

Symantec NetBackup™ Administrator's Guide, Volume I

Windows

Release 7.1

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: 7.1

PN: 21159721

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo 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 support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

www.symantec.com/business/support/

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/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support 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

Contents

Technical Support	4
Section 1 About NetBackup	29
Chapter 1 Introducing NetBackup interfaces	31
About NetBackup	31
Online documents	33
About NetBackup administration interfaces	33
About running the Windows-based NetBackup Administration Console	34
Starting the Java-based Windows Display Console	35
About administering remote servers	36
About using the NetBackup Administration Console	36
Standard and user toolbars	37
About customizing the administration console	38
NetBackup configuration wizards	38
Activity Monitor utility	39
NetBackup Management utilities	39
Media and Device Management utilities	41
Running the Troubleshooter	42
Access Management utility	43
Chapter 2 Administering NetBackup licenses	45
About administering NetBackup licenses	45
Accessing license keys for a NetBackup server	46
Adding new license keys	47
Printing license key lists	48
Deleting license keys	48
Viewing license key properties	49
Exporting license keys	49

Section 2	Configuring hosts	51
Chapter 3	Configuring Host Properties	53
	NetBackup Host Properties configuration methods	55
	About the Host Properties	55
	Viewing host properties	57
	Changing the host properties on multiple hosts at the same time	57
	Property states for multiple hosts	58
	Exporting host properties	60
	Standard host property dialog box options	60
	Access Control properties	61
	Authentication Domain tab	62
	Authorization Service tab	63
	Network Attributes tab	64
	Active Directory host properties	66
	Backup Exec Tape Reader properties	68
	Bandwidth properties	70
	Bandwidth limit usage considerations and restrictions	71
	Add Bandwidth Settings dialog box for Bandwidth properties	72
	Busy File Settings properties	72
	Activating the Busy File Settings in host properties	74
	Clean-up properties	75
	Client Name properties	77
	Client Attributes properties	78
	Add Client dialog box	80
	General tab of the Client Attributes properties	80
	Connect Options tab of the Client Attributes properties	84
	Windows Open File Backup tab of the Client Attributes properties	86
	Backlevel and upgraded clients that use Windows Open File Backup	90
	Client Settings properties for NetWare clients	91
	Client Settings (UNIX) properties	92
	VxFS file change log for incremental backups property	94
	Client Settings properties for Windows clients	96
	How to determine if change journal support is useful in your NetBackup environment	100
	Guidelines for enabling NetBackup change journal support	101
	Credential Access properties	102
	Data Classification properties	103

Creating a Data Classification	104
Default Job Priorities properties	105
Understanding the Job Priority setting	106
Distributed application restore mapping properties	107
Encryption properties	108
Enterprise Vault properties	111
Enterprise Vault Hosts properties	112
Exchange properties	113
Exclude Lists properties	115
About the Add to exclude list and Add to exceptions list dialog boxes	118
Syntax rules for exclude lists	120
Traversing excluded directories	122
Fibre Transport properties	122
Firewall properties	124
Enabling logging for vnetd	127
General Server properties	128
Forcing restores to use a specific server	130
Global Attributes properties	131
About constraints on the number of concurrent jobs	133
Setting up email notifications about backups	135
Configuring the nbmail.cmd script	136
Sending email notifications to the administrator about unsuccessful backups	137
Sending messages to the global administrator about unsuccessful backups	138
Sending messages to the administrator about successful and unsuccessful backups	139
Installing and testing the email utility	139
Logging properties	141
Login Banner Configuration properties	146
Removing login banner screen and text	149
Enabling the Auto log off timeout option	149
Lotus Notes properties	150
Media properties	153
Results when media overwrites are not permitted	158
Recommended use for Enable SCSI reserve property	159
NDMP Global Credentials properties	160
NetWare Client properties	161
Network properties	162
Network Settings Properties	163
Reverse Host Name Lookup property	164
IP Address Family Support property	166

Port Ranges properties	167
Registered ports and dynamically-allocated ports	168
Preferred Network properties	169
Add or Change Preferred Network Settings dialog box	170
How NetBackup uses the directives to determine which network to use	172
Configurations to use IPv6 networks	175
Configurations to use IPv4 networks	176
Order of directive processing in the Preferred Network properties	177
Order of directives can affect processing	178
bptestnetconn utility to display Preferred Network information	178
Configuration to prohibit using a specified address	180
Configuration that uses the same specification for both the network and the interface—less constrictive	180
Configuration that uses the same specification for both the network and the interface—more constrictive	181
Configuration that limits the addresses, but allows any interfaces	182
Resource Limit properties	182
Restore Failover properties	184
Assigning an alternate media server as a failover restore server	185
Retention Periods properties	186
Changing a retention period	187
Determining retention periods for volumes	188
Servers properties	189
Adding a server to the Additional servers list	190
Adding a server to the Media servers list	191
Removing a server from the Additional servers list or the Media servers list	192
Switching to another master server in the Servers properties dialog box	192
About sharing one Enterprise Media Manager (EMM) database across multiple master servers	193
SharedDisk properties	195
SharePoint properties	195
Consistency check options for SharePoint Server	197
Symantec Products properties	197
Throttle Bandwidth properties	197
Add Bandwidth Settings dialog box for Throttle Bandwidth properties	198

	Timeouts properties	199
	Universal Settings properties	201
	Logging the status of a redirected restore	203
	UNIX Client properties	204
	UNIX Server properties	205
	VMware Access Hosts properties	205
	VSP (Volume Snapshot Provider) properties	206
	Windows Client properties	207
Chapter 4	Configuring server groups	209
	About server groups	209
	Configuring a server group	210
	Server group properties	212
	Deleting a server group	212
Chapter 5	Configuring host credentials	213
	About configuring credentials	213
Chapter 6	Managing media servers	215
	Activating or deactivating a media server	215
	Adding a media server	216
	About decommissioning a media server	217
	About decommissioning limitations	218
	Before you decommission a media server	219
	Post decommission recommendations	219
	Decommission actions	220
	Previewing references to a media server	224
	Decommissioning a media server	225
	Registering a media server	226
	Deleting all devices from a media server	227
	Removing a device host from the EMM database	229
Section 3	Configuring storage	231
Chapter 7	Configuring robots and tape drives	233
	About optical device support in NetBackup 7.0	234
	About NetBackup robot types	234
	Device configuration prerequisites	235
	About the device mapping files	235
	Downloading the device mapping files	236

About configuring robots and tape drives	237
About device discovery	237
About device serialization	238
About adding devices without discovery	239
About robot control	239
Library sharing example	240
Configuring robots and tape drives	241
Configuring robots and tape drives by using the wizard	241
Adding a robot	241
Robot configuration options	243
Adding a tape drive	246
Adding a shared tape drive	248
Tape drive configuration options	248
About drive name rules	251
Configuring drive name rules	252
Adding a tape drive path	254
Correlating tape drives and SCSI addresses on Windows hosts	256
Updating the device configuration by using the wizard	257
Managing robots	257
Changing robot properties	258
Configuring a robot to operate in manual mode	258
Deleting a robot	258
Moving a robot and its media to a new media server	259
Managing tape drives	261
Changing a drive comment	261
About downed drives	261
Changing a drive operating mode	262
Changing a tape drive path	262
Changing a drive path operating mode	263
Changing tape drive properties	263
Changing a tape drive to a shared drive	264
Cleaning a tape drive from the Device Monitor	264
Deleting a drive	265
Resetting a drive	265
Resetting the mount time	266
Setting drive cleaning frequency	267
Viewing drive details	267
Performing device diagnostics	268
About device diagnostic tests	268
Running a robot diagnostic test	268
Running a tape drive diagnostic test	270

	Managing a diagnostic test step that requires operator intervention	271
	Obtaining detailed information for a diagnostic test step	271
	Verifying the device configuration	271
	About automatic path correction	272
	Enabling automatic path correction	272
	Replacing a device	273
	Updating device firmware	274
	About the NetBackup Device Manager	275
	Stopping and restarting the Device Manager	275
Chapter 8	Configuring tape media	277
	About tape volumes	277
	NetBackup media types	278
	Alternate NetBackup media types	279
	About WORM media	280
	How to use WORM media in NetBackup	281
	About adding volumes	283
	About adding robotic volumes	284
	About adding stand-alone volumes	284
	Adding volumes by using the wizard	285
	Adding volumes by using the Actions menu	285
	Add volume properties	286
	Managing volumes	289
	Changing the group of a volume	290
	About rules for moving volumes between groups	290
	Changing the owner of a volume	290
	Changing the pool of a volume	291
	Changing volume properties	291
	About assigning volumes	294
	About deassigning volumes	294
	Deleting a volume	295
	Erasing a volume	295
	About exchanging a volume	297
	About frozen media	299
	Freezing or unfreezing a volume	299
	About injecting and ejecting volumes	300
	Injecting volumes into robots	300
	Ejecting volumes	300
	Media ejection timeout periods	302
	About rescanning and updating bar codes	303
	Rescanning and updating bar codes	304

About labeling NetBackup volumes	305
Labeling a volume	305
About moving volumes	306
Moving volumes by using the robot inventory update option	307
Moving volumes by using the Actions menu	308
About recycling a volume	310
Suspending or unsuspending volumes	312
About volume pools	312
About scratch volume pools	313
Adding a volume pool	314
Volume pool properties	314
Managing volume pools	315
Changing the properties of a volume pool	315
Deleting a volume pool	315
About volume groups	316
About media sharing	317
Configuring unrestricted media sharing	318
Configuring media sharing with a server group	318

Chapter 9

Inventorying robots	321
About robot inventory	322
When to inventory a robot	323
About showing a robot's contents	325
About inventory results for API robots	326
Showing the media in a robot	328
About comparing a robot's contents with the volume configuration	329
Comparing media in a robot with the volume configuration	330
About updating the volume configuration	331
Volume update prerequisites	332
About previewing volume configuration changes	333
Updating the volume configuration with a robot's contents	333
Robot inventory options	336
Configuring media settings	337
Media settings - existing media	338
Media settings - new media	340
About bar codes	344
About bar code advantages	344
About bar code best practices	345
About bar code rules	346
About media ID generation rules	348

	Configuring bar code rules	349
	Bar code rules settings	350
	Configuring media ID generation rules	351
	Media ID generation options	352
	Configuring media type mappings	354
	About adding media type mapping entries	355
	Default and allowable media types	356
	About the vmphyinv physical inventory	361
	How vmphyinv performs a physical inventory	363
	Example volume configuration updates	369
	Volume Configuration Example 1: Removing a volume from a robot	370
	Volume Configuration Example 2: Adding existing stand-alone volumes to a robot	371
	Volume Configuration Example 3: Moving existing volumes within a robot	373
	Volume Configuration Example 4: Adding new volumes to a robot	374
	Volume Configuration Example 5: Adding cleaning tapes to a robot	376
	Volume Configuration Example 6: Moving existing volumes between robots	377
	Volume Configuration Example 7: Adding existing volumes when bar codes are not used	378
Chapter 10	Configuring disk storage	381
	Configuring BasicDisk storage	381
	Configuring NearStore storage	381
	About configuring disk pool storage	382
	About SharedDisk support in NetBackup 7.0 and later	382
Chapter 11	Configuring storage units	385
	About the Storage utility	385
	Using the Storage utility	386
	About storage units	386
	Creating a storage unit using the Device Configuration Wizard	388
	Creating a storage unit using the Actions menu	388
	Creating a storage unit by copying a storage unit	389
	Changing storage unit settings	389
	Deleting storage units	389
	Media Manager storage unit considerations	390

Disk storage unit considerations	392
NDMP storage unit considerations	398
About storage unit settings	400
Absolute pathname to directory or absolute pathname to volume setting for storage units	400
Density storage unit setting	401
Disk pool storage unit setting	401
Disk type storage unit setting	401
Enable block sharing storage unit setting	402
Enable multiplexing storage unit setting	402
High water mark storage unit setting	402
Low water mark storage unit setting	403
Maximum concurrent write drives storage unit setting	403
Maximum concurrent jobs storage unit setting	404
Maximum streams per drive storage unit setting	406
Media server storage unit setting	406
NDMP host storage unit setting	409
On demand only storage unit setting	409
Only use the following media servers storage unit setting	410
Properties option in the Change Storage Units dialog box	411
Reduce fragment size storage unit setting	413
Robot number storage unit setting	414
Robot type storage unit setting	415
Staging schedule option in Change Storage Units dialog	415
Storage device setting for storage units	415
Storage unit name setting	415
Storage unit type setting	415
Enable temporary staging area storage unit setting	416
Transfer throttle storage unit setting	416
Use any available media server storage unit setting	416

Chapter 12	Staging backups	419
	About staging backups	419
	About the two staging methods	420
	About basic disk staging	421
	Creating a basic disk staging storage unit	422
	Configuring multiple copies in a relocation schedule	424
	Disk staging storage unit size and capacity	426
	Finding potential free space on a BasicDisk disk staging storage unit	428
	Disk Staging Schedule dialog box	430
	Basic disk staging limitations	432

	Initiating a relocation schedule manually	432
Chapter 13	Configuring storage unit groups	435
	About Storage unit groups	435
	Creating a storage unit group	435
	Deleting a storage unit group	437
	Storage unit selection criteria within a group	438
	Media server load balancing	439
	Other load balancing methods	440
	Exception to the storage unit selection criteria	442
	About disk spanning within storage unit groups	442
Chapter 14	Configuring storage lifecycle policies	443
	About storage lifecycle policies	443
	Creating a storage lifecycle policy	444
	Storage Lifecycle Policy dialog box settings	445
	About associating backup data with a data classification	447
	Accessing the Data Classification host properties	448
	Deleting a storage lifecycle policy	448
	Adding a storage destination to a storage lifecycle policy	450
	New or Change Storage Destination dialog box settings	452
	Staged capacity managed retention type for storage destinations	455
	Staged capacity managed retention type and disk types that support SIS	457
	Use for: Backup, duplication, Snapshot, or Import destination	457
	Retention type mixing for storage destinations	459
	Hierarchical view of storage destinations in the Storage lifecycle policy dialog box	460
	Adding a hierarchical duplication destination	462
	Adding a non-hierarchical duplication destination	463
	Modifying the source of a hierarchical duplication destination	463
	Removing a destination from the storage destination list	465
	Example of storage destination hierarchical view	465
	About writing multiple copies using a storage lifecycle policy	467
	How destination order determines the copy order	467
	About ensuring successful copies using lifecycles	468
	About storage lifecycle policy versions	469
	Storage lifecycle changes and versioning	469

	When changes to storage lifecycle policies become effective	471
	About deleting old storage lifecycle policy versions	472
	LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration	472
	LIFECYCLE_PARAMETERS file example	478
	About batch creation logic in Storage Lifecycle Manager	479
	Lifecycle operation administration using the nbstlutil command	480
Chapter 15	Duplicating images to a remote master server domain	483
	Process overview to duplicate to a remote master	483
	Setup overview to duplicate to a remote master domain	485
	About defining the domain relationship	487
	Configuring a replication target using MSDP	487
	Configuring the storage lifecycle policies required to duplicate to a remote master server	488
	Customizing how nbstserv runs duplication and import jobs	491
	One-to-many duplication to remote master server model	492
	Cascading duplications to remote masters	492
	Restoring from a backup at a remote master domain	496
	Reporting on duplication to remote master jobs	496
Section 4	Configuring backups	499
Chapter 16	Creating backup policies	501
	About the Policies utility	502
	Navigating in the Policies utility	502
	Planning for policies	504
	Example of one client in multiple policies	506
	Policy attributes that affect how clients are grouped in policies	507
	Creating a policy using the Backup Policy Configuration Wizard	508
	Creating a policy without using the Backup Policy Configuration Wizard	508
	Adding or changing schedules in a policy	509
	Changing multiple policies at one time	510
	Copying or moving policy items to another policy or server	511
	Deleting schedules, backup selections, or clients from a policy	512
	Policy Attributes tab	513

Policy type (policy attribute)	514
Data classifications (policy attribute)	517
Policy storage (policy attribute)	518
Policy volume pool (policy attribute)	519
Take checkpoints every __ minutes (policy attribute)	521
Limit jobs per policy (policy attribute)	525
Job priority (policy attribute)	526
Media Owner (policy attribute)	527
Go into effect at (policy attribute)	527
Follow NFS (policy attribute)	528
Backup Network Drives (policy attribute)	529
Cross mount points (policy attribute)	533
Compression (policy attribute)	536
Encryption (policy attribute)	538
Collect disaster recovery information for Bare Metal Restore (policy attribute)	538
Collect true image restore information (policy attribute) with and without move detection	538
Allow multiple data streams (policy attribute)	542
Disable client-side deduplication (policy attribute)	546
Enable granular recovery (policy attribute)	547
Keyword phrase (policy attribute)	548
Snapshot Client (policy attributes)	548
Microsoft Exchange (policy attributes)	548
Schedules tab	549
Schedule Attributes tab	549
Name (schedule attribute)	549
Type of backup (schedule attribute)	550
Synthetic backup (schedule attribute)	559
Calendar (schedule attribute)	559
Frequency (schedule attribute)	560
Instant Recovery (schedule attribute)	562
Multiple copies (schedule attribute)	562
Override policy storage (schedule attribute)	568
Override policy volume pool (schedule attribute)	568
Override media owner (schedule attribute)	568
Retention (schedule attribute)	569
Media multiplexing (schedule attribute)	572
Start Window tab	579
Adding, changing, or deleting a time window in a schedule	579
Example of schedule duration	582
Excluding dates from a policy schedule	583
Calendar Schedule tab	584

Scheduling by specific dates	584
Scheduling by recurring days of the week	584
Scheduling by recurring days of the month	586
How NetBackup determines which schedule to run next	587
About schedule windows that span midnight	589
How open schedules affect calendar-based and frequency-based schedules	590
Creating an open schedule in the NetBackup Administration Console	593
Runtime considerations that affect backup frequency	594
Runtime considerations	594
About the Clients tab	595
Adding or changing clients in a policy	595
Browse for Hyper-V virtual machines	597
Backup Selections tab	598
Adding backup selections to a policy	599
Verifying the Backup Selections list	601
How to reduce backup time	603
Pathname rules for Windows client backups	604
Pathname rules for Windows disk image (raw) backups	605
Pathname rules for Windows registry backups	607
About hard links to files and directories	608
Pathname rules for UNIX client backups	610
Pathname rules for NetWare NonTarget clients	617
Pathname rules for NetWare Target clients	619
Pathname rules for the clients that run extension products	619
About the directives on the Backup Selections list	620
Files that are excluded from backups by default	629
About excluding files from automatic backups	630
Files that are excluded by Microsoft Windows Backup	631
Disaster Recovery tab	631
Adding policies to the Critical Policies list of a catalog backup policy	634
Creating a Vault policy	634
Performing manual backups	636
Active Directory granular backups and recovery	637
System requirements for Active Directory granular NetBackup backups and recovery	637
Creating a policy that allows Active Directory granular restores	638
Restoring Active Directory objects	639
Troubleshooting granular restore issues	641

Chapter 17	Synthetic backups	645
	About synthetic backups	645
	Recommendations for synthetic backups and restores	646
	Synthetic full backups	648
	Synthetic cumulative incremental backups	650
	Schedules that must appear in a policy for synthetic backups	652
	Adding clients to a policy for synthetic backups	652
	Change journal and synthesized backups	653
	True image restore and synthesized backups	653
	Displaying synthetic backups in the Activity Monitor	653
	Logs produced during synthetic backups	654
	Synthetic backups and directory and file attributes	654
	Using the multiple copy synthetic backups method	655
	Configuring multiple copy synthetic backups	656
	Configuration variables	657
	Configuration examples	658
	Optimized synthetic backups using OpenStorage	659
	Optimized synthetic backups for deduplication	659
Chapter 18	Protecting the NetBackup catalog	661
	About NetBackup catalogs	661
	Parts of the NetBackup catalog	662
	About the NetBackup image database	663
	About the NetBackup relational database	665
	Protecting the NetBackup catalog	668
	About online, hot catalog backups	668
	Recovering the catalog	678
	Disaster recovery emails and the disaster recovery file	678
	Archiving the catalog	679
	Creating a catalog archiving policy	680
	Catalog archiving commands	681
	When to catalog archive	683
	Extracting images from the catalog archives	683
	Estimating catalog space requirements	684
	NetBackup file size considerations	686
	About the binary catalog format	686
	Moving the image catalog	687
	Indexing the catalog for faster access to backups	688
	About image catalog compression	689

Chapter 19	About the NetBackup relational database	693
	About the NetBackup relational database (NBDB) installation	693
	About NetBackup master server installed directories and files	695
	About the NetBackup configuration entry	701
	Sybase SQL Anywhere server management	702
	Sybase SQL Anywhere and clustered environments	702
	Using the NetBackup Database Administration utility	703
	About the General tab of the NetBackup Database Administration utility	704
	About the Tools tab of the NetBackup Database Administration utility	711
	Post-installation tasks	721
	Changing the database password	721
	Moving NBDB database files after installation	722
	Adding a mirrored transaction log	723
	Creating the NBDB database manually	724
	About backup and recovery procedures	726
	Database transaction log	726
	About catalog recovery	727
	Commands for backing up and recovering the relational databases	727
	About the online, hot catalog backup process	728
	Unloading the NetBackup database	730
	Terminating database connections	730
	Moving the NetBackup database from one host to another	731
	Cluster considerations with the EMM server	735
	Moving the EMM server to a Windows cluster	735
	Removing the EMM server from a Windows cluster	736
Chapter 20	Managing backup images	737
	About the Catalog utility	737
	About searching for backup images	738
	Verifying backup images	740
	Viewing job results	741
	Promoting a copy to a primary copy	741
	Duplicating backup images	743
	About multiplexed duplication considerations	748
	Jobs that appear while making multiple copies	749
	Expiring backup images	750
	About importing backup images	750
	Importing backup images, Phase I	751

	Importing backup images, Phase II	753
	About importing expired images	753
	Initiating an import without the Import Wizard	754
	About importing Backup Exec media	755
	Differences between importing, browsing, and restoring Backup Exec and NetBackup images	758
Section 5	Monitoring and reporting	761
Chapter 21	Monitoring NetBackup activity	763
	About the Activity Monitor	763
	Activity Monitor topology	765
	About filtering topology objects	766
	About the Jobs tab	766
	Viewing job details	768
	Showing or hiding column heads	768
	Monitoring the detailed status of a selected job	769
	Deleting completed jobs	769
	Canceling a job that has not completed	769
	Restarting a completed job	770
	Suspending restore or backup jobs	770
	Resuming suspended or incomplete jobs	770
	Printing job list information	770
	Printing job detail information	771
	Copying Activity Monitor text to a file	771
	Changing the Job Priority dynamically	772
	About the Services tab	772
	Types of services	776
	Using the nbrbutil utility to configure the NetBackup Resource Broker	777
	Starting or stopping a service	782
	Monitoring NetBackup services	782
	About the Processes tab	782
	Monitoring NetBackup processes in the Process Details dialog box	787
	About the Drives tab	790
	Monitoring NetBackup tape drives	791
	Cleaning tape drives from the Activity Monitor	792
	About the jobs database	792
	About changing the default values	793
	About the BPDBJOBS_OPTIONS environment variable	793
	bpdjobs command line options	795

	Enabling the bpdjobs debug log	795
	About the Device Monitor	796
	About media mount errors	796
	About pending requests and actions	797
	About pending requests for storage units	797
	Managing pending requests and actions	798
	Resolving a pending request	798
	Resolving a pending action	799
	Resubmitting a request	800
	Denying a request	800
Chapter 22	Auditing NetBackup operations	801
	About NetBackup auditing	801
	Viewing the current audit settings	804
	Configuring auditing on a NetBackup master server	805
	Auditing configuration after upgrading to NetBackup 7.1	807
	User identity in the audit report	808
	Auditing host property changes	808
	Using the command line -reason or -r option	809
	Viewing the audit report	810
	nbaudit log behavior	814
	Retaining and backing up audit trail records	814
Chapter 23	Reporting in NetBackup	817
	About the Reports utility	818
	Running a report	819
	Copying report text to another document	819
	Saving or exporting a report	820
	Printing a report	820
	Status of Backups report	821
	Client Backups report	821
	Problems report	821
	All Log Entries report	821
	Images on Media report	821
	Media Logs report	821
	Images on Disk report	822
	Disk Logs report	822
	Disk Storage Unit Status report	822
	Disk Pool Status report	822
	Images on Tape report	822
	Tape Logs report	822
	Tape Contents report	823

	Tape Summary report	823
	Tape Written report	823
	Tape Lists report	823
Section 6	Administering NetBackup	825
Chapter 24	Management topics	827
	NetBackup naming conventions	827
	Wildcard use in NetBackup	828
	How to administer devices on other servers	830
	How to access media and devices on other hosts	831
	About the Enterprise Media Manager	832
	About Enterprise Media Manager domain requirements	832
	About sharing an EMM server	833
Chapter 25	Accessing a remote server	835
	Accessing remote servers	835
	About adding a NetBackup server to a server list	836
	Adding a server to a remote server list	837
	About choosing a remote server to administer	839
	Using the change server command to administer a remote server	840
	Indicating a remote system upon login	841
	About using the Remote Administration Console	842
	About using the Java Windows Administration Console	843
	About running the NetBackup Administration Console on a NetBackup client	844
	About troubleshooting remote server administration	845
Chapter 26	Using the NetBackup-Java administration console	847
	About the NetBackup-Java Administration Console	847
	About authorizing NetBackup-Java users	850
	Authorization file (auth.conf) characteristics	851
	About authorizing nonroot users for specific applications	853
	About authorizing specific tasks in jbpSA	854
	About authorizing NetBackup-Java users on Windows	855
	Restricting access to NetBackup-Java applications on Windows	856
	Runtime configuration options	856
	FIREWALL_IN	857
	FORCE_IPADDR_LOOKUP	858

	INITIAL_MEMORY, MAX_MEMORY	860
	MEM_USE_WARNING	860
	NBJAVA_CLIENT_PORT_WINDOW	860
	NBJAVA_CORBA_DEFAULT_TIMEOUT	861
	NBJAVA_CORBA_LONG_TIMEOUT	861
	PBX_PORT	862
	VNETD_PORT	862
	About logging the command lines that the NetBackup interfaces use	862
	About customizing jnbSA and jbpSA with bp.conf entries	863
	About improving NetBackup-Java performance	863
	About running the Java console locally	864
	About running a console locally and administering a remote server	864
	About enhancing console performance	865
	About determining better performance when console is run locally or uses remote display back	866
	NetBackup-Java performance scenario 1	866
	NetBackup-Java performance scenario 2	867
	About adjusting time zones in the NetBackup-Java console	867
	Adjusting the time zone in the NetBackup-Java console	868
	Configuring a custom time zone in the NetBackup-Java console	869
Chapter 27	Alternate server restores	871
	About alternate server restores	871
	About supported configurations for alternate server restores	872
	About performing alternate server restores	873
	About modifying the NetBackup catalogs	874
	Overriding the original server for restores	875
	About enabling automatic failover to an alternate server	877
	Expiring and importing media for alternate server restores	878
Chapter 28	Managing client restores	881
	About server-directed restores	881
	About client-redirected restores	882
	About restore restrictions	882
	About allowing all clients to perform redirected restores	883
	About allowing a single client to perform redirected restores	884
	About allowing redirected restores of a client's files	884
	Examples of redirected restores	885

	About restoring files and access control lists	889
	About restoring the files that have ACLs	889
	Restoring files without restoring ACLs	890
	How to improve search times by creating an image list	890
	About restoring the System State	891
	Restoring the System State	891
Chapter 29	Powering down and rebooting NetBackup servers	895
	Powering down and rebooting NetBackup servers	895
	Shutting down all NetBackup services on Windows	896
	Starting up all NetBackup services on Windows	896
	Rebooting a NetBackup server	896
	Rebooting a NetBackup media server	897
Chapter 30	About Granular Recovery Technology	899
	About installing and configuring Network File System (NFS) for Active Directory Granular Recovery	899
	About configuring Services for Network File System (NFS) on the Windows 2008 and Windows 2008 R2 NetBackup media server and NetBackup clients	900
	Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2	901
	Disabling the Client for NFS on the media server	905
	Disabling the Server for NFS	906
	About configuring Services for Network File System (NFS) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients	908
	Installing Services for NFS on the Windows 2003 R2 SP2 media server	909
	Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2	912
	Configuring a UNIX or Linux media server and Windows clients for backups and restores that use Granular Recovery Technology	915
	Configuring a different network port for NBFSD	915
	Configuring the log on account for the NetBackup Client Service for Windows	916
Index		917

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
- About NetBackup administration interfaces
- About using the NetBackup Administration Console
- NetBackup configuration wizards
- Activity Monitor utility
- NetBackup Management utilities
- Media and Device Management utilities
- Running the Troubleshooter
- 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.

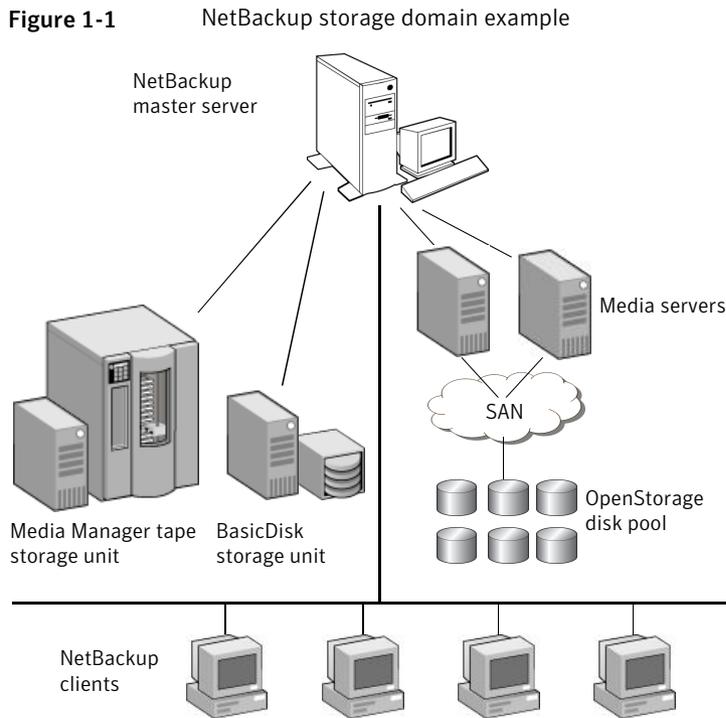
NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. An administrator can carefully schedule backups to achieve systematic and complete backups over a period of time, and optimize network traffic during off-peak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



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*.

About NetBackup administration interfaces

The NetBackup administrator has a choice of several interfaces to use to administer NetBackup. All the interfaces have similar capabilities. The best choice depends on personal preference and the workstation that is available to the administrator.

- **NetBackup Administration Console**
 - On Windows:

Select **NetBackup Administration Console** from the **Start** menu.
Or, install and use the Java Windows Administration Console. The Java Windows Administration Console is not automatically installed on the system. Installation is available on the main NetBackup for Windows Servers installation screen.

- On UNIX:
The **NetBackup Administration Console** is the recommended interface and is the interface referred to by most procedures and examples in the documentation. Start the Java-based, graphical user interface by running the `jnbSA` command.

Note: To log in to any **NetBackup Administration Console**, your login credentials must be authenticated from the connecting master or media server. This is true whether or not NetBackup Access Control (NBAC) is in use.

- Remote Administration Console
You can install the Remote Administration Console on a Windows computer to administer or manage any remote NetBackup server—Windows or UNIX. No license is required to install the Remote Administration Console. See “About using the Remote Administration Console” on page 842.
- Command line
You can enter NetBackup commands at the system prompt or use them in scripts.
All NetBackup administrator programs and commands require root or administrator user privileges by default.
For complete information on all NetBackup commands, see the *NetBackup Commands Reference Guide*.

About running the Windows-based NetBackup Administration Console

The **NetBackup Administration Console** is installed with the NetBackup server software.

The **NetBackup Administration Console** is the starting point for administering NetBackup. The left pane in the console contains a node for each major area of NetBackup administration. Click a node to display the information that is related to that node in the **Details** pane on the right. The menus contain commands relevant to the selected node.

Note: If there is more than one NetBackup server, the **NetBackup Administration Console** can be run on more than one server at one time. However, if more than one administrator makes changes to the configuration, the results are unpredictable.

Starting the Java-based Windows Display Console

The NetBackup-Java Windows Display Console is provided with NetBackup software. Use the Windows Display Console to administer UNIX NetBackup servers where a Java-capable UNIX system is not available.

See the *NetBackup Installation Guide* for information about how to install the Windows Display Console.

You can also use the Windows Display Console to administer a NetBackup UNIX or Windows server. Or, use a point-to-point (PPP) connection between the display console and other servers to perform remote administration.

The following procedure describes how to start the Windows display console.

To start the Windows display console

- 1 On a Windows system where the Windows Display Console is installed and configured, select **Start > Programs > Symantec NetBackup > NetBackup-Java Version 7.1**.
- 2 The login screen for the **NetBackup Administration Console** displays the host name. Log into another server by typing the name of another host in the **Host name** field. Or, select a host name from the drop-down list.
- 3 In the login screen, type your user name and password. To log into a Windows server, enter both the domain of the server and the user name as follows:

domain_name\user_name

The *domain_name* specifies the domain of the NetBackup host. If the host is not a member of a domain, the *domain_name* is not required.

- 4 Click *Login* to log into the NetBackup-Java application server program on the specified server. The interface program continues to communicate through the server that is specified in the login screen for the remainder of the current session.

The default host is the last host that was successfully logged into. The drop-down list contains the names of other hosts that have been logged into.

About administering remote servers

In a site that contains multiple master servers, you can configure the systems so that one **NetBackup Administration Console** can access remote servers. Indicate a remote server by using one of the following methods:

- Use the **File > Change Server** menu command.
- Use the **NetBackup-Java Administration Console**. Indicate a remote system upon NetBackup login.

Note: To log in to any **NetBackup Administration Console**, your login credentials must be authenticated from the connecting master or media server. This is true whether or not NetBackup Access Control (NBAC) is in use.

About 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

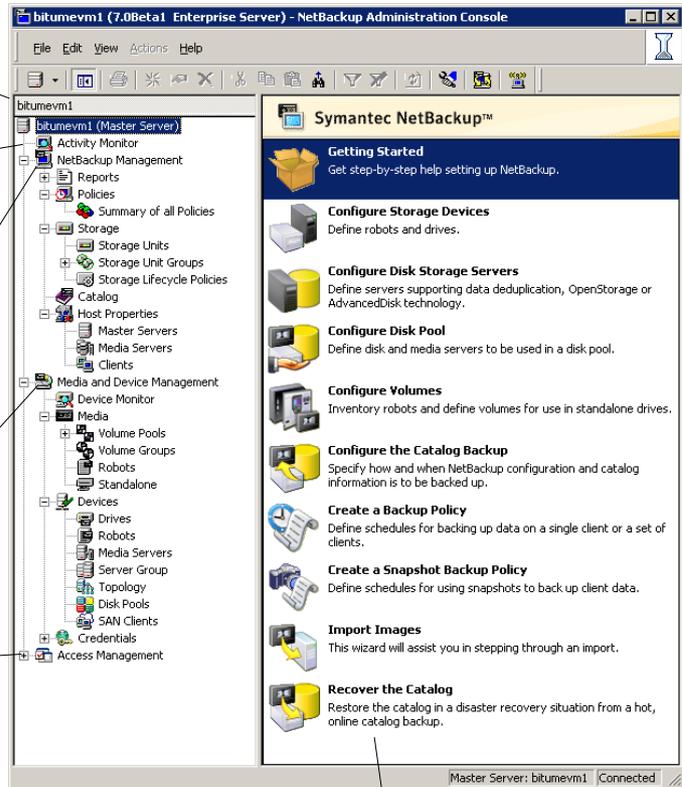
Master server
 The information in the NetBackup Administration Console applies to this server only.

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 the *NetBackup Commands Reference Guide*.

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**.

About 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.

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 which NetBackup utility is selected in the left portion of the screen.

The wizards help configure the basic properties of a NetBackup environment. After completing these basic wizards, you should be able to back up clients and perform a back up the NetBackup catalog.

Table 1-1 Configuration wizards

Wizard	Description
Getting Started Wizard	<p>Configures NetBackup for the first time. The wizard leads the user through the necessary steps to a working NetBackup configuration.</p> <p>The Getting Started Wizard is comprised of the following wizards, which can also be run separately, outside of the Getting Started Wizard:</p> <ul style="list-style-type: none">■ Device Configuration Wizard■ Volume Configuration Wizard■ Catalog Recovery Wizard■ Backup Policy and Configuration Wizard <p>Configure more advanced properties through the NetBackup Administration Console. You also can use the Administration Console if you prefer not to use the wizards.</p>
Device Configuration Wizard	Configures NetBackup to use robotic tape libraries and tape drives.

Table 1-1 Configuration wizards (*continued*)

Wizard	Description
Storage Server Configuration Wizard	Creates the servers that manage disk storage. The wizard appears if an Enterprise Disk Option license or NetBackup Deduplication Option license is installed.
Disk Pool Configuration Wizard	Creates pools of disk volumes for backup by one or more media servers. The wizard appears if an Enterprise Disk Option license or NetBackup Deduplication Option license is installed.
Volume Configuration Wizard	Configures removable media to use for backups.
Catalog Recovery Wizard	Sets up catalog backups. Catalog backups are essential to recover data in the case of a server failure or crash.
Backup Policy and Configuration Wizard	Adds a backup policy to the configuration.
Import Images Wizard	Imports NetBackup images in a two-part process.
Catalog Recovery Wizard	Specifies 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 “About the Activity Monitor” on page 763.

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 818.

- **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 “About the Policies utility” on page 502.

■ **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 385.

■ **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 737.

■ **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 55.

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.

Table 1-2 Media and device management utilities

Utility	Description
Device Monitor	Manages drives, device paths, and service requests for operators.
Media	Adds and manages removable media.
Devices	Adds, configures, and manages storage devices.

Table 1-2 Media and device management utilities (*continued*)

Utility	Description
Credentials	<p>Adds, removes, and manages log on credentials for the following:</p> <ul style="list-style-type: none"> ■ NDMP hosts (requires the NetBackup for NDMP license). ■ Storage servers (requires a NetBackup Deduplication Option or an Enterprise Disk Option license). <p>Credentials appears only if one of the previously mentioned license keys is installed.</p>

Running the Troubleshooter

When a NetBackup job returns a status code, use the **Troubleshooter** to find a description of the problem and a recommended solution. The **Troubleshooter** is particularly useful for understanding the status of a job in the **Activity Monitor** or in the **Reports** utility.

To run the Troubleshooter

- 1 In the **NetBackup Administration Console**, do one of the following:

- To understand the status of a job in the Activity Monitor
 - In the left pane, click **Activity Monitor**.
 - In the right pane, select the **Jobs** tab at the bottom of the pane.
 - Select a job from the list.

- To understand the status of a job in a report
 - In the left pane, expand **NetBackup Management > Reports**.
 - In the left pane, click the name of the report you want to run.
 For some reports, you must first expand a report group, and then click the name of the report.
 - In the right pane, click **Run Report**.
 - Select a job from the list that is generated.

To look up a status code Go to step 2.

2 Click **Help > Troubleshooter**.

The dialog box that appears describes the status code on the **Problem** tab. Possible solutions can be found on the **Troubleshoot** tab. The **Symantec Support** tab displays the Web address of Symantec Support or the URL to a technote that addresses the specific error code.

3 If no explanation appears, enter a status code and click **Lookup**.

The **Troubleshooter** provides assistance for NetBackup codes only. Assistance with Media and Device Management codes is available by using NetBackup online Help and searching for the particular status code.

See “Viewing job details” on page 768.

See “About the Jobs tab” on page 766.

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 NetBackup 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:

- About administering NetBackup licenses

About 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 NetBackup catalogs manually” on page 675.

A NetBackup capacity licensing utility is now available, which reports on the total amount of data that is protected by NetBackup.

For more information, see the “Capacity licensing” chapter of the *NetBackup Administrator’s Guide, Volume II*.

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

- Add a new license.
See “Adding new license keys” on page 47.
- Print a license.
See “Printing license key lists” on page 48.
- Delete a license.

See “Deleting license keys” on page 48.

- View the properties of one license.
See “Viewing license key properties” on page 49.
- Export the license list.
See “Exporting license keys” on page 49.

Restart the NetBackup Administration Console after any license updates.

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**, in the toolbar, click **Help > License Keys**.

To view the license keys of another server:

In the **NetBackup Administration Console**, in the toolbar, click **File > Change Server**, select another server, and click **OK**. In the toolbar, 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 47.
 - Print a license.
See “To print license key lists” on page 48.
 - Delete a license.
See “To delete license keys” on page 49.
 - View the properties of one license.
See “Viewing license key properties” on page 49.
 - Export the license list.
See “To export license keys” on page 49.

Adding new license keys

Use the following procedure to add new license keys.

To add new license keys

- 1 To add a license to the current server:
In the **NetBackup Administration Console**, in the toolbar, click **Help > License Keys**.
- To add a license to another server:
In the **NetBackup Administration Console**, in the toolbar, click **File > Change Server**, then select another server and click **OK**. Click **Help > License Keys** in the remote server.
- 2 In the **NetBackup License Keys** dialog box, click the **New** button.

- 3 In the **Add a new License Key** dialog box, 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 NetBackup catalogs manually” on page 675.

Printing license key lists

Use the following procedure to print license key lists.

To print license key lists

- 1 In the **NetBackup Administration Console**, in the toolbar, click **Help > License Keys**. 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 In the **NetBackup Administration Console**, in the toolbar, click **Help > License Keys**. In the **NetBackup License Keys** dialog box, 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 Administration Console**, in the toolbar, click **Help > License Keys**. 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 Administration Console**, in the toolbar, click **Help > License Keys**. In the **NetBackup License Keys** dialog box, click the **Export** button.
- 2 In the **Save As** dialog box, 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:

- NetBackup Host Properties configuration methods
- About the Host Properties
- Access Control properties
- Active Directory host properties
- Backup Exec Tape Reader properties
- Bandwidth properties
- Busy File Settings properties
- Clean-up properties
- Client Name properties
- Client Attributes properties
- Client Settings properties for NetWare clients
- Client Settings (UNIX) properties
- Client Settings properties for Windows clients
- 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
- Preferred Network properties
- Resource Limit properties
- Restore Failover properties
- Retention Periods properties
- Servers properties
- SharedDisk properties
- SharePoint properties
- Symantec Products properties
- Throttle Bandwidth properties
- Timeouts properties

- Universal Settings properties
- UNIX Client properties
- UNIX Server properties
- VMware Access Hosts properties
- VSP (Volume Snapshot Provider) properties
- Windows Client properties

NetBackup Host Properties configuration methods

The **Host Properties** and configuration options 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 55.

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**.

The `vm.conf` file contains media and device configuration options.

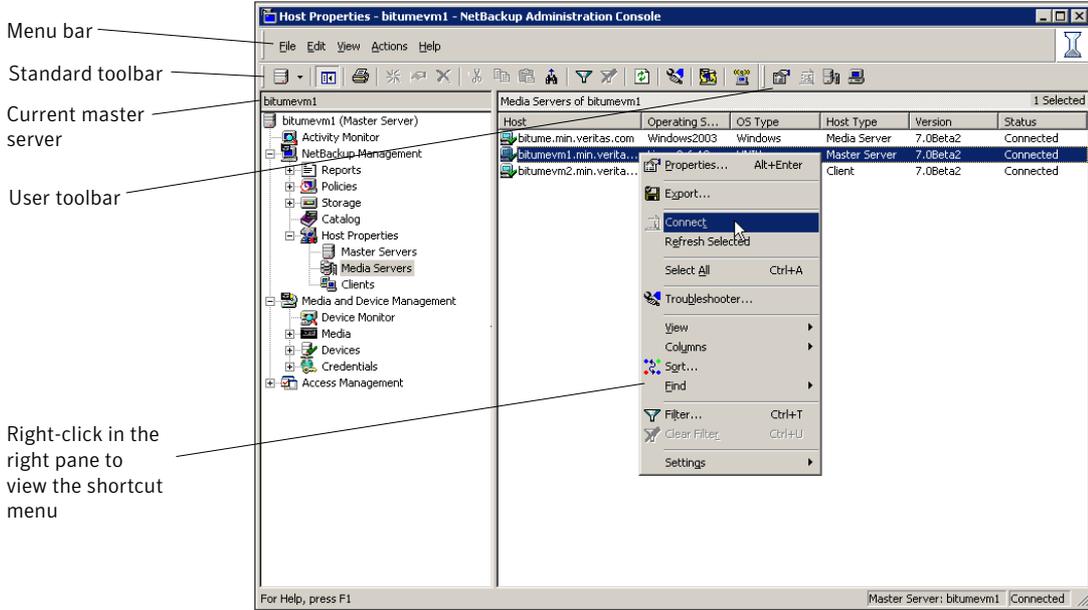
See the *NetBackup Administrator's Guide, Volume II* for more information.

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** in the **NetBackup Administration Console**.

Figure 3-1 Host Properties utility



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

The host properties can be changed in two ways:

- Use the **NetBackup Administration Console**.
- Use the `bpgetconfig` command to obtain a list of configuration entries, and then use `bpsetconfig` to change the entries in the registry.

The commands are described in the *NetBackup Commands Reference Guide*.

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 189.

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.

Note: All updates to a destination host fail if **Allow server file writes** is not enabled on the destination host. This property is located in the **Universal Settings** properties.

See “Universal Settings properties” on page 201.

See “About adding a NetBackup server to a server list” on page 836.

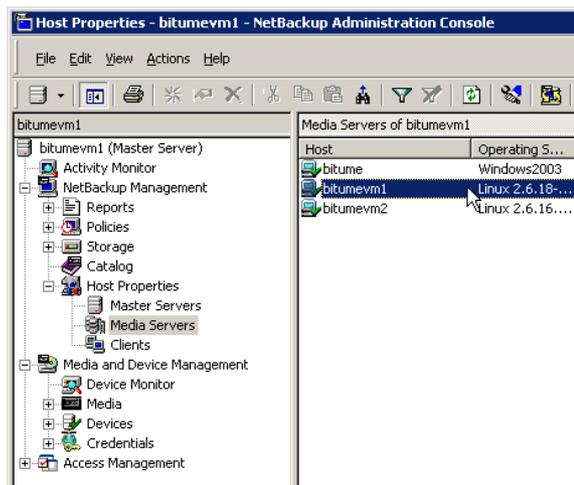
Viewing host properties

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 right pane, double-click the server or client to view the properties.

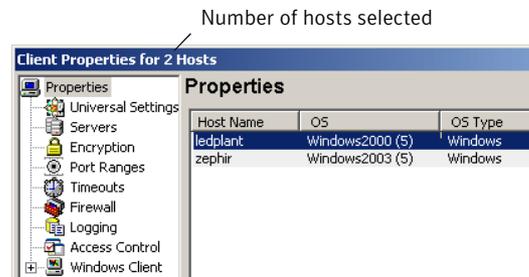
Changing the host properties on multiple hosts at the same time

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 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties**.
- 2 Select **Master Server**, **Media Server**, or **Clients**.
- 3 In the right pane, select a host. Hold down the **Shift** key and select another host.
- 4 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 computer in the configuration
 - Identifier
 - IP address
- 5 Make changes as necessary.
 - 6 Click **OK** to save the changes for all hosts and to close the dialog box.

Property states for multiple hosts

The **Host Properties** dialog boxes use the following conventions regarding multiple host selections:

Title of dialog box	<p>If a dialog box contains a Selected Host (or similarly named) box, all controls reflect the values for the host currently selected in the Selected Host box.</p> <p>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.</p>
Option selection	<p>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.</p>
Number spinners	<p>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.</p>
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.

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**. 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, check the associated option and type the value in the field.
- To leave the property unchanged, uncheck the associated option. The field changes to gray.

Exporting host properties

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

To export the properties of a host

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers, Media Servers, or Clients**.
- 2 Select a host. If you want to select multiple hosts, hold down the **Shift** key and select another host.
- 3 Click **File > Export**.
- 4 In the **Export Data** dialog box, enter the full path name or browse to the directory and click **Save**.

Standard host property dialog box options

The following options are available in every host property dialog box.

Default	Click Default to set all the properties in the current dialog box to the default values.
OK	Click OK to apply all changes since Apply was last clicked. OK also closes the dialog box.
Cancel	Click Cancel to cancel the changes that were made since the last time changes were applied.
Apply	Click Apply to save changes to all of the properties for the selected host(s). However, to apply changes click OK .

Help

Click **Help** for information on the properties that appear in the current dialog box.

Access Control properties

Use the **Access Control** host properties in the NetBackup Administration Console to configure NetBackup Authentication and Authorization. The properties apply to currently selected master servers, media servers, and clients.

The following tabs may display:

- Authentication Domain tab
See “Authentication Domain tab” on page 62.
- Authorization Service tab
See “Authorization Service tab” on page 63.
- Network Attributes
See “Network Attributes tab” on page 64.

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

The **NetBackup 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.

The **NetBackup Product Authentication and Authorization** property contains the following options.

Table 3-1 NetBackup Product Authentication and Authorization property options

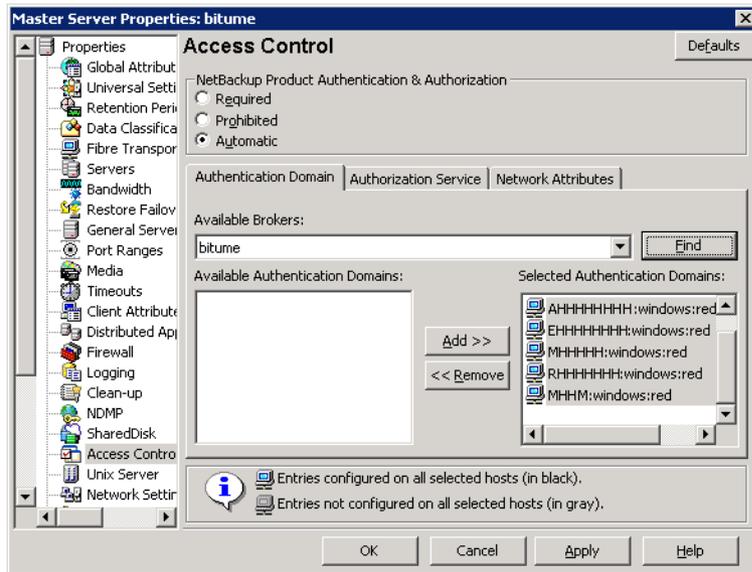
Option	Description
Required	Specifies that the local system should accept requests only from the remote systems that use NetBackup authentication and authorization. Connections from the remote systems that do not use NetBackup authentication and authorization are rejected. Select Required if all systems are at NetBackup 5.0 or later and maximum security is required.
Prohibited	Specifies that the local system should reject connections from any remote system that uses NetBackup authentication and authorization. Select Prohibited if the network is closed and maximum performance is required.
Automatic	Specifies that the local system should negotiate with the remote system about whether to use NetBackup 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*.

Authentication Domain tab

The **Authentication Domain** tab contains the properties that determine which authentication broker a computer uses. A master server that uses NetBackup authentication and authorization must have at least one authentication domain entry.

Figure 3-2 Authentication Domain tab



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

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

Table 3-2 Authentication Domain tab properties

Property	Description
Available Brokers	Select a broker, then click Find to list all of the available authentication domains.
Available Authentication Domains list	List of available authentication domains.

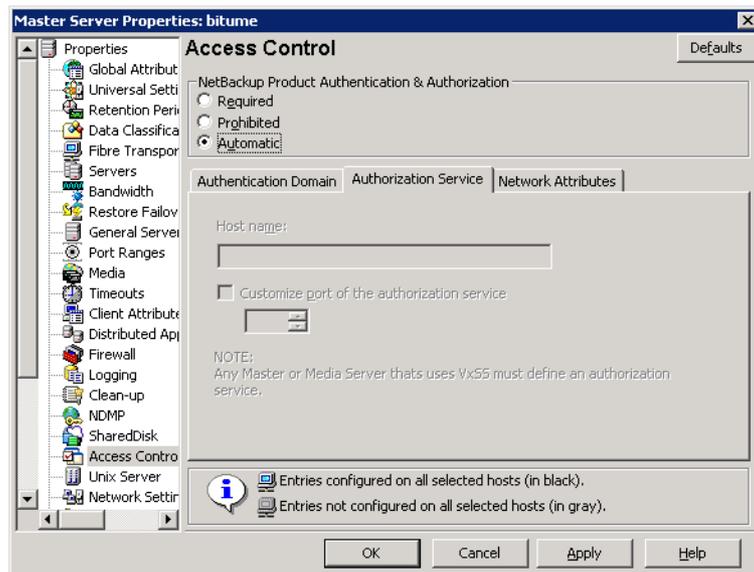
Table 3-2 Authentication Domain tab properties (*continued*)

Property	Description
Add button	Select the authentication domain(s) that this host can use and click Add .
Selected Authentication Domains list	List of the authentication domains selected for the host to use.
Remove button	Select the authentication domain(s) that you no longer want to use and click Remove .

Authorization Service tab

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

Figure 3-3 Authorization Service tab



The **Authorization Service** tab contains the following properties, which you can configure for a master or a media server.

Table 3-3 Authorization Service property options

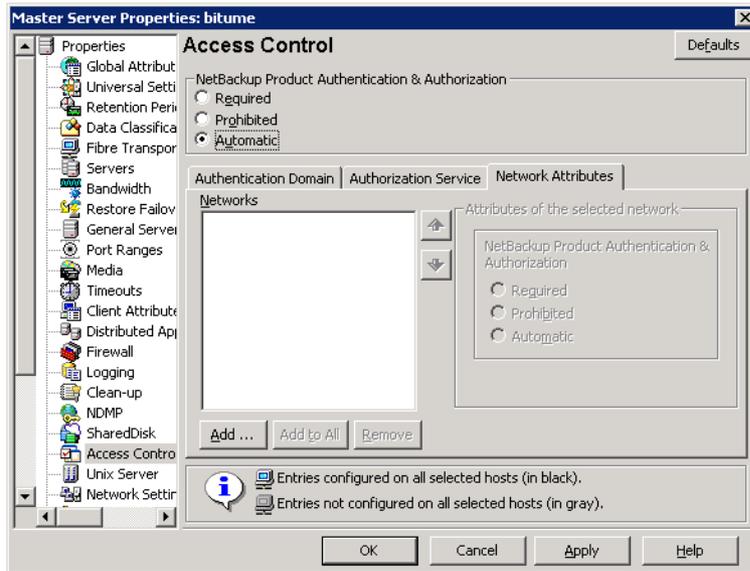
Option	Description
Host name	Specifies the host name or IP address of the authorization service.
Customize the port number of the authorization service	Specifies a nonstandard port number. Select Customize the port number and enter the port number of the authorization service.

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

Network Attributes tab

The **Network Attributes** tab contains a list of networks that are allowed (or not allowed) to use NetBackup authentication and authorization with the local system.

Figure 3-4 Network Attributes tab



The **Network Attributes** tab on the **Access Control** dialog box contains the following properties:

Networks

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

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

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

Click **Add** to add a network to the **Network** list.

Click **Add to All** to add a network to all currently selected hosts in the **Network** list.

Select a network name and click **Remove** to remove a network from the Network list.

NetBackup Product Authentication and Authorization property

The **NetBackup Product Authentication and Authorization property** in this tab determines whether the selected network uses access control and how the network uses it.

See “Access Control properties” on page 61.

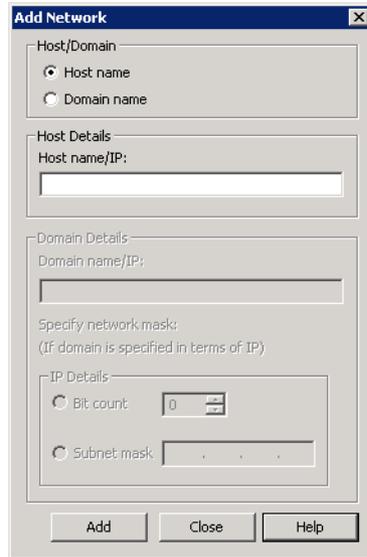
Add Network dialog box

The **Add Network** dialog box contains the following properties.

Table 3-4 Add Network dialog box properties

Property	Description
Host/ Domain	Indicates whether the network to be added is a Host name or a Domain name .
Host Details	Specifies that if the network is a host, one of the following items must be entered: <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.) ■ If the domain is specified by IP, select one of the following items: <ul style="list-style-type: none"> ■ Bit count Indicates 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 ■ Subnet mask Select to enter a subnet mask in the same format as the IP address.

Figure 3-5 Add Network dialog box

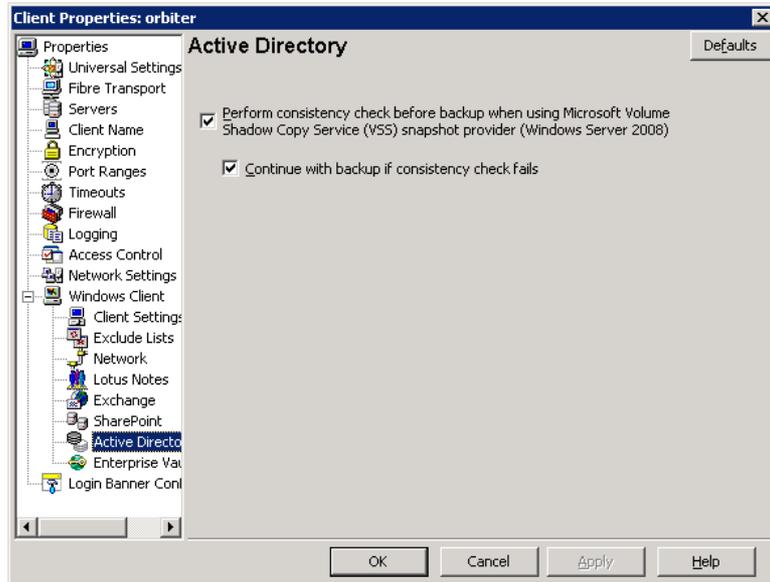


Active Directory host properties

The **Active Directory** properties in the NetBackup Administration Console 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 638.

Figure 3-6 Active Directory dialog box



The **Active Directory** dialog box contains the following properties.

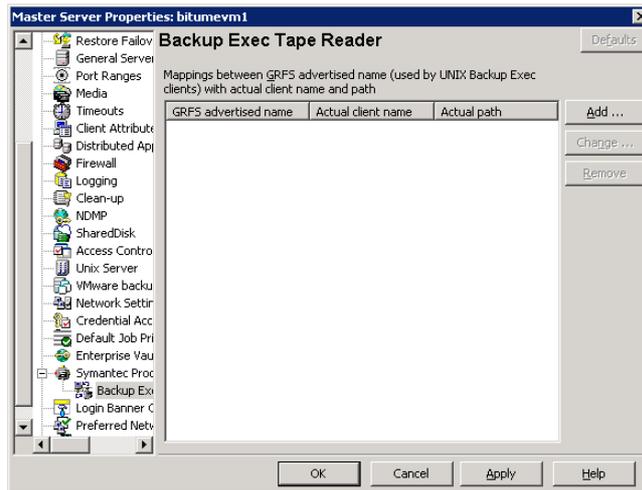
Table 3-5 Active Directory dialog box properties

Property	Description
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider	Checks snapshots for data corruption. 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 “Windows Open File Backup tab of the Client Attributes properties” on page 86.
Continue with backup if consistency check fails	Continues the backup job 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.

Backup Exec Tape Reader properties

The **Backup Exec Tape Reader** properties in the NetBackup Administration Console 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 3-7 Backup Exec Reader dialog box



The **Backup Exec Tape Reader** dialog box contains the following properties.

Table 3-6 Backup Exec Tape Reader dialog box properties

Property	Description
GRFS advertised name	<p>Specifies 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.</p> <p>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.</p> <p>The GRFS advertised name uses the following format:</p> <pre>ADVERTISED_HOST_NAME/advertised_path</pre> <p>where ADVERTISED_HOST_NAME is the advertised host name and advertised_path is the advertised path. Enter the ADVERTISED_HOST_NAME in capital letters.</p> <p>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.</p> <p>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.</p>
Actual client name	<p>Maps the advertised name to the real computer name.</p> <p>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.</p>
Actual path	<p>Maps the advertised path to the real path.</p> <p>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.</p>
Add	<p>Adds a GRFS entry. In the Backup Exec Tape Reader properties, click Add.</p>
Change	<p>Changes a selected GRFS entry. Select an entry in the Backup Exec Tape Reader properties list and click Change.</p>
Remove	<p>Removes a GRFS entry. Select an entry in the Backup Exec Tape Reader properties list and click Remove.</p>

See “About importing backup images” on page 750.

Bandwidth properties

Use the **Bandwidth** properties to specify network bandwidth limits for the NetBackup clients of the selected server.

