provide the location of the current NBDB data directory, `current_data_directory`, if the default location is not used.

- Drop the existing NBDB database and recreate it in the directories as specified by typing the following command:

```
create_nbdb -drop [current_data_directory] -data
data_directory-index index_directory -tlog log_directory
[-mloglog_mirror_directory]
```

If the NBDB database files were moved from the default location by using `nbdb_move`, use this command to recreate them in the same location. Specify `current_data_directory`.

If the location of `NBDB.db` changed from the default, `BMRDB.db` must also be recreated. The `BMRDB.db` files must reside in the same location as the NetBackup database files.

## About backup and recovery procedures

The online, hot catalog method can be performed while regular backup activity takes place.

It runs per a policy and is virtually transparent to the customer. Set up the policy by using either the Catalog Backup Wizard or the Policy Wizard.

Either wizard automatically includes all the necessary catalog files to include the database files (NBDB and BMRDB) and any catalog configuration files (`vxdbms.conf`, `server.conf`, `databases.conf`).

The online, hot catalog allows an administrator to recover either the entire catalog or pieces of the catalog. (For example, the databases separately from the image catalog.)

It offers an incremental backup. For Sybase SQL Anywhere, an incremental backup means a backup of the transaction log only. Transaction logs are managed automatically, truncated after each successful backup.

## Database transaction log

The transaction log for the NetBackup database is necessary for recovering the database. It is automatically truncated after a successful catalog backup.

The transaction log, `NBDB.log`, is located by default in the following directory:

```
Install_path\NetBackupDB\data\NBDB.log
```

The transaction log continues to grow until it becomes truncated. Catalog backups must run frequently enough so that the transaction log does not grow to fill the file system.

In addition to the default transaction log, a mirrored transaction log can be created for additional protection of NBDB by using:

The log is named in the following manner:

```
mirrored_log_directory\NBDB.m.log
```

The directory for the mirrored log should not be the same as the directory for the default transaction log. Ideally, the mirrored log should be located on a file system on a different physical disk drive.

If BMR is installed, a transaction log for BMRDB is also created by default in:

```
Install_path\NetBackupDB\data\BMRDB.log
```

It has an optional mirrored log in the following location:

```
mirrored_log_directory\BMRDB.m.log
```

The BMRDB transaction logs are backed up and truncated during the catalog backup along with the NBDB transaction logs.

---

**Note:** If a catalog backup is not run, the logs are not truncated. Truncation must be managed in this manner as it is critical to recovery of the database.

---

## About catalog recovery

Recovery scenarios include the following:

- A full recovery from a complete disaster  
Using the **Disaster Recovery** wizard, the databases are restored along with the image catalog to a consistent state.

- A recovery of the database files only  
 Using `bprecover`, the relational database files and configuration files can be restored and recovered.

Details about catalog recovery scenarios and procedures are available in the *NetBackup Troubleshooting Guide*.

## Backing up and recovering the relational databases

The recommended method to protect the relational databases is to use the catalog backup and recovery interfaces.

A temporary backup of the NBDB and BMRDB databases can be made for extra protection before database administration activities such moving or reorganizing the database files.

**Table 18-1** Commands to back up and recover relational databases

| Command                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>nbdb_backup.exe</code>  | <p>Use <code>nbdb_backup</code> to make either an online or an offline copy of the NBDB database files and the BMRDB database files in a directory. The transaction log is not truncated by using <code>nbdb_backup</code>. Transaction logs are managed only by using the catalog backup.</p> <pre><i>Install_path</i>\NetBackup\bin\nbdb_backup.exe [-dbn <i>database_name</i>] [-online   -offline] <i>destination_directory</i></pre> <p><code>-dbn <i>database_name</i></code> only backs up the specified database (NBDB or BMRDB).</p> <p><code>-offline</code> shuts down the database and access to the database. Connections to the database are refused at this time. The SQL Anywhere service does not shut down.</p> <p><b>Note:</b> Using this command (or the <b>Database Administration Tool</b>) to back up the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use this command (or the <b>Database Administration Tool</b>) to back up the NetBackup catalog only as a precautionary measure.</p> <p><b>Note:</b> The transaction logs are not truncated by using <code>nbdb_backup</code>. A catalog backup must be run to truncate the logs.</p> |
| <code>nbdb_restore.exe</code> | <p>Use <code>nbdb_restore</code> to recover from a database backup that was made using <code>nbdb_backup</code>.</p> <pre><i>Install_path</i>\NetBackup\bin\nbdb_restore.exe -recover <i>source_directory</i></pre> <p>Logs are recorded in the <code>\admin</code> directory.</p> <p><b>Note:</b> Using this command (or the <b>Database Administration Tool</b>) to restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use this command (or the <b>Database Administration Tool</b>) to restore the NetBackup catalog only as a precautionary measure.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Unloading the database

You can use the Database Administration tool or the `nbdb_unload` command line utility to dump the entire NetBackup or Bare Metal Restore databases, or individual tables (one `.dat` file is created for each table), or schema. Use either method to create a copy of the SQL Anywhere database that may be requested in some customer support situations.

There should be no active connections to the database when `nbdb_unload` is run.

See “[Terminating database connections](#)” on page 645.

When either method is used, a `reload.sql` script is generated. The script contains all the code that is required to recreate the database. Symantec Technical Support uses this script and the associated files to assist in support cases.

```
Install_path\NetBackup\bin\nbdb_unload.exe [-dbn database_name] [-t
table_list] [-s] destination_directory
```

In the script where:

- `-dbn database_name`  
`database_name` is NBDB (default) or BMRDB.
- `-t table_list`  
Must list the owner of the table, then the table name. For EMM, the account `EMM_MAIN` owns all tables.  
`nbdb_unload -t EMM_MAIN.EMM_Device, EMM_MAIN.EMM_Density`
- `-s`  
No data is dumped, only schema.
- `destination_directory`  
Specify the location where the dump is created.

## Terminating database connections

To eliminate concurrency problems, terminate all active connections to the database by shutting down NetBackup before running `nbdb_unload`.

**To terminate database connections**

- 1 Shut down all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 2 Start the SQL Anywhere service by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB
```

3 Use one of the following methods to terminate database connections:

- Use the Database Administration tool.  
See “[About the Database Administration tool](#)” on page 646.
- Run `nbdb_unload` and indicate the outputs (database name, table lists, or schema only) and the destination directory.

4 Stop the SQL Anywhere service by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdown -e SQLANYs_VERITAS_NB
```

5 Start up all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

Symantec does not recommend using `reload.sql` to make a copy of the relational databases in a production environment. Use the Database Administration tool or `nbdb_backup` to make a physical copy or use `nbdb_move` to relocate the database files.

## About the Database Administration tool

The Database Administration tool for NetBackup databases provides a way for administrators to perform the following administrative tasks more easily:

- Select, start, and stop the NetBackup relational database (NBDB) or the Bare Metal Restore database (BMRDB)
- Report on the status, consistency, and high-level space utilization
- Report on database space in full and table level reports
- Manage database space and memory cache
  - Full and table level defragmentation
  - Index compression
  - Add free space to the database
  - Adjust the database server memory cache size
- Manage the transaction log
  - Truncate the transaction log
  - Toggle the transaction log mode between full and minimum
- Check for database consistency using standard and full database validation
- Rebuild the database

- Move the database
  - Create or stop using a mirrored transaction log
- Unload the database
- Back up and restore the database
- Change the password
- Report on and change the database server settings

## Running the Database Administration Tool

The Database Administration tool is a stand-alone application (`NbDbAdmin.exe`) and is located in the following directory:

```
InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe
```

When the Database Administration tool starts, the administrator must enter the DBA password. If the DBA password is the default password that is used when NetBackup is installed, the administrator is encouraged to change the password. The administrator is not required to change the password, however.

The tool contains the **General** tab and the **Tools** tab. It also contains the following information when either tab is displayed:

- The database (displays NBDB information by default).
- The status of the database.
- The results from a validation check.
- Information on space utilization.
- **Drive Space** button
  - Shows the amount of free space and used space on a drive. If the database files are on multiple drives, this view is useful to see which drive has more free space available.
- **Help** button
  - Provides additional assistance in the console.

## Moving the NetBackup database from one host to another

The NetBackup database, NBDB, must always reside on the same host as the EMM server. If NBDB is moved, the EMM server must also be moved. The Bare Metal Restore database, BMRDB, must always reside on the master server. So, if NBDB

and EMM server are moved to a media server from a master server, BMRDB must remain on the master server.

Use the following procedure to move the NetBackup database (NBDB) from host A to host B. This procedure also reconfigures NetBackup so that the new database host becomes the EMM server.

If you move the NetBackup database and the EMM server to a different host in a cluster environment, see the following topic:

See [“Cluster considerations with the EMM server”](#) on page 650.

#### **To move the NetBackup database from one host to another**

- 1** Perform a catalog backup.
- 2** If NetBackup is currently installed on B, do the following:
  - Shut down NetBackup on B by typing the following command:  
*Install\_path/VERITAS/NetBackup/bin/bpdown*
  - Run the following command on B by typing the following command:  
*Install\_path/VERITAS/NetBackup/bin/nbdb\_relocate -make\_emmhost emmservername*  
This command is not internationalized.
  - Start the Sybase SQL Anywhere server on B by typing the following command:  
*Install\_path/VERITAS/NetBackup/bin/bpup -e SQLANYs\_VERITAS\_NB*
  - Create NBDB and associated files in the default location on B by typing the following command:  
*Install\_path/VERITAS/NetBackup/bin/create\_nbdb create\_nb -drop*  
If NetBackup is not installed on B, install NetBackup on B identifying B as the EMM server.
- 3** Set the database password on host B to match the password on A if the password has changed from the default. Use the Database Administration tool or type the following command:  
*Install\_path/VERITAS/NetBackup/bin/nbdb\_admin -dba password*
- 4** Shut down NetBackup on A and B and on all master servers and media servers using host A as the EMM server by typing the following command:  
*Install\_path/VERITAS/NetBackup/bin/bpdown*



- 5 Copy the following files from A to B to the final location (do not copy `vxdbms.conf`):

`NBDB.db`

`EMM_DATA.db`

`EMM_INDEX.db`

`NBDB.log`

`NBDB.m.log` (optional)

If the database files on B are in the default location

(`Install_path/VERITAS/NetbackupDB/data`) and server A is also running Windows, go to the following step:

11.

- 6 Change `databases.conf` on A and B so that the databases do not start automatically when the server is started by typing the following command:

```
Install_path/VERITAS/Netbackup/bin/ nbdb_admin -auto_start NONE
```

- 7 Start the Sybase SQL Anywhere server on B by typing the following command:

```
Install_path/VERITAS/NetBackup/bin/bpup -e SQLANYs_VERITAS_NB
```

- 8 Use the `nbdb_move` command on B to set the location of the database files by typing the following command:

```
nbdb_move -data dataDirectoryB -index indexDirectoryB-tlog
tlogDirectoryB [-mlog mlogDirectoryB] -config_only
```

- 9 Stop the Sybase SQL Anywhere server on B by typing the following command:

```
Install_path/VERITAS/NetBackup/bin/ bpdwn -e SQLANYs_VERITAS_NB
```

- 10 On B, delete the database files in the default directory if non-default locations are used for `dataDirectoryB`, `indexDirectoryB`, `tlogDirectoryB`, `mlogDirectoryB` by typing the following command:

`NBDB.db`

`EMM_DATA.db`

`EMM_INDEX.db`

`NBDB.log`

`NBDB.m.log` (optional)

- 11 Run the following command on A and on all master servers and media servers that used A as the EMM server:

```
Install_path/VERITAS/NetBackup/bin/ nbdb_relocate -change_emmhost
emmservername
```

This command is not internationalized.

- 12 On A, delete the following database files and configuration files:

```
NBDB.db
EMM_DATA.db
EMM_INDEX.db
NBDB.log
NBDB.m.log (optional)
```

- 13 On A do the following:

- If BMRDB does not exist on A, delete the configuration files by typing the following command:

```
Install_path/VERITAS/NetBackupDB/data/vxdbms.conf
Install_path/VERITAS/NetBackupDB/conf/databases.conf
Install_path/VERITAS/NetBackupDB/conf/server.conf
```

- If BMRDB exists on A, do the following:

Run the following command on A so that BMRDB starts automatically when the server is started:

```
Install_path/VERITAS/Netbackup/bin/ nbdb_admin -auto_start
BMRDB
```

- 14 Start NetBackup on B and on all master servers and media servers that use B as the EMM server.
- 15 Perform a catalog backup.

## Cluster considerations with the EMM server

If you move the NetBackup database and the EMM server to a different host in a Windows cluster environment, also be aware of the following:

See [“Moving the EMM server to a Windows cluster”](#) on page 651.

See [“About moving the EMM server from a Windows cluster”](#) on page 651.

## Moving the EMM server to a Windows cluster

If you move the NetBackup database and the EMM server to a different host in a Windows cluster environment, do the following:

- Use the virtual name of the EMM server when you configure NetBackup.
- Add the NetBackup Enterprise Media Manager service to the `ClusteredServices` entry in the following registry key:  

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Cluster\Instance1
```

This service must be included in the `ClusteredServices` entry so that it starts when a failover occurs.
- Add the NetBackup Enterprise Media Manager service to the `MonitoredServices` entry in the following registry key:  

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Cluster\Instance1
```

This service must be included in the `MonitoredServices` entry so that it is monitored. If the services fails, it is restarted. If it fails too many times, the NetBackup cluster group fails over to another node.
- Set the services to **Manual**.  
Windows then does not start the NetBackup services on the inactive node if the inactive node is rebooted.
- Update any paths to any shared drives to which the EMM server points.
- Change the server name to a virtual name and update any databases to reflect the name change.
- The database also needs to be moved (if it is with the EMM server).

## About moving the EMM server from a Windows cluster

If you move the EMM server to a different host in a Windows cluster environment, use the following process:

- Use the virtual name of the EMM server when you configure NetBackup
- Remove the NetBackup Enterprise Media Manager service from the `ClusteredServices` entry in the following registry key:  

```
(HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Cluster\Instance1)
```

Remove this service from the `ClusteredServices` entry so that it does not start when a failover occurs.

**Moving the NetBackup database from one host to another**

- Remove the NetBackup Enterprise Media Manager service from the `MonitoredServices` entry in the following registry key:  
(`HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Cluster\Instance1`)  
Remove this service from the `MonitoredServices` entry so that it does not get monitored.
- Set the services to **Manual** or remove them.  
Windows then does not start the NetBackup services on the inactive node if the inactive node is rebooted.
- Update or remove any paths to the shared drive that the EMM server points to.
- Change the server name to a non-virtual name and update any databases to reflect the name change.
- The database also needs to be moved (if it is with the EMM server).

# Using the Catalog utility

This chapter includes the following topics:

- [About the Catalog utility](#)
- [Searching for backup images](#)
- [Verifying backup images](#)
- [Viewing job results](#)
- [Promoting a copy to a primary copy](#)
- [Duplicating backup images](#)
- [Expiring backup images](#)
- [Importing backups](#)

## About the Catalog utility

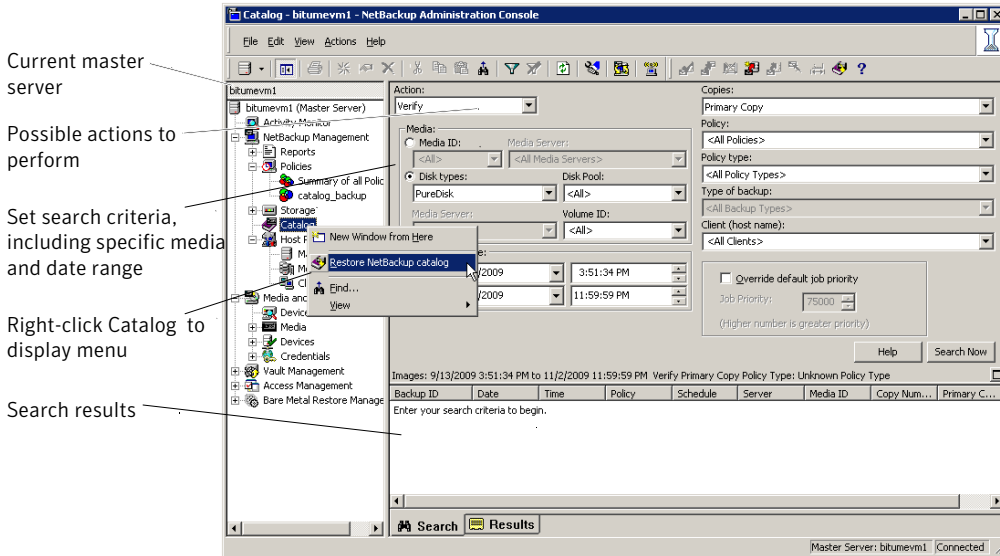
Use the **Catalog** utility to create and configure catalog backups. Catalog backups are required for NetBackup to protect NetBackup internal databases. The catalogs contain setup information as well as critical information about client backups. The catalog backups are tracked separately from other backups to ensure recovery in case of a server crash.

The **Catalog** utility is also used to perform the following actions:

- Search for backup images to verify the contents of media with what is recorded in the NetBackup catalog.
- Duplicate a backup image.
- Promote a backup image from a copy to the primary backup copy.
- Expire backup images.

- Import expired backup images or images from another NetBackup server.

Figure 19-1 Catalog utility options



- Current master server
- Possible actions to perform
- Set search criteria, including specific media and date range
- Right-click Catalog to display menu
- Search results

## Searching for backup images

Use the **Catalog** utility to search for a backup image to perform the following actions:

- Verify the backup contents with what is recorded in the NetBackup catalog.
- Duplicate the backup image to create up to 10 copies.
- Promote a copy of a backup to be the primary backup copy.
- Expire backup images.
- Import expired backup images or images from another NetBackup server.

NetBackup uses the specific search criteria to build a list of backups from which you can make your selections.

When you search for specific kinds of images, note the following:

- Verification image  
Backups that have fragments on another volume are included, as they exist in part on the specified volume.
- Import image

If a backup begins on a media ID that was not processed by the initiating backup procedure, the backup is not imported.

If a backup ends on a media ID that was not processed by the initiating backup procedure, the backup is incomplete.

See “[Importing backups](#)” on page 666.

[Table 19-1](#) lists the search criteria for backup images:

**Table 19-1** Catalog utility search criteria

|                  |                                                                                                                                                                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Action           | Select the action that was used to create the image: <b>Verify, Duplicate, Import</b> .<br>See “ <a href="#">Verifying backup images</a> ” on page 656.<br>See “ <a href="#">Duplicating backup images</a> ” on page 659.<br>See “ <a href="#">Expiring backup images</a> ” on page 665. |
| Media ID         | The media ID for the volume. Type a media ID in the box or select one from the scroll-down list. To search on all media, select <b>&lt;All&gt;</b> .                                                                                                                                     |
| Media Server     | The name of the media server that produced the originals. Type a media server name in the box or select one from the scroll-down list. To search through all media servers, select <b>All Media Servers</b> .                                                                            |
| Disk type        | The type of the disk storage unit on which to search for backup images.                                                                                                                                                                                                                  |
| Disk pool        | The name of the disk pool on which to search for backup images.                                                                                                                                                                                                                          |
| Volume ID        | The ID of the disk volume in the disk pool on which to search for backup images.                                                                                                                                                                                                         |
| NearStore Server | The name of the NearStore server to search for images. Type a server name in the box or select one from the scroll-down list. To search through all NearStore servers, select <b>All NearStore Servers</b> .                                                                             |
| Path             | To search for an image on a disk storage unit, enter the path to search. Or, select <b>All</b> to search all of the disk storage on the specified server. Appears if the disk type is BasicDisk or NearStore.                                                                            |
| Date/time range  | The range of dates and times that includes all the backups for which you want to search. The Global Attributes property <b>Policy Update Interval</b> determines the default range.                                                                                                      |
| Copies           | The source you want to search. From the scroll-down list, select either Primary or the copy number.                                                                                                                                                                                      |

**Table 19-1** Catalog utility search criteria (*continued*)

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy                        | The policy under which the selected backups were performed. Type a policy name in the box or select one from the scroll-down list. To search through all policies, select <b>All Policies</b> .                                                                                                                                                                                                                                                                                                                                                    |
| Client (host name)            | The host name of the client that produced the originals. Type a client name in the box or select one from the scroll-down list. To search through all hosts, select <b>All Clients</b> .                                                                                                                                                                                                                                                                                                                                                           |
| Type of backup                | The type of schedule that created the backup. Type a schedule type in the box or select one from the scroll-down list. To search through all schedule types, select <b>All Backup Types</b> .                                                                                                                                                                                                                                                                                                                                                      |
| Override default job priority | Select the job priority for verify, duplicate, and import actions.<br><br>To change the default for the selected action, enable <b>Override default job priority</b> . Then, select a value in the <b>Job Priority</b> field.<br><br>Changes in the catalog dialog box affect the priority for the selected job only.<br><br>If this option is not enabled, the job runs using the default priority as specified in the <b>Default Job Priorities</b> host properties.<br><br>See <a href="#">“Default Job Priorities properties”</a> on page 105. |

## Verifying backup images

NetBackup can verify the contents of a backup by reading the volume and comparing its contents to what is recorded in the NetBackup catalog.

This operation does not compare the data on the volume to the contents of the client disk. However, the operation does read each block in the image to verify that the volume is readable. (However, data corruption within a block is possible.) NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

### To verify backup images

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Set up the search criteria for the image you want to verify. Click **Search Now**.
- 3 Click the **Results** tab, then select the verification job to view the job results.  
See [“Viewing job results”](#) on page 657.



## Viewing job results

The results of verify, duplicate, or import jobs appear in the **Results** tab. The top portion of the dialog box displays all existing log files.

To view a log file, select the name of the log from the list. The current log file appears in the bottom portion of the **Results** dialog box. If an operation is in progress, the log file results refresh as the operation proceeds.

### To view job results

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Click the **Results** tab.
- 3 Select a log file.
- 4 Select **View > Full View** to display the entire log file in a screen editor.

Select **Edit > Delete** to delete the log.


You can also right-click the log file and select an action from the scroll-down menu.

## Promoting a copy to a primary copy

Each backup is assigned a primary copy. NetBackup uses the primary copy to satisfy restore requests. The first backup image that is created successfully by a NetBackup policy is the primary backup. If the primary copy is unavailable and a duplicate copy exists, select a copy of the backup and set it to be the primary copy.

NetBackup restores from the primary backup, and Vault duplicates from the primary backup. If your Vault profile performs duplication, you can designate one of the duplicates as the primary. In most circumstances, the copy remaining in the robot is the primary backup. When a primary backup expires, the next backup (if it exists) is promoted to primary automatically.

**Figure 19-2** Primary copy status



Primary Copy status indicates that the image is now the primary copy

**To promote a backup copy to a primary copy**

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Set up the search criteria for the image you want to promote to a primary copy. Be sure that you indicate a copy in the **Copies** field and not **Primary Copy**. Click **Search Now**.

See “[Searching for backup images](#)” on page 654.

- 3 Select the image you want to promote.
- 4 Click **Actions > Set Primary Copy**.

After the image is promoted to the primary copy, the Primary Status column immediately reads **Yes**.

**To promote a copy to a primary copy for many backups**

You can also promote a copy to be a primary copy for many backups using the `bpchangeprimary` command. For example, the following command promotes all copies on the media that belongs to the **SUN** volume pool. The copies must have been created after August 8, 2009:

```
bpchangeprimary -pool SUN -sd 08/01/2009
```

The following command promotes copy 2 of all backups of `client_a`. The copies must have been created after January 1, 2009:

```
bpchangeprimary -copy 2 -cl client_a -sd 01/01/2009
```

More information is available in *NetBackup Commands*.

### To use `bpduplicate` to promote a backup copy to a primary copy

- 1 Enter the following command:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpduplicate
-npc pcopy -backupid bid
```

Where:

*Install\_path* is the directory where NetBackup is installed.

*pcopy* is the copy number of the new primary copy.

*bid* is the backup identifier as shown in the Images on Media report.

Find the volume that contains the duplicate backup by using the Images on Media report.

- 2 Specify the backup ID that is known (and also the client name if possible to reduce the search time).

The `bpduplicate` command writes all output to the NetBackup logs. Nothing appears in the command window.

After the duplicate copy is promoted to the primary copy, use the client interface on the client to restore files from the backup.

For instructions, see the online Help in the Backup, Archive, and Restore client interface.

## Duplicating backup images

NetBackup does not verify in advance whether the storage units and the drives that are required for the duplicate operation are available for use. NetBackup verifies that the destination storage units exist. The storage units must be connected to the same media server.

[Table 19-2](#) lists the scenarios in which duplication is possible and scenarios in which duplication is not possible:

**Table 19-2** Backup duplication scenarios

| Duplication possible                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Duplication not possible                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>■ From one storage unit to another.</li> <li>■ From one media density to another.</li> <li>■ From one server to another.</li> <li>■ From multiplex to nonmultiplex format.</li> <li>■ From multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. The duplicate is created with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)</li> </ul> | <ul style="list-style-type: none"> <li>■ While the backup is created (unless making multiple copies concurrently).</li> <li>■ When the backup has expired.</li> <li>■ By using NetBackup to schedule duplications automatically (unless you use a Vault policy to schedule duplication)</li> <li>■ When it is a multiplexed duplicate of the following type:               <ul style="list-style-type: none"> <li>■ FlashBackup</li> <li>■ NDMP backup</li> <li>■ Backups from disk type storage units</li> <li>■ Backups to disk type storage units</li> <li>■ Nonmultiplexed backups</li> </ul> </li> </ul> |

An alternative to taking time to duplicate backups is to create up to four copies simultaneously at backup time. (This option is sometimes referred to as Inline Copy.) Another alternative is to use storage lifecycle policies.

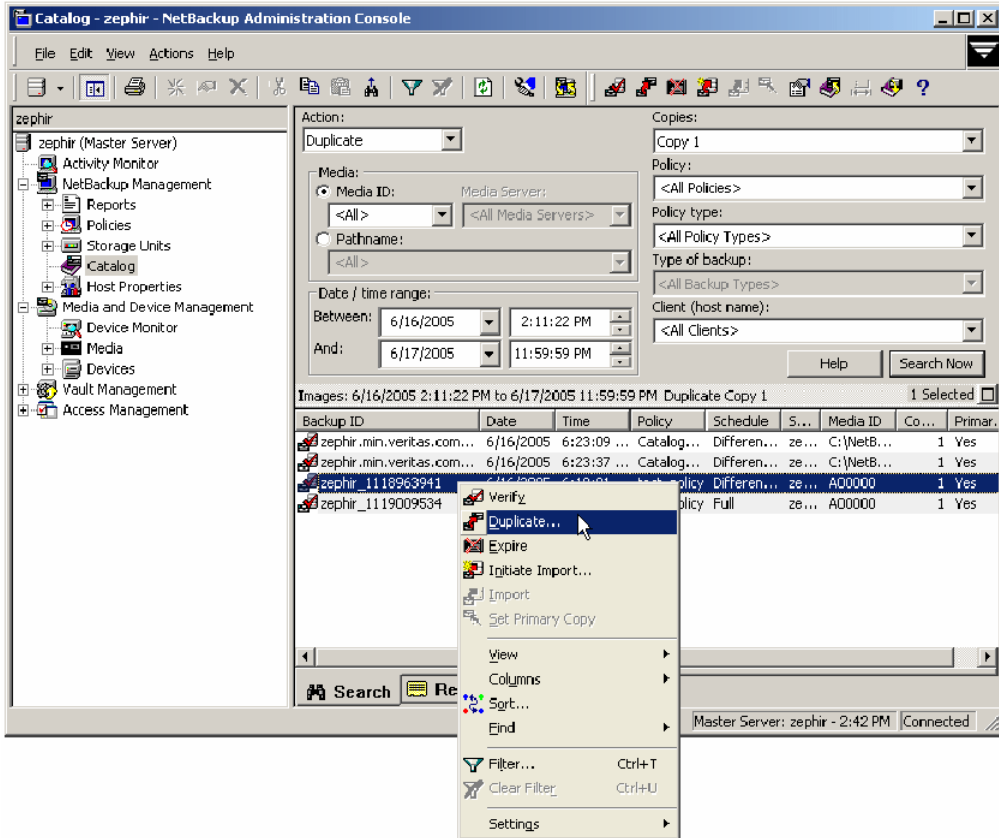
See [“Use only one method to create multiple copies”](#) on page 433.

**To duplicate backup images**

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Set up the search criteria for the image you want to duplicate. Click **Search Now**.

- 3 Right-click the image(s) you want to duplicate and select **Duplicate** from the shortcut menu.


If you duplicate an online, hot catalog backup, select all child jobs that were used to create the catalog backup. All jobs must be duplicated to duplicate the catalog backup.



**4** Specify the number of copies you want to create.

NetBackup can create up to 10 copies of unexpired backups. Indicate the number of backup copies in **Host Properties > Master Servers > Global Attributes > Maximum backup copies**.

See “[Maximum backup copies](#)” on page 135.



If enough drives are available, the copies are created simultaneously. Otherwise, the system may require operator intervention if four copies are to be created using only two drives, for example.

**5** The primary copy is the copy from which restores are done. Normally, the original backup is the primary copy.

If you want one of the duplicated copies to become the primary copy, check the appropriate check box, otherwise leave the fields blank.

When the primary expires, a different copy automatically becomes primary. (The copy that is chosen is the one with the smallest copy number. If the primary is copy 1, copy 2 becomes primary when it expires. If the primary is copy 5, copy 1 becomes primary when it expires.)

**6** Specify the storage unit where each copy is stored. If a storage unit has multiple drives, it can be used for both the source and destination.

All storage units must meet the criteria for creating multiple copies.

See “[Criteria for creating multiple copies](#)” on page 503.

**7** Specify the volume pool where each copy is stored.

The volume pool selections are based on the policy type setting that was used for the query:

- If the policy type was set to query for All Policy Types (default), all volume pools are included in the drop-down list. Both catalog and non-catalog volume pools are included.

- If the policy type was set to query for NBU-Catalog, only catalog volume pools are included in the drop-down list.
- If the policy type was set to query for a policy type other than **NBU-Catalog** or **All Policy Types**, only non-catalog volume pools are included in the drop-down list.

NetBackup does not verify that the media ID selected for the duplicate copy is different from the media ID that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure that a different volume is used.

**8** Select the retention level for the copy, or select No change.

The duplicate copy shares many attributes of the primary copy, including backup ID. Other attributes apply only to the primary. (For example, elapsed time.) NetBackup uses the primary copy to satisfy restore requests.

Consider the following items when selecting the retention level:

- If **No Change** is selected for the retention period, the expiration date is the same for the duplicate and the source copies. You can use the `bpeupdate` command to change the expiration date of the duplicate.
- If a retention period is indicated, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2009 and its retention period is one week, the new copy's expiration date is November 21, 2009.

**9** Specify whether the remaining copies should continue or fail if the specified copy fails.

**10** Specify who should own the media onto which you are duplicating images.

Select one of the following:

|                |                                                                                                                                                                                                                                            |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Any            | NetBackup chooses the media owner, either a media server or server group.                                                                                                                                                                  |
| None           | The media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.                                                                                           |
| A server group | Only those media servers in the group are allowed to write to the media on which backup images for this policy are written. All of the media server groups that are configured in your NetBackup environment appear in the drop-down list. |

- 11 If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, check **Preserve Multiplexing**. If you do not duplicate all the backups in a multiplexed group, the duplicate contains a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

By default, duplication is done serially and attempts to minimize media mounts and positioning time. Only one backup is processed at a time. If **Preserved Multiplexing** is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated before the multiplexed backups are duplicated.

The **Preserve Multiplexing** setting does not apply when the destination is a disk storage unit. However, if the source is a tape and the destination is a disk storage unit, selecting **Preserve Multiplexing** ensures that the tape is read in one pass.

- 12 Click **OK** to start duplicating.
- 13 Click the **Results** tab, then select the duplication job to view the job results.  
See “[Viewing job results](#)” on page 657.

## About multiplexed duplication

Consider the following items regarding multiplexed duplication:

- When multiplexed backups are duplicated, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups have a multiplexing factor that is no greater than the factor that was used during the original backup.
- When backups in a multiplexed group are duplicated to a storage unit, the duplicated group is identical as well. However, the storage unit must have the same characteristics as the unit where the backup was originally performed. The following items are exceptions:
  - If EOM (end of media) is encountered on either the source or the destination media.
  - If any of the fragments are zero length in the source backups, the fragments are removed during duplication. A fragment of zero length occurs if many multiplexed backups start at the same time.

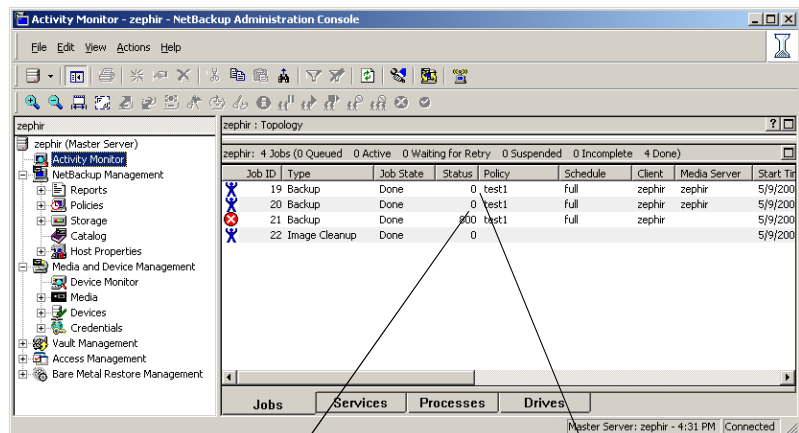


## Jobs that appear while making multiple copies

When multiple copies are made concurrently, a parent job appears, plus a job for each copy.

The parent job displays the overall status, whereas the copy jobs display the status of a single copy. Viewing the status of individual jobs allows you to troubleshoot jobs individually. For example, if one copy fails but the other copy is successful, or if each copy fails for different reasons. If at least one copy is successful, the status of the parent job is successful. Use the Parent Job ID filter to display the parent Job ID. Use the Copy filter to display the copy number for a particular copy.

The following example shows a backup that contains two copies. The parent job is 19, copy 1 is job 20, and copy 2 is job 21. Copy 1 finished successfully, but copy 2 failed with a 800 status (disk volume cannot be used for more than one copy in the same job). Since at least one copy successfully completed, the parent job displays a successful (0) status.



Copy 1 was successful, but  
Copy 2 failed

The parent job was successful because  
at least one copy was successful

## Expiring backup images

To expire a backup image means to force the retention period to expire. When the retention period expires, NetBackup deletes information about the backup. The files in the backups are unavailable for restores without first re-importing.

### To expire a backup image

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Set up the search criteria for the image you want to expire, then click **Search Now**.  
See [“Searching for backup images”](#) on page 654.
- 3 Select the image you want to expire and select **Actions > Expire**.
- 4 A message appears that announces that once the backups are expired, they cannot be used for restores. Select **Yes** to begin to expire the images.

## Importing backups

NetBackup can import the backups that have expired, the backups from another NetBackup server, or the backups written by Backup Exec for Windows.

See [“Importing Backup Exec media”](#) on page 671.

During an import operation, NetBackup recreates NetBackup catalog entries for the backups on the imported volume. The import capability is useful for moving volumes from one site to another and for recreating NetBackup catalog entries.

NetBackup supports the capability to import and restore the following Backup Exec backup types:

- Windows
- UNIX
- Exchange
- SQL
- NetWare

An image is imported in the following two phases:

- |          |                                                                                                                                                                                                        |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Phase I  | NetBackup creates a list of expired catalog entries for the backups on the imported volume. No actual import occurs in Phase I.<br>See <a href="#">“Importing backup images, Phase I”</a> on page 667. |
| Phase II | Images are selected for importing from the list of expired images that was created in Phase I.<br>See <a href="#">“Importing backup images, Phase II”</a> on page 668.                                 |

## Importing backup images, Phase I

Phase I of the import process creates a list of expired images from which to select to import in Phase II. No import occurs in Phase I.

Initiate an import by using either the Import Images Wizard or initiate it manually.

If tape is used, each tape must be mounted and read. It may take some time to read the catalog and build the list of images.

To import an online, hot catalog backup, import all of the child jobs that were used to create the catalog backup.

### To import backup images by using the Import Images Wizard, Phase I

**1** If you import Backup Exec media, run the `vmphyinv` physical inventory utility to update the Backup Exec media GUID in the NetBackup Media Manager database. Run the command only once after creating the media IDs in the NetBackup Media Manager database.

See [“About the physical inventory utility”](#) on page 343.

**2** Add the media IDs that contain the Media Manager backups to the server where the backups are to be imported.

**3** Select **Import Images** in the right pane to launch the wizard. **Import Images** is available when **Master Server** or **NetBackup Management** is selected.

**4** The wizard explains the 2-step import process and takes you through Phase I. Click **Next**.

**5** Type the name of the host that contains the volume to import. Click **Next**.

This media server becomes the media owner.

**6** Select whether the images to import are on tape or disk.

**7** Depending on whether the import is from tape or disk do one of the following:

- Type the Media ID for the volume that contains the backups to import.

- Enter the path from which the images are to be imported.

Click **Next**.

If the Backup Exec media is password-protected, the job fails without a correct password. The logs indicate that either no password or an incorrect password was provided. If the media is not password-protected and the user provides a password, the password is ignored.

To import Backup Exec media if the password contains non-ASCII characters do the following:

- Use the NetBackup Administration Console on Windows. (You cannot use the NetBackup-Java Administration Console.)

- Use the `bpimport` command.
- 8 Click **Finish**. The wizard explains how to check the progress as the media host reads the media.  
See “[Viewing job results](#)” on page 657.
- 9 Complete the import.  
See “[Importing backup images, Phase II](#)” on page 668.

## Importing backup images, Phase II

To import the backups that consist of fragments on multiple tapes, first run the Initiate Import (Import Phase I). The first phase reads the catalog to determine all the tapes that contain fragments. After Phase I, start the Import (Phase II). If Phase II is run before Phase I, the import fails with a message. For example, Unexpected EOF or Import of backup ID failed, fragments are not consecutive.

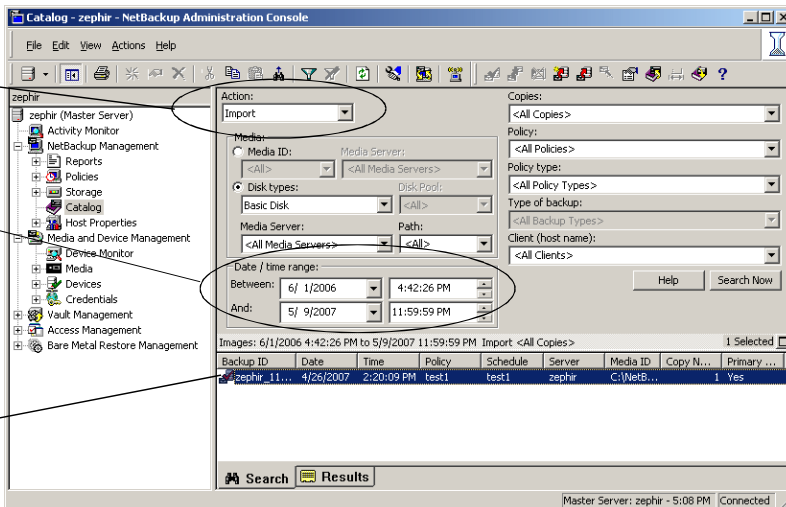
### To import backup images, Phase II

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 2 Set up the search criteria to find images available to import by setting the search action to **Import**. Be sure to select a date range that includes the images you want to import.

Select Import to search for imported images

Select the date range that includes the images to import

Images eligible for importing appear as a result



- 3 Select the image(s) you want to import and select **Actions > Import**.
- 4 To view the log, click the **Results** tab, then select the import job log.

## Importing expired images

The expiration date for the imported items is the current date plus the retention period. For example, if a backup is imported on November 14, 2009, and its retention period is one week, the new expiration date is November 21, 2009.

Consider the following items when importing backup images:

- NetBackup can import the disk images that NetBackup version 6.0 (or later) writes.
  - You cannot import a backup if an unexpired copy of it already exists on the server.
  - NetBackup does not direct backups to imported volumes.
  - If you import an online, hot catalog backup, import all the child jobs that were used to create the catalog backup. All jobs must be imported to import the catalog backup.
  - To import a volume with the same media ID as an existing volume on a server, use the following example where you want to import a volume with media ID A00001. (A volume with media ID A00001 already exists on the server.)
    - Duplicate the existing volume on the server to another media ID (for example, B00001).
    - Remove information about media ID A00001 from the NetBackup catalog by running the following command:

```
install_path \VERITAS\NetBackup\bin\admincmd\bpxpdate
-d 0 -m mediaID
```
    - Delete media ID A00001 from Media Manager on the server.
    - Add the other A00001 to Media Manager on the server.
- To avoid this problem in the future, use unique prefix characters for media IDs on all servers.

See [“Expiring backup images”](#) on page 665.

## Initiating an import without the Import Wizard

Use the following procedure to initiate an import without the Import Wizard.

### To initiate an import without the Import Wizard

- 1 To import Backup Exec media, run the `vmphyinv` physical inventory utility to update the Backup Exec media GUID in the NetBackup Media Manager database. Run the command only once after creating the media IDs in the NetBackup Media Manager database.

See “[About the physical inventory utility](#)” on page 343.

- 2 To import the images from tape, make the media accessible to the media server so the images can be imported.
- 3 In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
- 4 Select **Actions > Initiate Import**.
- 5 Enable the **Use Import Images Wizard** option to have the Import Wizard guide you through Phase I of the import process.
- 6 In the **Media Server** field, specify the name of the host that contains the volume to import. This media server becomes the media owner.
- 7 Indicate the location of the image. Under **Image type**, select whether the images to be imported are located on tape or on disk.
  - If images are on tape:
    - In the **Media ID** field, type the Media ID of the volume that contains the backups to import.
    - Check whether or not the images to import are password-protected Backup Exec images.
    - Validate the Backup Exec password by retyping the password in the field provided.
  - If images are on disk
    - In the **Disk type** field, select the type of the disk storage unit on which to search for backup images. The disk types depend on which NetBackup options are licensed.
    - If the disk type references a disk pool, enter or select the disk pool and the disk volume ID.

For a **BasicDisk** type, enter or browse to the path to the images in the field provided.

For a **NearStore** disk type, select or enter the name of the NearStore server and the NearStore volume.

Click **OK**.

- 8 Click **OK** to begin reading the catalog information from the source volume.
- 9 Click on the **Catalog Results** tab to watch as NetBackup looks at each image on the tape. NetBackup determines whether or not each image has expired and can be imported. The job also displays in Activity Monitor as an Import type. Select the import job log to view the job results.

## Importing Backup Exec media

Consider the following items when importing Backup Exec media:

- If the Backup Exec media is password-protected, the import job fails without a correct password. The logs indicate that either no password or an incorrect password, was provided. If the media is not password-protected and the user provides a password, the password is ignored.
- If the Backup Exec media uses a password that contains non-ASCII characters, use the NetBackup Administration Console on Windows. (The NetBackup-Java Administration Console cannot be used.) Or, use the `bpimport` command.
- Importing from Backup Exec media does not convert or migrate Backup Exec job history, job schedules, or job descriptions to NetBackup.
- Importing from Backup Exec media does not convert Backup Exec application setup or configuration information to NetBackup.
- Any Backup Exec backups that were created with the Intelligent Image Option cannot be restored.
- If Backup Exec hard link backups are redirected and restored to partitions or drives other than the source partition or drive, the hard links are not restored. The progress log may indicate that the hard links are restored successfully, but that is not the case.

### About the host properties for Backup Exec

The Backup Exec UNIX agent identifies itself to the Backup Exec server by using a GRFS-advertised name. The advertised name may not be the same as the real machine name and path.

NetBackup must know the advertised name, along with the actual client name and path to create accurate `.#` file paths. Set the **GRFS Advertised Name**, **Actual Client**, and **Actual Path** properties in the Backup Exec Tape Reader host properties. If no entries are indicated, NetBackup assumes that the advertised name is the real machine name and the advertised path is the real path.

See “[Backup Exec Tape Reader properties](#)” on page 69.

## Backup Exec Tape Reader limitations

The following are Backup Exec Tape Reader limitations:

- Support is limited to images residing on tape media supported by the NetBackup media server.
- Importing from disk backups is not supported.
- Importing encrypted images is not supported.
- Duplication after import is not supported.
- UNIX data cannot be restored to Windows systems, Windows data to UNIX systems, Windows data to NetWare systems, or UNIX data to NetWare systems.
- NetBackup does not read the Backup Exec media that Backup Exec for NetWare writes.

## Backup Exec Tape Reader support for Windows images

The Backup Exec Tape Reader provides support for all Windows versions that NetBackup currently supports.

The support includes the following:

- Importing Windows 2003 and 2008 images.
- Recovering files from full, incremental, and differential backups.
- Importing Windows 2003 and 2008 images from Backup Exec 7 through 12.
- Recovery of System State and Shadow Copy Components.
- Importing compressed images.

## Backup Exec Tape Reader support for Exchange Server images

The Backup Exec Tape Reader provides support for the following:

- Database recovery from full, incremental, and differential backups.
- Importing Exchange 2000 and 2003 images from Backup Exec 9.1 through 12.
- Importing Exchange 2007 images from Backup Exec 11 through 12.

The support for Backup Exec images of Exchange 2003 and 2007 is limited to recovering the backup image to the same storage group. This is supported for both VSS backups as well as non-VSS backups.

The following functionality is not available for Backup Exec images of Exchange 2003 and 2007:



- Restoring individual mailbox objects or public folder objects either to the same path or different path.
- Restoring to a different storage group or Recovery Storage Group for either VSS backups or Non-VSS backups.

## Backup Exec Tape Reader support for SQL images

The Backup Exec Tape Reader provides support for the following:

- Importing SQL Server 2005 images from Backup Exec 9.1 through 12.
- Database recovery from full, incremental, differential and transaction log backups.

## Differences between importing, browsing, and restoring Backup Exec and NetBackup images

The following topics describe differences between Backup Exec and NetBackup when importing, browsing, and restoring images:

- Run `vmphyinv` for Backup Exec media  
To import Backup Exec media requires `vmphyinv` to update the Backup Exec media GUID in the NetBackup Media Manager database. Create the media IDs in the NetBackup Media Manager database, run the command, then perform Phase I and Phase II import operations.  
See [“About the physical inventory utility”](#) on page 343.
- To import and restore QIC media  
Backup Exec Quarter Inch Cartridge (QIC) media that was written in tape block sizes more than 512 bytes must be imported and restored using a NetBackup Windows media server. A NetBackup UNIX media server cannot import and restore the media in this case.
- Spanned media: Importing differences  
To import a Backup Exec backup that spans multiple media, run a Phase I import on the first media of the spanned backup set. Then, run a Phase I import on the remaining media of the spanned backup set in any order.  
The Backup Exec import process differs from the NetBackup import process. In that NetBackup import process, Phase I can be run in any order in case the image spans multiple media.
- SQL: Browsing and restoring differences  
Backup Exec SQL images are browsed, then restored using the NetBackup Backup, Archive, and Restore client interface.

NetBackup SQL images are browsed, then restored using the NetBackup SQL interface.

- **File level objects: Browsing and restoring differences**

When a user selects a Backup Exec file to restore, the directory where that file is located is restored.

When a user selects a NetBackup file to restore, only the single file is restored.

- **NetWare: Restoring differences**

NetBackup does not support restoring Backup Exec NetWare non-SMS backups that were created using the NetWare redirector.

**Storage Management Services (SMS)** software allows data to be stored and retrieved on NetWare servers independent of the file system the data is maintained in.

- **Restoring NTFS hard links, NTFS SIS files, and Exchange SIS mail messages**

- When Backup Exec NTFS images are restored, any directory named SIS Common Store is restored. The directory named SIS Common Store is restored whether or not it is the actual NTFS single instance storage common store directory. The directory is restored even if the file was not specifically selected for restore.

- Under some circumstances, additional objects are sent to the client, even though the objects were not selected for restore. The items are sent to the client when objects are restored from any backups that contain NTFS hard links, NTFS SIS files, or Exchange SIS mail messages. These additional objects are skipped by the client and are not restored. The job is considered partially successful because some objects (though not selected by the user), are skipped.

- When NTFS hard links or SIS files, or Exchange SIS mailboxes are redirected for restore, all or some of the files should be redirected to any location on the source drive. Or, you also can redirect all files to a single location on a different drive.

For example, if the following hard link or SIS files are backed up:

```
C:\hard_links\one.txt
C:\hard_links\two.txt
C:\hard_links\three.txt
```

Upon restore, either the files can be redirected to any location on C:\, or all the files must be redirected to a different drive.

The following combination would be unsuccessful:

```
C:\hard_links\one.txt to a location on C:\
C:\hard_links\two.txt to a location on D:\
```

If all the files are to be redirected to a different drive, specify that `C:\` be replaced with `D:\` in the redirection paths.

**Unsuccessful:**

The redirection paths specify that `C:\hard_links` be replaced with `D:\hard_links`.

**Successful:**

The redirection paths specify that `C:\hard_links` be replaced with `C:\redir_hard_links`.



# Monitoring and reporting

- [Chapter 20. Monitoring NetBackup activity](#)
- [Chapter 21. Reporting in NetBackup](#)



# Monitoring NetBackup activity

This chapter includes the following topics:

- [Using the Activity Monitor](#)
- [Activity Monitor topology](#)
- [About the Jobs tab](#)
- [About the Services tab](#)
- [About the Processes tab](#)
- [About the Drives tab](#)
- [About the jobs database](#)
- [About the Device Monitor](#)
- [About media mount errors](#)
- [About pending requests and actions](#)
- [Managing pending requests and actions](#)

## Using the Activity Monitor

Use the Activity Monitor in the NetBackup Administration Console to monitor and control NetBackup jobs, services, processes, and drives.

The options on the Activity Monitor menu bar are described in the online Help.

---

**Note:** The **Filter** option on the **View** menu is useful for displaying in Activity Monitor only those jobs with specified characteristics. For example, the jobs that were started before a specific date; jobs in the queued state; jobs with status completion codes within a specified range.

---


The status bar appears at the top of the Activity Monitor list and displays the following information, depending on which tab is currently selected:

- The master server on which the jobs reside.
- The total number of jobs.
- The number of jobs in each of the job states: Active, Queued, Waiting for Retry, Suspended, Incomplete, and Done.
- The number of jobs currently selected.
- The number of NetBackup services that run.
- The number of drives and the state of each (Active, Down).

The numbers always reflect the actual number of jobs, even when the filter is used.



**Figure 20-1** Activity Monitor



## Activity Monitor topology

The Activity Monitor topology view displays the state and configuration of the entire NetBackup system being administered. The Activity Monitor displays only robots and the drives that have storage units configured. If a device host has no configured devices, the device host does not appear in the Activity Monitor.

The topology view shows master servers, media servers, clients, and NetBackup storage unit devices. The view displays backup and restore activity, job failures, the services that are down, and drive state status.


The topology view displays the physical and the logical connections between devices.

Lines appear between a drive in use and the server that uses it. For example, lines appear between a media server and a device that runs a job for the server.

Robots and drives are represented as connected to a media server. Drives that are physically located in a robot appear directly beneath the robot. Stand-alone drives are represented as individual drive objects.

Figure 20-2 shows some of the icons you may see in the Activity Monitor.

Figure 20-2 Example of Activity Monitor icons



## Filtering topology objects

To select an object in the topology pane is one method to filter the contents of the Activity Monitor list. To select multiple objects of the same type, press the **Ctrl** key and select another object. You cannot select the topology objects that are not alike.

Select an object to highlight the connecting lines from the object to all other objects to which it is connected. For example, click a server to highlight all attached robots, media, and drives configured to the server.

## About the Jobs tab

The **Jobs** tab displays all jobs that are in process or that have completed for the master server currently selected.

---

**Note:** Job selection preference is given to jobs from NetBackup 6.0 media servers over media servers of previous versions.

---

For some backup jobs, a parent job is used to perform pre- and post-processing. Parent jobs display a dash (-) in the Schedule column. A parent job runs the start

and end notify scripts (`PARENT_START_NOTIFY`, `PARENT_END_NOTIFY`) from the master server:

```
Install_path\VERITAS\NetBackup\bin
```

The role of the parent job is to initiate requested tasks in the form of children jobs.

The tasks vary, depending on the backup environment, as follows:

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Snapshot Client             | <p>The parent job creates the snapshot, initiates children jobs, and deletes the snapshot when complete.</p> <p>Children jobs are created if the Snapshot Client settings are configured to retain snapshots for Instant Recovery, then copy snapshots to a storage unit. (<b>Snapshots and copy snapshots to a storage unit</b> is selected in the policy <b>Schedule Attributes</b> tab.)</p> <p>Children jobs are not created if the Snapshot Client settings are configured to retain snapshots for Instant Recovery, but to create snapshots only. That is, the snapshot is not backed up to a storage unit, so no children jobs are generated. (<b>Snapshots only</b> is selected in the policy <b>Schedule Attributes</b> tab.)</p> |
| Bare Metal Restore          | <p>The parent job runs <code>brmsavecfg</code>, then initiates the backup as a child job. If multistreaming and BMR are used together, the parent job can start multiple children jobs.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Online, hot catalog backups | <p>The parent job for catalog backups works with <code>bpdbm</code> to initiate multiple children backup jobs:</p> <ul style="list-style-type: none"><li>■ A Sybase backup</li><li>■ A file system backup of the master server</li><li>■ A backup of the BMR database, if necessary</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Multiple copies             | <p>A multiple copies job produces one parent job and multiple child jobs. Child jobs that are part of a multiple copies parent job cannot be restarted individually. Only the parent job (and subsequently all the children jobs) can be restarted.</p> <p>See “<a href="#">Multiple copies attribute</a>” on page 503.</p>                                                                                                                                                                                                                                                                                                                                                                                                                |
| Multiple data streams       | <p>The parent job performs stream discovery and initiates children jobs. A parent job does not display a schedule in the Activity Monitor. Instead, a dash (-) appears for the schedule because the parent schedule is not used and the children schedules may be different. The children jobs display the ID of the parent job in the Activity Monitor.</p>                                                                                                                                                                                                                                                                                                                                                                               |

|            |                                                                                                                                                                                                                                                        |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SharePoint | The parent job runs a resolver process during which children jobs are started. This process is similar to the stream discovery for multiple data streams. If multiple data streams are enabled, some children jobs can be split into multiple streams. |
| Vault      | The parent job starts the Vault profile. Then, the Vault profile starts the duplicates as jobs. The duplicates do not appear as children jobs in the Activity Monitor.                                                                                 |

## Viewing job details

The following procedure describes how to view job details.

### To view job details

- ◆ To view the details for a specific job, double-click on the job in the **Jobs** tab. The **Job Details** dialog box appears that contains detailed job information on two tabs: a **Job Overview** tab and a **Detailed Status** tab.

Not all columns appear by default. Click **View > Columns > Layout** to show or hide columns.

## Setting job detail selections

The following procedure describes how to customize the jobs detail selections.

### To show or hide column heads

- 1 Open the Activity Monitor.
- 2 Click **View > Columns > Layout**. The **Set Column Layout** dialog box appears.
- 3 Select the heading you want to display or hide.
  - Select the **Show Column** button to display the heading.
  - Select the **Hide Column** button if you do not want to see the column head.
- 4 To change the order in which the columns appear, select the column head. Then, click the **Move Up** button or the **Move Down** button to reorder the columns.
- 5 Click **OK** to apply the changes.

## Monitoring the detailed status of a selected job

The following procedure describes how to monitor the detailed status of a job.

**To monitor the detailed status of selected jobs**

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job(s) for which you want to view details.
- 3 Select **Actions > Details**.

## Running the Troubleshooter from within the Activity Monitor

If a job fails, use the Troubleshooter on the **Help** menu. The Troubleshooter helps to explain the problem and provides the corrective actions that are based on the NetBackup status code that the job returns.

**To run troubleshooter within the Activity Monitor**

- 1 Select a job in the Activity Monitor.
- 2 Open the Troubleshooter:
  - Click the **Troubleshooter** icon.
  - Select **Help > Troubleshooter**.
  - Open the job details for a job and click the **Detailed Status** tab. Then click **Troubleshooter**.

An explanation of the problem appears on the **Problems** tab and a recommended action appears on the **Troubleshoot** tab.

If no status code is entered in the Troubleshooter status code field, enter the status code of the failed job. Click **Lookup** to locate the troubleshooting information. You can open the Troubleshooter at any time and enter a status code.

## Deleting completed jobs

The following procedure describes how to delete a completed job.

**To delete completed jobs**

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job(s) you want to delete.
- 3 Select **Edit > Delete**.

## Canceling a job that has not completed

The following procedure describes how to cancel a job that has not completed.

### To cancel a job that has not completed

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job that has not completed that you want to cancel. It may be a job that is in the Queued, Re-Queued, Active, Incomplete, or Suspended state.
- 3 Select **Actions > Cancel Job**.

If the selected job is a parent job, all the children of that parent job are canceled as well.

In most cases, a canceled child job cancels only that job and allows the other child jobs to continue. One exception is multiple copies created as part of a policy or storage lifecycle policy: canceling a child job cancels the parent job and all child jobs.

- 4 To cancel all jobs in the jobs list that have not completed, click **Actions > Cancel All Jobs**.

## Restarting a job

The following procedure describes how to restart a completed job.

### To restart a completed job

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the Done job you want to restart.
- 3 Select **Actions > Restart Job**. In this case, a new job ID is created for the job. The job details for the original job references the job ID of the new job.

## Suspending restore or backup jobs

The following procedure describes how to suspend restore or backup jobs.

### To suspend a restore or a backup job

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job you want to suspend.  
Only the backup and restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Suspend Job**.

## Resuming suspended jobs

The following procedure describes how to resume suspended jobs.

**To resume a suspended or an incomplete job**

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select the suspended or the incomplete job you want to resume.  
Only the backup and restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Resume Job**.

## Printing job list information

The following procedure describes how to print job list information.

**To print job detail information from a list of jobs**

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Select a job to print. Hold down the Control or Shift key to select multiple jobs. If no job is selected, all jobs print.
- 3 Select **File > Print**.

## Printing job detail information

The following procedure describes how to print job detail information.

**To print job detail information from a single job**

- 1 Open the Activity Monitor and select the **Jobs** tab.
- 2 Double-click on a job to open it.
- 3 In the **Job Details** dialog box, click **Print**. Then select a printer and set the printer options.

```

Job State Done
Job type: Backup
Backup type:
Policy type: MS-Windows-NT
Client: silk
Master Server: zephir
Priority: 0
Owner: root
Group: root
Retention: 2 weeks
Compression: No
Job Details:444
Started: 6/7/2007 6:55:00 PM
Elapsed: 00:02:59
Ended: 6/7/2007 6:57:59 PM

Job PID: 2220
Current kilobytes written: 30187
Current files written: 106
Storage unit: zephir-dlt-robot-tld-0
Media server: zephir
Status: the requested operation was successfully completed(0)
Attempt 1
Started: 6/7/2007 6:55:10 PM
Elapsed: 00:02:49
Ended: 6/7/2007 6:57:59 PM