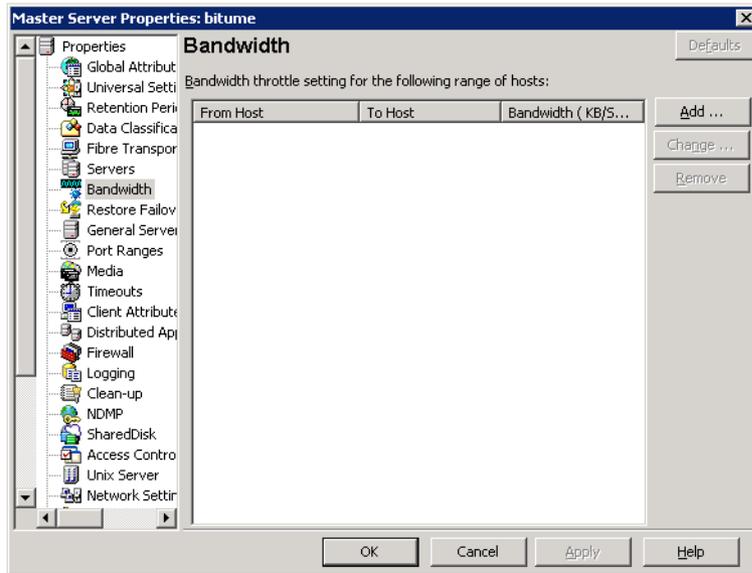
Note: The **Bandwidth** properties apply only to IPv4 networks. Use the **Throttle Bandwidth** properties to limit IPv6 networks.

See “Throttle Bandwidth properties” on page 197.

The actual limiting occurs on the client side of the backup connection. The bandwidth limits only restrict bandwidth during backups. By default, the bandwidth is not limited.

The **Bandwidth** properties apply to currently selected master servers.

Figure 3-8 Bandwidth dialog box



To manage entries in the **Bandwidth** dialog box, select one of the following buttons.

- Add** Adds an entry to the bandwidth table for each of the selected clients.
- Change** Changes an entry to the bandwidth table for each of the selected clients.
- Remove** Removes the selected entry from the bandwidth table.

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.

Bandwidth limit usage considerations and restrictions

Some usage restrictions apply to the bandwidth limit settings in the **Bandwidth** dialog box. The table below lists them and describes specific behavior that you may need to consider.

Table 3-7 Bandwidth limit usage considerations and restrictions

Client or operation	Bandwidth limit behavior or restrictions
NetBackup for Microsoft SQL-Server clients	Bandwidth limits are not supported
NetBackup for Oracle clients	Bandwidth limits are not supported
NetBackup for DataTools SQL-BackTrack clients	Bandwidth limits are not supported
local backups	If a server is also a client and data does not go over the network, bandwidth limits have no effect on local backups.
Setting required bandwidth	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.

Table 3-7 Bandwidth limit usage considerations and restrictions (continued)

Client or operation	Bandwidth limit behavior or restrictions
Distributing the workload of active backups	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.

Add Bandwidth Settings dialog box for Bandwidth properties

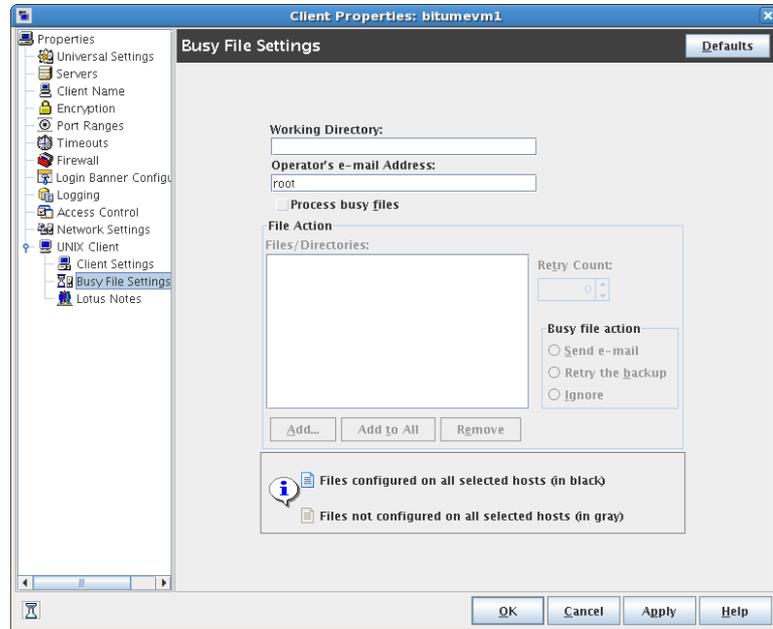
The **Add Bandwidth Settings** and the **Change Bandwidth Settings** dialog boxes contain the following properties.

From Host	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 Host	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 (KB/Sec)	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.

Busy File Settings properties

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

Figure 3-9 Busy File Settings dialog box



The **Busy File Settings** dialog box contains the following properties.

Table 3-8 Busy File Settings dialog box properties

Property	Description
Working directory	Specifies the path to the busy-files working directory. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, NetBackup creates the <code>busy_files</code> directory in the <code>/usr/opensv/netbackup</code> directory.
Operator's email address	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 <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, <code>BUSY_FILE_NOTIFY_USER</code> is not in any <code>bp.conf</code> file and the mail recipient is <code>root</code> .
Process busy files	Enables busy files to be processed 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	Specifies the absolute path and file name of the busy file. The metacharacters <code>*</code> , <code>?</code> , <code>[]</code> , <code>[-]</code> can be used for pattern matching of file names or parts of file names.
Add	Adds a new file entry. Enter the file and path directly, or browse to select a file.

Table 3-8 Busy File Settings dialog box properties (*continued*)

Property	Description
Add to All	Adds a new file entry for all of the clients currently selected. Enter the file and path directly, or browse to select a file.
Remove	Removes the selected file from the file action list.
Busy file action	The following options specify which action to take when busy-file processing is enabled. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. <ul style="list-style-type: none"> ■ Send email sends a busy file notification message to the user that is specified in Operator's email address. ■ Retry the backup retries the backup on the specified busy file. The Retry count value determines the number of times NetBackup tries a backup. ■ Ignore excludes 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	Specifies the number of times to try the backup. The default retry count is 1.

Activating the Busy File Settings in host properties

To activate the settings in the **Busy File Settings** host properties, use the following procedure.

To activate Busy File Settings

- 1 Copy the `bpend_notify_busy` script:

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

to the path:

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

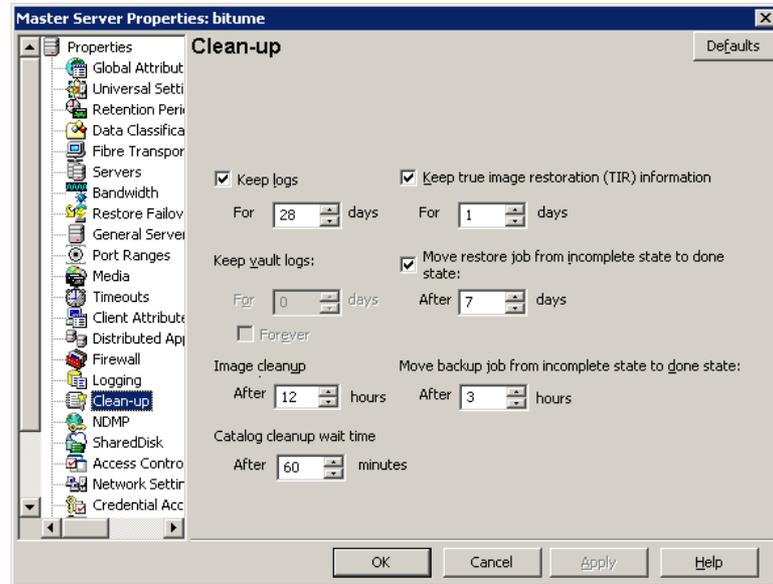
- 2 Set the file access permissions to allow group and others to run `bpend_notify`.
- 3 Configure a policy with a user backup schedule for the busy file backups.

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 in the **NetBackup Administration Console** refer to the retention of various logs and incomplete jobs. The **Clean-up** properties apply to currently selected master servers.

Figure 3-10 Clean-up dialog box



The **Clean-up** dialog box contains the following properties.

Table 3-9 Clean-up dialog box properties

Property	Description
Keep logs	<p>Specifies 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.</p> <p>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.</p>

Table 3-9 Clean-up dialog box properties (*continued*)

Property	Description
Keep vault logs	<p>If Vault is installed, the Keep vault logs option is enabled. It specifies the amount of time that the Vault session directories are kept.</p> <p>Session directories are found in the following location:</p> <pre data-bbox="515 467 1131 520">install_path\netbackup\vault\sessions\vaultname\session_x</pre> <p>where <i>x</i> is the session number. This directory contains vault log files, temporary working files, and report files.</p>
Image cleanup	<p>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.</p>
Catalog cleanup wait time	<p>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.</p>
Keep true image restoration information	<p>Specifies the number of days to keep true image restore information on disk. After the specified number of days, the images are pruned (removed). Applies to all policies for which NetBackup collects true image restore information. The default is one day.</p> <p>When NetBackup performs a true image backup, it stores the following images on the backup media:</p> <ul style="list-style-type: none"> ■ Backed up files ■ True image restore information <p>NetBackup also stores the true image restore information on disk in the <i>install_path\NetBackup\db\images</i> directory. NetBackup retains the information for the number of days that this property specifies.</p> <p>Keeping 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.</p>

Table 3-9 Clean-up dialog box properties (continued)

Property	Description
Move restore job from incomplete state to done state	<p>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.</p> <p>See “Universal Settings properties” on page 201.</p> <p>See “Checkpoint restart for restore jobs” on page 524.</p>
Move backup job from incomplete state to done state	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>

Client Name properties

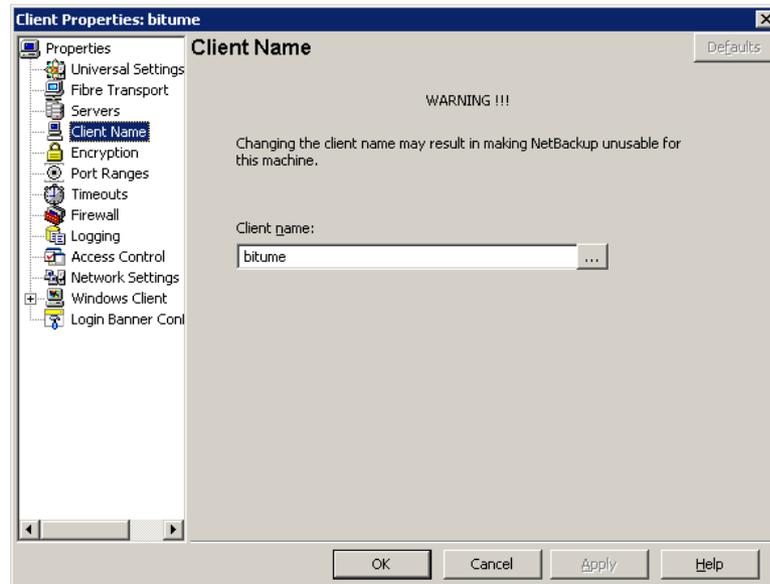
The **Client name** property in the **NetBackup Administration Console** specifies the NetBackup client name for the selected client. The name must match the name the policy uses to back up the client. The only exception is for a redirected restore, where the name must match that of the client whose 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 match, the client cannot browse for its own backups.

Note: Using an IPv6 address as a client name in a policy can cause backups to fail. Specify a hostname instead of an IPv6 address.

See “Client Attributes properties” on page 78.

Figure 3-11 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

In the **NetBackup Administration Console**, the **Client Attributes** properties apply to the clients of currently selected master servers.

The **Global client attributes** property applies to all clients, unless overridden as described in the following table.

Table 3-10 Global client attributes group box

Attribute	Description
Allow client browse	Allows 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	Allows 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	<p>Specifies the 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>If a client is not listed in the Clients list, click Add to display the Add Client dialog box and add a client to the client database. Type a client name in the text box or click the browse button (...) and select a client.</p> <p>See “Add Client dialog box” on page 80.</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. See “Client Name properties” on page 77.</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>
General tab	Specifies how to configure the selected Windows master servers (clients). See “General tab of the Client Attributes properties” on page 80.
Connect Options tab	Specifies how to configure the connection between a NetBackup server and a NetBackup client. See “Connect Options tab of the Client Attributes properties” on page 84.
Windows Open File Backup tab	Specifies whether a client uses Windows Open File Backup. Also, specifies whether Volume Snapshot Provider or Volume Shadow Copy Service is used as the snapshot provider. See “Windows Open File Backup tab of the Client Attributes properties” on page 86.

Add Client dialog box

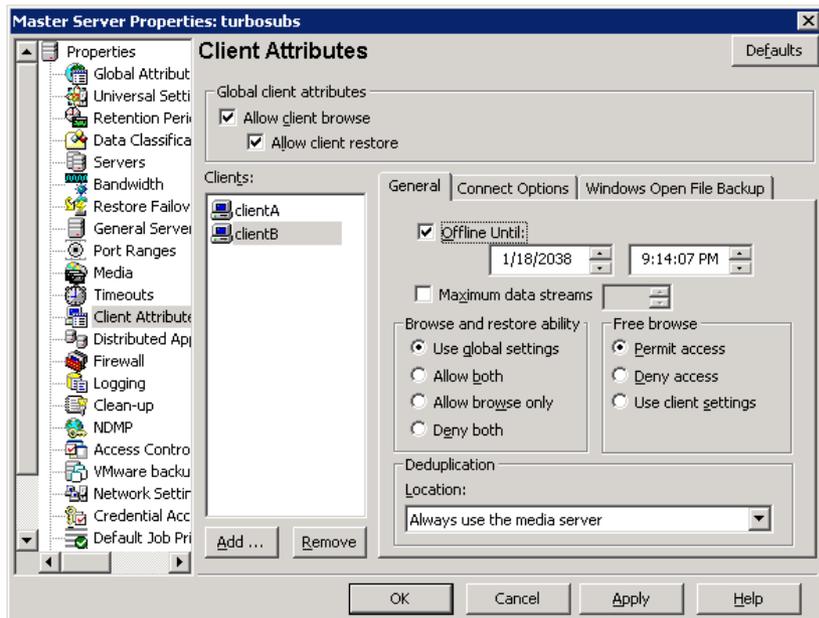
To add a client to the database, enter the name of a client, or browse to find a client. The **Add Client** dialog box contains the following properties.

- Enter client name** Specifies the name of the client to be added to the database. Type the name of the client to add.
- ... (browse)** Finds the list of current clients and displays them in the **Browse for computer** window. Select the client to add to the database and click **Add**.
- Add** Adds the specified client to the client database (client name displays in the **Clients** window).
- Close** Closes the **Add Client** dialog box.
- Help** Displays more information about how to add a client.

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 3-12 General tab of Client Attributes dialog box



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

Table 3-11 **General** tab properties

Property	Description
Offline Until:	<p>Makes the specified clients in the General tab unavailable for backups. By default, clients are online and included in the policies in which they are listed.</p> <p>When Offline Until: is selected for a client, no jobs are scheduled for that client. Since the client is not part of any job, no backup status will be listed for the client.</p> <p>After enabling the Offline Until: option, indicate the date and time when the clients are to be online again. The default setting is infinity, or approximately January 18, 2038, depending on the locale setting.</p> <p>Note: Changes to this property do not appear in the audit report.</p> <p>See “About NetBackup auditing” on page 801.</p> <p>The ability to take clients offline is useful in a number of situations.</p> <p>See “Offline option usage considerations and restrictions” on page 82.</p>
Maximum data streams	<p>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.)</p> <p>To change the setting, select Maximum data streams. Then scroll to or enter a value up to 99.</p> <p>The Maximum data streams property interacts with Maximum jobs per client and Limit jobs per policy as follows:</p> <ul style="list-style-type: none"> ■ 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. <p>See “Global Attributes properties” on page 131.</p> <p>See “Limit jobs per policy (policy attribute)” on page 525.</p>

Table 3-11 **General** tab properties (*continued*)

Property	Description
Browse and restore ability	<p>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.</p> <p>To use the Global client attributes settings, select Use global settings.</p> <ul style="list-style-type: none"> ■ 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	<p>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.</p> <p>Specifies whether the clients can list and restore from scheduled backups. (This setting does not affect user backups and archives.)</p> <p>Windows administrators can list and restore from scheduled backups as well as user backups regardless of the Free browse setting.</p>
Deduplication	<p>Specifies the deduplication action for clients if you use one of the following NetBackup deduplication options:</p> <ul style="list-style-type: none"> ■ NetBackup Deduplication Option ■ PureDisk Deduplication Option <p>For a description of the client direct deduplication options and their actions: See “Where deduplication should occur” on page 83.</p>

Offline option usage considerations and restrictions

The ability to take clients offline is useful in a number of situations. For example, in the event of planned outages or maintenance, client systems can be taken offline to avoid the unnecessary errors that administrators would then need to investigate. This option can also be used to anticipate new clients in the system; listing them in policies but configuring them as offline until they are in place and ready to be used.

The following actions can be performed if a client is offline.

Table 3-12 Offline option actions

Type of job or operation	Action or restriction
A client is offline and the job is already in progress	Offline clients continue to be included in any job.
A client is offline and job retries were started before the client was taken offline	Job retries continue as normal.
Any duplication job that is associated with a storage lifecycle policy and an offline client	Continues to run until complete.
LiveUpdate jobs for offline clients	Continues to run until complete.
Restore jobs	Can be run for offline clients.
The user attempts a manual backup for an offline client	The backup fails with a status code 1000, <code>Client is offline</code> . The user can either wait until the client is brought online again or bring the client online manually. Use either the NetBackup Administration Console or the <code>bpcclient</code> command to do so before resubmitting the manual job.
Archive backups	Not allowed for offline clients.
Administrators restarting or resuming jobs	Not allowed for offline clients.

Caution: If the master server is offline, hot catalog backups cannot run.

Where deduplication should occur

The **Deduplication** property specifies the deduplication action for clients if you use either the NetBackup Deduplication Option or the PureDisk Deduplication Option. The following table describes the client direct deduplication options.

Table 3-13 Client direct deduplication options

Option	Description
Always use the media server (the default)	Always deduplicates the data on the media server. The default. Jobs fail if one of the following are true: <ul style="list-style-type: none"> ■ The NetBackup Deduplication Engine on the deduplication storage server is inactive. ■ The PureDisk storage pool is inactive.

Table 3-13 Client direct deduplication options (*continued*)

Option	Description
Prefer to use client-side deduplication	Deduplicates data on the client and then send it directly to the storage server. 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.
Always use client-side deduplication	Always deduplicates the backup data on the client and then send it directly to the storage server. 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 (policy attribute)” on page 546.

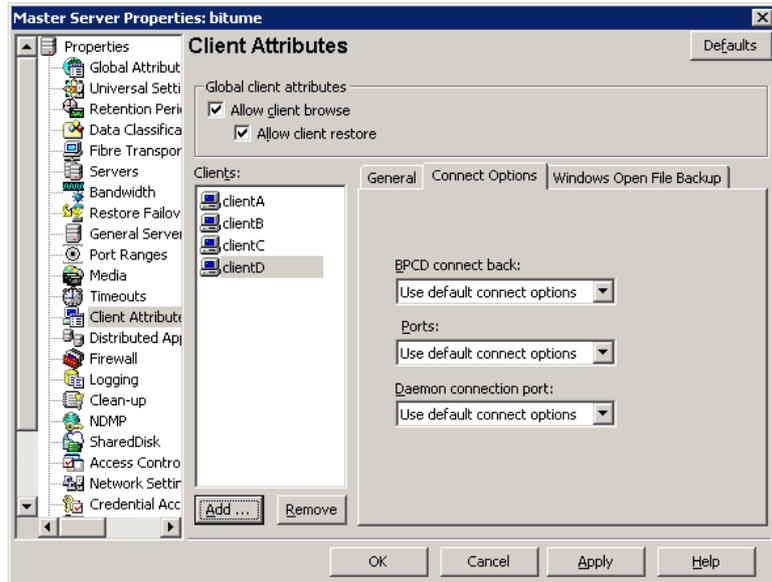
More information about client deduplication is available.

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 3-13 Connect Options tab of Client Attributes dialog box



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

Table 3-14 Connect Options tab properties

Property	Description
BPCD connect back	<p>Specifies how daemons are to connect back to the NetBackup Client daemon (BPCD) and contains the following options:</p> <ul style="list-style-type: none"> ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See "Firewall properties" on page 124. ■ Random port NetBackup randomly chooses a free port in the allowed range to perform the legacy connect-back method. ■ VNETD port NetBackup uses the <code>vnetd</code> port number for the connect-back method.

Table 3-14 Connect Options tab properties (*continued*)

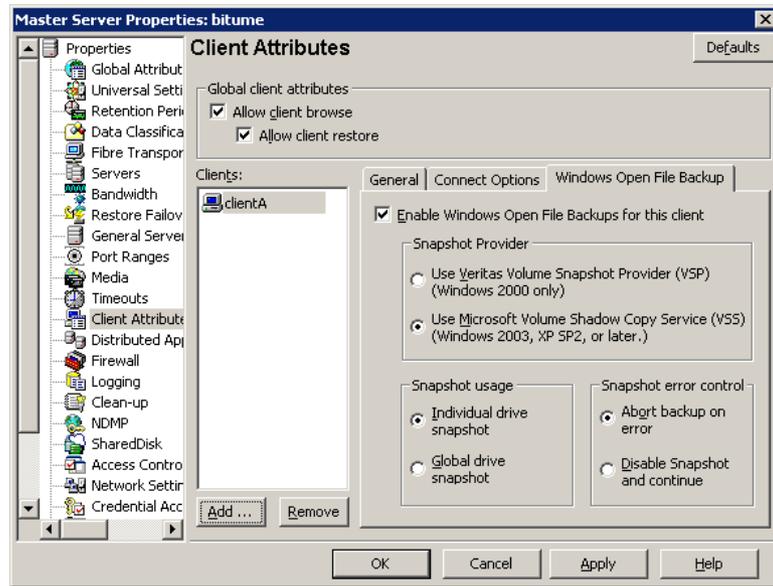
Property	Description
Ports	<p>Specifies the method that the selected clients should use to connect to the server and contains the following options:</p> <ul style="list-style-type: none"> ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See "Firewall properties" on page 124. ■ Reserved port Uses a reserved port number. ■ Non-reserved port Uses a non-reserved port number.
Daemon connection port	<p>Specifies the method that the selected clients should use to connect to the server and contains the following options.</p> <ul style="list-style-type: none"> ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. ■ Automatic Connects to the daemons on the server using <code>vnetd</code> if possible. If the daemons cannot use <code>vnetd</code>, the connection is made by using the daemon's legacy port number. ■ VNETD only Connects to the daemons on the server by using only <code>vnetd</code>. If the firewall rules prevent a server connection using the legacy port number, check this option. When selected, the BPCD connect back setting is not applicable. In addition, the Ports setting uses Non-reserved port, regardless of the value selected. ■ Daemon port only Connects to the daemons on the server by using only the legacy port number. This option only affects connections to NetBackup 7.0 and earlier. For connections to NetBackup 7.0.1 and later, the <code>veritas_pbx</code> port is used.

Windows Open File Backup tab of the Client Attributes properties

The **Windows Open File Backup** properties in the **NetBackup Administration Console** 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 3-14 Windows Open File Backup tab of Client Attributes dialog box



The **Windows Open File Backup** tab contains the following options.

Table 3-15 Windows Open File Backup tab properties

Property	Description
Add	<p>Adds NetBackup clients only if you want to change the default settings on the Windows Open File Backup tab.</p> <p>By default, no clients are listed in the Client Attributes dialog box. The server uses the following Windows Open File Backup defaults for all Windows clients:</p> <ul style="list-style-type: none"> ■ 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 90. ■ 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	<p>Deletes a client from the list by selecting the client and then clicking Delete.</p>

Table 3-15 Windows Open File Backup tab properties (*continued*)

Property	Description
Enable Windows Open File Backups	<p>Specifies that Windows Open File Backups be used for the selected clients. Adds clients to the list only if you want to change the default property settings.</p> <p>This option functions independently from the Perform Snapshot backups policy option that is available when the Snapshot Client is licensed.</p> <p>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.</p> <p>For more information, see the <i>NetBackup Snapshot Client Administrator's Guide</i>.</p>
Snapshot Provider	<p>Selects the snapshot provider for the selected clients:</p> <ul style="list-style-type: none"> ■ 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) <p>This option specifies that Microsoft VSS be used to create volume snapshots of volumes and logical drives for the selected clients.</p> <p>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.</p> <p>Configure VSS through the Microsoft VSS configuration dialog boxes.</p> <p>For information about performing Active Directory granular restores when using VSS, see the following topic:</p> <p>See "Active Directory host properties" on page 66.</p>

Table 3-15 Windows Open File Backup tab properties (*continued*)

Property	Description
Snapshot usage	<p>Selects how snapshots are made for the selected clients:</p> <ul style="list-style-type: none"> <p>■ Individual drive snapshot</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>■ Global drive snapshot</p> <p>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.</p> <p>In this situation, NetBackup takes a snapshot of C and D. Then NetBackup backs up C and backs up D.</p> <p>NetBackup then discards the C and D snapshots.</p> <p>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.</p> <p>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.</p>

Table 3-15 Windows Open File Backup tab properties (*continued*)

Property	Description
Snapshot error control	<p>Selects the processing instructions that NetBackup should follow if it encounters an error during processing:</p> <ul style="list-style-type: none"> <p>■ Abort backup on error</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>■ Disable snapshot and continue</p> <p>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.</p> <p>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.</p> <p>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.</p>

Backlevel and upgraded clients that use Windows Open File Backup

The following table shows the expected Open File Backup behavior based on the client version and the **Snapshot Provider** setting.

Table 3-16 Snapshot Provider behavior for clients in a 7.x 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.

Table 3-16 Snapshot Provider behavior for clients in a 7.x environment
(continued)

Client version	Snapshot Provider setting	Behavior
7.x	Windows VSS (7.0 default setting)	Using VSS for Open File Backup is a new default behavior in 7.x.
7.x	Veritas VSP	<p>Even if Veritas VSP is indicated, Windows VSS is used for Open File Backup.</p> <p>For upgraded clients:</p> <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.x	Windows VSS	Windows VSS is used for Open File Backup.

Client Settings properties for NetWare clients

The Client Settings properties apply to currently selected NetWare clients.

The **Client Settings** properties dialog box for NetWare clients includes the following options.

Table 3-17 NetWare Client Settings properties

Option	Description
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.

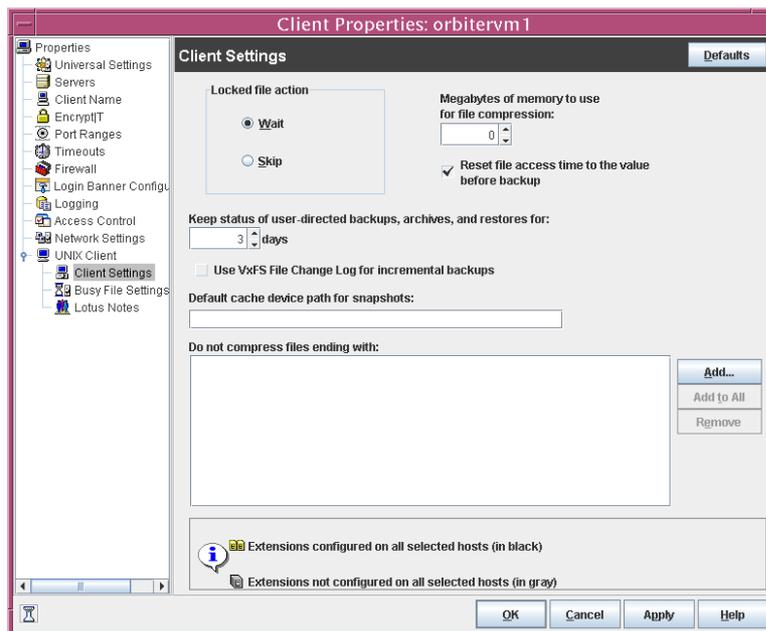
Table 3-17 NetWare Client Settings properties (*continued*)

Option	Description
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** properties in the **NetBackup Administration Console** apply to currently selected UNIX clients.

Figure 3-15 Client Settings (UNIX) dialog box



The UNIX **Client Settings** dialog box contains the following properties.

Table 3-18 UNIX Client Settings dialog box properties

Property	Description
Locked file action	<p>Determines what happens when NetBackup tries to back up a file with mandatory file locking enabled in its file mode.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> ■ 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 199. ■ 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	<p>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.</p> <p>Logs for user-directed operations are stored on the client system in the following directory:</p> <pre>install_path\NetBackup\logs\user_ops\ loginID\logs</pre>
Reset file access time	<p>Specifies that the access time (<code>atime</code>) 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.</p> <p>Note: This setting affects the software and the administration scripts that examine a file’s access time.</p>
Megabytes of memory to use for file compression	<p>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.</p> <p>The default is 0. This default is reasonable; change it only if problems are encountered.</p>
Use VxFS file change log for incremental backups	<p>Determines if NetBackup uses the File Change Log on VxFS clients.</p> <p>The default is off.</p> <p>See “VxFS file change log for incremental backups property” on page 94.</p>

Table 3-18 UNIX Client Settings dialog box properties (*continued*)

Property	Description
Default cache device path for snapshots	For additional information, see the <i>NetBackup Snapshot Client Administrator's Guide</i> .
Do not compress files ending with list	<p>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.</p> <p>Do not use wildcards to specify these extensions. For example, .A1 is allowed, but not .A* or .A[1-9]</p> <p>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.</p>
Add	Adds 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	Adds 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	Removes 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.

VxFS file change log for incremental backups property

The **Use VxFS file change log for incremental backups** property 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 “Backup Selections tab” on page 598.

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

- Enable the **Use VxFS file change log for incremental backups** 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 for incremental backups** 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 (policy attribute)” on page 533.

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 for incremental backups** property on the client.

The following table describes additional options that are available on the VxFS file change log feature.

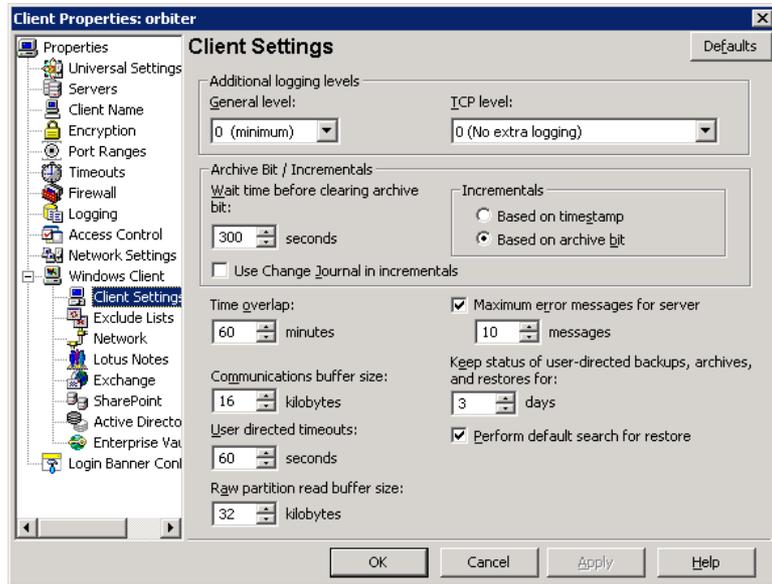
Table 3-19 VxFS file change log feature options

Option	Description
Activity Monitor messages	<p>Displays any messages that note when the file change log is used during a backup as follows:</p> <p>Using VxFS File Change Log for backup of <i>pathname</i></p> <p>Also notes when full and incremental backups are not synchronized.</p>
Keeping the data files synchronized with the FCL	<p>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.</p> <p>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.</p>
VxFS administration	<p>Additional VxFS commands are available to administrate the FCL in the <i>Veritas File System Administrator's Guide</i>.</p>

Client Settings properties for Windows clients

The Windows **Client Settings** properties apply to currently selected Windows clients.

Figure 3-16 Windows Client Settings dialog box



The **Client Settings** dialog box for Windows clients contains the following properties.

Table 3-20 Windows Client Settings properties

Property	Description
General level	Enables logs for <code>bpnetd</code> , <code>bpbkar</code> , <code>tar</code> , and <code>nbwin</code> . The higher the level, the more information is written. The default is 0.
TCP level	<p>Enables logs for TCP.</p> <p>Scroll to one of the following available log levels:</p> <ul style="list-style-type: none"> ■ 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 <p>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.</p>

Table 3-20 Windows Client Settings properties (*continued*)

Property	Description
Wait time before clearing archive bit	<p>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.</p> <p>This option applies only to differential-incremental backups. Cumulative-incremental backups do not clear the archive bit.</p>
Use change journal in incrementals	<p>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.</p> <p>When this property is enabled, it automatically enables the Incrementals are based on timestamp property.</p> <p>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.</p>
Incrementals based on timestamp	<p>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.</p>

Table 3-20 Windows Client Settings properties (*continued*)

Property	Description
Incrementals based on archive bit	<p>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.</p> <p>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.</p> <p>Disable this property to include a file in an incremental backup only if the date and time stamp for the file has changed since the last backup. For a differential-incremental backup, NetBackup compares the date/time stamp to the last full or incremental backup. For a cumulative-incremental backup, NetBackup compares the timestamp to the last full backup.</p> <p>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.</p> <p>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.</p>
Time overlap	<p>Specifies the number of minutes to add to the date range for incremental backups when you use 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.</p> <p>This value is used during incremental backups when you use the archive bit and when you examine the create time on folders. This comparison is done for archive bit-based backups as well as date-based backups.</p>
Communications buffer size	<p>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.</p>
User directed timeouts	<p>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.</p> <p>This property has no minimum value or maximum value. The default is 60 seconds.</p>

Table 3-20 Windows Client Settings properties (*continued*)

Property	Description
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 524.
- 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 you use the following options or products:

- True image restore (TIR) or True image restore with Move Detection
See “Collect true image restore information (policy attribute) with and without move detection” on page 538.
- Synthetic backups
See “About synthetic backups” on page 645.
- Bare Metal Restore (BMR)
For more information, see the *NetBackup Bare Metal Restore Administrator's Guide*.

See “How to determine if change journal support is useful in your NetBackup environment” on page 100.

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 3-17 Credential Access dialog box



The **Credential Access** dialog box contains the following properties.

Table 3-21 Credential Access dialog box properties

Property	Description
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	<p>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:</p> <ul style="list-style-type: none"> ■ 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. <p>Note: The credentials for the disk array or NDMP host are specified under Media and Device Management > Credentials.</p> <p>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.</p>

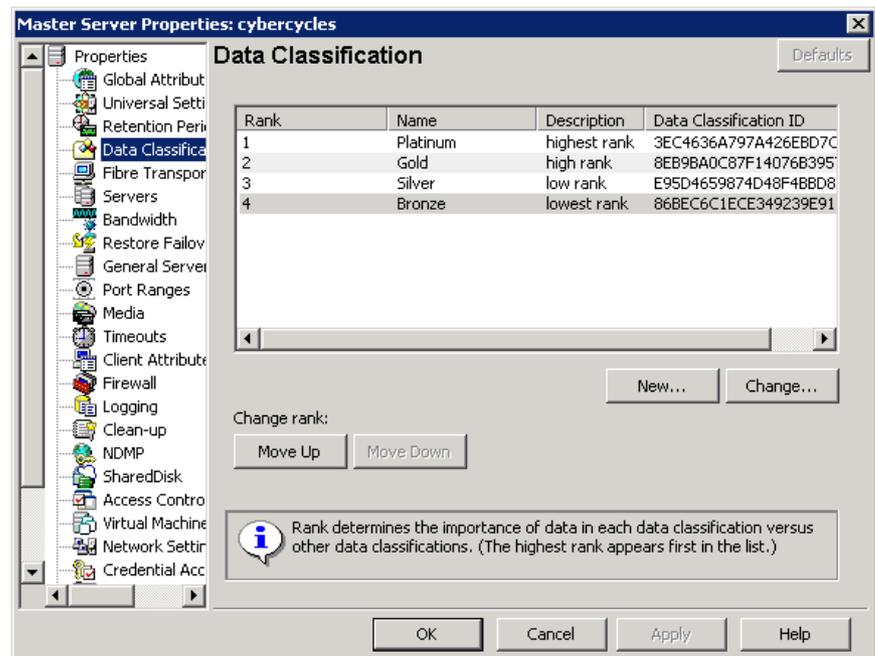
Data Classification properties

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

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

See “Data classifications (policy attribute)” on page 517.

Figure 3-18 Data Classification dialog box



The **Data Classification** dialog box contains the following properties.

Table 3-22 Data Classification dialog box properties

Property	Description
Rank column	<p>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.</p> <p>Use the Move Up and Move Down options to move the classification up or down in the list.</p> <p>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.</p>
Name column	<p>The Name column displays the data classification name. While data classifications cannot be deleted, the data classification names can be modified.</p> <p>NetBackup provides the following data classifications by default:</p> <ul style="list-style-type: none"> ■ Platinum (highest rank by default) ■ Gold (second highest rank by default) ■ Silver (third highest rank by default) ■ Bronze (lowest rank by default)
Description column	<p>In the Description column, enter a meaningful description for the data classification. Descriptions can be modified.</p>
Data Classification ID	<p>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.</p> <p>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 life cycles use the ID to identify the images that are associated with classification.</p> <p>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.</p>

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 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Host Properties**.
- 2 In the left pane, click **Data Classification**.
- 3 Click **New**.
- 4 Add the name and description in the **New Data Classification** dialog box.
- 5 Click **OK** to save the classification and close the dialog box.

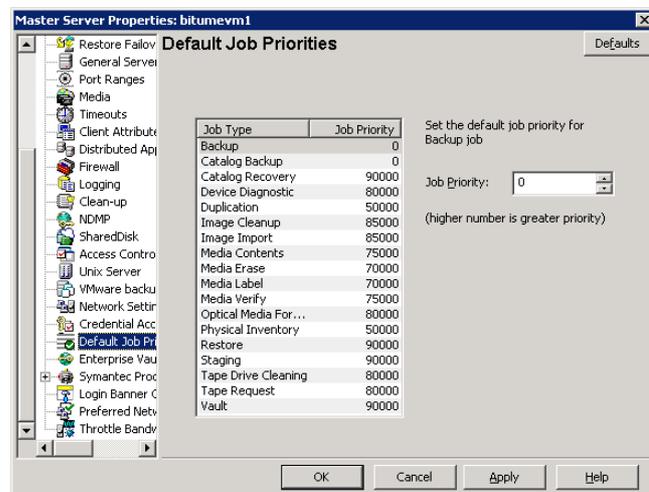
Note: Data classifications cannot be deleted.

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

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 3-19 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 772.
 - 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.
- The **Default Job Priorities** dialog box contains the following properties.

Table 3-23 Default Job Priorities dialog box properties

Property	Description
Job Type and Job Priority list	This listing includes 18 job types and the current configurable priority for each.
Job Priority	<p>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.</p> <p>A new priority setting affects all the policies that are created after the host property has been changed.</p> <p>A higher priority does not guarantee 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.</p> <p>However, the following factors can cause a job with a lower priority to run before a job with a higher priority:</p> <ul style="list-style-type: none"> ■ 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 (<code>nbrb</code>) receives a job request during an evaluation cycle, it does not consider the job until the next cycle, regardless of the job priority.

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, it 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 them. 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 properties.

Table 3-24 Distributed Application Restore Mapping dialog box properties

Property	Description
Add	<p>Adds a host that is authorized to run restores on SharePoint component hosts or Exchange hosts. You must provide the name of the Application host and the name of the Component host in the SharePoint farm or Exchange Database Availability Group (DAG).</p> <p>Note: For restores to be successful in an Exchange 2010 DAG environment, you must add the CAS server to the list.</p>
Change	Changes the application host or component host of the currently selected mapping.
Remove	Removes 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

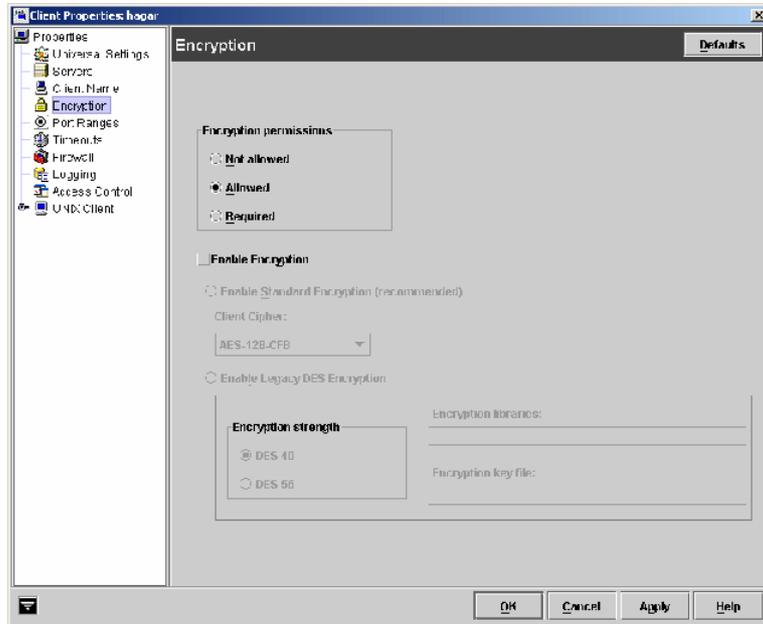
The **Encryption** properties control encryption on the currently selected client.

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 3-20 Encryption dialog box



The **Encryption permissions** property indicates the encryption setting on the selected NetBackup client as determined by the master server.

Table 3-25 Encryption permissions selections

Property	Description
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.

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

After **Enable Encryption** is selected, choose from the properties in Table 3-26.

Table 3-26 Encryption dialog box properties

Property	Description
Enable standard encryption	<p>Pertains to the 128-bit and the 256-bit options of NetBackup Encryption.</p> <p>If the selected client does not use Legacy encryption, Enable standard encryption is automatically selected.</p>
Client Cipher	<p>The following cipher types are available: BF-CFB, DES-EDE-CFB, AES-256-CFB, and AES-128-CFB. AES-128-CFB is the default.</p> <p>More information on the ciphers file is found in the NetBackup Security and Encryption Guide.</p>
Enable legacy DES encryption	<p>Pertains to the 40-bit and the 56-bit data encryption standard (DES) NetBackup encryption packages.</p>
Encryption strength	<p>Defines the encryption strength on the NetBackup client when Legacy encryption is used:</p> <ul style="list-style-type: none"> ■ 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	<p>Specify the folder that contains the encryption libraries on NetBackup clients.</p> <p>The default location is as follows:</p> <ul style="list-style-type: none"> ■ On Windows systems <code>install_path\netbackup\bin\</code> Where <i>install_path</i> is the directory where NetBackup is installed and by default is C:\Program Files\VERITAS. ■ On UNIX systems <code>/usr/opensv/lib</code> <p>If it is necessary to change the setting, specify the new name.</p>

Table 3-26 Encryption dialog box properties (*continued*)

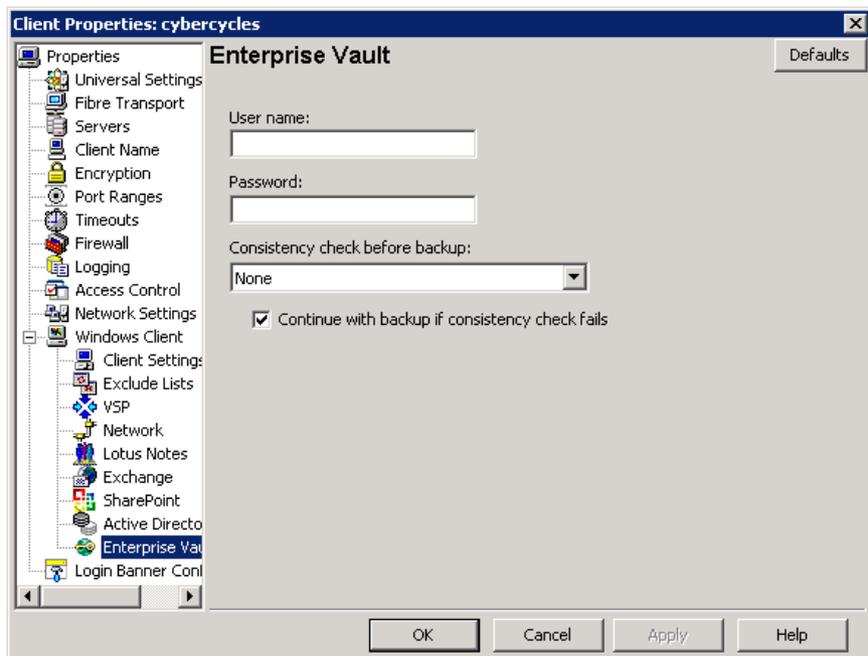
Property	Description
Encryption key file	<p>Specify the file that contains the encryption keys on NetBackup clients. The default location is as follows:</p> <ul style="list-style-type: none"> ■ On Windows systems <pre>install_path\NetBackup\bin\keyfile.dat</pre> <p>Where <i>install_path</i> is the folder where NetBackup is installed and by default is C:\Program Files\VERITAS.</p> ■ On UNIX systems <pre>/usr/opensv/netbackup/keyfile</pre> <p>If it is necessary to change the setting, specify the new name.</p>

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 3-21 Enterprise Vault dialog box



The Enterprise Vault dialog box contains the following properties.

Table 3-27 Enterprise Vault dialog box properties

Property	Description
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.

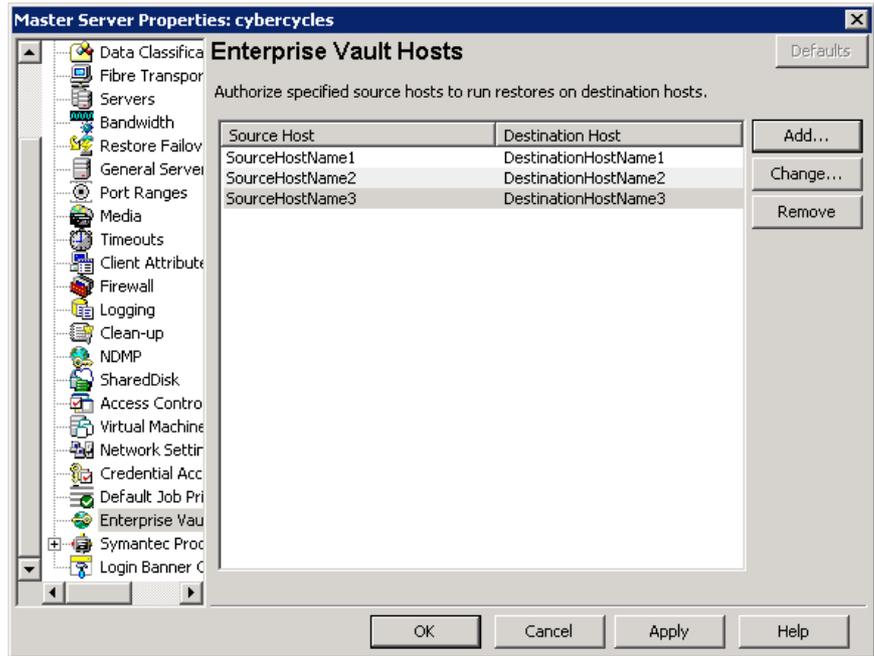
Enterprise Vault Hosts properties

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 3-22 Enterprise Vault Hosts master server properties



The Enterprise Vault Hosts dialog box contains the following properties.

Table 3-28 Enterprise Vault Hosts dialog box properties

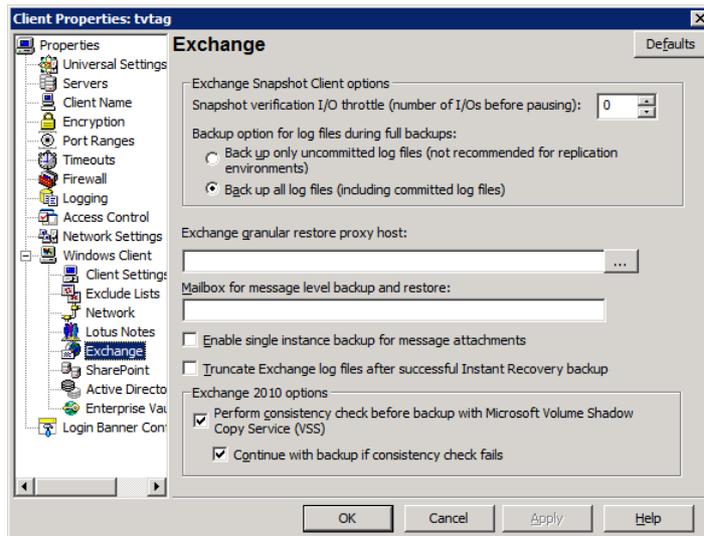
Option	Description
Add	Adds the source and the destination hosts within the Enterprise Vault configuration. You must provide the name of the Source host and the name of the Destination host .
Change	Changes the source host and the destination host, an entry that you select from the Enterprise Vault Hosts 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 3-23 Exchange dialog box



The **Exchange** dialog box contains the following properties.

Table 3-29 Exchange dialog box properties

Property	Description
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	<p>Choose which logs to include with snapshot backups:</p> <ul style="list-style-type: none"> ■ 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) <p>Note: In NetBackup 7.0, the default option is now Back up 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 Back up 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 Back up all log files (including committed log files)</p>

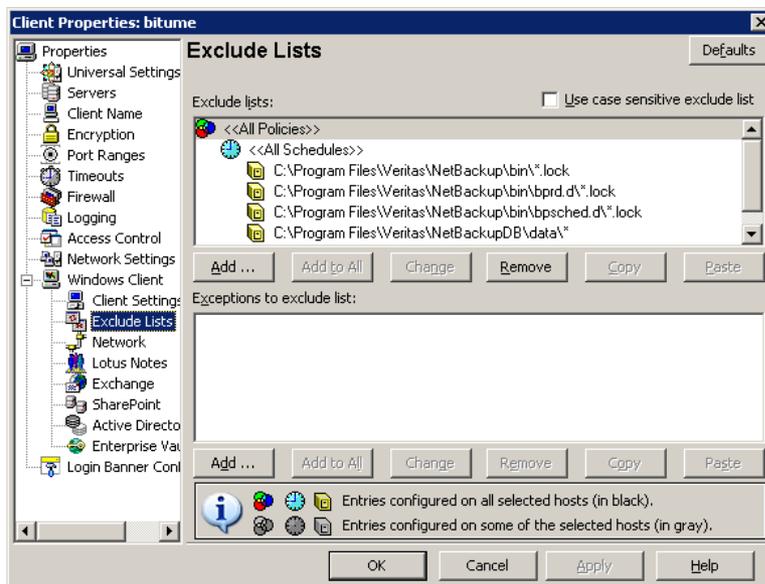
Table 3-29 Exchange dialog box properties (*continued*)

Property	Description
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 3-24 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.

Exclude and include lists that are set up for specific policies and schedules, are not used to determine if an entire drive is to be excluded when NetBackup determines if a backup job should be started.

Normally, this is not a problem. However, if a policy uses multistreaming, a drive which is excluded for a specific policy and schedule will have backup jobs started for it. Since no data will have needed to be backed up, this job reports an error

status when it completes. To avoid the situation, base the exclude list on the client and not on a policy and schedule.

The **Exclude Lists** dialog box contains the following properties.

Table 3-30 Excludes Lists dialog box properties

Property	Description
Use case sensitive exclude list property	Indicates that the files and directories to exclude are case-sensitive.
Exclude list	<p>Displays the policies that contain schedule, file, and directory exclusions as follows:</p> <ul style="list-style-type: none"> <p>■ Add</p> <p>Excludes 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 in this list are run, the files and directories that are specified on the list are backed up.</p> <p>■ Add to all</p> <p>Adds the selected list item to all currently selected clients. The item is excluded from the backup list on all selected clients. Add to all is enabled only when more than one client is selected for configuration and a list item is selected 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.</p> <p>■ Remove</p> <p>Removes the selected policy, schedule, or file from the Exclude list. The item is included in the backup.</p>

Table 3-30 Excludes Lists dialog box properties (*continued*)

Property	Description
Exceptions to exclude list	<p>Displays the policies, schedules, files, and 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.</p> <p>Options include the following:</p> <ul style="list-style-type: none"> ■ Add Creates 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. 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. ■ Add to all Adds 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. ■ Remove Removes the selected policy, schedule, or file from the Exceptions list. The item is excluded from the backup.

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:

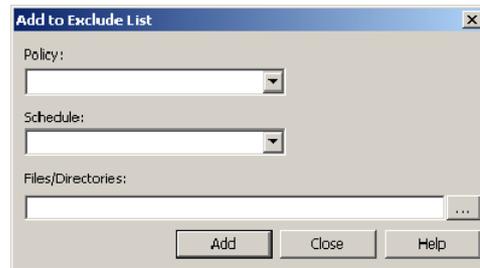
Table 3-31 Add to Exclude dialog box

Field	Description
Policy	The policy name that contains the files and the directories that you want 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>.
Schedule	The schedule name that is associated with the files and the directories that you want 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>.

Table 3-31 Add to Exclude dialog box (continued)

Field	Description
Files/Directories	The full path to the file(s) and the directories that you want to exclude or make exceptions for.

Figure 3-25 Add to Exclude List properties



Adding an entry to an exclude list

Use the following procedure to add an entry to an exclude list for a policy:

To add an entry to the exclude list

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Clients**. Double-click on a client.
- 2 Under the Exclude List, click **Add**.
- 3 In the **Policy** field, select a policy name from the drop-down menu or enter the name of a policy. Select **All Policies** to exclude these items from all policies.
- 4 In the **Schedule** field, select a schedule name from the drop-down menu or enter the name of a schedule. Select **All Schedules** to exclude the specified files and directories from all schedules in the policy.
- 5 In the **Files/Directories** field, enter the files or directories to be excluded from the backups that are based on the selected policy and schedule.
- 6 Click **Add** to add the specified files and directories to the exclude list.
- 7 Click **Apply** to accept the changes. Click **OK** to accept the changes and close the host properties dialog box.

Adding an exception to the exclude list

Use the following procedure to add an exception to the exclude list for a policy:

To add an exception to the exclude list

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Clients**. Double-click on a client.
- 2 Under the Exceptions to the Exclude List, click **Add**.
- 3 In the **Policy** field, select a policy name from the drop-down menu or enter the name of a policy. Select **All Policies** to add these items back into all policies. (In other words, these items are to be excluded from the exclude list.)
- 4 In the **Schedule** field, select a schedule name from the drop-down menu or enter the name of a schedule. Select **All Schedules** to add these items back into the schedules.
- 5 In the **Files/Directories** field, enter the files or directories to be added back into the backups that are based on the selected policy and schedule.
- 6 Click **Add** to add the specified files and directories to the Exceptions to the Exclude List.
- 7 Click **Apply** to accept the changes. Click **OK** to accept the changes and close the host properties dialog box.

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 “Wildcard use in NetBackup” on page 828.
See “NetBackup naming conventions” on page 827.
- 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 exclude list example

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** properties apply only when the SAN Client license is installed.

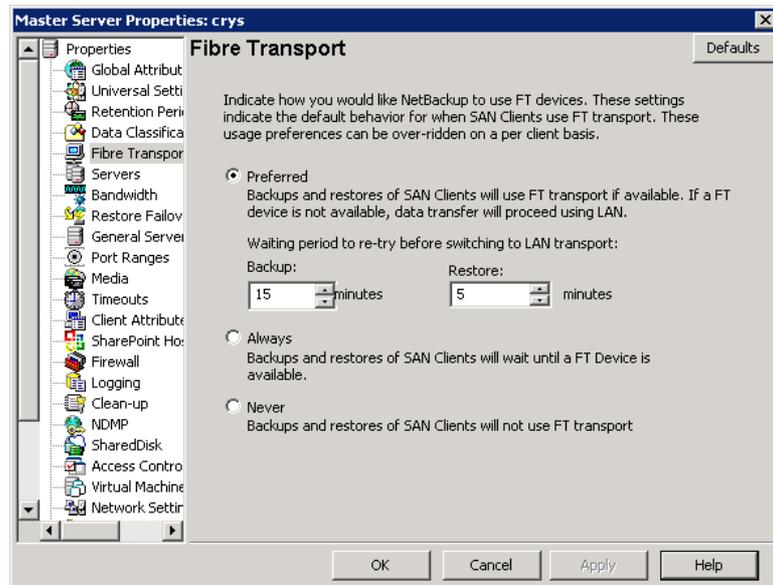
The **Fibre Transport** media server property applies to the SAN clients for selected media servers.

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 3-26 Master server Fibre Transport host properties



The master server **Fibre Transport** dialog box contains the following properties.

Table 3-32 Fibre Transport dialog box properties

Property	Description
Preferred	<p>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.</p> <p>If you select this option, also specify the wait period for backups and for restores.</p> <p>For the global property that is specified on the master server, the default is Preferred.</p>

Table 3-32 Fibre Transport dialog box properties (*continued*)

Property	Description
Always	<p>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.</p> <p>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.</p>
Never	<p>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.</p> <p>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.</p> <p>If you specify Never for a media server, Fibre Transport is disabled for the media server.</p> <p>If you specify Never for a SAN client, Fibre Transport is disabled for the client.</p>
Maximum concurrent FT connections	<p>This property applies to the media properties only.</p> <p>This property specifies the number of FT connections to allow to a media server.</p> <p>The default is four times the number of HBA target ports (maximum of 16).</p>
Use defaults from the master server configuration	<p>This property applies to the client properties only.</p> <p>This property specifies that the client follow the properties as they are configured on the master server.</p>

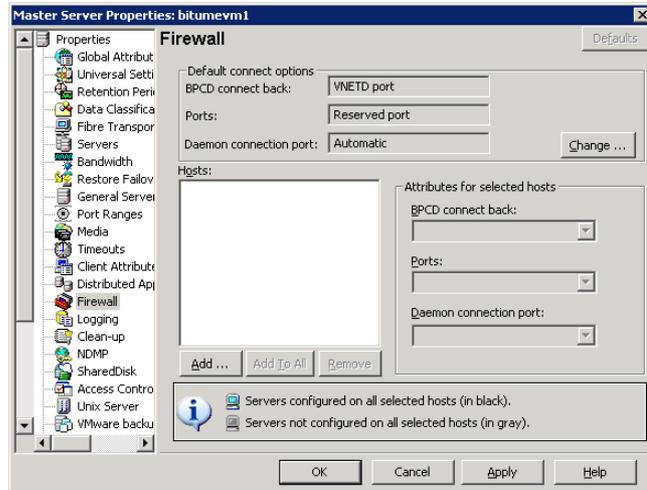
Firewall properties

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 78.

Figure 3-27 Firewall dialog box



The **Firewall** dialog box contains the following properties.

Table 3-33 Firewall dialog box properties

Property	Description
Default connect options	<p>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.</p> <p>By default, the firewall settings are configured to require the fewest possible ports to be open.</p> <p>To change the default connect options for the selected server, click Change.</p> <p>Click Change to change the Default connect options. Change the Firewall properties in the Default Connect Options dialog box.</p> <p>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.</p>

Table 3-33 Firewall dialog box properties (*continued*)

Property	Description
Hosts list	<p>To change the default connect options for any server, add the server to the host list. Servers do not automatically appear on the list.</p> <ul style="list-style-type: none"> ■ Add option 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 option 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 option 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	<p>This property specifies how daemons are to connect back to the NetBackup Client daemon (BPCD) as follows:</p> <ul style="list-style-type: none"> ■ 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 (<code>vnetd</code>) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. The server initiates all <code>bpcd</code> socket connections. Consider the example in which <code>bpbrm</code> on a media server initially connects with <code>bpcd</code> on a client. The situation does not pose a firewall problem because <code>bpbrm</code> uses the well-known <code>bpcd</code> port.

Table 3-33 Firewall dialog box properties (*continued*)

Property	Description
Ports	<p>Select whether a reserved or non-reserved port number should be used to connect to the server:</p> <ul style="list-style-type: none"> ■ 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 201.
Daemon connection port	<p>This option only affects connections to NetBackup 7.0 and earlier. For connections to NetBackup 7.0.1 and later, the <code>veritas_pbx</code> port is used.</p> <p>If configuring connections for NetBackup 7.0 and earlier, select the Daemon connection port method to use to connect to the server:</p> <ul style="list-style-type: none"> ■ 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 <code>vnetd</code> if possible. If it is not possible to use <code>vnetd</code>, the daemon’s traditional port number makes the connection. ■ VNETD only The daemons on the server are connected to by <code>vnetd</code> 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. <p>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.</p>
Defaults	Set property settings back to the defaults.

Enabling logging for vnetd

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

To enable logging for vnetd

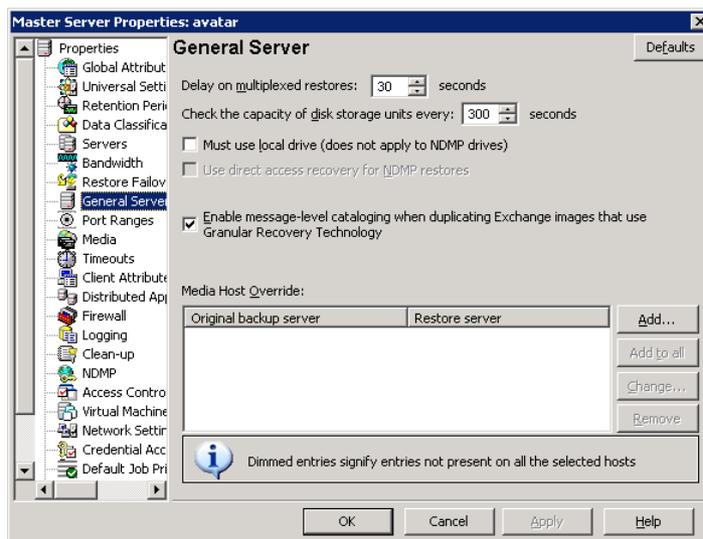
- ◆ Create a `vnetd` directory in the following location:

- 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`

General Server properties

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

Figure 3-28 General Server dialog box



The **General Server** dialog box contains the following properties.

Table 3-34 General Server dialog box properties

Property	Description
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.

Table 3-34 General Server dialog box properties (*continued*)

Property	Description
Check the capacity of disk storage units every	<p>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.</p> <p>The default is 300 seconds (5 minutes).</p> <p>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.</p>
Must use local drive	<p>This property appears for master servers only, but applies to all media servers as well. This property does not apply to NDMP drives.</p> <p>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.</p> <p>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.</p>
Use direct access recovery for NDMP restores	<p>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.</p> <p>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.</p>
Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology	<p>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.</p> <p>See the <i>NetBackup for Exchange Administrator's Guide</i>.</p>

Table 3-34 General Server dialog box properties (*continued*)

Property	Description
Media host override list	<p>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.</p> <p>The following items describe situations in which the capability to specify servers is useful:</p> <ul style="list-style-type: none"> ■ 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. <p>To add a host to the Media host override list, click Add.</p> <p>Click Add to All to add a host to the list for all of the hosts currently selected.</p> <p>To change an entry in the list, select a host name, then click Change.</p> <p>Configure the following options in the Add Media Override settings or Change Media Override settings dialog box:</p> <ul style="list-style-type: none"> ■ 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.
Defaults	Sets all properties back to the default settings.

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 Daemon (bprad) on the master server.

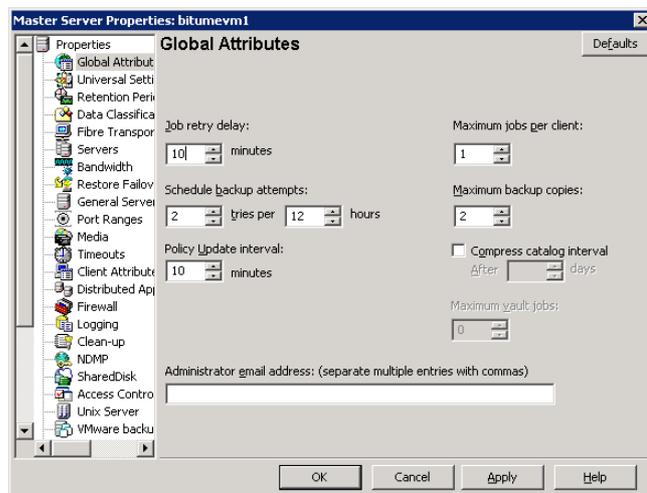
This process applies to all storage units on the original backup server. Restores for any storage unit on the **Original backup server** go to the server that is listed as the **Restore server**.

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

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 3-29 Global Attributes dialog box



The **Global Attributes** dialog box contains the following properties.

Table 3-35 Global Attributes dialog box properties

Property	Description
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	<p>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.</p> <p>This property determines the following characteristics for each policy:</p> <ul style="list-style-type: none"> ■ 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. <p>Consider the following example scenario using the default setting 2 tries every 12 hours:</p> <ul style="list-style-type: none"> ■ 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. <p>Note: This attribute does not apply to user backups and archives.</p>
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	<p>This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. The default is one job.</p> <p>NetBackup can process concurrent backup jobs from different policies on the same client only in the following situations:</p> <ul style="list-style-type: none"> ■ More than one storage unit available ■ One of the available storage units can perform more than one backup at a time. <p>See “About constraints on the number of concurrent jobs” on page 133.</p>

Table 3-35 Global Attributes dialog box properties (*continued*)

Property	Description
Maximum backup copies	<p>This property specifies the total number of backup copies that can exist in the NetBackup catalog (2 through 10)</p> <p>NetBackup creates one of the following, whichever is smaller:</p> <ul style="list-style-type: none"> ■ The number of copies that are specified under Multiple copies See “Multiple copies (schedule attribute)” on page 562. ■ The number of copies that are specified as the Maximum backup copies property <p>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.</p> <p>See “About configuring for multiple copies” on page 563.</p>
Compress catalog interval	<p>This property specifies how long NetBackup waits after a backup before it compresses the image catalog file.</p>
Maximum vault jobs	<p>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.</p> <p>See “About the Jobs tab” on page 766.</p>
Administrator email address property	<p>This property specifies the address(es) where NetBackup sends notifications of scheduled backups or administrator-directed manual backups.</p> <p>To send the information to more than one administrator, separate multiple email addresses by using a comma, as follows:</p> <p><i>useraccount1@company.com, useraccount2@company.com</i></p> <p>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.</p> <p>See “Disaster Recovery tab” on page 631.</p>

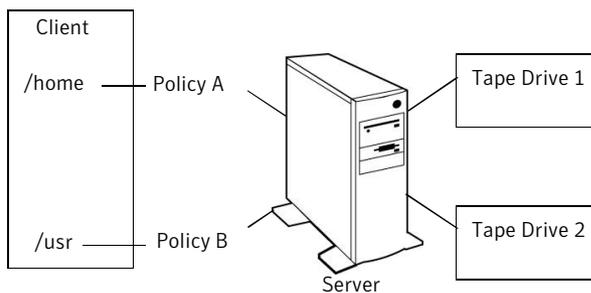
About constraints on the number of concurrent jobs

Specify any number of concurrent jobs within the following constraints.

Table 3-36 Constraints on concurrent jobs

Constraint	Description
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	<p>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.</p> <p>The Maximum jobs per client property applies to all clients in all policies.</p> <p>To accommodate weaker clients (ones that can handle only a small number of jobs concurrently), consider using one of the following approaches:</p> <ul style="list-style-type: none"> ■ 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 “General tab of the Client Attributes properties” on page 80. ■ 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 (policy attribute)” on page 525.
Network loading	<p>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.</p> <p>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.</p>

Figure 3-30 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.

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.

To set up email notifications about backups, 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 137.
- See “Sending messages to the global administrator about unsuccessful backups” on page 138.
- 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 and testing the email utility” on page 139.

See “About constraints on the number of concurrent jobs” on page 133.

Configuring the nbmail.cmd script

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

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

To configure the `nbmail.cmd` script

- 1 On a NetBackup master server, locate
`install_path\VERITAS\NetBackup\bin\goodies\nbmail.cmd`.
- 2 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.
- 3 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.

The following options are used in the script:

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

4 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 )
```

5 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
)
```

6 Save `nbmail.cmd`.

See “About constraints on the number of concurrent jobs” on page 133.

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 On the server, install and configure a mail client.
See “Installing and testing the email utility” on page 139.
- 2 On the server, edit the `nbmail.cmd` script.
See “Configuring the `nbmail.cmd` script” on page 136.
- 3 On the master server, in the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 4 In the right pane, double-click the master server you want to modify.
- 5 In the properties dialog box, in the left pane, click **Universal Settings**.

- 6 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 201.
- 7 Enable the **Server sends mail** option and click **Apply**.

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 On the server, install and configure a mail client.
See "Installing and testing the email utility" on page 139.
- 2 On the server, edit the `nbmail.cmd` script.
See "Configuring the nbmail.cmd script" on page 136.
- 3 On the master server, open the **NetBackup Administration Console**.
- 4 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Server**.
- 5 In the right pane, double-click the master server you want to modify.
- 6 Open the host properties of the master server.
- 7 In the properties dialog box, in the left pane, click **Global Attributes**.
- 8 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
```

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 On the client, install and configure a mail client.
See “Installing and testing the email utility” on page 139.
- 2 On the client, edit the `nbmail.cmd` script.
See “Configuring the `nbmail.cmd` script” on page 136.
- 3 On the master server, open the **NetBackup Administration Console**.
- 4 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Clients**.
- 5 In the right pane, double-click the client you want to modify. Multiple clients can also be selected.
- 6 In the properties dialog box, in the left pane, 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 201.
- 8 Enable the **Client sends mail** option and click **Apply**.

Installing and testing 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.

- 5 Test the email utility, following the To test the email utility procedure.

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.
 - 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 `-ti 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.

Table 3-37 Logging types

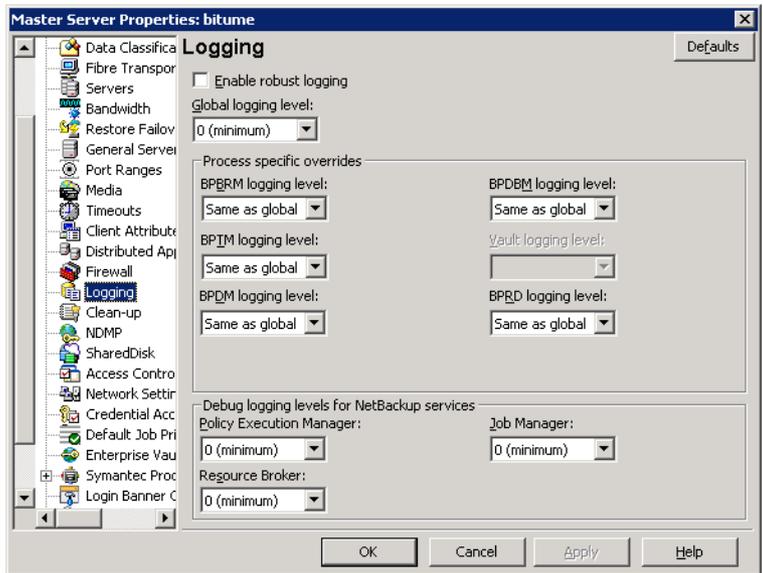
Type	Description
Unified logging	<p>Unified logging creates log file names and messages in a format that is standardized across Symantec products. Some NetBackup processes on the server use unified logging.</p> <p>Unified logging writes the logs into subdirectories in the following locations:</p> <ul style="list-style-type: none">■ UNIX: <code>/usr/opensv/logs</code>■ Windows: <code>install_path\NetBackup\logs</code> <p>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.</p> <p>Unlike legacy logging, subdirectories for the processes that use unified logging are created automatically.</p> <p>To control the size and number of unified logs, use the <code>vxlogcfg</code> command and the <code>vxlogmgr</code> command.</p>

Table 3-37 Logging types (*continued*)

Type	Description
Legacy logging	<p>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.</p> <p>Create the NetBackup legacy log directories in the following locations:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/usr/opensv/netbackup/logs/process_name</code> ■ Windows: <code>install_path\NetBackup\logs\process_name</code> <p>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.</p> <p>On a Windows server, you can create all of the NetBackup debug log directories at one time by double-clicking <code>mklogdir.bat</code> in the following directory:</p> <p><code>install_path\NetBackup\logs\</code></p> <p>Create the Media Manager legacy log directories in the following locations:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/usr/opensv/volmgr/debug</code> ■ Windows: <code>install_path\Volmgr\debug</code>

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

Figure 3-31 Logging dialog box



The **Logging** dialog box contains the following properties.

Table 3-38 Logging dialog box properties

Property	Description
<p>Enable robust logging</p>	<p>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.</p> <p>See the <i>NetBackup Troubleshooting Guide</i> for more information about controlling the log file size.</p> <p>If this property is enabled, the following processes produce log files:</p> <ul style="list-style-type: none"> ■ bprd ■ bpbkar ■ bpbrm ■ bpcd ■ bpdbrm ■ bptm ■ bpdm <p>The logs are named using the following convention:</p> <p><i>MMDDYY_NNNNN.log</i></p> <p>where <i>NNNNN</i> is an incrementing counter from 00001 - 99999</p> <p>If the Enable robust logging property is disabled, a single log file is produced each day:</p> <p><i>MMDDYY.log</i></p> <p>Whether Enable robust logging is selected or not, the log file is pruned by using <code>KEEP_LOGS_DAYS</code> and <code>DAYS_TO_KEEP_LOGS</code> settings.</p> <p>Note: If a NetBackup environment uses scripts depending on the <i>MMDDYY.log</i> naming convention, either update the scripts or disable Robust Logging.</p>
<p>Global logging level</p>	<p>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).</p> <p>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.</p> <p>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.</p>

Table 3-38 Logging dialog box properties (*continued*)

Property	Description
Process specific overrides	<p>The services that are listed under Process specific overrides use legacy logging. These services require that you first create a log directory in the following location:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/usr/opensv/netbackup/logs/process_name</code> ■ Windows: <code>install_path\NetBackup\logs\process_name</code> <p>Table 3-39 lists and describes the processes that use legacy logging.</p>
Debug logging levels for NetBackup services	<p>The Logging properties page offers configurable debug levels for the services that use unified logging.</p> <p>Each service creates a log automatically in the following directories:</p> <ul style="list-style-type: none"> ■ UNIX: <code>/usr/opensv/logs</code> ■ Windows: <code>install_path\NetBackup\logs</code> <p>You can also use the <code>vxlogcfg</code> command to change debug levels.</p> <p>Table 3-40 lists and describes the services that use unified logging.</p>

Table 3-39 Process specific overrides

Service	Description
BPBRM logging level	The NetBackup backup and restore manager.
BPTM logging level	The NetBackup tape manager.
BPDM logging level	The NetBackup disk manager.
BPDBM logging level	The NetBackup database manager.
Vault logging level	Select a logging level for <code>bpvault</code> .
BPRD logging level	The NetBackup request daemon or manager.

Table 3-40 Debug logging levels for NetBackup services

Service	Description
Policy Execution Manager	This property appears for EMM servers. 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. NBJM accepts the jobs that the Policy Execution Manager submits and acquires the necessary resources.
Resource Broker	NBRB makes the allocations for storage units, tape drives, client reservations.

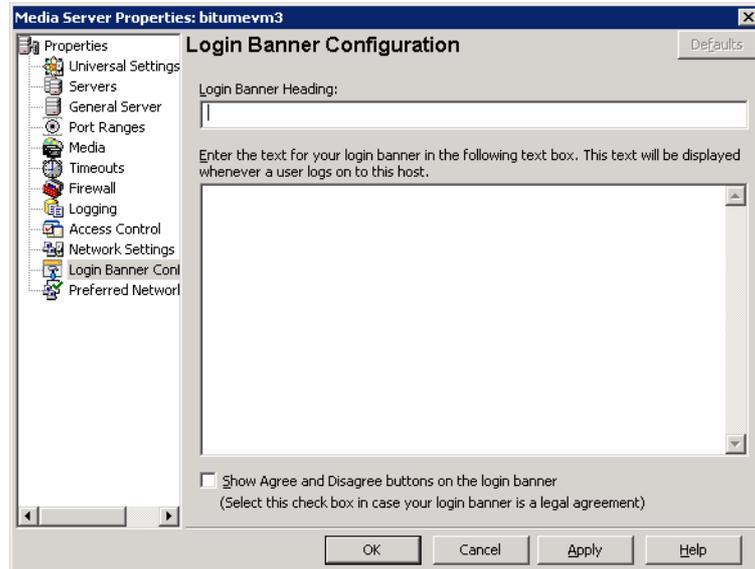
Login Banner Configuration properties

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 3-32 shows example banner text for a media server.

Figure 3-32 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. (On the **File** menu, click **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. (On the **File** menu, click the **New Console** option or the **New Window from Here** option.)

Table 3-41 Login Banner Configuration dialog box properties

Property	Description
Login Banner Heading	Enter the text that is to appear in the banner.

Table 3-41 Login Banner Configuration dialog box properties (continued)

Property	Description
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	<p>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.</p> <p>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.</p> <p>If the user chooses the Disagree option, the screen is closed.</p>

Figure 3-33 Login Banner with agreement option enabled

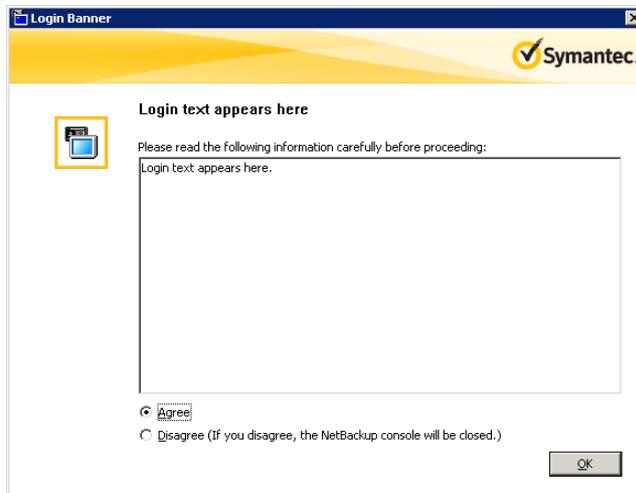
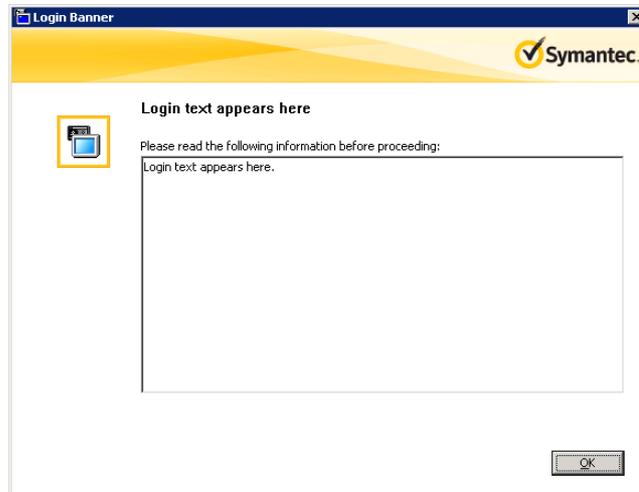


Figure 3-34 Login Banner without agreement option



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 Depending on the host that displays the login banner, select **Master Servers**, **Media Servers**, or **Clients**.
- 3 In the right pane, double-click the host name to display the properties.
- 4 In the properties dialog box, in the left pane, click the **Login Banner Configuration** host properties.
- 5 Clear the **Login Banner Heading** text and the login banner text.
- 6 Click **OK** to save the changes.

Enabling the 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 enable the Auto log off timeout option

- 1 Select **View > Options**. Then select the **Administration Console** tab.
- 2 Check the **Auto log off timeout** option.
- 3 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.
- 4 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 extends 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**

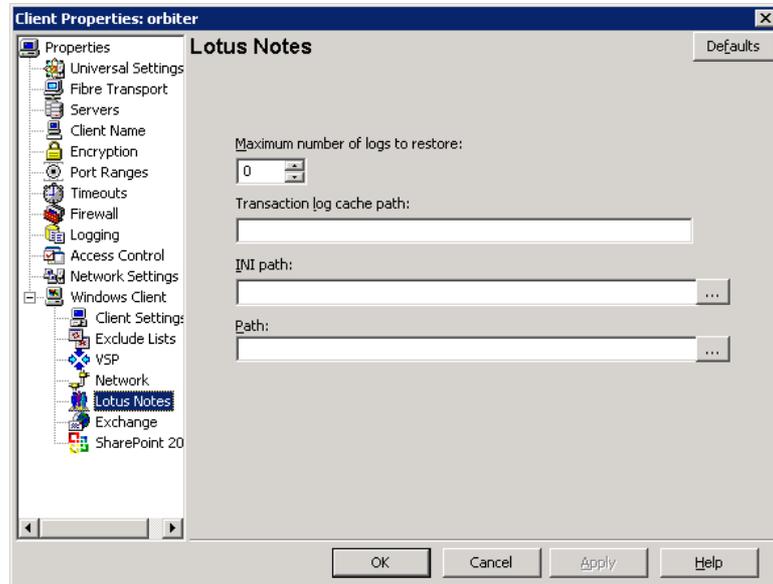
If the user selects this option, the session ends and NetBackup logs off the user immediately.

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 3-35 Lotus Notes dialog box



For UNIX or Linux servers: If you have multiple installations of Domino server, the values in the client properties or the `bp.conf` only apply to one installation. For other installations, specify the installation path and location of the `notes.ini` file with the `LOTUS_INSTALL_PATH` and `NOTES_INI_PATH` directives in the backup policy.

Table 3-42 Lotus Notes client host properties

Client host properties	Windows registry and <code>bp.conf</code> entries	Description
Maximum number of logs to restore	<p>LOTUS_NOTES_LOGCACHESIZE</p> <p>In the Windows registry, this value is a <code>DWORD</code> value.</p>	<p>The maximum number of logs that can be prefetched in a single restore job during recovery. Specify a value greater than 1.</p> <p>If this value is less than or equal to 1, NetBackup does not gather transaction logs during recovery. One transaction log extent per job is restored to the Domino server's log directory.</p> <p>LOTUS_NOTES_LOGCACHESIZE = 3</p>

Table 3-42 Lotus Notes client host properties (continued)

Client host properties	Windows registry and bp.conf entries	Description
Transaction log cache path	<p>LOTUS_NOTES_LOGCACHEPATH</p> <p>In the Windows registry, this value is a string value.</p>	<p>Specify a path where NetBackup can temporarily store the prefetched transaction logs during recovery.</p> <p>For example:</p> <ul style="list-style-type: none"> ■ UNIX: /tmp/logcache ■ Windows: D:\LogCache <p>If you do not specify a path, during recovery NetBackup restores the logs to the Domino server's transaction log directory.</p> <p>Note the following before specifying the Transaction log cache path:</p> <ul style="list-style-type: none"> ■ 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, NetBackup tries to restore only the number of logs that can be accommodated.
INI path	<p>LOTUS_NOTES_INI</p> <p>In the Windows registry, this value is a string value.</p>	<p>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.</p> <p>Specify the absolute path to the NOTES.INI file:</p> <ul style="list-style-type: none"> ■ 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 <p>Include the directory and the notes.ini file name.</p>

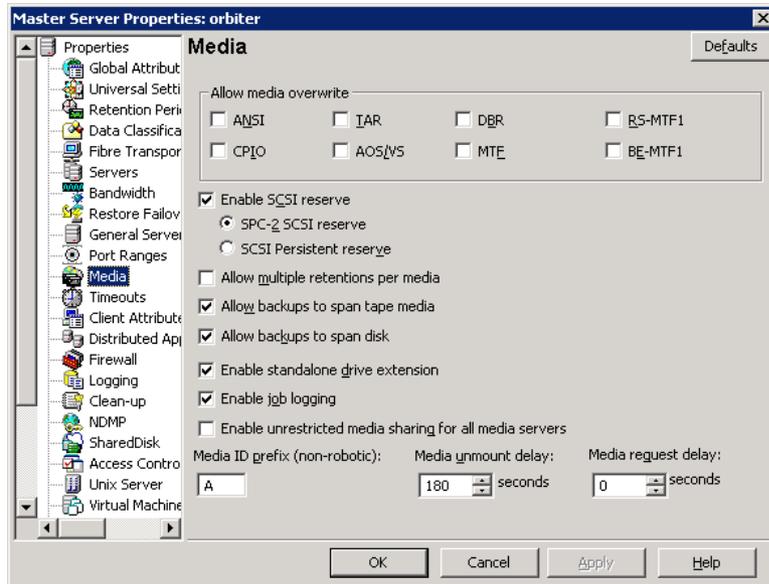
Table 3-42 Lotus Notes client host properties (continued)

Client host properties	Windows registry and bp.conf entries	Description
Path	<p>LOTUS_NOTES_PATH</p> <p>In the Windows registry, this value is a string value.</p>	<p>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.</p> <p>Specify the path where the Lotus Notes program files reside on the client:</p> <ul style="list-style-type: none"> ■ 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 <p>The Path value overrides the Lotus registry value, if both are defined.</p>

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 3-36 Media dialog box



The **Media** dialog box contains the following properties.

Table 3-43 Media dialog box properties

Property	Description
<p>Allow media overwrite property</p>	<p>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.</p> <p>For example, place a check in the CPIO check box to permit NetBackup to overwrite the cpio format.</p> <p>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.</p> <p>The following media formats on removable media can be selected to be overwritten:</p> <ul style="list-style-type: none"> ■ 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. <p>See "Results when media overwrites are not permitted" on page 158.</p>

Table 3-43 Media dialog box properties (*continued*)

Property	Description
<p>Enable SCSI reserve</p>	<p>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.</p> <p>SCSI reservations provide protection for NetBackup Shared Storage Option environments or any other multiple-initiator environment in which drives are shared.</p> <p>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.</p> <p>See “Recommended use for Enable SCSI reserve property” on page 159.</p> <p>See “Drive path options” on page 254.</p> <p>The following are the protection options:</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p>
<p>Allow multiple retentions per media</p>	<p>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.</p>
<p>Allow backups to span tape media</p>	<p>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.</p> <p>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.</p>

Table 3-43 Media dialog box properties (*continued*)

Property	Description
Allow backups to span disk	<p>This property allows backups to span disk volumes when one disk volume becomes full. The default is that this property is enabled.</p> <p>The Allow backups to span disk property does not apply to AdvancedDisk or OpenStorage storage units. Backups span disk volumes within disk pools automatically.</p> <p>The following destinations support disk spanning:</p> <ul style="list-style-type: none"> ■ 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. <p>For disk spanning to occur, the following conditions must be met:</p> <ul style="list-style-type: none"> ■ 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 storage unit setting” on page 402. ■ 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	<p>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.</p>
Enable job logging	<p>This property allows the logging of the job information. This logging is the same information that the NetBackup Activity Monitor uses. The default is that job logging occurs.</p>
Enable unrestricted media sharing for all media servers	<p>This property controls media sharing, as follows:</p> <ul style="list-style-type: none"> ■ 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. <p>The default is that media sharing is disabled. (The property is cleared and no server groups are configured.)</p> <p>See “About server groups” on page 209.</p>

Table 3-43 Media dialog box properties (*continued*)

Property	Description
Media ID prefix (non-robotic)	<p>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.</p> <p>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).</p>
Media unmount delay	<p>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.</p> <p>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.</p>
Media request delay	<p>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.</p> <p>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.</p> <p>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.</p>

Results when media overwrites are not permitted

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 | <ul style="list-style-type: none"> ■ Sets the volume's state to FROZEN ■ Selects a different volume ■ Logs an error |
|---|--|

<p>If the volume is in the NetBackup media catalog and was previously selected for backups</p>	<ul style="list-style-type: none"> ■ Sets the volume's state to SUSPENDED ■ Aborts the requested backup ■ Logs an error
<p>If the volume is mounted for a backup of the NetBackup catalog</p>	<p>The backup is aborted and an error is logged. The error indicates the volume cannot be overwritten.</p>
<p>If the volume is mounted to restore files or list the media contents</p>	<p>NetBackup aborts the request and logs an error. The error indicates that the volume does not have a NetBackup format.</p>

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 provides device status and correction and may be more effective in the following environments:

- Where 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.)
- Where the drive has 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.

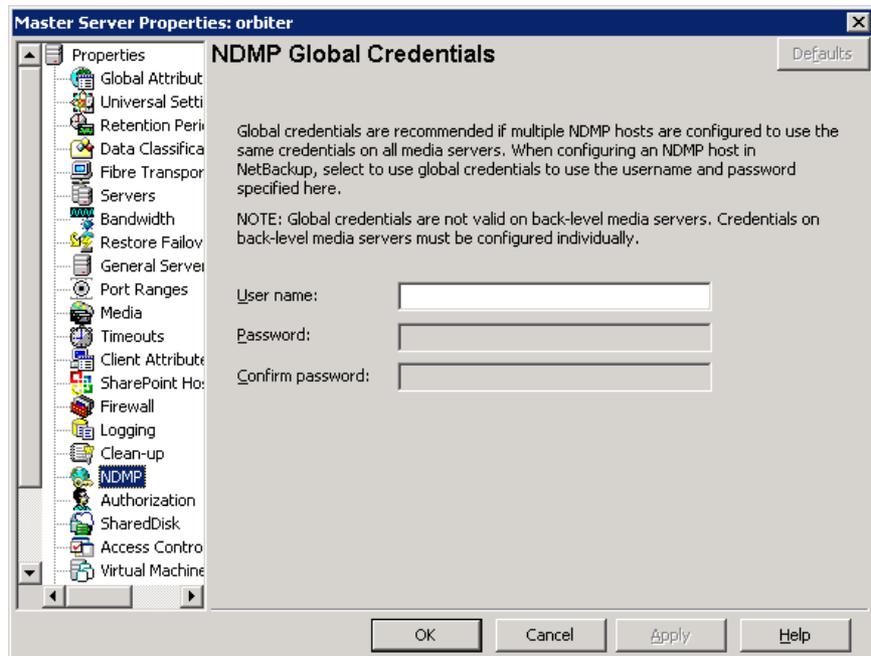
- The operating system tape drivers may require extensive configuration to use SCSI persistent reserve. For example, if the tape drives do not support SPC-3 Compatible Reservation Handling (CRH), ensure that the operating system does not issue SPC-2 reserve and release commands.

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 3-37 NDMP Global Credentials dialog box



The **NDMP Global Credentials** properties dialog box contains the following properties.

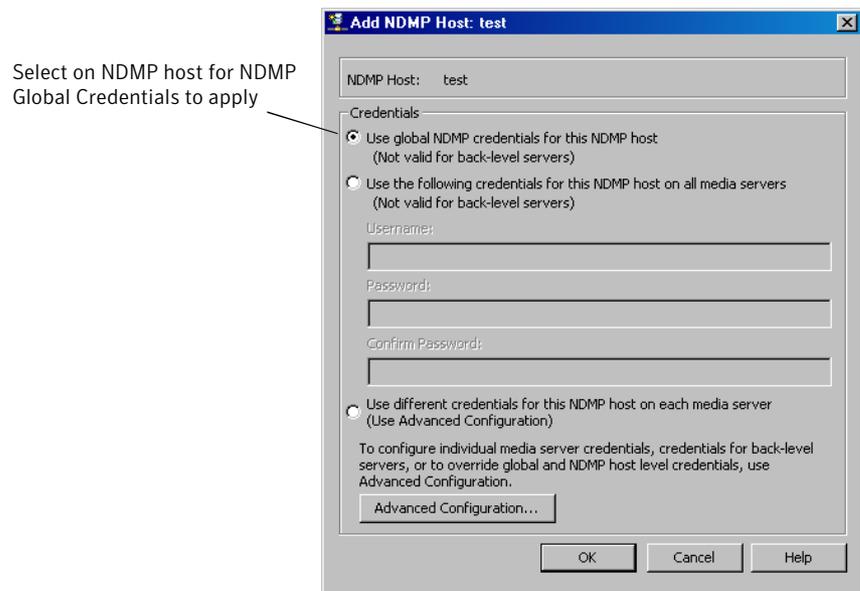
Table 3-44 NDMP Global Credentials dialog box properties

Property	Description
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 3-38 shows the **Add NDMP Host** dialog box. In the **Credentials** section, select **Use global NDMP credentials for this NDMP host** so that the **NDMP Global Credentials** apply to that host.

Figure 3-38 Add NDMP Host dialog box



NetWare Client properties

The **Netware Client** properties define NetBackup properties of NetWare clients.

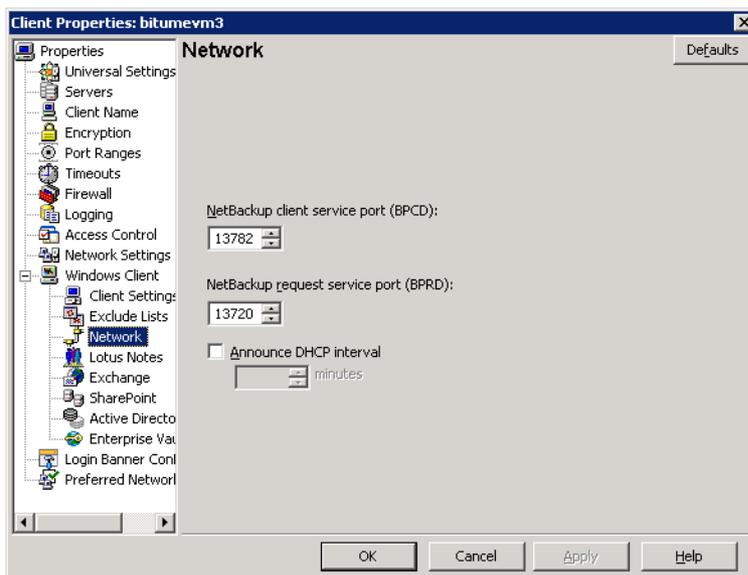
Netware Client properties include the **Client Settings** for NetWare clients as a subnode:

See “Client Settings properties for NetWare clients” on page 91.

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 3-39 Network dialog box



The **Network** dialog box contains the following properties.

Table 3-45 Network dialog box properties

Property	Description
NetBackup client service port (BPCD)	<p>This property specifies the port that the NetBackup client uses to communicate with the NetBackup server. The default is 13782.</p> <p>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.</p>

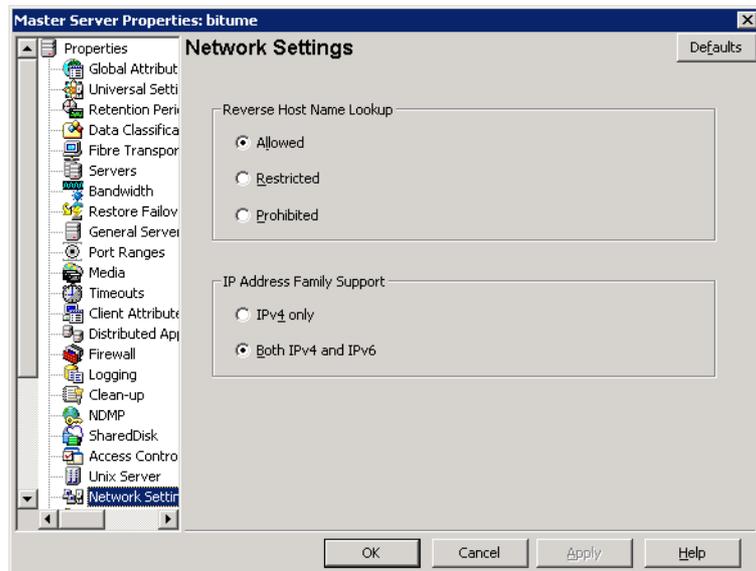
Table 3-45 Network dialog box properties (continued)

Property	Description
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 **Network Settings** host properties apply to master servers, media servers, and clients.

Figure 3-40 Network Settings dialog box



The **Network Settings** dialog box contains properties for **Reverse Host Name Lookup** and **IP Address Family Support**.

Reverse Host Name Lookup property

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.

Table 3-46 Reverse Host Name Lookup property settings

Property	Description
Allowed setting	<p>The Allowed property indicates that the host requires reverse host name lookup to work to determine that the connection comes from a recognizable server.</p> <p>By default, the host resolves the IP address of the connecting server to a host name by performing a reverse lookup.</p> <p>If the conversion of the IP address to host name fails, the connection fails.</p> <p>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.</p>
Restricted setting	<p>The Restricted property 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.</p> <p>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.</p> <p>If the comparison fails, the host rejects the connection from server and the connection fails.</p>
Prohibited setting	<p>The Prohibited property 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.</p> <p>The NetBackup host then compares the IP address of the connecting server to the list of known server IP addresses.</p> <p>If the comparison fails, the NetBackup host rejects the connection from the server and the connection fails.</p>

Reverse Host Name Lookup changes 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 it. Add the `REVERSE_NAME_LOOKUP` entry to the `bp.conf` file (UNIX) or to the Windows registry.

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 REVERSE_NAME_LOOKUP 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 set the `REVERSE_NAME_LOOKUP` property on UNIX hosts, 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 the *NetBackup Commands Reference Guide*.
- On UNIX clients
Edit the `bp.conf` directly to add the entry.

Setting the REVERSE_NAME_LOOKUP property on Windows hosts

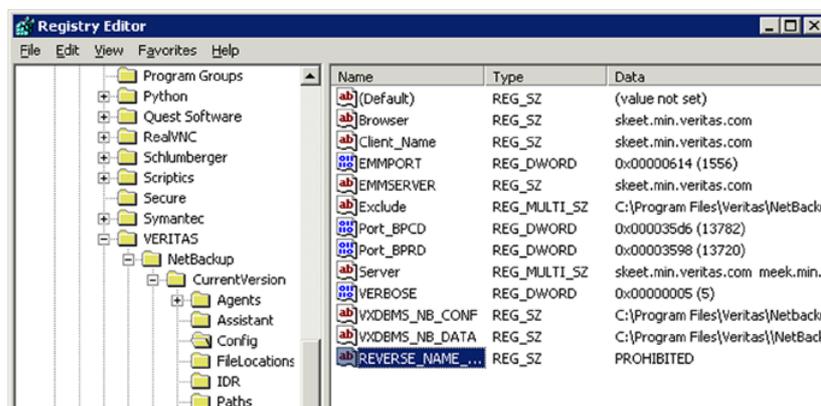
On master and media servers, the `bpsetconfig` command is available to add the `REVERSE_NAME_LOOKUP` entry to the registry. The `bpsetconfig` command is described in the *NetBackup Commands Reference Guide*.

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.

To set the Reverse Host Name Lookup property on a Windows client

- 1 From the command line, run `regedit` to open the registry editor.
- 2 Navigate to the following key directory:

```
My Computer\HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\  
NetBackup\CurrentVersion\Config
```
- 3 On the **Edit** menu, click **New > String Value**.
- 4 Name the String Value: `REVERSE_NAME_LOOKUP`.
- 5 Give `REVERSE_NAME_LOOKUP` the value data of either **PROHIBITED**, **RESTRICTED**, or **ALLOWED**.
- 6 Click **OK** and close the Registry Editor.



IP Address Family Support property

On hosts that use both IPv4 and IPv6 addresses, use the **IP Address Family Support** property to indicate which address family to use:

- **IPv4 only** (Default)
- **Both IPv4 and IPv6**

Upon installation or upgrade to NetBackup version 7.1, NetBackup defaults to IPv4. If any of the master servers do not support IPv4, NetBackup uses the configuration that supports both IPv4 and IPv6.

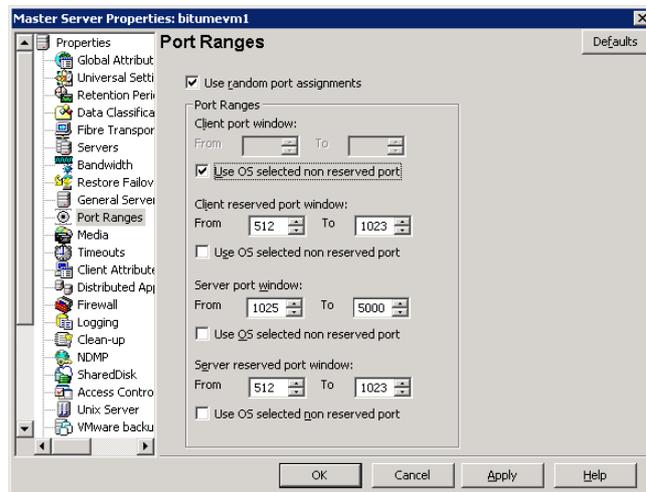
While the **IP Address Family Support** property controls how hostnames are resolved to IP addresses, the **Preferred Network** properties control how NetBackup uses the addresses.

See “Preferred Network properties” on page 169.

Port Ranges properties

Use the **Port Ranges** properties in the **NetBackup Administration Console** to determine how hosts connect to one another. These properties apply to selected master servers, media servers, and clients.

Figure 3-41 Port Ranges dialog box



The **Port Ranges** dialog box contains the following properties.

Table 3-47 Port Ranges dialog box properties

Property	Description
Use random port assignments	<p>Specifies how the selected computer chooses a port when it communicates with NetBackup on other computers. Enable this property to let NetBackup randomly select ports 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.</p> <p>If this property is not enabled, NetBackup chooses numbers sequentially, not randomly. 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 5000 is in use, port 4999 is chosen.</p> <p>This property is enabled by default.</p>

Table 3-47 Port Ranges dialog box properties (*continued*)

Property	Description
Client port window	Lets the administrator define the range of non-reserved ports on the selected computer. NetBackup can use any available port within this range to communicate with NetBackup on another computer.
Use OS selected non reserved port	Lets the operating system determine which non-reserved port to use.
Client reserved port window	This property no longer applies to NetBackup 7.0.1 and later. For information about this property, refer to documentation from a previous release.
Server port window	This property no longer applies to NetBackup 7.0.1 and later. For information about this property, refer to documentation from a previous release.
Server reserved port window	This property no longer applies NetBackup 7.0.1 and later. For information about this property, refer to documentation from a previous release.

See “Registered ports and dynamically-allocated ports” on page 168.

Registered ports and dynamically-allocated ports

NetBackup communicates between computers by using a combination of registered ports and dynamically-allocated ports.

Registered ports

These 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.

The following system configuration file can be used to override the default port numbers for each port:

```
%systemroot%\system32\drivers\etc\services
```

Dynamically-allocated ports

These 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 specify whether NetBackup selects a port number at random or starts at the top of the range and uses the first one available.

Preferred Network properties

Use the **Preferred Network** properties in the **NetBackup Administration Console** to specify to NetBackup which networks or interfaces to use for outgoing NetBackup traffic from the selected hosts. These properties apply to currently selected master servers, media servers, and clients.

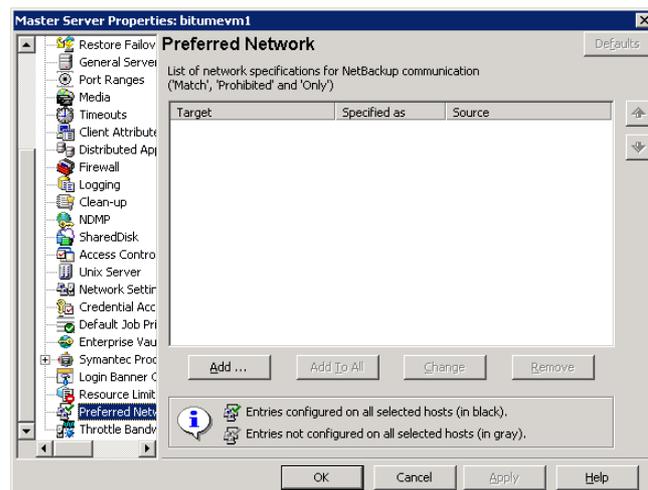
The **Preferred Network** properties are useful in NetBackup environments that include multihomed hosts—the hosts that are connected to two or more networks, or hosts that have two or more network addresses. The properties are especially helpful to administrators who must configure an environment that includes both Internet Protocol version 6 (IPv6) and IPv4 address families.

The **Preferred Network** properties compare to the **Use specified network interface** property in the **Universal Settings** properties. However, the **Use specified network interface** property can be used to specify only a single interface for NetBackup to use for outbound calls. The **Preferred Network** properties were introduced so that administrators can give more elaborate and constrictive instructions that apply to multiple individual networks, or a range of networks. For example, an administrator can configure a host to use any network except one.

Note: Do not inadvertently configure hosts so that they cannot communicate with any other host. Use the `bptestnetconn` utility to determine whether the hosts can communicate as you intend.

See “`bptestnetconn` utility to display Preferred Network information” on page 178.

Figure 3-42 Preferred Network dialog box



The **Preferred Network** dialog box contains a list of networks and the directive that has been configured for each.

Table 3-48 Preferred Network dialog box properties

Property	Description
List of network specifications for NetBackup communications	<p>The list of preferred networks contains the following information:</p> <ul style="list-style-type: none"> ■ The Target column lists the networks (or hostnames or IP addresses) that have been given specific directives. If a network is not specifically listed as a target, or if a range of addresses does not include the target, NetBackup considers the target to be available for selection. <p>Note that if the same network considerations apply for all of the hosts, the list of directives can be identical across all hosts in the NetBackup environment. If a directive contains an address that does not apply to a particular host, that host ignores it. For example, an IPv4-only host ignores IPv6 directives, and IPv6-only hosts ignore IPv4 directives. This lets the administrator use the same Preferred Network configurations for all the hosts in the NetBackup environment.</p> <ul style="list-style-type: none"> ■ The Specified as column indicates the directive for the network: Match, Prohibited, or Only. ■ The Source column lists source binding information to use to filter addresses. The Source property is an optional configuration property.
Ordering arrows	<p>Select a network in the list, then click the up or down arrow to change the order of the network in the list. The order can affect which network NetBackup selects.</p> <p>See “Order of directive processing in the Preferred Network properties” on page 177.</p>
Add	<p>Click Add to add a network to the Preferred Network properties. The directive for the network is configured in the Add Preferred Network Settings dialog box.</p> <p>See Table 3-49 on page 171.</p>
Add to all	<p>The Add to all button is active when multiple servers are selected.</p>
Change	<p>Select a network in the list, then click Change to change the Preferred Network properties. The directive is changed in the Change Preferred Network Settings dialog box.</p> <p>See “Add or Change Preferred Network Settings dialog box” on page 170.</p>
Remove	<p>Select a network in the list, then click Remove to remove the network from the list of preferred networks.</p>

Add or Change Preferred Network Settings dialog box

The **Add Preferred Network Settings** dialog box contains the following properties.

Table 3-49 Add or Change Preferred Network Settings dialog box properties

Property	Description
Target	<p>Enter a network address or a hostname:</p> <ul style="list-style-type: none"> ■ If an address is specified as the network, it is usually considered a remote or target address. NetBackup recognizes the following wildcard entries as addresses: <ul style="list-style-type: none"> ■ 0.0.0.0 Matches any IPv4 address. ■ 0::0 Matches any IPv6 address. ■ 0/0< /> Matches the address of any family. ■ If a hostname is specified as the network, then the address that is used is the first returned by the DNS resolver. <p>Note: Do not use the following malformed entries as wildcards: 0/32, 0/64, or 0/128. The left side of the slash must be a legitimate IP address. However, 0/0 may be used, as listed.</p>
Match	<p>The Match directive indicates that the specified network, address, or hostname is preferred for communication with the selected host.</p> <p>The Match directive does not reject other networks, addresses, or hostnames from being selected, even if they do not match. (The Only directive rejects unsuitable targets if they do not match.)</p> <p>The Match directive is useful following a Prohibited or a Only directive. When used with other directives, Match indicates to NetBackup to stop rule processing because a suitable match has been found.</p> <p>The Match directive can be used with the Source property to indicate source binding.</p>
Prohibited	<p>Use the Prohibited directive to exclude or prevent the specified network, address, or hostname from being considered. In a list of DNS addresses, addresses in these networks are avoided.</p>
Only	<p>The Only directive indicates that the specified network, address, or hostname that is used for communication with the selected host must be in the specified network.</p> <p>Use the Only directive to prevent any network from being considered other than those specified as Only.</p> <p>This directive replaces the <code>REQUIRED_NETWORK</code> entry in the <code>bp.conf</code> file or registry.</p> <p>The Only directive can be used with the Source property to indicate source binding.</p>

Table 3-49 Add or Change Preferred Network Settings dialog box properties
(continued)

Property	Description
Source	<p>Use this property with the Match or the Only directives to describe the local hostname, IP addresses, or networks that may be used for source binding.</p> <p>NetBackup matches the desired source interfaces, (backup networks, for example) with the target addresses described by the Source property.</p> <p>The corresponding <code>bp.conf</code> or registry entry for this property is <code>PREFERRED_NETWORK</code>. This property replaces the <code>REQUIRED_INTERFACE</code> entry.</p>

How NetBackup uses the directives to determine which network to use

Each host has an internal table of preferred network rules that NetBackup consults before it selects a network interface to use for communication with another host. The table includes every interface-IP address combination available to the selected host. Based on the **Preferred Network** directives, the table indicates to NetBackup whether or not the host is allowed to use a given network.

This topic uses the example of two multihomed servers (Server_A and Server_B) as shown in Figure 3-43. Server A is considering which addresses it can use to access Server_B, given the **Preferred Network** directives configured on Server_A.

When **Preferred Network** directives are used to place restrictions on targets, they are added from the perspective of the server making the connection. The directives on Server_A affect its preferences as to which Server_B addresses it can use.

Figure 3-43 Multihomed servers example

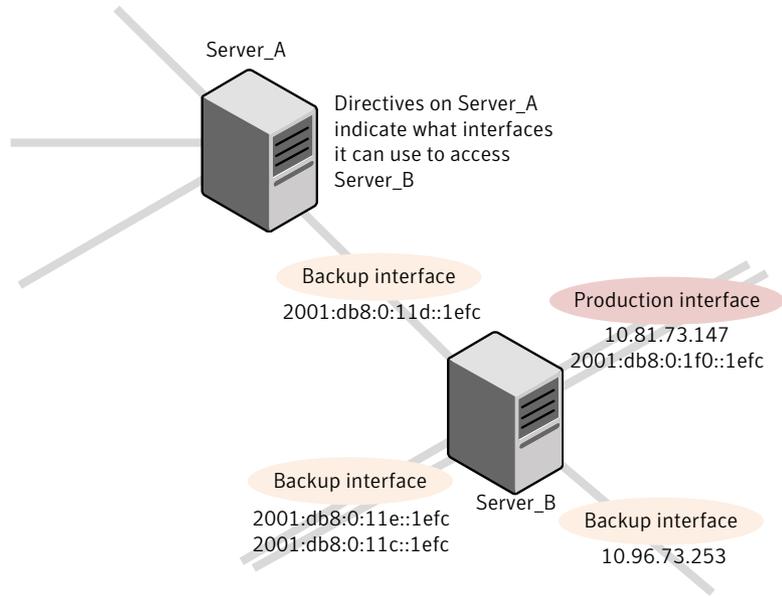


Figure 3-44 shows a table for Server_B. Server_B has multiple network interfaces, some of which have multiple IP addresses. In the table, *yes* indicates that NetBackup can use the network-IP combination as a source. In this example, no directives have been created for the host. Since no networks are listed in the **Preferred Network** properties, any network-IP combinations can be used for communication.

Note: Figure 3-51 shows the `bptestnetconn` output for this example configuration.

Figure 3-44 From Server_A's perspective: Available IP addresses on Server_B when no directives are indicated on Server_A

		IP addresses	
		IPv4	IPv6
Network interfaces	2001:0db8:0:1f0::1efc	---	Yes
	10.80.73.147	Yes	---
	2001:0db8:0:11c::1efc	---	Yes
	2001:0db8:0:11d::1efc	---	Yes
	2001:0db8:0:11e::1efc	---	Yes
	10.96.73.253	Yes	---

Figure 3-45 shows a table for the same host (Server_B). Now, the **Preferred Network** properties are configured so that all IPv4 addresses are excluded from selection consideration by NetBackup. All NetBackup traffic is to use only IPv6 addresses.

Figure 3-45 From Server_A's perspective: Available IP addresses on Server_B when directives to use IPv6 addresses only are indicated on Server_A

		IP addresses	
		IPv4	IPv6
Network interfaces	2001:0db8:0:1f0::1efc	---	Yes
	10.80.73.147	No	---
	2001:0db8:0:11c::1efc	---	Yes
	2001:0db8:0:11d::1efc	---	Yes
	2001:0db8:0:11e::1efc	---	Yes
	10.96.73.253	No	---

The following topics describe various configurations:

- See “Configurations to use IPv6 networks” on page 175.
- See “Configurations to use IPv4 networks” on page 176.
- See “Configuration to prohibit using a specified address” on page 180.
- See “Configuration that uses the same specification for both the network and the interface—less constrictive” on page 180.

- See “Configuration that uses the same specification for both the network and the interface—more constrictive” on page 181.
- See “Configuration that limits the addresses, but allows any interfaces” on page 182.

Configurations to use IPv6 networks

The following **Preferred Network** configurations instruct NetBackup to use only IPv6 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv6 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive (Figure 3-46) and one configuration uses the **Match** directive (Figure 3-47).

The more efficient method to specify one address family, (IPv6, in this case), is to prohibit IPv4. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

Figure 3-46 uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv4 addresses. In this situation, NetBackup must use an IPv6 address.

Note: The default configuration is for NetBackup to use only IPv4 addresses. Creating a directive that prohibits all IPv4 addresses renders the server mute unless you have IPv6 addresses and have them enabled.

See “IP Address Family Support property” on page 166.

Figure 3-46 Prohibit IPv4 addresses as targets

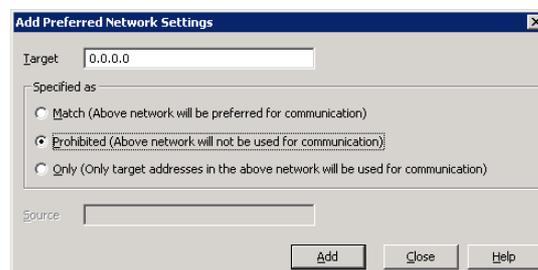


Figure 3-47 uses the **Match** directive with a wildcard to indicate to NetBackup to consider only IPv6 addresses. In this case, NetBackup tries to use an IPv6 address, but may consider IPv4 addresses if necessary.

Figure 3-47 Match IPv6 addresses as targets

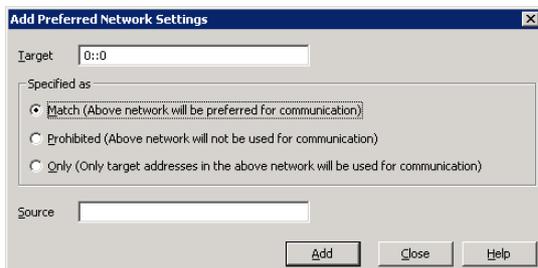
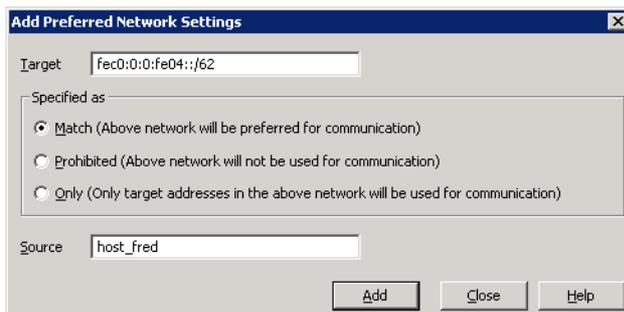


Figure 3-48 shows another configuration that allows NetBackup to choose from multiple IPv6 networks.

Given the multihomed example configuration in Figure 3-43, the directive indicates the following:

- Four IPv6 networks, from `fec0:0:0:fe04` to `fec0:0:0:fe07`, are described as targets.
- For all addresses in these networks, a source binding address that is derived from the IP addresses of hostname `host_fred` is used.

Figure 3-48 Indicating a range of IPv6 networks



Configurations to use IPv4 networks

The following **Preferred Network** configurations instruct NetBackup to use only IPv4 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv4 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive (Figure 3-49) and one configuration uses the **Match** directive (Figure 3-50).

The more efficient method to specify one address family, (IPv4, in this case), is to prohibit IPv6. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

Figure 3-49 uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv6 addresses. In this situation, NetBackup must use an IPv4 address.

Figure 3-49 Prohibit IPv6 addresses as targets

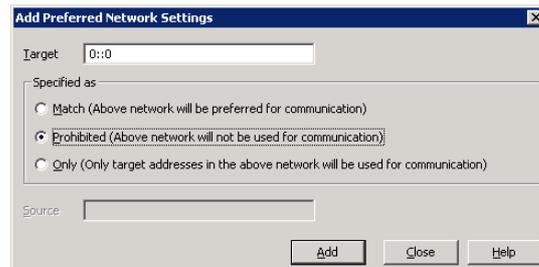
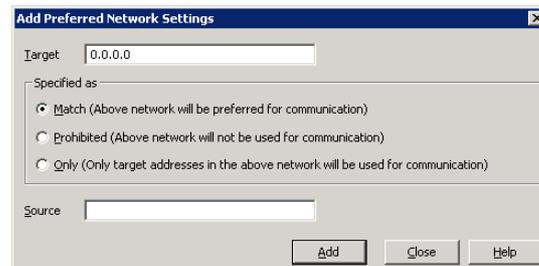


Figure 3-50 uses the **Match** directive with a wildcard to indicate to NetBackup to consider only IPv4 addresses. In this case, NetBackup tries to use an IPv4 address, but may consider IPv6 addresses if necessary.

Figure 3-50 Match IPv4 addresses as targets



Order of directive processing in the Preferred Network properties

NetBackup sorts all directives into decreasing order by subnet size so that the more specific network specifications, such as complete hostnames or IP addresses, match first. (For example, a /24 subnet matches before a /16 subnet.) In this way, NetBackup can honor host-specific overrides.

If NetBackup considers the directives of multiple networks to be equal in specificity (a tie), NetBackup looks at the order in which the networks are listed.

See “Order of directives can affect processing” on page 178.

NetBackup processes each resolved address in the network list according to specific rules. Directives that contain addresses that do not apply to the host are ignored.

Table 3-50 describes how NetBackup determines whether an address can be used for communication.

Table 3-50 Order of directive processing

Step	NetBackup considers the target	Target is selected or processing continues
1	<ul style="list-style-type: none"> ■ If the target is not a match for the directive, and ■ if the directive is an Only directive... 	<p>...then the target is treated as Prohibited, and processing stops for that target.</p> <p>NetBackup considers the next target.</p>
2	<ul style="list-style-type: none"> ■ If the target is a match for the directive, and ■ if the directive is a Prohibited directive... 	<p>...then the target is treated as Prohibited and processing stops for that target.</p> <p>NetBackup considers the next target.</p>
3	If the target is not a match...	<p>...then the processing continues.</p> <p>NetBackup considers the next directive in the list.</p>
4	If the target is a match...	<p>...then the directive is either Only or Match and further directive processing stops.</p> <p>An Only match is treated like a Match in terms of source binding computation. If no rules ever match, then the target is allowed, and no source binding is enforced.</p>

Order of directives can affect processing

The order of the networks in the list can affect which network NetBackup selects for communication for the selected hosts.

The strongest filters are **Prohibited** and **Only**.

Use the up or down arrows to the right of the list to change the order of the networks.

bptestnetconn utility to display Preferred Network information

The `bptestnetconn` utility is available to administrators to test and analyze host connections. Use the preferred network option (`--prefnet`) to display information about the preferred network configuration along with the forward lookup information of a host on the server list.

The `bptestnetconn` command is described in the *NetBackup Commands Reference Guide*.

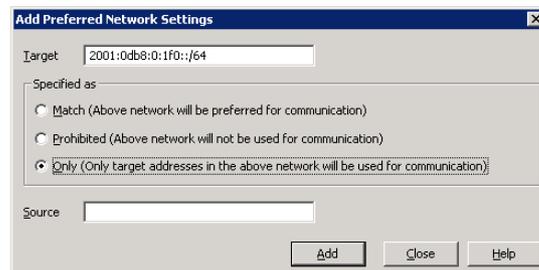
Figure 3-51 shows the `bptestnetconn` output when run on Server_A, for Server_B. That is, `bptestnetconn` is run from Server_A's perspective. Based on the directives configured on Server_A, for Server_B, `bptestnetconn` shows the available IP addresses on Server_B. In this example, no directives are configured on Server_A. (See Figure 3-43.)

Figure 3-51 `bptestnetconn` for Server_B with no directives listed

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
-----
FL: Server_B -> 10.81.73.147           :      11 ms SRC: ANY
FL: Server_B -> 10.96.73.253          :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11d::1efc  :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11e::1efc  :      11 ms SRC: ANY
FL: Server_B -> 2001:d8b:0:1f0::1efc  :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11c::1efc  :      11 ms SRC: ANY
-----
Total elapsed time: 0 sec
```

Host for which lookup is performed
List of networks available to Server_B
Any source is available to use for a connection

The following directive is added to the **Preferred Networks** properties on Server_B:



In the `bp.conf` file or the registry, the directive appears as follows:

```
PREFERRED_NETWORK = 2001:0db8:0:11c::/62 ONLY
```

This directive provides NetBackup with the information to filter the addresses and choose to communicate with only those that match the `:11c`, `:11d`, `:11e`, and `:11f` networks. The addresses that do not match the **Only** directive are prohibited, as shown in the `bptestnetconn` output.

Figure 3-52 shows the `bptestnetconn` output for Server_B, given this directive.