File list:
C:\Documents and Settings

```

## Copying Activity Monitor text to another document

The following procedure describes how to copy Activity Monitor text to a file.

### To copy Activity Monitor text to a file

- 1 Open the Activity Monitor and select a job.
- 2 Select **Edit > Copy**.
- 3 Paste the selected text into the file (for example, an Excel document).

## Changing the Job Priority dynamically

To dynamically change the priority of a job, select one or more queued or active jobs that wait for resources. Then, either from the **Actions** menu or by right-clicking the job, select **Change Job Priority**.

Select one of the following methods to change the priority:

**Set Job Priority to** Enter the specific job priority for the selected jobs.

**Increment the Job Priority by** Raise the priority of the job by the selected internal.

**Decrement the Job Priority by** Lower the priority of the job by the selected internal.

Changes in the **Change job priority** dialog box affect the priority for the selected job only, and not all other jobs of that type.

To change the job priority defaults, use the Default Job Priorities host properties.

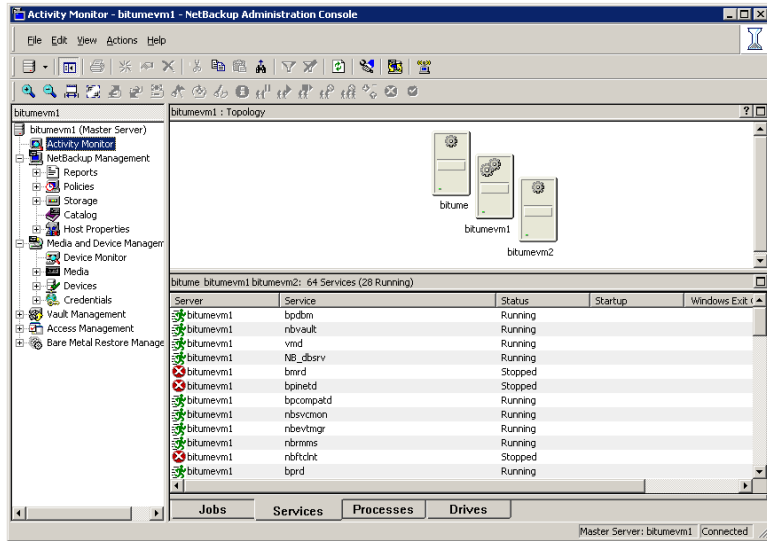
See “[Default Job Priorities properties](#)” on page 105.

## About the Services tab

The **Services** tab displays the status of NetBackup services on the master server and all media servers that the selected master server uses.



**Figure 20-3** Services tab in the Activity Monitor



**Note:** To see any services or processes on another machine, the other machine must be running on a Microsoft platform. The user must be authenticated on the Microsoft platform.

Not all columns appear by default. Click **View > Columns > Layout** to show or hide columns.

**Table 20-1** NetBackup services

| Service                                                | Description                                 |
|--------------------------------------------------------|---------------------------------------------|
| NetBackup Bare Metal Restore Master Server (bmrtd.exe) | Appears if Bare Metal Restore is installed. |

**Table 20-1** NetBackup services (*continued*)

| Service                                                                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>NetBackup Client Service<br/>(<code>bpinetd.exe</code>)</p>          | <p>Listens for connections from NetBackup servers in the network and when an authorized connection is made, starts the necessary NetBackup process to service the connection.</p> <p><b>Note:</b> The Client Service must be run as either an Administrator or Local System account. Problems arise if the Client Service logon account differs from the user that is logged on to use NetBackup. When NetBackup tries to contact the Client Service, a message appears that states the service did not start because of improper logon information. The event is recorded in the Windows System event log. The log notes that the account name is invalid, does not exist, or that the password is invalid.</p> <p>The service cannot be stopped from the Activity Monitor because it receives data that appears in the Administration Console. If it is stopped, the console cannot display the data.</p> |
| <p>NetBackup Compatibility Service<br/>(<code>bpcompatd.exe</code>)</p> | <p>Service that is used to communicate with legacy NetBackup services.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <p>NetBackup Database Manager<br/>(<code>bpdbm.exe</code>)</p>          | <p>Manages the NetBackup internal databases and catalogs. <code>BPDBM</code> must be running on the NetBackup master server during all normal NetBackup operations.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <p>NetBackup Deduplication Engine<br/>(<code>spoold.exe</code>)</p>     | <p>Service that runs on the NetBackup deduplication media server host. This service deduplicates client data. The file name <code>spoold.exe</code> is short for storage pool daemon; do not confuse it with a print spooler daemon.</p> <p>Active only if the NetBackup Deduplication Option is licensed and the media server is configured as a deduplication media server.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <p>NetBackup Deduplication Manager<br/>(<code>spad.exe</code>)</p>      | <p>Service that runs on the NetBackup deduplication media server host. This service maintains the NetBackup deduplication configuration, controls deduplication internal processes, controls replication, controls security, and controls event escalation.</p> <p>Active only if the NetBackup Deduplication Option is licensed and the media server is configured as a deduplication media server.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

**Table 20-1** NetBackup services (*continued*)

| Service                                                              | Description                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetBackup Device Manager<br>( <i>ltid.exe</i> )                      | Starts the Volume Manager ( <i>vmd</i> ), the automatic volume recognition process ( <i>avrd</i> ), and any robotic processes. Processes the requests to mount and dismount tapes in robotically controlled devices through the robotic control processes. Mounts the volumes on the tape devices in response to user requests.                                            |
| NetBackup Enterprise Media Manager ( <i>nbemm.exe</i> )              | Accesses and manages the database where media and device configuration information is stored ( <i>EMM_DATA.db</i> ). <i>nbemm.exe</i> must be running in order for jobs to run.<br><br>The service cannot be stopped from the Activity Monitor because it receives data that appears in the Administration Console. If it is stopped, the console cannot display the data. |
| NetBackup Event Manager Service ( <i>nbvmtmgr.exe</i> )              | Provides the communication infrastructure to pass information and events between distributed NetBackup components. Runs on the same system as the NetBackup Enterprise Media Manager.                                                                                                                                                                                      |
| NetBackup Job Manager<br>( <i>nbjm.exe</i> )                         | Accepts the jobs that the Policy Execution Manager ( <i>nbpem.exe</i> ) submits and acquires the necessary resources. The Job Manager then starts the job and informs <i>nbpem.exe</i> that the job is completed.                                                                                                                                                          |
| NetBackup Policy Execution Manager ( <i>nbpem.exe</i> )              | Creates Policy/Client tasks and determinate when jobs are due to run. If a policy is modified or if an image expires, <i>nbpem</i> is notified and the Policy/Client task objects are updated.                                                                                                                                                                             |
| NetBackup Relational Database Manager<br>( <i>dbsrv11.exe</i> )      | Manages the NetBackup relational database. The service must be running on the NetBackup Enterprise Media Manager server during all normal NetBackup operations. The display name on Windows is <i>SQLANYs_VERITAS_NB</i> .                                                                                                                                                 |
| NetBackup Remote Manager and Monitor Service<br>( <i>nbmms.exe</i> ) | Discovers and monitors disk storage on NetBackup media servers. Also discovers, monitors, and manages Fibre Transport (FT) connections on media servers and clients for the NetBackup SAN Client option. Runs on NetBackup media servers.                                                                                                                                  |
| NetBackup Request Manager<br>( <i>bprd.exe</i> )                     | Processes the requests from NetBackup clients and servers. <i>bprd</i> also prompts NetBackup to perform automatically scheduled backups. <i>bprd</i> must be running on the NetBackup master server to perform any backups or restores.                                                                                                                                   |

**Table 20-1** NetBackup services (*continued*)

| Service                                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NetBackup Resource Broker<br>( <i>nbrb.exe</i> )               | Allocates the storage units, tape drives, and client reservations for jobs. <i>nbrb</i> works with the Enterprise Media Manager (NBEMM).                                                                                                                                                                                                                                                                                                                                                                |
| NetBackup Service Layer<br>( <i>nbsl.exe</i> )                 | Facilitates the communication between the NetBackup graphical user interface and NetBackup logic. NBSL is required to run Symantec OpsCenter, an application that manages and monitors multiple NetBackup environments.<br><br>The service cannot be stopped from the Activity Monitor because it receives data that appears in the Administration Console. If it is stopped, the console cannot display the data.                                                                                      |
| NetBackup Service Monitor<br>( <i>nbsvcmon.exe</i> )           | Monitors the NetBackup services that run on the local machine. If a service unexpectedly terminates, the service tries to restart the terminated service. If <i>nbsvcmon</i> determines that NetBackup is configured for a cluster, the service shuts down, and the monitoring is taken over by the cluster.<br><br>The service cannot be stopped from the Activity Monitor because it receives data that appears in the Administration Console. If it is stopped, the console cannot display the data. |
| NetBackup Storage Lifecycle Manager<br>( <i>nbstserv.exe</i> ) | The NetBackup Storage Lifecycle Manager manages lifecycle operations including duplication, staging, and image expiration.                                                                                                                                                                                                                                                                                                                                                                              |
| NetBackup Vault Manager<br>( <i>nbvault.exe</i> )              | Manages NetBackup Vault. <i>NBVAULT</i> must be running on the NetBackup Vault server during all NetBackup Vault operations.                                                                                                                                                                                                                                                                                                                                                                            |
| NetBackup Volume Manager<br>( <i>vmd.exe</i> )                 | Manages the volumes (tapes) needed for backup or restore and starts local device management daemons and processes.                                                                                                                                                                                                                                                                                                                                                                                      |

## Types of services

The following items describe additional information about NetBackup services:

|                      |                                                                                                                                                       |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| Stand-alone services | These NetBackup services always run and listen to accept connections. Examples include <i>bpdbm</i> , <i>bprd</i> , <i>bpjobjd</i> , and <i>vmd</i> . |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|

|                                     |                                                                                                                                                                                   |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multiprocess stand-alone services   | These NetBackup services "fork" a child process to handle requests. Examples include <code>bpd</code> and <code>bprd</code> .                                                     |
| Single-process stand-alone services | These NetBackup services accept connections and handle requests in the same process.                                                                                              |
| <code>inetd</code> services         | <code>inetd</code> (1m) or <code>bpinetd</code> usually launch these NetBackup services. Examples include <code>bpcd</code> , <code>bpjava-msvc</code> , and <code>vnetd</code> . |

## Other Symantec services

Several Symantec services do not appear in the Activity Monitor:

|                                                                       |                                                                                                                                                                                           |
|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Symantec Private Branch Exchange<br>( <code>pbx_exchange.exe</code> ) | The Symantec Private Branch Exchange provides single-port access to clients outside the firewall that connect to Symantec product services. Service name: <code>VRTSspb</code> .          |
| Symantec Product Authentication Service<br>( <code>vxat.exe</code> )  | The Symantec Product Authentication Service validates, identities, and forms the basis for authorization and access control in Symantec applications. Service name: <code>VRTSat</code> . |
| Symantec Product Authorization Service<br>( <code>vxazd.exe</code> )  | The Symantec Product Authorization Service provides access control in Symantec applications. Service name: <code>VRTSaz</code> .                                                          |

## Starting or stopping a service

The following procedure describes how to start or stop a service.

### To start or stop a service

- 1 Open the Activity Monitor and select the **Services** tab.
- 2 Select the service(s) you want to start or stop.
- 3 Select **Actions > Stop Selected** or **Actions > Start Selected**.

To start or stop services requires the necessary permissions on the system where the service is running.

## Monitoring NetBackup services

The following procedure describes how to monitor NetBackup services.

**To monitor NetBackup services**

- 1 Open the Activity Monitor and select the **Services** tab.
- 2 Double-click a service from the service list to view a detailed status.

To view the status of the previous service or the next service, click the up or down arrow.

To view the details of a service, double-click the process in the **Services** tab. For a description of the service details, click **Help** in the **Service Details** dialog box.

## About the Processes tab

The **Processes** tab displays the NetBackup processes that run on the master server.

---

**Note:** To view services on another system, the system must be a Microsoft platform and the user must be authenticated on the Microsoft platform.

---

Not all columns display by default. Click **View > Columns > Layout** to show or hide columns.

See [Table 20-2](#) on page 694. lists and describes the NetBackup processes.

**Table 20-2** NetBackup processes

| Process | Port | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| acsd    | 1370 | The <code>acsd</code> (Automated Cartridge System) daemon runs on the NetBackup media server and communicates mount and unmount requests to the host that controls the ACS robotics.                                                                                                                                                                                                                                                                                                                                                             |
| avrd    |      | The Automatic Volume Recognition process handles automatic volume recognition and label scans. The process allows NetBackup to read labeled tapes and assign the associated removable media requests to drives.                                                                                                                                                                                                                                                                                                                                  |
| bmrtd   | 8362 | The process for the NetBackup Bare Metal Restore Master Server service.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| bpcd    | 1370 | The NetBackup Client daemon, this process issues requests to and from the master server and the media server to start programs on remote hosts.<br><br>On UNIX clients, <code>bpcd</code> can only be run in stand-alone mode.<br><br>On Windows, <code>bpcd</code> always runs under the supervision of <code>bpnetd.exe</code> . NetBackup has a specific configuration parameter for <code>bpcd</code> : if the port number is changed within the NetBackup configuration, the software updates the port number in the services file as well. |

Table 20-2 NetBackup processes (continued)

| Process      | Port | Description                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bpcompatd    |      | The process for the NetBackup Compatibility service.                                                                                                                                                                                                                                                                                                                                                                 |
| bpdbm        | 1321 | The process for the NetBackup Database Manager service.<br>The process that responds to queries that are related to the NetBackup catalog.                                                                                                                                                                                                                                                                           |
| bpinetd      |      | The process for the NetBackup Client service.<br>The process that provides a listening service for connection requests.                                                                                                                                                                                                                                                                                              |
| bpjava-msvc  | 1322 | The NetBackup-Java application server authentication service program. <code>bpinetd</code> starts the program during startup of the NetBackup-Java GUI applications and authenticates the user that started the NetBackup-Java GUI application.                                                                                                                                                                      |
| bpjava-susvc |      | The NetBackup-Java application server user service program on NetBackup servers. <code>bpjava-msvc</code> starts the program upon successful login with the NetBackup-Java applications login dialog box. <code>bpjava-susvc</code> services all requests from the NetBackup-Java GUI applications for administration and end-user operations on the host on which the NetBackup-Java application server is running. |
| bpjobd       | 1323 | The NetBackup Jobs Database Management daemon. This process queries and updates the jobs database.                                                                                                                                                                                                                                                                                                                   |
| bprd         | 1320 | The process for the NetBackup Request Manager service.<br>The process that starts the automatic backup of clients and responds to client requests for file restores and user backups and archives.<br>NetBackup has a specific configuration parameter for <code>bprd</code> : if the port number changes within the NetBackup configuration, the software updates the port number in the services file as well.     |
| ltid         |      | The process for the NetBackup Device Manager service.                                                                                                                                                                                                                                                                                                                                                                |
| NBConsole    |      | The NetBackup Administration Console on the Windows platform.                                                                                                                                                                                                                                                                                                                                                        |
| nbemm        |      | The process for the NetBackup Enterprise Media Manager service.<br>The process that accesses and manages the database where media and device configuration information is stored ( <code>EMM_DATA.db</code> ). <code>nbemm.exe</code> must be running in order for jobs to run.                                                                                                                                      |
| nbEvtMgr     |      | The process for the NetBackup Event Manager service.<br>The process that creates and manages event channels and objects for communication among NetBackup daemon. The Event Manager daemon runs with the Enterprise Media Manager ( <code>nbemm</code> ) only on master servers.                                                                                                                                     |

**Table 20-2** NetBackup processes (*continued*)

| Process  | Port | Description                                                                                                                                                                                                                                                                                                   |
|----------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| nbfdrv64 |      | The process that controls the Fibre Transport target mode drivers on the media server. <code>nbfdrv64</code> runs on media servers configured for NetBackup Fibre Transport.                                                                                                                                  |
| nbftsrvr |      | The Fibre Transport (FT) server process that runs on media servers configured for NetBackup Fibre Transport. It does the following for the server side of the FT connection: controls data flow, processes SCSI commands, manages data buffers, and manages the target mode driver for the host bus adaptors. |
| nbjm     |      | The process for the NetBackup Job Manager service.<br><br>The process that accepts the jobs that the Policy Execution Manager (NBPEM) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.                             |
| nbpem    |      | The process for the NetBackup Policy Execution Manager service.<br><br>It creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, NBPEM is notified and the appropriate Policy/Client tasks are updated.                                         |
| nbproxy  |      | The process that safely allows multi-threaded NetBackup processes to use existing multi-threaded unsafe libraries.                                                                                                                                                                                            |
| nbrb     |      | This process allocates storage units, tape drives, and client reservations for jobs. <code>nbrb</code> works with the Enterprise Media Manager (NBEMM).                                                                                                                                                       |
| nbrmms   |      | The process for the NetBackup Remote Manager and Monitor service. Enables NetBackup to remotely manage and monitor resources on a system that are used for backup (or affected by backup activity).                                                                                                           |
| nbsl     |      | The process for the NetBackup Service Layer service.<br><br><code>nbsl</code> facilitates the communication between the graphical user interface and NetBackup logic.                                                                                                                                         |
| nbstserv |      | The process for the NetBackup Storage Lifecycle Manager. Manages storage lifecycle policy operations and schedules duplication jobs. Monitors disk capacity on capacity managed volumes and removes older images when required.                                                                               |
| nbsvcmon |      | The process for the NetBackup Service Monitor. Monitors the NetBackup services. When a service unexpectedly terminates, <code>nbsvcmon</code> attempts to restart the terminated service.                                                                                                                     |
| nbvault  |      | If Vault is installed, the process for the NetBackup Vault Manager service.                                                                                                                                                                                                                                   |
| ndmp     | 1000 | NDMP is the acronym for Network Data Management Protocol. NDMP servers are designed to adhere to this protocol and listen on port 10000 for NDMP clients to connect to them.                                                                                                                                  |



Table 20-2 NetBackup processes (continued)

| Process                                 | Port | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| opr <sub>d</sub>                        |      | The NetBackup Volume Manager (vmd) starts the opr <sub>d</sub> operator request daemon. This process receives requests to mount and unmount volumes and communicates the requests to the NetBackup Device Manager tl <sub>td</sub> . The NetBackup Device Manager communicates the requests to the robotics through SCSI interfaces.                                                                                                                                                             |
| spoold                                  |      | The process for the NetBackup Deduplication Engine service.<br>Active only if the NetBackup Media Server Deduplication option is licensed.                                                                                                                                                                                                                                                                                                                                                       |
| tl4 <sub>d</sub>                        | 1373 | The tl4 <sub>d</sub> process runs on the host that has a Tape Library 4mm. This process receives NetBackup Device Manager requests to mount and unmount volumes and communicates these requests to the robotics through SCSI interfaces.                                                                                                                                                                                                                                                         |
| tl8 <sub>d</sub><br>tl8 <sub>cd</sub>   | 1375 | The tl8 <sub>d</sub> process runs on a NetBackup media server that manages a drive in a Tape Library 8mm. This process receives NetBackup Device Manager requests to mount and unmount volumes, and sends these requests to the robotic-control process tl8 <sub>cd</sub> .<br>The tl8 <sub>cd</sub> process communicates with the TL8 robotics through SCSI interfaces.<br>To share the tape library, tl8 <sub>cd</sub> runs on the NetBackup server that provides the robotic control.         |
| tl1 <sub>dd</sub><br>tl1 <sub>dcd</sub> | 1371 | The tl1 <sub>dd</sub> process runs on a NetBackup server that manages drive in a Tape Library DLT. This process receives NetBackup Device Manager requests to mount and unmount volumes and sends these requests to the robotic-control process tl1 <sub>dcd</sub> .<br>The tl1 <sub>dcd</sub> process communicates with the Tape Library DLT robotics through SCSI interfaces.<br>To share the tape library, tl1 <sub>dcd</sub> runs on the NetBackup server that provides the robotic control. |
| tlh <sub>d</sub><br>tlh <sub>cd</sub>   | 1377 | The tlh <sub>d</sub> process runs on each NetBackup server that manages a drive in a Tape Library Half-inch. This process receives NetBackup Device Manager requests to mount and unmount volumes and sends these requests to the robotic-control process tlh <sub>cd</sub> .<br>The tlh <sub>cd</sub> process runs on the NetBackup server that provides the robotic control and communicates with the TLH robotics through SCSI interfaces.                                                    |
| tlm <sub>d</sub>                        | 1376 | The tl <sub>m</sub> d Tape Library Multimedia (TLM) daemon runs on a NetBackup server. It communicates mount, unmount, and robot inventory requests to a NetBackup media server that hosts ADIC DAS/SDLC software and controls the TLM robotics.                                                                                                                                                                                                                                                 |
| vmd                                     | 1301 | The process for the NetBackup Volume Manager service.                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 20-2** NetBackup processes (*continued*)

| Process        | Port | Description                                                                                                                                                                                                          |
|----------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vnetd          | 1374 | Veritas Network Daemon allows all socket communication to take place while connecting to a single port. Legacy NetBackup services that were introduced before NetBackup 6.0 use the vnetd port number.               |
| vrts-auth-port | 4032 | The Veritas Authorization Service verifies that an identity has permission to perform a specific task.                                                                                                               |
| vrts-at-port   | 2821 | The Veritas Authentication Service validates, identifies, and forms the basis for authorization and access.                                                                                                          |
| veritas_pbx    | 1556 | The Symantec Private Branch Exchange allows all socket communication to take place while connecting through a single port. NetBackup services that were introduced in NetBackup 6.0 use the veritas_pbx port number. |

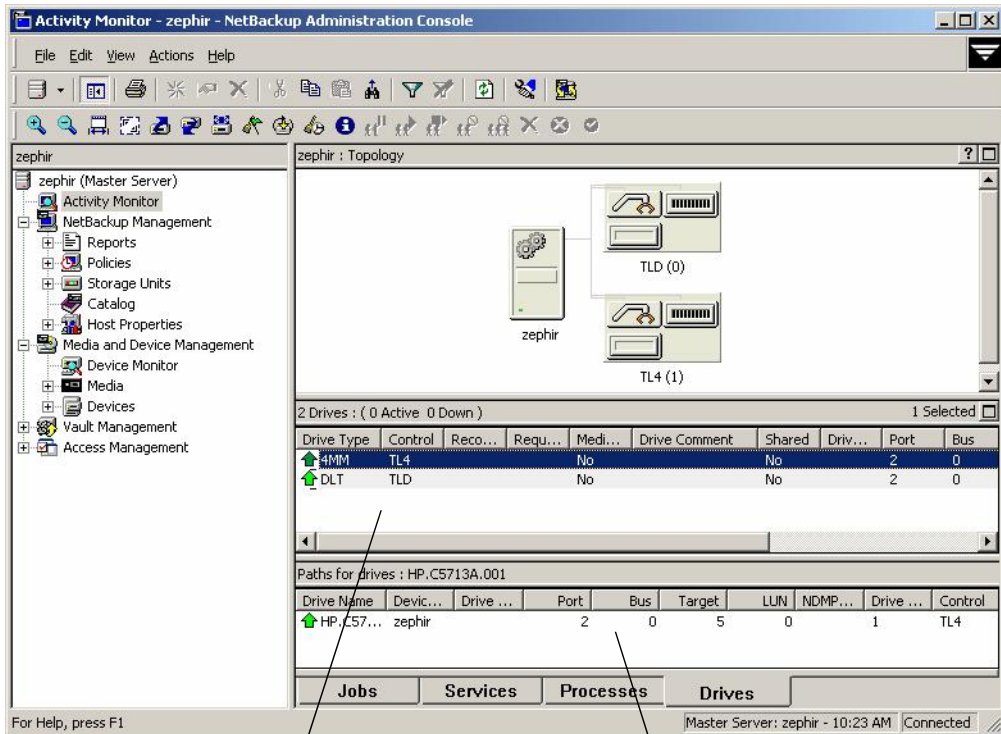
## Monitoring NetBackup processes

To view the details for a process, double-click the process in the Processes tab. For a description of the process details, click **Help** in the **Process Details** dialog box.

## About the Drives tab

The **Drives** tab displays the status of NetBackup drives on the master server being monitored.

Figure 20-4 Activity Monitor Drives tab



Drives pane

Drives Paths pane

The Drives Paths pane appears if a drive is configured as a shared drive, or if there are multiple paths to a drive configured. The Drive Paths pane lists path information for drives.

## Monitoring NetBackup tape drives

The following procedure describes how to monitor NetBackup drives.

### To monitor NetBackup drives

- 1 Open the Activity Monitor.
- 2 Select the **Drives** tab. Double-click a drive from the drive list to view a detailed status.
- 3 A **Drives Details** dialog box appears for the drive you selected. To view the status of the previous drive or the next drive, click the up or down arrow.

## Cleaning tape drives from the Activity Monitor

Drive cleaning functions can also be performed from the Device Monitor.

### To clean a tape drive

- 1 In the NetBackup Administration Console, select **Activity Monitor**. Then, select the Drives tab in the **Details** pane.
- 2 Select the drive that you want to clean.
- 3 Select **Actions > Drive Cleaning**, then select one of the following drive cleaning actions:
  - **Clean Now**  
Start an operator-initiated cleaning of the selected drive, regardless of the cleaning frequency or accumulated mount time. If the drive is a stand-alone drive, it must contain a cleaning tape for a mount request to be issued. **Clean Now** resets the mount time to zero, but the cleaning frequency value remains the same.
  - **Reset Mount Time**  
Reset the mount time for the selected drive to zero. Use Reset Mount Time to reset the mount time after doing a manual cleaning of a drive.
  - **Set Cleaning Frequency**  
Set the number of mount hours between drive cleanings.

## About the jobs database

NetBackup uses the `install_path\NetBackup\bin\admincmd\bpdjobs -clean` command to delete done jobs periodically.

By default, the `bpdjobs` process deletes all completed jobs that are more than three days old. By default, the `bpdjobs` process retains more recent done jobs until the three-day retention period expires.

If the `bprd` NetBackup request daemon is active, `bprd` starts the `bpdjobs` process automatically when it performs other cleanup tasks. The process starts the first time `bprd` wakes up after midnight. The automatic startups occur regardless of whether you choose to run `bpdjobs` at other times by using `cron` or alternate methods.

## Retaining job information in the database

You may want to keep jobs in the jobs database longer than the default of three days.

## Changing the default

To change the default values on a permanent basis, use the following method to add new registry key(s) to HKEY\_LOCAL\_MACHINE\SOFTWARE\VERITAS\NetBackup\

```
CurrentVersion\Config
```

To add the key(s) safely, run the following commands. For example:

```
install_path\VERITAS\NetBackup\bin\admincmd\
echo KEEP_JOBS_HOURS = 192 | bpsetconfig
```

Where 192 is the number of hours that unsuccessful jobs are kept in the jobs database or Activity Monitor display.

For example, run:

```
echo KEEP_JOBS_SUCCESSFUL_HOURS = 192 | bpsetconfig
```

Where 192 is the number of hours that successful jobs are kept in the jobs database or Activity Monitor display.

Consider the following notes when changing the default:

- The default values for `KEEP_JOBS_SUCCESSFUL_HOURS` and `KEEP_JOBS_HOURS` is 78 hours.
- The retention period values are measured against the time the job ended.
- Information about successful jobs cannot be kept longer than information about unsuccessful jobs. If `KEEP_JOBS_SUCCESSFUL_HOURS` is greater than `KEEP_JOBS_HOURS`, `bpdbjobs` sets `KEEP_JOBS_SUCCESSFUL_HOURS` to equal `KEEP_JOBS_HOURS`.
- If `KEEP_JOBS_SUCCESSFUL_HOURS` is set to 0, `bpjobd` uses the `KEEP_JOBS_HOURS` `bpdbjobs` value instead for successful jobs.  
If the `KEEP_JOBS_SUCCESSFUL_HOURS` value is greater than 0 but less than `KEEP_JOBS_HOURS`, `KEEP_JOBS_HOURS` is used for unsuccessful jobs only.

## About the BPDBJOBS\_OPTIONS environment variable

The `BPDBJOBS_OPTIONS` environment variable provides a convenient method to set job retention options with a script. The `bpdbjobs` process determines how long to retain a job by checking for the `BPDBJOBS_OPTIONS` environment variable. If present, `BPDBJOBS_OPTIONS` overrides the registry key settings.

The following options can be used to determine the length of time NetBackup retains jobs.

The options should be entered in lower case in the `BPDBJOBS_OPTIONS` environmental variable:

■ `-keep_hours` *hours*

Use with the `-clean` option to specify how many hours `bpdbjobs` keeps unsuccessfully completed jobs. Default: 78 hours.

To keep both successful and both failed jobs longer than the default of 78 hours, `keep_successful_hours` must be used with `keep_hours`

■ `-keep_successful_hours` *hours*

Use with the `-clean` option to specify how many hours `bpdbjobs` keeps successfully completed jobs. The number of hours must be less than or equal to `keep_hours`.

Values outside the range are ignored. Default: 78 hours.

■ `-keep_days` *days*

Use with the `-clean` option to specify how many days `bpdbjobs` keeps completed jobs. Default: 3 days.

■ `-keep_successful_days` *days*

Use with the `-clean` option to specify how many days `bpdbjobs` keeps successfully completed jobs. Default: 3 days.

This value must be less than the `-keep_days` value.

A batch file (`cleanjobs.bat`) was used in the following example. You can copy the script directly from this document and changed as needed.

- The first line specifies how long to keep unsuccessful jobs (24 hours) and successful jobs (five hours).
- The second line specifies the path to the `bpdbjobs` command. Indicate the correct location of `bpdbjobs` in the `.bat` file. In this example, NetBackup was installed in the default location:

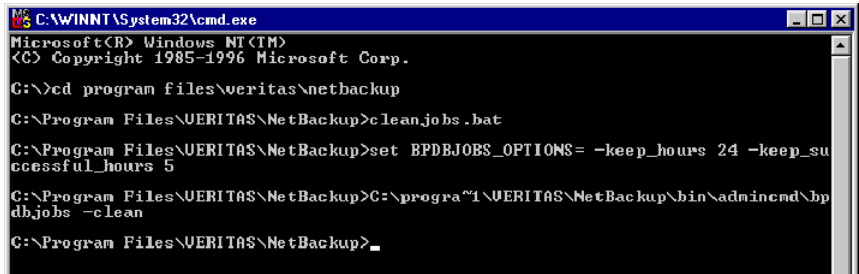
```
set BPDBJOBS_OPTIONS= -keep_hours 24 -keep_successful_hours 5
C:\progra~1\VERITAS\NetBackup\bin\admincmd\bpdbjobs -clean
```

You can store the `.bat` file anywhere, as long as it is run from the appropriate directory.

In the following example, the administrator created and stored `cleanjobs.bat` in `C:\Program Files\VERITAS\NetBackup`.

Figure 20-5 is a screen capture of `cleanjobs.bat` being run:

Figure 20-5 Running cleanjobs.bat



```
C:\WINNT\System32\cmd.exe
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

C:\>cd program files\veritas\netbackup

C:\Program Files\VERITAS\NetBackup>cleanjobs.bat

C:\Program Files\VERITAS\NetBackup>set BPDBJOBS_OPTIONS= -keep_hours 24 -keep_successful_hours 5

C:\Program Files\VERITAS\NetBackup>C:\program files\VERITAS\NetBackup\bin\admincmd\bpdbjobs -clean

C:\Program Files\VERITAS\NetBackup>_
```

## bpdbjobs command line options

The `bpdbjobs` command interacts with the jobs database to delete or move completed job files. The command line options supersede all other job retention instructions.

The `-clean` option causes **bpdbjobs** to delete the done jobs that are older than a specified time period as follows:

```
bpdbjobs -clean [-M <master servers>]
[-keep_hours <hours>] or [-keep_days <days>]
[-keep_successful_hours <hours>] or
[-keep_successful_days <days>]
```

For example, the following command deletes unsuccessful jobs older than 72 hours.

```
bpdbjobs -clean -keep_hours 72
```

More information is available in *NetBackup Commands*.

## About the bpdbjobs debug log

If you need detailed information on `bpdbjobs` activities, enable the `bpdbjobs` debug log by creating the following directory:

```
install_path\NetBackup\logs\bpdbjobs
```

---

**Note:** Before you use a debug log, read the guidelines in the Debug Logs section of the *NetBackup Troubleshooting Guide for UNIX and Windows*.

---

## About the Device Monitor

Use the **Device Monitor** to manage device paths, disk pools, service requests for operators, and tape drives.

## About media mount errors

Errors can occur when media is mounted for NetBackup jobs. Depending on the type of error, the request queues or it is canceled.

When the mount request is queued, an operator-pending action is created and appears in the **Device Monitor**.

A queued mount request leads to one of the following actions:

- The mount request is suspended until the condition is resolved.
- The operator denies the request.
- The media mount timeout is reached.

When a mount request is automatically canceled, NetBackup tries to select other media to use for backups. (Selection applies only in the case of backup requests.)

Many conditions lead to a mount request being automatically canceled instead of queued. When a media mount is canceled, different media is selected so that the backup is not held up.

The following conditions can lead to automatic media reselection:

- The requested media is in a DOWN drive.
- The requested media is misplaced.
- The requested media is write protected.
- The requested media is in a drive not accessible to the media server.
- The requested media is in an offline ACS LSM (Automated Cartridge System Library Storage Module). (ACS robot type only.)
- The requested media has an unreadable barcode. (ACS robot type only.)
- The requested media is in an ACS that is not accessible. (ACS robot type only.)
- The requested media is determined to be unmountable.



## About pending requests and actions

If a tape mount requires a specific volume, the request appears in the **Pending Requests** pane of the **Device Monitor** window. For example, if NetBackup requires a specific volume for a restore operation, NetBackup loads or requests the volume.

**Table 20-3** Pending states

| Pending state   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pending request | <p>A pending request is for a tape mount that NetBackup cannot service automatically. Operator assistance is required to complete the request. NetBackup displays the request in the <b>Pending Requests</b> pane.</p> <p>NetBackup assigns pending status to a mount request when it cannot determine the following:</p> <ul style="list-style-type: none"> <li>■ Which stand-alone drive to use for a job.</li> <li>■ Which drive in a robot is in Automatic Volume Recognition (AVR) mode.</li> </ul> |
| Pending action  | <p>A tape mount request becomes a pending action when the mount operation encounters problems, and the tape cannot be mounted. Operator assistance is required to complete the request, and NetBackup displays an action request in the <b>Pending Requests</b> pane. Pending actions usually occur with drives in robotic libraries.</p>                                                                                                                                                                |

The **Pending Requests** pane appears only if requests await action or when NetBackup acts on a request. After all requests are resolved (automatically by NetBackup or manually by operator intervention), the **Pending Requests** pane disappears.

If NetBackup cannot service a media-specific mount request automatically, it changes the request or action to a pending state.

### About pending requests for storage units

The following tape mount requests do not appear in the **Device Monitor Pending Requests** pane:

- Requests for backups
- Requests for a tape that is required as the target of a duplication operation

Such requests are for resources in a storage unit and therefore are not for a specific volume. NetBackup does not assign a mount request for one storage unit to the drives of another storage unit automatically. Also, you cannot reassign the mount request to another storage unit.

If the storage unit is not available, NetBackup tries to select another storage unit that has a working robot. If NetBackup cannot find a storage unit for the job, NetBackup queues the job (a **Queued** state appears in the **Activity Monitor**).

You can configure NetBackup so that storage unit mount requests are displayed in the **Device Monitor** if the robot or drive is down. Pending requests appear in the **Device Monitor**, and you can assign these mount requests to drives manually.

See [“Configuring a robot to operate in manual mode”](#) on page 239.

## Managing pending requests and actions

You can perform various actions to resolve or deny pending requests and actions.

### Resolving a pending request

Use the following procedure to resolve a pending request.

For ACS robots: If a request pends because the Library Storage Module (LSM) in which the media resides is offline, no operator action is required. NetBackup retries such requests hourly until the LSM is online. NetBackup reports the LSM offline status in the **Job Details** dialog box. Open the **Job Details** dialog box from the **Jobs** tab in the **Activity Monitor**.

#### To resolve a pending request on Windows (Enterprise Server only)

- 1 If the drive and the request are on the same host, select the request in the **Pending Requests** pane.
- 2 Drag it to the **Drive Status** pane and then drop it on the wanted drive.

#### To resolve a pending request

- 1 Insert the requested volume in a drive that matches the density of the volume that was requested.
- 2 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 3 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 4 In the **Pending Requests** pane, select the request and note the contents of the following columns of the request:
  - Density
  - External Media ID
  - Mode

- 5 In the **Drive Status** pane, find a drive type that matches the density for the pending request.
- 6 Verify that the drive is up and not assigned to another request.
- 7 Select the drive.
- 8 The following applies only to NetBackup Enterprise Server: Ensure that the drive and the pending request are on the same host.
- 9 If necessary, get the media, write-enable it, and insert it into the drive.
- 10 Wait for the drive to become ready, as explained in the vendor's drive equipment manual.
- 11 On the **Actions** menu, select **Assign Request**.
- 12 Verify that the request was removed from the Pending Requests pane.
- 13 In the **Drive status** pane, verify the following:
  - The job request ID appears in the Request ID column for the drive
  - The User column is not blank

## Resolving a pending action

Use the following procedure to resolve a pending action.

For a pending action, NetBackup determines the cause of the problem and issues instruction to the operator to resolve the problem.

A pending action is similar to a pending request. An asterisk identifies a pending action; the asterisk appears to the left of the request ID.

### To resolve a pending action

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Pending Requests** pane, select the pending action.
- 4 On the **Actions** menu, select **Display Pending Action**.
- 5 In the message box that describes the problem, review the list of possible corrective actions. The message box also shows other information, such as user name, recorded media ID, external media IDs, and drive number.

- 6 Click **OK**.
- 7 Correct the error condition and either resubmit the request or deny the request.  
See [“Resubmitting a request”](#) on page 708.  
See [“Denying a request”](#) on page 708.

## Resubmitting a request

After you correct the problem with a pending action, you can resubmit the request.

Use the following procedure to resubmit a request.

If the problem is a volume missing from a robot, first locate the volume, insert it into the robot, and then update the volume configuration. Usually, a missing volume was removed from a robot and then requested by NetBackup.

See [“Robot inventory options”](#) on page 316.

### To resubmit a request

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Pending Requests** pane, select the request.
- 4 On the **Actions** menu, select **Resubmit Request**.

## Denying a request

Some situations may require that you deny requests for service. For example, when a drive is not available, you cannot find the volume, or the user is not authorized to use the volume. When you deny a request, NetBackup sends an appropriate status message to the user.

Use the following procedure to deny a request.

### To deny a request

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Device Monitor**.
- 2 If an Enterprise Disk Option license is installed, select the **Drives** tab.
- 3 In the **Pending Requests** pane, select the request.
- 4 On the **Actions** menu, select **Deny Request**.

# Reporting in NetBackup

This chapter includes the following topics:

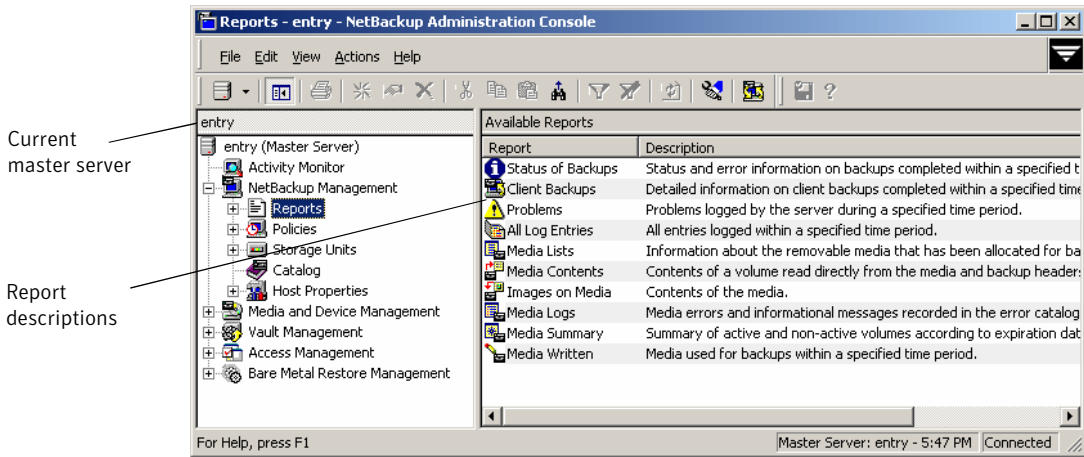
- [About the Reports utility](#)
- [About the Reports window](#)
- [Running a report](#)
- [Running the Troubleshooter within reports](#)
- [Copying report text to another document](#)
- [Saving or exporting a report](#)
- [Printing a report](#)

## About the Reports utility

Use the Reports utility to generate reports on many aspects of the NetBackup environment. The reports serve to verify, manage, and troubleshoot NetBackup operations. NetBackup reports display information according to job status, client backups, and media contents. The Troubleshooter is available within the Reports utility to help analyze the cause of errors that can appear in a NetBackup report.

Expand Reports in the NetBackup Administration Console to display a description of all possible reports in the **Details** pane.

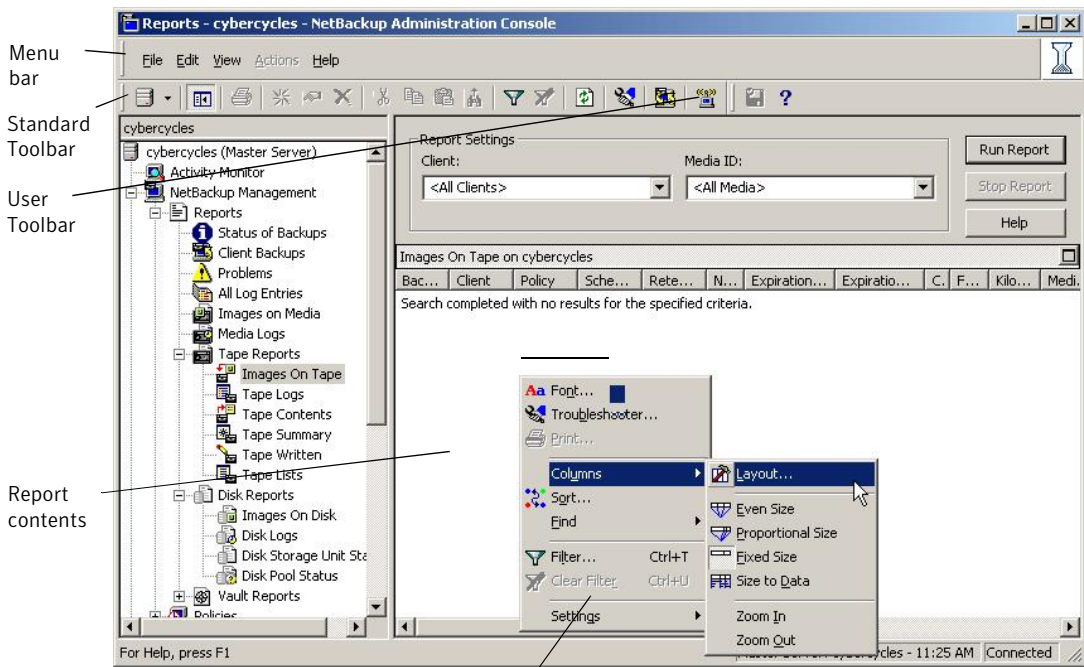
Figure 21-1 NetBackup Reports utility



## About the Reports window

The Reports window contains multiple ways to view report listings and manage report data.

**Figure 21-2** NetBackup Reports window



Right-click in contents area to display shortcut menus

## About the Reports shortcut menus

To display a list of commands that apply to a list, right-click on a report.

Depending on which report is viewed, the shortcut list may include:

- Font** Use to change the typeface and point size of the display text in the report pane.
- Troubleshooter** Launches the Troubleshooter. The Troubleshooter is available only when a line in a report is selected that displays a NetBackup status code.

See [“Running the Troubleshooter within reports”](#) on page 715.

|              |                                                                                                                                                                                                                 |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Print        | Opens the system Print dialog box to print the contents of the generated report.<br>See <a href="#">“Printing a report”</a> on page 716.                                                                        |
| Columns      | Opens a submenu that contains commands for changing the order and size of columns. Includes: <b>Layout</b> options, <b>Even Size</b> , <b>Proportional Size</b> , <b>Fixed Size</b> , and <b>Size to Data</b> . |
| Sort         | Use to specify sort criteria for the columns.                                                                                                                                                                   |
| Find         | Use to find text within the report.                                                                                                                                                                             |
| Filter       | Use the Filter option to narrow in on specific data in a table. Use the controls on the <b>Filter</b> dialog box to list the rows that match specified criteria.                                                |
| Clear Filter | Clears the filter if a filter is currently in effect.                                                                                                                                                           |
| Settings     | Options to reload settings from default, save the current settings as the default settings, or save settings on exit.                                                                                           |

## Reports settings

Use the report settings to specify the following criteria for building a report. Not all settings are available for every report type.

## Report types

Select a report type in the NetBackup Administration Console under **NetBackup Management > Reports**. Reports may present information about images storage on disk or tape media.

Select the Report Settings for the specific report, then, click **Run Report**. Click **Help** within the report window for a description of each column.

The following table describes the contents of the NetBackup reports, tape reports, and disk reports:

|                          |                                                                                                                                                                                   |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status of Backups report | Displays the status and the error information about the jobs that completed within the specified time period. If an error occurred, a short explanation of the error is included. |
| Client Backups report    | Displays the detailed information on the backups that complete within the specified time period.                                                                                  |



|                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Problems report        | Lists any problems that the server has logged during the specified time period. The information in this report is a subset of the information from the All Log Entries report.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| All Log Entries report | Lists all log entries for the specified time period. This report includes the information from the Problems report and Media Logs report. This report also displays the transfer rate. The transfer rate is useful to determine and predict rates and backup times for future backups. (The transfer rate does not appear for multiplexed backups.)                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Images on Media report | Lists the contents of the media as recorded in the NetBackup image catalog. You can generate this report for any type of media (including disk) and filter it according to client, media ID, or path.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Media Logs             | Displays the media errors or the informational messages that are recorded in the NetBackup error catalog.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Images on Tape report  | Generates the image list present on the tape storage units that are connected to the media server. The report is a subset of the Images on Media report and displays only tape-specific columns.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Tape Logs report       | Displays the media errors or the informational messages that are recorded in the NetBackup error catalog. The report is a subset of the Media Logs report and displays only tape-specific columns.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Tape Contents report   | <p>Displays the contents of a volume as read directly from the media header and backup headers. This report lists the backup IDs (not each individual file) that are on a single volume. If a tape must be mounted, the delay is longer before the report appears.</p> <p>Before running the Tape Contents report, you can choose to override the default job priority for the job.</p> <p>To change the priority for this job, enable <b>Override default job priority</b>. Then, select a value in the <b>Job Priority</b> field.</p> <p>If this option is not enabled, the job runs using the default priority as specified in the <b>Default Job Priorities</b> host properties.</p> <p>See <a href="#">“Default Job Priorities properties”</a> on page 105.</p> |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tape Summary report             | <p>Summarizes the volumes that are active and inactive for the specified media owner according to expiration date. It also displays how many volumes are at each retention level. In verbose mode, the report displays each media ID and the expiration date.</p> <p>Inactive media are those with a status of FULL, FROZEN, SUSPENDED, or IMPORTED. Other volumes are considered active.</p> <p>Only FROZEN expired volumes appear in the report. NetBackup deletes other expired volumes from the media catalog when backups are run. Expired, non-FROZEN volumes appear if the report is run between the time the volumes expire and the time that the next backup is run.</p> |
| Tape Written report             | <p>Identifies the volumes that were used for backups within the specified time period. The report also does not display the volumes that were used for duplication if the original was created before the specified time period.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Tape Lists report               | <p>Displays the information for the volumes that have been allocated for backups for the selected media owner or media ID. This report does not show media for disk type storage units.</p> <p>For information about the backups that are saved to disk storage units, use the Images on Media report.</p>                                                                                                                                                                                                                                                                                                                                                                        |
| Images on Disk report           | <p>Generates a list of images that are expected to be present on the disk storage units, according to the NetBackup catalog. The report is a subset of the Images on Media report and displays only disk-specific columns.</p> <p>The report provides a summary of the storage unit contents. If a disk becomes bad or if a media server crashes, this report can let you know what data is lost.</p>                                                                                                                                                                                                                                                                             |
| Disk Logs report                | <p>Displays the media errors or the informational messages that are recorded in the NetBackup error catalog. The report is a subset of the Media Logs report and displays only disk-specific columns.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Disk Storage Unit Status report | <p>Displays the state of disk storage units in the current NetBackup configuration. For example, the total capacity and the used capacity of the disk storage unit.</p> <p>Storage units that reference disk pools do not display capacity values. To view these values, expand <b>Media and Device Management &gt; Devices &gt; Disk Pools</b>.</p>                                                                                                                                                                                                                                                                                                                              |
| Disk Pool Status report         | <p>Displays the state of disk pool storage units. This report displays only when an Enterprise Disk Option is installed.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Running a report

The following procedure describes how to run a report.

### To run a report

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.

See “[Report types](#)” on page 712.

The report information is for the master server that is currently selected. To run a report on a different master server, click **File > Change Server**.

See “[Accessing remote servers](#)” on page 727.

- 2 Double-click the name of the report you want to run.
- 3 Select the media servers and clients on which to run the report, then select the time period for which the report runs.
- 4 Click **Run Report**.

## Running the Troubleshooter within reports

Use the Troubleshooter within Reports to find explanations and the corrective actions that are based on the NetBackup status code that the job returns.

### To run the Troubleshooter within reports

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.
- 2 Run a report.
- 3 Right-click a line in the report. Then select **Troubleshooter** from the shortcut menu.
- 4 The Troubleshooter dialog box appears with an explanation of the problem on the **Problem** tab. A recommended action appears on the **Troubleshoot** tab.

Open the Troubleshooter at any time (**Help > Troubleshooter**), enter a status code, and click **Lookup**.

## Copying report text to another document

The following procedure describes how copy report text to another document.

#### To copy report text to another document

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.
- 2 Run a report, and then select the report text to copy.
- 3 Click **Edit > Copy** or press Ctrl+C.
- 4 Paste the selected text into a document (for example, an Excel document).

## Saving or exporting a report

The following procedure describes how to save or export a report.

#### To save or export a report

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.
- 2 Run a report, and then click the **Export** button or click **File > Export**.
- 3 In the **Save As** dialog box, select the drive and directory where you want to save the report.
- 4 Specify the file name and file type.
- 5 Click **Save**.

## Printing a report

The following procedure describes how to print a report.

#### To print a report

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.
- 2 Run a report, and then click the **Print** button or click **File > Print**.
- 3 In the Print dialog box, specify the settings. You can also specify font and column settings for the report.
- 4 Click **OK**.

# Administering NetBackup

- [Chapter 22. Management topics](#)
- [Chapter 23. Accessing a remote server](#)
- [Chapter 24. Using the NetBackup-Java administration console](#)
- [Chapter 25. Alternate server restores](#)
- [Chapter 26. Managing client restores](#)
- [Chapter 27. Powering down and rebooting NetBackup servers](#)
- [Chapter 28. About Granular Recovery Technology](#)



# Management topics

This chapter includes the following topics:

- [NetBackup naming conventions](#)
- [Wildcards in NetBackup](#)
- [How to administer devices on other servers](#)
- [How to access media and devices on other hosts](#)
- [About the Enterprise Media Manager](#)

## NetBackup naming conventions

The following set of characters can be used in user-defined names, such as storage units and policies:

- Alphabetic (A-Z a-z) (names are case sensitive)
- Numeric (0-9)
- Period (.)
- Plus (+)
- Minus (-)  
Do not use a minus as the first character.
- Underscore (\_)

These characters are used for foreign languages as well. Spaces are only allowed in a drive comment.

# Wildcards in NetBackup

NetBackup recognizes the following wildcard characters in areas where wildcards can be used. (For example, paths in the backup selections list and exclude file lists.)

**Table 22-1** shows the wildcards that can be used in various NetBackup dialog boxes and lists.

**Table 22-1** Wildcard use in NetBackup

| Wildcard | Use                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| *        | <p>Serve as a wildcard for zero or more characters.</p> <p>An asterisk can be used in the backup selection list, the include list, and the exclude list for Windows, UNIX, and Exchange clients.</p> <p>For example:</p> <p><code>r*</code> refers to all files that begin with <code>r</code></p> <p><code>r*.doc</code> refers to all files that begin with <code>r</code> and end with <code>.doc</code>.</p> <p>To back up all files that end in <code>.conf</code>, specify:</p> <p><code>/etc/*.conf</code></p>                                                                                                                              |
| ?        | <p>Serves as a wildcard for any single character (A through Z; 0 through 9).</p> <p>A question mark can be used in the backup selection list, the include list, and the exclude list for Windows, UNIX, and Exchange clients.</p> <p>For example:</p> <p><code>file?</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code></p> <p><code>file??</code> refers to <code>file12</code>, <code>file28</code>, <code>file89</code></p> <p>To back up all files named <code>log01_03</code>, <code>log02_03</code>, specify:</p> <p><code>c:\system\log??_03</code></p>                                                            |
| [ ]      | <p>Use a pair of square brackets to indicate any single character or range of characters that are separated with a dash.</p> <p>Square brackets can be used in the backup selection list, the include list, and the exclude list for Windows, UNIX, and Exchange clients.</p> <p>For example:</p> <p><code>file[2-4]</code> refers to <code>file2</code>, <code>file3</code>, and <code>file4</code></p> <p><code>file[24]</code> refers to <code>file2</code>, <code>file4</code></p> <p><code>*[2-4]</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code>, <code>name2</code>, <code>name3</code>, <code>name4</code></p> |



**Table 22-1** Wildcard use in NetBackup (*continued*)

| Wildcard | Use                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| { }      | <p>Curly brackets can be used in the backup selection list, the include list, and the exclude list for UNIX clients only.</p> <p>Use a pair of curly brackets (or braces) to indicate multiple file name patterns. Separate the patterns by commas only; no spaces are permitted. A match is made for any or all entries.</p> <p>For example:</p> <p>{*1.doc, *.pdf} refers to file1.doc, file1.pdf, file2.pdf</p> <p><b>Note:</b> Curly brackets are valid characters for Windows file names and cannot be used as wildcards on Windows platforms. Backslashes cannot be used as escape characters for curly bracket characters.</p> |

To use wildcard characters literally, precede the character with a backslash (\).

A backslash (\) acts as an escape character only when it precedes a special or a wildcard character. NetBackup normally interprets a backslash literally because a backslash is a legal character to use in paths.

Assume the brackets in the following are to be used literally:

C:\abc\fun[ny]name

In the exclude list, precede the brackets with a backslash:

C:\abc\fun\[ny\]name

## How to administer devices on other servers

The following applies only to NetBackup Enterprise Server.

The **NetBackup Administration Console** on the master server is the central management console for NetBackup servers, NetBackup clients, and storage devices in the environment. You can configure and manage the storage devices on all of the media servers from an **Administration Console** that is connected to the master server.

Alternatively, you can administer the devices on a specific media server from an **Administration Console** connected to that media server. To do so, change to or log into the media server by using one of the following methods:

- On the **File** menu, select **Change Server** in an existing instance of the **NetBackup Administration Console** and change to the media server.

- Start the **NetBackup Administration Console** on the media server.
- See “[Choosing a remote server to administer](#)” on page 731.

For device discovery, configuration, and management to occur, the following must be true:

- The devices must be configured correctly in the operating system of the media server host.
- The media server must be in the additional servers list on the NetBackup master server and the EMM server. Normally, the EMM server resides on the same machine as the NetBackup master server.
- The EMM server must be up and running, both when you install the media server software and when you configure the devices.

If the EMM server is not running when you install a media server, the media server is not registered. You cannot discover, configure, and manage the devices of that media server. You must register the media server with the EMM server.

The following procedure assumes that all other steps to add a media server are accomplished.

Information on how to add a media server is available.

See the *NetBackup Administrator's Guide, Volume II*.

## How to access media and devices on other hosts

For NetBackup to access media and device management functionality on a remote NetBackup host, you may need to add a `SERVER` entry to the `vm.conf` file on the remote host.

`SERVER` entries are used in the NetBackup `bp.conf` and `vm.conf` files for security. You can add entries that allow only specific hosts to access those capabilities remotely.

If the `vm.conf` file on a remote host contains no `SERVER` entries, a host can manage media and devices on the remote host if it is added to the `bp.conf` file of the server you logged into. You do not need to add a `SERVER` entry to the `vm.conf` file.

If the `vm.conf` file on a remote host contains any `SERVER` entries, add a `SERVER` entry for the host on which the NetBackup Administration Console is running (the server you logged into) to that `vm.conf` file.

## Example SERVER entries

Assume that you have three hosts named eel, yak, and shark. You want to centralize device management on host shark and also permit each host to manage its own devices.

The following example scenerio applies:

- The **vm.conf** file on shark contains the following:

```
SERVER = shark
```

The **vm.conf** file on shark does not require any additional **SERVER** entries, because all device management for shark are performed from shark.

- The **vm.conf** file on eel contains the following, which lets eel manage its own devices and permits shark to access them:

```
SERVER = eel
SERVER = shark
```

- The **vm.conf** file on yak contains the following, which lets yak manage its own devices and permits shark to access them:

```
SERVER = yak
SERVER = shark
```

## About the Enterprise Media Manager

The Enterprise Media Manager (EMM) is a NetBackup service that manages the device and the media information for NetBackup. The Enterprise Media Manager stores its managed information in a database, and the database resides on the EMM host.

See [“About the Enterprise Media Manager \(EMM\) database”](#) on page 603.

NetBackup is based on a static configuration of devices. These configurations are persistent for robotic libraries and tape drives in the NetBackup EMM database.

The Enterprise Media Manager manages the following:

- All media servers and their current status (online, offline).
- All drive allocations
- All configured devices

A NetBackup master server can have only one EMM server. However, an EMM server can manage device and media information for more than one NetBackup

master server. An EMM domain comprises all of the master and the media servers for which it manages device and media information.

NetBackup configures the EMM server when you install NetBackup.

Usually, the EMM service runs on the master server host. However, you can install and run the EMM service on a NetBackup media server.

## Enterprise Media Manager domain requirements

Applies only to NetBackup Enterprise Server.

An Enterprise Media Manager domain includes all of the servers in the Enterprise Media Manager database and the devices, media, and storage they manage. The Enterprise Media Manager can manage more than one NetBackup master server. That is, multiple NetBackup master server domains can share one Enterprise Media Manager domain.

The following are the rules for an EMM domain:

- The Enterprise Media Manager must be installed on a system that hosts a NetBackup master or media server. Symantec recommends that you install the EMM on the same system as a NetBackup master server.
- Host names must be consistent throughout an EMM domain. Do not use a fully qualified name and an unqualified name to refer to the same host. Do not use a physical name and a virtual host name to refer to the same host.
- All hosts in the same NetBackup domain must use the same EMM server.
- Robot numbers must be unique within an EMM domain.
- Media IDs must be unique within an EMM domain.
- Bar codes must be unique within an EMM domain.
- Drive names must be unique within an EMM domain and should be descriptive.
- Users cannot share devices or volumes between EMM domains.

## Sharing an EMM server

Although multiple domains can share an EMM server, Symantec does not recommend this configuration. The only situation that merits a shared EMM server is a configuration where multiple NetBackup domains share storage devices. However, there is no performance advantage to this type of configuration.

Care must be taken when you implement a catalog backup and recovery strategy, since all domains create backups of the central EMM database. Restoring any

catalog backup can result in inconsistencies in the catalogs of other domains that share the same EMM server.

If you use one EMM domain for multiple master server domains, observe the following:

- The EMM should reside on one of the NetBackup master servers. Only one EMM server should exist per EMM domain.
- Each master server must be allowed access to the EMM host. Use the **Servers** host property on the EMM host to allow access.
- All names and numbers for devices and all media IDs and bar codes should remain unique across the entire enterprise.



# Accessing a remote server

This chapter includes the following topics:

- [Accessing remote servers](#)
- [Adding a NetBackup server to a server list](#)
- [Choosing a remote server to administer](#)
- [Using the Remote Administration Console](#)
- [Using the Java Windows Administration Console](#)
- [Running the Administration Console on a NetBackup client](#)
- [Troubleshooting remote server administration](#)

## Accessing remote servers

If a NetBackup site has multiple master servers, you can configure the systems so that multiple servers can be accessed from one NetBackup Administration Console.

A host running NetBackup Enterprise Server or NetBackup Server may use the **Change Server** command to access another host. The other host must run either NetBackup Enterprise Server or NetBackup Server.

To access remote servers perform the following actions:


- First, make the remote server accessible to the local server.  
See [“Adding a NetBackup server to a server list”](#) on page 728.
- Second, indicate the remote server you want to administer.  
See [“Choosing a remote server to administer”](#) on page 731.

## Adding a NetBackup server to a server list

For a local host to administer a remote server, the name of the local host must appear in the server list of the remote server.


Figure 23-1 assumes that server\_1 wants to administer server\_2.

Figure 23-1 Server accessing a remote server



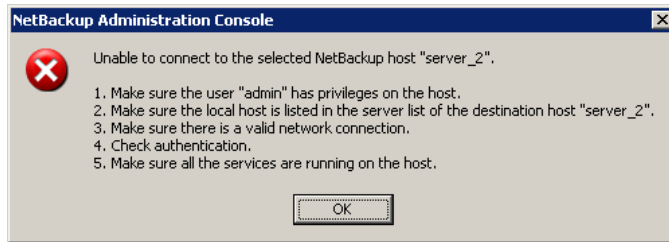
On server\_1, in the NetBackup Administration Console, select **File > Change Server** and type **server\_2** as the host name.

Figure 23-2 Changing the host name



If server\_1 is not listed on the server list of server\_2, server\_1 receives an error message after it tries to change servers to server\_2.





To add server\_1 to the server list of server\_2, see the following topics:

See [“Adding a server to a remote server list”](#) on page 729.

Other reasons may exist why a remote server is inaccessible:

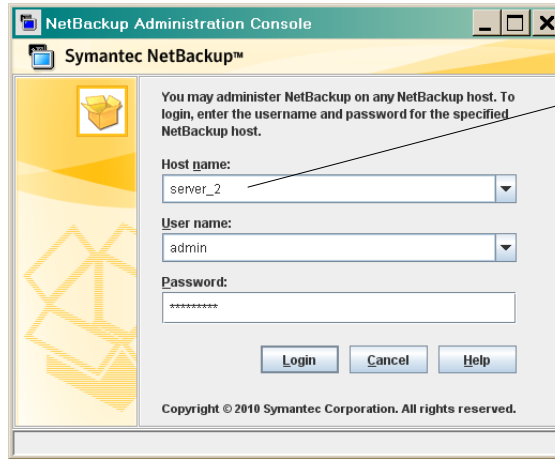
See [“Troubleshooting remote server administration”](#) on page 736.

## Adding a server to a remote server list

Use the following procedure to add a server to the server list of a remote server. This procedure is necessary to allow remote access to the server.

### To add a server to the server list of a remote server

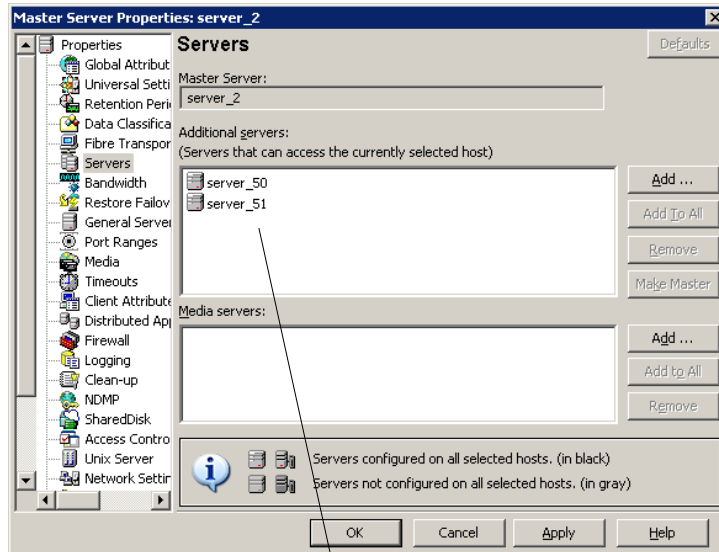
- 1 Access the server properties of the remote server in one of the following ways:
  - Physically go to the Windows destination host (server\_2) and start the NetBackup Administration Console.
  - Start the Java Windows Administration Console, if installed, on the local Windows host. Indicate the destination host (server\_2) on the login dialog box.
  - Physically go to the UNIX destination host (server\_2) and start jnbSA. Indicate server\_2 on the logon dialog box.
  - Start the NetBackup-Java Administration Console (jnbSA) on the local UNIX server (server\_1). Indicate the destination host server\_2 on the login dialog box.



Log in to server\_2 from server\_1. The user name must have sufficient privileges. Or, log in at server\_2.


- 2 Expand **Host Properties > Master Server**.
- 3 Double-click the server name (server\_2) to view the properties.
- 4 Select the **Servers** tab to display the server list.

Since the server list does not include server\_1, server\_2 considers server\_1 to be an invalid server.

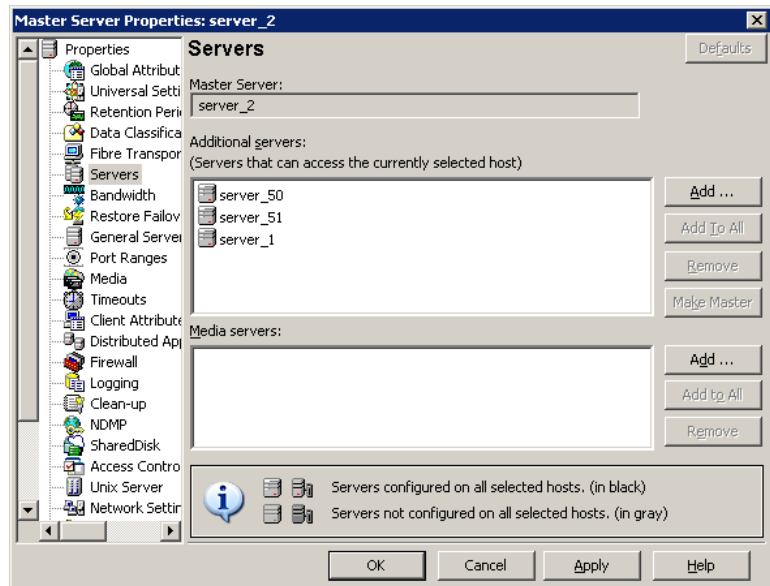


Currently, server\_2 allows remote access to two servers: server\_50 and server\_51

- 5 To add a server to the server list, click **Add**.
- 6 In the **Add New Server Entry** dialog box, type the server name (server\_2) in the field.



- 7 Click **Add** to add the server to the list. Then, click **Close** to close the dialog box. The server appears in the server list.



- 8 Click **OK** to save the changes.

## Choosing a remote server to administer

To indicate a remote server, use one of the following methods:

- Select the **File > Change Server** menu command in the NetBackup Administration Console.

See [“Using the change server command to administer a remote server”](#) on page 732.

- Specify the remote server in the host name field to start the NetBackup-Java console.

See [“Indicating a remote system upon login”](#) on page 733.

For a local host to administer a remote server, the name of the local host must appear in the server list of the remote server.

See [“Adding a server to a remote server list”](#) on page 729.

## Using the change server command to administer a remote server

Use the following procedure to change the NetBackup Administration Console to a different (or remote) server.

### To use the change server command to administer a remote server

- 1 Start the NetBackup Administration Console on the system:
  - To start the console on a Windows NetBackup server, select **Start > Programs > Symantec NetBackup > NetBackup Administration Console**.
  - To start the console on a Windows system with the NetBackup Remote Administration Console installed, select **Start > Programs > Symantec NetBackup > NetBackup Administration Console**.  
See [“Using the Remote Administration Console”](#) on page 734.
  - To start the console on the Windows system where the Java Windows Administration Console is installed, select **Start > Programs > Symantec NetBackup > NetBackup-Java Version 7.0**.
- 2 Select **File > Change Server**.
- 3 Enter or select the host name and click **OK**.

If the user has the necessary permissions on both servers, the user can transition from one to another without setting up trust relationships.

See [“Adding a server to a remote server list”](#) on page 729.

If the user has administrative privileges on one server and different privileges on another server, the user is required to reauthenticate.

Select **File > Login as New User** to reauthenticate from the NetBackup Administration Console. Or, close and reopen the NetBackup-Java Administration Console, then log in as a different user.

## Indicating a remote system upon login

Use the following procedure to indicate a remote system upon logging in to NetBackup.

This procedure requires that the administrator has one of the following available:

- A Windows system with the Java Windows Administration Console installed.
- A NetBackup-Java capable machine.

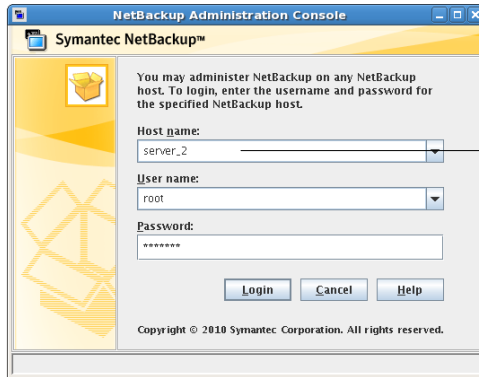
### To indicate a remote system upon login

- 1 Log in to the NetBackup client or server where you want to start the NetBackup Administration Console:
  - To start the console on the Windows system where the Java Windows Administration Console is installed:  
Select **Start > Programs > Symantec NetBackup > NetBackup-Java Version 7.0**.
  - To start the NetBackup Administration Console on a NetBackup-Java capable machine, run `jnbSA` as follows:

```
/usr/opensv/java/jnbSA
```

- 2 In the NetBackup Administration Console log in screen, specify the remote server to manage.

Type the user name and password for an authorized NetBackup administrator, then click **Login**.



To log in to a remote server, specify the name of the remote host in the login screen

This process logs you in to the NetBackup-Java application server program on the specified server.

The console program continues to communicate through the server you specified for the remainder of the current session.

See [“Using the NetBackup-Java administration console”](#) on page 739.

See [“Restricting access to NetBackup-Java applications on Windows”](#) on page 748.

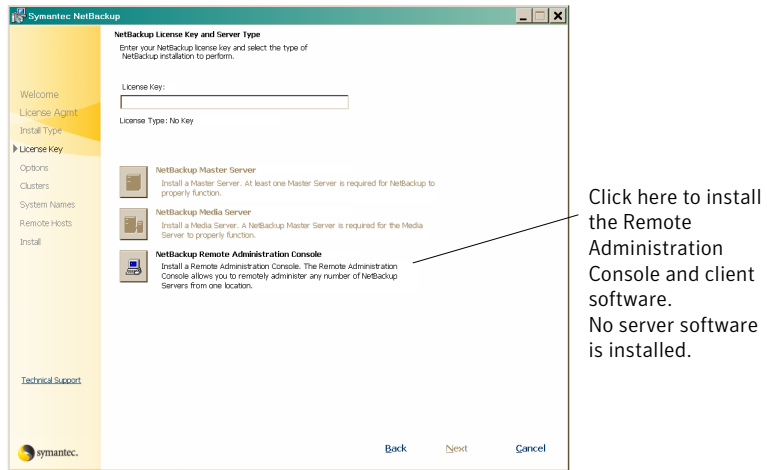
## Using the Remote Administration Console

Install the NetBackup Remote Administration Console on a Windows machine to remotely manage a Windows or UNIX server. No license is required to install only the console.

Installing the NetBackup Remote Administration Console installs the administration console and the client software. The presence of the client software enables the machine to be backed up like any other client. No master server software or media server software is installed.

[Figure 23-3](#) shows how to install the Remote Administration Console.

Figure 23-3 Remote Administration Console selection on the installation screen



Start the NetBackup Administration Console, then select **File > Change Server** to change to a NetBackup server.

See [“Adding a server to a remote server list”](#) on page 729.

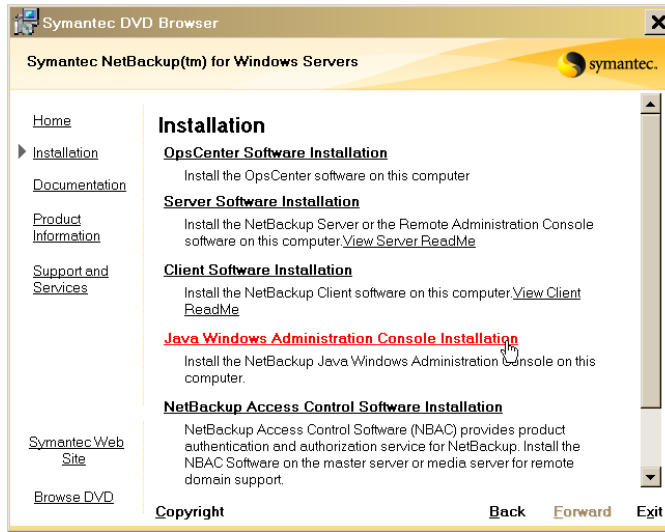
See [“Choosing a remote server to administer”](#) on page 731.

## Using the Java Windows Administration Console

No license is required to install the Java Windows Administration Console. Installing the Java Windows Administration Console installs the administration console only. No NetBackup master server, media server, or client software is installed.

[Figure 23-4](#) shows how to install the Java Windows Administration Console.

**Figure 23-4** Java Windows Administration Console selection on the installation screen



After it is installed, select **Start > Symantec NetBackup > NetBackup-Java Version 7.0** to start the Java Windows Administration Console.

See [“Using the NetBackup-Java administration console”](#) on page 739.

## Running the Administration Console on a NetBackup client

The NetBackup Administration Console on a client is useful to administer a NetBackup server remotely. (No NetBackup server software is installed.)

Run the NetBackup Administration Console on a client under the following conditions:

- On a Windows client if the Java Windows Administration Console is installed.
- On a UNIX client if the client is NetBackup-Java capable.

## Troubleshooting remote server administration

To administer a server from another master server, make sure that the following conditions are met:

- The destination server is operational.



- NetBackup services are running on both hosts.
- The network connection is valid.
- The user has administrative privileges on the destination host.
- The current host is listed in the server list of the destination host.  
 See “[Adding a NetBackup server to a server list](#)” on page 728.  
 The host does not need to be listed if the host is a media server or a client. Or, it does not need to be listed if only media and device management or monitoring is to take place.  
 To ensure that all appropriate NetBackup processes use the new server entry, stop and restart the following processes:
  - The NetBackup Database Manager and NetBackup Request Manager services on the remote server if it is Windows.
  - The NetBackup Database Manager (`bpdbm`) and NetBackup Request Manager (`bpfd`) on the remote server if it is UNIX.
- Authentication is set up correctly, if used.
- For problems changing servers to configure media or devices or monitor devices, verify that the NetBackup Volume Manager is running on that server.
- If you cannot access devices on the remote host, it may be necessary to add a `SERVER` entry to the `vm.conf` file on that host.  
 See the *NetBackup Administrator's Guide, Volume II* for instructions.
- If you cannot start or stop processes or services through the Activity Monitor, verify the following:
  - The remote server is a Windows system. Only on other Windows systems can processes be monitored and controlled.
  - You have the required permissions on the remote server. Windows security must allow access to the user that is running the Activity Monitor.



# Using the NetBackup-Java administration console

This chapter includes the following topics:

- [Using the NetBackup-Java administration console](#)
- [Authorizing NetBackup-Java users](#)
- [Authorization file \(auth.conf\) characteristics](#)
- [Authorizing nonroot users for specific applications](#)
- [Authorizing specific tasks in jbpSA](#)
- [Authorizing NetBackup-Java users on Windows](#)
- [Restricting access to NetBackup-Java applications on Windows](#)
- [Runtime configuration options](#)
- [How to log the command lines that the NetBackup interfaces use](#)
- [How to customize jnbSA and jbpSA with bp.conf entries](#)
- [How to improve NetBackup-Java performance](#)
- [Adjusting time zones in the NetBackup-Java console](#)

## Using the NetBackup-Java administration console

The following topics contain information about the NetBackup-Java Administration Console. On Windows systems, the console is also referred to as the Java Windows Administration Console.

The NetBackup-Java Administration Console is a distributed application that consists of separate system processes:

- The NetBackup Administration Console graphical user interface
  - Available on UNIX by running `jnbSA`
  - Available on Windows by installing the Java Windows Administration Console  
See “Using the Java Windows Administration Console” on page 735.
- The application server (`bpjava` processes)

These processes can be run on two different NetBackup hosts. This distributed application architecture holds true for the UNIX Backup, Archive, and Restore client graphical user interface (`jbpsa`) as well.

The administrator first starts the NetBackup-Java Administration Console interface using one of the following methods:

- Run the `jnbSA` command on UNIX
- Select **Start > Symantec NetBackup > NetBackup-Java Version 7.0** on a Windows system on which the Java Windows Administration Console is installed

Then the administrator logs on to the application server on the host that is specified in the logon dialog box.

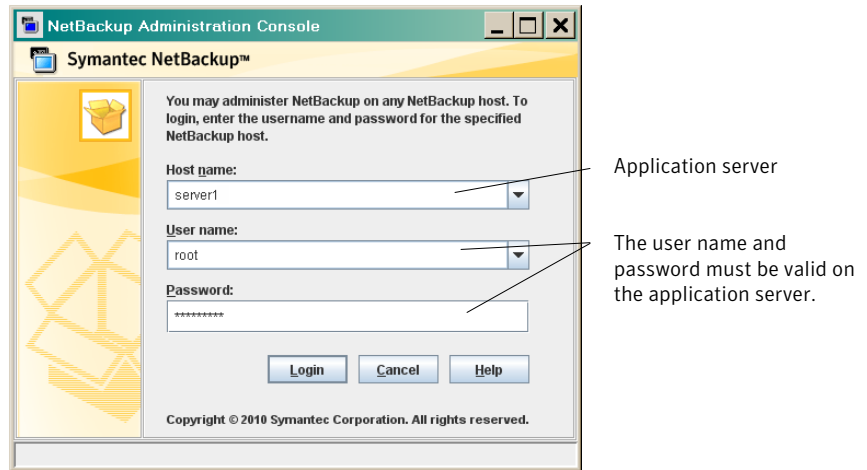
---

**Note:** The host that is specified in the logon dialog box and the system that runs the NetBackup Administration Console must run the same NetBackup version.

---

The application server of the host that is specified in the NetBackup Administration Console logon dialog box authenticates the logon credentials of the user. The credentials are authenticated by using standard UNIX user account data and associated APIs. The logon credentials must be valid on the host that is specified in the logon dialog box.


Figure 24-1 NetBackup logon dialog box



The server that is usually the object of all administrative tasks is the host specified in the NetBackup Administration Console logon dialog box.

An exception is the use of the **File > Change Server** capability in the NetBackup Administration Console. The **Change Server** capability allows administration of a remote server (a server other than the one specified in the NetBackup Administration Console logon dialog box).

Figure 24-2 Change Server dialog box



Regardless of which server is administered, all administrative tasks that are performed in the NetBackup Administration Console make requests of the application server. All tasks are run on the application server host, whether the server is remote or whether the server is specified on the logon dialog box.

However, regardless of which NetBackup authorization method is configured, authorization for tasks in the Administration Console is specific to the server being administered. For example, NetBackup-Java authorization capabilities are in use on Host\_A. Use **Change Server** to change to Host\_B. The permissions are honored as configured in the `auth.conf` on Host\_B.

To administrate from a remote server, the application server host must be included in the server list of the remote server.

See [“Adding a NetBackup server to a server list”](#) on page 728.

See [“Indicating a remote system upon login”](#) on page 733.

## Authorizing NetBackup-Java users

NetBackup offers access control through the Access Management utility in the NetBackup Administration Console.

Instructions on how to install the necessary components to use Access Management are available in the *NetBackup Security and Encryption Guide*.

If NetBackup Access Control is not configured, you may still authorize users of the NetBackup-Java administration console for specific applications. NetBackup Access Control always takes precedence over the capabilities authorization of NetBackup-Java.

If a user is not an authorized administrator by NetBackup Access Control, the actions that the user can perform in the Backup, Archive, and Restore application are limited. The user can perform the actions that are defined in the `auth.conf` file on the host that is specified in the NetBackup-Java logon dialog box. NetBackup-Java users must log on to the NetBackup-Java application server that is on the NetBackup host where they want to perform administrator or user operations.

The `/usr/opensv/java/auth.conf` file contains the authorization data for accessing NetBackup-Java applications. This file exists only on NetBackup-Java capable machines where the NetBackup-Java interface software is installed.

The default `auth.conf` file provides the following authorizations:

On NetBackup servers      Administration capabilities for the root user and user backup and restore capabilities for all other users.

On NetBackup clients      User backup and restore capabilities for all users.

On all other UNIX NetBackup systems, the file does not exist but the NetBackup-Java application server provides the same default authorization. To change these defaults on other UNIX systems, create the `/usr/opensv/java/auth.conf` file.

To perform remote administration or user operations with `jbpSA`, a user must have valid accounts on the NetBackup UNIX server or client machine.

Nonroot or non-administrator users can be authorized to administer Windows NetBackup servers remotely from the NetBackup-Java Console. Do so by setting up authorization in the `auth.conf` file on the Windows server.

The `auth.conf` file must contain entries for the UNIX user names that are used in the logon dialog box of the NetBackup-Java Console. The `auth.conf` file must reside in `install_path\VERITAS\java` on each Windows server you want to provide nonroot administration capability. Without an `auth.conf` file, the user has the same privileges on the remote server as on the server that is specified in the logon screen. User privileges are the same if `auth.conf` does not contain an entry for the user name even though host authorization between the two is configured. (`SERVER` entries in the configuration of each.)

## Authorization file (auth.conf) characteristics

The `/usr/opensv/java/auth.conf` file is installed on all NetBackup-Java capable hosts and contains only the following entries:

```
root ADMIN=ALL JBP=ALL
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

The first field of each entry is the user name that is granted access to the rights that the entry specifies. In the released version, the first field allows root users to use all of the NetBackup-Java applications.

An asterisk in the first field indicates that any user name is accepted and the user is allowed to use the applications as specified. If the `auth.conf` file exists, it must have an entry for each user. Or, the `auth.conf` file must have an entry that contains an asterisk (\*) in the user name field; users without entries cannot access any NetBackup-Java applications. Any entries that designate specific user names must precede a line that contains an asterisk in the user name field.

---

**Note:** The asterisk specification cannot be used to authorize all users for any administrator capabilities. Each user must be authorized by using individual entries in the `auth.conf` file.

---

To deny all capabilities to a specific user, add a line that indicates the user before a line that starts with an asterisk.