Figure 3-52 bptestnetconn for Server_B with directive

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
-----
FL: Server_B -> 10.81.73.147           :      11 ms TGT PROHIBITED
FL: Server_B -> 10.96.73.253           :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11d::1efc   :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11e::1efc   :      11 ms SRC: ANY
FL: Server_B -> 2001:d8b:0:1f0::1efc   :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11c::1efc   :      11 ms SRC: ANY
-----
Total elapsed time: 0 sec
```

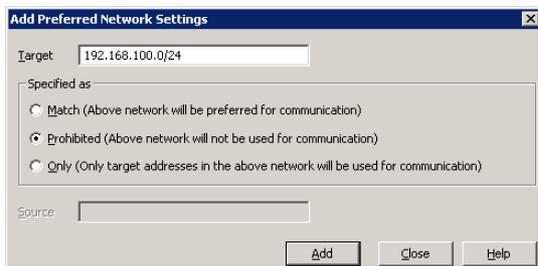
List of networks available to Server_B

Directives make some targets unavailable to Server_B

Configuration to prohibit using a specified address

Figure 3-53 shows a configuration that prohibits NetBackup from using the specified address.

Figure 3-53 Prohibited target example



Configuration that uses the same specification for both the network and the interface—less constrictive

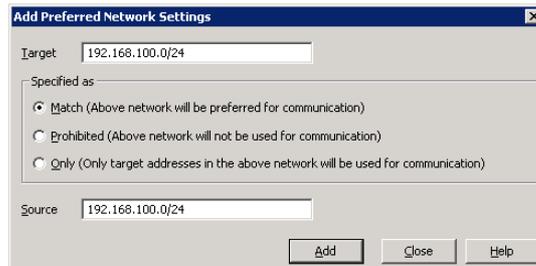
Figure 3-54 shows a configuration that uses the same specification for both the network and the interface.

For all target addresses in the specified network, a source binding in the same network is selected. This directive is considered generic since the identical directive applies to all NetBackup hosts on the network. The closest preferred source address that matches a remote address is used for source binding.

A production network outside this range can then be **Prohibited**, thereby preferring these addresses from both a remote and source binding perspective.

Additional **Match** directives may be used to indicate additional backup networks that are allowed.

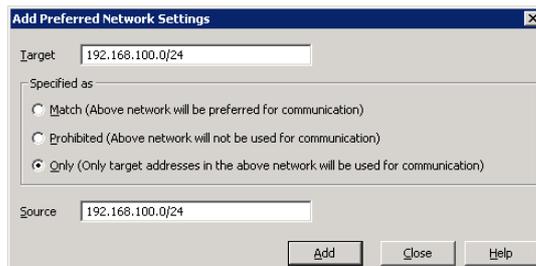
Figure 3-54 Match network selection with the source



Configuration that uses the same specification for both the network and the interface—more constrictive

Figure 3-55 also uses the same specification for both target and source binding, however this example is more restrictive. With the **Only** property specified, this configuration does not allow multiple backup networks to be specified.

Figure 3-55 Only network selection with the same source binding address



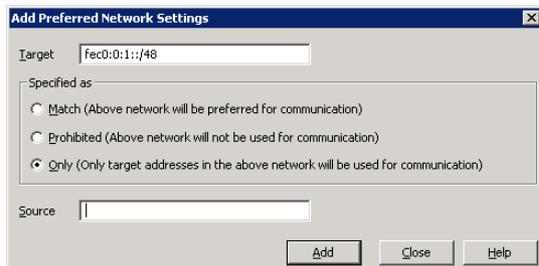
A host with the **Only** directive configured considers only those target addresses in the 192.168.100.0 subnet. Additionally, source binding to the local interface must be done on the 192.168.100.0 subnet.

On hosts that have a 192.168.100.0 interface but no :1b0 interface, source binding to the :1f0 interface is the default of the operating system.

Configuration that limits the addresses, but allows any interfaces

Figure 3-56 shows a configuration that allows only addresses that start with the specified prefix to be considered. No source binding is specified, so any interface may be used.

Figure 3-56 Limiting the addresses, without any source binding

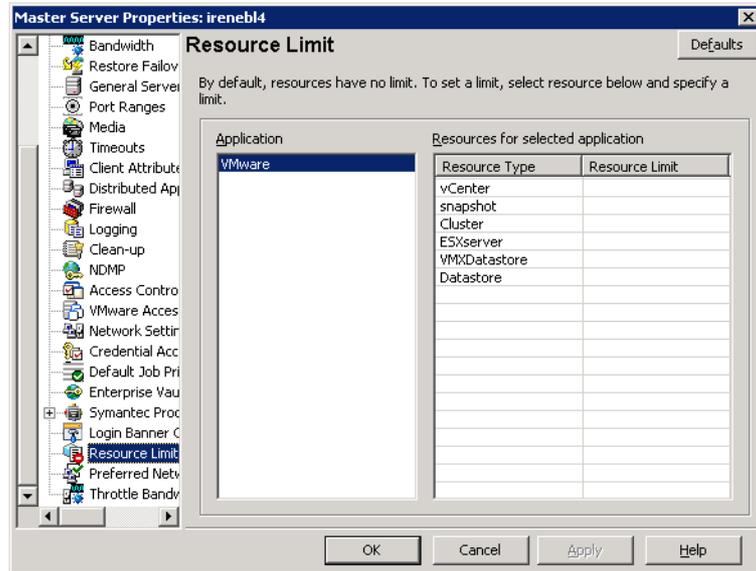


Resource Limit properties

The **Resource Limit** properties in the **NetBackup Administration Console** control the number of simultaneous backups that can be performed on a VMware resource type. These settings apply to all policies for the currently selected master server.

Note: The **Resource Limit** dialog applies only to policies that use automatic selection of virtual machines (the policy's Query Builder). If you select virtual machines manually on the **Browse for Virtual Machines** dialog box, the **Resource Limit** properties have no effect.

Figure 3-57 Resource Limit dialog box



The **Resource Limit** dialog box contains the following properties.

Table 3-51 Resource Limit dialog box properties

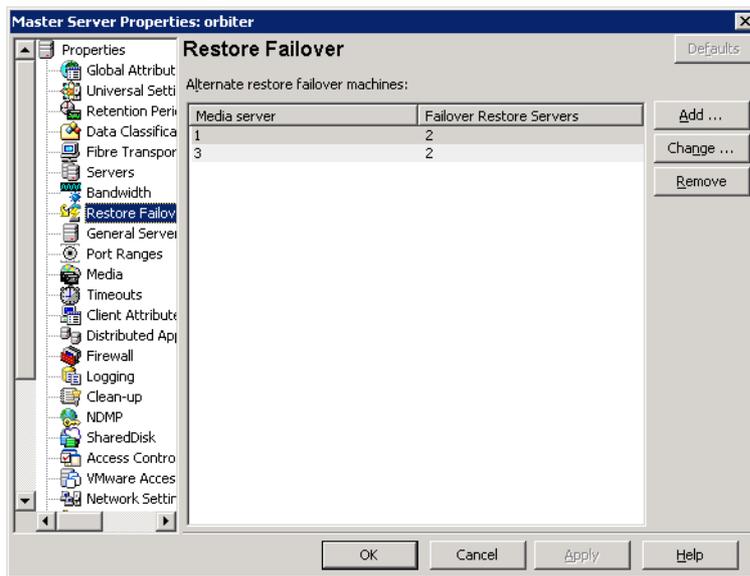
Property	Description
vCenter	The maximum number of simultaneous backups per vCenter server.
snapshot	The maximum number of simultaneous snapshot operations (create or delete) per vCenter.
Cluster	The maximum number of simultaneous backups per VMware cluster.
ESXserver	The maximum number of simultaneous backups per ESX server.
VMXDatastore	The maximum number of simultaneous backups per VMX datastore.
Datastore	The maximum number of simultaneous backups per Datastore.

For example, a **Resource Limit** of two for Datastore means that NetBackup policies can perform no more than two simultaneous backups on any particular datastore.

Restore Failover properties

The **Restore Failover** properties in the **NetBackup Administration Console** 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. These properties apply to currently selected master servers.

Figure 3-58 Restore Failover dialog box



The **Restore Failover** dialog box contains the following properties.

Table 3-52 Restore Failover dialog box properties

Property	Description
Media server	Displays the NetBackup media servers that have failover protection for restores.
Failover restore server	Displays the servers that provide the failover protection. NetBackup searches from top to bottom in the 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.

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, `bptm` cannot find the required tape.)

Assigning an alternate media server as a failover restore server

You can assign another media server to act as a failover restore server for your media server. If your media server is unavailable during a restore, the failover restore server takes its place.

To assign an alternate media server as a failover restore server

- 1 In the **NetBackup Administration Console**, in the left panel, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the right pane, double-click on the master server you want to modify.
- 3 In the properties dialog box, in the left pane, click **Restore Failover**.
- 4 Click **Add**.
- 5 In the **Media server** field, specify the media server for failover protection.
- 6 In the **Failover restore servers** field, specify the media servers to try if the server that is designated in the **Media server** field is unavailable. Separate the names of multiple servers with a single space.
- 7 Click **Add**. The dialog box remains open for another entry.

8 Click **Close**.

9 From the **Restore Failover** dialog box, click **Apply** to accept the changes.

Before the change takes effect, you must stop and restart the NetBackup Request daemon on the master server where the configuration was changed.

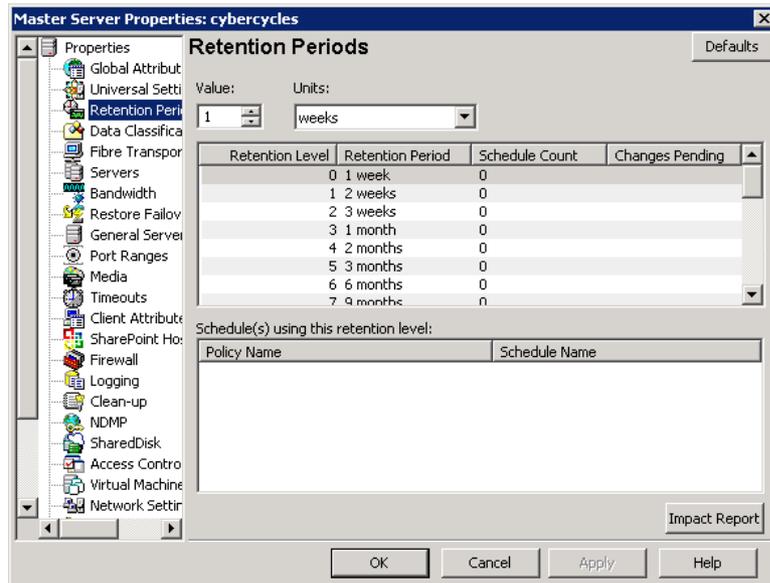
See “About enabling automatic failover to an alternate server” on page 877.

Retention Periods properties

Use the **Retention Periods** properties in the **NetBackup Administration Console** to define a duration for each retention level. You can select from 25 retention levels.

In a policy, the retention period determines how long NetBackup retains the backups or the archives that are created according to the schedule. These properties apply to selected master servers.

Figure 3-59 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, some backups that share the same volume may have different retention periods.

For example, if the retention level 3 is changed from one month to six 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 similar values. However, before a major change is made to a retention period, suspend the volumes that were previously used for that retention level.

See “Determining retention periods for volumes” on page 188.

The **Retention Periods** dialog box contains the following properties.

Table 3-53 Retention Periods dialog box properties

Property	Description
Value	Assigns a number to the retention level setting.
Units	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	<p>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.</p> <p>With the default, there is no difference between a retention level of 12 and a retention level of 20, for example.</p> <p>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.</p> <p>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 accepts or applies the changes.</p>
Schedules list	Lists the schedules that use the currently selected retention level, and the policy to which each schedule belongs.
Impact Report	Displays 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**, in the left pane, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the right pane, double-click on the master server you want to modify.
- 3 In the properties dialog box, in the left pane, click **Retention Periods**.
- 4 Select the retention level to change.

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. Level 9 cannot be changed and remains at a setting of infinite.

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.

- 5 Type the new retention period in the **Value** box.
- 6 From the **Units** drop-down list, select a unit of measure (days, weeks, months, years, Infinite, or Expires immediately).

After you change the value or unit of measure, 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 accepts or applies the changes.

- 7 Click **Impact Report**.

The policy impact list displays the policies and the schedule names where the new retention period is less than the frequency period. To prevent a potential gap in backup coverage, redefine the retention period for the schedules or change the retention or frequency for the schedule.

- 8 Do one of the following:
 - To discard your changes, click **Cancel**.
 - To save your changes and leave the dialog box open to make further changes, click **Apply**.
 - To save your changes and close the dialog box, click **OK**.

Determining retention periods for volumes

Use the following procedure to determine retention periods for volumes.

To determine retention periods for volumes

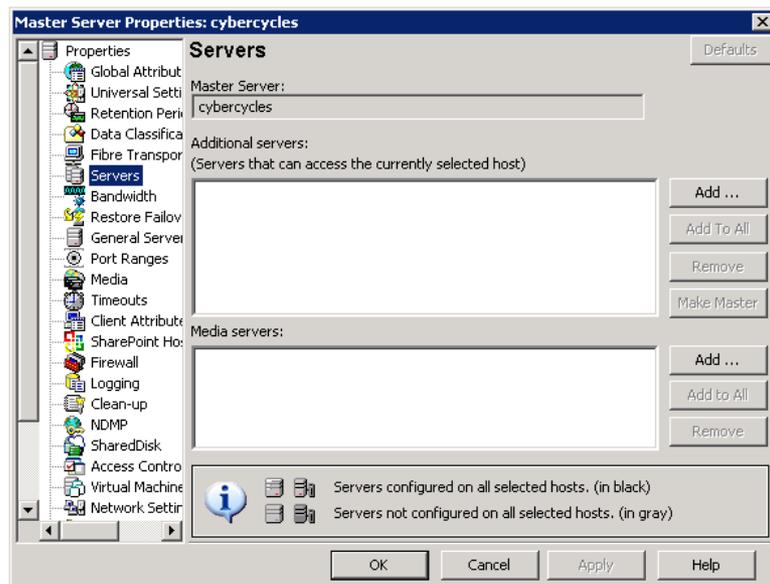
- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**
- 2 In the right pane, find the volume on the list and examine the value in the **Retention Period** column.

To see all volumes that have the same retention period, click the **Retention Period** column header to sort the volumes by retention period. This column heading is hidden by default.

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 3-60 Servers dialog box



The **Servers** dialog box contains the following properties.

Table 3-54 Servers dialog box properties

Property	Description
Master server	Specifies the master server for the selected host. (The name of the selected host appears in the title bar.)
Additional servers list	<p>Lists the additional servers that can access the server that is specified as Master server.</p> <p>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.</p> <p>To configure access to a remote server, add to the server list the name of the host seeking access.</p> <p>See “Accessing remote servers” on page 835.</p>
Media servers list	<p>Lists the hosts that are media servers only. Hosts that are listed as media servers can back up and restore clients, but have limited administrative privileges.</p> <p>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.</p>

If the server appears only in the **Media servers** list, the server is a media server. Media servers can back up and restore clients, but they have limited administrative privileges.

A server that appears in the **Media servers** list and the **Additional servers** list may introduce unintended consequences. A computer that is defined as both a master server and a media server gives the administrator of the media server full master server privileges. By listing the media server in both places, you may inadvertently give the media server administrator more privileges than intended.

See “Adding a server to a remote server list” on page 837.

Adding a server to the Additional servers list

You can add a master server, media server, or client to the **Additional servers** list in the **Servers** properties dialog box.

To add a server to the Additional servers list

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**.
- 2 Depending on the host to be configured, select **Master Servers**, **Media Servers**, or **Clients**.
- 3 In the right pane, double-click the master server, media server, or client you want to modify.
- 4 In the properties dialog box, in the left pane, click **Servers**.
- 5 From the **Additional servers** list, click **Add**.
To add multiple hosts, select more than one media server or client in step 2 and click **Apply To All** in step 5. However, you can add only one master server to the list at a time.
- 6 In the **Add a New Server Entry** dialog box, type the name of the new server.
- 7 Click **Add**. The dialog box remains open for another entry.
- 8 Click **Close**.

Note: If you add a media server, run `nbeemmcmd -addhost` to add the media server to the Enterprise Media Manager (EMM) database of the existing master server.

See “About sharing one Enterprise Media Manager (EMM) database across multiple master servers” on page 193.

Adding a server to the Media servers list

You can add a master server, media server, or client to the **Media servers** list in the **Servers** properties dialog box.

To add a server to the Media servers list

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**.
- 2 Depending on the host to be configured, select **Master Servers**, **Media Servers**, or **Clients**.
- 3 In the right pane, double-click the master server, media server, or client you want to modify.
- 4 In the properties dialog box, in the left pane, click **Servers**.

- 5 From the **Media servers** list, click **Add**.
To add multiple hosts, select more than one media server or client in step 3 and click **Apply To All** in step 5. However, you can add only one master server to the list at a time.
- 6 In the **Add a New Server Entry** dialog box, type the name of the new server.
- 7 Click **Add**. The dialog box remains open for another entry.
- 8 Click **Close**.

Note: If you add a media server, run `nbemmcmd -addhost` to add the media server to the Enterprise Media Manager (EMM) database of the existing master server.

See “About sharing one Enterprise Media Manager (EMM) database across multiple master servers” on page 193.

Removing a server from the Additional servers list or the Media servers list

You can remove a master server or a media server from the **Additional servers** list. You can also remove a media server from the **Media servers** list.

To change the Master Server

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**.
- 2 Depending on the host to be configured, select **Master Servers**, **Media Servers**, or **Clients**.
- 3 In the right pane, double-click the master server, media server, or client you want to modify.
- 4 In the properties dialog box, in the left pane, click **Servers**.
- 5 From the **Additional servers** list or the **Media servers** list, select a server.
- 6 Click **Remove**.

Switching to another master server in the Servers properties dialog box

You can switch to view the properties of another master server in the Servers properties dialog box.

To switch the master server in the Servers properties dialog box

- 1** In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**.
- 2** Depending on the host to be configured, select **Master Servers**, **Media Servers**, or **Clients**.
- 3** In the details pane, double-click the master server, media server, or client you want to modify.
- 4** In the dialog box, click **Servers**.
- 5** From the **Additional servers** list, select a server.
- 6** Click **Make Master**.

About sharing one Enterprise Media Manager (EMM) database across multiple master servers

Multiple master servers can share one Enterprise Media Manager (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 host that contains the EMM database.

Use the `bpgetconfig` command to obtain a list of configuration entries. Then, use the `bpsetconfig` command to change the entries as needed. For information about these commands, see *NetBackup Commands Reference Guide*.

The following table shows example registry entries from three master servers (*Meadow*, *Havarti*, and *Study*) that share one EMM database. One of the servers (*Meadow*) hosts the EMM database.

Table 3-55 Example registry entries from three master servers that share an EMM database

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
EMMSERVER = meadow		

Use the following conventions when making entries like those in the example:

- The first `SERVER` entry must be the name of the master server. The table shows that the first `SERVER` entry matches the name of each master server.
- The host server must have a `SERVER` entry for each server that shares the EMM database. This entry allows the **NetBackup Administration Console** to administer the other servers. The table shows *Havarti* and *Study* listed under *Meadow*.
See “About choosing a remote server to administer” on page 839.
- If the EMM database is hosted on another master server, that server must be listed. The table shows *Meadow* listed under *Havarti* and *Study*.
- The `CLIENT_NAME` entry must match the name of the master server.
- The `EMMSERVER` entry must be present on all master servers that share the EMM host. The table shows *Meadow* listed as the `EMMSERVER` for all three servers.

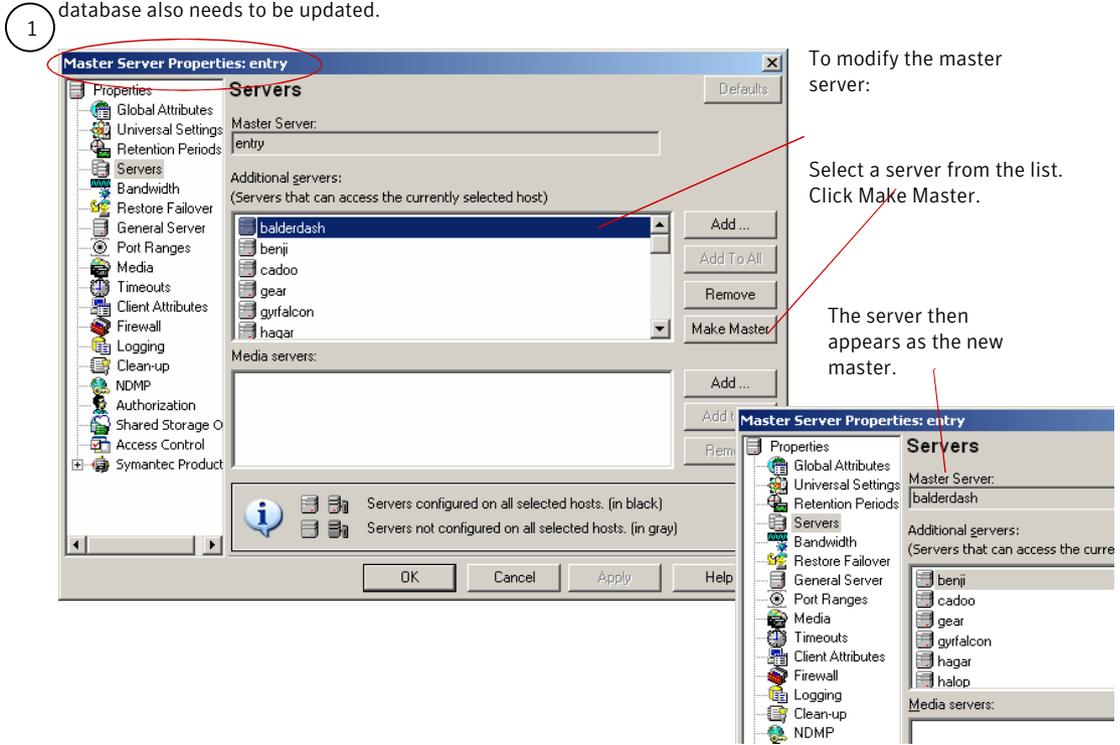
If you assign the media server to a different master, the Enterprise Media Manager database also needs to be updated. To update the EMM database, run the following command:

See “Switching to another master server in the Servers properties dialog box” on page 192.

```
/usr/opensv/netbackup/bin/admincmd/nbemcmd -updatehost
```

Figure 3-61 A shared EMM database that is located on a media server

If the master server is changed on a media server, the EMM database also needs to be updated.



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

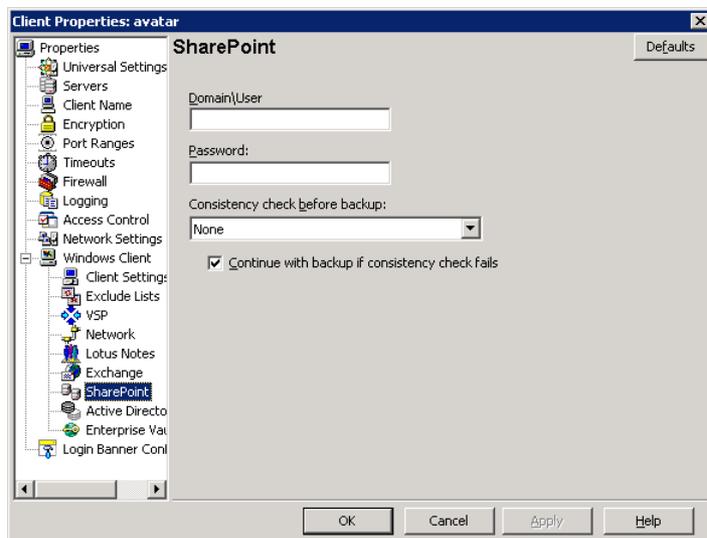
The **SharedDisk** properties specify the SharedDisk storage option properties for a NetBackup configuration. These properties apply to currently selected master servers.

See “About SharedDisk support in NetBackup 7.0 and later” on page 382.

SharePoint properties

The **SharePoint** properties apply to currently selected Windows clients to protect SharePoint Server installations.

Figure 3-62 SharePoint dialog box



The **SharePoint** dialog box contains the following properties.

Table 3-56 SharePoint dialog box properties

Property	Description
Domain\User	Specifies the domain and the user name for the account you want to use to log on to SharePoint (DOMAIN\user name).
Password	Specifies the password for the account.
Consistency check before backup	Specifies 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. If you choose to perform a consistency check, you can select Continue with backup if consistency check fails . NetBackup then continues to perform the backup if the consistency check fails.

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

Consistency check options for SharePoint Server

The following consistency checks can be performed before a SharePoint Server backup.

Table 3-57 Consistency check options

Option	Description
None	Do not perform consistency checking.
Full check, excluding indexes	Select this option to exclude indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the non-clustered index pages is not checked.
Full check, including indexes	Include indexes in the consistency check. Any errors are logged.
Physical check only (SQL 2000 only)	Select this option to perform a low overhead check of the physical consistency of the SQL Server 2000 database. This option only checks the integrity of the physical structure of the page headers and record headers. It also checks the consistency between the pages' object ID and index ID and the allocation structures.

Symantec Products properties

The **Symantec Products** properties encompass properties for other Symantec products. This property includes properties for the Backup Exec Tape Reader.

The Symantec Products properties include the subnode, Backup Exec Tape Reader properties.

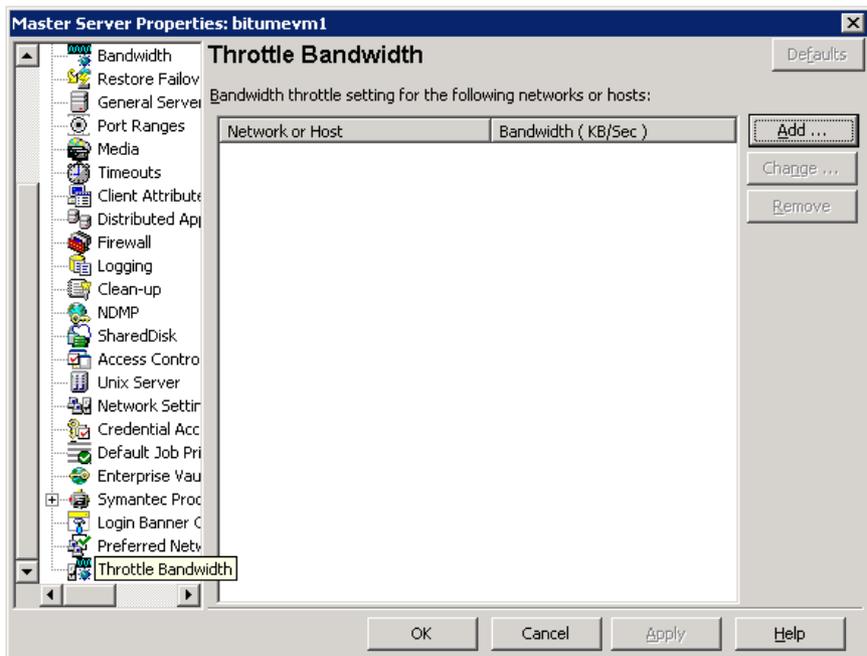
See “Backup Exec Tape Reader properties” on page 68.

Throttle Bandwidth properties

Use the **Throttle Bandwidth** properties to specify a limit for the network bandwidth or transfer rate that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. These properties limit only backups. Restores are unaffected. The default is that the bandwidth is not limited.

The **Throttle Bandwidth** properties are similar to the **Bandwidth** host properties, but offer greater flexibility in IPv6 environments.

Figure 3-63 Throttle Bandwidth dialog box



To manage entries in the **Throttle Bandwidth** dialog box, select one of the following buttons:

- Add** Add a network or host to the **Network or Host** list using the **Add Bandwidth Settings** dialog box.
- Change** Change the selected network or host property using the **Change Bandwidth Settings** dialog box.
- Remove** Removes the selected network or host from the **Network or Host** list.

Add Bandwidth Settings dialog box for Throttle Bandwidth properties

The **Add Bandwidth Settings** and the **Change Bandwidth Settings** dialog boxes contain the following properties.

- Network or Host** The network or host to which the throttle applies.

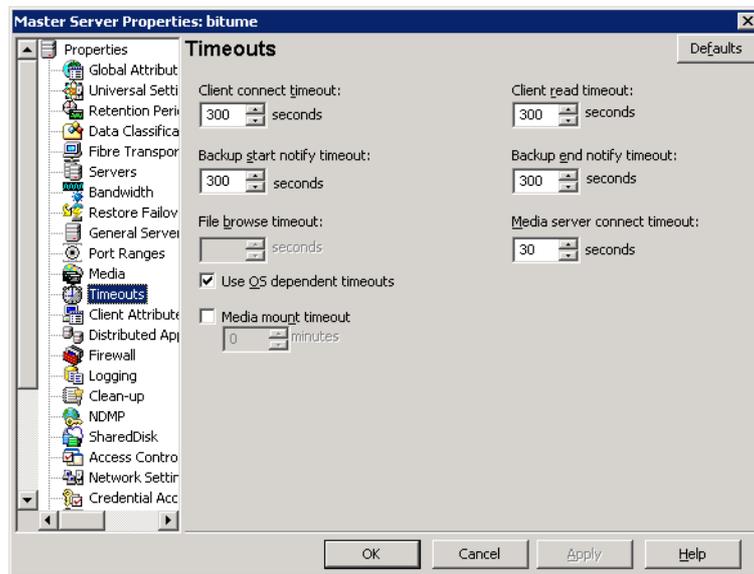
Bandwidth (KB/Sec)

The bandwidth or the transfer rate in kilobyte per second for the network or host indicated. A value of zero disables throttling IPv6 addresses.

Timeouts properties

The **Timeouts** properties apply to selected master servers, media servers, and clients.

Figure 3-64 Timeouts dialog box



The **Timeouts** dialog box contains the following properties.

Table 3-58 Timeouts dialog box properties

Property	Description
Client connect timeout	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	Specifies the number of seconds the server waits for the <code>bpstart_notify</code> 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.

Table 3-58 Timeouts dialog box properties (*continued*)

Property	Description
File browse timeout	<p>Specifies how long the client can wait for a response from the NetBackup master server while it lists files.</p> <p>Note: If it exists, the value in a UNIX client's <code>\$HOME/bp.conf</code> file takes precedence to the property here.</p> <p>If the limit is exceeded, the user receives a socket read failed error. The timeout can be exceeded even while the server processes the request.</p>
Use OS dependent timeouts	<p>Specifies that the client waits for the timeout period as determined by the operating system when it lists files, as follows:</p> <ul style="list-style-type: none"> ■ Windows client: 300 seconds ■ UNIX client: 1800 seconds
Media mount timeout	<p>Specifies how long NetBackup waits for the requested media to be mounted, positioned, and ready on backups, restores, and duplications.</p> <p>This property applies to currently selected master servers.</p> <p>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.)</p>
Client read timeout	<p>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.</p> <p>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.</p> <p>Note: For database-extension clients, Symantec suggests that the Client read timeout be set to a value greater than 5 minutes. 15 minutes are adequate for many installations. For other clients, change this property only if the client encounters problems.</p> <p>The sequence on a database-extension client is as follows:</p> <ul style="list-style-type: none"> ■ 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. <p>See "Client Settings (UNIX) properties" on page 92.</p>

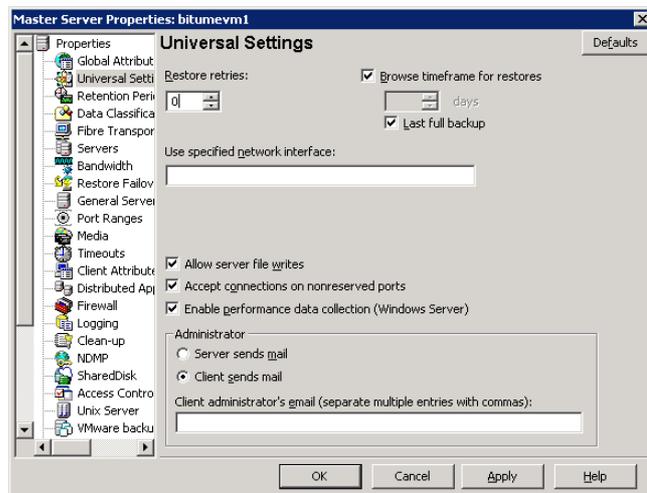
Table 3-58 Timeouts dialog box properties (*continued*)

Property	Description
Backup end notify timeout	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	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

Use the **Universal Settings** properties in the NetBackup Administration Console to configure basic backup and restore settings. These properties apply to selected master servers, media servers, and clients.

Figure 3-65 Universal Settings dialog box



The **Universal Settings** dialog box contains the following options.

Table 3-59 Universal Settings dialog box properties

Property	Description
Restore retries	<p>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.</p> <p>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.</p> <p>See “Clean-up properties” on page 75.</p> <p>A checkpointed job is retried from the start of the last checkpointed file rather than at the beginning of the job.</p> <p>Checkpoint restart for restore jobs allows a NetBackup administrator to resume a failed restore job from the Activity Monitor.</p> <p>See “Take checkpoints every __ minutes (policy attribute)” on page 521.</p>
Browse timeframe for restores	<p>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.</p> <p>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.</p> <p>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.</p>
Last full backup	<p>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.</p>
Allow server file writes	<p>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.</p> <p>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.</p>
Accept connections on nonreserved ports	<p>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.</p> <p>This property no longer applies. For information about this property, refer to NetBackup 6.5 documentation.</p>

Table 3-59 Universal Settings dialog box properties (continued)

Property	Description
Enable performance data collection (Windows server only)	Specifies whether NetBackup updates disk and tape performance object counters. (Applies only to Windows master and media servers. Use the Windows Performance Monitor utility (<code>perfmon</code>) to view the NetBackup performance counters. The default is that this property is enabled. See the <i>NetBackup Administration Guide, Volume II</i> for more information about using the System Monitor with NetBackup.
Client sends mail	Specifies whether the client sends an email to the address that is specified in the Universal Settings properties. If the client cannot send email, use Server sends mail . The default is that this property is enabled.
Server sends mail	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	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.

Logging the status of a redirected restore

A redirected restore may not produce a progress log. The name of the requesting server must appear in the server list for the server that performs the restore. Otherwise, no progress log is produced for a redirected restore. (A progress log is an entry in the **Task Progress** tab of the Backup, Archive, and Restore client interface.)

Without the entry in the server list, the restoring server has no access to write the log files to the requesting server. Add the requesting server to the server list and log into the requesting server

To produce a progress log

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the right pane, double-click the master server you want to modify.
In the properties dialog box, in the left pane, click **Servers**.

3 Perform one of the following actions:

To add the restoring server to the **Additional servers** list From the **Media servers** list, click **Add**.

To add the restoring server to the **Media servers** list From the **Additional servers** list, click **Add**.

4 In the **Add a New Server Entry** dialog box, type the name of the new server.

5 Click **Add**. The dialog box remains open for another entry.

6 Click **Close**.

7 Log on to the restoring server.

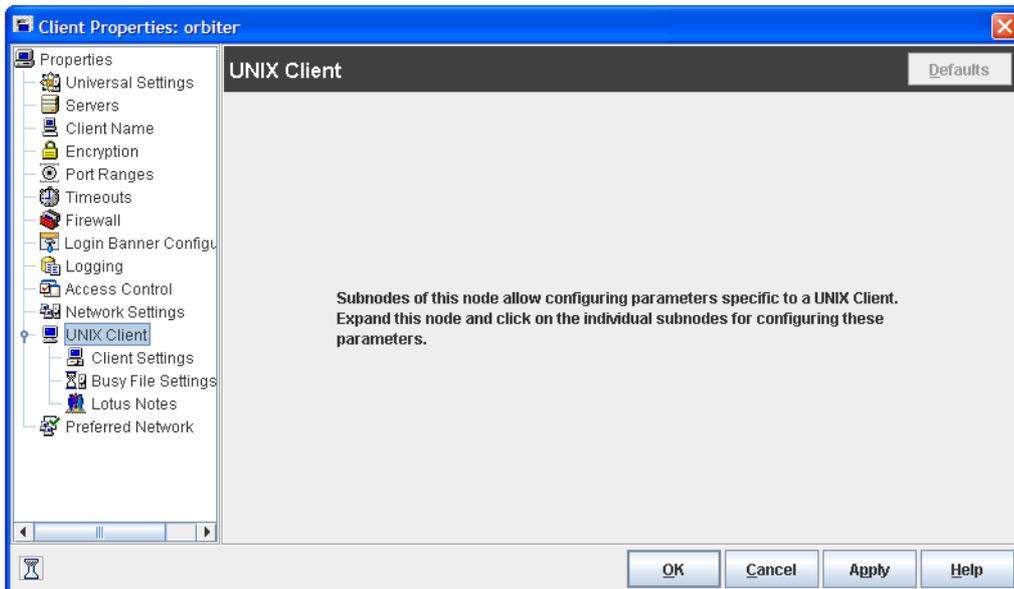
Check the **Activity Monitor** to determine the success of the restore operation.

See “About the Jobs tab” on page 766.

UNIX Client properties

Use the **UNIX Client** properties in the **NetBackup Administration Console** to define properties of UNIX clients.

Figure 3-66 UNIX Client dialog box



See “Client Settings (UNIX) properties” on page 92.

See “Busy File Settings properties” on page 72.

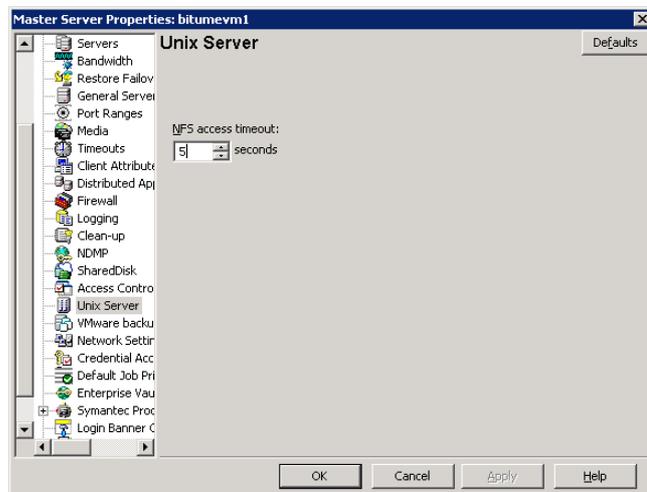
See “Lotus Notes properties” on page 150.

UNIX Server properties

Use the **UNIX Server** properties in the **NetBackup Administration Console** to change the **NFS access timeout** property. This property specifies how long the backup waits to process the mount table before it considers an NFS file system unavailable. The default is 5 seconds.

These properties apply to selected UNIX master servers.

Figure 3-67 UNIX Server dialog box

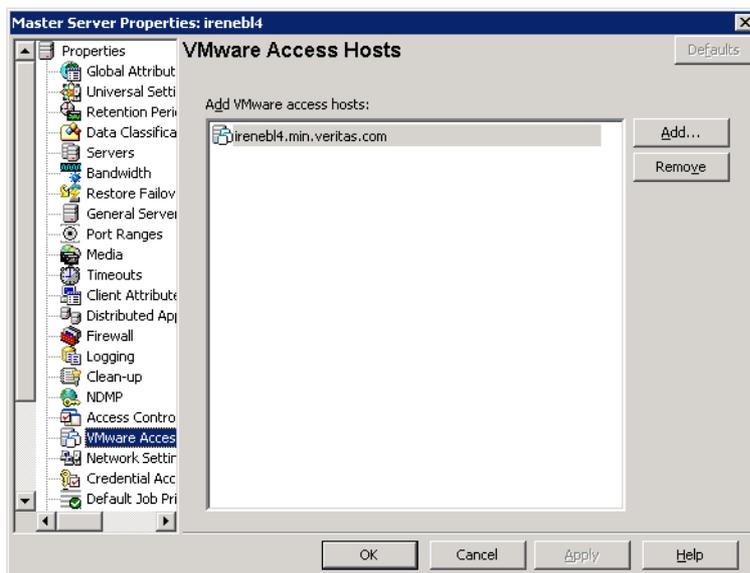


See “Follow NFS (policy attribute)” on page 528.

VMware Access Hosts properties

Use the **VMware backup hosts** properties in the **NetBackup Administration Console** to add or remove VMware backup hosts. These properties appear when the NetBackup Enterprise Client license is installed. These properties apply to currently selected master servers.

Figure 3-68 VMware Access Hosts dialog box



You can add servers to and remove servers from the backup hosts list.

A VMware backup host is a server on the same SAN as a VMware ESX server. The VMware ESX server must be able to access the snapshot of the VMware 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

Use the **Volume Snapshot Provider** properties in the **NetBackup Administration Console** to change the way NetBackup manages snapshots. These properties are displayed when the selected client is running NetBackup 6.x. The VSP properties do not appear for 7.x clients.

See the following topic for information about selecting VSP for backlevel and upgraded clients:

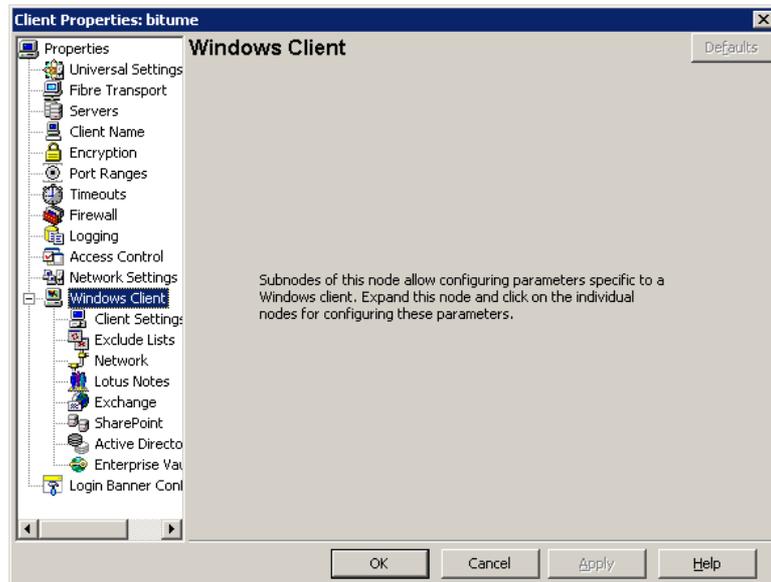
See “Backlevel and upgraded clients that use Windows Open File Backup” on page 90.

For information about VSP settings, see the *6.5 NetBackup Administrator's Guide, Volume I*.

Windows Client properties

Use the **Windows Client** properties in the **NetBackup Administration Console** to define NetBackup properties for Microsoft Windows clients.

Figure 3-69 Windows Client dialog box



Windows Client properties include specific host properties for configuring Windows clients.

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 “About media sharing” on page 317.

See “Configuring media sharing with a server group” on page 318.

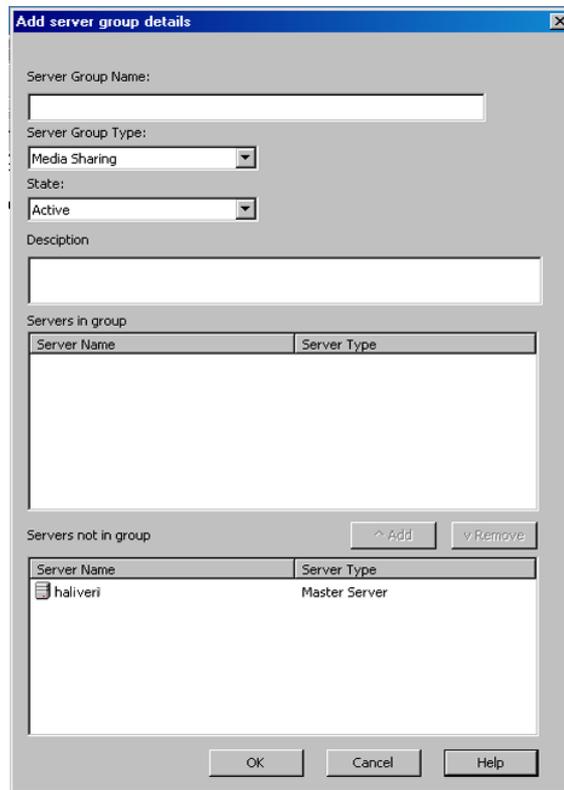
Configuring a server group

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.

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 > New Server Groups**.



- 3 In the **Add Server Group Details** dialog box, enter or select the appropriate information.

See “Server group properties” on page 212.

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 following table describes server group properties that include the following options.

Table 4-1 Server group properties

Property	Description
Server group name	<p>Specifies 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>Specifies the type of server group.</p> <p>See “About server groups” on page 209.</p> <p>Other server group types (such as Alternate Restore) are reserved for future use.</p>
State	<p>Specifies 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	<p>Describes the media server group.</p>
Servers in group	<p>Specifies the servers (and the server type) that belong to the group.</p>
Servers not in group	<p>Specifies the servers (and the server type) that do not belong to the group.</p>

Deleting a server group

Use the following procedure to delete a server group.

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
- About decommissioning a media server
- Previewing references to 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

The following table describes an overview of how to 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 189.
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 782.

Table 6-1 Adding a media server (continued)

Procedure	Section
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</i> .
On the master server, configure the robots and drives that are attached to the media server.	See “Configuring robots and tape drives” on page 241.
On the master server, configure the volumes.	See “About adding volumes” on page 283.
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 388.
On the master server, configure the NetBackup policies and schedules to use the storage units that are configured on the media server.	See “About the Policies utility” on page 502.
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 636.

About decommissioning a media server

New with this release is a command to decommission a media server, `nbdecommission`. The command launches a text-based wizard that guides you through the decommission process. The wizard removes the references to a media server from a NetBackup domain. (You may have to remove some references manually; the wizard provides instructions to do so in most cases.)

The `nbdecommission` command helps in the following scenarios:

- You add a new media server and new storage to your environment. You direct all backup jobs that went to the old server to the new server. After all of the backup images on the old server expire, you run `nbdecommission` to retire the old server.

- You replace an old server with a new server and keep the same storage. You want to access all of the old server storage and backup images from the new server.
- The old server fails, and you need to replace it with a new server.

You also can use the wizard if you try to decommission a media server manually and references to it still remain. The wizard may clean up any references that remain.

Throughout this documentation, the media server to be decommissioned is referred to as the old server.

Warning: Be careful when you use the `nbdecommission` command. Because the command may expire images, data loss may occur. Therefore, you should completely understand what the command does before you use it. Symantec recommends that you first preview all of the references to a media server before you decommission it.

See “Previewing references to a media server” on page 224.

About decommissioning limitations

The following are the limitations of the `nbdecommission` command:

- Does not decommission media servers at release levels earlier than 6.0.
- Does not decommission clustered media servers. Those include NetBackup failover media servers or application clusters.
- Does not process the media server deduplication pools.
- Does not update the `vm.conf` files on the NetBackup servers in your environment. Therefore, the old server may remain in the `vm.conf` files on the NetBackup servers.
- Does not update the configuration files on the clients. Therefore, the old server may remain in the server lists on the clients. If you replace an old server with a new server, the new server is not added to the client server lists.
- Does not process the NetBackup Vault profiles. If NetBackup Vault profiles exist that refer to the storage units on the old server, update the Vault profiles manually.
- Does not notify you about orphaned resources.
- Does not restart the daemons and services on other servers that the decommissioning affects.

- Requires that you shut down all daemons and services on the old server after it is decommissioned.
- Requires that you reconfigure devices on the new server manually (if required).
- Requires that you know which jobs are running on the old server. You must kill them or let them run to completion before you run the decommission process.
- The `-list_ref` option only reports on the references that it removes explicitly. The command removes some items implicitly and it does not report them. For example, host aliases and host credentials are removed but not reported.
- Requires that you move any media ID generation rules that exist on the old server. You must move them manually to the media server that performs robot inventory.
- Moves the old server to an Administrative Pause state so that no new jobs are started. However, NetBackup still can start backup and restore jobs for basic disk; they obtain resources differently than do jobs for other storage destinations. Also, the `nbdecommission` command may clear the Administrative Pause to expire images (depending on your responses to the wizard). Jobs may start during this period.

Before you decommission a media server

Before you decommission a media server, Symantec recommends that you do the following:

- Preview the actions of the `nbdecommission` command.
 See “Previewing references to a media server” on page 224.
 Analyze the output of the preview operation to ensure that the command captures all references to the old server. If it did not, make a list of the items that the command does not cover and fix them manually later.
- Back up the NetBackup catalog before you begin. You can use it to return your environment to the pre-decommission state if something goes wrong or you have to abort the decommission.
- Run the command during a maintenance window when the load on the NetBackup environment is minimal.

Post decommission recommendations

After you run the `nbdecommission` command, the following actions are recommended:

- Follow all of the instructions the command provides.

The command may provide instructions for performing the actions that it cannot perform. For example, it may provide instructions to cancel the backup jobs that are active on the old server.

- Move the physical storage (if needed) and then reconfigure and reinventory those devices.
- Examine the `vm.conf` files on all of the NetBackup servers in your environment. Remove references to the old server and add references to the new server where necessary.
- Remove the old server from the server lists on the clients and add the new server where necessary.
 The `nbdecommission` command outputs a list of clients that refer to old server.
- Verify that the old server was removed correctly. Examine the various logical components (backup policies, storage units, and so on) to make sure that the old server references have been removed.
- Back up the NetBackup catalog as soon as possible.

Decommission actions

The `nbdecommission` command deletes the configuration for the old server from the EMM database, the NetBackup image catalog, and configuration files on servers.

The following table shows the actions it performs for the components that reference the media server. The table is organized in the order in which the command processes the component.

Table 6-2 `nbdecommission` command actions

Component	Action
Storage unit - Tape	Deletes the following tape storage units: <ul style="list-style-type: none"> ■ Those in which the Storage device attribute specifies a robot for which the old server is the robot control host. ■ Those in which the Media server attribute specifies the old server. ■ Those in which the Media server attribute specifies Any Available and the old server is the only server that can access the storage unit.

Table 6-2 `nbdecommission` Command actions (continued)

Component	Action
Tape drive	<p>Deletes the tape drive path for each tape drive that is attached to the old server. If the path on the old server is the only path, it also deletes the tape drive.</p> <p>If a path to a drive exists on more than one media server, the tape drive may become unusable. You may have to connect the tape drive to a different media server and then reconfigure it in NetBackup. For example, if the old server is a scan host for a shared drive, NetBackup cannot use the drive if no other host can scan.</p>
Robotic library	<p>Deletes all of the robotic libraries that are attached to the old server.</p> <p>If the old server is the robot control host for a shared library, the drives and media become stand-alone and unusable. You must reconfigure and re-inventory the library.</p>
Tape media	<p>Specifies if you want to expire the following tape media or move them to another media server:</p> <ul style="list-style-type: none"> ■ Those assigned to the old server. ■ Those owned by a media sharing group in which the old server is the only member of the group. ■ Those that have no specific Media owner and the last write host is same as the old server.
Storage unit - BasicDisk	<p>Deletes the storage unit if no images exist on it. If images exist, the wizard lets you choose one of the following options:</p> <ul style="list-style-type: none"> ■ Expire the images and delete the storage unit. ■ Move the images to the new server. The wizard also updates the Media server field in the storage unit. The BasicDisk storage must be shared, and the same disk path must be available on the new server.
Storage unit - Nearstore	<p>Deletes the storage unit if no images exist on it. If images exist, the wizard lets you choose one of the following options:</p> <ul style="list-style-type: none"> ■ Expire the images and delete the storage unit. ■ Move the images to a new server. The wizard also transfers the credentials to the new server and updates the Media server field in the storage unit.

Table 6-2 `nbdecommission` Command actions (continued)

Component	Action
Storage unit - SnapVault	<p>Deletes the storage unit if no images exist on it. If images exist, the wizard lets you choose one of the following options:</p> <ul style="list-style-type: none"> ■ Expire the images and delete the storage unit. ■ Move the images to a new server. The wizard also transfers the credentials to the new server and updates the Media server field in the storage unit.
Storage unit - AdvancedDisk and SharedDisk	<p>Specifies that if more than one media server can access the disk pool that is the destination of the storage unit, it does the following:</p> <ul style="list-style-type: none"> ■ Removes the old server from the Media Servers list of the storage unit. ■ Deletes the old server as a storage server. <p>If the old server is the only server that can access the disk pool, the wizard lets you choose to do one of the following:</p> <ul style="list-style-type: none"> ■ Move the storage and images to the new server and delete the old server as a storage server. The disk volumes must be available on the new server at the same path as the old server. ■ Expire the images (if any), delete any storage units that reference the disk pool, delete the disk pool, and delete the storage server. (A reference is when the disk pool appears in the Disk pool setting of a storage unit.)
Storage unit - OpenStorage	<p>Specifies that if more than one media server can access the disk pool that is the destination of the storage unit, it does the following:</p> <ul style="list-style-type: none"> ■ Removes the old server from the Media Servers list of the storage unit. ■ Deletes the media server as an OpenStorage storage server. <p>If the old server is the only server that can access the disk pool, the wizard lets you choose to do one of the following:</p> <ul style="list-style-type: none"> ■ Transfer the credentials to the new server and update the Media server field in the storage unit if required. ■ Expire the images (if any), delete any storage units that reference the disk pool, and delete the disk pool. (A reference is when the disk pool appears in the Disk pool setting of a storage unit.)

Table 6-2 `nbdecommission` command actions (continued)

Component	Action
Storage unit group	<p>Specifies that if the <code>nbdecommission</code> command deletes all of the storage units in a storage unit group, it also deletes the storage unit group. Deleting the storage unit group also may affect backup policies and storage lifecycle policies.</p> <p>See “Backup policy and schedule” and “Storage lifecycle policy” in this table.</p>
Backup policy and schedule	<p>Deactivates any backup policy in which the storage destination (directly or indirectly) is a storage unit that the command deletes. Specifically, deactivates any backup policy that meets any of the following conditions:</p> <ul style="list-style-type: none"> ■ The destination is a storage unit that the <code>nbdecommission</code> command deleted. ■ The destination is a storage unit group that contains only one storage unit and the <code>nbdecommission</code> command deleted that storage unit. ■ The destination is a storage lifecycle policy and the <code>nbdecommission</code> command deleted the storage unit that is a Backup destination of the storage lifecycle policy.
Storage lifecycle policy	<p>Specifies that for each storage lifecycle policy in which one or more destinations is a storage unit the command deleted, it does the following:</p> <ul style="list-style-type: none"> ■ If images under the SLP control are in-process or yet to be processed, displays the commands to cancel the SLP jobs and then exits. After you cancel the jobs (or wait until the jobs complete), rerun the <code>nbdecommission</code> command to continue with the decommissioning. ■ If all of the images under SLP control are processed, deactivates the storage lifecycle policy. ■ If a deleted storage unit is a Backup or Snapshot destination, deactivates all backup policies with the storage lifecycle policy as the destination.
Fibre Transport media server	<p>Displays the commands necessary to delete the old server as an FT media server and then exits.</p> <p>After you delete the old server as an FT media server, rerun the <code>nbdecommission</code> command to continue with the decommissioning.</p>

Table 6-2 `nbdecommission` command actions (continued)

Component	Action
<code>bp.conf</code> file	<p>On UNIX NetBackup servers, removes the old server from the following <code>bp.conf</code> file entries:</p> <ul style="list-style-type: none"> ■ <code>SERVER</code> ■ <code>MEDIA_SERVER</code> ■ <code>CLIENT_NAME</code> ■ <code>BROWSER</code> <p>On UNIX master servers, also removes the old server from the <code>FORCE_RESTORE_MEDIA_SERVER</code> and <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> entries.</p>
Windows registry	<p>On Windows NetBackup servers, removes the old server from the following registry keys:</p> <ul style="list-style-type: none"> ■ <code>SERVER</code> ■ <code>MEDIA_SERVER</code> ■ <code>CLIENT_NAME</code> ■ <code>BROWSER</code> <p>On Windows master servers, also removes the old server from the <code>FORCE_RESTORE_MEDIA_SERVER</code> and <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> keys.</p>
Clients	<p>Lists the clients on which the old server appears in their server lists. You must remove the references to the old server manually.</p>

Previewing references to a media server

Use the following procedure to preview the associations and references to a media server that you want to decommission. Symantec recommends that you preview the references to a media server before you decommission it.

The old server does not have to be up and responsive.

See “About decommissioning a media server” on page 217.

See “Decommissioning a media server” on page 225.

The `nbdecommission` command resides in the following directories:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\Veritas\NetBackup\bin\admincmd`

To preview references to a media server

- 1 Run the `nbdecommission` command on the master server or on a media server. The following is the command syntax:

```
nbdecommission -list_ref -oldserver OldServer > file.txt
```

Replace *OldServer* with the name of the host to be decommissioned. Replace *file* with a name that denotes its contents or purpose.

- 2 Analyze the output of the preview operation to ensure that the command captures all references to the old server. If it did not, make a list of the items that the command does not cover and fix them manually later.

Decommissioning a media server

Use the `nbdecommission` text-based wizard to decommission a media server. The wizard guides you through the decommission process. Your path through the wizard depends on how you respond to the wizard prompts. Depending on your environment and how you respond to prompts, the wizard may advise you to perform an action and then exit. To continue in the wizard, you must run the wizard again after you perform the advised action. You may have to exit and rerun the wizard several times.

If active jobs exist on the media server, you must cancel them before the command can begin to decommission the media server. Alternatively, you can wait until they finish.

The *OldServer* does not have to be up and responsive.

Symantec recommends that you preview the media server references before you decommission a media server.

See “About decommissioning a media server” on page 217.

See “Previewing references to a media server” on page 224.

The `nbdecommission` command resides in the following directories:

- UNIX: `/usr/opensv/netbackup/bin/admincmd`
- Windows: `install_path\Veritas\NetBackup\bin\admincmd`

The `nbdecommission` command logs to the standard NetBackup administrator commands log directory.

To replace an old media server with a new media server

- 1 Run the `nbdecommission` command on the master server or on a media server that is not the object of this operation. The following is the command syntax:

```
nbdecommission -oldserver OldServer [-newserver NewServer] [-file  
decom_ops.txt]
```

Replace *OldServer* with the name of the host to be decommissioned.

`-newserver` is optional. If you specify a new server, the new server becomes the default media server for the replacement operations. If you do not specify a new server, the wizard prompts you for the new server for each storage type that contains valid backup images. This method is useful if you want to move backup images to different media servers. For example, you can move backup images from tape storage to one media server and backup images from disk storage to another media server.

`-file` is optional. It writes the command operations to the specified file. Replace *decom_ops.txt* with a name that denotes its purpose or contents. Symantec recommends that you use the `-file` option to maintain a record of the command operations.

- 2 Follow the prompts and perform the requested actions.

For example, the command may make changes on the master server and on multiple media servers. You may be required to restart the NetBackup services on those servers so that the changes take effect.

To decommission a media server

- 1 Run the following command on the master server or on a media server that is not the object of this operation. The *OldServer* does not have to be up and responsive.

```
nbdecommission -oldserver OldServer
```

Replace *OldServer* with the name of the host to be decommissioned.

- 2 Follow the prompts and perform the requested actions.

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):

```
nbemmcmd -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 `nbemmcmd` command usage is available.

See the *NetBackup Commands Reference 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\nbemmcmd -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\bpimmedia -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**, in the left pane, 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 tape 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
- Chapter 15. Duplicating images to a remote master server domain

Configuring robots and tape 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 files
- Downloading the device mapping files
- About configuring robots and tape drives
- About device discovery
- About robot control
- 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
- About automatic path correction
- Enabling automatic path correction

- 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.

The following table 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 determined 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	32000	SCSI control.
TLH	Tape library Half-inch	256	No limit	API control.
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 189. 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 `nbeemmcmd -addhost` command to add the host. See the *NetBackup Commands Reference Guide*.

About the device mapping files

NetBackup uses several files to determine which protocols and settings to use to communicate with storage devices. NetBackup also uses the files during device discovery and configuration.

The device mapping files are available for download from the Symantec support site. The download packages contain the following files:

- external_robotics.txt
- external_types.txt
- Readme.txt

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 files package from the Symantec support Web site and configure NetBackup to use that file. For instructions, see the `Readme.txt` file that is supplied with the device mapping file package.

Note: The contents of the device mapping files do not indicate support for any of the devices, only the ability to recognize and automatically configure them.

See “Downloading the device mapping files” on page 236.

See “About device discovery” on page 237.

Downloading the device mapping files

Use the following procedure to download the current device mapping files and update the NetBackup Enterprise Media Manager database with their information.

See “About the device mapping files” on page 235.

To download the current device mapping files

- 1 Open the following location in your Web browser:

`http://entsupport.symantec.com`

- 2 In the **Knowledge Base Search** box, enter the following string (include the quotation marks) and then press Enter:

`"device mappings package"`

- 3 Select the package for your NetBackup release level and operating system.
- 4 Download the archive file, either a `.tar` or `.zip` depending on operating system.
- 5 Follow the instructions in the `Readme.txt` file to update the device mappings. The `Readme.txt` file contains instructions for both Windows and UNIX operating systems.

About configuring robots and tape drives

You can configure robots and tape drives in NetBackup as follows:

Device Configuration Wizard

Symantec recommends that you use the **Device Configuration Wizard** to add, configure, and update the following types of devices in NetBackup:

- 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 that are attached to the media servers and helps you configure them.

See “About device discovery” on page 237.

See “Configuring robots and tape drives by using the wizard” on page 241.

Manually

Alternatively, you can add robots and drives manually as follows:

- Use menu options in the **NetBackup Administration Console**.
 See “Adding a robot” on page 241.
 See “Adding a tape drive” on page 246.
- Use NetBackup commands.
 See *NetBackup Commands Reference Guide*.

Manual methods do not use device discovery.

If you add a robot and drives, 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*.

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.

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.

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

See “Enabling automatic path correction” on page 272.

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 device:

- 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 256.

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 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, select **NetBackup Administration Console > Media and Device Management** or use the `tpconfig` command.

For the devices that NetBackup cannot discover or that do not have serial numbers, automatic device path correction is limited.

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 (robot configuration options)” on page 243.

Table 7-2 lists the information that is required to configure the three robot control types (local, NDMP, and remote). The information that is 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	ACSL host
TL4	UNIX	Local	Robotic device file
TL4	Windows	Local	Robot device or SCSI coordinates

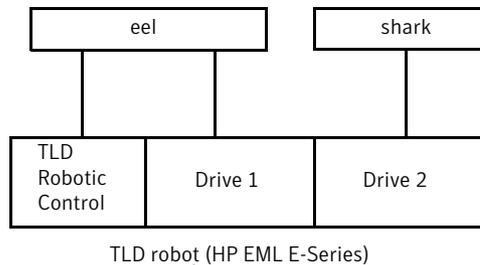
Table 7-2 Robot control information (continued)

Robot type	Media server type	Robot control	Information required for configuration
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

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

Configuring robots and tape drives

Symantec recommends that you use the **NetBackup Device Configuration Wizard** to configure robots and drives. However, you can add robots and drives manually.

Configuring robots and tape 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**, in the left pane, click **Media and Device Management**.
- 2 In the right pane, click the **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.

Adding a robot

When you add a robot manually, you must specify how the robot is controlled.

See “About NetBackup robot types” on page 234.

See “About robot control” on page 239.

After you add a robot, you should add the robot's drives.

See “Adding a tape drive” on page 246.

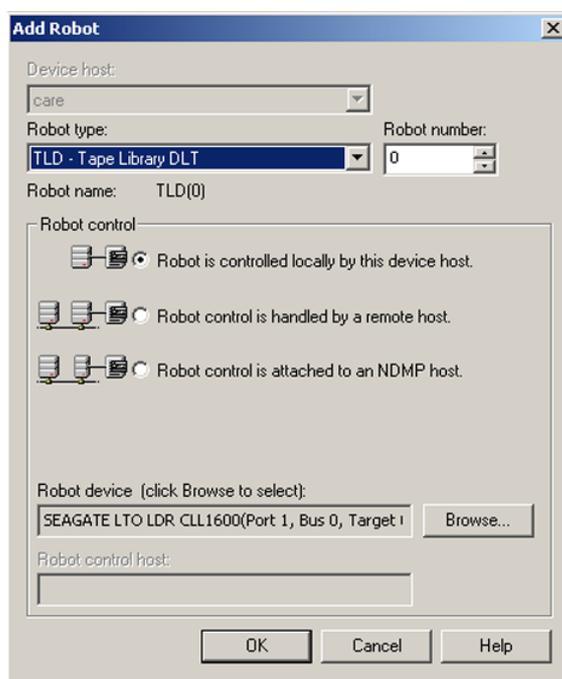
Note: Symantec recommends that you use the **Device Configuration Wizard** to add and update tape storage devices.

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 243.



- 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 that you can configure. The properties that you can configure depend on the robot type, host type, and robot control selections that you make in the dialog box.

Device host (robot configuration option)

Specifies the host to which the device is attached.

Robot type (robot configuration option)

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>

Robot number (robot configuration option)

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 (robot configuration options)

The **Robot control** section of the dialog box specifies the type of control for the robot. The options that you configure depend on the robot type and the media server type.

Table 7-3 Robot configuration properties

Property	Description
Robot control is attached to an NDMP host	Specifies that an NDMP host controls the robot. You must configure other options (depending on the robot type and device host type).
Robot is controlled locally by this device host	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).
Robot control is handled by a remote host	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).
ACSLS host	<p>Specifies 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.</p> <p>The ACS library software component can be any of the following:</p> <ul style="list-style-type: none"> ■ Automated Cartridge System Library Software (ACSLS) Examples are available. See the <i>NetBackup Device Configuration Guide</i>. ■ 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). <p>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.</p> <p>An overview of ACS robots is available. See the <i>NetBackup Device Configuration Guide</i>.</p>
DAS server	<p>Specifies the name of the ADIC DAS/SDLC server that controls TLM robots.</p> <p>This server is an OS/2 workstation near or within the robot cabinet or a Windows server near the ADIC Scalar library.</p> <p>An overview of TLM robots is available. See the <i>NetBackup Device Configuration Guide</i>.</p>

Table 7-3 Robot configuration properties (*continued*)

Property	Description
Library name	<p>The following applies only to a TLH robot on NetBackup Enterprise Server only.</p> <p>For UNIX device hosts (except AIX), specifies the library name that is configured on the UNIX host.</p> <p>For Windows devices hosts, do the following:</p> <ul style="list-style-type: none"> ■ Determine the library name by viewing the <code>C:\winnt\ibmat1.conf</code> file. For example, in the following example entry in that file, 3494AH is the library name: <code>3494AH 176.123.154.141 ibmpc1</code> ■ Enter the library name. <p>An overview of TLH robots is available.</p> <p>See the <i>NetBackup Device Configuration Guide</i>.</p>
LMCP device file	<p>Applies to NetBackup Enterprise Server on an AIX device host only.</p> <p>Specifies 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.</p>
NDMP host name	<p>Specifies the name of the NDMP host to which the robot is attached.</p>
Robot control host	<p>Specifies the host that controls the robot.</p> <p>The name of the host on which the robot information is defined for TL8, TLD, or TLH robots.</p>
Robot device	<p>The following applies to a Windows device host only. Specifies the name of the robot device.</p> <p>Click Browse and then select a robot from the list that appears in the Devices dialog box.</p> <p>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.</p> <p>You can find Port, Bus, Target, and LUN numbers by using Windows management tools.</p> <p>If the browse operation does not find attached robots, an error dialog box appears.</p>

Table 7-3 Robot configuration properties (*continued*)

Property	Description
Robotic device file	<p>UNIX device host only. Specifies the device file that is used for SCSI connections. The device files are located in the <code>/dev</code> directory tree on the device host.</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>If the browse operation does not find attached robots, an error dialog box appears.</p> <p>Information about how to add device files is available.</p> <p>See the <i>NetBackup Device Configuration Guide</i>.</p>
Robot device path	<p>NDMP host only. Specifies the name of the robotic device that is attached to the NDMP host.</p>
Port, Bus, Target, LUN	<p>Windows systems only. The Port, Bus, Target, and LUN are the SCSI coordinates for the robotic device. To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN.</p>

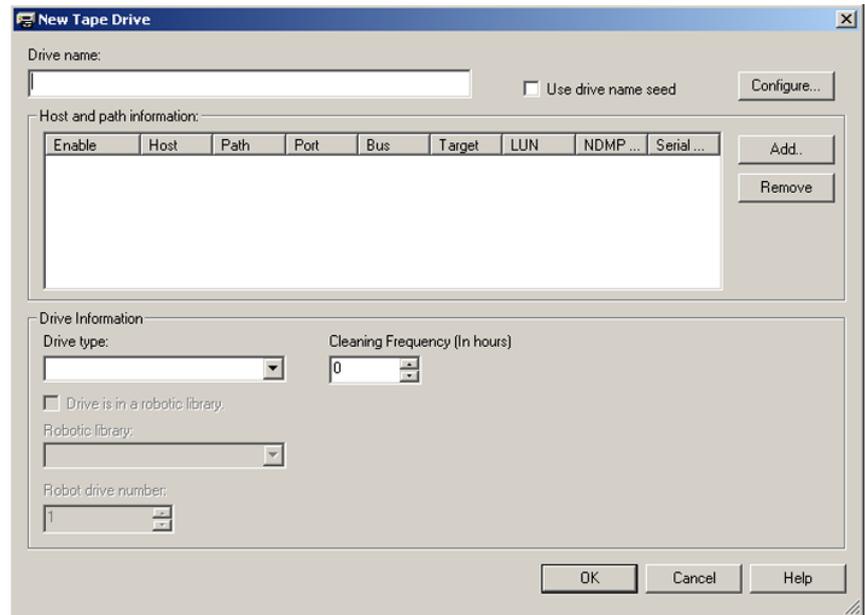
Adding a tape drive

Use the following procedures to add a tape drive manually.

Note: Symantec recommends that you use the **Device Configuration Wizard** to add and update tape storage devices.

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 For the drive name, do one of the following:
 - Enter a name for the drive in the **Drive name** field.
See “Drive name (tape drive configuration option)” on page 248.
 - Select **Use drive name seed**. This option uses rules to name the drive automatically.
See “About drive name rules” on page 251.
See “Configuring drive name rules” on page 252.
- 4 To configure the host and path information, click **Add** in the **Host and path information** area of the dialog box.
See “Host and path information (tape drive configuration options)” on page 249.
- 5 In the **Drive information** area of the dialog box, configure the drive properties. The properties depend on the drive type and host server type.
See “Drive information (tape drive configuration options)” on page 249.

- 6 After you configure all of the properties, click **OK**.
- 7 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

Symantec recommends that you use the **Device Configuration Wizard** to add, configure, and update shared drives. The **NetBackup Device Configuration Wizard** is the easiest method for adding shared drives in a Shared Storage Option configuration.

See the *NetBackup Shared Storage Guide*.

Tape drive configuration options

You can specify properties when you add a tape drive or change the properties of a drive. The properties that you can specify depend on the drive type, server platforms, or NetBackup server types.

Drive name (tape drive configuration option)

Specifies 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 the drive name seed to create a unique drive name.

Use drive name seed (tape drive configuration option)

Adds 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 251.

See “Configuring drive name rules” on page 252.

Host and path information (tape drive configuration options)

Use the **Host and path information** group box 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**.

To change a drive path, click **Change**.

To delete a drive path, click **Remove**.

See “About SCSI reserve on drive paths” on page 254.

See “Drive path options” on page 254.

Drive information (tape drive configuration options)

The **Drive information** group box includes drive properties. The properties that you can specify depend on the drive type, server platforms, and NetBackup server types.

Table 7-4 describes the tape drive configuration options.

Table 7-4 Tape drive configuration options

Option	Description
Drive type	<p>Specifies the type of drive. The following are the valid drive types:</p> <ul style="list-style-type: none"> ■ 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	<p>Specifies that the drive is in a robot. If the drive is a stand-alone drive (it is not in a robot), do not select this option.</p> <p>If you select this option, configure the Robotic library and Robot drive number fields.</p>

Table 7-4 Tape drive configuration options (*continued*)

Option	Description
Cleaning Frequency	<p>Specifies the frequency-based cleaning for the drive. NetBackup does not support drive cleaning in some robot types.</p> <p>If you want to configure 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.</p> <p>When the accumulated mount time exceeds the time you specify for cleaning frequency, drive cleaning occurs if the following are true:</p> <ul style="list-style-type: none"> ■ 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 <p>NetBackup resets the mount time when the drive is cleaned.</p> <p>Drives can also be cleaned from the Device Monitor.</p> <p>If you do not specify a cleaning frequency, you can still use automated drive cleaning with the TapeAlert feature. Information about TapeAlert drive cleaning is available.</p> <p>See the <i>NetBackup Administrator's Guide for Windows, Volume II</i>.</p>
Serial Number	<p>A read-only field that shows the serial number of the drive.</p>
Robotic library	<p>Specifies a robot that controls the drive. You can select any configured robot that can control the drive.</p>
Robot drive number	<p>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.</p> <p>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.</p> <p>You must determine which physical drive in the robot is identified by the logical device name.</p> <p>See “Correlating tape drives and SCSI addresses on Windows hosts” on page 256.</p> <p>NetBackup does not detect incorrect drive number assignment during configuration; however, an error occurs when NetBackup tries to mount media on the drive.</p> <p>Note: The Robot drive number property does not apply when you add drives to API robots. API robots are ACS, TLH, and TLM type in NetBackup.</p>

Table 7-4 Tape drive configuration options (*continued*)

Option	Description
ACS, LSM, Panel, Drive	<p>Specify the drive locations within an ACS robot.</p> <p>The following information applies only to the ACS robot drive. The ACS property specifies 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.</p> <p>The drive location properties are as follows:</p> <ul style="list-style-type: none"> ■ ACS Number - specifies the index (in ACS library software terms) that identifies the robot that has this drive. ■ LSM Number - specifies the Library Storage Module that has this drive. ■ Panel Number - specifies the robot panel where this drive is located. ■ Drive Number - specifies the physical number of the drive (in ACS library software terms).
IBM device number	<p>Specifies the IBM device number of the drive within the robot. This property applies only to the TLH robot drive.</p>
DAS drive name	<p>Specifies the DAS/SDLC drive name of the drive within the robot. This property applies only to the TLM robot drive.</p>

About drive name rules

The 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.

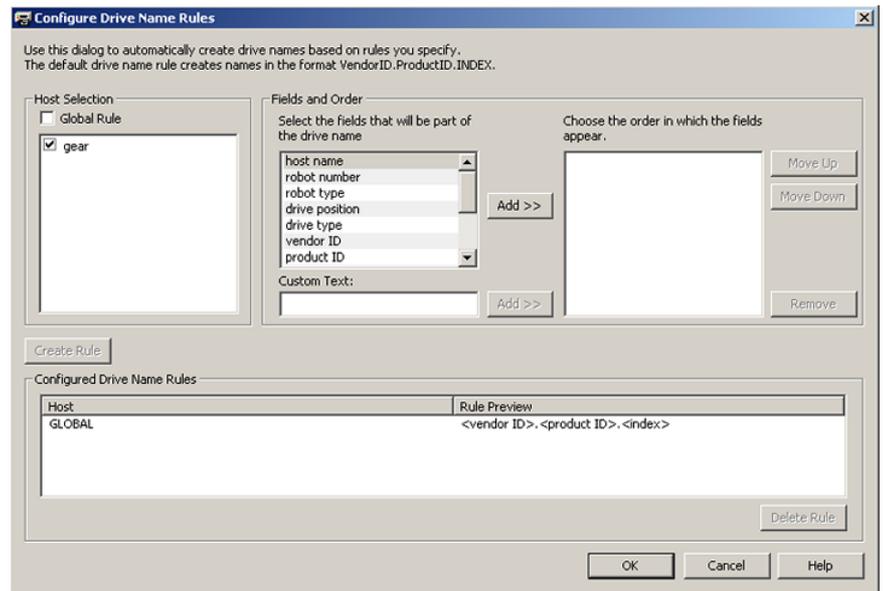
See “Configuring drive name rules” on page 252.

Configuring drive name rules

Use the following procedure to configure the drive name rules.

To configure drive name rules

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Drives**. Expand the **Actions > New > New Tape Drive** menu. See “Adding a tape drive” on page 246.
- 2 In the **New Tape Drive** dialog box, click **Configure**.
 Alternatively, if you use the **NetBackup 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

Usually, you add a tape drive path when you add a drive to NetBackup. Use the following procedure to add a drive path.

To add a tape drive path

1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**. In the **Actions** menu, select **New > New Tape Drive**. 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 that you can specify depend on drive type, server platform, or NetBackup server type.

See “About SCSI reserve on drive paths” on page 254.

See “Drive path options” on page 254.

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, Volume II*.

Drive path options

The following table describes the options to add a drive path.

Table 7-5 Add drive path options

Option	Description
Host name	Specifies the device host for the drive.
Enable host path	Specifies that the path is active and that NetBackup can use it for backups and restores.

Table 7-5 Add drive path options (*continued*)

Option	Description
NDMP host	<p>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>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>To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN.</p> <p>The device attributes on Windows systems cannot change during a NetBackup operation.</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 that you can change depend on the robot type, the host type, and the robot control.

See “Robot configuration options” on page 243.

- 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 797.

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 295.

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 227.

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 7-6 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: <pre>bpmedialist -mlist -l -h old_server</pre> 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 308.

Table 7-6 Move a robot and media to a new server overview (continued)

Task	Procedure
Move the media logically from the <i>old_server</i> to the <i>new_server</i> .	<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>bpmmedia -movedb -allvolumes -oldserver old_server -newserver new_server</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>bpmmedia -movedb -ev media_ID -oldserver old_server -newserver new_server</pre> <p>For the media that has active images, see the <code>bpmmedialist</code> command output from the first step of this process.</p>
Configure NetBackup so that restore requests are directed to the <i>new_server</i> .	See “Forcing restores to use a specific server” on page 130.
Shut down both the <i>old_server</i> and the <i>new_server</i> .	See the vendor's documentation.
Disconnect the robot from the <i>old_server</i> .	See the vendor's documentation.
Connect the robot to the <i>new_server</i> . Verify that the operating system on the new media server recognizes the robots.	See the vendor's documentation.
Use the NetBackup Device Configuration Wizard to add the robots and drives to the media servers.	See “Configuring robots and tape drives by using the wizard” on page 241.
Create the appropriate NetBackup storage units.	See “Creating a storage unit using the Actions menu” on page 388.
Inventory the robots that are attached to the <i>new_server</i> . The inventory updates the location of all tapes in the robot.	See “Updating the volume configuration with a robot's contents” on page 333.

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 827.
- 7 Click **OK**.

About downed drives

NetBackup downs a drive automatically when there are read or write errors that 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 `nbeMMCcmd` command with the `--drive_error_threshold` and `-time_window` options to change the default values.

Additional information about `nbeMMCcmd` is available.

See *NetBackup Commands Reference Guide*.

To reverse a down action, in the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor** to set the device to Up.

See “Changing a drive operating mode” on page 262.

Changing a drive operating mode

Usually you do not need to change the operating mode of a drive. When you add a 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 263.

To change a drive path

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**. Double-click on the drive that you want to change. In the **Change Tape Drive** dialog box, select the drive path.
- 2 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 254.

See “Drive path options” on page 254.

Changing a drive path operating mode

In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**. In the right pane of the **Device Monitor** dialog box, the **Drive Paths** pane 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 248.

- 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 266.

See the *NetBackup Administrator’s Guide, Volume II*.

Drive cleaning functions can also be performed from the **Activity Monitor**.

See “Cleaning tape drives from the Activity Monitor” on page 792.

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, expand **Drive Cleaning > Clean Now**. NetBackup initiates drive cleaning regardless of the cleaning frequency or accumulated mount time.

The **Clean Now** option 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 227.

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**, click on **Activity Monitor**. In the right pane of the **Activity Monitor** dialog box, select the **Jobs** tab and 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 > Drive Cleaning > 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. In the **NetBackup Administration Console**, expand **Media and Device Management > 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, expand **Drive Cleaning > Set Cleaning Frequency**.
- 5 Enter a time (hours) or use the arrow controls to select the number of mount hours between drive cleaning.

The **Cleaning Frequency** option 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

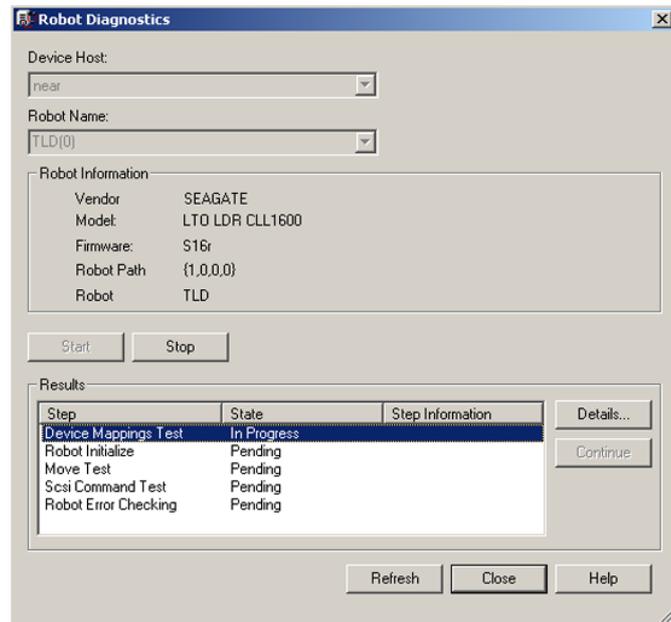
Use this procedure to run diagnostic tests on TLD or TL8 robotic libraries.

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 pre-labeled 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.

About automatic path correction

NetBackup automatic path correction 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.

By default, Windows and Linux systems are configured for automatic path correction. On other operating systems, you must enable it.

See “Enabling automatic path correction” on page 272.

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 **NetBackup Device Configuration Wizard**.

See “Configuring robots and tape drives by using the wizard” on page 241.

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.

Enabling automatic path correction

You can configure NetBackup to automatic device path correction. To do so, use the following procedure.

See “About automatic path correction” on page 272.

To configure automatic path correction

- 1 Use a text editor to open the following file:

```
install_path\VERITAS\Volmgr\vm.conf
```

- 2 Add the following `AUTO_PATH_CORRECTION` entry to the file:

```
AUTO_PATH_CORRECTION = YES
```

If it already exists but is set to **NO**, change the value to **YES**.

- 3 Save the file and exit the text editor.

Replacing a device

Table 7-7 describes the process to replace a device on a single host.

Table 7-8 describes the process to replace a shared device.

Table 7-7 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 262.
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 262.
If either of the following are true, configure the new device by using the NetBackup 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 tape drives by using the wizard” on page 241.

Table 7-8 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 262.

Table 7-8 To replace a shared device (*continued*)

Task	Instructions
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.	<p>The following command scans for new hardware and produces a report that shows the new and the replaced hardware:</p> <pre data-bbox="803 525 1225 583">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 782.
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 data-bbox="803 790 1225 874">install_path\Veritas\Volmgr\bin\tpautoconf -replace_robot robot_number -path robot_path</pre> <p>If the device is a drive, run the following commands:</p> <pre data-bbox="803 966 1225 1051">install_path\Veritas\Volmgr\bin\tpautoconf -replace_drive drive_name -path path_name</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 tape drives by using the wizard” on page 241.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 262.

Updating device firmware

By default, NetBackup recognizes if you update the firmware of a device.

The following table describes an overview of how to update device firmware.

Table 7-9 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 262.
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 262.

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 (`avr.d.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.

See “Stopping and restarting the Device Manager” on page 275.

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.
- 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 tape volumes
- NetBackup media types
- About WORM media
- About adding volumes
- 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 unrestricted media sharing
- Configuring media sharing with a server group

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 666.

NetBackup uses two volume types, as follows:

Robotic volumes	Volumes that are located in a robot.
Stand-alone volumes	Volumes that are 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 adding volumes, use the NetBackup Catalog Backup wizard to configure a catalog backup policy.

See “About NetBackup catalogs” on page 661.

WORM media can be used with NetBackup.

See “About WORM media” on page 280.

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 8-1 describes the NetBackup media types.

Table 8-1 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

Table 8-1 NetBackup media types (*continued*)

Media type	Description
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, Volume II*.

Alternate NetBackup 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, 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 `bplabel` 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.

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 281.

See “About using unique drive and media types to manage WORM media” on page 283.

Supported WORM drives

NetBackup requires a 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 314.

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 backups 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.

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.

About adding robotic volumes

Robotic volumes are the volumes that are located in a robotic tape library.

Table 8-2 Methods for adding robotic volumes

Method	Description
The Volume Configuration Wizard	See “Adding volumes by using the wizard” on page 285.
Robot inventory	See “Updating the volume configuration with a robot’s contents” on page 333.
The Actions menu	See “Adding volumes by using the Actions menu” on page 285.
NetBackup commands	See <i>NetBackup Commands Reference Guide</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 8-3 Methods for adding stand-alone volumes

Method	Description
The Volume Configuration Wizard	See “Adding volumes by using the wizard” on page 285.
The Actions menu	See “Adding volumes by using the Actions menu” on page 285.
NetBackup commands	See <i>NetBackup Commands Reference Guide</i> .

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**, in the left pane, expand **Media and Device Management > Devices**.
- 2 From the list of wizards in the right 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**, in the left pane, expand **Media and Device Management > Media**.
- 3 On the **Actions** menu, select **New > New Volumes**.

The screenshot shows the 'New Volumes' dialog box with the following configuration:

- Media type: DLT cartridge tape
- Robot: TLD(0) - ga
- Number of volumes: 1
- Media ID: A00001
- First slot number: 1
- Volume group: (empty)
- Media ID naming style: (empty)
- Media Description: (empty)
- Maximum mounts/cleanings: 0
- Volume pool: NetBackup

- 4 In the **Add Volumes** dialog box, specify the attributes for the volumes.
 See “Add volume properties” on page 286.
- 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.

Add volume properties

Table 8-4 describes the properties to configure when you add volumes. The topics are arranged alphabetically.

Table 8-4 Volume properties when adding volumes

Property	Description
Device host	The name of the host to which the robot is attached.
First media ID	<p>This property appears only if the number of volumes is more than one.</p> <p>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.</p> <p>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.</p> <p>NetBackup allows specific characters in names.</p> <p>See “NetBackup naming conventions” on page 827.</p>
First slot number	<p>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.</p> <p>Note: You cannot enter slot information for volumes in an API robot. The robot vendor tracks the slot locations for API robot types.</p>

Table 8-4 Volume properties when adding volumes (*continued*)

Property	Description
Maximum cleanings	<p>The maximum number of times NetBackup should mount the volume or use the cleaning tape.</p> <p>When a volume reaches the mount limit, the volume can be read, but not written. Zero (0) indicates unlimited mounts. If you enter a value larger than 99999, NetBackup may display it as 0 although it uses the actual value. For example, the output of the <code>vmrule</code> command displays 0 for values larger than 99999.</p> <p>To determine the maximum mount limit to use, consult the vendor documentation for information on the expected life of the volume.</p>
Media description	<p>A description of the media, up to 25 character maximum.</p> <p>NetBackup allows specific characters in names.</p> <p>See “NetBackup naming conventions” on page 827.</p>
Media ID	<p>This property appears only if the number of volumes is one.</p> <p>The ID for the new volume. Media IDs can be from 1 to 6 characters in length.</p> <p>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.</p> <p>NetBackup allows specific characters in names.</p> <p>See “NetBackup naming conventions” on page 827.</p>

Table 8-4 Volume properties when adding volumes (*continued*)

Property	Description
<p>Media ID naming style</p>	<p>This property appears only if the number of volumes is more than one.</p> <p>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.</p> <p>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.</p> <p>NetBackup allows specific characters in names.</p> <p>See "NetBackup naming conventions" on page 827.</p>
<p>Media type</p>	<p>The media type for the volume to add.</p> <p>Select the type from the drop-down list.</p> <p>See " NetBackup media types" on page 278.</p>
<p>Number of volumes</p>	<p>The number of volumes to add. For a robotic library, enough slots must exist for the volumes.</p>
<p>Robot</p>	<p>The robotic library to add the volumes to.</p> <p>To add volumes for a different robot, select a robot from the drop-down list . The list shows robots on the selected host that can contain volumes of the selected media type.</p> <p>To find a robot that does not appear in the Robot box, click Find Robots to open the Find Robot dialog box.</p> <p>To add volumes to a stand-alone drive, select Standalone.</p>

Table 8-4 Volume properties when adding volumes (*continued*)

Property	Description
Volume group	<p>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.</p> <p>If you do not specify a volume group (you leave the volume group blank), the following occurs:</p> <ul style="list-style-type: none"> ■ 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. <p>See “About volume groups” on page 316.</p>
Volume is in a robotic library	<p>To specify that the volume is in a robot, select Volume is in a robotic library. If the volume is a stand-alone volume, do not select this option.</p>
Volume pool	<p>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:</p> <ul style="list-style-type: none"> ■ 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. <p>When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool.</p> <p>See “About volume pools” on page 312.</p>

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 rules for moving volumes between groups” on page 290.

To change the group of a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, 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 left pane.

About rules for moving volumes between groups

The following are the rules for moving volumes between groups:

- The target volume group must contain the same type of media as the source volume group. If the target volume group is empty: The successive volumes that you add to it must match the type of media that you first add to it.
- All volumes in a robotic library must belong to a volume group. If you do not specify a group, NetBackup generates a new volume group name by using the robot number and type.
- 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 stand-alone volume group.
- All members of a group must be in the same robotic library or be stand-alone. That is, if volume group already exists in another robotic library, you cannot add it (or part of it) to a robotic library.

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 209.

See “About media sharing” on page 317.

To change the owner of a volume

- 1 In the **NetBackup Administration Console**, in the left pane, 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 291.

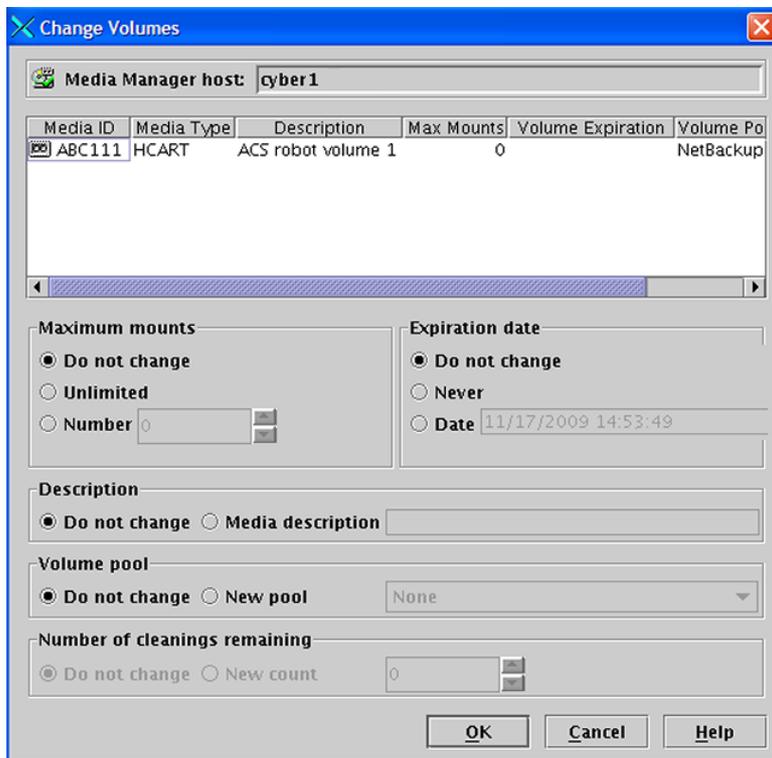
Changing volume properties

You can change some of the properties of a volume, including the volume pool.

To change volume properties

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, 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 “Change volume properties” on page 292.
- 5 Click **OK**.

Change volume properties

Table 8-5 describes the volume properties that you can change.

Table 8-5 Volume properties when changing volumes

Property	Description
Description	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 827.

Table 8-5 Volume properties when changing volumes (*continued*)

Property	Description
Expiration date	<p>The following does not apply to cleaning tapes.</p> <p>The date after which the volume is too old to be reliable.</p> <p>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.</p> <p>See “About exchanging a volume” on page 297.</p> <p>When you add a new volume, NetBackup does not set an expiration date.</p> <p>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.</p>
Maximum mounts	<p>The following topic does not apply to cleaning tapes.</p> <p>The Maximum mounts property specifies the number of times that the selected volumes can be mounted.</p> <p>When the limit is reached, NetBackup reads data on the volume but does not mount and write to the volume.</p> <p>A value of zero (the default) is the same as Unlimited.</p> <p>To help determine the maximum mount limit, consult the vendor documentation for information on the expected life of the volume.</p>
Number of cleanings remaining	<p>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.</p> <p>Additional information about drive cleaning is available.</p> <p>See the <i>NetBackup Administrator's Guide, Volume II</i>.</p>
Volume pool	<p>The following topic does not apply to cleaning tapes.</p> <p>The pool to which the volume or volumes should be assigned.</p> <p>Select a volume pool you created or one of the following standard NetBackup pools:</p> <ul style="list-style-type: none"> ■ 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. <p>When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool.</p> <p>See “About volume pools” on page 312.</p>

About assigning 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.

To determine which application currently uses a volume, see the **Application** column of the right pane, labeled **Volumes**.

See “About deassigning volumes” on page 294.

About deassigning volumes

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 the volume for catalog backups.

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 750.

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 “About importing backup images” on page 750.

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.

See “About assigning volumes” on page 294.

Deleting a volume

You can delete volumes from the NetBackup configuration.

Note: 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 “About deassigning volumes” on page 294.

To delete volumes

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, 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 “About deassigning volumes” on page 294.

Table 8-6 Types of erase

Type of erase	Description
SCSI long erase	Rewinds the media and the data is overwritten with a known data pattern. A SCSI long erase is also called a secure erase because it erases the recorded data completely. Note: A long erase is a time-consuming operation and can take as long as two hours to three hours. For example, it takes about 45 minutes to erase a 4-mm tape on a standalone drive
SCSI quick erase	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. 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.

Note: NetBackup does not support erase functions on NDMP drives.

To erase a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, 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.

About 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).

Depending on whether you want to reuse the old media ID or not, follow one of the exchange volumes processes in the following subsections.

Exchanging a volume and using a new media ID

Use this procedure when the following 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 8-7 Exchange a volume and using a new media ID

Step	Task	Instructions
Step 1	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 306.

Table 8-7 Exchange a volume and using a new media ID *(continued)*

Step	Task	Instructions
Step 2	<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 283.
Step 3	<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 8-8 Exchange a volume and use the old media ID

Step	Task	Instructions
Step 1	Delete the volume.	See “Deleting a volume” on page 295.
Step 2	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 300.
Step 3	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 283.
Step 4	Set a new expiration date for the volume.	See “Changing volume properties” on page 291.
Step 5	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 305.

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 “About importing backup images” on page 750.

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 *NetBackup Commands Reference Guide*.

To reverse a freeze action, use the `bpmedia` command to unfreeze the volume.

See “Freezing or unfreezing a volume” on page 299.

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**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select the volume that you want to freeze or unfreeze.
- 3 On the **Actions** menu, select **Freeze** or **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 into robots

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 333.
- 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.

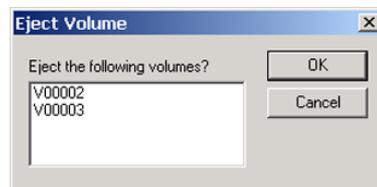
See “Media ejection timeout periods” on page 302.

To eject volumes

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, 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.



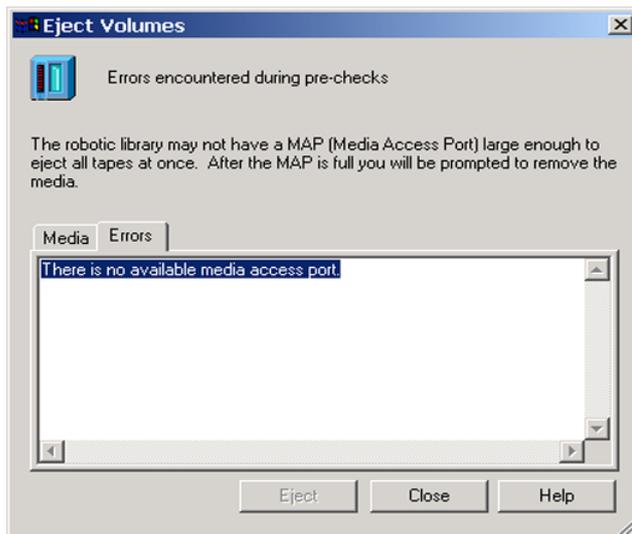
If you select multiple volumes, operator action is required to remove each volume after each eject (a dialog box prompts you to remove volumes).

- 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** option 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.

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 8-9 shows the ejection timeout periods for robots.

Table 8-9 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.

When not to rescan and update bar codes

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 333.
- 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 328.

When to rescan and update bar codes

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.

See “About bar codes” on page 344.

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 303.

To rescan bar codes and update the EMM database

- 1 In the **NetBackup Administration Console**, in the left pane, 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 right pane, in the **Volumes** list, 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 351.
- 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, Volume II*.

Media is 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 types, the following must be true:

- NetBackup has not assigned the media
- The media contains no valid NetBackup images

Labeling a volume

If a volume contains valid NetBackup images, deassign the volume so that it can be labeled.

See “About deassigning volumes” on page 294.

If you want to label media and assign specific media IDs (rather than allow NetBackup to assign IDs), use the `bplabel` 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 305.

To label a volume

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, 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.
- Move 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:

- When 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 322.

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**, in the left pane, expand **Media and Device Management > Media**.
- 3 In the right pane, in the **Volumes** list, 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 308.
- 5 In the **Move Volumes** dialog box, specify the properties for the move.

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

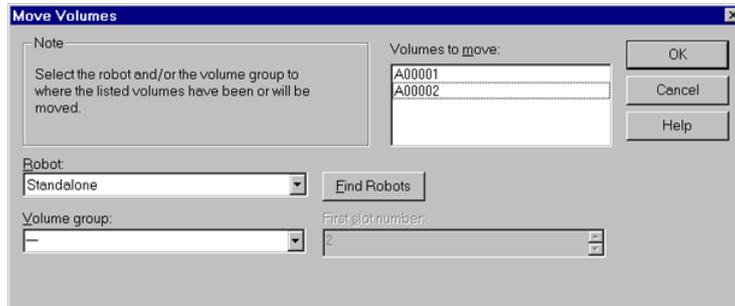
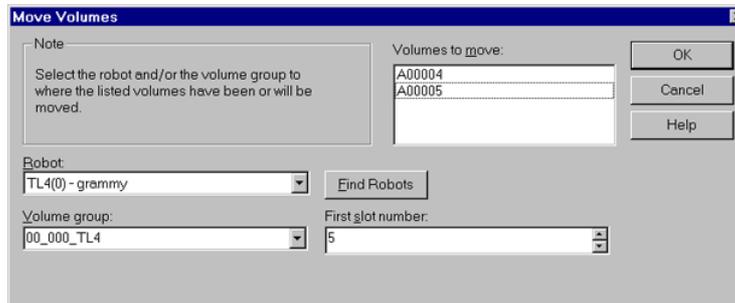


Figure 8-2 Move volumes to the robot



Move volumes properties

Table 8-10 describes the properties to configure in the **Move Volumes** dialog box.

Table 8-10 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>

Table 8-10 Move volumes properties (*continued*)

Property	Description
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. <p>See “About rules for moving volumes between groups” on page 290.</p>
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>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>

About 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 “About deassigning volumes” on page 294.

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 bar codes on the volume.

Table 8-11 Recycling a volume and using a new media ID

Step	Action	Description
Step 1	Physically remove the volume from the storage device.	See “Ejecting volumes” on page 300.
Step 2	If the volume is in a robotic library, move it to stand alone.	See “About moving volumes” on page 306.
Step 3	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).
Step 4	Delete the volume entry.	See “Deleting a volume” on page 295.
Step 5	Add a new volume entry.	<p>See “Adding volumes by using the Actions menu” on page 285.</p> <p>Because NetBackup sets the mount value to zero for new volume entries, you must adjust the value to account for previous mounts.</p> <p>Set the maximum mounts to a value that is equal to or less than the following value:</p> <p>The number of mounts that the manufacturer recommends minus the value that you recorded earlier.</p>
Step 6	Physically add the volume to the storage device.	See “Injecting volumes into robots” on page 300.

Table 8-11 Recycling a volume and using a new media ID (*continued*)

Step	Action	Description
Step 7	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>
Step 8	Set the expiration date to the value you recorded earlier.	See “Changing volume properties” on page 291.

Suspending or unsuspending volumes

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**, in the left pane, select **Media and Device Management > Media**.
- 2 In the right pane, in the **Volumes** list, select the volume or volumes that you want to suspend or unsuspend.
- 3 On the **Actions** menu, select **Suspend** or **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, 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**, in the left pane, 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 314.

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.

Table 8-12

Property	Description
Catalog backup pool	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	Provides a brief description of the volume pool.
Maximum number of partially full media	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: <ul style="list-style-type: none"> ■ 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.

Table 8-12 (continued)

Property	Description
Pool name	The Pool name is the name for the new volume pool. Volume pool names are case-sensitive and can be up to 20 characters.
Scratch pool	Specifies that the pool should be a scratch pool. Symantec recommends that you use a descriptive name for the pool and use the term <code>scratch pool</code> 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 volume pools” on page 312.

See “Adding a volume pool” on page 314.

See “Changing the properties of a volume pool” on page 315.

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**, in the left pane, select **Media and Device Management > Media > Volume Pools**.
- 2 Select a pool 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 314.

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**, in the left pane, 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, Volume II*.

About media sharing

Media sharing allows media servers to share media for write purposes (backups).

Media sharing provides the following benefits:

- Increases the utilization of media by reducing the number of partially full media.
- Reduces media-related expenses because fewer tape volumes are required and fewer tape volumes are vaulted (NetBackup Vault option).
- Reduces administrative overhead because you inject fewer scratch media into the robotic library.
- Increases the media life because tapes 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.

You can configure the following media sharing:

- Unrestricted media sharing.
See “Configuring unrestricted media sharing” on page 318.
- Media media sharing with server groups.
See “Configuring media sharing with a server group” on page 318.\

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 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**, in the left pane, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the right 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 209.

Table 8-13 outlines the process for configuring media sharing with a server group.

Table 8-13 Configuring media sharing with a server group process overview

Step	Action	Description
Step 1	Ensure the appropriate connectivity between and among the media servers and robots and drives.	Beyond the scope of the NetBackup documentation.

Table 8-13 Configuring media sharing with a server group process overview
(continued)

Step	Action	Description
Step 2	Configure the media sharing server group.	See “Configuring a server group” on page 210.
Step 3	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 314. See “Changing the properties of a volume pool” on page 315.
Step 4	Configure the 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 508.

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 `vmphyinv` physical inventory utility
- Example volume configuration updates

About robot inventory

Robot inventory is a logical operation that verifies the presence of media. (Robot inventory does not inventory the data on the media.)

After you physically add, remove, or move volumes in a robot, use a robot inventory to update the NetBackup volume configuration.

See “When to inventory a robot” on page 323.

See “Example volume configuration updates” on page 369.

See “How to access media and devices on other hosts” on page 831.

Table 9-1 describes the **NetBackup Administration Console** robot inventory options for the robotic libraries that contain bar code readers and contain bar coded media.

Table 9-1 Robot inventory options

Inventory option	Description
Show contents	<p>Displays the media in the selected robotic library; does not check or change the EMM database.</p> <p>See “About showing a robot's contents” on page 325.</p> <p>For the robotic libraries without bar code readers (or that contain media without bar codes), you can only show the contents of a robot. However, more detailed information is required to perform automated media management. Use the <code>vmphyinv</code> physical inventory utility to inventory such robots.</p> <p>See “About the <code>vmphyinv</code> physical inventory utility” on page 361.</p>
Compare contents with volume configuration	<p>Compares the contents of a robotic library with the contents of the EMM database but does not change the database.</p> <p>See “About comparing a robot's contents with the volume configuration” on page 329.</p>
Preview changes	<p>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.</p> <p>See “About previewing volume configuration changes” on page 333.</p>

Table 9-1 Robot inventory options (*continued*)

Inventory option	Description
Update volume configuration	<p>Updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur.</p> <p>See “About updating the volume configuration” on page 331.</p>

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 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 328.</p>
To determine if volumes were moved physically within a robot	<p>For robots with bar code readers and robots 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 330.</p>
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 333.</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 333.</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 333.</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 306.</p> <p>See “About the vmphyinv physical inventory utility” on page 361.</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 333.</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 333.</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 306.</p> <p>See “About the vmphyinv physical inventory utility” on page 361.</p> <p>See “Volume Configuration Example 7: Adding existing volumes when bar codes are not used” on page 378.</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 333.</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 306.</p> <p>See “About the vmphyinv physical inventory utility” on page 361.</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 “Volume Configuration Example 6: Moving existing volumes between robots” on page 377.</p>
<p>To remove existing volumes from a robot (an existing volume is one that already has a NetBackup media ID)</p>	<p>For any robot NetBackup supports, use the Update volume configuration option to update the NetBackup volume configuration information.</p> <p>See “Updating the volume configuration with a robot's contents” on page 333.</p>

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.

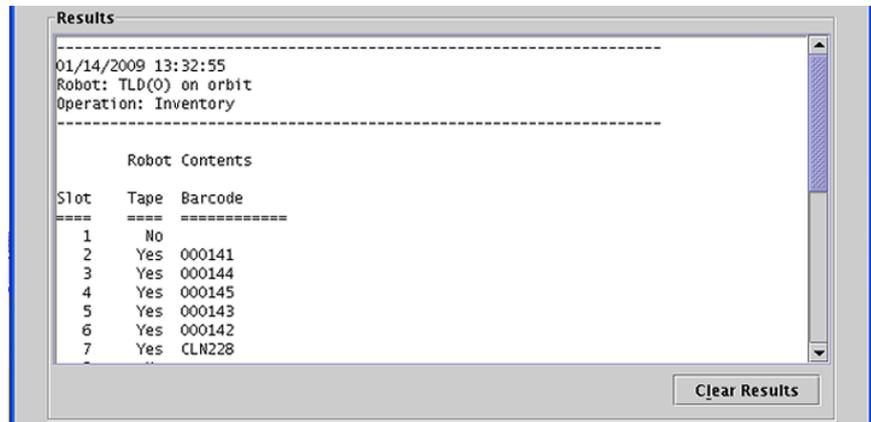
Table 9-3 describes the report contents.

Table 9-3 Show contents description

Robot and media	Report contents
The robot has a bar code reader and the robot contains media with bar codes.	Shows if each slot has media and lists the bar code for the media.
The robot does not have a bar code 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 326.

Figure 9-1 is an example of the report.

Figure 9-1 Show contents report



See "Showing the media in a robot" on page 328.

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).
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 9-2 shows the results for an ACS robot; the results for other API robots are similar.

Figure 9-2 Show contents report (API robot)

Media ID	ACS Media Type	MM Media Type
000026	STK1R	HCART
000042	STK1R	HCART
000043	STK1R	HCART
000044	STK1R	HCART
000045	STK1R	HCART
000046	STK1R	HCART
000047	STK1R	HCART
000049	STK1R	HCART
000051	STK1R	HCART

Showing the media in a robot

Use the following procedure to show the media that is in a robot.

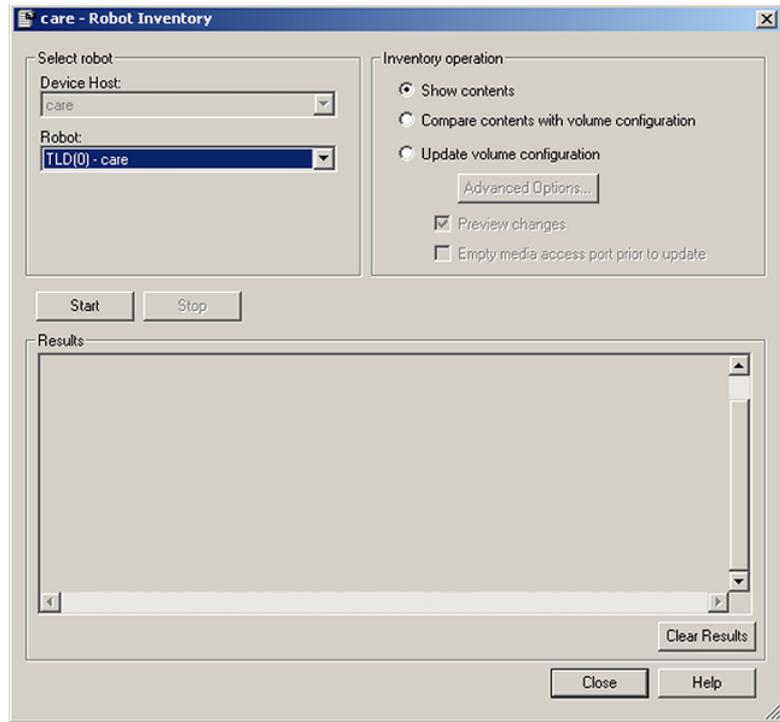
See “About robot inventory” on page 322.

See “Robot inventory options” on page 336.

To show the media in a robot

- 1 In the **NetBackup Administration Console**, in the left pane, 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.

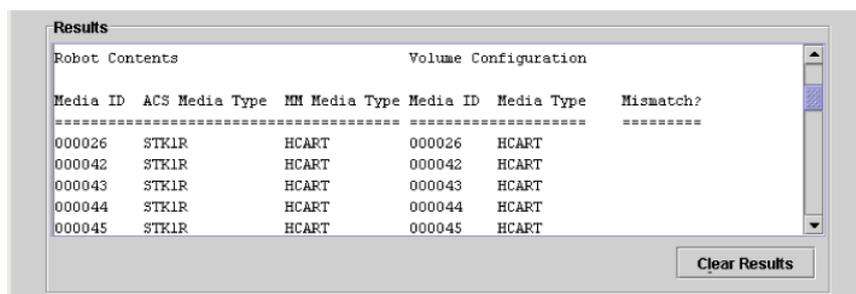
If the results show that the EMM database does not match the contents of the robotic library, perform the following actions:

- 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 331.

Figure 9-3 shows a sample compare report.

Figure 9-3 Compare contents report (API robot)



See “Comparing media in a robot with the volume configuration” on page 330.

Comparing media in a robot with the volume configuration

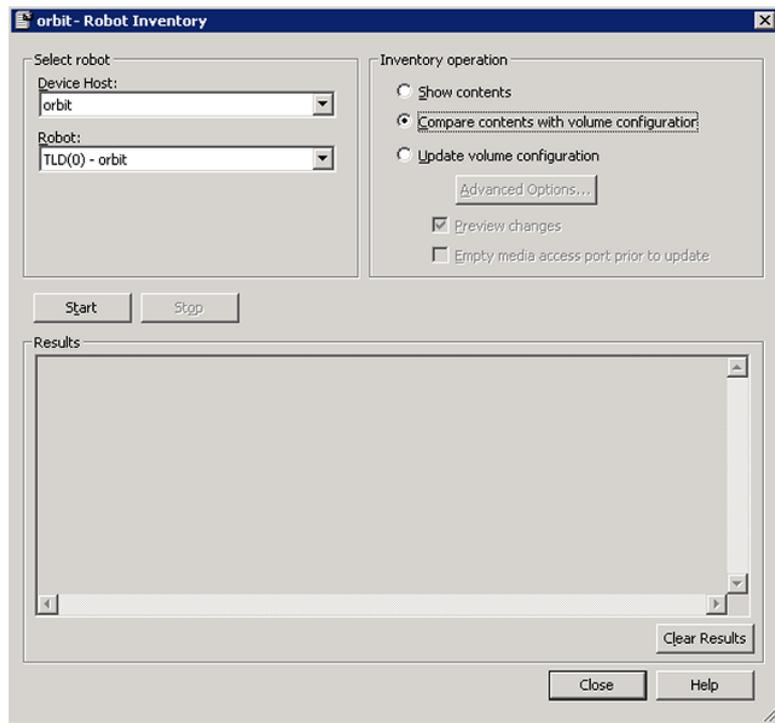
Use the following procedure to compare the media in a robot with the EMM database.

See “About robot inventory” on page 322.

See “Robot inventory options” on page 336.

To compare media in a robot with the volume configuration

- 1 In the **NetBackup Administration Console**, in the left pane, 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 336.

For API robots, the update returns an error if the volume serial number or the media ID contain unsupported characters.

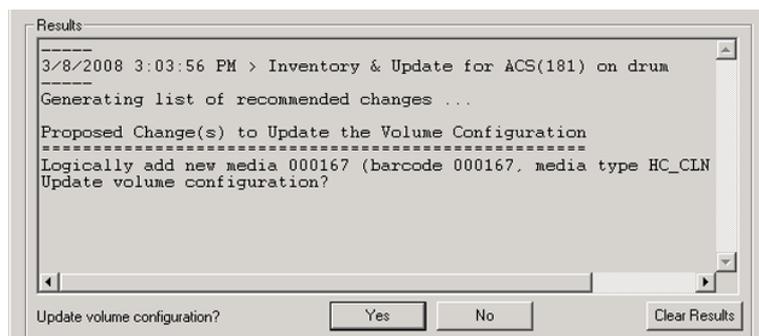
For robots without bar code readers, the new media IDs are based on a media ID prefix that you specify. Similarly, for volumes without readable bar codes, the new media IDs are based on a media ID prefix that you specify

Figure 9-4 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 “Volume update prerequisites” on page 332.

Figure 9-4 Update volume configuration for API robot report



See “Updating the volume configuration with a robot's contents” on page 333.

Volume update prerequisites

The following are the robot prerequisites and media prerequisites for updating the volume configuration:

- The robotic library must read bar codes.
- Volumes in the library must have readable bar codes.

You can check the bar code capabilities of the robotic library and the volumes by comparing the robot contents with the NetBackup volume configuration.

See “Comparing media in a robot with the volume configuration” on page 330.

If the robotic library does not support bar codes or the volumes do 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.

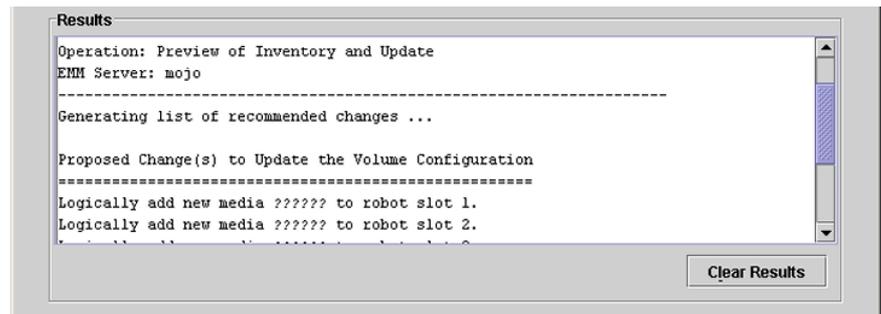
About previewing volume configuration changes

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 and 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 9-5 Preview volume configuration changes (non-API robot)



See “Updating the volume configuration with a robot's contents” on page 333.

Updating the volume configuration with a robot's contents

Use the procedure in this topic to update the EMM database with the contents of a robot.

See “About robot inventory” on page 322.

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.

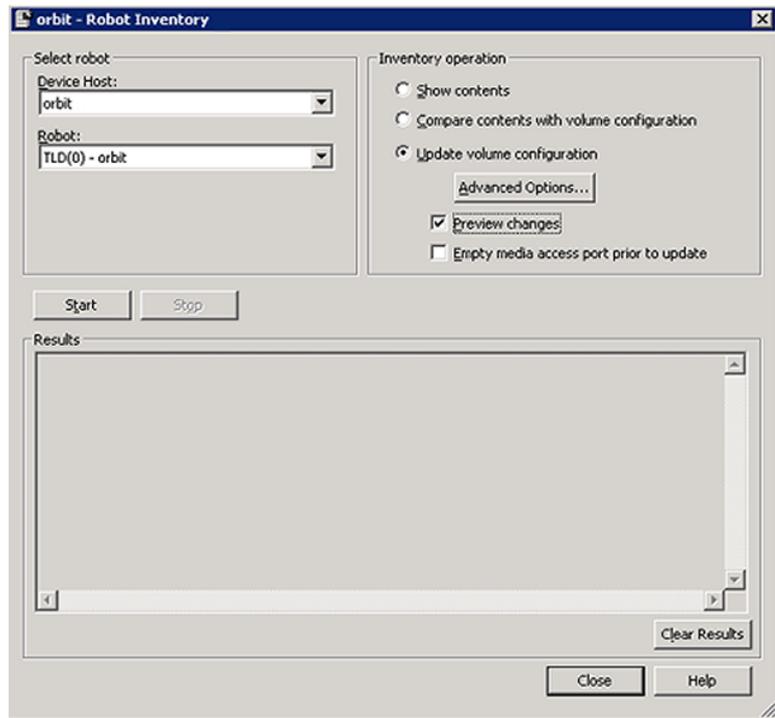
Table 9-6 Attributes for new media

What	Where
Media settings	See “Configuring media settings” on page 337.
Bar code rules	See “Configuring bar code rules” on page 349.
Media ID generation rules	See “Configuring media ID generation rules” on page 351.
Map media for API robots	See “Configuring media type mappings” on page 354.

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**, in the left pane, 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 robot inventory options are available 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 337.
- See “Configuring bar code rules” on page 349.
- See “Configuring media ID generation rules” on page 351.
- See “Configuring media type mappings” on page 354.
- 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 325.

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 329.
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 333.
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 331.

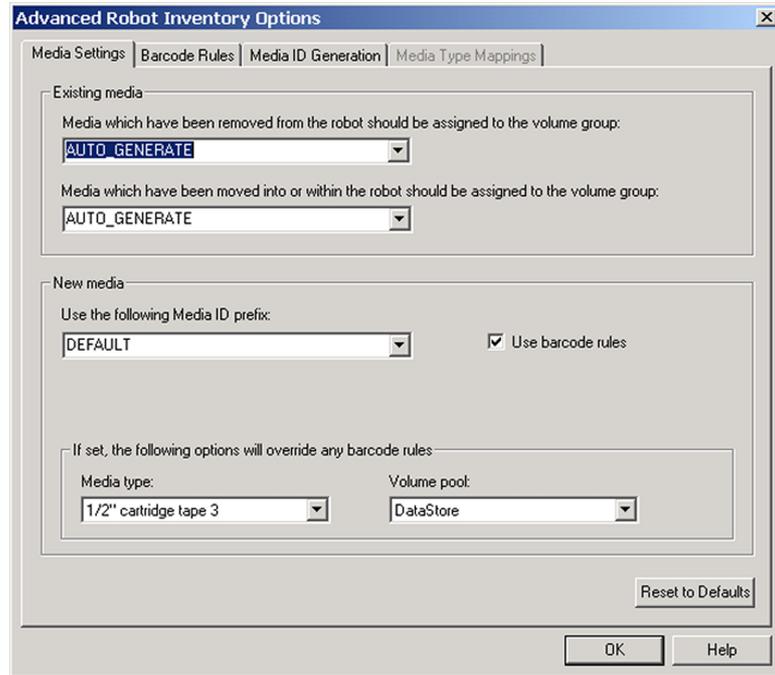
Configuring media settings

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 338.
See “Media settings - new media” on page 340.
- 4 Click **OK**.

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 that 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.

■ **Media that 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 settings - media type” on page 341.

Media settings - new media

For new media in the robot to add to your volume configuration, specify the attributes for the new media.

Media settings - 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 <code>DEFAULT</code> 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 <code>A</code> if no <code>MEDIA_ID_PREFIX</code> entries are defined.
NOT USED	If <code>NOT USED</code> is selected, the operation succeeds only if the robot supports bar codes and the volume has readable bar codes. <code>NOT USED</code> 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, Volume II*.

Media settings - use bar code 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 344.

See “Configuring bar code rules” on page 349.

Information about the `vm.conf` file is available.

See the *NetBackup Administrator's Guide, Volume II*.

Media settings - 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 354.

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 that 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
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 “Volume Configuration Example 5: Adding cleaning tapes to a robot” on page 376.

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 349.

Media type when not using bar code rules

Choose one of the following if bar code rules are not used:

- | | |
|-----------------------|--|
| DEFAULT | <p>NetBackup uses the media type that is configured for the drives if:</p> <ul style="list-style-type: none"> ■ 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 <p>If the drives are not the same type, NetBackup uses the default media type for the robot.</p> |
| A specific media type | <p>If the robot supports multiple media types and you do not want to use the default media type, select a specific type.</p> <p>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.</p> |

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.

Table 9-8 Default media types for non-API robots (*continued*)

Robot type	Default media type
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

Media settings - volume pool

The volume pool for the new media. The actions depend on whether you use bar code rules to assign media attributes, as follows:

DEFAULT	<p>DEFAULT. If you select DEFAULT and:</p> <ul style="list-style-type: none"> ■ Use bar code rules, the bar code rules determine the volume pool to which new volumes are assigned ■ Do not use bar code 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 bar code 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)
- Uses 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

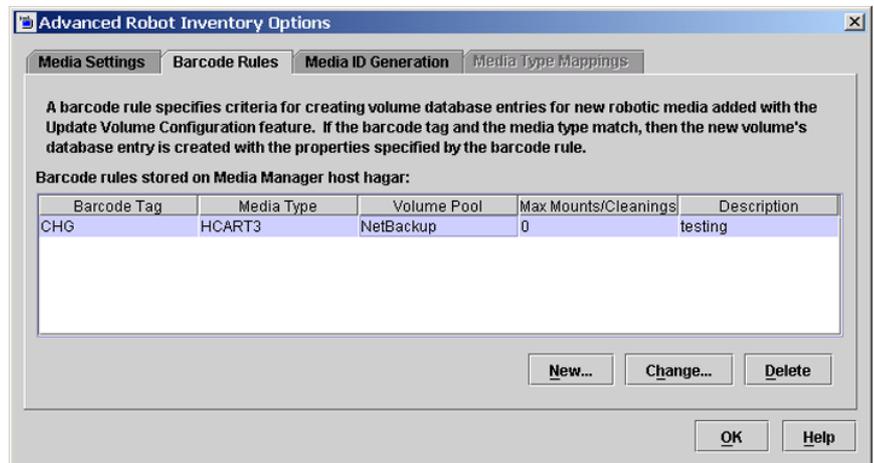
Use the **Barcode Rules** tab of the **Advanced Robot Inventory Options** dialog box to configure rules for assigning attributes to the new volumes that are added to a robot. NetBackup assigns bar codes when you select **Use bar code rules** on the **Media Settings** tab.

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 344.

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 “Bar code rules settings” on page 350.

- 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 bar code tag of a bar code rule. You first must delete the old rule and then add a rule with a new bar code tag.

See “Bar code rules settings” on page 350.

- 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.

Bar code rules settings

See Table 9-10 on page 350. describes the settings you can configure for bar code rules. NetBackup uses these rules to assign bar codes to new media.