For example:

```
mydomain\ray ADMIN= JBP=
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

The remaining fields specify the access rights.

- The `ADMIN` keyword specifies the applications that the user can access. `ADMIN=ALL` allows access to all NetBackup-Java applications and the related administrator-related capabilities.  
See “[Authorizing nonroot users for specific applications](#)” on page 745.
- The `JBP` keyword specifies what the user can do with the Backup, Archive, and Restore client application (`jbpSA`). `JBP=ALL` allows access to all Backup, Archive, and Restore capabilities, including those for administration.  
See “[Authorizing specific tasks in jbpSA](#)” on page 746.
- An asterisk in the first field indicates that any user name is accepted and the user is allowed to use the applications as specified. The second line of the released version contains an asterisk in the first field. The asterisk means that NetBackup-Java validates any user name for access to the Backup, Archive, and Restore client application `jbpSA`. `JBP=ENDUSER+BU+ARC` allows users to back up, archive, and restore files only.

The user name and password that is entered in the logon screen must be valid on the machine that is specified in the host field. (True for starting the NetBackup-Java administration console or the Backup, Archive, and Restore application (`jbpSA`).) The NetBackup-Java application server authenticates the user name and password by using the system password file data for the specified machine. The password must be the same password that was used upon logon at that machine.

For example, assume you log on with the following information:

```
username = joe
password = access
```

Here you must use the same user name and password to log into NetBackup-Java.

---

**Note:** The NetBackup-Java logon box accepts passwords greater than eight characters. However, only the first eight are significant upon logon to a NetBackup-Java application server on a UNIX system.

---

You can log on to the NetBackup-Java application server under a different user name than the name used to log on to the operating system. For example, if you log on to the operating system with a user name of `joe`, you can subsequently log on to `jnbSA` as `root`.

Upon exit, some application state information is automatically saved in the directory of `joe` `$HOME/.java/.userPrefs/vrts` directory. (For example, table column order.) The information is restored the next time you log on to the operating system under account `joe` and initiate the NetBackup-Java application.



This logon method of is useful if there is more than one administrator because it saves the state information for each administrator.

---

**Note:** NetBackup-Java creates a user's `$HOME/.java/.userPrefs/vrts` directory the first time an application is exited. Only NetBackup-Java applications use the `.java/.userPrefs/vrts` directory.

---

If the user name is not valid as determined by the contents of the `auth.conf` file, an error message appears. All applications are inaccessible to the user:

```
No authorization entry exists in the auth.conf file for username
name_specified_in_login_dialog. None of the NB-Java applications are
available to you.
```

To summarize, the following types of entries are contained in the `auth.conf` file, as follows:

- The defaults allow anyone with any valid user name to use the Backup, Archive, and Restore client application (`jbpsA`). Only root users can access the administrator applications and the administrator capabilities in `jbpsA`.
- Specify entries for valid user names.

---

**Note:** The validated user name is the account the user can back up, archive or restore files from or to. The Backup, Archive, and Restore application (`jbpsA`) relies on system file permissions of when to browse directories and files to back up or restore.

---

## Authorizing nonroot users for specific applications

Nonroot users can be authorized for a subset of the NetBackup-Java administrator applications.

To authorize users for a subset of the NetBackup-Java administrator applications, use the following identifiers for the `ADMIN` keyword in the `auth.conf` file:

|     |                                                                                                                  |
|-----|------------------------------------------------------------------------------------------------------------------|
| ALL | Indicates that the user has administrative privileges for all of the applications that are listed in this table. |
| AM  | Activity Monitor                                                                                                 |
| BMR | Bare Metal Restore                                                                                               |
| BPM | Backup Policy Management                                                                                         |

|            |                              |
|------------|------------------------------|
| BAR or JBP | Backup, Archive, and Restore |
| CAT        | Catalog                      |
| DM         | Device Monitor               |
| HPD        | Host Properties              |
| MM         | Media Management             |
| REP        | Reports                      |
| SUM        | Storage Unit Management      |
| VLT        | Vault Management             |

For example, to give a user (`user1`) access only to the Device Monitor and Activity Monitor, add the following entry to the `auth.conf` file:

```
user1 ADMIN=DM+AM
```

In order for a nonroot user to modify the files that the NetBackup-Java Administration Console uses, run the `nonroot_admin_nbjava` script. The script changes permissions on the following files:

```
/usr/opensv/java/auth.conf
/usr/opensv/java/Debug.properties
/usr/opensv/java/nbj.conf
```

---

**Note:** `nonroot_admin_nbjava` is located in  
`/usr/opensv/java/nonroot_admin_nbjava.`

---

## Authorizing specific tasks in jbpSA

The Backup, Archive, and Restore interface can be configured to allow only a user to perform certain tasks. Not all tasks can be performed successfully without some additional configuration.

The following require additional configuration and are documented elsewhere:

- Redirected restores.  
See “[Server-directed restores](#)” on page 771.  
See “[Client-redirected restores](#)” on page 772.
- User backups or archives require a policy schedule of these types and the task to be submitted within the time window of the schedule.

To authorize users for a subset of Backup, Archive, and Restore capabilities, use the following identifiers for the `JBP` keyword in the `auth.conf` file:

- `ENDUSER`  
Allows the users to perform restore tasks from true image, archive, or regular backups plus redirected restores.
- `BU`  
Allows the users to perform backup tasks.
- `ARC`  
Allows the users to perform archive tasks. The capability to perform backups (`BU`) is required to allow archive tasks.
- `RAWPART`  
Allows the users to perform raw partition restores.
- `ALL`  
Allows the users to perform all actions, including server-directed restores. (Restores to a client that is different from the client that is logged into.) Server-directed restores can only be performed from a NetBackup master server.

For example, to allow a user (`user1`) to restore but not backup up or archive files:

```
user1 ADMIN=JBP JBP=ENDUSER
```

## Authorizing NetBackup-Java users on Windows

To use the Java Windows Administration Console, first log on to the NetBackup-Java application server. The application server is on the NetBackup host where you want to perform NetBackup administration or user operations.

To log on to the application server, log on to the dialog box that appears when the console is started. Provide a valid user name and password for the system that is specified in the **Host name** field of the log in dialog box.

[Figure 24-1](#) shows the log in dialog box.

The user name for Windows must be of the form: *domainname\username*

*domainname* specifies the domain of the NetBackup host. The domain is not required if the NetBackup host is not a member of a domain.

The NetBackup-Java application server authenticates the user name and password by using standard Windows authentication capabilities for the specified computer.

If NetBackup Access Control is not configured for the users, by default the NetBackup-Java application server provides authorization data. The authorization

data allows all users that are members of the administrator group for the host's domain to use all the NetBackup-Java applications. Other users are allowed to access only Backup, Archive, and Restore.

To restrict access to NetBackup-Java or some of its applications, create a `nbservice_install_path\java\auth.conf` authorization file.

See [“Using the NetBackup-Java administration console”](#) on page 739.

## Restricting access to NetBackup-Java applications on Windows

To restrict access to one or more of the NetBackup-Java applications, create the following file on the Windows system:

```
nbservice_install_path\java\auth.conf
```

Add an entry in `auth.conf` for each user that accesses NetBackup-Java applications. The existence of this file, along with the entries it contains, prohibits unlisted users from accessing NetBackup-Java applications on the Windows system. The following is a sample `auth.conf` file on a Windows system:

```
mydomain\Administrator ADMIN=ALL JBP=ALL
mydomain\joe ADMIN=ALL JBP=ALL
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

## Runtime configuration options

On UNIX systems, file `/usr/openv/java/nbj.conf` contains configuration options for the NetBackup-Java Administration Console. Enter one option per line, following the same syntax rules as exist for the `bp.conf` file.

On Windows systems, the analogous file containing configuration options for the Java Windows Administration Console is `nbservice_install_path\java\setconf.bat`

`nbj.conf` and `setconf.bat` contain commands for each of the configuration options that are described in the following topics. To make changes, change the value after the equal sign in the relevant set command.

### BPJAVA\_PORT, VNETD\_PORT

The following ports are the configured ports for the `bpjava-msvc` and `vnetd` daemon processes. These ports are registered with the Internet Assigned Numbers Authority (IANA).

**Table 24-1** Port numbers

| Port        | Process and Registered Default Port Number |
|-------------|--------------------------------------------|
| bpjava-msvc | BPJAVA_PORT=13722                          |
| vnetd       | VNETD_PORT=13724                           |

Symantec recommends that these ports are not changed. If changes are necessary, make the change on all NetBackup hosts in the relevant NetBackup cluster.

See the *NetBackup Installation Guide*.

The value must be set in the corresponding `nbj.conf` (UNIX) or `setconf.bat` (Windows) configuration option.

## FIREWALL\_IN

The `FIREWALL_IN` configuration option provides a method to use a Java Administration Console that is outside of a trusted network to administer the NetBackup master servers that are within a trusted network.

This option uses the following format.

On UNIX:

```
FIREWALL_IN= HOST1:PORT1=HOST2:PORT2[; ...;HOSTn:PORTn=HOSTm:PORTm]
```

On Windows:

```
SET FIREWALL_IN=
HOST1:PORT1=HOST2:PORT2;IP_ADDR1:PORT3=IP_ADDR2:PORT4
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

Where *HOST* is a host name or an IP address.

This configuration option provides a way to allow administrators to bypass the firewall by using one of the following methods:

- Enter the port number of the `bpjava` service in the trusted internal network. Then, map the private interface where the `bpjava` service runs to a public interface that can be reached from outside the firewall.
- Set up a Secure Shell (SSH) tunnel from the local host to the system inside the firewall.

In the following example:

- Master server `NBUMaster.symc.com` is in a trusted network, behind a firewall.
- The IP address of `NBUMaster.symc.com` is `10.221.12.55`.

- The NetBackup Java Administration Console is installed on localhost.
- SSH tunnels exist from localhost to NBUMaster.symc.com as follows:

```
bpjava-msvc port (default 13722) localhost:port1
vnetd port (default 13724) localhost:port2
pbx port (default 1556) localhost:12345
```

Where **localhost** is the host name and port1 is the IP port.

To make relevant changes for connections to `bpjava-msvc` and `vnetd`, see the following topic:

See “[BPJAVA\\_PORT, VNETD\\_PORT](#)” on page 748.

On UNIX systems, add the following line to the `nbj.conf` file:

```
FIREWALL_IN=NBUMaster.symc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
```

The entry indicates the following:

- The connection to NBUMaster.symc.com:1556 is to be redirected to localhost:12345.
- The connection to 10.221.12.55:1556 is to be redirected to localhost:12345.

On Windows systems, use `setconf.bat` to add the option:

```
SET FIREWALL_IN=
NBUMaster.symc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

---

**Note:** The same options are used if NBUMaster.symc.com has a public interface (NBUMasterpub.symc.com) that can be reached from the Internet. In this case, the administrator replaces localhost with NBUMasterPub.symc.com.

---

## FORCE\_IPADDR\_LOOKUP

The `FORCE_IPADDR_LOOKUP` configuration option specifies whether NetBackup performs an IP address lookup to determine if two host name strings are indeed the same host. This option uses the following format:

```
FORCE_IPADDR_LOOKUP = [0 | 1]
```

Where:

0 = Indicates that no IP address lookup is performed to determine if two host name strings are indeed the same host. They are considered to be the same host if the host name strings compare equally. Or, if a short name compares equally to the short name of a partially or fully qualified host name.

1 = Indicates that an IP address lookup is performed if the two host name strings do not match. The lookup determines if they have the same host. The default is to perform an IP address lookup if necessary to resolve the comparison. The IP address lookup is not performed if the host name strings compare equally.

---

**Note:** Use a value of 1 for this option if you have the same host name in two different domains. For example, `eagle.abc.xyz` and `eagle.def.xyz` or by using host name aliases.

---

Many places in the NetBackup Administration Console compare host names to determine if the two are the same host. For example, the **File > Change Server** command.

The IP address lookup can consume time and result in slower response time. However, accurate comparisons are important.

No IP address lookup is necessary if the host name is specified consistently in the NetBackup Administration Console logon dialog box. It must match how the host names are configured in NetBackup. Host names are identified in the server list that is found in the Servers host properties. On UNIX systems, the host names also appear in the `bp.conf` file.

Using host names `eagle` and `hawk`, the following describes how this option works:

- `FORCE_IPADDR_LOOKUP = 0`

Comparisons of the following result in no IP address lookup. The hosts are considered to be the same host.

```
eagle and eagle
eagle.abc.def and eagle.abc.def
eagle.abc and eagle.abc.def
eagle and eagle.abc.def
eagle and eagle.anything
```

The hosts are considered to be different for any comparisons of short, partially, or fully qualified host names of `eagle` and `hawk` regardless of aliases.

- `FORCE_IPADDR_LOOKUP = 1`

Comparisons of the following result in no IP address lookup. The hosts are considered to be the same host.

```
eagle and eagle
eagle.abc and eagle.abc
eagle.abc.def and eagle.abc.def
```

In addition to all comparisons of eagle and hawk, the following result in an IP address lookup. The comparison determines if the hosts are indeed the same host.

```
eagle.abc and eagle.abc.def
eagle and eagle.abc.def
eagle and eagle.anything
```

## INITIAL\_MEMORY, MAX\_MEMORY

Both `INITIAL_MEMORY` and `MAX_MEMORY` allow configuration of memory usage for the Java Virtual Machine (JVM).

Symantec recommends that the NetBackup-Java Administration Console, the Java Windows Administration Console, or the NetBackup, Archive, and Restore user interface run on a system that contains at least 1 gigabyte of physical memory. Make sure that 256 megabytes of memory are available to the application.

`INITIAL_MEMORY` specifies how much memory is allocated for the heap when the JVM starts. The value probably does not require changing. The default is sufficient for quickest initialization of `jnbSA`, the Java Windows Administration Console, or `jbpSA` on a system with the recommended amount of memory.

On UNIX systems, the initial memory allocation can also be specified as part of the `jnbSA` or `jbpSA` command. For example:

```
jnbSA -ms 36M
```

Default = 36M (megabytes).

`MAX_MEMORY` specifies the maximum heap size that the JVM uses for dynamically allocated objects and arrays. If the amount of data is large, consider specifying the maximum heap size. For example, a large number of jobs in the Activity Monitor.

On UNIX systems, the maximum memory allocation can also be specified as part of the `jnbSA` or `jbpSA` command. For example:

```
jnbSA -mx 512M
```

Default = 256M (megabytes).



## MEM\_USE\_WARNING

The `MEM_USE_WARNING` configuration option specifies the percent of memory used compared to `MAX_MEMORY`, at which time a warning dialog box appears to the user. Default = 80%. This option uses the following format:

```
MEM_USE_WARNING=80
```

## NBJAVA\_CLIENT\_PORT\_WINDOW

The `NBJAVA_CLIENT_PORT_WINDOW` configuration option specifies the range of non-reserved ports on this computer to use for connecting to the NetBackup-Java application server. It also specifies the range of ports to use to connect to the `bpjobjd` daemon from the NetBackup-Java Administration Console's Activity Monitor.

This option uses the following format:

```
NBJAVA_CLIENT_PORT_WINDOW = n m
```

Where:

- *n* indicates the first in a range of non-reserved ports that are used for connecting to the `bpjava` processes on the NetBackup-Java application server. It also specifies the range of ports to use to connect to the `bpjobjd` daemon or Windows service from the Activity Monitor of the Java Windows Administration Console.  
If *n* is set to 0, the operating system determines the non-reserved port to use (default).
- *m* indicates the last in a range of non-reserved ports that are used for connecting to the NetBackup-Java Administration Console or the Java Windows Administration Console.  
If *n* and *m* are set to 0, the operating system determines the non-reserved port to use (default).

The minimum acceptable range for each user is 120. Each additional concurrent user requires an additional 120. For example, the entry for three concurrent users might look as follows:

```
NBJAVA_CLIENT_PORT_WINDOW = 5000 5360
```

If the range is not set wide enough, `jnbSA` exits with an error message that states an invalid value has occurred during initialization.

---

**Note:** Performance is reduced with the use of `NBJAVA_CLIENT_PORT_WINDOW`.

---

## NBJAVA\_CONNECT\_OPTION

The `NBJAVA_CONNECT_OPTION` configuration option specifies how the NetBackup-Java application server is connected to. It may be done using the `vnetd` daemon (`VNETD_PORT`) or directly using the application server's port (`BPJAVA_PORT`). The option also specifies the callback method that the server or the client uses when it communicates with the NetBackup-Java consoles (`jnbSA`, `jbpSA`).

The default for `NBJAVA_CONNECT_OPTION` requires only that the `vnetd` port is accessible through any firewall.

```
NBJAVA_CONNECT_OPTION = [0 | 1]
```

Where:

0 = Indicates a direct connection to the application server and the traditional callback method.

1 = Indicates a connection to the application server using `vnetd` and the no callback method (default).

## NBJAVA\_CORBA\_DEFAULT\_TIMEOUT

The `NBJAVA_CORBA_DEFAULT_TIMEOUT` configuration entry specifies the default timeout that is used for most CORBA operations that the Java Administration Console performs.

This option is present by default and uses the following format:

```
NBJAVA_CORBA_DEFAULT_TIMEOUT=60
```

The default is 60 seconds.

## NBJAVA\_CORBA\_LONG\_TIMEOUT

The `NBJAVA_CORBA_LONG_TIMEOUT` configuration entry specifies the timeout value that the Java Administration Console uses in the following areas:

- Device Configuration Wizard
- Disk Pool Configuration Wizard
- Disk Pool Inventory

This option is present by default and uses the following format:

```
NBJAVA_CORBA_LONG_TIMEOUT=1800
```

The default is 1800 seconds.

## How to log the command lines that the NetBackup interfaces use

At times it may be helpful to see which command lines the NetBackup-Java Administration Console or the NetBackup, Archive, and Restore user interface uses. Use option `-lc` to log to a log file the command lines that `jnbSA` or `jbpSA` uses. No value is necessary. For example:

```
/usr/opensv/java/jbpSA -lc
```

---

**Note:** `jnbSA` and `jbpSA` do not always use the command lines to retrieve or update data. The interfaces have protocols that instruct the application server to perform tasks using NetBackup and Media Manager APIs.

---

## How to customize `jnbSA` and `jbpSA` with `bp.conf` entries

The `INITIAL_BROWSE_SEARCH_LIMIT` and `KEEP_LOGS_DAYS` options in the `/usr/opensv/netbackup/bp.conf` file allow the administrator and users to customize the following aspects of `jbpSA` operation, as follows:

- `INITIAL_BROWSE_SEARCH_LIMIT` limits the start date of the search for restores and can improve performance when large numbers of backups are done.
- `KEEP_LOGS_DAYS` specifies how long job and progress log files are kept that the NetBackup-Java Backup, Archive, and Restore application (`jbpSA`) generates. The files are written into the following directories:

```
/usr/opensv/netbackup/logs/user_ops/_username_/jobs
```

```
/usr/opensv/netbackup/logs/user_ops/_username_/logs
```

A directory exists for each user that uses the NetBackup-Java applications. The default is three days.

This option also controls how long the NetBackup-Java GUI log files are kept in `/usr/opensv/netbackup/logs/user_ops/nbjlogs`.

## How to improve NetBackup-Java performance

The most important factor to consider concerning performance issues while using the following interfaces is the platform on which the console is running:

- NetBackup-Java Administration Console
- Java Windows Administration Console
- NetBackup Backup, Archive, and Restore user interface

Regardless of the platform, you can run the administration console from one of the following locations:

- Run it locally on a desktop host (on supported Windows and UNIX platforms)
- Run it remotely and display it back to a desktop host (from supported UNIX platforms)

To provide the best performance, the recommended method for using these consoles is to run the consoles locally on a desktop host. When the consoles are run locally, they do not exhibit the font and the display issues that can be present in some remote display-back configurations.

## Running the Java console locally on a UNIX platform

On supported UNIX platforms, the console is run locally if `j_nbSA` or `j_bpSA` is entered on the same host on which the console is appears. That is, your display environment variable is set to the host on which the `j_nbSA` or `j_bpSA` commands were entered.

Improvements in Java technology have made remote X-display back potentially viable on some platforms. However, problems continue with certain controls in the consoles. For example, incorrect combo box operations, sluggish scrolling, and display problems in tables with many rows. More serious issues have also occurred. Consoles can abort and hang because of a Java Virtual Machine (JVM) failure when run in this mode on some platforms. These JVM failures are most often seen on the AIX platform. Therefore, Symantec cannot recommend running the consoles in a remote X-display back configuration.

## Running the console locally on a Windows platform

On Windows platforms, select **Start > Symantec NetBackup > NetBackup-Java Version 7.0** to start the Java Windows Administration Console. The **Start** menu item appears if you install the optional Java Windows Administration Console available on the main NetBackup for Windows installation screen.

See [“Using the Java Windows Administration Console”](#) on page 735.

## How to run a console locally and administer a remote server

The NetBackup Administration Console and the Backup, Archive, and Restore user console are distributed applications. Both applications consist of two major and separate system processes that can run on different machines. For example: the NetBackup Administration Console on one machine and the console's application server – `bpjava` processes on another machine.

The NetBackup Administration Console does not need to run on a NetBackup server host. However, the application server must run on this host in order for you to be able to administer NetBackup.

Although the NetBackup-Java Administration Console does not run on all NetBackup-supported platforms, the application server for the console does run on all supported platforms. The distributed application architecture enables direct administration of all NetBackup platforms, even though the consoles themselves run only on a subset of the NetBackup-supported platforms.

To log into the NetBackup-Java Administration Console, specify a host name. The host name is the machine where the application server (`bpjava`) runs. (For example, a NetBackup master server.) All requests or updates that are initiated in the console are sent to its application server that runs on this host.

## How to enhance console performance

Performance of the NetBackup-Java applications depends on the environment where the applications are running, including available resources and network throughput. The NetBackup-Java default configuration, specifically the `INITIAL_MEMORY` and `MAX_MEMORY` configuration options, assumes sufficient memory resources on the machine where the console is running. For example, where the `jnbSA` command is run or the NetBackup-Java Administration Console is started.

Following are guidelines for improving performance:

- Consider the network communication speed and the amount of data being transferred.
- Consider the amount of work being performed on the relevant machines. Run NetBackup-Java on a machine that has a low level of activity. For example, there can be large differences in response time when other memory-intensive applications are running on the machine. (For example, Web browsers.) Multiple instances of NetBackup-Java on the same machine have the same effect.
- Run NetBackup-Java on a 1-gigabyte machine that has at least 256 MB of RAM available to the application. In some instances, the application does not initiate due to insufficient memory. A number of messages identify these failures in

the xterm window where the `jnbSA` command was run. Or, the messages appear in the application log file. Possible messages include the following:

```
Error occurred during initialization of VM
Could not reserve enough space for object heap
Out of Memory
```

- Consider the amount of physical memory on the relevant machines. Possibly add memory on the host being administered (the console's application server host).
- Consider increasing the swap space to relevant machines:
  - The console host (the host where the console is started)
  - The host being administeredIncrease the amount of swap space available to the system where you are running the applications can increase performance. Especially if there is a great deal of other activity on the machine. More swap space can alleviate hangs or other problems that relate to insufficient memory for the applications.
- Consider additional or faster CPUs to relevant machines:
  - The console host (the host where the console is started)
  - The host being administered
- To save startup time, allow NetBackup-Java to run rather than exit and restart. Startup of the Java Virtual Machine can take longer than other applications.
- Consider limiting the amount of NetBackup data that is retained for long periods of time to only that which is necessary. For example, do not retain successfully completed jobs for more than a few hours.

## Determining better performance when run locally or using remote display back

Performance depends on the following:

- The speed of the network
- The console and the application server machine resources
- The workloads on the console
- The application server hosts
- The amount of NetBackup data (Data is the number of jobs in the Activity Monitor or number of NetBackup policies.)

The console may perform better if started on the console's application server host, then displayed back to the desktop host. However, Symantec is not aware of a situation where that configuration produces better console performance. As previously mentioned, the configuration is not recommended due to problems unrelated to performance issues.

Consider the following scenarios to determine what would provide the best performance for your configuration.

## Scenario 1

Assume no deficiency in either the console host's resources or the application server host's resources. Assume that the amount of NetBackup configuration data being transferred to the console host far exceeds the X-Windows pixel display data. That is, the actual console screen being sent from the remote host.

Unfortunately, the only way to determine the viability of this situation is to try it. Network capabilities and the proximity of the two hosts influences each NetBackup configuration.

## Scenario 2

Assume that the available resources of the application server host far exceed that of the console host.

Assume that the console host has a very limited CPU and memory as compared to the NetBackup master server being administered. (The console host is the machine on which the console is started.) If the console is run on the master server and displayed back to the desktop host, performance may be enhanced.

If the desktop host is a Windows machine, X-terminal emulation or remote display tools such as Exceed and VNC are required.

These scenarios address the performance aspect of using the NetBackup-Java console. There may be other reasons that require you to display back remotely to your desktop, however, it is not recommended. Review the Release Notes for additional issues of relevance to the NetBackup-Java Administration Console and Backup, Archive, and Restore client console.

[Table 24-2](#) shows the files that contain configuration entries.

**Table 24-2** Files containing configuration entries

| File                                    | Description            |
|-----------------------------------------|------------------------|
| <code>/usr/opensv/java/auth.conf</code> | Authorization options. |

**Table 24-2** Files containing configuration entries (*continued*)

| File                                       | Description                                            |
|--------------------------------------------|--------------------------------------------------------|
| <code>/usr/opensv/netbackup/bp.conf</code> | Configuration options (server and client).             |
| <code>/usr/opensv/java/nbj.conf</code>     | Configuration options for the NetBackup-Java Console   |
| <code>/usr/opensv/volmgr/vm.conf</code>    | Configuration options for media and device management. |
| <code>\$HOME/bp.conf</code>                | Configuration options for user (on client).            |

## Adjusting time zones in the NetBackup-Java console

Sites in a geographically dispersed NetBackup configuration may need to adjust the time zone in the NetBackup-Java Console for administration of remote NetBackup hosts. (In this context, a remote NetBackup host may either be the host that is specified in the console logon dialog box or one referenced by the **File > Change Server** capability in the console.)

The default time zone for the console is that of the host on which the console is started, not the host that is specified (if different) in the console logon dialog box.

For backup, restore, or archive operations from within the NetBackup-Java Console (`jnbSA`) or the Backup, Archive, and Restore application when run on a client (`jbpSA`), set the time zone relative to the NetBackup server from which the client restores files.

Set the time zone in separate instances of the NetBackup-Java Console when servers in different time zones are administered.

For example, open a NetBackup-Java Console to set the time zone for the local server in the Central time zone. To set the time zone for a server in the Pacific time zone as well, open another NetBackup-Java Console.

Do not open a new window in the first NetBackup-Java Console. Change servers (**File > Change Server**), and then set the time zone for the Pacific time zone server. Doing so changes the time zone for the Central time zone server as well.

### Adjusting the time zone

Use the following procedure to adjust the time zone or to use daylight savings time.



### To adjust the time zone

- 1 In the NetBackup Administration Console, or in the Backup, Archive, and Restore client interface, select **File > Adjust Application Time Zone**.
- 2 Select the **Standard** tab.
- 3 Clear the **Use custom time zone** check box.
- 4 Select the time zone.
- 5 For daylight savings time, select **Use daylight savings time**.
- 6 To have administrative capabilities and to apply the settings to the current session and all future sessions, select **Save as default time zone**.
- 7 Click **OK**.

## Configuring a custom time zone

Use the following procedure to configure a custom time zone.

### To configure a custom time zone

- 1 In the NetBackup Administration Console, or in the Backup, Archive, and Restore client interface, select **File > Adjust Application Time Zone**.
- 2 Select the **Use custom time zone** check box.
- 3 Select the Custom tab.
- 4 Select the time zone on which to base the Backup, Archive, and Restore interface time.
- 5 For the **Offset from Greenwich Mean Time** setting, adjust the time to reflect how many hours and minutes the server's time zone is either behind or ahead of Greenwich Mean Time.
- 6 Select the **Use daylight savings time** checkbox.
- 7 In the Daylight savings time start section of the dialog, do the following:
  - To begin DST on a specific date, select **Absolute date** and indicate the month and day.  
To begin DST on April 5, set as follows:
  - To begin DST on the first occurrence of a day in a month, select **First day of week in month**. Indicate the day of the week and the month.  
To begin DST on the first Monday in April, set as follows:
  - To begin DST on the first occurrence of a day in a month and after a specific date, select **First day of week in month after date**. Indicate the day of the week and the month and day.

To begin DST on the first Monday after April 5, set as follows:

- To begin DST on the last occurrence of a day in a month, select **Last day of week in month**. Indicate the day of the week and the month.

To begin DST on the last Thursday in April:

- To begin DST on the last occurrence of a day in a month and before a specific date, select **Last day of week in month before date**. Indicate the day of the week and the month and day.

To begin DST before April 30, set as follows:

- 8 Indicate when DST should end by using one of the methods in the previous step.
- 9 To have administrative capabilities and apply the settings to the current session and all future sessions, select **Save as default time zone**.
- 10 Click **OK**.

# Alternate server restores

This chapter includes the following topics:

- [About alternate server restores](#)
- [Supported configurations for alternate server restores](#)
- [Performing alternate server restores](#)

## About alternate server restores

This topic explains how to restore files by using a NetBackup server other than the one that was used to write the backup. This type of restore operation is called an alternate server restore or server independent restore. It allows easier access to data for restores in master and media server clusters and provides better failover and disaster recovery capabilities.

The architecture of NetBackup allows storage devices to be located on multiple servers (either separate storage devices or a shared robot). The NetBackup image catalog on the master server contains an entry that defines the server (master or media server) to which each backup was written. Information specific to the backup media is contained within the master server image catalog (in the attribute file for each backup). The information is also contained in the Enterprise Media Manager (EMM) database, generally located on the master server.

To restore data through a device on another server is more involved than other restores. Use the methods that are described in this topic to restore the backups. Although the methods do not require you to expire and import backup images, in some instances it is useful.

The information in this topic is also pertinent in the case of restoring from a backup copy. If you created multiple copies of a backup, it is possible to restore from a specific backup copy other than the primary copy. To do so, use the `bprestore` command.

More information is available in *NetBackup Commands*.


See “[Expiring and importing media for alternate server restores](#)” on page 770.

## Supported configurations for alternate server restores

All of the methods for alternate server restores require that the server that is used for the restore be in the same cluster as the server that performed the original backup. It must also share the same Enterprise Media Manager database.

[Figure 25-1](#) and [Figure 25-2](#) show configurations where NetBackup supports alternate server restores. All methods require that the server that is used for the restore be in the same cluster as the server that performed the original backup. The server must also share the same Enterprise Media Manager database.

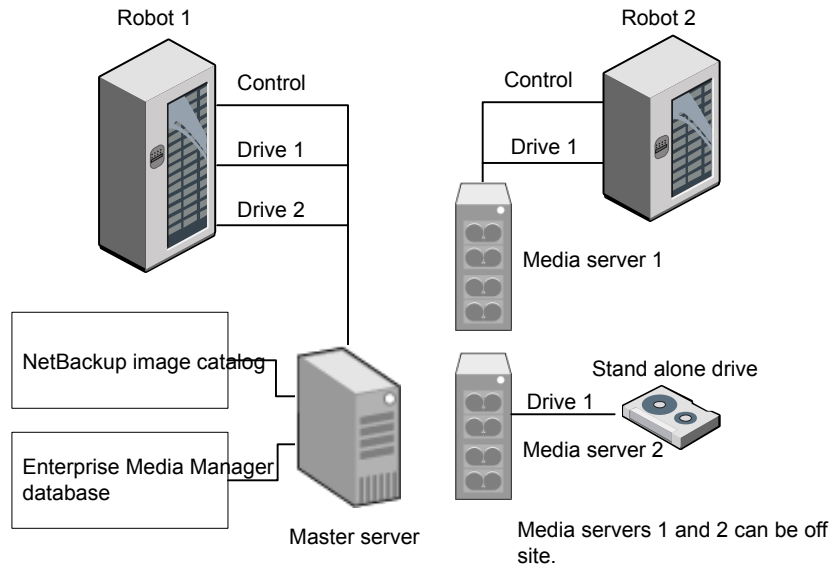
**Figure 25-1** NetBackup servers that share robotic peripherals



Assume the following in [Figure 25-1](#):

- A single, shared Enterprise Media Manager database exists on the NetBackup master server.
- The NetBackup master server is available at time of restore.
- Robotic control is on a NetBackup server that is available at the time of the restore.

**Figure 25-2** NetBackup servers with separate non-shared peripherals



Assume the following in See [Figure 25-2](#) on page 765.:

- The media is made physically accessible through an available NetBackup server. The Enterprise Media Manager database is updated to reflect this move.
- A single, shared Enterprise Media Manager database exists on the NetBackup master server.
- The NetBackup master server is available at time of restore
- Robotic control (if applicable) is on a NetBackup server that is available at the time of the restore.

## Performing alternate server restores

The method that NetBackup administrators can use to perform alternate server restores depends on the configuration and the situation. The method can include one or more of the following:

- See [“Modifying the NetBackup catalogs”](#) on page 766.
- See [“Overriding the original server for restores”](#) on page 767.
- See [“Enabling automatic failover to an alternate server”](#) on page 769.

## Modifying the NetBackup catalogs

This method changes the contents of NetBackup catalogs and thus requires administrator intervention. Use this method only when the server reassignment is permanent.

Some examples of when to use this method are as follows:

- Media is moved to an off-site location, where a media server exists.
- A robot was moved from one server to another.
- Two (or more) servers share a robot, each with connected drives and one of the servers is to be disconnected or replaced.
- Two (or more) servers each have their own robots. One of the server's robots has run out of media capacity for future backups, while several empty slots exist on another server's robot.

The actual steps that are used vary depending on whether the original server is still available.

### Modifying NetBackup catalogs when the server that wrote the media is available

Use the following procedure to modify catalogs when the server that wrote the media is available.

#### To modify NetBackup catalogs when the server that wrote the media is available

- 1 If necessary, physically move the media.
- 2 Update the Enterprise Media Manager database by using move volume options in the Media Manager administration utilities.
- 3 Update the NetBackup image catalog on the master server.
- 4 Update the NetBackup media catalogs on both the original NetBackup server (*oldserver*) and the destination NetBackup server (*newserver*).

Use the following command, which can be run from any one of the NetBackup servers.

Enter the `admincmd` command on one line:

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpmedia -movedb -m media_id -newserver hostname
-oldserver hostname
```

- As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -m media_id
-newserver hostname -oldserver hostname
```

## Modifying NetBackup catalogs when the server that wrote the media is unavailable

Use the following procedure to modify catalogs when the server that wrote the media is unavailable.

### To modify NetBackup catalogs when the server that wrote the media is unavailable

- 1 If necessary, physically move the media.
- 2 Update the Enterprise Media Manager database by using the move volume options in the **Media and Device Management** window.
- 3 Update only the NetBackup image catalog on the master server.

Use the following commands from the NetBackup master server.

Enter the `admincmd` command on one line:

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpimage -id media_id -newserver hostname
-oldserver hostname
```

- As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpimage.exe -id media_id -newserver hostname
-oldserver hostname
```

## Overriding the original server for restores

NetBackup allows the administrator to force restores to a specific server, regardless of where the files were backed up. For example, if files were backed up on server A, a restore request can be forced to use server B.

Examples of when to use this method are as follows:

- Two (or more) servers share a robot, each with connected drives. A restore is requested while one of the servers is either temporarily unavailable or is busy doing backups.

- A server was removed from the NetBackup configuration, and is no longer available.

Use the following procedure to override the original server for restores.

#### To override the original server for restores

- 1 In the NetBackup Administration console, open the **General Server** host properties dialog box.  
See “[General Server properties](#)” on page 127.
- 2 Add an entry in the **Media Host Override** list that lists the original backup server and the restore server.
- 3 Click **OK**.