Table 9-10 Bar code rule settings

Bar code rule setting	Description
Barcode tag	<p>A unique string of bar code characters that identifies the type of media.</p> <p>For example, use DLT as the bar code tag for a bar code rule if the following is true:</p> <ul style="list-style-type: none"> ■ You use DLT on the bar codes to identify DLT tapes ■ DLT is not used on any other bar codes in the robot <p>Similarly, if you use CLND for DLT cleaning media, use CLND as the bar code tag for the rule for DLT cleaning media.</p> <p>The bar code tag can have from 1 to 16 characters but cannot contain spaces.</p> <p>The following are the special bar code rules that can match special characters in the bar code tags:</p> <ul style="list-style-type: none"> ■ NONE Matches when rules are used and the volume has an unreadable bar code or the robot does not support bar codes. ■ DEFAULT For volumes with bar codes, this tag matches when none of the other bar code tags match. However, the following must be compatible: the media type in the DEFAULT rule and the media type on the Media Settings tab. <p>You cannot change the bar code tag of a bar code rule. Instead, first delete the old rule, then add a rule with a new bar code tag.</p> <p>Use the Media Settings tab to set up the criteria for a robot update.</p> <p>See “Configuring media settings” on page 337.</p>

Table 9-10 Bar code rule settings (*continued*)

Bar code rule setting	Description
Description	A description of the bar code rule. Enter from 1 to 25 characters.
Maximum mounts	<p>The maximum number of mounts (or cleanings) that are allowed for the volume.</p> <p>For data volumes, a value of zero means the volume can be mounted an unlimited number of times.</p> <p>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.</p>
Media type option	<p>The media type to assign to the media.</p> <p>The media type that is specified on the Media Settings tab always overrides the media type of the bar code rule. If you specify a value other than <code>DEFAULT</code> on the Media Settings tab, the bar code rule media type must be the same as the media or be <code>DEFAULT</code>. If not, the bar code rule does not match the media (except for cleaning media).</p> <p>See “Media type when using bar code rules” on page 341.</p> <p>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.</p> <p>See “NetBackup media types” on page 278.</p>
Volume pool	<p>The volume pool for the new media. The actions depend on whether you use bar code rules to assign media attributes.</p> <p>Select from the following:</p> <ul style="list-style-type: none"> ■ DEFAULT <p>If <code>DEFAULT</code> is selected, NetBackup performs the following actions:</p> <ul style="list-style-type: none"> ■ If you use bar code rules, the bar code rules determine the volume pool to which new volumes are assigned. ■ If you do not use bar code rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool. <ul style="list-style-type: none"> ■ A specific volume pool <p>This volume pool setting always overrides any bar code rules.</p>

Configuring media ID generation rules

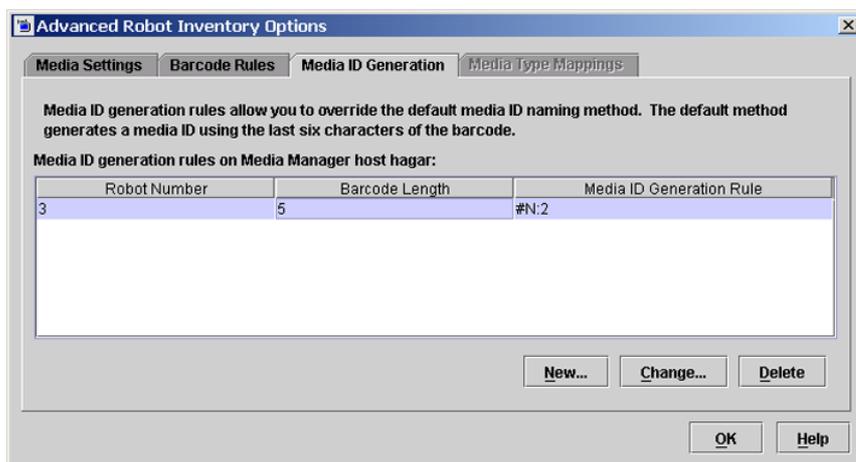
For non-API robots only.

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 348.

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 352.
- 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 bar code 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.

The following list describes the media ID generation rule options:

■ **Bar code length**

The **Barcode length** is the number of characters in the bar code for tapes in the robot.

You cannot change the bar code 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-separate 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 idea, 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.

The table shows some examples of rules and the resulting media IDs.

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 355.

See “Default and allowable media types” on page 356.

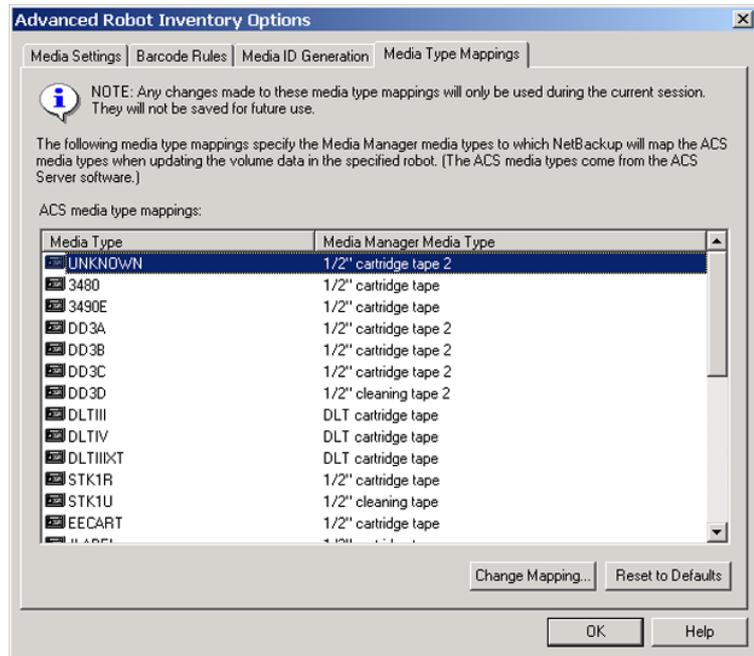
See “NetBackup media types” on page 278.

Note: You can write a bar code 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 bar code rules by media type.

Use the following procedure to change media type mappings.

To configure 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, Volume II*.

Table 9-11 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

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 9-12 shows the default media types and the allowable media types for ACS robots.

Table 9-12 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
DLTIIIXT	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 9-12 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 9-12 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
T10000T1	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000TS	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 9-13 shows the default and allowable media types for TLH robots.

Table 9-13 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

Table 9-13 Default and allowable media types for TLH robots (*continued*)

TLH media type	Default Media Manager media type	Allowable media types through mappings
3592JR	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
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 9-14 shows the default and allowable media types for TLM robots.

Table 9-14 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 `vmphyinv` physical inventory utility

For the following robotic libraries, the **NetBackup Administration Console** reports only the presence of media:

- For the robots without bar code readers
- For the robots that contain media without bar codes

More detailed information is required to perform automated media management. For such robots, use the `vmphyinv` physical inventory utility.

The `vmphyinv` physical inventory utility inventories nonbar coded 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

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.

`vmphyinv` is a command-line utility. Additional information about the syntax of the `vmphyinv` command is available.

See *NetBackup Commands Reference Guide*.

Table 9-15 `vmphyinv` features, requirements, restrictions, and when to use

<code>vmphyinv</code> features	<p>The <code>vmphyinv</code> utility has the following features:</p> <ul style="list-style-type: none">■ 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 <code>vmphyinv</code> 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 <code>-u</code> option or the <code>-n</code> option to specify the drive; the drive must contain media and it must be ready.
<code>vmphyinv</code> requirements and restrictions	<p>The <code>vmphyinv</code> utility has the following requirements and restrictions:</p> <ul style="list-style-type: none">■ 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.

Table 9-15 `vmphyinv` features, requirements, restrictions, and when to use
(continued)

When to use <code>vmphyinv</code>	<p>Use <code>vmphyinv</code> to update the EMM database for NetBackup and Backup Exec media in the following cases:</p> <ul style="list-style-type: none"> ■ You want to inventory a robot that does not have a bar code reader or that contains nonbar coded media. ■ You insert new media into a robotic library and no NetBackup volume records correspond to the media. Use the <code>slot range</code> or <code>list</code> option of <code>vmphyinv</code> 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 <code>slot range</code> or <code>list</code> option of <code>vmphyinv</code> to perform the inventory operation. The <code>vmphyinv</code> 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 <code>vmphyinv</code>.
-----------------------------------	--

See “How `vmphyinv` performs a physical inventory” on page 363.

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 `vmphyinv` list of drives” on page 364.
- Obtains a list of media to mount
 See “About the media `vmphyinv` mounts” on page 364.
- Mounts the media and reads the tape headers
 See “How `vmphyinv` mounts the media and reads the tape header” on page 365.
- Updates the EMM database
 See “How `vmphyinv` updates the EMM database” on page 366.

About the `vmphyinv` 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 `vmphyinv` mounts

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.

```
vmphyinv -rn 4 -pn bear           Mounts the media only in robot 4 and in the  
                                volume pool bear.
```

```
vmphyinv -rn 2 -v moon           Mounts the media in robot 2 and in the  
                                volume group moon.
```

```
vmphyinv -rn 1 -rc1 2 -number    Mounts the media in robot 1 and slot range 2  
3                                to 4.
```

```
vmphyinv -rn 5 -pn NetBackup     Mounts the media in robot 5, slot range 2 to  
-v mars -rc1 2 -number 6        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 -d dlt      Mounts the media in slot range 2 to 10 in
                                             robot 1.

vmphyinv -rn 0 -slot_list 3:4:5 -d 8mm      Mounts the media in slots 3, 4, and 5 in
                                             robot 0.

vmphyinv -rn 2 -slot_range 2 4 -slot_list 5:6:7 -d dlt      Mounts the media in slots 2, 3, 4, 5, 6, and
                                                             7 in robot 2.
```

How `vmphyinv` 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.

See “About media that `vmphyinv` does not recognize” on page 366.

See “How `vmphyinv` processes cleaning media” on page 366.

About media that `vmphyinv` does not recognize

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.

How `vmphyinv` processes 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.

How `vmphyinv` 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.

Table 9-16 vmphyinv criteria and actions

Criteria or action	Description
The <code>vmphyinv update</code> criteria	<p>For valid media types, <code>vmphyinv</code> performs the following actions:</p> <ul style="list-style-type: none"> ■ 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.
How <code>vmphyinv</code> updates NetBackup media	<p>The <code>vmphyinv</code> 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, <code>vmphyinv</code> updates the NetBackup volume record that corresponds to the media ID. If the media ID does not exist, <code>vmphyinv</code> creates a new NetBackup volume record that corresponds to the NetBackup media.</p>

Table 9-16 `vmphyinv` criteria and actions (*continued*)

Criteria or action	Description
<p>How <code>vmphyinv</code> updates Backup Exec media</p>	<p>The <code>vmphyinv</code> utility searches the EMM database. It checks if the media GUID from the tape is present in the <code>ADAMM_GUID</code> field of any record in the EMM database. If the media GUID exists, <code>vmphyinv</code> updates the NetBackup record that contains the GUID. If a media GUID does not exist, <code>vmphyinv</code> creates a new NetBackup record that corresponds to the Backup Exec media. <code>vmphyinv</code> may use an existing NetBackup volume record if the record does not correspond to any media in the tape library.</p> <p>For each NetBackup volume record, <code>vmphyinv</code> does the following:</p> <ul style="list-style-type: none"> ■ In the NetBackup record, updates the <code>ADAMM_GUID</code> 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, <code>vmphyinv</code> creates it. <p>Note: If a <code>MEDIA_ID_PREFIX</code> entry is not specified in the <code>vm.conf</code> file, NetBackup uses BE as the default prefix for Backup Exec media.</p>

Table 9-16 vmphyinv criteria and actions (*continued*)

Criteria or action	Description
vmphyinv error cases	<p>The <code>vmphyinv</code> utility may not be able to update the EMM database correctly in the following cases. These cases are reported as errors.</p> <p>If any of the following cases are encountered, you must intervene to continue:</p> <ul style="list-style-type: none"> ■ 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 different types of volume configuration updates. The examples include only the relevant volume attributes.

See “Volume Configuration Example 1: Removing a volume from a robot” on page 370.

See “Volume Configuration Example 2: Adding existing stand-alone volumes to a robot” on page 371.

See “Volume Configuration Example 3: Moving existing volumes within a robot” on page 373.

See “Volume Configuration Example 4: Adding new volumes to a robot” on page 374.

See “Volume Configuration Example 5: Adding cleaning tapes to a robot” on page 376.

See “Volume Configuration Example 6: Moving existing volumes between robots” on page 377.

See “Volume Configuration Example 7: Adding existing volumes when bar codes are not used” on page 378.

Volume Configuration 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	tl8 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	tl8 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.

Volume Configuration 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 “Volume Configuration Example 6: Moving existing volumes between robots” on page 377.

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 9-17 exist.

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
<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 0 (unlimited)
 allowed

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.

Volume Configuration 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 0 (unlimited)
allowed

Assume that you move the volume to empty slot 10, specify the following on the **Media Settings** tab, then run the update.

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 0 (unlimited)
allowed

The updated volume attributes show the new slot number, but all other information is unchanged.

Volume Configuration 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 9-18 contains the example bar code rules.

Table 9-18 Example bar code rules

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
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 EXB220
Maximum mounts allowed 0 (unlimited)

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 the 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 are 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.

Volume Configuration Example 5: Adding cleaning tapes to a robot

A special case exists when you add cleaning tapes. For example, assume that 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 9-19 contains the example bar code rules:

Table 9-19 Example bar code rules

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
CLN	DLT_CLN	None	30	dlt cleaning

Table 9-19 Example bar code rules (continued)

Bar code tag	Media type	Volume pool	Max mounts/ cleanings	Description
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).

Volume Configuration 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 "Volume Configuration Example 7: Adding existing volumes when bar codes are not used" on page 378.

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 action updates the volume attributes to show the volume as stand-alone.
- Perform an Update volume configuration on robot 2. This action updates the configuration to show the volume in robot 2.

Volume Configuration 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 that 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 vmphyinv physical inventory utility” on page 361.

Configuring disk storage

This chapter includes the following topics:

- Configuring BasicDisk storage
- Configuring NearStore storage
- About configuring disk pool storage
- About SharedDisk support in NetBackup 7.0 and later

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. The directory is specified when the storage unit is configured.

See “About storage units” on page 386.

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 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*.

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 or later 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 or later, 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 and later 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 `bpxpdate` commands:

```
bpxpdate -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*.

```
bpxpdate -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 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. 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 386.

- **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 435.

- **Storage Lifecycle Policies**

Storage lifecycle policies let you apply the same behavior to all the backup images in the lifecycle.
See “About storage lifecycle policies” on page 443.

Using the Storage utility

To use the storage utility

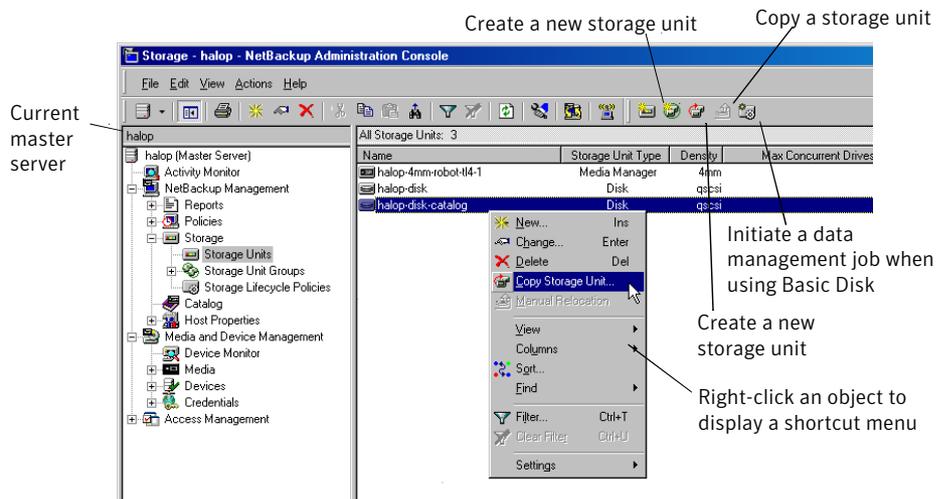
- ◆ In the **NetBackup Administration Console**, expand **Storage > Storage Units**, **Storage Unit Groups**, or **Storage Lifecycle Policies**.

The storage destinations that were created for the selected server are displayed in the right pane.

The storage configuration can be displayed for other master servers.

See “Accessing remote servers” on page 835.

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:

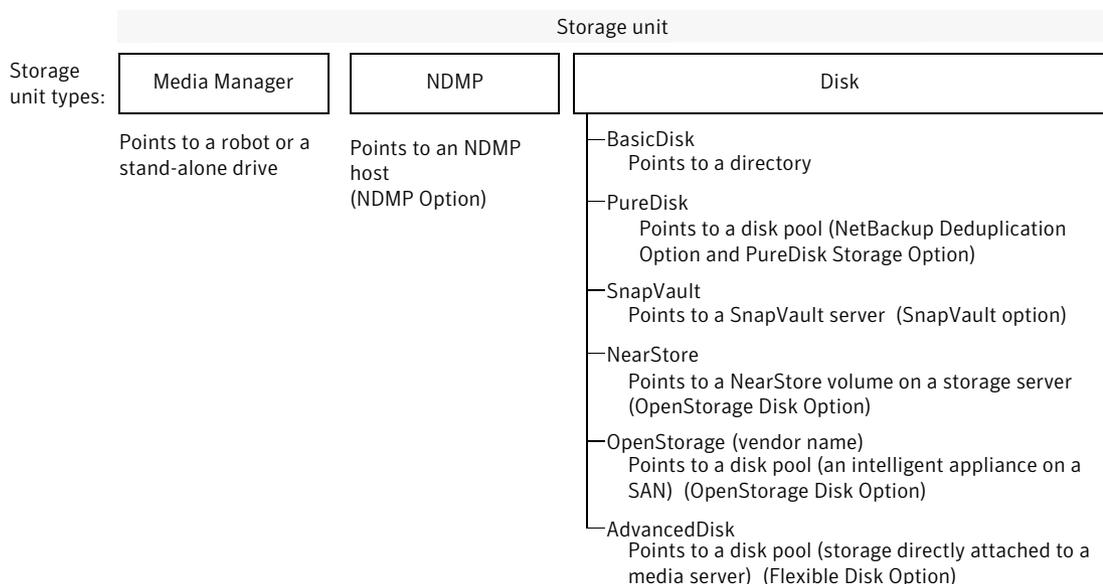
- Name the storage unit. A configured storage unit indicates to NetBackup the existence of physical storage.

- Choose the storage unit type: Media Manager, disk, or NDMP. See Figure 11-2 on page 388.
- Select a media server. The selection indicates that the media server has permission to write to the storage unit. You can select multiple servers if you like.
- 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 382.
SnapVault storage	The destination is a SnapVault server.
NDMP storage	The destination is an NDMP host. The NDMP protocol is used to perform backups and recoveries.

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 right pane, click **Configure Storage Devices** and follow the wizard instructions.

For help while running the wizard, click the **Help** option 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 400.
- 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 In the right pane, select a storage unit.
- 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 827.
- 5 Complete the fields in the **Copy Storage Unit** dialog box.
See “About storage unit settings” on page 400.

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 In the right pane, double-click the storage unit you want to change.
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 400.

Deleting storage units

To delete a storage unit from a NetBackup configuration means to delete the label that NetBackup associates with the physical storage.

Deleting 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 action removes the image from the NetBackup catalog.

See “Expiring backup images” on page 750.

- 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 “About importing backup images” on page 750.

NetBackup automatically deletes any image fragments from a disk storage unit or a disk pool. This deletion 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 **Storage > Storage Units**.
- 3 In the right 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 affecting the disk pool.

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 400.

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.

Take the following items into consideration when adding a Media Manager storage unit:

- Where to add the storage unit depends on which version of NetBackup you are using.

- 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.
 - 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 storage unit setting" on page 403.
 - 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 (`nbemm`) 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.

Table 11-1 describes the different disk types that NetBackup can use as disk media.

Table 11-1 Disk media descriptions

Type of disk storage unit	Description
BasicDisk	<p>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.</p> <p>Notes about the BasicDisk type storage unit:</p> <ul style="list-style-type: none"> ■ 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	<p>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.</p> <p>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.</p> <p>For AdvancedDisk, the NetBackup media servers function as both data movers and storage servers.</p> <p>See the <i>NetBackup Shared Storage Guide</i>.</p>
NearStore	<p>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.</p> <p>For NearStore, the NetBackup media servers function as the data movers. The NearStore host is the storage server.</p> <p>Note: NearStore storage units cannot be used as part of a storage unit group or used in a storage lifecycle policy.</p> <p>Information about configuring NearStore storage units is described in the <i>NetBackup Administrator's Guide, Volume II</i>.</p>

Table 11-1 Disk media descriptions (*continued*)

Type of disk storage unit	Description
OpenStorage	<p>An OpenStorage disk type storage unit is used for disk storage on an intelligent disk appliance or on the Nirvanix Storage Directory Network. The actual name of the disk type depends on the vendor. An OpenStorage selection is available only when the OpenStorage Disk Option is licensed.</p> <p>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.</p> <p>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.</p> <p>See the <i>NetBackup Shared Storage Guide</i>.</p>
PureDisk	<p>A PureDisk disk type storage unit is used to store deduplicated data for the following NetBackup options:</p> <ul style="list-style-type: none"> ■ Media server deduplication pool. NetBackup deduplication must be configured. See the <i>NetBackup Deduplication Guide</i>. ■ PureDisk deduplication pool (PureDisk 6.6 and later). NetBackup deduplication must be configured. See the <i>NetBackup Deduplication Guide</i>. ■ PureDisk Deduplication Option (PDDO) storage pool (PureDisk 6.5 and later). PureDisk Deduplication Option (PDDO) must be configured. See the <i>NetBackup PureDisk Deduplication Option Guide</i>. <p>Note: PDDO storage units cannot be used as part of a storage unit group.</p> <p>PureDisk appears as a selection when the NetBackup Deduplication Option or the PureDisk Storage Option is licensed.</p>
SnapVault	<p>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.</p> <p>SnapVault storage units cannot be used in a storage unit group or as part of a staging operation.</p> <p>For SnapVault, the NetBackup media servers function as the data movers. The SnapVault host is the storage server.</p>

Not all settings are available on each disk storage unit type.

See “About storage unit settings” on page 400.

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 BasicDisk 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
 - For Nirvanix cloud storage, nirvanix.com.
- **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.

Configuring credentials for CIFS and disk storage units

For Common Internet File System (CIFS) storage and AdvancedDisk and BasicDisk storage units, two NetBackup services on Windows computers require matching credentials. The following NetBackup services on Windows media servers to which the CIFS storage is attached must use the same credentials:

- **NetBackup Client Service**
The NetBackup Client Service is either `bpcd.exe` or `bpinetd.exe`, depending on NetBackup release level. Regardless of the binary file name, the service requires the credentials.

CIFS storage is supported for AdvancedDisk on NetBackup 7.0.1 and later.

- **NetBackup Remote Manager and Monitor Service**

The NetBackup Remote Manager and Monitor Service binary file name is `nbrmms.exe`.

The credentials must be valid Windows credentials that allow read and write access to the storage. Configure the credentials on the media server to which the CIFS storage is attached.

If credentials are not configured, NetBackup marks all CIFS AdvancedDisk and BasicDisk storage units that use the UNC naming convention as DOWN.

To configure credentials for CIFS with AdvancedDisk and BasicDisk

1 In Windows, in the **Microsoft Management Console**, open **Services**.

How you open **Services** depends on the Windows version.

2 Open the **Properties** dialog box for the service.

The **General** tab should be displayed.

3 Click **Stop** to stop the service.

4 Select the **Log On** tab.

5 Select **This account** and then enter the credentials.

6 Select the **General** tab.

7 Click to start the service.

8 Repeats the steps 2 to 7 for each service.

Disk storage units in storage lifecycle policies

Figure 11-4 is an example of how storage lifecycle policies can interact with volumes in a disk pool that a storage unit references.

Two policies are created as follows:

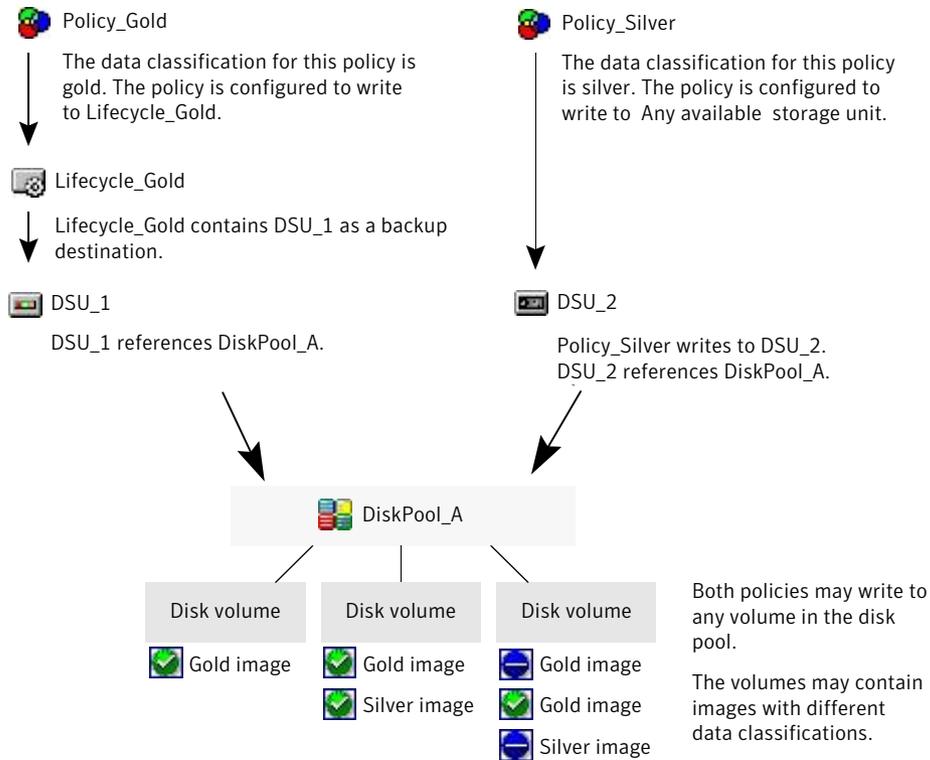
- Policy_gold has a gold classification. It is configured to use Lifecycle_Gold, which has a gold data classification.
- Policy_silver has a silver classification. It is configured to use Any Available storage unit. That means it can use any available storage unit or any lifecycle that has a silver classification.

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 “About storage lifecycle policies” on page 443.

Maintaining available disk space on disk storage units

Disk storage units can be managed so that they do not become entirely full and cause backups to fail.

Create space for more images on a disk storage unit in the following ways:

- 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 storage unit setting” on page 402.

Maintain space on basic disk staging storage units in the following ways:

- 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.

- Run the `nb_updatedssu` script.

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
```

- Determine the potential free space.
See “Finding potential free space on a BasicDisk disk staging storage unit” on page 428.
- Monitor disk space by enabling the **Check the capacity of disk storage units** host property.
This General Server host 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 “General Server properties” on page 128.

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 400.

Figure 11-5 NDMP storage unit settings

New Storage Unit

Storage unit name:

Storage unit type: NDMP On demand only

Disk type: BasicDisk

Properties

Storage device:

Robot type:	Static
Density:	Static
Robot number:	Static

NDMP Host:

Media Server: <Any Available>

Maximum concurrent write drives: 0 Reduce fragment size to: 1048576 Megabytes

OK Cancel Help

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 absolute pathname to volume setting for storage units

Absolute pathname to directory or **Absolute pathname to volume** is available for any storage unit that is not based on disk pools.

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.

Use platform-specific file path separators (/ and \) and colon (:) within a drive specification.

The **Properties** button displays properties for the directory or volume.

See “Properties option in the Change Storage Units dialog box” on page 411.

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 the disk pools that are composed of the disk volumes that are dedicated file systems for disk backup.

See “NetBackup naming conventions” on page 827.

See “Low water mark storage unit setting” on page 403.

Directory can exist on the root file system or system disk setting for storage units

When checked, this setting allows the user to specify a directory on the root file system (UNIX) or on a system drive (Windows) in the **Absolute pathname to directory** field.

When this setting is checked, the directory is created automatically. If a storage unit is configured on C drive and this option is not checked, backups fail with error code 12.

Note: With this setting checked, the system drive can fill up.

A job fails under the following conditions:

- If the setting is not checked, and if the directory already exists on a system drive.
- If the setting is not checked, and the requested directory is to be created on a system drive.

See “Absolute pathname to directory or absolute pathname to volume setting for storage units” on page 400.

Density storage unit setting

The **Storage device** selection determines the media **Density**. This setting appears for Media Manager and NDMP storage units only.

Disk pool storage unit setting

Denali content:

The **Disk pool** storage unit setting specifies whether the storage unit uses a disk pool that is configured for snapshots.

The following table describes 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 storage unit setting

The **Disk type** storage unit setting identifies the type of 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 382.
- SnapVault (NetBackup Snapshot Client option needed).
For information on SnapVault storage units, see the *NetBackup Snapshot Client Administrator’s Guide*.

Enable block sharing storage unit setting

The **Enable block sharing** storage unit setting sets the data blocks that have not changed from one backup to the next be shared. Sharing data blocks can significantly save disk space in the storage unit.

Enable multiplexing storage unit setting

The **Enable multiplexing** storage unit setting allows multiple backups to multiplex onto a single drive in a storage unit.

High water mark storage unit setting

The **High water mark** storage unit setting applies to **BasicDisk** storage units only.

A **High water mark** setting also applies to disk pools; that setting is described in the online Help or a separate guide.

See the *NetBackup Deduplication Guide*.

See the *NetBackup Shared Storage Guide*.

The **High water mark** setting (default 98%) is a threshold that triggers the following actions:

- When an individual disk volume of the underlying storage reaches the **High water mark**, NetBackup considers the volume full. NetBackup chooses a different volume in the underlying storage to write backup images to.
- When all volumes in the underlying storage reach the **High water mark**, the **BasicDisk** storage is considered full. NetBackup fails any backup jobs that are assigned to a storage unit in which the underlying storage is full. NetBackup also does not assign new jobs to a **BasicDisk** storage unit in which the underlying storage is full.
- NetBackup begins image cleanup when a volume reaches the **High water mark**; image cleanup expires the images that are no longer valid. NetBackup again assigns jobs to the storage unit when image cleanup reduces any disk volume's capacity to less than the **High water mark**.

If the storage unit is in a capacity-managed storage lifecycle policy, other factors affect image cleanup.

See “Staged capacity managed retention type for storage destinations” on page 455.

See “Maximum concurrent jobs storage unit setting” on page 404.

Information about the disk pool **High water mark** setting is available.

See the *NetBackup Deduplication Guide*.

See the *NetBackup Shared Storage Guide*.

Low water mark storage unit setting

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.

Once the **High Water Mark** is reached, space is created on the disk storage unit until the **Low Water Mark** is met. The default setting is 80%.

See “Staged capacity managed retention type for storage destinations” on page 455.

The **Low water mark** setting cannot be greater than the **High water mark** setting.

For the disk storage units that reference disk pools, the **Low water mark** applies to the disk pool.

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 more information, see the following guides:

- See the *NetBackup Deduplication Guide*.
- See the *NetBackup Shared Storage Guide*.

Maximum concurrent write drives storage unit setting

The **Maximum concurrent write drives** storage unit 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.

When selecting the number of **Maximum concurrent write drives**, use the following guidelines:

- 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 storage unit setting

The **Maximum concurrent jobs** storage unit setting specifies the maximum number of jobs that NetBackup can send to a disk storage unit at one time. The default setting is one job. The job count can range from 0 to 256.

Note: To specify a **Maximum concurrent jobs** setting of 0 disables the 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: 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.

The **Maximum concurrent jobs** setting uses and dependencies are as follows:

- 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.
- 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.

See "Impact when two disk storage units reference one disk pool" on page 405.

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.

In the example, DSU_1 is configured as follows:

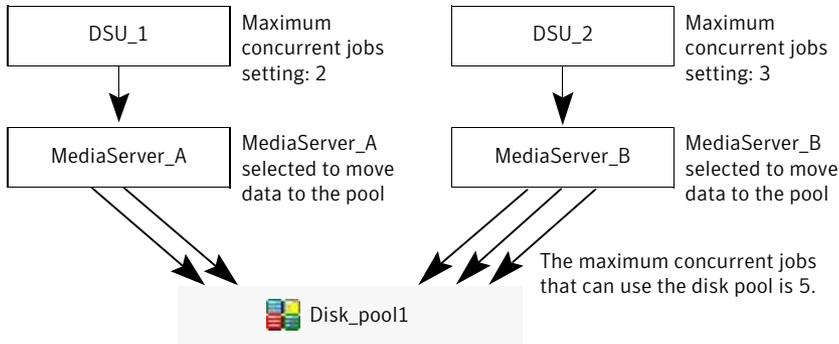
- To use MediaServer_A
- To have a **Maximum concurrent jobs** setting of two
- To reference Disk_pool1

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.

See “About storage unit settings” on page 400.

Maximum streams per drive storage unit setting

The **Maximum streams per drive** storage unit 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 573.

See “Media multiplexing (schedule attribute)” on page 572.

Media server storage unit setting

The **Media server** storage unit 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 write data to and read data from the disk pool.
- 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, depending on the type of storage.

Table 11-2 Media server setting details

Storage unit type	Considerations
BasicDisk	To configure a disk storage unit, select a single media server.
AdvancedDisk	<p>The Media server setting specifies the NetBackup media servers that can write data to and read data from the disk pool.</p> <p>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.</p> <p>NetBackup selects a media server when the policy runs.</p>
NDMP	<p>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.</p> <p>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.</p>

Table 11-2 Media server setting details (*continued*)

Storage unit type	Considerations
OpenStorage	<p>The Media server setting specifies the NetBackup media servers that can move data to or from the storage server.</p> <p>To allow any media server in the media server list to move data to the storage server, check Use Any Available Media Server.</p> <p>To restrict the media servers that can move data to the storage server, check Only Use The Following Media Servers. Then select the media servers that are allowed to move the data.</p> <p>Any media server in the list can receive data from the storage server; it does not have to be selected. A media server receives data for restore jobs and for storage monitoring purposes.</p> <p>Each media server that moves the data must meet the following requirements:</p> <ul style="list-style-type: none"> ■ The vendor’s software plug-in is installed. ■ The login credentials to the storage server are configured. <p>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.</p> <p>Note: Run the <code>tpconfig</code> command line utility directly on the media server to configure and verify credentials.</p> <p>NetBackup selects a media server when the policy runs.</p>
SharedDisk	<p>See “About SharedDisk support in NetBackup 7.0 and later” on page 382.</p>
PureDisk (Media Server Deduplication Pool and PureDisk Deduplication Pool)	<p>To allow any media server in the list to deduplicate data, select Use Any Available Media Server.</p> <p>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.</p> <p>Each media server must be configured as a deduplication media server.</p> <p>See the <i>NetBackup Deduplication Guide</i>.</p>

Table 11-2 Media server setting details (*continued*)

Storage unit type	Considerations
PureDisk (PureDisk Deduplication Option storage pool)	<p>To allow any media server in the list to access the storage (default), select Use Any Available Media Server.</p> <p>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.</p> <p>NetBackup selects a media server when the policy runs.</p> <p>Each media server that accesses the storage must meet the following requirements:</p> <ul style="list-style-type: none"> ■ The PureDisk agent is installed. ■ The logon credentials to the PureDisk server are configured on the media server. <p>See the <i>NetBackup PureDisk Remote Office Edition Administrator's Guide</i> for the media server requirements.</p>

See “Use any available media server storage unit setting” on page 416.

See “Only use the following media servers storage unit setting” on page 410.

NDMP host storage unit setting

The **NDMP host** storage unit 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 storage unit setting

The **On demand only** storage unit 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. Uncheck **On demand only** 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 storage unit setting

The **Only use the following media servers** storage unit setting restricts the media servers earmarked for storage. Check this setting and select the media servers that you want to use.

The following table describes the media server functionality for each type of storage.

Table 11-3 Media server functionality

Media server type	Functionality
AdvancedDisk storage media server	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.
OpenStorage media server	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. Each media server that accesses the storage must meet the following requirements: <ul style="list-style-type: none"> ■ The vendor’s software plug-in is installed. ■ The login credentials to the storage server are configured.
PureDisk media server (media server deduplication pool and PureDisk deduplication pool)	The media servers function as deduplication servers. NetBackup deduplication must be configured.
PureDisk media server (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.

See “Use any available media server storage unit setting” on page 416.

See “Only use the following media servers storage unit setting” on page 410.

Properties option in the Change Storage Units dialog box

Click **Properties** to display information about the volume or the disk pool, as follows:

Table 11-4 Storage Units Properties

Property	Description
Available storage or Available	<p>This value reflects the space that remains for storage on a disk storage unit. The following equation determines the available space:</p> $\text{Available space} = \text{free space} + \text{potential free space} - \text{committed space}$ <p>The <code>df</code> 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 <code>nbdevquery</code> command:</p> <pre>nbdevquery -listdv -stype server_type -dp disk_pool</pre> <p>The available space that the <code>df</code> command lists does not include the space that the operating system reserves. Since NetBackup runs as <code>root</code>, the <code>nbdevquery</code> command includes the reserved space in the available space equation.</p>
Capacity	<p>The Capacity value reflects the total amount of space that the disk storage unit or pool contains, both used and unused.</p>
Disk pool comments	<p>Comments that are associated with the disk pool.</p>

Table 11-4 Storage Units Properties (*continued*)

Property	Description
High water mark	<p>The high water mark for the disk pool applies to both the individual disk volumes in the pool and the disk pool:</p> <ul style="list-style-type: none"> ■ Individual volumes When a disk volume reaches the high water mark, new jobs are not assigned to the volume. This behavior happens for all disk types except BasicDisk staging storage units. The high water mark event triggers the deletion of images that have been relocated, attempting 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. The default setting is 99%.
Low water mark	<p>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.</p> <p>Note: The Low water mark setting has no effect unless backups are written through a storage lifecycle policy, using the capacity-managed retention type.</p>
Name	The name of the disk pool.
Number of volumes	The number of disk volumes in the disk pool.
% full	<p>The percentage of storage that is currently in use on the volume.</p> <p>The <code>df</code> 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.)</p>
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 storage unit setting

The **Reduce fragment size** storage unit setting specifies the largest fragment size that NetBackup can create to store backups.

Table 11-5 Maximum fragment size

Storage unit type	Fragment size
Media Manager storage units	<p>The default maximum fragment size for a Media Manager storage unit is 1000 GB. To specify a maximum fragment size other than the default, check Reduce fragment size. Then enter a value from 50 megabytes to 1,048,575 megabytes.</p> <p>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.</p>

Table 11-5 Maximum fragment size (*continued*)

Storage unit type	Fragment size
Disk storage units	<p>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.</p> <p>For media server deduplication pools and PureDisk deduplication pools, you can enter a value from 20 megabytes to 51200 megabytes.</p> <p>Backups to disk are usually fragmented to ensure that the backup does not exceed the maximum size that the file system allows.</p> <p>The Reduce fragment size setting is intended primarily for storing large backup images on a disk type storage unit.</p> <p>Note: OpenStorage vendors may have special requirements for the maximum fragment size. Consult the vendor's documentation for guidance.</p> <p>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.</p>

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.)

Robot number storage unit setting

The **Robot number** storage unit setting indicates the number of robots the storage unit contains. The **Storage device** selection determines the **Robot number**. It is the same robot number used in the Media Manager configuration.

Robot type storage unit setting

The **Robot type** storage unit setting 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.

See “Storage device setting for storage units” on page 415.

Staging schedule option in Change Storage Units dialog

Click the **Staging Schedule** option 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 “Disk Staging Schedule dialog box” on page 430.

See “Enable temporary staging area storage unit setting” on page 416.

See “About basic disk staging” on page 421.

See “About staging backups” on page 419.

Storage device setting for storage units

The **Storage device** list contains all possible storage devices available. Storage units can be created for the listed devices only.

The **Storage device** selection determines the media **Density**. This setting appears for Media Manager and NDMP storage units only.

Storage unit name setting

The **Storage unit name** setting defines a unique name for the new storage unit. The name can 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 when changing settings for a storage unit.

See “NetBackup naming conventions” on page 827.

Storage unit type setting

The **Storage unit type** setting specifies the type of storage that this storage unit uses, as follows:

Disk	See “Disk storage unit considerations” on page 392.
Media Manager	See “Media Manager storage unit considerations” on page 390.
NDMP	See “NDMP storage unit considerations” on page 398.

Enable temporary staging area storage unit setting

The **Enable temporary staging area** storage unit setting allows this storage unit to be used as a temporary staging area. Check **Enable Temporary Staging Area** and then configure the staging schedule.

See “Staging schedule option in Change Storage Units dialog” on page 415.

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. Not all columns display by default.

See “About basic disk staging” on page 421.

See “Staging schedule option in Change Storage Units dialog” on page 415.

Transfer throttle storage unit setting

The **Transfer throttle** setting appears for SnapVault storage units only.

This setting allows the user 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. The range is 0 to 9999999.

A value greater than 0 indicates a transfer speed for SnapVault in kilobytes per second. For example, a value of one sets a transfer speed limit for SnapVault of 1 kilobyte per second, which is a very slow transfer rate.

Use any available media server storage unit setting

When checked, the **Use any available media server** storage unit setting allows any media server in the media server list to access the storage (default).

The following table describes the media server functionality for each type of storage.

Table 11-6 Media server functionality

Storage unit type	Functionality
AdvancedDisk storage media server	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.
OpenStorage media server	<p>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.</p> <p>The following is required on each media server that accesses the storage:</p> <ul style="list-style-type: none"> ■ The vendor's software plug-in is installed. ■ The login credentials to the storage server are configured.
PureDisk media server (media server deduplication pool and PureDisk deduplication pool)	<p>The media servers function as deduplication servers.</p> <p>NetBackup deduplication must be configured.</p>
PureDisk media server (PureDisk Deduplication Option storage pool)	<p>The NetBackup media servers function as the data movers. The PureDisk Linux servers function as the storage servers.</p> <p>PureDisk Deduplication Option (PDDO) must be configured.</p>

Staging backups

This chapter includes the following topics:

- About staging backups
- About the two staging methods
- About basic disk staging
- Creating a basic disk staging storage unit
- Configuring multiple copies in a relocation schedule
- Disk staging storage unit size and capacity
- Finding potential free space on a BasicDisk disk staging storage unit
- Disk Staging Schedule dialog box
- Basic disk staging limitations
- Initiating a relocation schedule manually

About staging backups

In the staged backups process, NetBackup writes a backup to a storage unit and then duplicates it 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:

- Allows for faster restores from disk.
- Allows the backups to run when tape drives are scarce.

- Allows the data to be streamed to tape without image multiplexing.

About the two staging methods

NetBackup offers the following methods for staging backups.

Table 12-1 Methods for staging backups

Staging method	Description
Basic disk staging	<p>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.</p> <p>See “About basic disk staging” on page 421.</p> <p>The following storage unit types are available for basic disk staging: BasicDisk, NearStore, and tape.</p>
Staging using the Storage Lifecycle Policies utility	<p>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.</p> <p>No BasicDisk, SnapVault, or disk staging storage units can be used as destinations in a lifecycle.</p> <p>See “About storage lifecycle policies” on page 443.</p>

The staging method 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.

About basic disk staging

Basic disk staging is conducted in the following stages.

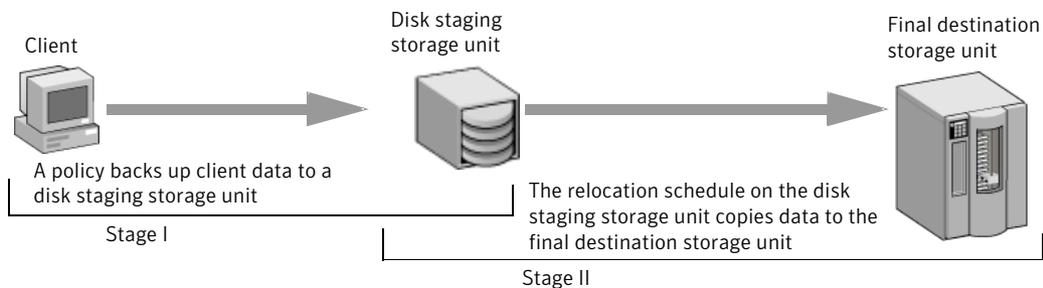
Table 12-2 Basic disk staging

Stage	Description
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 in 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.

Note: The basic disk staging method does not support 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 “Take checkpoints every __ minutes (policy attribute)” on page 521.

Creating a basic disk staging storage unit

Use the following processes to create a basic disk staging storage unit.

To add a new storage unit

- 1 In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Units**.
- 2 Click **Actions > New > New Storage Unit**.
- 3 Configure the storage unit.
See “Creating a basic disk staging storage unit” on page 422.
- 4 Define the disk staging schedule.
See “To define the disk staging schedule” on page 423.
- 5 Click **OK** to add the storage unit.

To configure a new storage unit

- 1 In the **New Storage Unit** dialog box, name the storage unit.
See “Storage unit name setting” on page 415.
- 2 Select Disk as the **Storage unit type**.
See “Storage unit type setting” on page 415.
- 3 Select the **Disk type** of disk storage unit that is to be a disk staging storage unit: BasicDisk or NearStore.
- 4 Select a media server.
See “Media server storage unit setting” on page 406.

- 5 Enter the absolute path to the directory to be used for storage.
 See “Absolute pathname to directory or absolute pathname to volume setting for storage units” on page 400.
- 6 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 setting for storage units” on page 400.
- 7 Enter the maximum concurrent jobs that are allowed to write to this storage unit at one time.
 See “Maximum concurrent jobs storage unit setting” on page 404.
- 8 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.

Note: The **Low water mark** setting does not apply to disk staging storage units.

- 9 Check the **Enable temporary staging area** option. Once the option is enabled, the **Staging Schedule** option is enabled.

The Disk Staging Schedule is similar to the scheduling dialog box used to configure policies. The differences appear on the **Attributes** tab.

To define the disk staging schedule

- 1 Click **Staging Schedule**.
- 2 In the **Disk Staging Schedule** dialog box, select the priority that the relocation jobs that are started from this schedule have compared to other types of jobs.
 The schedule name defaults to the storage unit name.
 See “Disk Staging Schedule dialog box” on page 430.

- 3 Select whether to create multiple copies. When the **Multiple copies** attribute is checked, NetBackup can create up to four copies of a backup simultaneously. See “Multiple copies (schedule attribute)” on page 562.
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 “Global Attributes properties” on page 131.
- 4 Select a storage unit to contain the images from this storage unit upon relocation.
- 5 Select a volume pool to contain the images from this storage unit upon relocation.
- 6 Select a media owner to own the images from this storage unit upon relocation.
- 7 Select whether to use an alternate server for the images from this storage unit upon relocation.
- 8 Click **OK** to accept the disk staging schedule.

Configuring multiple copies in a relocation schedule

To configure a relocation schedule for basic disk staging to create multiple copies, use the following procedure.

To configure a relocation schedule for basic disk staging to create multiple copies

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Storage > Storage Units**.
- 2 Perform one of the following actions:

To change an existing basic disk storage unit	<ul style="list-style-type: none"> ■ Select the storage unit to change. ■ On the Edit menu, click Change.
To create a new basic disk storage unit	<ul style="list-style-type: none"> ■ On the Actions menu, click New > New Storage Unit. ■ Name the storage unit. ■ From the Storage unit type list, select Disk. ■ Check Enable Temporary Staging Area. ■ Configure the other storage unit settings as necessary. ■ See “Creating a basic disk staging storage unit” on page 422.

See “About staging backups” on page 419.

- 3 Click **Staging Schedule**.
- 4 In the dialog box that appears, on the **Attributes** tab, specify a priority in the field **Priority of relocation jobs started from this schedule** (0 to 99999).
- 5 Select a schedule type and schedule when the policy should run.
- 6 Check **Use alternate read server**, and select an alternate server from the list. The alternate server can read a backup image originally written by a different media server.
- 7 Select **Multiple copies** and click **Configure**.

If **Multiple copies** is grayed out, make sure that the **Maximum backup copies** host property is set to at least 3. This host property is in the **Global Attributes** properties.

See “Global Attributes properties” on page 131.

- 8 In the **Copies** field, specify the number of copies to create simultaneously. The number must be between 1 and 4.

The maximum is four, or the number of copies that the **Maximum backup copies** setting specifies, whichever is fewer.

The **Maximum backup copies** 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 **Configure Multiple Copies** dialog box, set the **Maximum backup copies** property to five or more.

Copy 1 is the primary copy. If **Copy 1** fails, the first successful copy is the primary 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 “Configure Multiple Copies dialog box” on page 564.

See “About configuring for multiple copies” on page 563.

- 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 one of the following from the **If this copy fails** list:

- | | |
|------------------------|---|
| continue | Continues making the remaining copies.
Note: Note: If Take checkpoints every __ minutes is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.
See “Take checkpoints every __ minutes (policy attribute)” on page 521. |
| fail all copies | Fails the entire job. |

12 For tape media, specify who should own the media onto which NetBackup writes the images:

- | | |
|----------------|---|
| Any | NetBackup selects 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 that are configured in the NetBackup environment appear in the list. |

These settings do not affect images residing on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images.

13 Click **OK**.

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.

The following table gives more information about determining disk size for a basic disk staging storage unit.

Table 12-3 Size considerations for a basic disk staging storage unit

Disk size	Description
Minimum disk size	<p>The minimum disk size is the smallest size that is required for the successful operation of the disk staging logic.</p> <p>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.)</p> <p>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.</p> <p>To determine the minimum size for the storage unit in this example, use the following formula:</p> <p>Minimum size = Max data per cycle × (1 cycle + 1 cycle for safety)</p> <p>For example: 1.4 GB = 700 MB × (1+1)</p>

Table 12-3 Size considerations for a basic disk staging storage unit (*continued*)

Disk size	Description
Average disk size	<p>The average disk size represents a good compromise between the minimum and the maximum sizes.</p> <p>In this example, the average nightly backup is 400 MB and the NetBackup administrator wants to keep the images for one week.</p> <p>To determine the average size for the storage unit in this example, use the following formula:</p> <p>Average size = Average data per cycle × (number of cycles to keep data + 1 cycle for safety)</p> <p>2.8 GB = 400 MB × (6 + 1)</p>
Maximum disk size	<p>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.</p> <p>To determine the maximum size for the storage unit in this example, use the following formula:</p> <p>Maximum size = Max data per cycle × (# of cycles to keep data + 1 cycle for safety)</p> <p>For example: 4.9 GB = 700 MB × (6 + 1)</p>

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 as follows:

- Run `bpstulist -label` to find the disk pool name.
- Note that the name of the storage unit and disk pools are case sensitive. 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:

<code>-stype</code> server_type	Specifies the vendor-specific string that identifies the storage server type. For a BasicDisk storage unit, enter <code>BasicDisk</code> .
<code>-dp</code>	Specifies the disk pool name. For a basic disk type, the disk pool name is the name of the BasicDisk storage unit.

So the complete command might look like the following.

```
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>
```

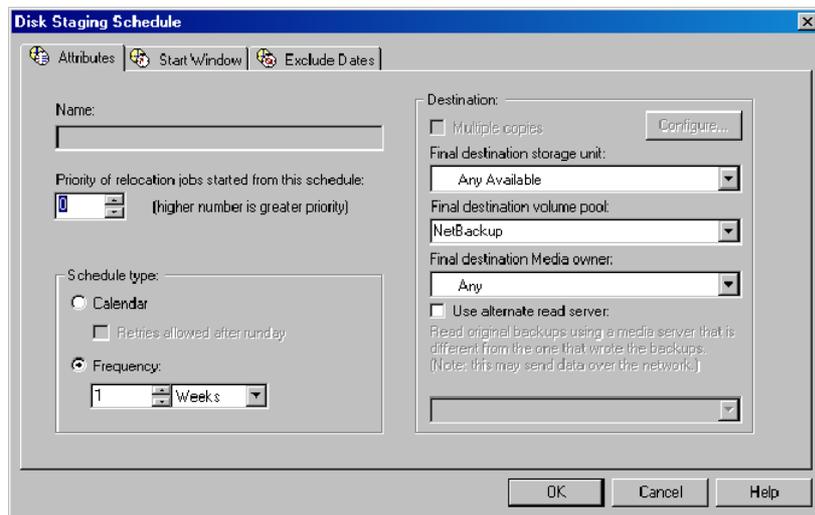
Disk Staging Schedule dialog box

Click **Staging Schedule** 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 table.

Table 12-4 Attributes tab settings

Attributes tab setting	Description
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).

Table 12-4 Attributes tab settings (continued)

Attributes tab setting	Description
Final destination storage unit	<p>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.</p> <p>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.</p> <p>See “Low water mark storage unit setting” on page 403.</p> <p>See “Maximum concurrent write drives storage unit setting” on page 403.</p> <p>See “About staging backups” on page 419.</p>
Final destination volume pool	<p>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.</p> <p>See “About staging backups” on page 419.</p> <p>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.</p>
Final destination media owner	<p>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.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"> ■ 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.

Table 12-4 Attributes tab settings (continued)

Attributes tab setting	Description
Use alternate read server	<p>The Use alternate read server attribute applies to NetBackup Enterprise Server only.</p> <p>An alternate read server is a server allowed to read a backup image originally written by a different media server.</p> <p>The path to the disk or directory must be identical for each media server that is to access the disk.</p> <p>If the backup image is on tape, the media servers must share the same tape library or the operator must find the media.</p> <p>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 <code>bpmedia -oldserver -newserver</code> or assign a failover media server.</p> <p>To avoid sending data over the network during duplication, specify an alternate read server that meets the following conditions:</p> <ul style="list-style-type: none"> ■ 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. <p>If the final destination storage unit is not connected to the alternate read server, data is sent over the network.</p>

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 “Take checkpoints every __ minutes (policy attribute)” on page 521.

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** In the right pane, select a basic disk staging storage unit.
- 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 on both storage units until the disk staging (Stage I) storage unit becomes full and the oldest images are deleted.

See “Maintaining available disk space on disk storage units” on page 397.

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
- About 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 827.



The storage unit group name is case-sensitive .

- 4 Add or remove storage units to and from the group:
 - To add storage units to the group, select the storage units from the **Storage units not in the group** list and click **Add**.
 - To remove storage units from the group, select the storage units from the **Storage units in group** list and click **Remove**.
 - To change the priority of a storage unit, select the storage unit and click **Move Up** or **Move Down**. The units at the top of the list have the highest priority in the group.

Note: OpenStorage, SnapVault, NearStore, and PureDisk storage units cannot be included in storage unit groups.

- 5 Choose how storage units are selected within the group:
 - **Prioritized.** Choose the first storage unit in the list that is not busy, down, or out of media.
 - **Failover.** Choose the first storage unit in the list that is not down or out of media.
 - **Round Robin.** Choose the least recently selected storage unit in the list.
 - **Media server load balancing.** Choose a storage unit based on a capacity-managed approach.
Symantec recommends the **Media server load balancing** criteria for disk staging storage units within a storage unit group.
See “Media server load balancing” on page 439.

See “Storage unit selection criteria within a group” on page 438.

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 442.
- 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 In the right pane, from the list of storage unit groups, select the storage unit group you want to delete. 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.

The only difference between the selection criteria options is the order in which the storage units are selected.

Choose from one of the following selection criteria.

Selection	Description
Prioritized	<p>If the Prioritized option is selected, NetBackup chooses the next available storage unit in the list. Prioritized is the default selection.</p> <p>If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.</p>
Failover	<p>If the Failover option is selected, when a job must queue for a storage unit, the job queues rather than try another storage unit in the group.</p>
Round robin	<p>If the Round robin option is selected, NetBackup chooses the least recently selected storage unit in the list as each new job is started.</p> <p>If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.</p>
Media server load balancing	<p>If the Media server load balancing option is selected, NetBackup selects a storage unit based on a capacity-managed approach. In this way, NetBackup avoids sending jobs to busy media servers.</p> <p>If a storage unit is unavailable, NetBackup examines the next storage unit until it finds one that is available.</p> <p>See “Media server load balancing” on page 439.</p>

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 storage unit setting” on page 404.

See “Exception to the storage unit selection criteria” on page 442.

Media server load balancing

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:

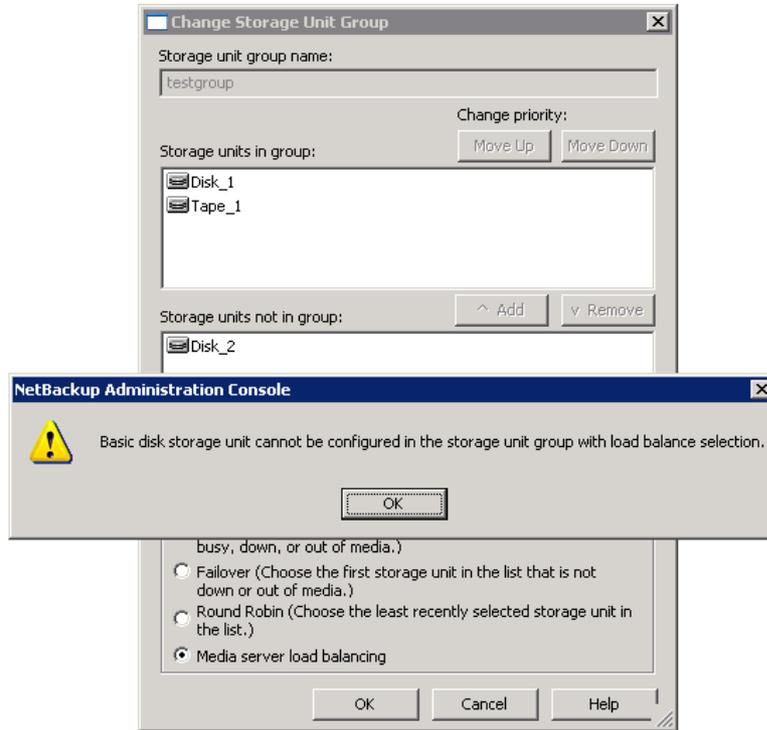
- The rank of the media server.
NetBackup considers the number of processes that are running on each CPU along with the memory thresholds on each server to determine the rank of a media server. If the free memory drops below a determined threshold, or if the number of running processes per CPU rises 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



See "Other load balancing methods" on page 440.

Other load balancing methods

Using the **Media server load balancing** option to balance the storage load requires a license.

The following methods to distribute the backup workload do not require additional licenses:

Adjust the backup load on a media server.	<ul style="list-style-type: none"> ■ Change the Limit jobs per policy policy attribute for one or more of the policies that are sent to a media server. Specifying a lower limit reduces the workload on a media server on a specific network segment. See “Limit jobs per policy (policy attribute)” on page 525. ■ 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 438.
Distribute the backup load on media servers during peak periods.	Reconfigure policy schedules so that they write backups 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 the client.	<p>Change the Maximum jobs per client global attribute. For example, raising the Maximum jobs per client limit increases the number of concurrent jobs that any one client can process and therefore increases the load.</p> <p>See “Storage unit selection criteria within a group” on page 438.</p>
Reduce the time needed 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.	<p>Increase the Limit jobs per policy attribute for the preferred policy relative to other policies. Or, increase the priority for the policy.</p> <p>See “Limit jobs per policy (policy attribute)” on page 525.</p>
Adjust the load between fast and slow networks.	<p>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.</p> <p>See “Limit jobs per policy (policy attribute)” on page 525.</p> <p>See “Storage unit selection criteria within a group” on page 438.</p>
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.	<ul style="list-style-type: none"> ■ 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 storage unit setting” on page 409.

See “Storage unit selection criteria within a group” on page 438.

About 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 438.

Configuring storage lifecycle policies

This chapter includes the following topics:

- About storage lifecycle policies
- Creating a storage lifecycle policy
- Storage Lifecycle Policy dialog box settings
- About associating backup data with a data classification
- Deleting a storage lifecycle policy
- Adding a storage destination to a storage lifecycle policy
- Hierarchical view of storage destinations in the Storage lifecycle policy dialog box
- About writing multiple copies using a storage lifecycle policy
- About storage lifecycle policy versions
- LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration
- Lifecycle operation administration using the nbstlutil command

About storage lifecycle policies

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 “Creating a Data Classification” on page 104.
- 5 Select the **Duplication job priority**. This number represents 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 “Storage Lifecycle Policy dialog box settings” on page 445.
- 6 Click **Add** to add storage destinations to the lifecycle.
 - See “Adding a storage destination to a storage lifecycle policy” on page 450.
 - See “Adding a hierarchical duplication destination” on page 462.
- 7 Click **OK** to create the storage lifecycle. After they are created, data classifications cannot be deleted.

Storage Lifecycle Policy dialog box settings

A storage lifecycle policy consists of one or more storage destinations.

A data classification can be selected for the storage lifecycle, which applies to all of the storage destinations in the lifecycle.

The **New Storage Lifecycle** dialog box and the **Change Storage Lifecycle** dialog box contains the following settings.

Table 14-1 New or Change Storage Lifecycle settings

Setting	Description
Storage lifecycle policy name	The Storage lifecycle policy name describes the storage lifecycle. The name cannot be modified after the SLP is created.
Duplication job priority	<p>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).</p> <p>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.</p> <p>The priority of the backup job is set in the backup policy on the Attributes tab.</p> <p>See “Job priority (policy attribute)” on page 526.</p>

Table 14-1 New or Change Storage Lifecycle settings (*continued*)

Setting	Description
<p>Data classification</p>	<p>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.</p> <p>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.</p> <p>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.</p> <p>The Data classification setting 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.</p> <p>For example, email backup data can be assigned to the silver data classification and financial data backup may be assigned to the platinum classification.</p> <p>A backup policy associates backup data with a data classification. Policy data can be stored only in a storage lifecycle policy with the same data classification.</p> <p>Once data is backed up in a storage lifecycle policy, 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 has expired.</p> <p>See “About associating backup data with a data classification” on page 447.</p> <p>See “Accessing the Data Classification host properties” on page 448.</p>
<p>Storage destination list</p>	<p>The Storage destination list contains all of the destinations where the backup is to be created or copied to. Multiple destinations implies a multiple copies operation.</p> <p>See “About writing multiple copies using a storage lifecycle policy” on page 467.</p> <p>The list also contains columns that display information about each destination. Note that not all columns display by default.</p> <p>For column descriptions, see the following topic:</p> <p>See “New or Change Storage Destination dialog box settings” on page 452.</p>

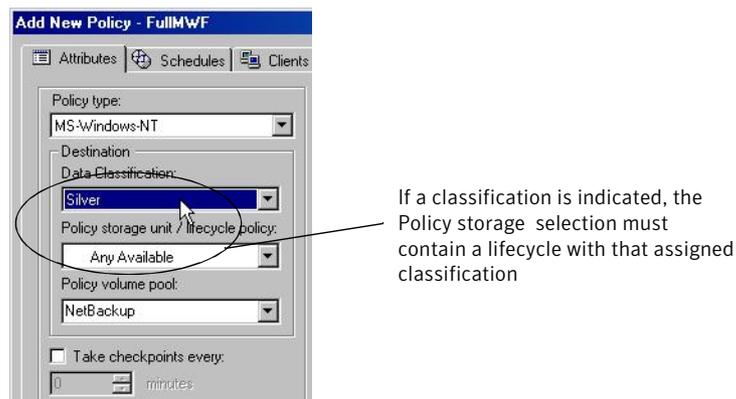
Table 14-1 New or Change Storage Lifecycle settings (*continued*)

Setting	Description
Arrows	<p>Use the arrows to indicate the indentation (or hierarchy) of the source for each copy. One copy can be the source for many other copies.</p> <p>See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.</p> <p>The destination can be hierarchical or non-hierarchical:</p> <ul style="list-style-type: none"> ■ See “Modifying the source of a hierarchical duplication destination” on page 463. ■ See “Adding a non-hierarchical duplication destination” on page 463.

About associating 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



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.

Accessing the Data Classification host properties

NetBackup contains four default data classifications. The name, the description, and the rank of each can be changed in the Data Classification host properties.

To access the Data Classification host properties

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers**.
- 2 In the right pane, double-click a master server.
- 3 In the **Master Server Properties** dialog box, in the left list of **Properties**, click **Data Classifications**.

New data classifications can also be created. However, data classifications cannot be deleted.

See “Data Classification properties” on page 103.

See “Creating a Data Classification” on page 104.

Deleting a storage lifecycle policy

To delete a storage lifecycle policy, use the following procedure:

To delete a storage lifecycle policy

- 1 Remove the storage lifecycle policy from all backup policies.
This step prevents new backup jobs from writing to the storage lifecycle policy.
- 2 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.
This step addresses in-process backup jobs writing to the storage lifecycle policy.
- 3 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.

This step prevents new duplication job from writing to the storage lifecycle policy.

- 4 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.

This step addresses in-process duplication jobs writing to the storage lifecycle policy.

- 5 Once all of the operations are complete, delete the storage lifecycle policy using the To delete a storage lifecycle policy using the Administration Console or To delete a storage lifecycle policy using the `nbstl` command, enter the following procedure below. 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 delete a storage lifecycle policy using the Administration Console

- 1 Expand **Storage > Storage Lifecycle Policies**.
- 2 Select the storage lifecycle policy name.
- 3 Select **Edit > Delete**.
- 4 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 delete a storage lifecycle policy using the `nbstl` command, 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 462.

The screenshot shows the 'New Storage Destination' dialog box. The 'Use for:' dropdown is set to 'Backup'. The 'Retention type' dropdown is set to 'Duplication'. The 'Retention period' is set to '2 weeks (level 1)'. The 'Local storage' radio button is selected. The 'Volume pool' is set to 'NetBackup'. The 'Media Owner' is set to 'Any'. The 'Alternate read server' is empty. The 'Preserve multiplexing' and 'Override job priority' checkboxes are unchecked. The 'Override job priority' value is 0. The dialog box has 'OK', 'Cancel', and 'Help' buttons at the bottom.

- 4 In the **New Storage Destination** dialog box, under **Use for**, 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. Snapshot destinations cannot be used as the source for any duplication.

- Select **Import** if the destination is part of the process to duplicate to a remote master.
See “Process overview to duplicate to a remote master” on page 483.
- 5 Indicate where the backups are to be written:
 - **Remote Master**
This is available when the **Duplication** destination type is selected. Together, they indicate that the duplication is to take place in another domain. This is called duplicating to a remote master.
 - **Local Storage**
Selections include storage units or storage unit groups.

No BasicDisk, SnapVault, or disk staging storage units can be used as destinations in a lifecycle.
 - 6 Indicate the **Volume pool** where the backups (or copies) are to be written.
If this volume pool is a Remote Master Duplicate destination, the **Volume pool** selection is disabled.
 - 7 Indicate the **Media owner** if the storage unit is a Media Manager type and server groups are configured.

By specifying a **Media owner**, you allow only those media servers to write to the media on which backup images for this policy are written.
 - 8 Select the retention type for the destination:
 - **Fixed.**
Fixed is selected automatically when the storage is **Remote master**.
 - **Staged capacity managed.**
 - **Expire after duplication.**
If a policy is configured to back up to a lifecycle, the retention that is indicated in the lifecycle is the value that is used. The **Retention** attribute in the schedule is not used.
See “Retention (schedule attribute)” on page 569.
 - 9 Indicate an **Alternate read server** that is allowed to read a backup image originally written by a different media server.
 - 10 Select whether to **Preserve multiplexing**. This option is available for duplication destinations that use tape media.

If this is a Remote Master Duplicate destination, the **Preserve multiplexing** option is disabled.
 - 11 Click **OK** to create the storage destination.

See “New or Change Storage Destination dialog box settings” on page 452.

New or Change Storage Destination dialog box settings

The **New Storage Destination** and **Change Storage Destination** dialog boxes contain the following settings.

Table 14-2 New or Change Storage Destination dialog box settings

Setting	Description
Use for	<p>Select how the destination is to be used:</p> <ul style="list-style-type: none"> ■ Backup ■ Duplication ■ Snapshot ■ Import <p>An Import destination (or Import SLP) is used on the remote master as part of the duplication to remote master process to import NetBackup images into a remote domain. See “Process overview to duplicate to a remote master” on page 483.</p> <p>See “Use for: Backup, duplication, Snapshot, or Import destination” on page 457.</p> <p>If a storage lifecycle contains multiple destinations, a multiple copies operation is implied. A storage lifecycle policy does not need to contain a duplication destination if staging behavior or creating multiple copies is not the objective.</p> <p>See “About writing multiple copies using a storage lifecycle policy” on page 467.</p> <p>The storage destinations for a storage lifecycle policy must meet the following requirements:</p> <ul style="list-style-type: none"> ■ At least one destination must be a backup destination, unless the SLP is an Import SLP used in the duplication to remote master domain process. (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.

Table 14-2 New or Change Storage Destination dialog box settings (*continued*)

Setting	Description
Retention type	<p>Select a Retention type from the following options:</p> <ul style="list-style-type: none"> <p>■ Fixed</p> <p>Indicates that the data on the storage destination is retained for the specified length of time, after which the backups are expired.</p> <p>■ Staged managed capacity</p> <p>Indicates that the disk space of the volumes in the storage destination is automatically managed by NetBackup, based on the High water mark setting for each volume. This retention type is not available to tape storage units since tape capacity is considered to be infinite.</p> <p>See “Staged capacity managed retention type for storage destinations” on page 455.</p> <p>■ Expire after duplication</p> <p>Indicates that after the data is duplicated to other storage, the data on this storage destination is expired. The last destination in the lifecycle cannot use the Expire after duplication retention type because no subsequent copy is configured.</p> <p>■ Remote retention</p> <p>This setting is used in a duplication to remote master configuration in an Import storage lifecycle policy. It indicates that the data at the remote master shall use the expiration date that was imported with the image. This is actually a fixed date because the copy must have a fixed retention.</p> <p>See “Process overview to duplicate to a remote master” on page 483.</p> <p>■ Maximum snapshot limit</p> <p>Indicates that the snapshot copy is to be deleted based on the maximum snapshot limit defined in the policy.</p>
Remote master	<p>Indicates that the copy of the image is being created in a different master server domain. The remote master server manages the storage where the image is to be copied.</p> <p>If Remote master is selected for a Duplication destination, the destination becomes a Replication Destination for use in a duplication-to-remote-master-configuration.</p>

Table 14-2 New or Change Storage Destination dialog box settings (*continued*)

Setting	Description
Local storage	<ul style="list-style-type: none"> ■ Local storage Indicate the storage unit where the backups are to be written. Select from the following destinations: <ul style="list-style-type: none"> ■ 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. <p>Note: The storage destination list cannot contain other storage lifecycles.</p> 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 Select a Volume pool. The Volume pool option is enabled for tape storage units. ■ Media owner Select a Media owner. A Media owner is a group of NetBackup servers that are used for a common purpose. ■ Alternate read server 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.
Storage unit	<p>Indicate the storage unit where the backups are to be written.</p> <p>Select from the following destinations:</p> <ul style="list-style-type: none"> ■ 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. <p>Note: The storage destination list cannot contain other storage lifecycles.</p> <p>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.</p>
Volume pool	<p>Select a Volume pool. The Volume pool option is enabled for tape storage units.</p>

Table 14-2 New or Change Storage Destination dialog box settings (*continued*)

Setting	Description
Media owner	Select a Media owner . A Media owner is a group of NetBackup servers that are used for a common purpose.
Alternate read server	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.
Preserve multiplexing	<p>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.</p> <p>To preserve multiplexing significantly improves performance of duplication jobs because it eliminates the need to request the write-side duplication media for every image.</p>
Override job priority	The Override job priority option is available when configuring an Import storage destination. The job priority indicated is the job priority for any import jobs which use this storage lifecycle policy.

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 storage unit setting” on page 402.

See “Low water mark storage unit setting” on page 403.

If space is needed for new images, NetBackup removes expired backup images from a capacity managed disk volume in two passes, as follows:

Pass one NetBackup removes any backup images that are past the **Desired cache period** setting. NetBackup removes images until the low water mark is reached or all images that are past the **Desired cache period** are removed.

Pass two Pass two processing is initiated if the outcome of the pass one processing is one of the following:

- The disk pool remains over the high water mark.
- The number of volumes in the disk pool under the high water mark is less than the number of media servers that access the disk pool.

NetBackup removes images until the low water mark is reached or all images that are not past the **Desired cache period** are removed.

An image may be deleted if it has not been duplicated to all destinations in a storage lifecycle policy. If the operating system time is past the data that matches the longest retention period for an image, the image is eligible for deletion.

If the disk pool remains at or over the high water mark after this expiration processing, NetBackup processes the next volume in the disk pool.

See “About writing multiple copies using a storage lifecycle policy” on page 467.

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 in mind the following points 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.

Staged capacity managed retention type and disk types that support SIS

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.

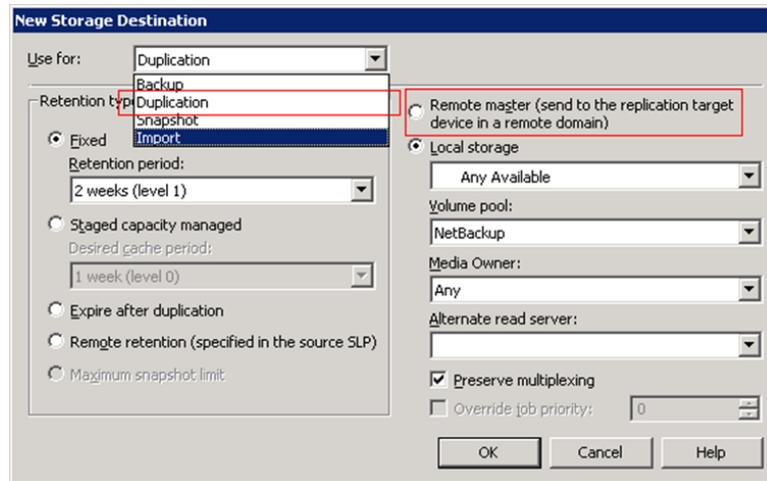
The following storage unit configurations use sis:

- 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.

Use for: Backup, duplication, Snapshot, or Import destination

Select how the destination in the storage lifecycle policy is to be used. If a storage lifecycle contains multiple destinations, a multiple copies operation is implied.

Figure 14-2 New or Change Storage Destination dialog box



A storage destination can be configured for the following purposes:

Table 14-3 Storage destination types

Selection	Description of use
Backup	<p>For images to be written to the destination as part of a backup operation.</p> <p>This is the first destination in a storage lifecycle policy if the SLP is not part of a duplication-to-remote-master configuration.</p>
Duplication	<p>For images to be written to the destination as part of a duplication operation.</p> <p>During the duplication operation, one of the following events occurs:</p> <ul style="list-style-type: none"> ■ If Local storage is selected, backups are duplicated to secondary storage. ■ If Remote master is selected, the destination becomes a Replication Destination for use in a duplication-to-remote-master-configuration. Remote master indicates that the copy is to be created in another master server domain. No storage destination is explicitly specified. Instead, the storage destination is configured through the replication properties of the source copy's storage.
Duplication	<p>For Duplication images to be written to the destination as part of a duplication operation. During the duplication operation, backups are duplicated to secondary storage.</p>

Table 14-3 Storage destination types (*continued*)

Selection	Description of use
Snapshot	<p>For images to be written to the destination as part of the snapshot operation.</p> <p>Snapshot destinations cannot be used as the source for any duplication.</p>
Import	<p>Use in the SLP on the remote master. This option indicates that the SLP will automatically import replicated images into the remote master.</p> <p>If the Import option is selected:</p> <ul style="list-style-type: none"> ■ This must be the first destination in a storage lifecycle policy and the only Import destination. ■ Other destinations in the storage lifecycle policy must be Duplication destinations. ■ At least one destination in the storage lifecycle policy must use the Remote retention retention type. <p>The Override job priority option can be selected. This allows administrators to specify a job priority for any import jobs which use this SLP.</p>

Retention type mixing for storage destinations

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 in mind the following points 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.

Hierarchical view of storage destinations in the Storage lifecycle policy dialog box

The storage destination list contains all the destinations 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 14-3 shows how after the first copy is created, all subsequent copies can be made locally from that source, without tying up network resources.

Figure 14-3 Hierarchical destinations

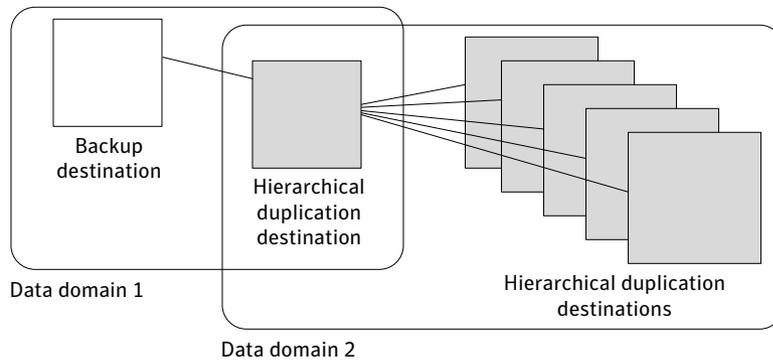
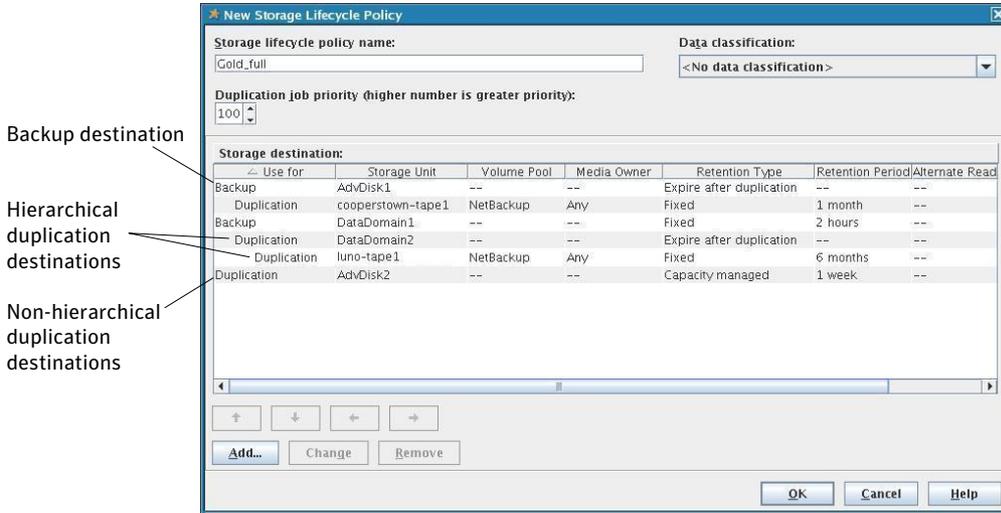


Figure 14-4 shows how the hierarchy of destinations is represented graphically in the **Storage Lifecycle Policy** dialog box. The indentation (or hierarchy) indicates the backup source for each copy. One copy can be the source for many other copies.

Hierarchical view of storage destinations in the Storage lifecycle policy dialog box

Figure 14-4 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>

See “Adding a hierarchical duplication destination” on page 462.

See “Adding a non-hierarchical duplication destination” on page 463.

See “Modifying the source of a hierarchical duplication destination” on page 463.

See “Removing a destination from the storage destination list” on page 465.

See “Example of storage destination hierarchical view” on page 465.

Adding a hierarchical duplication destination

A hierarchical duplication destination means that the destination uses a specific source.

To add a hierarchical duplication destination

- 1 In the **Change Storage Lifecycle Policy** dialog box, select a backup or a duplication storage destination to become the source of the destination to be added.
- 2 Click **Add**.
- 3 In the **New Storage Destination** dialog box, select the **Duplication** type. 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.

See “Adding a non-hierarchical duplication destination” on page 463.

See “Modifying the source of a hierarchical duplication destination” on page 463.

See “Removing a destination from the storage destination list” on page 465.

See “Example of storage destination hierarchical view” on page 465.

See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.

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 box.
- 2 Click **Add**.
- 3 In the **New Storage Destination** dialog box, select the **Duplication** type. 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.

See “Adding a hierarchical duplication destination” on page 462.

See “Modifying the source of a hierarchical duplication destination” on page 463.

See “Removing a destination from the storage destination list” on page 465.

See “Example of storage destination hierarchical view” on page 465.

See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.

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 In the **Change Storage Lifecycle Policy** dialog box, select the hierarchical duplication destination.
- 2 Click the arrows to move the destination into the new position.
 - 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.

See “Adding a hierarchical duplication destination” on page 462.

See “Adding a non-hierarchical duplication destination” on page 463.

See “Removing a destination from the storage destination list” on page 465.

See “Example of storage destination hierarchical view” on page 465.

See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.

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 In the **Change Storage Lifecycle Policy** dialog box, select the destination.
- 2 Click **Remove**. The destination is removed from the list of destinations. The children shift left in the hierarchy.

See “Adding a hierarchical duplication destination” on page 462.

See “Adding a non-hierarchical duplication destination” on page 463.

See “Modifying the source of a hierarchical duplication destination” on page 463.

See “Example of storage destination hierarchical view” on page 465.

See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.

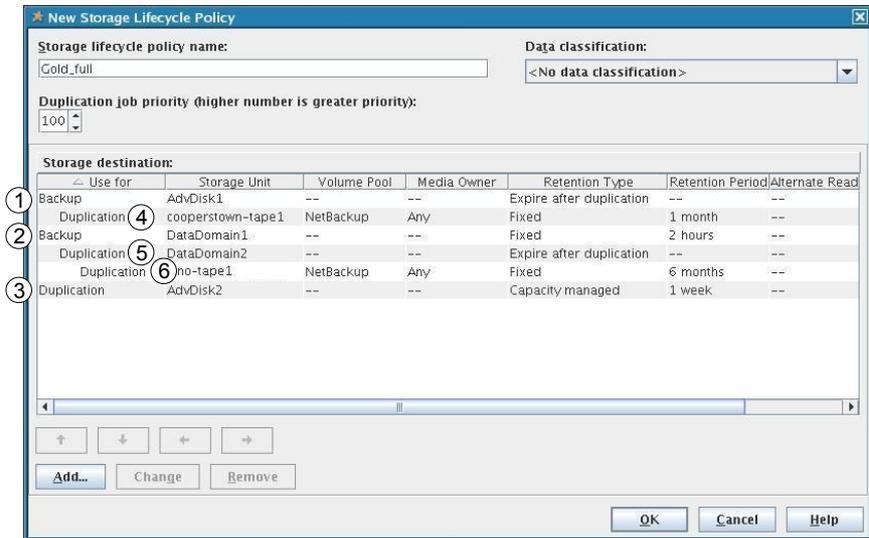
Example of storage destination hierarchical view

Figure 14-5 shows a configuration with hierarchical and non-hierarchical storage destinations.

Notice the following aspects of this example:

- The backup source for each copy.
- The order in which the copies are created.

Figure 14-5 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 14-4 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.
6	Hierarchical duplication	Indented under 5; the copy uses duplication destination 5 as the source.

See “Adding a hierarchical duplication destination” on page 462.

See “Adding a non-hierarchical duplication destination” on page 463.

See “Modifying the source of a hierarchical duplication destination” on page 463.

See “Removing a destination from the storage destination list” on page 465.

See “Hierarchical view of storage destinations in the Storage lifecycle policy dialog box” on page 460.

About 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.

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 (schedule attribute)” on page 562.
- 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 450.

The same criteria for creating copies applies to both methods.

The following topics are considerations when using storage lifecycle policies to create multiple copies.

How destination order determines the copy order

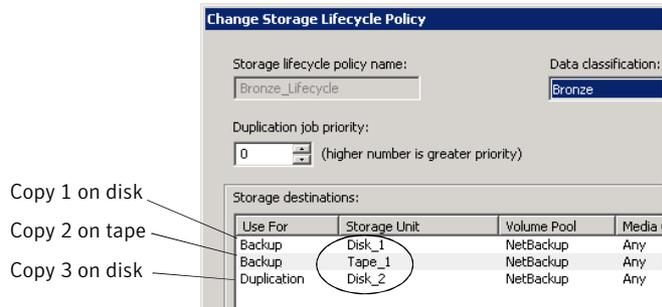
The order in which the destinations appear in a lifecycle determines the copy number of the backup.

For example, in Figure 14-6 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 14-6 Destination order determines copy order



About 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 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 Infinity 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 “Lifecycle operation administration using the `nbstlutil` command” on page 480.

About 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 associated with the 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.

Storage lifecycle changes and versioning

Administrators can make changes to a storage lifecycle policy in one of the following ways:

- Using the NetBackup Administration Console or `bpadm` command.
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.
- Using the `nbstl` command.
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. The options are as follows:

<code>-all_versions</code>	Use to display all versions of a storage lifecycle policy definition. Without specifying this option, only the most recent version is displayed by default.
<code>-version <i>number</i></code>	Use to display a specific version.
<code>-modify_current</code>	Use with most <code>nbstl</code> 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.
<code>-modify_version</code> <code>-version <i>number</i></code>	Use with most <code>nbstl</code> 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:

<code>-dp</code>	The duplication priority
<code>-residence</code>	The storage unit to be used for each destination
<code>-pool</code>	The volume pool for each destination
<code>-server_group</code>	The server group for each destination
<code>-rl</code>	The retention level for each destination
<code>-as</code>	The alternate read server for each destination
<code>-mpx</code>	The preserve multiplexing option for duplication copies

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:

<code>-uf</code>	The type of the destination: Backup, duplication, snapshot
<code>-managed</code>	The retention type for the destination: Fixed, capacity managed, expire after duplication
<code>-source</code>	The source of a destination, used primarily in hierarchical storage lifecycle policy configurations
<code>-dc</code>	The data classification of an existing version The number of destinations. You cannot add a destination or remove a destination from the storage lifecycle policy definitions.

See “Creating a storage lifecycle policy” on page 444.

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 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 the job is canceled, `nbstserv` re-forms and re-submits new duplication jobs for these image copies, using the changes to the appropriate version of the 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.

About 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:

- CLEANUP_SESSION_INTERVAL_HOURS
- VERSION_CLEANUP_DELAY_HOURS

See “LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration” on page 472.

LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration

The NetBackup administrator can customize how the NetBackup Storage Lifecycle Manager (`nbstserv`) runs duplication and import jobs.

Both the Duplication Manager service and the Import Manager service run within `nbstserv`. Table 14-5 describes the role of each service.

Table 14-5 Role of the nbstserv services

nbstserv service	Purpose of service	Location of LIFECYCLE_PARAMETER file
Duplication Manager	<p>Manages the duplication jobs within storage lifecycle policies.</p> <p>In the duplication to remote master process, the Duplication Manager duplicates images and creates batches of the images to be imported in the target domain.</p>	<p>Configure Duplication Manager parameters in the source domain to tune the duplication to remote master jobs.</p> <p>Note: The Duplication Manager parameters affect all SLP duplications, even those that are not duplicated to a remote master. Exercise caution when tuning for one case or the other.</p>
Import Manager	<p>In the duplication to remote master process, the Import Manager monitors a worklist in EMM for images to be imported and initiates <code>bpimport</code> jobs for those images.</p> <p>If the NetBackup environment is not configured for duplication to a remote master, the Import Manager does not perform any actions.</p>	<p>If the NetBackup environment is configured for duplication to a remote master, configure Import Manager parameters in the target domain.</p> <p>See “Process overview to duplicate to a remote master” on page 483.</p>

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 in Table 14-6 can appear in the `LIFECYCLE_PARAMETERS` file in any order. If the file does not exist, NetBackup uses the defaults as indicated.

Table 14-6 Lifecycle parameters

Parameter	Description
<p>AUTO_CREATE_IMPORT_SLP</p> <p>Affects: Import Manager</p>	<p>Indicates to the Import Manager how NetBackup should handle notifications from storage about images for which there is no matching Import storage lifecycle policy.</p> <p>The entry is Boolean, where a non-zero value directs NetBackup to create a storage lifecycle policy definition that uses the name provided in the import notification to the storage device.</p> <p>Syntax: AUTO_CREATE_IMPORT_SLP 0 1</p> <p>0 = Select 0 to indicate that NetBackup should not automatically create an Import SLP if a notification is received for an Import SLP which does not exist.</p> <p>To remove the parameter from LIFECYCLE_PARAMETERS file has the same effect as 0.</p> <p>1 = Select 1 to automatically create an Import SLP if a notification is received for an Import SLP which does not exist.</p> <p>The SLP that is automatically created has the following characteristics:</p> <ul style="list-style-type: none"> ■ The SLP is always a data classification of None. ■ The SLP always uses the default import priority. ■ The SLP always specifies any storage unit which includes the device from which the event was received. <p>Default: 0; storage lifecycle policies are not created automatically.</p> <p>Note: This parameter is used primarily to duplicate to a remote master using NetBackup appliances.</p> <p>See “Configuring the storage lifecycle policies required to duplicate to a remote master server” on page 488.</p>
<p>CLEANUP_SESSION_INTERVAL_HOURS</p> <p>Affects: Duplication Manager</p>	<p>Concerns the deletion of storage lifecycle policy versions where a more recent version exists.</p> <p>Controls how often <code>nbstserv</code> looks for the versions that have been deleted.</p> <p>Syntax: CLEANUP_SESSION_INTERVAL_HOURS <i>nn_hours</i></p> <p>Default: 24 (24 hours).</p> <p>See “About deleting old storage lifecycle policy versions” on page 472.</p>

Table 14-6 Lifecycle parameters (*continued*)

Parameter	Description
<p>DUPLICATION_GROUP_CRITERIA</p> <p>Affects:</p> <p>Duplication Manager</p>	<p>Indicates how batches are created. The entry applies to the use of tape and disk.</p> <p>Syntax: DUPLICATION_GROUP_CRITERIA 0 1</p> <p>0 = Select 0 to indicate that batches be created based on the storage lifecycle policy name.</p> <p>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.</p> <p>Default: 1; use the storage lifecycle policy name.</p>
<p>DUPLICATION_SESSION_INTERVAL_MINUTES</p> <p>Affects:</p> <p>Duplication Manager</p>	<p>Indicates how frequently the Duplication Manager 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.</p> <p>Syntax: DUPLICATION_SESSION_INTERVAL_MINUTES 5</p> <p>Default: 5 (five minutes). Minimum: one minute.</p>
<p>IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS</p> <p>Affects:</p> <p>Duplication Manager</p>	<p>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.)</p> <p>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 <code>nbstlutil</code>.</p> <p>Syntax: IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS 2</p> <p>Default: 2 (two hours). Minimum: one hour.</p>
<p>IMPORT_EXTENDED_RETRY_SESSION_TIMER</p> <p>Affects:</p> <p>Import Manager</p>	<p>After four failed attempts, the Import Manager retries at the extended retry interval indefinitely or until the number of days specified by the REPLICA_METADATA_CLEANUP_TIMER parameter has elapsed.</p> <p>Syntax: IMPORT_EXTENDED_RETRY_SESSION_TIMER 360</p> <p>Default: 360 (360 minutes; six hours).</p>

Table 14-6 Lifecycle parameters (*continued*)

Parameter	Description
<p>IMPORT_SESSION_TIMER</p> <p>Affects: Import Manager</p>	<p>Indicates to the Import Manager how many minutes to wait between import sessions.</p> <p>During an import session, the Import Manager performs the following tasks:</p> <ul style="list-style-type: none"> ■ Cleans up image import status and increments the retry count of failed imports. The Import Manager retries in the current session. After four failed attempts to import an image, NetBackup enters an extended retry state. (See the IMPORT_EXTENDED_RETRY_SESSION_TIMER parameter description in this table.) ■ Deletes old, unimported import records. (By default, the records are not deleted and the Import Manager retries indefinitely. For the Import Manager to delete old import records, configure the REPLICATA_METADATA_CLEANUP_TIMER parameter.) ■ Creates batches of images for import and starts import jobs. <p>Syntax: IMPORT_SESSION_TIMER 5</p> <p>Default: 5 (five minutes).</p>
<p>MIN_GB_SIZE_PER_DUPLICATION_JOB</p> <p>Affects: Duplication Manager</p>	<p>Indicates the size that the batch of images should reach before one duplication job is run for the entire batch.</p> <p>The lifecycle does not request a duplication job until either:</p> <ul style="list-style-type: none"> ■ 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. <p>Syntax: MIN_GB_SIZE_PER_DUPLICATION_JOB <i>GB_value</i></p> <p>Default: 7 (7 gigabytes).</p>
<p>MAX_GB_SIZE_PER_DUPLICATION_JOB</p> <p>Affects: Duplication Manager</p>	<p>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.</p> <p>Syntax: MAX_GB_SIZE_PER_DUPLICATION_JOB <i>GB_value</i></p> <p>Default: 25 (25 gigabytes).</p>

Table 14-6 Lifecycle parameters (*continued*)

Parameter	Description
<p><code>MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB</code></p> <p>Affects: Duplication Manager</p>	<p>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.</p> <p>The <code>MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB</code> entry working differently in this release than it did in previous releases.</p> <p>A very small batch is not submitted to <code>nbstserv</code> until one source job in the batch has finished at least 30 minutes ago.</p> <p>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.</p> <p>This parameter helps to ensure a balance between the following conditions:</p> <ul style="list-style-type: none"> ■ Submitting many small duplication jobs too soon or too frequently. On the one hand, <code>nbstserv</code> 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, <code>nbstserv</code> should not wait too long to submit a small job. <p>Syntax: <code>MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB 30</code></p> <p>Default: 30 (30 minutes).</p>
<p><code>REPLICA_METADATA_CLEANUP_TIMER</code></p> <p>Affects: Import Manager</p>	<p>Indicates the number of days after which the Import Manager stops trying to import the image. After the number of days indicated, the record is deleted.</p> <p>How frequently the Import Manager tries to import the images depends on the setting of the extended retry timer and session timer. The first four attempts occur at the regular session interval and the remaining attempts occur at the extended retry interval.</p> <p>Syntax: <code>REPLICA_METADATA_CLEANUP_TIMER 0 n</code></p> <p>Default: 0 (off).</p>

Table 14-6 Lifecycle parameters (*continued*)

Parameter	Description
<p>TAPE_RESOURCE_MULTIPLIER</p> <p>Affects: Duplication Manager</p>	<p>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.</p> <p>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.</p> <p>This parameter helps administrators ensure a balance in the following situation:</p> <ul style="list-style-type: none"> ■ To overload the Resource Broker with jobs 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. <p>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.</p> <p>Syntax: TAPE_RESOURCE_MULTIPLIER <i>n</i></p> <p>Default: 2 (multiplier of two).</p>
<p>VERSION_CLEANUP_DELAY_HOURS</p> <p>Affects: Duplication Manager</p>	<p>Concerns the deletion of storage lifecycle policy versions where a more recent version exists.</p> <p>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.</p> <p>Syntax: VERSION_CLEANUP_DELAY_HOURS <i>nn_hours</i></p> <p>Default: 336 (336 hours; 14 days).</p> <p>See "About deleting old storage lifecycle policy versions" on page 472.</p>

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
    
```

See “About batch creation logic in Storage Lifecycle Manager” on page 479.

See “LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration” on page 472.

About batch creation logic in Storage Lifecycle Manager

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 “Storage Lifecycle Policy dialog box settings” on page 445.

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, do not duplicate 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 “`LIFECYCLE_PARAMETERS` file for optional lifecycle-managed job configuration” on page 472.

Lifecycle operation administration using the nbstlutil command

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.

Table 14-7 nbstlutil details

nbstlutil information	Details
Where to find	The command is found in the following location: <code>install_path\NetBackup\bin\admincmd\nbstlutil</code>
How to use	Use <code>nbstlutil</code> to perform the following administrative actions: <ul style="list-style-type: none"> ■ 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. It does not attempt to create any more copies of the backup image. ■ Deactivate (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 the <i>NetBackup Commands Reference Guide</i> for a description of all the options available for <code>nbstlutil</code> .

Table 14-7 nbstlutil details (*continued*)

nbstlutil information	Details
When to use	<p>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 <code>DUPLICATION_SESSION_INTERVAL_MINUTES</code> parameter.)</p> <p>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 <code>IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS</code> parameter.)</p> <p>Use the <code>nbstlutil</code> 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.</p> <p>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 <code>nbstlutil</code> command to inactivate the lifecycle while the library is repaired. When ready, the lifecycle can be activated and duplication jobs can begin.</p> <p>Note: Once the job is reactivated, the administrator may want to temporarily change the <code>IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS</code> parameter to one hour to begin duplication jobs sooner.</p>

Duplicating images to a remote master server domain

This chapter includes the following topics:

- Process overview to duplicate to a remote master
- Setup overview to duplicate to a remote master domain
- About defining the domain relationship
- Configuring the storage lifecycle policies required to duplicate to a remote master server
- One-to-many duplication to remote master server model
- Cascading duplications to remote masters
- Restoring from a backup at a remote master domain
- Reporting on duplication to remote master jobs

Process overview to duplicate to a remote master

The backups that are generated in one NetBackup domain can be replicated to storage in one or more NetBackup domains. This process is referred to as duplication to a remote master.

The ability to duplicate backups to storage in other NetBackup domains, often across various geographical sites, helps facilitate the following disaster recovery needs:

- **One-to-many model**
 A single production datacenter can back up to multiple disaster recovery sites. See “One-to-many duplication to remote master server model” on page 492.
- **Many-to-one model**
 Remote offices in multiple domains can back up to a storage device in a single domain.
- **Many-to-many model**
 Remote datacenters in multiple domains can back up multiple disaster recovery sites.

Note: Although duplicating to a remote master domain is a disaster recovery solution, the administrator cannot directly restore to clients in the primary (originating) domain from the remote master domain.

Table 15-1 is an overview of the process, generally describing the events in the originating domain and in the target domain.

Table 15-1 Process overview to duplicate to a remote master domain

Event	Domain in which event occurs	Event description
1	Originating master (Domain 1)	Clients are backed up according to a policy that indicates a storage lifecycle policy as the Policy storage selection. At least one of the duplication destinations in the SLP must be configured for duplication to one of the following types of devices: <ul style="list-style-type: none"> ■ An OpenStorage (OST) appliance on a remote master. ■ An Media Server Deduplication Pool (MSDP) on a remote master. See “Configuring the storage lifecycle policies required to duplicate to a remote master server” on page 488.
2	Remote master (Domain 2)	The storage server (that represents the OpenStorage appliance or MSDP) in the remote domain recognizes that a replication event has occurred and notifies the NetBackup master server in that domain.
3	Remote master (Domain 2)	NetBackup imports the image immediately. The image import is referred to as OST Optimized Duplication. NetBackup can import the image quickly because the metadata is duplicated as part of the image. (This import process is not the same as the import process available in the Catalog utility.)

Table 15-1 Process overview to duplicate to a remote master domain (*continued*)

Event	Domain in which event occurs	Event description
4	Remote master (Domain 2)	After the image is imported into the remote domain, NetBackup continues to manage the copies in that domain. Depending on the configuration, the master server in Domain 2 can replicate the images to a master server in Domain 3.

Setup overview to duplicate to a remote master domain

Table 15-2 is an overview of the setup process to duplicate to a remote master domain, describing the actions that are required.

Table 15-2 Setup overview to duplicate to a remote master domain

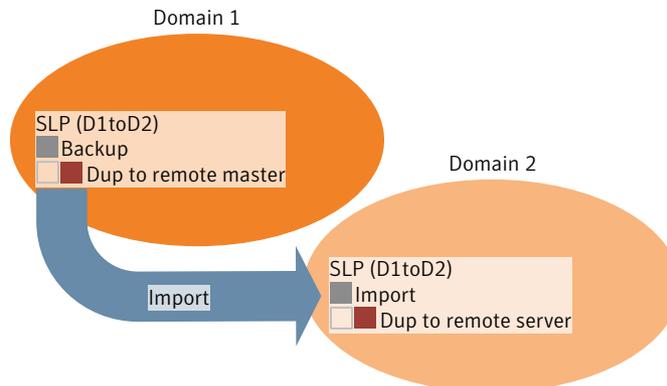
Step	Action	Description
Step 1	Install or upgrade NetBackup.	All master servers and media servers must be at NetBackup version 7.1 or later. See the <i>NetBackup Installation Guide</i> .
Step 2	Configure the storage devices.	To send images from one domain to another requires that suitable storage be configured in each domain. The storage in the originating domain and the storage in the target domain must be of the same type. The storage can be either of the following types: <ul style="list-style-type: none"> ■ OpenStorage (OST) appliances whose plug-ins conform to version v11.1 of the OpenStorage API. See the <i>NetBackup Shared Storage Guide</i>. ■ Media Server Deduplication Pools (MSDP). For MSDP, the plug-in is installed with NetBackup; no separate installation is required. To use MSDP, a Media Server Deduplication Pool must be configured in both domains. When you configure the disk pool in the target domain, consider using the Limit I/O streams setting in the Maximum I/O Streams section. Doing so can reduce the load on the target storage server. See the <i>NetBackup Deduplication Guide</i>. <p>Note: For successful import and duplication, make sure that the storage appliances are working properly in each domain.</p>

Table 15-2 Setup overview to duplicate to a remote master domain (continued)

Step	Action	Description
Step 3	Configure the storage units.	Configure the storage units in both the source domain and the target domain. The storage unit in the source domain should not be used for backups other than duplication to remote master.
Step 4	Define the relationship between the domains.	Define the relationship between the domains so that the source domain knows where to send the data. See “About defining the domain relationship” on page 487.
Step 5	Configure the storage lifecycle policies.	Configure a pair of storage lifecycle policies; one in each master server domain. The storage lifecycle policy pair includes: <ul style="list-style-type: none"> ■ An SLP in the originating domain that contains a duplication-to remote-master destination. (The Duplication-to remote-master SLP.) ■ An SLP in the target domain that contains an import destination. (The Import SLP.) The following topic describes how the SLPs must be named and the storage destinations and retention type that each SLP must contain. See “Configuring the storage lifecycle policies required to duplicate to a remote master server” on page 488.
Step 6	Configure and run the backup policy.	The backup policy must indicate the configured SLP as the Policy storage selection.

Figure 15-1 represents the process of importing images using storage lifecycle policies.

Figure 15-1 Duplicating from one master server domain to another



About defining the domain relationship

The following items describe important configuration differences depending on which method is used for duplication to a remote server.

- Using media server deduplication pools:
The relationship between the source and the target domains is established by setting the properties in the source storage server.
Specifically, in the **Replication** tab of the **Change Storage Server** dialog box when configuring the MSDP storage server.
See “Configuring a replication target using MSDP” on page 487.
- Using OpenStorage appliances:
For OpenStorage duplication to a remote master server, the source NetBackup domain has no knowledge of the remote domain or its storage server. The relationship between the source and the target domains is configured using the disk appliance vendor's tools. When the appliances are configured properly, NetBackup images on the source disk appliance are replicated automatically to the target disk appliance. That disk appliance uses the OpenStorage API to notify NetBackup that a replication event occurred. NetBackup then imports those images.

Caution: Choose the target storage server or servers carefully. A target storage server must not also be a storage server for the source domain.

Information about how to configure media server deduplication pools is available in the *NetBackup Deduplication Guide*.

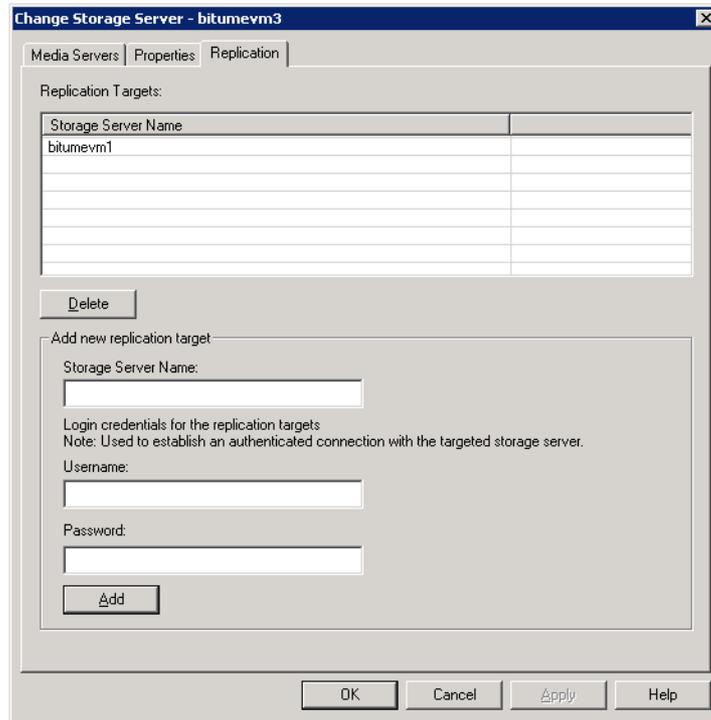
Configuring a replication target using MSDP

When using **Media Server Deduplication Pool**, use the following procedure to establish the relationship between the source and the target domains.

To configure a replication target

- 1 In the NetBackup Administration Console, expand **Media and Device Management > Credentials > Storage Server**
- 2 Select the MSDP storage server.
- 3 On the **Edit** menu, select **Change**.

- 4 In the **Change Storage Server** dialog box, select the **Replication** tab.



- 5 To add a replication target:
 - Enter the **Storage Server Name**.
 - Enter **Username** and **Password** credentials for the NetBackup Deduplication Engine.
 - Click **Add** to add the storage server to the **Replication Targets** list.
All targets are considered for duplication, depending on the rules of the storage lifecycle policies that control the duplication.
- 6 When finished adding replication targets, click **OK**.

Configuring the storage lifecycle policies required to duplicate to a remote master server

To duplicate images from one master server domain to a remote domain requires that two storage lifecycle policies be configured:

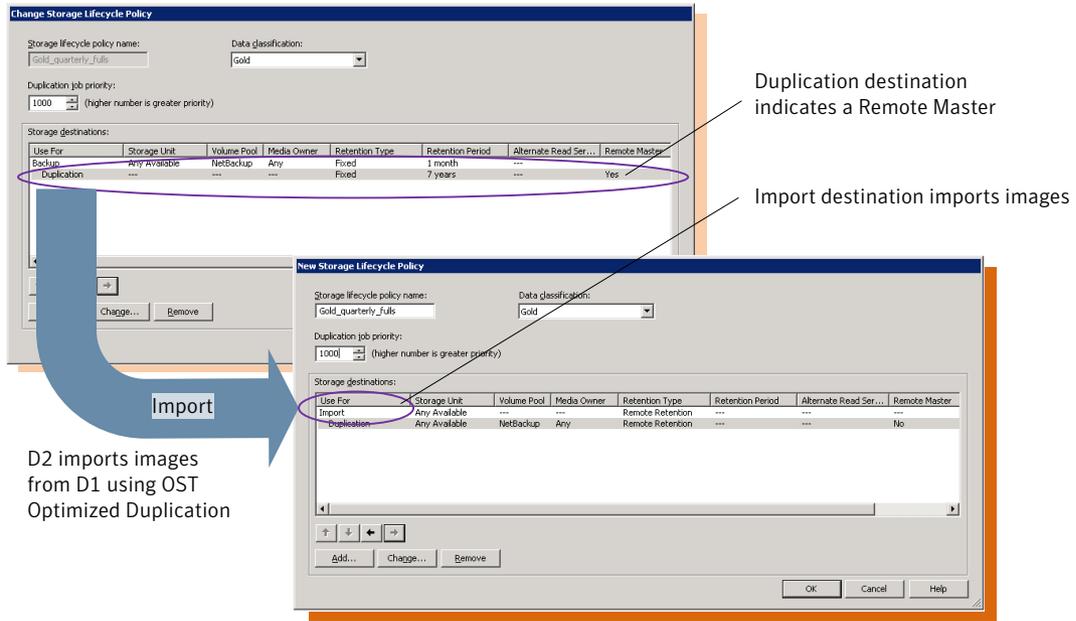
Configuring the storage lifecycle policies required to duplicate to a remote master server

- One SLP in the originating domain that is set up to duplicate to a remote server. (The Duplication-to remote-master SLP.)
- One SLP in the target domain that is set up to import the duplication. (The Import SLP.) The Import SLP can be configured to create additional copies in that domain or another.

Figure 15-2 shows how the SLP in the target domain is set up to duplicate the images from the originating master server domain.

Figure 15-2 Storage lifecycle policy pair required to duplicate to a remote master

Duplication-to-remote-master SLP
Domain 1



Import SLP
Domain 2 (the remote master)

Table 15-3 describes the requirements for each SLP in the pair.

Table 15-3 SLP requirements for duplication to remote master

Domain	Storage lifecycle policy requirements
Domain 1 (Originating domain)	<p>The Duplication-to remote-master SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The SLP must have the same name as the Import SLP in Domain 2. ■ The SLP must be of the same data classification as the Import SLP in Domain 2. ■ The backup destination must be to an OpenStorage (OST) appliance or Media Server Deduplication Pool (MSDP). Indicate the exact storage unit from the drop-down list. Do not select Any Available. <p>Note: The target domain must contain the same type of storage to import the image.</p> <ul style="list-style-type: none"> ■ At least one destination must be a duplication destination with the Remote Master option selected. <p>See Figure 15-3 on page 491.</p> <p>Multiple duplication destinations can be configured for a remote master server. The master server in Domain 1 does not know which remote master will be selected. If multiple SLPs at target domains meet the criteria, NetBackup will import copies in all qualifying domains. See “New or Change Storage Destination dialog box settings” on page 452.</p>
Domain 2 (Target domain)	<p>The Import SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The SLP must have the same name as the SLP in Domain 1 that indicates a Remote Master for the duplication destination. The matching name indicates to the SLP which images to process. ■ The SLP must be of the same data classification as the SLP in Domain 1 that indicates a Remote Master for the duplication destination. Matching the data classification keeps a consistent meaning to the classification and facilitates global reporting by data classification. ■ The first destination in the SLP must be an Import destination. <p>Indicate the exact storage unit from the drop-down list. Do not select Any Available</p> <p>See Figure 15-4 on page 491.</p> <ul style="list-style-type: none"> ■ The SLP must contain at least one duplication destination. One of the duplication destinations must have the Remote retention specified.

The following topic describes useful reporting information about duplication and import jobs.

See “Reporting on duplication to remote master jobs” on page 496.

Figure 15-3 Duplication destination with Remote master selected in Domain 1 storage lifecycle policy

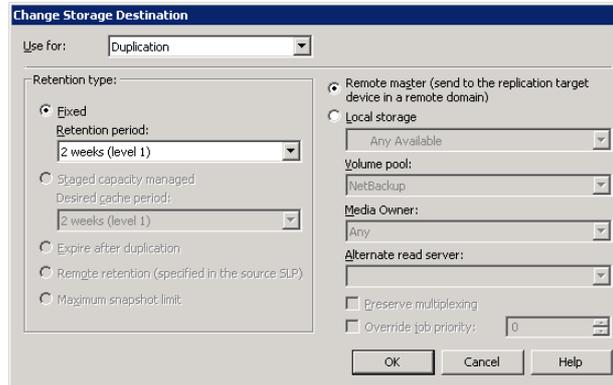
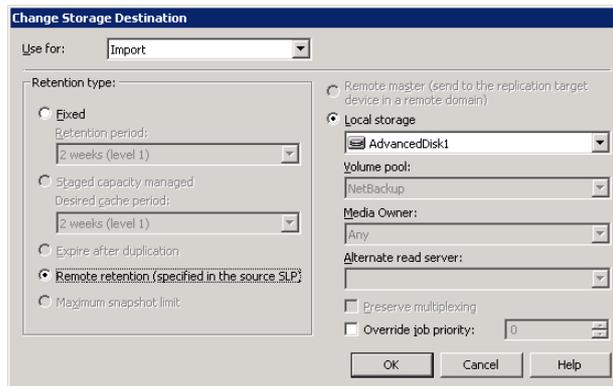


Figure 15-4 Import destination in Domain 2 storage lifecycle policy



Customizing how nbstserv runs duplication and import jobs

The NetBackup Storage Lifecycle Manager (`nbstserv`) runs duplication and import jobs. Both the Duplication Manager service and the Import Manager service run within `nbstserv`.

The NetBackup administrator can customize how `nbstserv` runs duplication and import jobs by adding parameters to the `LIFECYCLE_PARAMETERS` file.

For example, the administrator can indicate how NetBackup should handle the situation in which an appropriate Import SLP does not exist in the target domain. The administrator can also indicate how many attempts NetBackup makes to retry a failed import job.

See “LIFECYCLE_PARAMETERS file for optional lifecycle-managed job configuration” on page 472.

One-to-many duplication to remote master server model

A one-to-many duplication to remote master server model is configured on the OpenStorage storage device. Do this by configuring a logical unit of space (referred to as an LSU or disk volume) with multiple replication targets for OpenStorage.

In this configuration, all copies are made in parallel. The copies are made within the context of one NetBackup job and simultaneously within the source storage server context. If one target storage server fails, the entire job fails and is retried later.

All copies have the same **Remote Retention**. To achieve different **Remote Retention** settings in each target master server domain, either create multiple source copies or cascade duplication to remote servers.

See “Cascading duplications to remote masters” on page 492.

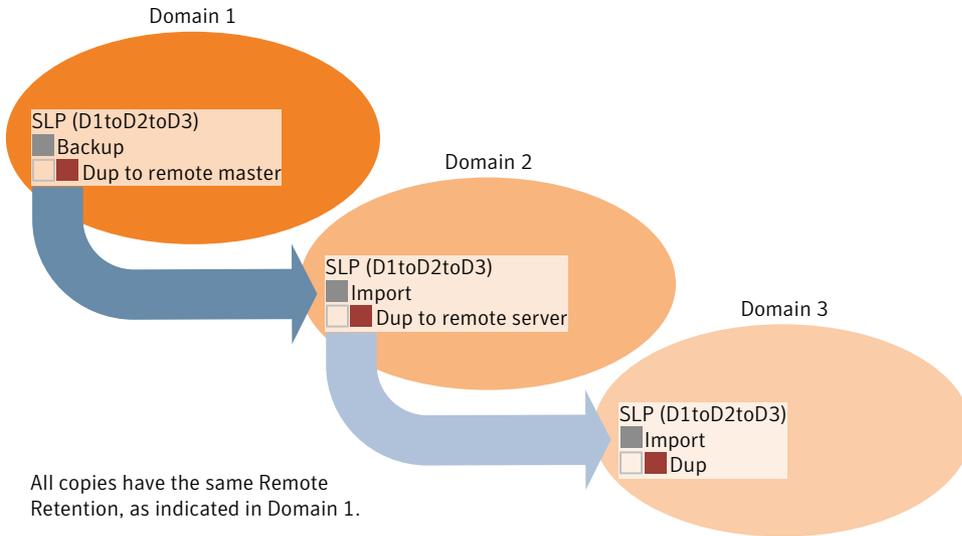
Cascading duplications to remote masters

Duplications can be cascaded from one master server domain to multiple domains. To do so, storage lifecycle policies are set up in each domain to anticipate the source image, import it and then duplicate it to the next remote master.

Figure 15-5 represents the following cascading configuration across three domains.

- The image is created in master server Domain 1, and then duplicated to remote master server Domain 2.
- The image is imported in Domain 2, and then duplicated to remote master server Domain 3.
- The image is then imported in Domain 3.

Figure 15-5 Cascading duplications to remote master servers



In the cascading model, the originating master server for Domain 2 and Domain 3 is the master server in Domain 1.

Note: When the image is replicated in Domain 3, the replication notification event initially indicates that the master server in Domain 2 is the originating master server. However, when the image is successfully imported into Domain 3, this information is updated to correctly indicate that the originating master server is in Domain 1.

The cascading model presents a special case for the Import SLP that will duplicate the imported copy to a remote master. (This is the master server that is neither the first nor the last in the string of remote servers.)

As discussed previously, the requirements for an Import SLP include at least one destination that uses a **Fixed** retention type and at least one destination that uses a **Remote Retention** type. So that the Import SLP can satisfy these requirements, the import destination must use a **Remote retention**.

Table 15-4 shows the difference in the import destination setup.

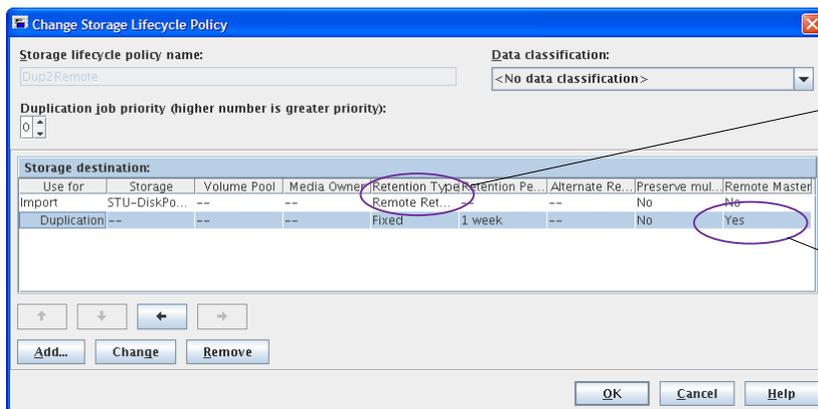
Table 15-4 Import destination difference in an SLP configured to duplicate the imported copy

Import destination criteria	Import destination in a cascading model
The first destination must be an import destination.	Same; no difference.
A duplication to remote master must use a Fixed retention type	Same; no difference.
At least one destination must use the Remote retention .	Here is the difference: To meet the criteria, the import destination must use Remote retention .

The remote retention is embedded in the source image of the duplication to the remote master server domain.

Because the imported copy is the copy being duplicated to a remote master server domain, the fixed retention (three weeks in this example) on the duplication to remote master destination is ignored. The remote retention is used instead. (See Figure 15-6.)

Figure 15-6 Storage lifecycle policy configured to duplicate the imported copy



Remote retention of image is used

Duplication goes to domain

In the cascading model that is represented in Figure 15-5, all copies have the same **Remote Retention**—the **Remote Retention** indicated in Domain 1.

For the copy in Domain 3 to have a different remote retention, add an intermediary duplication destination to the Domain 2 storage lifecycle policy. The intermediary duplication destination acts as the source for the duplication to remote master. Since the remote retention is embedded in the source image of the duplication to

remote master, the copy in Domain 3 honors the retention level that is set for the intermediary duplication destination.

Figure 15-7 Cascading duplications to remote master servers, with various remote retentions

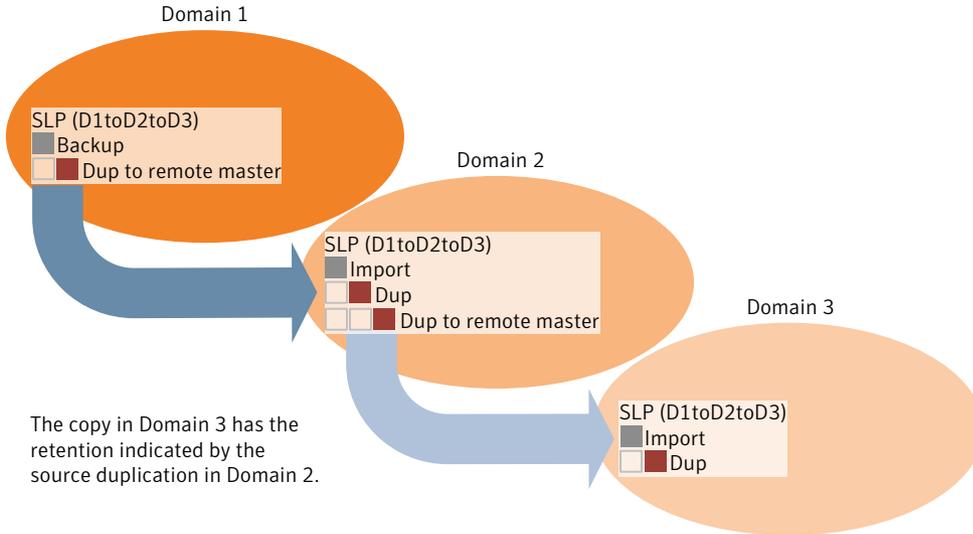
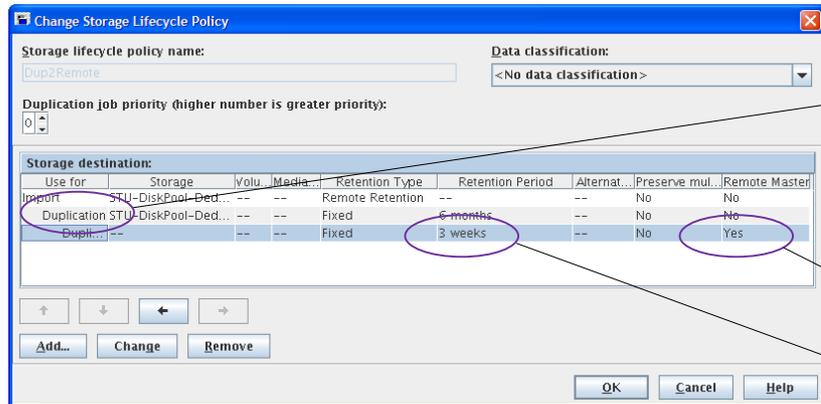


Figure 15-8 shows the storage lifecycle policy in Domain 2. The remote retention for the copy in Domain 3 is three weeks.

Figure 15-8 Storage lifecycle policy (in Domain 2) with intermediary duplication destination



Intermediary d
 destination ser
 for next duplic
 remote master

Remote master
 Domain 3

Different reten
 in Domain 3

Restoring from a backup at a remote master domain

While it is possible to restore a client directly by using the images in the remote master domain, do so only in a disaster recovery solution. In this discussion, a disaster recovery situation is one in which the originating domain no longer exists and clients must be recovered from the remote domain.

Table 15-5 Client restores in disaster recovery scenarios

Disaster recovery scenario	Does client exist?	Description
Scenario 1	Yes	Configure the client in another domain and restore directly to the client.
Scenario 2	No	Create the client in the recovery domain and restore directly to it. This is the most likely scenario.
Scenario 3	No	Perform an alternate client restore in the recovery domain.