### Overriding the original server for restores manually

Use the following procedure to manually override the original server for restores.

#### To manually override the original server for restores

- 1 If necessary, physically move the media and update the Enterprise Media Manager database Media Manager volume database to reflect the move.
- 2 Modify the NetBackup configuration on the master server as follows:
  - By using the NetBackup Administration Console:  
Open the **General Server** host properties dialog box of the master server. Add an entry in the **Media Host Override** list that lists the original backup server and the restore server.

- By modifying the `bp.conf` file on a UNIX NetBackup server:

As `root` add the following entry to the

```
/usr/openv/netbackup/bp.conf file:
FORCE_RESTORE_MEDIA_SERVER = fromhost tohost
```

The *fromhost* is the server that wrote the original backup and the *tohost* is the server to use for the restore.

To revert to the original configuration for future restores, delete the changes made in this step.

- 3 Click **OK**.
- 4 Stop and restart the NetBackup Request daemon on the master server.

The override applies to all storage units on the original server. This means that restores for any storage unit on *fromhost* go to *tohost*.



## Enabling automatic failover to an alternate server

NetBackup allows the administrator to configure automatic restore failover to an alternate server if the original server is temporarily inaccessible. Once it is configured, this method does not require administrator intervention.

See “[Restore Failover properties](#)” on page 169.

Some examples of when to use this method are as follows:

- Two or more servers share a robot, each with connected drives.  
When a restore is requested, one of the servers is temporarily inaccessible.
- Two or more servers have stand-alone drives of the same type.  
When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the master server and `bptm` on the original server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The original server is down.
- The original server is up but `bpcd` on that server does not respond. (For example, if the connection is refused or access is denied.)
- The original server is up and `bpcd` is fine, but `bptm` has problems. (For example, if `bptm` cannot find the required tape.)

---

**Note:** The failover uses only the failover hosts that are listed in the NetBackup configuration. By default, the list is empty and NetBackup does not perform the automatic failover.

---

## Failing over to an alternate server

Use the following procedure to enable automatic failover to an alternate server.

### To enable automatic failover to an alternate server

- 1 Modify the NetBackup configuration on the master server as follows:
  - By using the NetBackup Administration Console:  
Open the **Restore Failover** host properties dialog box of the master server.  
Add an entry in the **Alternate Restore Failover Machines** list that lists the media server and failover restore server(s).
  - By modifying the `bp.conf` file on a UNIX NetBackup server:  
As `root`, add the following entry to the `/usr/openv/netbackup/bp.conf` file:

```
FAILOVER_RESTORE_MEDIA_SERVERS =
failed_host host1 host2 ... hostN
```

Where:

*failed\_host* is the server that is not operational.

*host1 ... hostN* are the servers that provide failover capabilities.

When automatic failover is necessary for a given server, NetBackup searches through the relevant `FAILOVER_RESTORE_MEDIA_SERVERS` list. NetBackup looks from left to right for the first server that is eligible to perform the restore.

There can be multiple `FAILOVER_RESTORE_MEDIA_SERVERS` entries and each entry can have multiple servers. However, a NetBackup server can be a *failed\_host* in only one entry.

- 2 Stop and restart the NetBackup Request daemon on the master server.

## Expiring and importing media for alternate server restores

Regarding expiring and importing media, even with the alternate server restore capabilities, it may be necessary to expire media and then import it.

Regarding identifying media spanning groups, an alternate server restore operation can include media IDs that contain backup images that span media. It may be necessary to identify the media IDs that contain fragments of the spanned images. The group of related media is called a media spanning group.

To identify the media in a specific media spanning group, run the following command as administrator from the command prompt on the NetBackup master server:

```
cd install_path\NetBackup\bin
bpimmedia.exe -spangroups -U -mediaid media_id
```

To display all media in all spanning groups, omit `-mediaid media_id` from the command.

# Managing client restores

This chapter includes the following topics:

- [Server-directed restores](#)
- [Client-redirected restores](#)
- [Restoring files and access control lists](#)
- [How to improve search times by creating an image list](#)
- [How to restore System State](#)

## Server-directed restores

By default, NetBackup clients are configured to allow NetBackup administrators on a master server to direct restores to any client.

To prevent server-directed restores, configure the client accordingly as follows:

- **Windows clients**  
Open the Backup, Archive, and Restore interface on the client.  
Select **File > NetBackup Client Properties > General tab > Clear the Allow server-directed restores** checkbox.
- **UNIX clients**  
Add `DISALLOW_SERVER_FILE_WRITES` to the following file on the client:

```
/usr/opensv/netbackup/bp.conf
```

---

**Note:** On UNIX systems, the redirected restores can incorrectly set UIDs or GIDs that are too long. The UIDs and GIDs of files that are restored from one platform to another may be represented with more bits on the source system than on the destination system. If the UID or the GID name in question is not common to both systems, the original UID or GID may be invalid on the destination system. In this case, the UID or GID is replaced with the UID or GID of the user that performs the restore.

---

Consider the following solutions:

- To produce a progress log, add the requesting server to the server list. To do so, log into the requesting server. In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers** > Double-click on the master server > **Servers**. Add the restoring server to the server list.
- Log on to the restoring server. Check the Activity Monitor to determine the success of the restore operation.

To restore a UNIX backup that contains soft and hard links, run the Backup, Archive, and Restore client interface from a UNIX machine. Only the Java version of the client interface contains the **Rename hard links** and **Rename soft links** restore options. Windows users can install the Windows display console to access the Java version of the Backup, Archive, and Restore interface from a Windows machine.

## Client-redirected restores

The Backup, Archive, and Restore client interface contains options for allowing clients to restore the files that were backed up by other clients. The operation is called a redirected restore.

## Restore restrictions

By default, NetBackup permits only the client that backs up files to restore those files. NetBackup ensures that the client name of the requesting client matches the peer name that was used to connect to the NetBackup server.

Unless clients share an IP address, the peer name is equivalent to the client's host name. (Clients can share an IP address due to the use of a gateway and token ring combination, or multiple connections.) When a client connects through a gateway, the gateway can use its own peer name to make the connection.

The NetBackup client name is normally the client's short host name, such as `client1` rather than a longer form such as `client1.null.com`.

The client name is found in the following locations:

- Windows clients (including NetWare NonTarget):  
Open Backup, Archive, and Restore and select **File > Specify NetBackup Machines and Policy Type**. The client name that is selected as **Source Client for Restores** is the source of the backups to be restored.
- On NetWare target clients:  
Specify the client name in the `bp.ini` file.
- UNIX clients:  
Open Backup, Archive, and Restore and select the client name as the **Source client for restore**.

## To allow all clients to perform redirected restores

The NetBackup administrator can allow clients to perform redirected restores. That is, allow all clients to restore the backups that belong to other clients. Place an empty `No.Restrictions` file on the NetBackup master server where the policy that backed up the other clients resides.

---

**Note:** The information in this topic applies to restores made by using the command line, not the Backup, Archive, and Restore client interface.

---

Create an `altnames` directory in the following location, then place the empty file inside of the directory:

```
Install_path\NetBackup\db\altnames\No.Restrictions
```

The NetBackup client name setting on the requesting client must match the name of the client for which the backup was created. The peer name of the requesting client does not need to match the NetBackup client name setting.

---

**Note:** Do not add a suffix to the files in the `altnames` directory.

---

---

**Note:** The `Install_path\NetBackup\db\altnames` directory can present a potential breach of security. Users that are permitted to restore files from other clients may also have local permission to create the files that are found in the backup.

---

## To allow a single client to perform redirected restores

The NetBackup administrator can permit a single client to restore the backups that belong to other clients. Create a *peername* file on the NetBackup master server where the policy that backed up the other client(s) resides.

---

**Note:** The information in this topic applies to restores made by using the command line, not the Backup, Archive, and Restore client interface.

---

Create an `altnames` directory in the following location, then place the empty file inside of the directory:

```
Install_path\NetBackup\db\altnames\peername
```

Where *peername* is the client to possess restore privileges.

In this case, the requesting client (*peername*) can access the files that are backed up by another client. The NetBackup client name setting on *peername* must match the name of the other client.

## To allow redirected restores of a client's files

The NetBackup administrator can permit a single client to restore the backups that belong to another client. Create a *peername* file on the NetBackup master server of the requesting client as described here.

---

**Note:** The information within this topic applies to restores made using the command line, not the Backup, Archive, and Restore client interface.

---

Create an `altnames` directory in the following location, then place the *peername* file inside of the directory:

```
Install_path\NetBackup\db\altnames\peername
```

Where *peername* is the client to possess restore privileges. Add to the *peername* file the names of the client(s) whose files the requesting client wants to restore.

The requesting client can restore the files that were backed up by another client if:

- The names of the other clients appear in the *peername* file, and
- The NetBackup client name of the requesting client is changed to match the name of the client whose files the requesting client wants to restore.

## Examples of redirected restores

This topic provides some example configurations that allow clients to restore the files that were backed up by other clients. These methods may be required when a client connects through a gateway or has multiple Ethernet connections.

In all cases, the requesting client must have access to an image database directory on the master server (*Install\_path*\NetBackup\db\images\*client\_name*). Or, the requesting client must be a member of an existing NetBackup policy.

---

**Note:** Not all file system types on all machines support the same features. Problems can be encountered when a file is restored from one file system type to another. For example, the S51K file system on an SCO machine does not support symbolic links nor does it support names greater than 14 characters long. You may want to restore a file to a machine that doesn't support all the features of the machine from which the restore was performed. In this case, all files may not be recovered.

---

In the following examples, assume the following conditions:

- *client1* is the client that requests the restore.
- *client2* is the client that created the backups that the requesting client wants to restore.
- *Install\_path* is the path where you installed the NetBackup software. By default, this path is C:\Program Files\VERITAS.

---

**Note:** The information in this topic applies to restores made by using the command line, not the Backup, Archive, and Restore client interface.

---

---

**Note:** You must have the necessary permissions to perform the following steps.

---

### Example 1: Redirected client restore

Assume you must restore files to *client1* that were backed up from *client2*. The *client1* and *client2* names are those specified by the NetBackup client name setting on the clients.

In the nominal case, do the following:

- Log on on the NetBackup server.

Add *client2* to the following file and perform one of the following::

- Edit `Install_path\NetBackup\db\altnames\client1` to include the name of `client2`.
- Create the following empty file:

```
Install_path\NetBackup\db\altnames\No.Restrictions
```

- Log on on `client1` and change the NetBackup client name to `client2`.
- Restore the file.
- Undo the changes that were made on the server and client.

### Example 2: Redirected client restore using the altnames file


This example explains how `altnames` provides restore capabilities to clients that do not use their own host name when they connect to the NetBackup server.

By default, the NetBackup client name of the requesting client must match the peer name that is used in the connection to the NetBackup server. When the NetBackup client name is the host name for the client and matches the peer name (normal case), this requirement is met.

However, problems arise when clients connect to multiple ethernet or connect to the NetBackup server through a gateway.

Consider the configuration in [Figure 26-1](#).

**Figure 26-1** Example restore from token ring client



In this example, restore requests from `client1`, `client2`, and `client3` are routed through the TCP gateway. Because the gateway uses its own peer name rather than the client host names for connection to the NetBackup server, NetBackup refuses the requests. Clients cannot restore even their own files.

To correct the situation, do the following.



Determine the peer name of the gateway:

- Try a restore from the client in question. In this example, the request fails with an error message similar to the following:

```
client is not validated to use the server
```

- Examine the NetBackup problems report and identify the peer name that is used on the request. Entries in the report may be similar to the following:

```
01/29/07 08:25:03 bpserver - request from invalid
server or client client1.dvlp.null.com
```

In this example, the peer name is `client1.dvlp.null.com`.

Determine the peer name, then create the following file on the NetBackup master server:

```
Install_path\NetBackup\db\altnames\peername
```

In this example, the file is:

```
Install_path\NetBackup\db\altnames\client1.dvlp.null.com
```

Edit the *peername* file so that it includes the client names.

For example, if you leave the file

```
Install_path\NetBackup\db\altnames\client1.dvlp.null.com
```

empty, *client1*, *client2*, and *client3* can all access the backups that correspond to their NetBackup client name setting.

See [“To allow a single client to perform redirected restores”](#) on page 774.

If you add the names *client2* and *client3* to the file, you give these two clients access to NetBackup file restores, but exclude *client1*.

See [“To allow redirected restores of a client’s files”](#) on page 774.

Note that this example requires no changes on the clients.

Restore the files.

See [“To allow redirected restores of a client’s files”](#) on page 774.

See [“To allow a single client to perform redirected restores”](#) on page 774.

### Example 3: Troubleshoot redirected client restore using the altnames file

If you cannot restore files with a redirected client restore by using the `altnames` file, troubleshoot the situation, as follows:

- On the master server, in the NetBackup Administration Console, select **NetBackup Management > Host Properties > Master Server > Double-click on master server > Universal Settings**. Enable the **Enable Performance Data Collection** property.

- Create the debug log directory for the NetBackup Request Manager service:

```
Install_path\NetBackup\logs\bprd
```

- On the master server, stop and restart the NetBackup Request Manager service. Restart the service to ensure that this service is running in verbose mode and logs information regarding client requests.

- On *client1* (the requesting client), try the file restore.

- On the master server, identify the peer name connection that *client1* uses.

- Examine the failure as logged on the All Log Entries report. Or, examine the debug log for the NetBackup Request Manager service to identify the failing name combination:

```
Install_path\NetBackup\logs\bprd\mmdyy.log
```

- On the master server, do one of the following:

- Create an *Install\_path\NetBackup\db\altnames\No.Restrictions* file. The file allows any client to access *client2* backups if the client changes its NetBackup client name setting to *client2*.

- Create an *Install\_path\NetBackup\db\altnames\peername* file. The file allows *client1* to access *client2* backups if *client1* changes its NetBackup client name setting to *client2*.

- Add *client2* name to the following file:

```
Install_path\NetBackup\db\altnames\peername.
```

- *client1* is allowed to access backups on *client2* only.

- On *client1*, change the NetBackup client name setting to match what is specified on *client2*.

- Restore the files from *client1*.

- Perform the following:

- Delete `Install_path\NetBackup\logs\bprd` and the contents.
- On the master server, select **NetBackup Management > Host Properties > Master Server > Double-click on master server > Clean-up**. Clear the **Keep Logs** property.
- If you do not want the change to be permanent, do the following:
  - Delete `Install_path\NetBackup\db\altnames\No.Restrictions` (if existent)
  - Delete `Install_path\NetBackup\db\altnames\peername` (if existent)
  - On `client1`, change the NetBackup client name to its original value.

## Restoring files and access control lists

An access control list (ACL) is a table that conveys the access rights users need to a file or directory. Each file or directory can have a security attribute that extends or restricts users' access.

### Restoring the files that possess ACLs

By default, the NetBackup-modified GNU tar (`/usr/opensv/netbackup/bin/tar`) restores ACLs along with file and directory data.

However, in some situations the ACLs cannot be restored to the file data, as follows:

- Where the restore is cross-platform. (Examples: An AIX ACL restored to a Solaris client or a Windows ACL restored to an HP client.)
- When a tar other than the NetBackup modified tar is used to restore files.

In these instances, NetBackup stores the ACL information in a series of generated files in the `root` directory using the following naming form:

```
.SeCuRiT.y.nnnn
```

These files can be deleted or can be read and the ACLs regenerated by hand.

More information is available in the *NetBackup Administrator's Guide for Windows, Volume II*.

### Restoring files without restoring ACLs

The NetBackup client interface on Windows is available to administrators to restore data without restoring the ACLs. Both the destination client and the source of the backup must be Windows systems.

To restore files without restoring ACLs, the following conditions must be met:

- The policy that backed up the client is of policy type MS-Windows.
- An administrator performs the restore and is logged into a NetBackup server (Windows or UNIX). The option is set at the server by using the client interface. The option is unavailable on stand-alone clients (clients that do not contain the NetBackup server software).
- The destination client and the source of the backup must both be systems running supported Windows OS levels. The option is disabled on UNIX clients.

Use the following procedure to restore files without restoring ACLs.

#### To restore files without restoring ACLs

- 1 Log on to the NetBackup server as administrator.
- 2 Open the Backup, Archive, and Restore client interface.
- 3 From the client interface, initiate a restore.
- 4 Select the files to be restored, then select **Actions > Start Restore of Marked Files**.
- 5 In the **Restore Marked Files** dialog box, place a check in the **Restore without access-control attributes** check box.
- 6 Make any other selections for the restore job.
- 7 Click **Start Restore**.

## How to improve search times by creating an image list

Create an image list to improve searching among many small backup images.

Run the following command on the master server while logged on as administrator.

Enter the following as one line:

```
install_path\netbackup\bin\admincmd\bpimage
-create_image_list -client name
```

Where *name* is the name of the client with small backup images.

The command creates files in the following location:

```
install_path\netbackup\db\images\clientname
```

IMAGE\_LIST: List of images for this client

IMAGE\_INFO: Information about the images for this client

`IMAGE_FILES`: The file information for small images

Do not edit these files. The files contain offsets and byte counts that are used to seek and read the image information.

The files require 35 to 40% more space in the client directory. The files improve search performance only if thousands of small backup images for a client exist.

## How to restore System State

The System State includes the registry, the COM+ Class Registration database, and boot and system files. If the server is a domain controller, the data also includes the Active Directory services database and the SYSVOL directory.

---

**Note:** The best recovery procedure depends on many hardware and software variables that pertain to the server and its environment. For a complete Windows recovery procedure, refer to the Microsoft documentation.

---

Read the following notes carefully before you restore the System State:

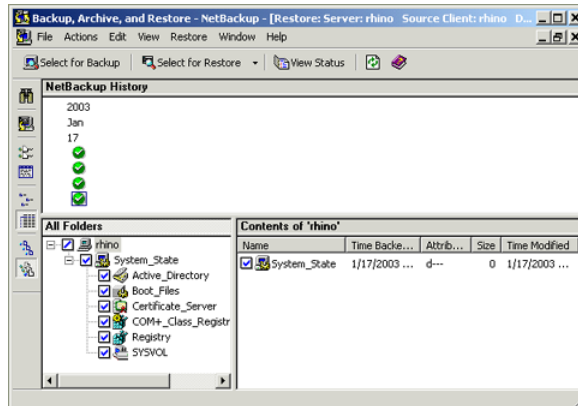
- The System State should be restored in its entirety. Do not restore selected files.
- Although incremental backups of the System State can be configured, NetBackup always performs a full backup. Therefore, only the most recent backup of the System State must be restored.
- Do not redirect a System State restore. System State is computer-specific and to restore it to an alternate computer can result in an unusable system.
- Do not cancel a System State restore operation. To cancel the operation may leave the system unusable.
- To restore the System State to a domain controller, the Active Directory must not be running.

## Restoring the system state

Use the following procedure to restore the system state.

### To restore the system state

- 1 To restore the Active Directory, restart the system, and press F8 during the boot process. F8 brings up a startup options menu. Press F8 upon restart if the system to which you are to restore is a Windows domain controller. Otherwise, begin with step 4.
- 2 From the startup options, select **Directory Services Restore Mode** and continue the boot process.
- 3 Ensure that the **NetBackup Client Service**, `bpinetd`, has started. Use the Activity Monitor or the Services application in the Windows Control Panel.
- 4 Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and place a checkmark next to **System State**.



- 5 From the **Actions** menu, select **Start Restore of Marked Files**.
- 6 From the **Restore Marked Files** dialog box, select **Restore everything to its original location** and **Overwrite the existing file**.  
Do not redirect the System State restore to a different host. System State is computer-specific . To restore it to a different computer can result in an unusable system.
- 7 Click **Start Restore**.

- 8 The network may contain more than one domain controller. To replicate Active Directory to other domain controllers, perform an authoritative restore of the Active Directory after the NetBackup restore job completes.

To perform an authoritative restore of the Active Directory, run the Microsoft `ntdsutil` utility after you restored the System State data but before the server is restarted. An authoritative restore ensures that the data is replicated to all of the servers.

Additional information about an authoritative restore and the `ntdsutil` utility is available.

Refer to the Microsoft documentation.

- 9 Reboot the system before performing subsequent restore operations.

If you booted into **Directory Services Restore Mode** on a domain controller, reboot into normal mode when the restore is complete.





# Powering down and rebooting NetBackup servers

This chapter includes the following topics:

- [Powering down and rebooting NetBackup servers](#)

## Powering down and rebooting NetBackup servers

To close and restart NetBackup servers, use the following recommended procedures.

### To power down a server

- 1 In the NetBackup Administration Console, click **Activity Monitor**, then select the Jobs tab to make sure no backups or restores are running.
- 2 Use the NetBackup Administration Console or the command line to stop the NetBackup Request service, `bprd`. Stop `bprd` to stop additional backup and restore activity and to allow current activity to end.
- 3 From the NetBackup Administration Console, click **Activity Monitor**, then select the Services tab. Right-click the services that are running and select **Stop Service**.
- 4 From the command line, run:

```
Install_path\NetBackup\bin\bpdwn.exe
```

- 5 From the command line, enter:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 6 Power down the server.

## Shutting down all NetBackup services

Use the following procedure to shut down all NetBackup services.

To shut down all NetBackup services

From a command line, enter the following:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

## Starting up all NetBackup services

Use the following procedure to start up all NetBackup services.

To start up all NetBackup services

From a command line, enter the following:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

## Rebooting a NetBackup server

Use the following procedure to reboot a NetBackup server.

**To reboot a NetBackup master server**

- 1 Restart the system.
- 2 If the required NetBackup services are not set up to start automatically, do the following:
  - From the Windows desktop, start the Windows Services applet.
  - Start the NetBackup Client service.
  - Start the NetBackup Device Manager service. The NetBackup Volume Manager service automatically starts as well.
  - Start the NetBackup Request Manager service to start the NetBackup Database Manager service.

## Rebooting a NetBackup media server

Use the following procedure to reboot a NetBackup media server.

### To reboot a NetBackup media server

- 1 Restart the system.
- 2 The required NetBackup services start automatically if they are set up to do so.

If they are not set to start automatically, do the following:

- From the Windows desktop, start the Windows Services applet.
- Start the NetBackup Client service.
- Start the NetBackup Device Manager service (`ltid`). The NetBackup Volume Manager service (`vmd`) starts as well.

## About displaying robotic processes with `vmops`

The `vmops` script shows the Media Manager daemons and robotic processes that are active on a UNIX system.

You can run this script by using the following command:

```
/usr/opensv/volmgr/bin/vmops
```

In the following sample, the second column contains the process IDs for the processes.

```
root 303 0.0 0.2 136 264 ? S Feb 11 4:32 ltid -v
root 305 0.0 0.0 156 0 ? IW Feb 11 0:54 vmd -v
root 306 0.0 0.0 104 0 ? IW Feb 11 0:15 tl8d -v
root 307 0.0 0.0 68 56 ? S Feb 11 12:16 avrd
root 310 0.0 0.0 116 0 ? IW Feb 11 0:07 tl8cd -v
```

Status for the `nbemm` command is not shown in the output of `vmops`. The `nbemm` status is shown in the output of the `bpps` script.



# About Granular Recovery Technology

This chapter includes the following topics:

- [About installing and configuring Network File System \(NFS\) for Active Directory Granular Recovery](#)
- [About configuring Services for Network File System \(NFS\) on the Windows 2008 and Windows 2008 R2 NetBackup media server and NetBackup clients](#)
- [About configuring Services for Network File System \(NFS\) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients](#)
- [Configuring a UNIX or Linux media server and Windows clients for backups and restores that use Granular Recovery Technology](#)
- [Configuring a different network port for NBFSD](#)
- [Configuring the log on account for the NetBackup Client Service](#)

## About installing and configuring Network File System (NFS) for Active Directory Granular Recovery

NetBackup Granular Recovery leverages Network File System, or NFS, to read individual objects from a database backup image. Specifically, the NetBackup client uses NFS to extract data from the backup image on the NetBackup media server. The NetBackup client uses “Client for NFS” to mount and access a mapped drive that is connected to the NetBackup media server. The NetBackup media server handles the I/O requests from the client through NBFSD.

NBFSD is the NetBackup File System (NBFS) service that runs on the media server. NBFSD makes a NetBackup backup image appear as a file system folder to the NetBackup client over a secure connection.

Network File System, or NFS, is a widely recognized, open standard for client and server file access over a network. It allows clients to access files on dissimilar servers through a shared TCP/IP network. NFS is typically bundled with the host operating system. NetBackup uses Granular Recovery Technology (GRT) and NFS to recover the individual objects that reside within a database backup image, such as:

- A user account from an Active Directory database backup
- Email messages or folders from an Exchange database backup
- A document from a SharePoint database backup

Multiple NetBackup agents that support GRT (for example, Exchange, SharePoint, and Active Directory) can use the same media server.

## About configuring Services for Network File System (NFS) on the Windows 2008 and Windows 2008 R2 NetBackup media server and NetBackup clients

To configure NFS in a Windows 2008 or Windows 2008 R2 environment, perform the following configuration:

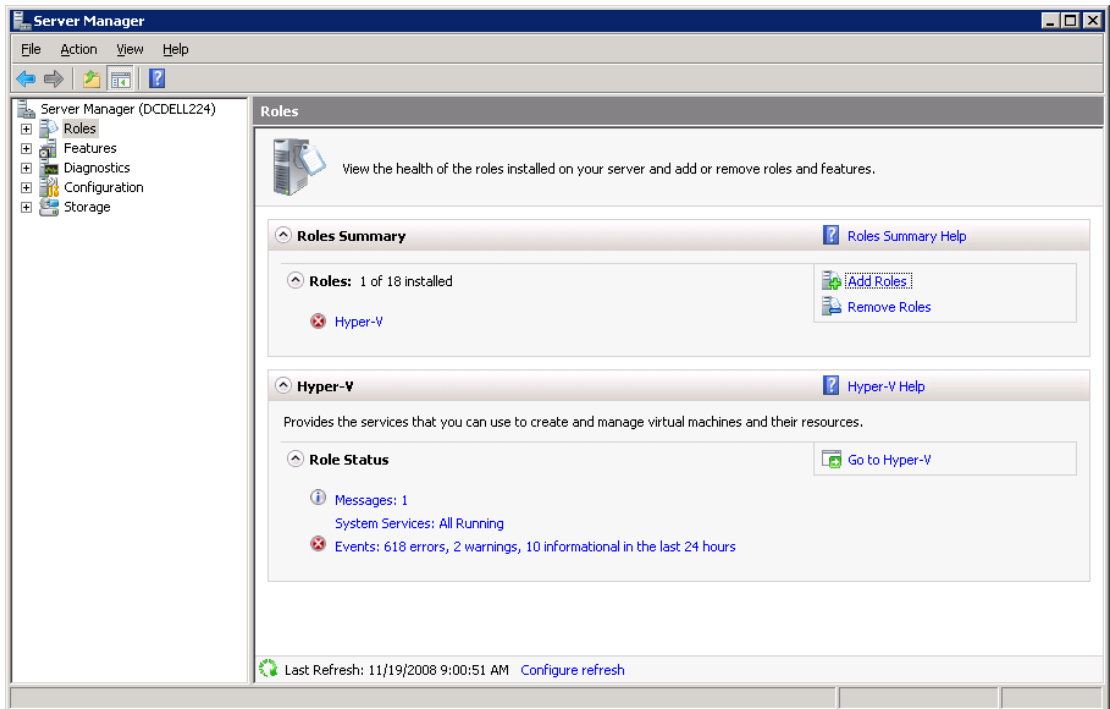
- Enable NFS on the following:
  - The NetBackup media server
  - All Active Directory domain controllers or ADAM/LDS hosts.  
See [“Enabling Services for Network File System \(NFS\) on Windows 2008 or Windows 2008 R2”](#) on page 791.
- You can disable the Server for NFS on the following:
  - The NetBackup media server
  - All Active Directory domain controllers or ADAM/LDS hosts.  
See [“Disabling the Server for NFS”](#) on page 795.
- You can disable the Client for NFS on the NetBackup media server.  
See [“Disabling the Client for NFS on the media server”](#) on page 794.  
If the Active Directory domain controller or ADAM/LDS host resides on the media server, do not disable the Client for NFS.

## Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2

To restore individual items from a backup that uses Granular Recovery Technology (GRT), you must enable Services for Network File System. When this configuration is completed on the media server and all Active Directory domain controllers or ADAM/LDS hosts, you can disable any unnecessary NFS services.

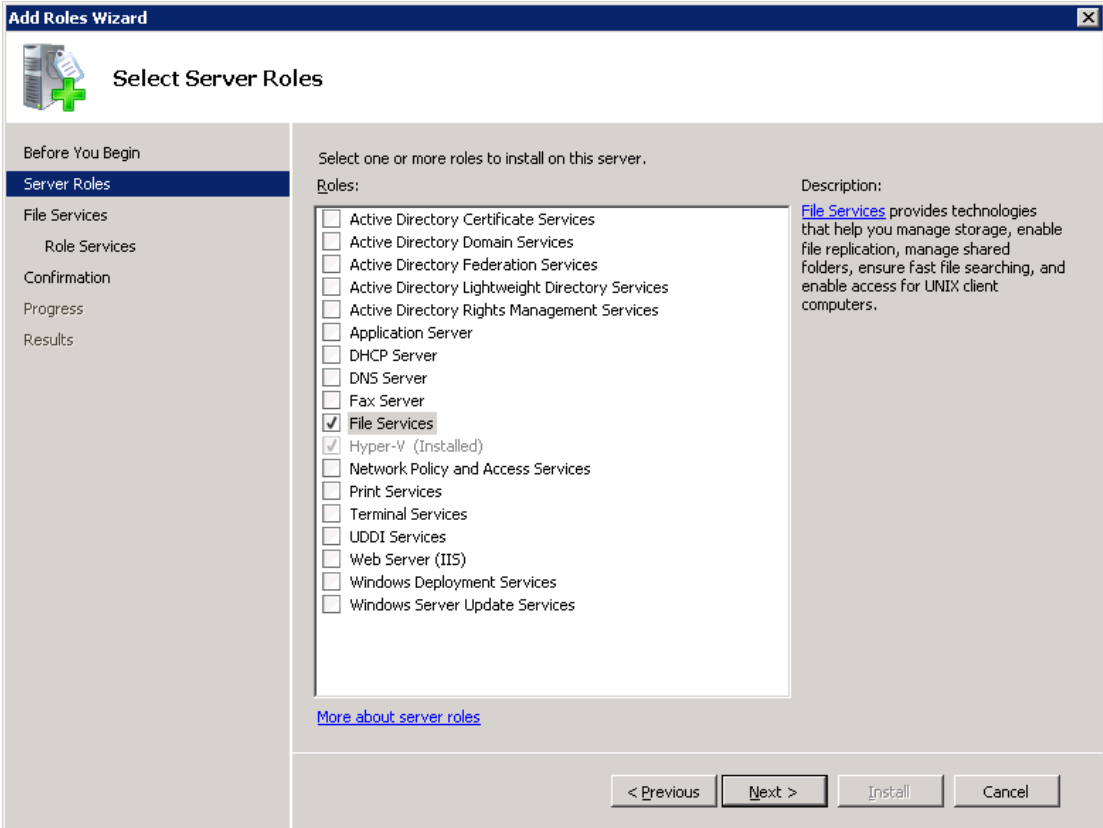
**To enable Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2**

- 1 Open the Server Manager.
- 2 In the left pane, click **Roles** and, in the right pane, click **Add Roles**.



- 3 In the Add Roles Wizard, on the **Before You Begin** page, click **Next**.

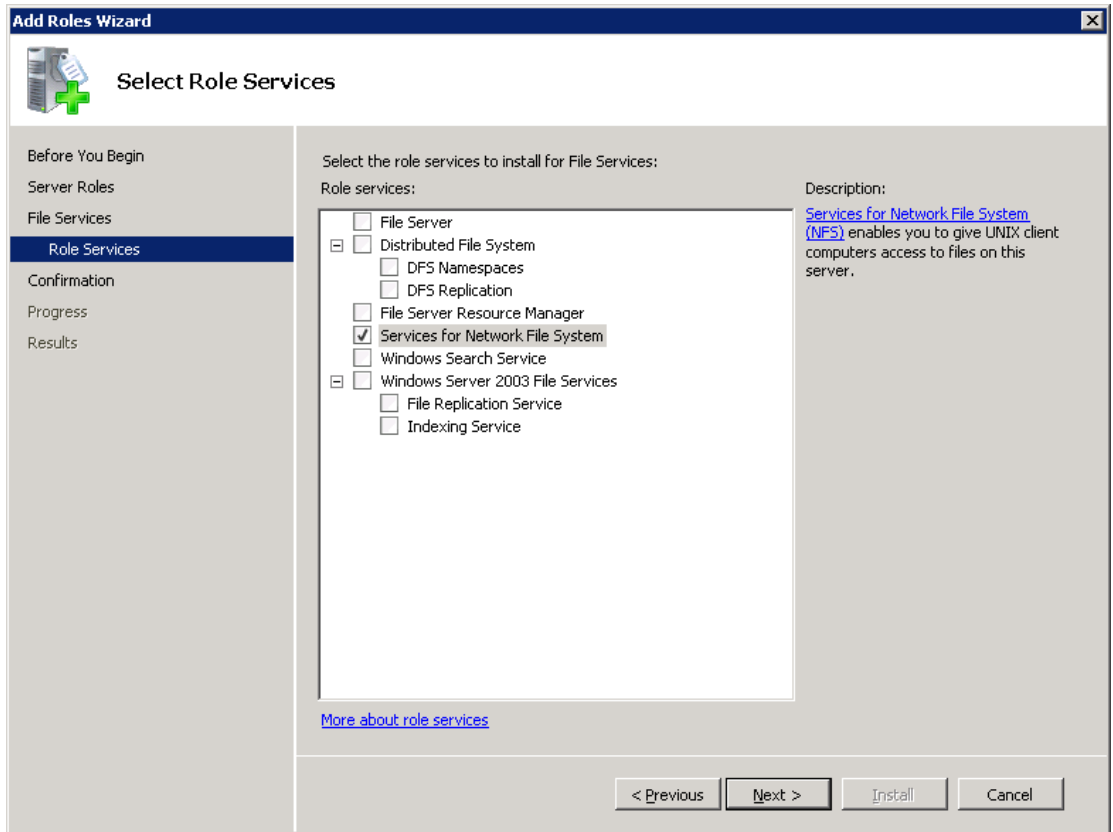
- 4 On the **Select Server Roles** page, under **Roles**, check the **File Services** check box.



- 5 Click **Next**.
- 6 On the **Files Services** page, click **Next**.
- 7 On the **Select Role Services** page, uncheck **File Server**.



## 8 Check Services for Network File System.



9 Click **Next** and complete the wizard.

10 On the media server, configure the portmap service to start automatically at server restart.

Issue the following from the command prompt:

```
sc config portmap start= auto
```

This command should return the status [SC] ChangeServiceConfig SUCCESS.

11 For each host in your configuration, choose from one of the following:

- If you have a single host that functions as both the media server and the Active Directory domain controllers or ADAM/LDS host, you can disable the Server for NFS.

- For a host that is only the NetBackup media server, you can disable the Server for NFS and the Client for NFS.
- For a host that is only an Active Directory domain controllers or ADAM/LDS host, you can disable the Server for NFS.

See “Disabling the Server for NFS” on page 795.

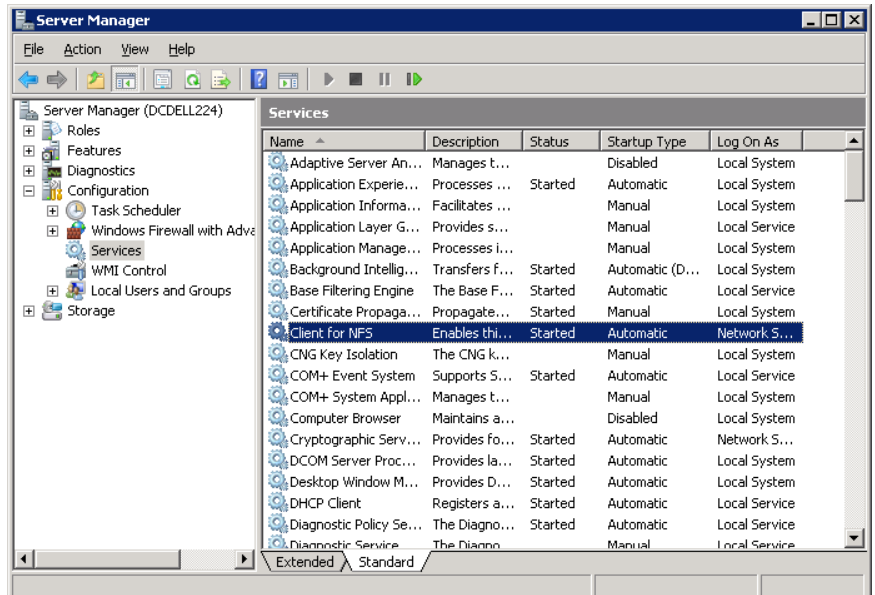
See “Disabling the Client for NFS on the media server” on page 794.

## Disabling the Client for NFS on the media server

After you enable Services for Network File System (NFS) on a host that is only a NetBackup media server, you can disable the Client for NFS.


To disable the Client for NFS on the NetBackup media server

- 1 Open the Server Manager.
- 2 In the left pane, expand **Configuration**.
- 3 Click **Services**.



- 4 In the right pane, right-click on **Client for NFS** and click **Stop**.
- 5 In the right pane, right-click on **Client for NFS** and click **Properties**.

- 6 In the **Client for NFS Properties** dialog box, from the **Startup type** list, click **Disabled**.



- 7 Click **OK**.

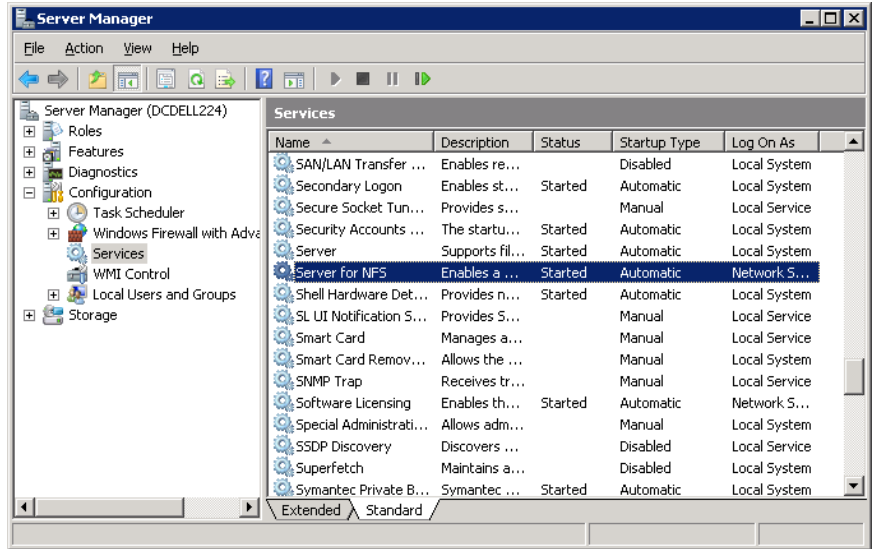
## Disabling the Server for NFS

After you enable Services for Network File System (NFS) on the media server and on the Active Directory domain controllers or ADAM/LDS hosts, you can disable Server for NFS.

### To disable the Server for NFS


- 1 Open the Server Manager.
- 2 In the left pane, expand **Configuration**.

3 Click **Services**.



- 4 In the right pane, right-click on **Server for NFS** and click **Stop**.
- 5 In the right pane, right-click on **Server for NFS** and click **Properties**.

- 6 In the **Server for NFS Properties** dialog box, from the **Startup type** list, click **Disabled**.



- 7 Click **OK**.
- 8 Repeat this procedure for the media server and for all Active Directory domain controllers or ADAM/LDS hosts.

## About configuring Services for Network File System (NFS) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients

---

**Note:** NetBackup does not support Granular Recovery Technology (GRT) with Windows Server 2003 R1 or earlier versions.

---

For a Windows 2003 R2 SP2 environment, perform the following configuration:

- Install the necessary NFS components on the NetBackup media server.

See [“Installing Services for NFS on the Windows 2003 R2 SP2 media server”](#) on page 798.

- Install the necessary NFS components on all Active Directory domain controllers or ADAM/LDS hosts.

See [“Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2”](#) on page 801.

**Table 28-1** NFS components required for Windows 2003 R2 SP2

| NFS component                             | NetBackup client | NetBackup media server |
|-------------------------------------------|------------------|------------------------|
| Client for NFS                            | X                |                        |
| Microsoft Services for NFS Administration | X                |                        |
| RPC External Data Representation          | X                | X                      |
| RPC Port Mapper                           |                  | X                      |

**Note:** If the Active Directory domain controllers or ADAM/LDS host resides on the media server, install all the components on the media server.

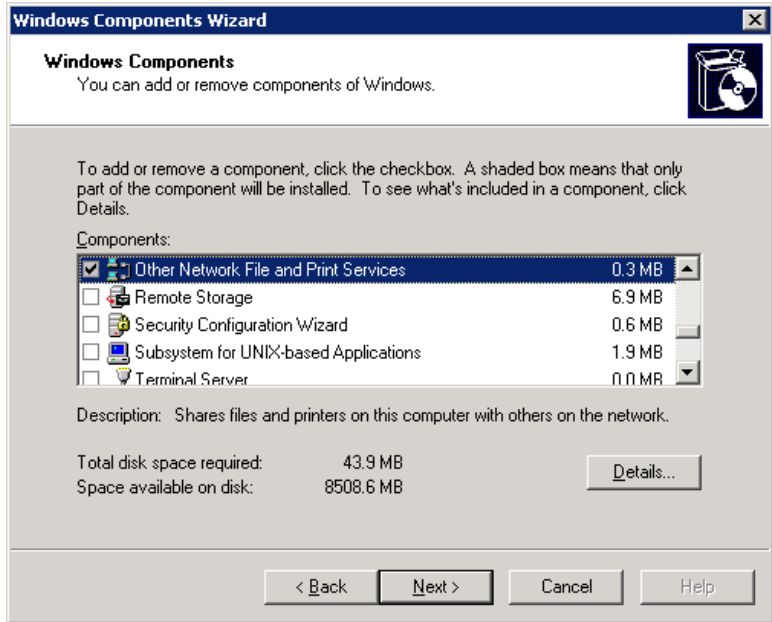
## Installing Services for NFS on the Windows 2003 R2 SP2 media server

This topic describes how to install Services for NFS on a Windows 2003 R2 SP2 media server.

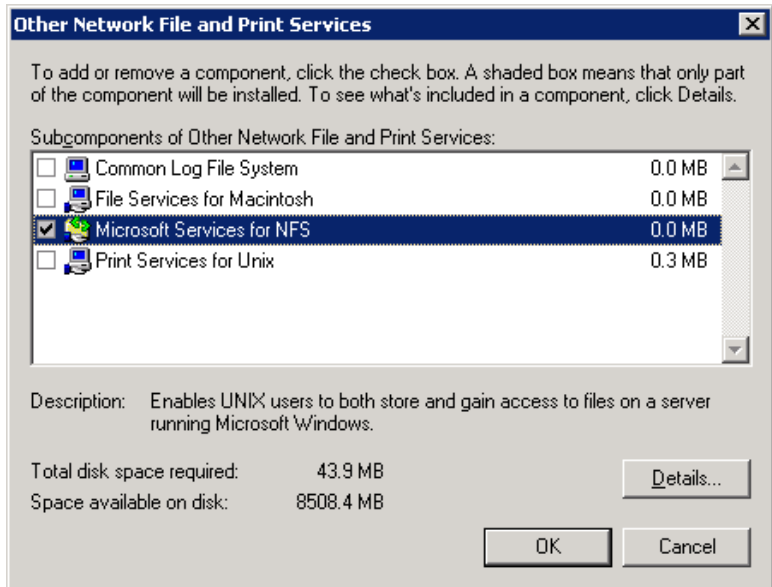
To install Services for NFS on the Windows 2003 R2 SP2 media server

- 1 Click **Start > Control Panel > Add or Remove Programs**.
- 2 Click **Add/Remove Windows Components**.

**3 Check Other Network File and Print Services and click Details.**



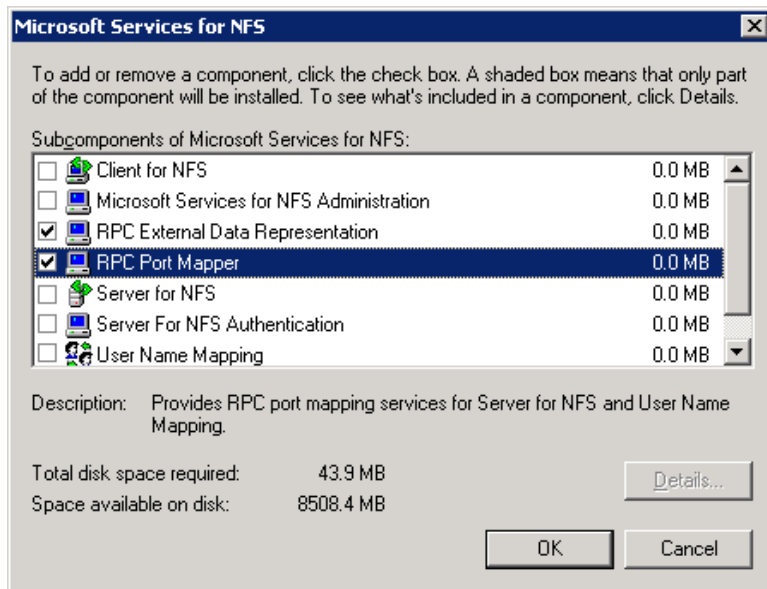
**4 Check Microsoft Service for NFS and click Details.**



- 5 Install the components that apply to your configuration.
  - If the host is only a NetBackup media server, check the following components:
    - RPC External Data Representation
    - RPC Port Mapper
  - If you have a single host that functions as both the media server and the Active Directory domain controllers or ADAM/LDS host, check the following components:
    - Client for NFS
    - Microsoft Services for NFS Administration
    - RPC External Data Representation
    - RPC Port Mapper

Media server  
and client

Media  
server only



- 6 Click **OK**.
- 7 Click **OK**.
- 8 Click **Next** and complete the Windows Components Wizard.
- 9 After the installation is complete, open Services in the Control Panel.



- 10 Depending on configuration of the host, verify that Client for NFS is running or is stopped and disabled:
  - For a single host that has both the media server and the Active Directory domain controller or ADAM/LDS, ensure Client for NFS is running.
  - For a host that is only a NetBackup media server, Client for NFS can be stopped and disabled.
- 11 Configure the portmap service to start automatically at server restart.

Issue the following from the command prompt:

```
sc config portmap start= auto
```

This command should return the status [SC] ChangeServiceConfig SUCCESS.

## Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2


This topic describes how to install NFS on the NetBackup clients with Windows 2003 R2 SP2. Only the clients that are Active Directory domain controllers or ADAM/LDS hosts require NFS. If an Active Directory domain controllers or ADAM/LDS host is also a media server, you must follow a different procedure.

See [“Installing Services for NFS on the Windows 2003 R2 SP2 media server”](#) on page 798.


**To install Services for NFS on the NetBackup clients with Windows 2003 R2 SP2**

- 1 Click **Start > Control Panel > Add or Remove Programs**.
- 2 Click **Add/Remove Windows Components**.

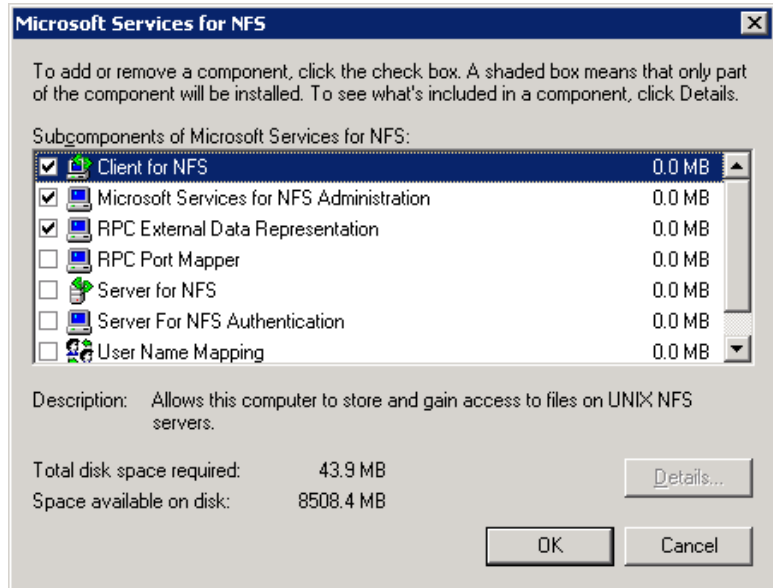
### 3 Check **Other Network File and Print Services** and click **Details**.



### 4 Check **Microsoft Service for NFS** and click **Details**.



- 5 Check the following components:
  - Client for NFS
  - Microsoft Services for NFS Administration
  - RPC External Data Representation



- 6 Click **OK**.
- 7 Click **OK**.
- 8 Click **Next** and complete the Windows Components Wizard.
- 9 After the installation is complete, open Services in the Control Panel.
- 10 Ensure the following that the Client for NFS service is running.
- 11 Repeat this procedure for all Active Directory domain controllers or ADAM/LDS hosts.

## Configuring a UNIX or Linux media server and Windows clients for backups and restores that use Granular Recovery Technology

To perform backups and restores that use Granular Recovery Technology, perform the following configuration if you use a UNIX or Linux media server and Windows clients:

- Confirm that your media server is installed on a platform that supports granular recovery.  
See the *NetBackup Enterprise Server and Server 7.x OS Software Compatibility List*.
- No other configuration is required for the UNIX or Linux media server.
- Enable or install NFS on all Active Directory domain controllers or ADAM/LDS hosts.  
See [“Enabling Services for Network File System \(NFS\) on Windows 2008 or Windows 2008 R2”](#) on page 791.  
See [“Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2”](#) on page 801.
- You can configure a different network port for NBFSD.  
See [“Configuring a different network port for NBFSD”](#) on page 804.

## Configuring a different network port for NBFSD

NBFSD runs on port 7394. If another service uses the standard NBFSD port in your organization, you can configure the service on another port. The following procedures describe how to configure a NetBackup server to use a network port other than the default.

### To configure a different network port for NBFSD (Windows server)

- 1 Log on as administrator on the computer where NetBackup server is installed.
- 2 Open Regedit.
- 3 Open the following key.:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config
```

- 4 Create a new DWORD value named **FSE\_PORT**.
- 5 Right-click on the new value and click **Modify**.

- 6 In the **Value data** box, provide a port number between 1 and 65535.
- 7 Click **OK**.

**To configure a different network port for NBFSD (UNIX or Linux server)**

- 1 Log on as root on the computer where NetBackup server is installed.
- 2 Open the `bp.conf` file.
- 3 Add the following entry, where `XXXX` is an integer and is a port number between 1 and 65535.

```
FSE_PORT = XXXX
```

## Configuring the log on account for the NetBackup Client Service

By default, the NetBackup Client Service uses “Local System” as the account on which to log on. To perform operations using Granular Recovery Technology, change the service account to a domain-privileged account.

**To configure the log on account for the NetBackup Client Service**

- 1 Open the Windows Services application.
- 2 Double-click on the **NetBackup Client Service** entry.
- 3 Click on the **Log On** tab.
- 4 Provide the name of an account that has domain privileges.
- 5 Type the password.
- 6 Stop and start the NetBackup Client Service.
- 7 Close the Services control panel application.



# Index

## Symbols

- 240, 256. *See* moving a robot and its media
- .ds files 378
- .f files in catalog 601
- .SeCuRiTy.nnnn files 779

## A

- Absolute pathname
  - to directory/volume storage unit setting 380
- Accept connections on non reserved ports
  - property 187
- Access Control
  - authorizing users 748
  - host properties
    - Authentication domain tab 62
    - Authorization service tab 65
    - Networks list 60
    - Symantec Product authentication and authorization 59
  - NetBackup 42, 742
- access control lists (ACLs) 550, 779
- Access Management 42
- ACS. *See* Automated Cartridge System
- Active Directory
  - granular recovery 575–577
  - host properties 66–67
  - restoring objects 578
- Active Directory ApplicationMode (ADAM) 576–577, 580
- Activity Monitor
  - (continued)*
  - resuming suspended jobs 687
  - set column heads to view 684
  - suspending a job 686
  - topology 681
  - using the Troubleshooter 685
- Actual client property (Backup Exec Tape Reader) 70
- Actual path property (Backup Exec Tape Reader) 70
- ADAM (Active Directory ApplicationMode) 576–577, 580
- adding volumes 257
- adjust time zone 760
- administering remote systems 733
- Administration Console options 150
- administrator
  - email address property 135
  - nonroot 745
- AdvancedDisk disk storage units 374
- AFS policy type 463
- All log entries report 76, 713
- ALL\_LOCAL\_DRIVES directive 558
- Allow backups to span tape media property 156
- Allow client
  - browse property 80
  - restore property 80
- Allow media overwrite property 154
- Allow multiple data streams
  - directives 563
  - set policy attribute 486
  - when to use 487
- Allow multiple retentions per media property 156, 512
- Allow server file writes property 56, 187
- alternate media types
  - defined 261
  - example 261
- Alternate read server for storage destinations 426
- Alternate restore failover machines host
  - properties 170
- Always property in Fibre Transport host
  - properties 121
- Announce DHCP interval property 163

- ANSI format 154
- AOS/VS format 154
- API robots 289, 326, 336
- application backups 493
- archive bit 97–98, 492, 585
- archive jobs, keeping progress reports 92
- asterisk as wildcard 720
- atime 92, 498
- auth.conf file
  - capabilities identifiers 747
  - description 743
  - entries for specific applications 745
  - overview 742
- Authentication
  - NetBackup Access Control 42
  - service 693
- Authorization
  - host properties
    - Doman\Group 68
    - Group/Domain 69
    - Host 68
    - User 68
    - User must be an OS administrator 69
  - NetBackup Access Control 42
  - Service 693
- Auto log off timeout option 37, 150
- auto-discovery streaming mode 565
- Automated Cartridge System
  - drive information 232
- automatic
  - backups 493
  - cumulative incremental
    - backups 493
  - differential incremental backups 494
  - failover to an alternate server 769
  - full backups 494
  - Vault policy type 494
  - Volume Recognition service (avrd) 694
- automounted directories 478
- avrd (Automatic Volume Recognition process) 694

## B

- Back up all log files 113
- Back up only uncommitted log files 113
- Backup end notify timeout property 185
- Backup Exec QIC media 673
- Backup Exec Tape Reader
  - Exchange Server support 672

- Backup Exec Tape Reader (*continued*)
  - host properties
    - Actual client 70
    - Actual path 70
    - adding a GRFS entry 69
    - changing a GRFS entry 70
    - GRFS advertised name 69
    - removing a GRFS entry 70
  - limitations 672
  - SQL support 673
  - Windows 2003 support 672
  - Windows 2008 support 672
- Backup migrated files property 90
- Backup option for log files during full backups
  - property 113
- Backup Policy Wizard 608
- Backup start notify timeout property 183
- Backup status report 76
- backups
  - activating policy 474
  - application 493
  - archive 493
  - automatic 493
    - cumulative incremental 492–493
    - differential incremental 493–494
    - full 494
    - Vault 494
  - Client backups report 712
  - creating copies 433, 504
  - deactivating policy 474
  - duplicating 659
  - expiring 665
  - frequency
    - effect on priority 502
    - guidelines for setting 501
    - setting 500
  - full 492
  - importing 666
  - network drives 474
  - NFS mounted files 461, 476
  - off-site storage 510
  - raw partitions on Windows 469, 545
  - registry on Windows clients 547
  - selections list
    - verifying 542
  - send email notification 138
  - send email notification about 136, 139–140
  - Status of Backups report 712
  - types of 492



- backups (*continued*)
  - user directed
    - schedules 525
    - type of backup 493
  - verifying 656
  - windows
    - duration
    - See also* examples
    - specifying 517–518
- Bandwidth
  - host properties 70, 72
    - Bandwidth 72
    - From IP address 72
    - To IP address 72
  - limiting 72
- bar codes 325–326, 329, 331
- Bare Metal Restore (BMR) 100, 483, 602, 629, 683, 689
- Bare Metal Restore daemon 694
- basic disk staging
  - creating a storage unit 397
  - Final destination media owner 403
  - Final destination storage unit 403
  - Final destination volume pool 403
  - priority of duplication jobs 507
  - relocation schedule 391, 402, 491–492
  - schedule 402
  - storage units
    - size recommendations 399
    - using checkpoint restart 470
  - Use alternate read server 507
  - Use alternate read server attribute 404, 507
- BasicDisk
  - spanning storage units 414
- BasicDisk storage units 374, 423
- BE-MTF1 format 155
- BLAT mail utility 140
- Block level incremental backups 470
- BMRD (Bare Metal Restore process) 694
- BMRD (NetBackup Bare Metal Restore Master Server) 689
- BMRDB.db
  - configuration entry 637
  - in catalog 602
  - relocating 632, 639
- bp.conf file
  - configuring to use ctime 499
  - customizing jnbSA and jbpSA 755
- bp.conf file (*continued*)
  - NetBackup-Java Administration Console
    - configuration entries 748
    - when master servers share EMM database 178
  - BPARCHIVE\_POLICY 526
  - BPARCHIVE\_SCHED 526
  - bpbackup command 613
  - BPBACKUP\_POLICY 526
  - BPBACKUP\_SCHED 526
  - BPBRM logging property 146
  - bpcatarc command 619
  - bpcatlist command 619
  - bpcatres command 620
  - bpcatrm command 620
  - BPCD connect back 126
  - BPCD connect-back property 124
  - bpcd daemon 694
  - BPCD port setting on client 162
  - bpchangeprimary command 658
  - BPCOMPATD (NetBackup Compatibility Service) 690
  - bpcompatd process 695
  - bpconfig command 566
  - bpdjobs
    - adding registry key 701
    - command 703
    - debug log 703
  - BPDBJOBS\_OPTIONS environmental variable 701
  - BPDBM
    - NetBackup Database Manager, description 690
  - BPDBM logging property 146
  - bpdbm process 695
  - BPDM logging property 146
  - bpend 185
  - bpexpdate command 663
  - bpgetconfig command 55
  - BPINETD (NetBackup Client Service) 690
  - bpinetd client process 695
  - bpjava-msvc service 695
  - bpjava-susvc service 695
  - BPJAVA\_PORT 748
  - bpjobd process 695
  - BPRD
    - logging property 146
    - NetBackup Request Manager, description 691
    - port setting on client 162
    - process 695
  - bpsetconfig command 55
  - bpstart 183
  - bpsynth log 593

- BPTM logging property 146
- bpvault 146
- Browse and restore ability property 82
- buffer size 98
- Busy action property 75
- Busy file
  - host properties
    - Busy file action 75
    - File action file list 74
    - how to activate settings 75
    - Operator's email Address 74
    - Process busy files 74
    - Retry count 75
    - Working directory 73

## C

- cachefs file systems, excluding from backup 568
- Calendar schedule type 500
- calendar scheduling
  - using 521
- canceling uncompleted jobs 686
- capacity-based licenses 44
- catalog archiving
  - bpcatarc command 619
  - bpcatlist command 619
  - bpcatres command 620
  - bpcatrm command 620
  - deactivating policy for 474
- Catalog Backup Wizard 605
- catalog backups 605
  - adding critical policies to 572
  - archiving 616
    - bpcatarc 619
    - bpcatlist 619
    - bpcatres 620
    - bpcatrm 620
    - catarc policy 618
    - deactivate policy 618
    - extracting images 621
    - overview 616
    - retention level setting 618
    - type of backup indicated 618
  - compressing image catalog 625
  - image files 601
  - manual backup 612
  - Maximum concurrent jobs setting 384
  - moving client images 623
  - multiple file layout 601
  - offline, cold method 605
  - catalog backups (*continued*)
    - overview 599
    - parent and child jobs 683
    - policy type 569
    - running concurrently with other backups 134
    - schedules for 613
    - single file layout 601
    - space required 621
    - uncompressing 627
    - volume pool 467
  - catalog recovery 615
  - Catalogbackup volume pool 467
  - cdrom file system, excluding from backup 568
  - change journal 100
    - and synthetic backups 592
    - determining if enabling is useful 99
    - using in incremental backups 97
  - Change server option 732
  - changing
    - policies
      - multiple 460
    - robot properties 239
    - volume expiration date 272
    - volume group name 269–270, 279, 291
    - volume pool attributes 296
    - volume pool for a volume 270
  - changing to another server 732
  - Check the capacity of disk storage units
    - property 128, 378
  - checkpoint restart
    - and synthetic backups 592
    - Move job from incomplete state to done state
      - property 78
    - Move Restore Job from Incomplete State to Done State 471
    - multiple copies 470
    - NearStore storage units 470
    - NetWare clients 470
    - restore retries 471
    - resuming a restore job 471
    - suspending a restore job 471
    - with Basic disk staging 470
    - with Snapshot Client 470
    - with Windows clients 469
  - cipher types for NetBackup encryption 109
  - Clean-up
    - host properties 76
      - Catalog cleanup wait time 77
      - Image cleanup 77

Clean-up (*continued*)

- host properties (*continued*)
  - Keep logs 76
  - Keep true image restoration information 77
  - Keep vault logs 77
  - Move backup job from incomplete state to done state 78, 469
  - Move job from incomplete state to done state 78

## cleaning

- drives 245, 700
- frequency 231
- tape
  - change cleanings allowed 267, 272
  - set count 267

CLEANUP\_SESSION\_INTERVAL\_HOURS 438

Client administrator's email property 188

## Client Attributes

- host properties 80
  - adding a client to the client database 80
  - adding and removing clients 85
  - Allow client browse 80
  - Allow client restore 80
  - BPCD connect back 84
  - Browse and restore ability 82
  - Clients list 80
  - Connect Options tab 83
  - Daemon connection port 85
  - Free browse 82
  - Maximum data streams 81
  - Ports 84
  - Windows Open File Backup tab 85

Client backups report 712

Client cipher property 109

Client connect timeout property 183

Client name property 79

Client port window property 168–169

Client read timeout property 91, 184

Client sends mail setting 188

## clients

- BPCD port 162
- BPRD port 162
- choosing policy type 462
- database 80
- deleting from policy 461
- DHCP interval property 163
- exclude and include lists 119
- exclude file list 114–115, 117, 568
- list property 80

clients (*continued*)

- maximum jobs 133
- moving image catalog 623
- name 773
- peername 772
- setting host names 535
- clustering 57, 637
- Collect disaster recovery information for Bare Metal Restore 483
- Collect true image restoration (TIR) with move detection property 585
- Collect true image restore (TIR) information with move detection property 484
- Collect true image restore information 483
- collecting disaster recovery information 483
- column heads
  - selecting to view 684
- Communications buffer size property 98
- Compress Catalog Interval property 625
- Compress catalog interval property 135
- compression
  - by software
    - advantages 482
    - disadvantages 482
    - specifications 481
- concurrent jobs
  - on client 133
  - per policy 472
- configuring
  - drives and robots 216
  - media 257
  - storage devices 216
- Consistency check before backup host property 181
- control path
  - robotic 226
- copies
  - creating using Catalog duplicating option 659
  - creating using storage lifecycle policies 433, 504
  - third party 504
- copy
  - primary 663
- Copy on write snapshots 546
- copy window 520
- cpio format 154
- create media ID generation rules 314
- Credential Access
  - host properties 101
- credentials
  - about configuring 201

credentials (*continued*)

- about NDMP 201
- about NetBackup Deduplication Engine 201
- OpenStorage 201
- Critical Policies list 572, 608, 610
- Critical Policies list 612
- cross mount points
  - effect with UNIX raw partitions 478
  - examples 479
  - interaction with Follow NFS policy attribute 479
  - policy attribute 552
  - separate policies for 478
  - setting 478
- ctime 554
- cumulative incremental backups 492, 495
- curly brackets as wildcards 721

**D**

- Daemon connection port property 125
- Daemon port only property (for selection of ports) 125
- daemons
  - check with vmps 787
  - tlmd 697
- data
  - deduplication 482
  - movers 373, 376
  - streams 563
- Data Classification setting 417
- Data Classifications
  - creating 103
  - in storage lifecycle policies 416, 418
  - selection in policy 465, 508
- Database Administration tool 634, 646–647
- database manager process (bpdbm) 695
- database-extension clients, adding file paths for 557
- DataStore
  - volume pool 467
- DataStore policy type 463
- DataTools-SQL-BackTrack policy type 463
- datetime stamp 498
- Daylight savings time 760–761
- DB2 policy type 463–464
- DBA password, changing 638
- DBR format 155
- debug level, changing 146
- Debug logging levels for NetBackup services 146
- decommission a media server 205
- deduplication disk pool
  - configuring 363
- deduplication pool
  - configuring 363
- deduplication storage server. *See* NetBackup Deduplication Guide
  - credentials for 201
- Default cache device path for Snapshots property 92
- Default Job Priorities host properties 105, 473
- Delay on multiplexed restores property 128
- delete all devices for a media server 207
- deleting
  - drive 246
  - storage unit groups 409
  - storage units 370
  - volume pools 296
- deleting a device host 209
- Density storage unit setting 381
- denying requests 708
- description
  - for new volume 268, 271
- detailed job status 685, 694
- device
  - configuration wizard 229, 238
  - discovery 217
  - file
    - robotic 226
    - mapping file 215
- Device Configuration Wizard 368
- device host
  - for move volume 289
  - removing 209
- device management
  - remote 721
- Device Monitor 704
  - about 704
  - add drive comment 242
  - assigning requests 706
  - display pending requests 705
  - resubmit request 708
- devices
  - administer on other servers 721
- devpts file system, excluding from backup 568
- DHCP setting on client 163
- differential incremental backups 493
- Direct Access Recovery (DAR) 129
- Directory can exist on the root file system or system disk setting 380
- DISABLE\_STANDALONE\_DRIVE\_EXTENSIONS 258

- Disaster recovery
  - file, sending 607, 610, 616
  - sending emails 607, 610
- disaster recovery
  - collect information for 483
  - file
    - sending 571
  - information 136
  - sending e-mails 571
  - tab 569
- disk
  - logs report 714
  - pool status report 714
  - pools 373
  - spanning 157, 414, 466
  - staging storage units
    - selection within a storage unit group 411
  - storage unit status report 714
  - storage units 392
- Disk image backups 545
- disk pool
  - AdvancedDisk 363
  - deduplication 363
  - OpenStorage 363
- Disk pool storage unit setting 381
- disk staging 395
- Disk type storage unit setting 381
- disk-image backups 469
- Distributed Application Restore Mapping
  - host properties 107
- Do not compress files ending with property 93
- down a device 242
- drive
  - add comment 242
  - adding 227
  - cleaning 231, 245, 248
  - initial state 228, 245
  - initial status 228, 245
  - running diagnostics 251
  - servicing requests 705
  - type 230
- drives
  - cleaning 700
  - monitoring 699
  - replacing 253
  - updating firmware 255
- duplicate backups
  - becoming a primary copy 662
  - creating 659
  - duplicate backups (*continued*)
    - restoring from 657
  - duplicate window 520
  - Duplication job priority 419
  - DUPLICATION\_GROUP\_CRITERIA 439, 443
  - DUPLICATION\_SESSION\_INTERVAL\_MINUTES 439, 444
  - duration of backup window
    - examples 519
- E**
- E-mail
  - disaster recovery 571
- EFI System partitions 559
- Email
  - address for administrator of this client 188
  - send from client 188
  - send from server 188
- EMM. *See* Enterprise Media Manager
- EMM database 603
  - removing a device host from 209
  - shared 178
- empty media access port prior to update 316
- Enable block sharing storage unit setting 382
- Enable encryption property 109
- Enable granular recovery 577
- Enable granular recovery attribute 489
- Enable job logging property 157
- Enable multiplexing storage unit setting 382
- Enable performance data collection property 188, 778
- Enable robust logging property 144
- Enable SCSI reserve property 155
- Enable single instance backup for message
  - attachments property 114
- Enable standalone drive extensions property 157
- Enable standard encryption property 109
- Encryption
  - host properties
    - Client cipher 109
    - Enable encryption 109
    - Enable standard encryption 109
    - Encryption key file 110
    - Encryption libraries 110
    - Encryption permissions 109
    - Encryption strength 109
    - UseLegacy DES encryption 109
    - use with synthetic backups 590
  - encryption method for SQL Anywhere 633, 636

- end time
  - schedules 518
- Enterprise Disk license key 598
- Enterprise Disk Options 373
- Enterprise Media Manager
  - about 723
  - domain 724
  - shared 178
- Enterprise Media Manager (EMM) 176, 371, 602–603, 629, 691, 764–768
- Enterprise Media Manager server
  - about 724
  - sharing 631, 724
- Enterprise Vault Hosts properties 112
- Enterprise Vault properties 110
- erasing media 274
- errors
  - media mount 704
- escape character
  - on UNIX 721
- Exceptions to the exclude list host property 116
- Exchange granular restore proxy host property 114
- Exchange Server images, importing with BETR 672
- exclude
  - dates from schedule 521
  - files and directories from backup 114–115, 117
  - list syntax 117
- exclude file list 569
- exclude files list
  - on client 119
  - overview 568
  - Windows example 118
- Exclude list
  - host properties
    - Exceptions to the exclude list 116
    - Use case sensitive exclude list 115
- exclude\_list 569
- Expire after duplication retention type 426
- expiring backups 665
- export
  - a report 716
  - host properties 56
  - license key 46
- extended attribute files
  - disabling the restore of 555
  - Solaris 9 550

## F

- fail all copies when creating multiple copies 506

- failover
  - media server to alternate media server(s) 169
  - storage unit selection in group 410
- failover to an alternate server 769
- Fibre Transport
  - host properties
    - Always 121
    - Maximum concurrent fibre transport connections 121
    - Never 121
    - Preferred 121
    - Use defaults from the master server configuration 121
- File browse timeout property 183
- File Change Log 92
- File Change Log (FCL) 93
- file lists
  - disk image on Windows 545
  - extension clients 557
  - links on UNIX 551
  - NetWare clients
    - nontarget 555
    - target 557
  - raw partitions 545, 551
  - standard clients 549
  - UNIX files not backed up 544, 550, 567
- File system backup coverage report 542
- File System Export option 484
- files
  - .SeCuRiT.y.nnnn 779
  - catalog space requirements 621
  - excluding from backup 114–115, 117
  - linked
    - UNIX 550
  - NFS mounted 461, 476
  - No.restrictions 773
  - NOTES.INI 152
  - peername 774
  - redirected restores 775
  - restrictions on restores 772
- FilesNotToBackup list 569
- filters, applying in the Activity Monitor 684
- Final destination
  - media owner 403
  - storage unit 403
  - volume pool 403
- Firewall properties
  - example setup 126

**Firewalls**

- host properties
  - BPCD connect back 124
  - BPCD connect-back 124
  - Daemon connection port 125
  - Default connect options 122
  - Hosts list 123
  - Ports 124
  - using vnetd with 124
- first slot number
  - add volumes 267
  - for move volumes 289
- Fixed retention type 424
- FlashBackup 463, 550, 552
- FlashBackup-Windows policy type 463
- Flexible Disk Option 374–375, 381, 415
- Follow NFS 552
- Follow NFS mounts
  - notes on use
    - with cross mount points 477
    - with raw partitions 477
- Follow NFS setting 476, 479
- FORCE\_IPADDR\_LOOKUP 750
- Free browse property 82
- Frequency schedule type 500
- From IP address property 72
- frozen media 278
- full backups 492, 494, 586

**G**

- General level logging property 96
- General server
  - host properties
    - Add Media Override dialog box 130
    - Check the capacity of disk storage units 128, 378
    - Delay on multiplexed restores 128
    - Media host override 129–130
    - Must use local drive 128
    - Use direct access recovery for NDMP restores 129
- Generic policy type 463
- Global attributes
  - host properties
    - Administrator's email address 135
    - Compress catalog interval 135
    - Job retry delay 132
    - Maximum backup copies 135
    - Maximum jobs per client 133, 525

**Global attributes** *(continued)*

- host properties *(continued)*
  - Maximum vault jobs 135
  - Policy update interval 133
  - Schedule backup attempts 132, 534
- Global logging level property 145
- Go into effect at Policy attribute 474
- granular recovery 489
- granular recovery of Active Directory objects 575
- Granular Recovery Technology (GRT) 129
- GRFS advertised name property 69
- Group Policy Objects 581

**H**

- hard links
  - NTFS volumes 547
  - UNIX directories 551
- High water mark storage unit setting 382
- HKEYS
  - backing up 547
- host
  - device 35
  - properties
    - changing in a clustered environment 57
    - exporting 56
    - permission to change 55
- host credentials. *See* credentials

**I**

- image catalog file, compressing 135
- IMAGE\_EXTENDED\_RETRY\_PERIOD\_IN\_HOURS 439, 444
- images
  - changing primary copy 657
  - duplicating 659
  - moving client catalog 623
  - on disk report 714
  - on media report 713
  - on tape report 713
  - restoring from duplicate 657
  - verifying 656
- importing backups 666
- inactive media 714
- include
  - files list 568
  - list, on client 119
- include file list 569
- include\_list 569

- Incrementals based on
  - archive bit property 97
  - timestamp property 97
- Informix extension (license) 463
- Informix policy type 463
- INI file, for Lotus Notes 152
- Initial browse search limit property 187
- INITIAL\_BROWSE\_SEARCH\_LIMIT 755
- INITIAL\_MEMORY 752, 757
- inject volume into robot
  - multiple volumes 316
  - robot inventory 279
- Inline copy option 503, 660, 665
- Instant Recovery
  - Advanced backup method 470
  - Backups to disk only setting 502
- Intelligent Disaster Recovery (IDR) 483
- Internet Assigned Numbers Authority (IANA) 166, 748
- inventory and compare robot contents 307

**J**

- Java
  - auth.conf file 743
  - authorizing users 742
  - directory 745
  - jbpSA configuration options 755
  - performance improvement hints 757
  - Virtual Machine (JVM) 752
- Java Windows Administration Console 729, 732–733, 736, 739, 747–748
  - improving performance 756–757
  - installing 735
- jbpSA
  - customizing 755
  - logging 755
- jnbSA 739
  - customizing 755
  - logging 755
- Job Manager logging property 146
- Job retry delay property 132
- jobs
  - Concurrent per disk storage unit 383
  - filters
    - specifying 684
  - maximum
    - per client 133
    - per policy 471
  - priority for policy 473

- jobs (*continued*)
  - setting default priority 105
  - SLP\_MultipleLifecycles 442
  - viewing in the Activity Monitor 682
- JVM (Java Virtual Machine) 752

## K

- Keep logs property 76
- Keep status of user-directed backups
  - archives
    - and restores property 90, 92, 99
- Keep true image restoration information property 77
- Keep vault logs property 77
- KEEP\_LOGS\_DAYS 755
- KeysNotToRestore list 569
- keyword phrase 489
- killing jobs 686

## L

- label
  - media tapes 284
  - new media 285
- legacy logging 145
- library sharing 220
- license keys
  - accessing 44
  - adding 43, 45
  - deleting 46
  - export 46
  - printing 45
  - viewing the properties of one key 46
- LIFECYCLE\_PARAMETERS
  - CLEANUP\_SESSION\_INTERVAL\_HOURS 438
  - DUPLICATION\_GROUP\_CRITERIA 439
  - IMAGE\_EXTENDED\_RETRY\_PERIOD\_IN\_HOURS 439
  - LIFECYCLE\_PARAMETERS 439
  - MAX\_GB\_SIZE\_PER\_DUPLICATION\_JOB 440
  - MAX\_MINUTES\_TIL\_FORCE\_SMALL\_DUPLICATION\_JOB 440
  - MIN\_GB\_SIZE\_PER\_DUPLICATION\_JOB 440
  - TAPE\_RESOURCE\_MULTIPLIER 441
  - VERSION\_CLEANUP\_DELAY\_HOURS 441
- LIFECYCLE\_PARAMETERS file 438
- Limit jobs per policy setting 471, 516, 525
- links
  - UNIX hard-linked directories 551
  - UNIX symbolic 550
- load balancing methods 412
- Locked file action property 91



- log off NetBackup 150
  - logging
    - bpsynth 593
    - deleting after a set time 76
    - jbpSA 755
    - jnbSA 755
    - legacy 143
    - unified 142
  - Logging host properties
    - Debug logging levels for NetBackup services 146
    - Enable robust logging 144
    - Global logging level 145
    - NetBackup logging types 142
    - Process specific overrides property 145
  - Login Banner Configuration host properties 147
  - login banner text, removing 150
  - long erase 275
  - Lotus Notes
    - host properties
      - INI file 152
      - Maximum number of logs to restore 151
      - Path 153
      - Transaction log cache path 152
    - policy type 463
    - properties 151
  - Lotus Notes extension (license) 463
  - Low water mark storage unit setting 380, 382
  - ltid (NetBackup Device Manager) 691, 695
- M**
- Mac OS X 464
  - mail notifications
    - administrator email address 188
    - Disaster Recovery attachment
      - sending 571
    - email address for administrator 135
    - Windows nbmail.cmd script 135
  - mail\_dr\_info.cmd 616
  - Mailbox for message level backup and restore
    - property 114
  - manual backups
    - NetBackup catalogs 612
    - policy for 575
  - master servers
    - rebooting 787
  - MAX\_GB\_SIZE\_PER\_DUPLICATION\_JOB 440
  - MAX\_MEMORY 752, 757
  - MAX\_MINUTES\_TIL\_FORCE\_SMALL\_DUPLICATION\_JOB 440
  - maximum
    - concurrent FT connections property 121
    - concurrent jobs storage unit setting 383
    - concurrent write drives storage unit setting 503
    - data streams property 81
    - error messages for server property 99
    - jobs per client 133
    - jobs per policy 472
    - vault jobs property 135
  - Maximum backup copies property 135
  - maximum bar code lengths 326
  - Maximum concurrent write drives setting 383
  - Maximum number of logs to restore property 151
  - Maximum streams per drive storage unit setting 385, 513
  - Media
    - host properties 153
  - media
    - active 714
    - formats 261
    - freeze 278
    - frozen 278
    - host override property 129
    - host properties
      - Allow backups to span disk 157
      - Allow backups to span tape media 156
      - Allow media overwrite 154
      - Allow multiple retentions per media 156
      - Enable job logging 157
      - Enable SCSI reserve/release 155–156
      - Enable standalone drive extensions 157
      - Media ID prefix (non-robotic) 158
      - Media request delay 159
      - Media unmount delay 158
    - ID generation rules 329
    - ID prefix (non-robotic) property 158
    - inactive 714
    - log entries report 76, 713
    - mount
      - errors. *See* canceled
      - errors, queued 704
      - mount timeout property 184
      - pools (see volume pools) 292
      - request delay property 159
      - server connect timeout property 185
      - server register 206
      - suspend 291
      - type when not an API robot 321
      - unfreeze 278

- media (*continued*)
  - unmount delay property 158
  - unsuspend 291
- media ejection timeout period 282
- media ID
  - prefix for update robot 320
- media server
  - activate 203
  - deactivate 203
  - decommission 205
  - delete all devices from 207
- Media server copy advanced backup method 470
- Media server load balancing storage unit selection
  - in group 411–412
- Media server storage unit setting 385
- media servers
  - adding a media server to the Alternate restore failover machine list 171
  - moving a robot and its media 240
  - rebooting 787
  - registering with the EMM server 722
  - Restore failover host properties 169
- media sharing
  - about 297
  - configuring 298
  - configuring unrestricted 298
  - configuring with a server group 299
- media type
  - 4MM 260
  - 4MM\_CLN 260
  - 8MM 260
  - 8MM2 260
  - 8MM2\_CLN 260
  - 8MM3 260
  - 8MM3\_CLN 260
  - 8MM\_CLN 260
  - DLT 260
  - DLT2 260
  - DLT2\_CLN 260
  - DLT3 260
  - DLT3\_CLN 260
  - DLT\_CLN 260
  - DTF 260
  - DTF\_CLN 260
  - HC2\_CLN 260
  - HC3\_CLN 260
  - HC\_CLN 260
  - HCART 260
  - HCART2 260
- media type (*continued*)
  - HCART3 260
  - QCART 260
- Megabytes of memory property 92
- MEM\_USE\_WARNING 753
- Microsoft Volume Shadow Copy Service (VSS) 67, 87
- Microsoft Windows Backup 569
- MIN\_GB\_SIZE\_PER\_DUPLICATION\_JOB 440
- mirrored transaction log, creating 640
- mixing retention levels on tape volumes 512
- mklogdir.bat 143
- mntfs file system, excluding from backup 568
- monitoring
  - NetBackup drives 699
  - NetBackup processes 698
- monthly backups, scheduling 524
- mount
  - points 478
  - requests, pending 705
- move
  - backup job from incomplete state to done state
    - property 78, 469
  - detection 484
  - job from incomplete state to done state
    - property 78
  - Restore Job from Incomplete State to Done State 471
  - restore job from incomplete state to done state
    - property 78
  - volumes
    - logical move 286
    - overview 286
    - physical move 286
    - update volume configuration 287
- moving NBDB database files 639
- MS Exchange extension (license) 463
- MS SQL Server extension (license) 463
- MS-Exchange-Server policy type 463
- MS-SharePoint policy type 463
- MS-SQL-Server policy type 463
- MS-Windows policy type 464
- MTF format 155
- mtime 554
- multiple copies
  - checkpoint restart 470
  - creating using a policy schedule 504
  - creating using storage lifecycle policies 433, 504
  - criteria for creating 503
  - fail all copies 506

- multiple copies *(continued)*
  - parent and child jobs 683
  - setting 503
- Multiple copy synthetic backups method 594
- multiple data streams 563
  - allowing 487
  - parent and child jobs 683
- multiple file layout for NetBackup catalogs 601
- multiplexing (MPX)
  - and synthetic backups 589
  - demultiplexing 517
  - Maximum jobs per client property 516
  - preserving 427
  - set for schedule 512
  - use with Enable block sharing 382
- multistreaming and synthetic backups 589
- Must use local drive property 128
  
- N**
- named data streams
  - disabling the restore of 555
- naming conventions 719
- nb\_updatedssu script 378
- NBDB.db
  - configuration entry 637
  - creating manually 640
  - in catalog 602
  - installation overview 630
  - moving from one host to another 648
  - relocating 631, 639
- NbDbAdmin.exe (Database Administration tool) 647
- NBEMM (NetBackup Enterprise Media Manager) 691, 695
- nbemmcmd command 179
- nbEvtMgr process 695
- nbfsd port 804
- nbftsvr process 696
- nbj.conf 748
- NBJAVA\_CLIENT\_PORT\_WINDOW 753
- NBJAVA\_CONNECT\_OPTION 754
- NBJAVA\_CORBA\_DEFAULT\_TIMEOUT 754
- NBJAVA\_CORBA\_LONG\_TIMEOUT 754
- NBJM (NetBackup Job Manager) 146, 696
- nbmail.cmd script 135, 616
- NBPEM (NetBackup Policy Execution Manager) 146, 691, 696
- nbproxy process 696
- NBRB (NetBackup Resource Broker) 147, 692
- nbrb process 696
- NBRB\_CLEANUP\_OBSOLETE\_DBINFO 194
- NBRB\_ENABLE\_OPTIMIZATIONS 194
- NBRB\_FORCE\_FULL\_EVAL 195
- NBRB\_MPX\_GROUP\_UNLOAD\_DELAY 195
- NBRB\_REEVAL\_PENDING 195
- NBRB\_REEVAL\_PERIOD 195
- NBRB\_RETRY\_DELAY\_AFTER\_EMM\_ERR 195
- NBRMMS (NetBackup Remote Management and Monitor Service) 691, 696
- NBSL (NetBackup Service Layer) 692
- nbsl process 696
- nbstlutil (lifecycle utility) command 443
- nbstsvr process 696
- nbsvcmon process 696
- NBU-Catalog policy type 467, 569
- NBVAULT (NetBackup Vault Manager) 692, 696
- NCR-Teradata policy type 464
- NDMP
  - clients list 101
  - credentials for 41, 201
  - Direct Access Recovery for restores 129
  - drives 128
  - global credentials 160
  - host storage unit setting 387
  - hosts 160, 216
  - policy type 464
  - storage units 378, 392, 504
- NearStore storage units 361, 374, 388, 409, 426, 470, 484
- NetApp 374
- NetBackup
  - client service 162
  - request service port (BPRD) 162
  - volume pools 296
- NetBackup Access Control (NBAC) 42, 742
- NetBackup Client Service (BPINETD) 690
- NetBackup Client Service log on account, configuring 805
- NetBackup Compatibility Service (BPCOMPATD) 690
- NetBackup Database Manager (BPDBM) 690
- NetBackup Device Manager 256, 691
- NetBackup for MS-Exchange 557
- NetBackup Job Manager (NBJM) 146, 691
- NetBackup media kit 35
- NetBackup Monitor Service 692
- NetBackup Policy Execution Manager (NBPEM) 146, 691
- NetBackup Remote Administration Console 732, 734

- NetBackup Remote Management and Monitor Service (NBRMMS) 691
- NetBackup Request Manager (BPRD) 691
- NetBackup Request Service Port (BPRD) property 162
- NetBackup Resource Broker (NBRB) 106, 147, 692
- NetBackup Service Layer (NBSL) 692, 696
- NetBackup Storage Lifecycle Manager 692
- NetBackup support Web site 216
- NetBackup Vault Manager (NBVAULT) 692
- NetBackup Volume Manager (VMD) 692
- NetBackup volume pool 467
- NetBackup-Java 742
- NetBackup-Java Administration Console
  - improving performance 756
- NetBackup-Java Version 7.0 732–733
- NetWare client
  - host properties 90
    - Back up migrated files 90
    - Keep status of user-directed backups
      - See also* and restores
      - See also* archives
    - Uncompress files before backing up 90
  - target and nontarget 115
- NetWare policy type 464
- network
  - drives, backing up 474
  - host properties
    - Announce DHCP interval 163
    - NetBackup client service port (BPCD) 162
    - NetBackup request service port (BPRD) 162
- Network Attached Storage (NAS) 376, 379
- Network Setting host properties 163–165
- Never property in Fibre Transport host
  - properties 121
- NEW\_STREAM
  - file list directive 563
- NFS (Network File System)
  - Follow NFS policy attribute 476, 479
  - NFS access timeout property 190
  - no disk spanning 157
- non reserved ports 187
- none of the files in the file list exist (NetBackup status message) 537
- None volume pool 467
- nonroot administration for specific applications 745

## O

- offline, cold catalog backup method 605
- On demand only storage unit setting 387, 414

- open schedules 530
- OpenStorage. *See* NetBackup Shared Storage Guide
  - storage server. *See* NetBackup Shared Storage Guide
- OpenStorage Disk Option 374–375, 381, 415
- OpenStorage disk storage unit 375
- OpenStorage optimized synthetic backup method 598
- operating mode of tape drive
  - changing 243
- Operator's email address property 74
- OpsCenter 692
- optical devices
  - support in NetBackup 7.0 214
- Oracle extension (license) 464
- Oracle policy type 464
- Oracle\_RMAN 540
- Oracle\_XML\_Export 540
- OS/2 policy type 464
- Override default job priority
  - for Catalog jobs 105, 656
  - for Media Contents report 105
  - for Media contents report 713
  - for queued or active jobs 688
- Override policy
  - storage selection setting 508
  - volume pool setting 509
- Overwrite existing files 555

## P

- pagefile.sys 546
- parent jobs 486, 682
  - in Activity Monitor Jobs tab 682
  - Limit jobs per policy setting 472
  - parent\_end\_notify script 683
  - parent\_start\_notify script 683
- parent\_end\_notify script 683
- parent\_start\_notify script 683
- password, changing 638
- path
  - separators 380
- PBX (Symantec Private Branch Exchange) 698
- PC NetLink files 550
- peername
  - files 774
  - of client 772
- pending actions
  - overview 705
  - resolving 707

- pending requests
    - resolving 706
    - resubmitting 708
  - Perform
    - default search for restore property 99
  - Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS) property 114
  - Perform default search for restore property 99
  - Perform incrementals based on archive bit 497
  - permissions
    - to change NetBackup properties 55
  - physical inventory utility 343
  - policies
    - activating 474
    - changing properties 456–458, 461
    - configuration wizard 455
    - creating policy for Vault 573
    - deactivating 474
    - for Active Directory granular restores 576
    - overview 448
    - planning 449
    - setting priority 105, 473
    - user schedules 525
    - volume pool setting 467
  - Policy Execution Manager
    - Logging property 146
  - Policy storage policy attribute 465, 611
  - policy type
    - attribute 462
    - MS-Windows 464
    - NBU-Catalog 464
    - NCR-Teradata 464
    - OS/2 464
    - Vault Catalog Backup 494
  - Policy update interval property 133, 527, 655
  - Port Ranges
    - host properties
      - Client port window 168–169
      - Server port window 168
      - Use OS selected non reserved port 168–169
      - Use random port Assignments 167
  - ports
    - allow operating system to select non reserved port 168–169
    - non reserved 187
  - power down NetBackup servers 785
  - Preferred property in Fibre Transport host properties 121
  - prelabel media 284–285
  - preprocess interval 565
  - Preserve multiplexing 427
  - preview volume configuration update 303
  - primary copy
    - becoming a 662
    - changing 657
    - definition 663
    - promoting to 658
  - print
    - job detail information 687
    - job list information 687
    - license key 45
    - report 716
  - Prioritized storage unit selection in group 410
  - priority
    - of a job 105, 473
    - of duplication jobs 507
    - of relocation jobs started from this schedule setting 402
  - Private Branch Exchange 693, 698
  - Problems report 76, 713
  - proc file system
    - excluding from backups 568
  - Process busy files property 74
  - processes
    - check with vmps 787
    - monitoring 698
  - properties
    - changing on multiple hosts 57
    - exporting 56
    - overview 55
    - viewing 55
  - PureDisk
    - PureDisk-Export policy type 464
    - Storage Option 375, 381
    - Storage Pool Authority (SPA) 381
    - storage units 426
- ## Q
- question mark as wildcard 720
  - quick erase 275
  - quotas on file systems 373
- ## R
- random ports, setting on server 167
  - raw partitions
    - backing up 469, 492, 545

- raw partitions (*continued*)
  - backups on UNIX 551–553
  - Follow NFS policy attribute 477
  - restoring 546
- rebooting NetBackup servers 787
- recommended method of configuring devices 216
- redirected restores 552, 772
- Reduce fragment size storage unit setting 390
- register a media server 206
- registry
  - backup/restore 547
- reload.sql 645–646
- relocation schedule 403, 491–492, 500, 506
  - initiating manually 405
- remote
  - access, allowing 728–729
  - device management 721
  - systems
    - administering 733
- Remote Administration Console 732, 734
- removing a device host 209
- replacing a drive 253
- reports
  - All log entries report 713
  - Client backups report 712
  - copying from 716
  - description of utility in Administration
    - Console 709
  - Disk logs report 714
  - Disk pool status report 714
  - Disk storage unit status report 714
  - Images on disk report 714
  - Images on media report 713
  - Images on tape report 713
  - Media log entries report 713
  - printing 716
  - Problems report 713
  - running a report 715
  - saving 716
  - settings for building a report 712
  - Status of backups report 712
  - Tape contents report 105, 713
  - Tape lists report 714
  - Tape logs report 713
  - Tape summary report 714
  - Tape written report 714
  - using the Troubleshooter 715
- requests
  - assigning 706
- requests (*continued*)
  - denying 708
  - display pending 705
  - overview 705
- REQUIRED\_INTERFACE 196
- reset
  - file access time property 92
  - mount time 248
- residence, updating volume configuration 305
- Resource Broker logging property 147
- restarting jobs 686
- Restore Failover
  - host properties
    - Alternate restore failover machines list 170
- Restore job
  - resuming 471
  - suspending 471
- Restore retries
  - interaction with checkpoint restart 471
  - property 186
- restores
  - adjust time zone for 760
  - alternate server 763
  - directed from the server 771
  - from a specific backup copy 763
  - from a specific backup copy 508
  - keeping progress reports 92
  - raw partition 546
  - redirected 169, 772–773
  - reducing search time 625
  - registry on Windows clients 547
  - server independent 763
  - symbolic links on UNIX 550
  - System State 781
  - using a specific server 130
- resuming suspended jobs 687
- retention levels
  - default 512
  - for archiving catalogs 618
- retention periods
  - caution for setting 525
  - changing 173
  - expiration 525
  - guidelines for setting 510
  - lifecycle and policy-based 422
  - mixing on tape volumes 156, 512
  - precautions for setting 511
  - redefining 172
  - setting 510

- retention periods (*continued*)
    - user schedule 525
  - Retention types for storage lifecycle policies
    - Expire after duplication 426
    - Fixed 424
    - Staged capacity managed 425
  - retire a media server. *See* decommission a media server
  - Retries allowed after runday policy setting 500
  - Retry count property 75
  - retry restores, setting 186
  - Reverse Host Name Lookup host property 163–165
  - REVERSE\_NAME\_LOOKUP entry 165
  - robot
    - adding 221
    - compare contents 307
    - control host 226
    - destination for move volume 289
    - device file 226
    - device host 223
    - for new volume 268, 289
    - inventory 302, 307
    - moving to new media server 240
    - number 223
    - number storage unit setting 391
    - running diagnostics 249
    - type 223
    - type storage unit setting 391
  - robot type
    - ACS 214
    - TL4 214
    - TL8 214
    - TLD 214
    - TLH 214
    - TLM 215
  - Round robin storage unit selection in group 410
  - RS-MTF1 format 155
- S**
- SAP extension (license) 464
  - SAP policy type 464
  - save a report 716
  - Schedule backup attempts property 132, 487, 534
  - schedules
    - backups on specific dates 522
    - considerations 526
    - duplicating 520
    - excluding dates 521
    - frequency 500
  - schedules (*continued*)
    - monthly backups 524
    - naming 491
    - not combining calendar-based and frequency-based 501
    - overview 490
    - priority 502
    - recalculating 527
    - retention level defaults 512
    - retention periods
      - guidelines 510
      - setting 510
    - setting backup times 517–518
    - specify multiplexing 512
    - storage unit/storage lifecycle policy 508
    - type of backup 492
    - user backup or archive 525
    - volume pool 509
    - windows that span midnight 529
  - schedules, creating weekly backups 523
  - scratch
    - pool and WORM media 263
    - pool, adding 295
    - volume pool 467
  - scratch pool
    - description 292
  - scripts 683
    - bpdbjobs example 702
    - vmps 787
  - SCSI
    - Long Erase 275
    - pass-through command 217
    - persistent reserve 156
      - drive path override 236
    - Quick Erase 275
    - reserve, configuring 155
    - reserve/release
      - drive path override 236
  - SeCuRiT<sub>y</sub>.nnnn files 779
  - SERVER
    - vm.conf entry 722
  - server
    - directed restores 747
      - allowing access 728–729
    - alternate server restores 763
    - directed restore 771
    - EMM server 603
    - host properties 175
      - media servers 176

- server (*continued*)
  - host properties (*continued*)
    - using 728–729
  - independent restores 169, 763
  - list definition 175
  - list, adding a server 728–729
  - port window property 168
  - power down 785
  - rebooting 785
  - sends mail property 188
- server group, configuring 198
- Service Manager 637
- Services for NFS
  - installing on Windows 2003 R2 SP2 801
- setconf.bat file 748
- Shadow Copy Components 672
- Shadow Copy Components directive 559
- Shadow Copy Service 67, 87
- shared drives
  - configuration wizards 216
  - drive operating mode 244, 246
- shared tape drives
  - operating mode 243
- SharedDisk storage units 375
- SharePoint 2003 684
- SharePoint Server 180
- show robot contents 307
- shut down NetBackup services 786
- single file
  - layout for NetBackup catalogs 601
  - restore program
    - FlashBackup 550
- Single-Instance Storage (SIS) 114, 426, 469, 482
- slot number
  - add volume 267
  - for move volumes 289
- SLP\_MultipleLifecycles.job 442
- Snapshot Client 87, 191, 376, 381, 463, 470, 490, 502, 557, 683
- Snapshot verification I/O throttle property 113
- SnapVault storage units 376, 388, 392, 409, 423–424
- Solaris 9 extended attributes 550
- SPC-2 SCSI reserve 156
- SQL Anywhere
  - encryption method 633, 636
  - in NetBackup installation 602
- SQL images, importing with BETR 673
- SQL-BackTrack extension (license) 463
- SQLANYs\_VERITAS\_NB 637, 691
- square brackets as wildcards 720
- Staged capacity managed retention type 425
- staging
  - backups 395
  - schedule storage unit setting 391
  - using BasicDisk storage unit 380
  - using Storage Lifecycle Policies 415
- Standard policy type 464
- start time
  - for schedules 518
- start up NetBackup services 786
- Start Window tab 517
- startup text, removing 150
- status
  - drive initial configuration 228, 245
- status codes
  - NetBackup
    - 71 537
- Status of backups report 712
- Storage device storage unit setting 392
- Storage Lifecycle Manager service (nbstserv) 442
- Storage Lifecycle Policies
  - Alternate read server for destination 426
  - and the Multiple copies configuration dialog 509
  - copy number 433
  - Data classification setting 417
  - data classifications 418
  - deleting 419
  - Duplication job priority setting 419
  - Expire after duplication retention type 426
  - Fixed retention type 424
  - hierarchy 427, 429–431
  - Media owner for destination 424
  - optional LIFECYCLE\_PARAMETERS
    - configuration 438
  - Preserve multiplexing for destination 427
  - retention type 422, 510
  - Staged capacity managed retention type 425
  - storage destination list requirements 423
  - storage destinations 421, 423
  - Storage lifecycle policy name 417
  - Storage unit for destination 423
  - using nbstlutil to administrate lifecycle
    - operations 443
    - utility 415
    - versions of 435, 437–438
    - volume pool for destination 424
    - writing multiple copies 433



- storage server
  - AdvancedDisk. *See* NetBackup Shared Storage Guide
  - credentials for deduplication 201
  - deduplication. *See* NetBackup Deduplication Guide
  - OpenStorage. *See* NetBackup Shared Storage Guide
- storage servers 373
- storage unit
  - groups 407, 409
  - name setting 392
  - selection within a storage unit group 409, 413
  - type setting 392
- storage units
  - AdvancedDisk disk type 374
  - available storage property of volume 389
  - BasicDisk type 374
  - capacity property of volume 389
  - changing server to manage 727
  - creating 368–369
  - creating a basic disk staging unit 397
  - creation overview 366
  - deleting 370
  - disk pool comment property 389
  - disk storage units 372
  - for policy 465
  - for schedule 508
  - high water mark property of volume 389
  - low water mark property of volume 389
  - maintaining space on disk storage units 377
  - Media Manager type 370
  - name property 390
  - naming conventions 719
  - NDMP disk type 378
  - NearStore disk type 374, 388, 409
  - number of volumes property 390
  - OpenStorage disk type 375
  - percent full property on volume 390
  - PureDisk disk type 375, 409
  - QIC drive type 504
  - raw size property on volume 390
  - SharedDisk disk type 375
  - SnapVault disk type 376, 388, 409
  - storage lifecycle policies 376
  - usable size property of volume 390
  - vendor-specific 375
- subnets 72
- Sun PC NetLink 550
- suspend
  - backups and restores 687
  - suspended jobs 78, 686
  - Sybase extension (license) 464
  - Sybase policy type 464
  - Sybase SQL Anywhere
    - dbsrv11.exe 691
    - default password 636
    - management of 637
    - starting/stopping the service 637
    - use in NetBackup 629
  - Symantec OpsCenter 633, 636, 692
  - Symantec Private Branch Exchange 693, 698
  - Symantec Product Authentication and Authorization 42, 693
  - Symantec products properties 181
  - Symantec support Web site 216
  - symbolic links
    - included in backup selection list 542
    - UNIX 550
  - synthetic backups
    - and encryption 590
    - checkpoint restart 470
    - component images 586
    - cumulative incremental 587
    - full 586
    - logs produced during 593
    - multiple copy backups method 594
    - no multiple copy support 504
    - no NetBackup change journal support 100
    - OpenStorage optimized method 598
    - recommendations for running 589
    - schedules 499
  - System State
    - backups 469
    - directive 558
    - restoring 781
- T**
- Take checkpoints every 468
- tape
  - assigning requests 706
  - contents report 713
  - lists report 714
  - logs report 713
  - Media contents report 713
  - summary report 714
  - written report 714

## tape drive

- changing operating mode 243

TAPE\_RESOURCE\_MULTIPLIER 441

TapeAlert 231

tar format 155

TCDebug\_TCPP level logging property 96

temporary staging area 383, 391–392, 506

third-party copies 504

Third-Party Copy Device Advanced Backup method 470

Time overlap property 98

time zones

- adjustment for restores 760

- setting Daylight savings time 760–761

Timeouts

- host properties

- Backup end notify timeout 185

- Backup start notify timeout 183

- Client connect timeout 183

- Client read timeout 184

- File browse timeout 183

- Media mount timeout 184

- Media server connect timeout 185

- Use OS dependent timeouts 184

tlmd daemon 697

tmpfs file system, excluding from backup 568

To IP address property 72

tpext utility 641

Transaction log cache path property 152

transaction log, creating 640

Transfer throttle storage unit setting 392

traversing directories to back up a file 119

Troubleshooter

- using in Activity Monitor 685

Troubleshooter, using in Reports application 715

True Image Restoration (TIR)

- configuration 483

- Error code 136 592

- length of time to keep information 77

- move detection 484

- no NetBackup change journal support 100

- pruning information 592

- with Move Detection 592

- with move detection 100

Truncate log after successful Instant Recovery backup property 114

**U**

UNC path 469

## uncompress

- files before backing up property 90

- NetBackup catalogs 627

unified logging 142, 146

Universal

- host properties

- Accept connections on non reserved ports 187

- Allow server file writes 187

- Browse timeframe for restores 187

- Client administrator's email 188

- Client sends mail 188

- Enable performance data collection 188

- Last full backup 187

- Restore retries 186

- Server sends mail 188

- Use specified network interface 187

- settings properties 185

UNIX client

- host properties 91

- Add to all 93

- Do not compress files ending with 93

- Do not reset file access time 92

- Keep status of user-directed backups 92

- Megabytes of memory 92

- primary node in tree 190

- Use VxFS file change log for Incremental backups property 92–93

UNIX server properties 190

UNSET file list directive 567

UNSET\_ALL file list directive 567

unsupported characters 313

update

- robot procedure 313

- volume configuration 304, 307

updating drive firmware 255

usbdevfs file system, excluding from backup 568

Use alternate read server 404, 507

Use case sensitive exclude list host property 115

Use change journal in incrementals property 97

Use defaults from the master server configuration property 121

Use Direct Access Recovery for NDMP restores property 129

Use legacy DES encryption property 109

Use non reserved ports property 124

Use OS dependent timeouts property 184

Use random port assignments properties 167

Use reserved ports property 124

- Use specified network interface property 187
- Use VxFS file change log for Incremental backups property 92
- user
  - archive backups 493
  - backups 493
  - schedules, planning 525
- User directed timeouts property 98

## V

### Vault

- backup type 494
- catalog archiving 620
- designating duplicate as the primary 657
- license 464
- Logging property 146
- Maximum vault host property 135
- parent and child jobs 684
- policy
  - creating 573
  - type 464
- vendor-specific storage units 375
- verifying backup
  - images 656
  - selections list 542
- Veritas Volume Manager (VxVM) 552
- Veritas Volume Snapshot Provider 87, 192
- veritas\_pbx (Symantec Private Branch Exchange) 698
- VERSION\_CLEANUP\_DELAY\_HOURS 441
- view properties of a license key 46
- vm.conf file, adding SERVER entries 722
- VMD (NetBackup Volume Manager) 692
- vmd process 697
- vmops script 787
- VMware backup hosts host properties 191
- vnetd
  - enabling logging for 126
  - example setup 126
  - Only property (for selection of ports) 125
  - setting up between a server and a client 125
  - setting up between two servers 126
  - Veritas Network Daemon 124
- VNETD\_PORT 748
- volume groups
  - about 296
  - changing name 269–270, 279, 291
  - for move volume 289
  - rules for assigning 297

### volume pools

- about 292
- add volume 269
- adding 293
- and WORM media 263
- CatalogBackup 467
- changing attributes 296
- changing for a volume 270
- configuring 292
- DataStore 467
- DataStore pool 269, 272
- deleting 296
- for schedule 509
- indicating one for use by a policy 467
- NetBackup 467
- None 467
- overview 292
- properties 293
- scratch 467
- Volume Shadow Copy Service (VSS) 67, 87, 558–559
- Volume Snapshot Provider (VSP) 87, 192
- volumes
  - adding 257
  - allocation 467
  - assignments 467
  - cleaning count 267, 272
  - description for new volume 268, 271
  - moving 286
    - actions menu 287
    - scratch 467
- VRTSAt (Symantec Product Authentication Service) 693
- VRTSaz (Symantec Product Authorization Service) 693
- VRTSspb (Symantec Private Branch Exchange) 693
- VXDBMS\_NB\_DATA registry entry 637
- VxFS
  - file change log 92
  - named data streams 554
- vxlogcfg command 142, 146
- vxlogmgr command 142

## W

- Wait time before clearing archive bit property 97–98
- weekly backups
  - scheduling 523

wildcard characters 720

UNIX

escape character 721

file paths 549

Windows clients 544

windows 529

*See also* see schedules

Windows client

host properties

Communications buffer size 98

General level logging 96

Incrementals based on archive bit 97

Incrementals based on timestamp 97

Keep status of user-directed backups

*See also* and restores

*See also* archives

Maximum error messages for server 99

Perform default search for restore 99

TCP level logging 96

Timeout overlap 98

Use change journal in incrementals 97

User directed timeouts 98

Wait time before clearing archive bit 97–  
98

Windows Disk-Image (raw) backups 469, 545

Windows Open File Backups

host properties

Abort backup on error 89

Disable snapshot and continue 90

Enable Windows Open File backups for this  
client 87

Global drive snapshot 89

Individual drive snapshot 88

Use Microsoft Volume Shadow Copy Service  
(VSS) 87

Use Veritas Volume Snapshot Provider  
(VSP) 87

Windows Service Manager 637

wizards

backup policy 455

Device Configuration 368

device configuration 238

shared drive configuration 229

Working directory property 73

WORM media

about 261

and drive types 264

and media types 264

and Quantum drives 264

WORM media (*continued*)

and scratch pool 263

and volume pools 263

limitations 262

supported drives 262