The steps to recover the client are the same as any other client recovery. The actual steps depend on the client type, the storage type, and whether this is an alternate client restore.

For restores that use Granular Recovery Technology (GRT), an application instance must exist in the recovery domain. The application instance is required so that NetBackup has something to recover to.

For information on granular recovery, see the following topics and guides:

- See “Active Directory granular backups and recovery” on page 637.
- See “Enable granular recovery (policy attribute)” on page 547.
- See “Configuring a UNIX or Linux media server and Windows clients for backups and restores that use Granular Recovery Technology” on page 915.
- *NetBackup for Microsoft SharePoint Server Administrator's Guide*
- *NetBackup for Microsoft Exchange Server Administrator's Guide*

Reporting on duplication to remote master jobs

The Activity Monitor displays both the duplication job and the import job when duplicating to a remote master server domain.

Table 15-6 Duplication to remote master jobs in the Activity Monitor

Job type	Description
Duplication	<p>The job that duplicates a backup image to a remote master displays in the Activity Monitor as a duplication job. The Remote Master label displays in the Storage Unit column for this type of job.</p> <p>Similar to other duplication jobs, the job that replicates images to a remote master can work on multiple backup images in one instance.</p> <p>The detailed status for this job contains a list of the backup IDs that were duplicated.</p>
Import	<p>The job that imports an image into the remote master domain displays in the Activity Monitor as an import job. An import job can import multiple backup images in one instance. The detailed status for this job contains a list of processed backup IDs and a list of failed backup IDs.</p> <p>Note that a successful duplication does not confirm that the image was imported at the remote master.</p> <p>If the SLP names or data classifications are not the same in both domains, the import fails and NetBackup does not attempt to import the image again.</p> <p>Failed imports fail with a status 191 and appear in the Problems report when run on the target master server.</p> <p>The image is expired and deleted during catalog cleanup. Note that the originating domain (Domain 1) does not track failed imports.</p> <p>See “Running a report” on page 819.</p>

Configuring backups

- Chapter 16. Creating backup policies
- Chapter 17. Synthetic backups
- Chapter 18. Protecting the NetBackup catalog
- Chapter 19. About the NetBackup relational database
- Chapter 20. Managing backup images

Creating backup policies

This chapter includes the following topics:

- About the Policies utility
- Navigating in 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
- Adding or changing schedules in a policy
- Changing multiple policies at one time
- Copying or moving policy items to another policy or server
- Deleting schedules, backup selections, or clients from a policy
- Policy Attributes tab
- Schedules tab
- Schedule Attributes tab
- Start Window tab
- Excluding dates from a policy schedule
- Calendar Schedule tab
- How NetBackup determines which schedule to run next
- About schedule windows that span midnight
- How open schedules affect calendar-based and frequency-based schedules

- Runtime considerations
- About the Clients tab
- Backup Selections tab
- Disaster Recovery tab
- Creating a Vault policy
- Performing manual backups
- Active Directory granular backups and recovery

About the Policies utility

Backup policies provide the instructions that NetBackup follows to back up clients. Use the **Policies** utility to create NetBackup backup polices.

Backup policies provide the following instructions for a backup:

What type of client to back up.	See “Policy Attributes tab” on page 513.
Where to store the backup.	See “Policy Attributes tab” on page 513.
When and how frequently to perform the backup.	See “Schedules tab” on page 549.
Which clients to back up.	See “About the Clients tab” on page 595.
Which client files and directories to back up.	See “Backup Selections tab” on page 598.

Navigating in the Policies utility

The **Policies** utility offers a number of methods to view the configuration information for one or multiple policies.

To navigate the Policies utility

1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.

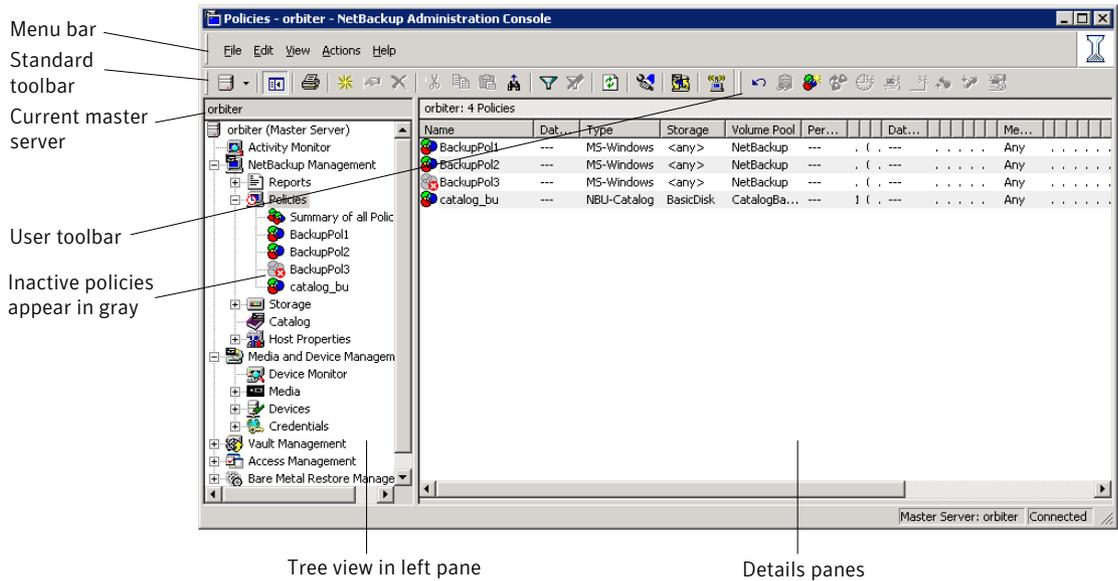
2 To see a hierarchical view of the policies on the selected master server: Select **Policies** in the left pane.

To display the policy details of a single policy: In the left pane, select a policy name. The policy details display in the right pane.

To open a policy: In the left or right pane, double-click on the policy name. The **Change Policy** dialog box opens.

To display information about all policies on the current master server: In the left pane, select **Summary of All Policies**. Click on the title bar of each horizontal pane to expand or collapse it.

Figure 16-1 NetBackup Policies utility



Planning for policies

Policy configuration is flexible enough to meet the various needs of all the clients in a NetBackup environment. To take advantage of this flexibility, take time to plan before starting to configure the policies in the **Policies** utility.

The following table outlines the steps to take to ensure that you get optimal results from your policy configurations.

Table 16-1 Steps for planning policies

Step	Action	Description
Step 1	Gather client information	<p>Gather the following information about each client:</p> <ul style="list-style-type: none"> ■ The client name. ■ The approximate number of files on each client to be backed up. ■ The typical file size of the files. <p>One client may be a file server that contains a large amount of data while the other clients are workstations. To avoid long backup times, include the file server in one policy and the workstations in another policy. It may be beneficial to create more than one policy for the file server.</p>
Step 2	Group the clients based on backup requirements	<p>Divide the clients into groups according to the different backup and archive requirements.</p> <p>The groups can be based on the type of work that the clients perform. Clients that are used for similar tasks generally have similar backup requirements. For example, most clients in an engineering department create the same types of files at similar levels of importance. In some instances, create a single policy for each group of clients. In other cases, subdivide the clients and include them in the separate policies that are based on their backup requirements.</p> <p>A backup policy can apply to one or more clients. Every client must be in at least one backup policy so that it can be backed up.</p>
Step 3	Consider the storage requirements	<p>The NetBackup environment may have some special storage requirements that the backup policies must accommodate.</p> <p>The storage unit and volume pool settings apply to all the files that are backed up by a policy. If files have special storage requirements, create separate policies for the files, even if other factors are the same, such as schedules.</p> <p>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.</p> <p>See “Example of one client in multiple policies” on page 506.</p>

Table 16-1 Steps for planning policies (*continued*)

Step	Action	Description
Step 4	Consider the backup schedule	<p>Create additional backup policies if the schedules in one policy do not accommodate all clients and files.</p> <p>Consider the following factors when deciding to create additional policies:</p> <ul style="list-style-type: none"> ■ Best times for backups to occur. To back up different clients on different schedules may require additional policies with different time schedules. For example, create different policies for night-shift and day-shift clients. ■ How frequently the files change. If some files change more frequently than others, the difference may be enough to warrant creating another policy with a different backup frequency. ■ How long backups need to be retained. Each schedule includes 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.
Step 5	Group clients by common attributes	<p>Create separate policies for the clients that require similar policy attribute settings.</p> <p>See “Policy attributes that affect how clients are grouped in policies” on page 507.</p>
Step 6	Maximize multiplexed backups	<p>Create separate policies as necessary to maximize the benefits of multiplexed backups.</p> <p>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.</p> <p>See “Media multiplexing (schedule attribute)” on page 572.</p>

Table 16-1 Steps for planning policies (*continued*)

Step	Action	Description
Step 7	Evaluate backup times	<p>Evaluate total backup times for each schedule and further subdivide policies to reduce backup times to an acceptable level.</p> <p>For example, if the backup of <code>D:\User</code>, <code>D:\h001</code>, and <code>E:\h002\Projects</code> on <code>client1</code> takes too much time, create a new policy for <code>E:\h002\Projects</code>.</p> <p>In addition to reducing the backup time for each policy, separate policies can reduce the total backup time for the server. NetBackup processes files within a backup selection list 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 per client host property is set to allow it.</p> <p>See “Global Attributes properties” on page 131.</p> <p>The Multiplexing and Allow multiple data streams policy attributes also allow backup policies to be processed in parallel.</p> <p>See “About multiplexing” on page 573.</p> <p>See “Allow multiple data streams (policy attribute)” on page 542.</p>

See “About the Policies utility” on page 502.

See “Policy Attributes tab” on page 513.

Example of one client in multiple policies

The following table shows that the files in two different subdirectories on one client can be stored in two different locations.

- Policy1 sends backups of `E:\h002\projects` to 8mm storage.
- Policy2 sends backups of `E:\h002\DevExp` and `E:\h002\DesDoc` to DLT storage.

Table 16-2 One client in multiple policies

Policies	Client	Files	Storage
Policy1	client1	<code>C:\</code> <code>D:\User</code> <code>D:\h001</code> <code>E:\h002\Projects</code>	8mm
Policy2	client1 client1	<code>E:\h002\DevExp</code> <code>E:\h002\DesDoc</code>	DLT

Policy attributes that affect how clients are grouped in policies

The following table lists the attributes that may determine which clients are grouped in the same policy.

Table 16-3 Policy attributes that affect how clients are grouped in policies

Attribute	Description
Policy Type	<p>Each client must be in a policy of the correct policy type. For example, Windows clients must be in a policy of a MS-Windows policy type.</p> <p>See “Policy type (policy attribute)” on page 514.</p>
Destination	<p>All of the data that the policy generates is sent to the same destination that is indicated in the policy. The data must share the same Data Classification, Policy storage, and Policy volume pool.</p> <p>See “Data classifications (policy attribute)” on page 517.</p> <p>See “Policy storage (policy attribute)” on page 518.</p> <p>See “Policy volume pool (policy attribute)” on page 519.</p>
Job Priority	<p>This attribute determines the priority for the backups of all of the clients in the policy.</p> <p>See “Job priority (policy attribute)” on page 526.</p>
Follow NFS	<p>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.</p> <p>See “Follow NFS (policy attribute)” on page 528.</p>
Cross mount points	<p>This attribute lets NetBackup cross file system boundaries for all clients in the policy.</p> <p>See “Cross mount points (policy attribute)” on page 533.</p>
Backup Network Drives	<p>This attribute lets NetBackup back up the files that all clients in the policy store on network drives. (Applies only to the MS-Windows policy type.)</p> <p>See “Backup Network Drives (policy attribute)” on page 529.</p>

Table 16-3 Policy attributes that affect how clients are grouped in policies
(continued)

Attribute	Description
Compression	This attribute indicates that all clients in the policy are to compress their backups before they send them to the server. Note that the time to compress can increase backup time and make it unsuitable to use for all clients. Consider creating a different policy for those clients. See “Compression (policy attribute)” on page 536.

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 by automatically choosing the best values for most configurations.

The **Backup Policy Configuration Wizard** uses a frequency-based schedule. It cannot be used to configure a calendar-based schedule. However, you can change the schedule to a calendar-based schedule after running the wizard.

See “Calendar Schedule tab” on page 584.

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**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Create a Backup Policy** to begin the wizard.
- 3 Click **Next** to start the wizard and follow the on-screen prompts.

Click **Help** on any wizard screen for assistance while running the wizard.

See “NetBackup naming conventions” on page 827.

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**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **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 827.
- 4 If necessary, uncheck **User Backup Policy Configuration Wizard**.
- 5 Click **OK**.
- 6 Configure the attributes, the schedules, the clients, and the backup selections for the new policy.
 - See “Policy Attributes tab” on page 513.
 - See “Schedules tab” on page 549.
 - See “About the Clients tab” on page 595.
 - See “Backup Selections tab” on page 598.

Adding or changing schedules in a policy

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.

Use the following procedure to add or change schedules in an existing NetBackup policy.

To add or change schedules in a policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 Perform one of the following actions:

To add a schedule

On the **Actions** menu, click **New > New Schedule**.

To change an existing schedule

In the right pane, double-click the schedule name.

- 4 Complete the entries in the **Attributes** tab, **Start Window** tab, **Exclude Dates** tab, and **Calendar Schedule** tab (when applicable).
See “Schedule Attributes tab” on page 549.
See “Start Window tab” on page 579.
See “Excluding dates from a policy schedule” on page 583.
See “Calendar Schedule tab” on page 584.
- 5 Click **OK**.
- 6 If this schedule is the last schedule, click **OK**.
To add more schedules, click **Add** and repeat step 4.

Changing multiple policies at one time

Use the following procedure to change more than one NetBackup policy at the same time.

To change multiple policies

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies > Summary of All Policies**.
- 2 In the right pane, hold down the Shift key and select either multiple policy names, multiple schedules, or multiple clients.
- 3 On the **Edit** menu, click **Change**.

In the dialog box that appears, the settings display in one of the following states:

A value	The setting has the same value for all selected policies, schedules, or clients.
No value	The attribute does not have the same value for all selected policies, schedules, or clients.
Checked	The attribute is active for all the selected policies, schedules, or clients.
Unchecked	The attribute is inactive on all the selected policies, schedules, or clients.
Gray checked	The attribute is set differently on all the selected policies, schedules, or clients.

- 4 Specify a value, enable or disable an attribute, or enter text for the attributes you want to change. Any change is applied to the field for every selected policy.

To enable an attribute for all selected policies Check the box.

To disable an attribute for all selected policies Uncheck the box.

To leave an attribute unchanged for all selected policies Set (or leave) the box to a gray check.

See “Policy Attributes tab” on page 513.

See “Schedule Attributes tab” on page 549.

- 5 See “About the Clients tab” on page 595.
- 6 Click **Cancel** to cancel changes, or click **OK** to apply all changes and close the dialog box.

Copying or moving policy items to another policy or server

You can copy or move entire policies, attributes, schedules, clients, and backup selections from one policy to another. You can also copy or move policy items from one server to another. The following is a description of which policy items can be copied or moved.

Use the following procedure to copy or move items from one policy to another.

To copy or move items from one policy to another

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the left pane, select the policy you want to copy or move information from.
- 3 In the right pane, select the attributes, schedules, clients, or backup selections you want to move or copy.
- 4 Do one of the following:

To copy an item On the **Edit** menu, click **Copy**.

- To move an item
- On the **Edit** menu, click **Cut**.
 - Click **Yes** when asked if you want to delete the selected item from the policy.
- 5 To copy or move the information to another server, do the following. Otherwise, continue to step 6.
- On the **File** menu, click **Change Server**.
 - In the **Change Server** dialog box, enter a host name or chose a host name from the list.
 - Click **OK**.
 - After the settings load for the other server, expand **NetBackup Management > Policies**.
- See “Accessing remote servers” on page 835.
- 6 In the left pane, select the policy where you want to copy or move the items to.
- 7 In the right pane, click the horizontal pane where you want to paste the contents of the clipboard: Attributes, Clients, Schedules, or Selections.
- To view the contents of the clipboard, on the **Edit** menu, click **Clipboard**.
- 8 On the **Edit** menu, click **Paste**.
- Any items with the same name are replaced with the contents of the clipboard. If the schedules do not match the policy type, the schedules are deleted or renamed. The action is indicated in a dialog box.

Note: If you copied items to another server, you must complete the rest of the configuration on the destination server for the configuration to work. For example, you must select a storage unit and volume pool at the destination server.

Deleting schedules, backup selections, or clients from a policy

Use the following procedure to delete schedules, backup selections, or clients from a NetBackup policy.

To delete a schedule, backup selections, or clients from a policy

- 1 In the NetBackup Administration Console, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the left pane, select the policy name.
- 3 In the right pane, select the item you want to delete.
- 4 On the **Edit** menu, click **Delete**.
- 5 Click **Yes** when asked if you want to delete the selected item from the policy.

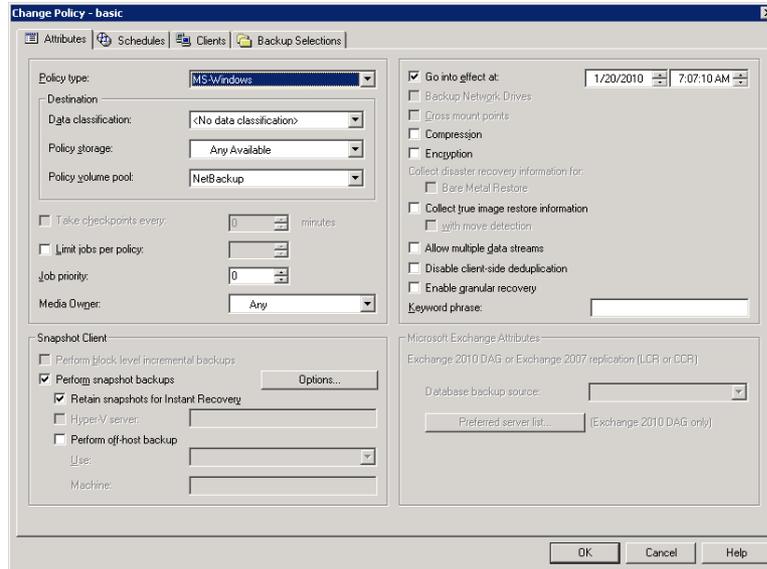
When a client is deleted from the client list, the NetBackup client software is not deleted or uninstalled from the client. Backups for the client can be recovered until the backups expire. Also, when a file is deleted from a backup selection list, the actual file is not deleted from the client.

Policy Attributes tab

Use the policy **Attributes** tab to configure backup settings when you add a new policy or change an existing policy. When you create a policy, you give the policy a name and select a policy type. The policy type you select typically depends on the type of client you want to back up. The number of policy types available varies depending on which NetBackup options are installed. Each policy type has a unique combination of attributes. Not all attributes apply to every policy type. When you select a policy type, the attributes that apply to that policy type are active. The unavailable attributes are grayed out.

Figure 16-2 shows the Attributes tab of a NetBackup policy.

Figure 16-2 Policy Attributes tab



The following topics describe the settings on the policy **Attributes** tab.

Policy type (policy attribute)

The **Policy type** attribute determines the purpose of the policy. Select a policy type from the list. The policy type you select typically depends on the type of client to be backed up. Some policy types are not used for client backups. NBU-Catalog is an example.

The list of policy types changes depending on the NetBackup options that have been installed. Each policy type offers a unique combination of attributes. When you select a policy type, only the attributes that apply to that policy type are active.

You can change the policy type of an existing policy. However, the schedules for the policy may become invalid. If the schedules become invalid, NetBackup displays a warning message and then deletes the invalid schedules or changes the schedules to an equivalent type.

Table 16-4 describes all the types of NetBackup policies.

Table 16-4 NetBackup policy types

Policy type	Description
AFS (UNIX only)	Use for the policies that back up only AFS file systems on clients. For information on setting up these policies, see "Using NetBackup with AFS," in the <i>NetBackup Administrator's Guide, Volume II</i> .
DataTools-SQL-BackTrack (UNIX only)	Use for the policies that contain only clients with the NetBackup SQL-BackTrack extension. For information on setting up this policy type, see the guide for this option.
DataStore	This policy type is reserved for use by Symantec or its partners to provide agents for new applications or databases.
DB2	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.
FlashBackup (UNIX only)	Applies only to NetBackup Enterprise Server. Use for the policies that contain only NetBackup FlashBackup clients on UNIX. This policy is available only when the Snapshot Client is installed. For information on setting up this policy type, see the <i>Snapshot Client Administrator's Guide</i> .
FlashBackup- Windows (Windows only)	Applies only to NetBackup Enterprise Server. Use for the policies that contain only FlashBackup-Windows clients on Windows. This policy is available only when the Snapshot Client is installed. For information on setting up this policy type, see the <i>NetBackup Snapshot Client Administrator's Guide</i> .
Informix-On-BAR (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.
Lotus-Notes	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.
MS-Exchange-Server	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.
MS-SharePoint (Windows only)	Use to configure a policy for NetBackup for SharePoint Portal Server.
MS-SQL-Server	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 16-4 NetBackup policy types (continued)

Policy type	Description
MS-Windows	<p>Use for the policies that contain only Windows clients of supported Windows operating system levels.</p> <p>Standard and MS-Windows policy types are the only policy types that support the following options:</p> <ul style="list-style-type: none"> ■ Checkpoint restart for backups or restores See “Take checkpoints every __ minutes (policy attribute)” on page 521. ■ Synthetic backups See “Policy type (policy attribute)” on page 514. ■ Collect disaster recovery information for Bare Metal Restore See “Collect disaster recovery information for Bare Metal Restore (policy attribute)” on page 538.
NBU-Catalog	Use for catalog backup jobs. Allows for a catalog backup while other jobs are running.
NCR-Teradata	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.
NDMP	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.
NetWare	Use for the policies that contain only NonTarget NetBackup Novell NetWare clients. (This version uses a Microsoft Windows interface.)
Oracle	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.
OS/2	Use for the policies that contain only OS/2 clients.
PureDisk-Export	Use for the policies that export data from PureDisk to NetBackup.
SAP	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.

Table 16-4 NetBackup policy types (*continued*)

Policy type	Description
Standard	<p>Use for the policies that contain any combination of the following:</p> <ul style="list-style-type: none"> ■ UNIX clients (including supported Mac clients), except those covered by specific products such as Oracle. ■ NetBackup Novell NetWare clients that have the target version of NetBackup software. <p>Standard and MS-Windows policy types are the only policy types that support the following options:</p> <ul style="list-style-type: none"> ■ Checkpoint restart for backups or restores ■ Synthetic backups ■ Collect disaster recovery information for Bare Metal Restore
Sybase	<p>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.</p>
Vault	<p>Use as a policy type to schedule and run a Vault job. This policy type is available only when Vault is licensed.</p>

For more details on off-host backups, refer to the *NetBackup Snapshot Client Administrator's Guide*.

Data classifications (policy attribute)

The **Data Classification** attribute specifies the classification of the storage lifecycle policy that stores the backup. For example, a backup with a gold classification must go to a storage unit with a gold data classification. By default, NetBackup provides four data classifications: platinum, gold, silver, and bronze.

This attribute is optional and applies only when the backup is to be written to a storage lifecycle policy. If the list displays **No data classification**, the policy uses the storage selection that is displayed in the **Policy storage** list. If a data classification is selected, all the images that the policy creates are tagged with the classification ID.

See “Storage Lifecycle Policy dialog box settings” on page 445.

See “Data Classification properties” on page 103.

See “Creating a Data Classification” on page 104.

See “About storage lifecycle policies” on page 443.

See “About associating backup data with a data classification” on page 447.

Policy storage (policy attribute)

The **Policy storage** attribute specifies the storage destination for the policy's data. Select a specific storage unit, storage lifecycle policy, or storage unit group from the list.

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. If a local device is not found or **Must use local drive** is not selected, NetBackup tries to find an available storage unit alphabetically.

When NetBackup looks for an available storage unit, it selects 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.

However, NetBackup makes an exception when a client is also a media server with locally-attached storage units. In that case, NetBackup selects the locally-attached storage units first.

If it is configured to do so, the storage unit or lifecycle policy determines which type of disk staging is used for the policy.

See “About staging backups” on page 419.

The list includes only those lifecycles that have 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 lacks a classification. Data classification is optional.

See “Global Attributes properties” on page 131.

Note: The storage unit that is selected on the **Schedule** tab, overrides the **Policy storage** attribute.

See “Override policy storage (schedule attribute)” on page 568.

See “Considerations for selecting a destination for Policy storage” on page 518.

Considerations for selecting a destination for Policy storage

Consider the following scenarios before selecting a destination from the **Policy storage** list on the policy **Attributes** tab.

Table 16-5

Scenario	Action
The site contains one storage unit, or there is no storage unit preference.	<p>Do one of the following:</p> <ul style="list-style-type: none"> ■ Specify Any Available for the Policy storage attribute. ■ Do not specify a storage unit at the schedule level. See “Override policy storage (schedule attribute)” on page 568. ■ Do not set all storage units to On demand only. NetBackup may not find an available storage unit for the backups. See “Changing storage unit settings” on page 389. See “On demand only storage unit setting” on page 409.
A specific storage unit is designated but the unit is unavailable.	Consider changing the destination to Any Available since backups cannot run for those policies and the schedules that require the unit.
Any Available is selected.	<p>Be aware that any basic disk storage unit that is not assigned to a storage group is considered available for disk spanning.</p> <p>See “Media properties” on page 153.</p>
You want to limit the storage units available to a policy.	<p>Do one of the following:</p> <ul style="list-style-type: none"> ■ Select a storage unit group that contains only the units you want the policy to use. ■ Limit the storage units by doing the following: <ul style="list-style-type: none"> ■ Create a volume pool that contains the volumes that are available only to the specific storage units. Disable Scratch pool for the volume pool. If Scratch pool is enabled, any storage unit has access to the volumes in the volume pool. See “Adding a volume pool” on page 314. See “About scratch volume pools” on page 313. ■ 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. ■ If the policy specifies a storage unit group, set the storage units within the group to On demand only to satisfy the policy requirement. See “Changing storage unit settings” on page 389. See “On demand only storage unit setting” on page 409.

Policy volume pool (policy attribute)

The **Policy volume pool** attribute specifies the default volume pool where the backups for the policy are stored. A volume pool is a set of media that is grouped for use by a single application. The volume pool is protected from access by other applications and users.

The available volume pools appear on the list. Whenever a new volume is required, it is allocated from the volume pool indicated.

If you select a volume pool on the **Schedule** tab, that selection overrides the **Policy volume pool** selection on the **Attributes** tab.

See “Override policy storage (schedule attribute)” on page 568.

See “Example of overriding the policy volume pool” on page 521.

The following table describes the default volume pools that NetBackup defines.

Table 16-6 Default volume pools defined by NetBackup

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.
CatalogBackup	This pool is selected by default for the NBU-Catalog policy type. It is used exclusively for online catalog backups. Catalogs are directed to a single, dedicated pool to facilitate faster catalog restores.

The following table describes the additional volume pools that are useful to create.

Table 16-7 Additional volume pools

Volume pool	Description
Scratch volume pool	Allows NetBackup to automatically transfer volumes when another volume pool does not have media available.
Auto volume pool	Used by automatic backups.
User volume pool	Used 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 312.

See “Adding a volume pool” on page 314.

See “About scratch volume pools” on page 313.

Example of overriding the policy volume pool

The following example shows how to override the policy volume pool from the policy **Schedule** tab. In this example, you change a policy named *Backup-Archive*. Until now, all schedules in the policy have used the *Backups* volume pool. Change the policy so that the user-archive schedule uses the *Archive* pool instead.

To override the Policy volume pool attribute

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**
- 2 In the left pane, select the *Backup-Archive* policy and on the **Edit** menu, click **Change**.
- 3 In the policy **Attributes** tab, on the **Policy volume pool** list, select *Backups*.
- 4 Click the **Schedules** tab.
- 5 Select the schedules that use the *Backups* volume pool, and click **Properties**.
- 6 Make sure that **Override policy volume pool** is unchecked, and click **OK** to save the change in the schedule..
- 7 Select the user-archive schedule that you want assigned to the *Archive* volume pool, and click **Properties**.
- 8 Check **Override policy volume pool**.
- 9 Underneath the check box, select *Archive* from the list.
- 10 Click **OK** to save the change in the schedule.
- 11 Click **OK** to save the change in the policy.

Take checkpoints every __ minutes (policy attribute)

By taking checkpoints during a backup, you can save time if the backup fails. 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.

The checkpoint frequency indicates how often NetBackup takes a checkpoint during a backup. The default is 15 minutes. The administrator determines checkpoint frequency on a policy-by-policy basis. When you select the checkpoint frequency, balance the loss of performance due to frequent checkpoints with the possible time lost when failed backups restart. If the frequency of checkpoints affects performance, increase the time between checkpoints.

Checkpoints are saved at file boundaries and point to the next file in the list. Checkpoint restart is only available after choosing the **MS-Windows** or **Standard** policy type. Check **Take checkpoints every __ minutes** to enable checkpoint

restart. When the box is checked, NetBackup takes checkpoints during a backup job at the frequency you specify. If the box is not checked, no checkpoints are taken and a failed backup restarts from the beginning of the job. Checkpoint restart can also be used for restore jobs.

See “Checkpoint restart for restore jobs” on page 524.

The **Global Attributes** property, **Schedule backup attempts**, indicates the number of times that NetBackup tries to restart a failed backup.

See “Global Attributes properties” on page 131.

Note: Checkpoints are saved at file boundaries and point to the next file in the list to be backed up. Checkpoints cannot occur in the middle of a file. After the file is backed up, the checkpoint is saved.

Note: Checkpoints are not taken for a user-archive backup. If a user-archive backup resumes, it restarts from the beginning.

In the following situations, NetBackup starts a new job instead of resuming an incomplete job:

- If a new job is due to run, or, for calendar-based scheduling, another run day has arrived.
- 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.

The following table describes the level of support for various policy attributes, storage, and clients for checkpoint restart. For an agent or option not listed, refer to the manual for that agent or option.

Table 16-8 Support for checkpoint restart

Item	Description
Basic disk staging	Checkpoint restart is supported for Stage I. Checkpoint restart is not supported for Stage II. See “About basic disk staging” on page 421. See “About staging backups” on page 419.

Table 16-8 Support for checkpoint restart (*continued*)

Item	Description
MS-Windows (policy type)	The following pertain to Windows clients: <ul style="list-style-type: none"> ■ Checkpoint restart is not supported for the backup selections that are 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 resumes. 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 (schedule attribute)	Checkpoint restart is supported for the policies that are configured to create multiple backup copies. See “Multiple copies (schedule attribute)” on page 562. The last failed copy that contains a checkpoint can be resumed if a copy is configured to allow other copies to continue the job if the copy fails and subsequent checkpoints occur.
NearStore storage units	Checkpoint restart is not supported. See “Disk storage unit considerations” on page 392. See the <i>NetBackup Administrator's Guide, Volume II</i> .
NetWare (policy type)	Checkpoint restart is not supported. NetWare clients can also use the Standard policy type, but that policy type does not support NetWare clients.
Snapshot Client (policy attribute)	Checkpoint restart is supported for use with local or alternate client backups. However, the following policy attributes are not supported: <ul style="list-style-type: none"> ■ Block Level Incremental Backups ■ Media Server Copy ■ Third-Party Copy Device ■ Instant Recovery backup See “Snapshot Client (policy attributes)” on page 548.
Standard (policy type)	Checkpoint restart is supported for UNIX clients.
Synthetic backups (schedule attribute)	Checkpoint restart is not supported. See “Synthetic backup (schedule attribute)” on page 559.

Checkpoint restart for restore jobs

Checkpoint restart for restore jobs saves time by letting NetBackup resume a failed restore job. The job resumes automatically from the start of the file that was last checkpointed rather than starting from the beginning of the entire restore job. NetBackup automatically takes checkpoints once every minute during a restore job.

The following host properties affect checkpoint restart for restore jobs.

Move restore job from incomplete state to done state This Clean-up host property indicates the number of days that a failed restore job can remain in an Incomplete state.

See “Clean-up properties” on page 75.

Restore retries This Universal Setting host property specifies the number of attempts that a client has to restore after a failure.

See “Universal Settings properties” on page 201.

Checkpoint restart for restore jobs has the following limitations:

- The restore restarts at the beginning of the last checkpointed file, not within the file.
- Only the backups that are created using **MS-Windows** or **Standard** policy types are supported.
- 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.

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.

Consider a situation in which a checkpointed restore that has no end date is suspended and then resumed. If a new backup occurs before the resume is initiated, the files from the new backup are included in the restore. For example, a user request the restore of a directory. The restore begins, but is suspended. The request is resumed the next day after another backup of the directory is performed. The files that are restored are from the latest backup.

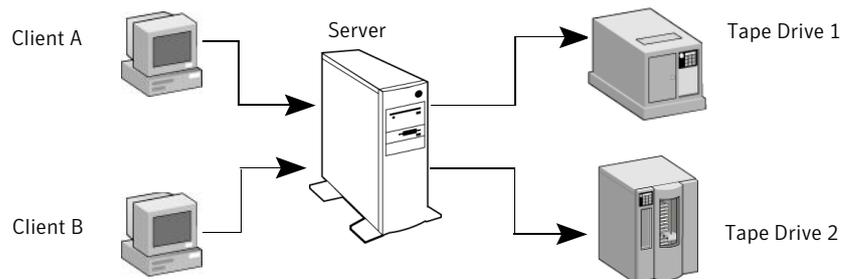
See “Take checkpoints every __ minutes (policy attribute)” on page 521.

Limit jobs per policy (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 box is unchecked, 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, check the box and specify a value from 1 to 999.

Figure 16-3 Limit jobs per policy attribute



Client A and Client B backups can occur concurrently and back up to different devices

Table 16-9 describes the factors that affect the number of concurrent backup jobs that NetBackup can perform.

Table 16-9 Factors affecting the number of concurrent backup jobs

Item	Description
Jobs from different policies	<p>The limit does not apply to concurrent jobs if the jobs are from different policies.</p> <p>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.</p>
Multiplexing	<p>If multiplexing is used, set the limit high enough to support the specified level of multiplexing.</p> <p>Lower values can limit multiplexing within a policy if jobs from different schedules exist within the policy. For example, the limit 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.</p>

Table 16-9 Factors affecting the number of concurrent backup jobs (*continued*)

Item	Description
Network load	<p>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 the Compression attribute.</p> <p>See “Compression (policy attribute)” on page 536.</p> <p>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.</p>
Number of storage devices available and multiplexing limits	<p>To process more than one backup job at a time, the configuration must include one of the following:</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p>
Parent job and children jobs	<p>Parent jobs do not count toward the limit. Only the children jobs count toward the limit.</p> <p>The following jobs produce a parent job and children jobs:</p> <ul style="list-style-type: none"> ■ Multistreamed jobs ■ Catalog backups ■ Snapshot Client snapshots ■ Bare Metal Restore jobs <p>See “About the Jobs tab” on page 766.</p>
Server speed	<p>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.</p>

Job priority (policy attribute)

The **Job priority** attribute specifies the priority that a policy has as it competes with other policies for resources. Enter a value from 0 to 99999. The higher the number, the greater the priority of the job. NetBackup assigns the first available resource to the policy with the highest priority.

In the **Default Job Priorities** host properties, you can set a job priority default for a job type.

See “Default Job Priorities properties” on page 105.

Media Owner (policy attribute)

The **Media Owner** attribute specifies which media server or server group should own the media that backup images for this policy are written to.

This attribute is active under the following conditions:

- A Media Manager storage unit is used.
- The **Policy storage** attribute is set to **Any Available**

You can specify the following for the **Media Owner**:

Any (default)	Allows NetBackup to select the media owner. NetBackup selects 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	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 list.

See “About media sharing” on page 317.

See “Configuring a server group” on page 210.

Go into effect at (policy attribute)

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 A.M., the policy does not run until that time or later. Use this attribute to configure a series of policies in advance of when the policies need to become active.

To activate the policy, check **Go into effect at**. The policy must be active for NetBackup to use the policy.

To deactivate a policy, uncheck the box. Inactive policies appear are unavailable in the Administration Console. Inactive policies remain on the list of policies in the left pane of the NetBackup Administration Console. To resume backups, recheck the box. Make sure that the date and time are 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. Uncheck the box to deactivate the policy.

See “Creating a catalog archiving policy” on page 680.

Follow NFS (policy attribute)

The **Follow NFS** (Network File System) attribute specifies whether NetBackup is to back up or archive any NFS-mounted files. These files are named in the backup selection list or by the user, in the case of a user backup or archive. Uncheck the box to prevent the backup or archive of NFS-mounted files.

Note: This attribute applies only to UNIX clients in certain policy types. NetBackup allows it to be selected in those instances only.

This attribute 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 205.

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.

Consider the following before enabling this attribute:

Table 16-10 Issues that affect Follow NFS

Item	Description
Cross mount points (policy attribute)	<p>The behavior of Follow NFS can vary depending on how it is used in combination with Cross mount points.</p> <p>See “Examples of using Cross mount points and Follow NFS in combination” on page 534.</p> <p>See “Cross mount points (policy attribute)” on page 533.</p>
Raw partitions	<p>This attribute 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 computers that use NFS mounts to access the raw partitions. The devices are not accessible on other computers through NFS.</p> <p>Note: NetBackup does not support raw partition backups on unformatted partitions.</p>
Automounted directories	<p>This attribute 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.</p>

See “Examples of using Cross mount points and Follow NFS in combination” on page 534.

Backup Network Drives (policy 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. For a computer that is not a NetBackup client, the preferred method for backing up data is to use UNC paths. UNC paths are more precise and indicate exactly what should be backed up.

When you use **Backup Network Drives** or UNC paths, 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.

This attribute must be enabled for the policies that 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.

Example of using UNC paths to back up a shared folder

The following example gives the steps for backing up a shared folder using a UNC path. The procedure backs up the folder *TestData* on *win_PC* through *win_client*. Consult the following descriptions before you review the example.

- master1* NetBackup master server
- win_client* Windows NetBackup client
- win_PC* Windows computer (not necessarily a NetBackup client)
- TestData* A shared folder on *win_PC*

Table 16-11 Using UNC paths to back up a shared folder on *win_PC*

Step	Action	Description
Step 1	Create a policy	On <i>master1</i> create a policy for <i>win_client</i> . See “Creating a policy using the Backup Policy Configuration Wizard” on page 508. See “Creating a policy without using the Backup Policy Configuration Wizard” on page 508.
Step 2	Add the folder name to the policy	Add <code>\\win_PC\TestData</code> to the file list of the policy. This step is not necessary if the policy is only used for user-directed backups. See “Adding backup selections to a policy” on page 599.
Step 3	Configure the NetBackup Client Service	Do the following: <ul style="list-style-type: none"> ■ On <i>win_client</i>, change the NetBackup Client Service 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.

Table 16-11 Using UNC paths to back up a shared folder on *win_PC* (continued)

Step	Action	Description
Step 4	Perform a backup	<p>Do the following to perform a user backup:</p> <ul style="list-style-type: none"> ■ In the NetBackup Administration Console, on the File menu, click Backup, Archive, and Restore. ■ In the Backup, Archive, and Restore client interface, on the File menu, click Select Files and Folders to Backup. ■ In the screen that appears, in the left pane, expand Network. ■ In the left pane, locate the computer <i>win_PC</i> and expand it until you can select the directory <i>TestData</i>. ■ On the Actions menu, click Backup. ■ In the Backup Files dialog box, specify the necessary options. ■ Click Start Backup. <p>For more information, see the online Help in the Backup, Archive, and Restore client interface.</p> <p>Backups run as scheduled or when a manual backup is performed.</p>

Example of using Backup Network Drives (policy attribute) to back up a shared folder

The following example gives the steps for backing up a shared folder using the **Backup Network Drives** policy attribute. The procedure backs up the folder *share* on *win_PC* through *win_client*. Consult the following descriptions before you review the example.

- master1* NetBackup master server
- win_client* Windows NetBackup client
- win_PC* Windows computer (not necessarily a NetBackup client)
- share* A shared folder on *win_PC*

Table 16-12 Using Backup Network Drives to back up a shared folder on *win_PC*

Step	Action	Description
Step 1	Create a policy	<p>On <i>master1</i> create a policy for <i>win_client</i>, and check Backup network drives in the policy attributes tab.</p> <p>See “Creating a policy using the Backup Policy Configuration Wizard” on page 508.</p> <p>See “Creating a policy without using the Backup Policy Configuration Wizard” on page 508.</p> <p>See “Backup Network Drives (policy attribute)” on page 529.</p>
Step 2	Configure the NetBackup Client Service	<p>Do the following:</p> <ul style="list-style-type: none"> ■ On <i>win_client</i>, change the NetBackup Client Service 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.
Step 3	Create a batch file	<p>Create a batch file <code>bpstart_notify.bat</code> that does the following:</p> <ul style="list-style-type: none"> ■ Maps a drive on <i>win_client</i> to <code>\\win_PC\share</code>. ■ Includes the following command (where X: is the mapped drive letter): <pre>net use X: \\win_PC\share</pre>
Step 4	Perform a backup	<p>Do the following to perform a user backup:</p> <ul style="list-style-type: none"> ■ In the NetBackup Administration Console, on the File menu, click Backup, Archive, and Restore ■ In the Backup, Archive, and Restore client interface, on the File menu, click Select Files and Folders to Backup. ■ In the screen that appears, in the left pane, expand Network. ■ In the left pane, locate the computer <i>win_PC</i> and expand it until you can select the directory <i>share</i>. ■ On the Actions menu, click Backup. ■ In the Backup Files dialog box, specify the necessary options. ■ Click Start Backup. <p>For more information, see the online Help in the Backup, Archive, and Restore client interface.</p> <p>Backups run as scheduled or when a manual backup is performed.</p>

Cross mount points (policy attribute)

The **Cross mount points** attribute controls whether NetBackup crosses file system boundaries to back up or archive all files and directories in the selected path. 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. This attribute is supported on computers running UNIX or Windows 2003 and later.

When this attribute is disabled, only the files that are in the same file system as the selected file path are backed up. By disabling, you also 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.)

In some cases, consider creating separate policies for the backups that cross mount points and those that do not. For example, in one policy, disable **Cross mount points** and include `root (/)` in the backup selection list. As a result, 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** and include `root (/)` in the backup selection list. As a result, all the data on the client is backed up.

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.

The following table lists items to consider when you use this policy attribute.

Table 16-13 Considerations for Cross mount points (policy attribute)

Item	Description
Follow NFS (policy attribute)	<p>The behavior of Cross mount points can vary depending on how it is used in combination with Follow NFS.</p> <p>See “Examples of using Cross mount points and Follow NFS in combination” on page 534.</p> <p>See “Follow NFS (policy attribute)” on page 528.</p>
Backup selection entries	<p>The following backup selection entries behave in the same manner on both UNIX and Windows systems when the Cross mount points attribute is used::</p> <p>/</p> <p>:\</p> <p>*:\</p> <p>Note: Do not use the Cross mount points attribute in policies on UNIX systems where you use the ALL_LOCAL_DRIVES directive in the backup selection list.</p>

Table 16-13 Considerations for Cross mount points (policy attribute) (continued)

Item	Description
UNIX raw partitions	This attribute 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 when this attribute is enabled.
ALL_LOCAL_DRIVES directive	Do not use this attribute in policies on UNIX systems where you use the ALL_LOCAL_DRIVES directive in the backup selection list.
Mount points to disk storage	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.

Examples of using Cross mount points and Follow NFS in combination

By using **Cross mount points** and **Follow NFS** in combination, you can get a variety of results. Table 16-14 summarizes the possible results.

Table 16-14 Results of using Cross mount point and Follow NFS in combination

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.

Example 1 and Example 2 assume that the client disks are partitioned as shown in Figure 16-4.

Figure 16-4 Example configuration of client disks

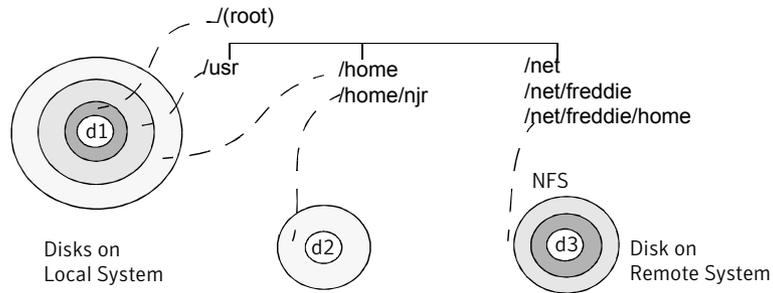


Table 16-15 Legend

Disks	Description
d1	Contains / (root), /usr, and /home in separate partitions.
d2	Contains a file system named /home/njr. Mounted on /home, which is a partition on d1.
d3	Contains a directory named /net/freddie/home that is NFS-mounted on /net/freddie

Example 1:

- **Cross mount points** and **Follow NFS** are not selected.
- The backup selection list contains the following entry:
`//usr/home`
- In this case, NetBackup considers only the directories and files that are in the same file system as the backup selection list entry. It does not back up `/home/njr` or `/net/freddie/home`.

Example 2:

- **Cross mount points** and **Follow NFS** are selected.
- The backup selection list only contains a forward slash:
`/`
- 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 (policy attribute)

The **Compression** attribute specifies that the backups use the software compression that is possible, based on the policy type. Check the box to enable compression. By default, compression is disabled.

Compression reduces the size of a backup by reducing the size of files in the backup. In turn, the smaller backup size decreases the number 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. However, 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.

See “About multiplexing” on page 573.

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 the space that 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.)

Table 16-16 describes factors to consider when you choose to use **Compression**.

Table 16-16 Considerations for Compression

Item	Description
Data types that compress well	<p>Programs, ASCII files, and unstripped binaries (typically 40% of the original size).</p> <p>Best-case compression: Files that are composed of the strings that repeat can sometimes be compressed to 1% of their original size.</p>
Data types that do not compress well	<p>Stripped binaries (usually 60% of original size).</p> <p>Worst-case compression: Files that are already compressed become slightly larger if compressed again.</p>
Effect of file size	<p>File size has no effect on the amount of compression. However, it takes longer to compress many small files than a single large one.</p>
Client resources that are required	<p>Compression requires client computer processing unit time and as much memory as the administrator configures.</p>
Effect on client performance	<p>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.</p>
Files that are not compressed	<p>NetBackup does not compress the following files:</p> <ul style="list-style-type: none"> ■ Files that are equal to or less than 512 bytes, because that is the tar block size. ■ On UNIX clients, files with the following suffixes: <pre style="margin-left: 40px;"> .arc .gz .iff .sit.bin .arj .hqx .pit .tiff .au .hqx.bin .pit.bin .Y .cpt .jpeg .scf .zip .cpt.bin .jpg .sea .zom .F .lha .sea.bin .zoo .F3B .lzh .sit .z .gif .pak </pre> <ul style="list-style-type: none"> ■ 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 using with storage units with SIS capabilities	<p>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.</p>

Encryption (policy 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 (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.

Collect disaster recovery information for Bare Metal Restore (policy 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.

Only policy types **MS-Windows** (for Windows clients) and **Standard** (for UNIX clients) support this policy attribute. This attribute is enabled by default when one of these policy types is used to create a policy on a master server that is licensed for BMR.

For more information, see the *Bare Metal Restore Administrator's Guide for UNIX, Windows, and Linux*.

Collect true image restore information (policy attribute) with and without move detection

The **Collect true image restore information** attribute specifies whether the policy collects the information necessary to perform a true image restore. A true image restore (TIR) restores the contents of a directory to reflect the contents of the

directory 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.

For true image incremental backups, enable **With move detection** to include the files that were moved, renamed, or newly installed in the directories. These files may be 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.

NetBackup detects changes by comparing path names and inode numbers with those from the previous full or incremental backup. If either the name or an inode number is new or changed, the file or directory is backed up. NetBackup begins to collect the information 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.

Note: With move detection must be enabled to create a synthetic backup. It must also be enabled to back up data to the NearStore disk storage units that use the File System Export option.

See “Synthetic backup (schedule attribute)” on page 559.

See the *NetBackup Administrator’s Guide, Volume II* for more information about configuring NearStore storage units.

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`. 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\`. 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.

Example of true image restores

The following table lists the files that were backed up in the `C:\user\doc` directory during a series of backups between 12/01/2009 and 12/04/2009. **Collect true image restore information** was turned on for the policy that performed the backups.

Table 16-17 Sample backups taken before a true image restore

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.

A restore of the 12/04/2009 version of the `C:\user\doc` directory produces the following results:

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
```

After a true image restore The restored directory 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.

Consider the following points to use either **Collect true image restore** or **Collect true image restore with move detection**:

- NetBackup collects additional information for the incremental backups that collect true image restore information. Policies that use move detection require even more space.
- Incremental backups are slower for a policy in which true image restore information is collected.
- Configure the period of time that NetBackup retains the true image restore information. Set the **Keep true image restoration (TIR) information** property in the **Clean-up** properties dialog box. See “Clean-up properties” on page 75.
- Only directories can be listed and selected. In true image restore mode, the client interface does not display individual files. Refer to the online Help in

the **Backup, Archive, and Restore** client interface for more information on true image restores.

- A true image restore preserves the files that are currently in the directory but were not present when the backup was completed. If you created a file `file5` after an incremental backup on 12/04/2009 but before a restore, the contents of the restored directory would be as follows:

```
file1  
file2  
file4  
file5
```

Allow multiple data streams (policy 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. Because 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.

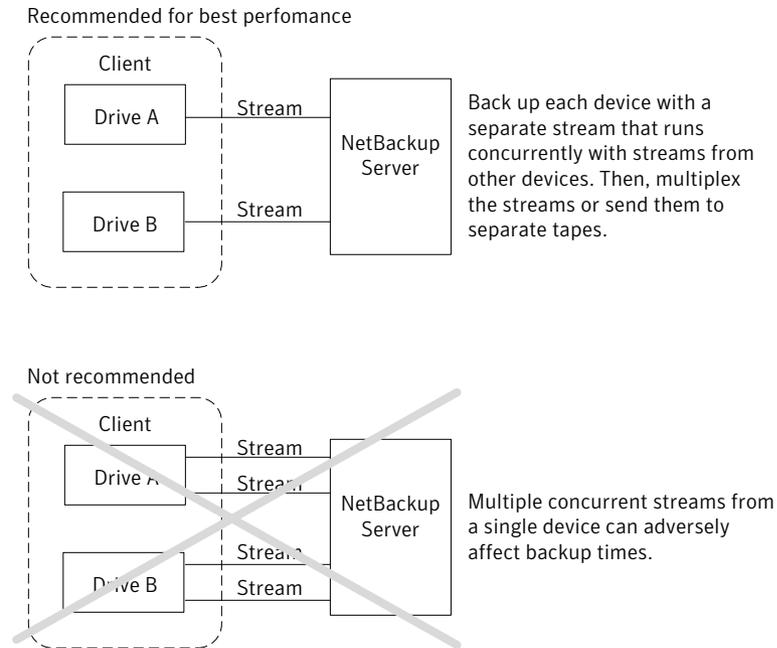
Note: If this attribute is enabled, and a file system is in a client's exclude list, a NetBackup job appears in the **Activity Monitor** for the excluded file system. However, no files in the excluded file system are backed up by the job.

The following table describes the reasons to use multiple data streams.

Table 16-18 Reasons to use multiple data streams

Reason	Description
To reduce backup time	<p>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.</p> <p>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.</p> <p>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.</p> <p>Figure 16-5 shows why multiple concurrent streams from a single device are not recommended.</p>
To reduce retry time for backup failures	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
To reduce administration by running more backups with fewer policies	<p>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.</p>

Figure 16-5 Multiple stream recommendations



The following table describes the aspects of multiple data streams that are adjustable.

Table 16-19 Adjustable aspects of multiple data streams

Item	Description
The total number of streams	<p>The backup selection list determines the total number of streams that are started. The NEW_STREAM directive lets you configure a fixed number of streams, or you can allow the client dynamically define the streams.</p> <p>See “About the directives on the Backup Selections list” on page 620.</p> <p>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. Backup times are affected because the device heads must move between the tracks that contain files for the respective streams.</p>

Table 16-19 Adjustable aspects of multiple data streams (*continued*)

Item	Description
The number of streams that run concurrently	<p>The following factors determine the number of streams that can run concurrently for a policy or client:</p> <ul style="list-style-type: none"> ■ Number of the drives that are available. ■ Maximum concurrent jobs settings for the policy and client. ■ Storage unit and schedule multiplexing limit. <p>Each storage unit and each schedule have a maximum multiplexing setting. The lower of the two settings is the limit for a specific schedule and storage unit. The maximum streams are limited to the sum of the multiplexing limits for all drives available in the storage unit and schedule combinations.</p> <p>For example, assume that two storage units have one drive in each. Multiplexing on storage unit 1 is set to 3 and multiplexing on storage unit 2 is set to 5. If multiplexing is set to 5 or greater in the schedules, then 8 streams can run concurrently.</p> <p>See “Media multiplexing (schedule attribute)” on page 572.</p>

The maximum jobs settings limit the maximum number of streams as follows:

Table 16-20 Job settings that limit the maximum number of streams

Item	Access method
<p>Maximum jobs per client (host property)</p>	<ul style="list-style-type: none"> ■ In the left pane, expand NetBackup Management > Host Properties. ■ Select Master Servers, and in the right pane, double-click the master server you want to modify. ■ In the properties dialog box, in the left pane, click Global Attributes. <p>See “Global Attributes properties” on page 131.</p> <p>See “Media multiplexing (schedule attribute)” on page 572.</p>
<p>Limit jobs per policy (policy attribute)</p>	<ul style="list-style-type: none"> ■ In the left pane, expand NetBackup Management > Policies. ■ In the right pane, double-click a policy you want to modify. <p>See “Limit jobs per policy (policy attribute)” on page 525.</p>

Table 16-20 Job settings that limit the maximum number of streams (continued)

Item	Access method
Maximum data streams (host property)	<ul style="list-style-type: none"> ■ In the left pane, expand NetBackup Management > Host Properties. ■ Select Master Servers, and in the right pane, double-click the master server you want to modify. ■ In the properties dialog box, in the left pane, click Client Attributes. <p>See “General tab of the Client Attributes properties” on page 80.</p>

Job settings also affect the maximum number of streams. The following table describes the interdependency of these settings.

Table 16-21 Interdependency of job settings

Item	Description
Maximum data streams property is disabled.	NetBackup uses the value that is specified by either Maximum jobs per client or Limit jobs per policy , whichever is lower.
Maximum data streams property is enabled.	NetBackup ignores Maximum jobs per client . Instead, NetBackup uses the value that is specified by either Maximum data streams or Limit jobs per policy , whichever is lower.

See “About the directives on the Backup Selections list” on page 620.

Disable client-side deduplication (policy attribute)

The **Disable client-side deduplication** attribute appears only if the NetBackup Deduplication Option license key is active.

The clients that are configured for client direct backup behave as follows when this attribute is enabled or disabled:

- | | |
|----------|--|
| Enabled | 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. |
| Disabled | 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 are bypassed. |

The **Deduplication** property configures clients for client direct deduplication. The **Disable client-side deduplication** policy attribute overrides the **Deduplication** property. The **Deduplication** property is found on the **General** tab of the **Client Attributes** host properties.

See “Where deduplication should occur” on page 83.

See the *NetBackup Deduplication Guide*.

Enable granular recovery (policy 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*

Note: In IPv6-enabled NetBackup 7.1 environments, granular recovery is not supported for Exchange Server or SharePoint Server.

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 637.

Keyword phrase (policy 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. By default, the keyword phrase is blank.

Clients can also specify a keyword phrase for a user backup or archive. A user keyword phrase overrides the policy phrase.

Snapshot Client (policy attributes)

The **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 (policy attributes)

The **Microsoft Exchange** attributes let you indicate the database backup source to use for the following:

- Exchange 2010 Database Availability Group
- Exchange 2007 replication backup

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

Schedules tab

The schedules that are defined on the **Schedules** tab determine when backups occur for the selected policy. 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:

Attributes tab	Schedule the time and frequency at which a task runs, along with other scheduled attributes. See “Schedule Attributes tab” on page 549.
Start Window tab	Schedule the time of each day that a task runs. See “Start Window tab” on page 579.
Exclude Dates tab	Indicate the dates that a task should not run. See “Excluding dates from a policy schedule” on page 583.
Calendar Schedule tab	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 .) See “Calendar Schedule tab” on page 584.

Schedule Attributes tab

The schedule **Attributes** tab contains both schedule information and other configuration options, beyond when the job is to run.

The following topic describes the options on the **Attributes** tab for schedules.

Name (schedule attribute)

Specify a name for the schedule by typing it in the **Name** attribute. The schedule name appears on screens and messages about the schedule.

See “NetBackup naming conventions” on page 827.

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 419.

Type of backup (schedule attribute)

The **Type of backup** attribute specifies the type of backup that the schedule controls. Select a backup type from the 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.

Table 16-22 and Table 16-23 describe the types of backups available in NetBackup. Table 16-22 describes the types of backups that come standard with NetBackup.

Table 16-22 Standard backup types

Item	Description
Full Backup	<p>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. Use this option if you configure a policy for a raw partition backup (formatted partitions only).</p>
Cumulative Incremental Backup	<p>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 “About incremental backups” on page 552.</p>

Table 16-22 Standard backup types (continued)

Item	Description
Differential Incremental Backup	<p>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 “About incremental backups” on page 552.</p>
User Backup	<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 Start Window tab.</p> <p>Use this backup type for a catalog archive.</p> <p>See “Considerations for user schedules” on page 557.</p> <p>See “Creating a catalog archiving policy” on page 680.</p>
User Archive	<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 Start Window tab.</p> <p>Note: The NetBackup administrator should make sure that a full backup of the client exists before a user archives files from the client.</p>

Table 16-23 describes the types of backups that are available when you install additional agents and options.

Table 16-23 Additional backup types

Item	Description
Application Backup	<p>Applies to all database agent clients.</p> <p>For more information, see the NetBackup guide that came with the product.</p>
Automatic Backup	<p>Applies to all database agent clients, except NetBackup for Informix and NetBackup for Oracle.</p> <p>For more information, see the NetBackup guide for the database product.</p>

Table 16-23 Additional backup types (continued)

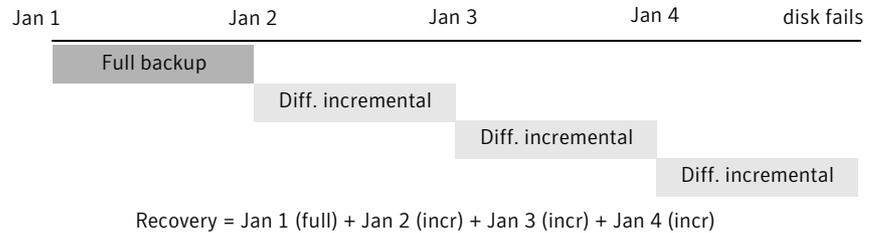
Item	Description
Automatic Incremental Backup	Applies only to NetBackup for Informix clients. For more information, see the <i>NetBackup for Informix Administrator's Guide</i> .
Automatic Cumulative Incremental Backup	Applies only to NetBackup for Oracle clients. For more information, see the <i>NetBackup for Oracle Administrator's Guide</i> .
Automatic Differential Incremental Backup	An automatic differential incremental backup applies only to NetBackup for Oracle clients. For more information, see the <i>NetBackup for Oracle Administrator's Guide</i> .
Automatic Full Backup	Applies only to NetBackup for Informix and NetBackup for Oracle clients. For more information, see the <i>NetBackup for Informix Administrator's Guide</i> or the <i>NetBackup for Oracle Administrator's Guide</i> .
Automatic Vault	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. See "Creating a Vault policy" on page 634.
Vault Catalog Backup	Use when the schedule is for a catalog backup policy that Vault uses. Available only when Vault is licensed. If this type is selected, you must configure one of the two schedule attribute combinations or the schedule cannot be saved: <ul style="list-style-type: none"> ■ Check and configure Multiple copies, or ■ Check Override policy storage selection, Override policy volume pool, and specify the Retention. <p>Note: 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 16-6 shows how data is included in a series of full and differential incremental backups between January 1 and January 4.

Figure 16-6 Full and differential incremental example



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 16-7 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 16-7 Full and cumulative incremental example

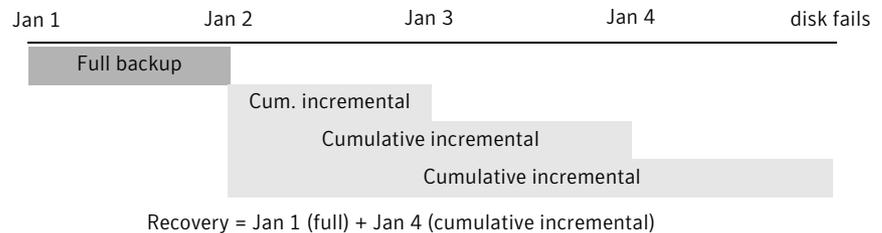


Table 16-24 describes how to determine the retention of differential and cumulative incremental backups to prevent a gap in backup coverage.

Table 16-24 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 16-25 compares the advantages and disadvantages of using differential or cumulative incremental backups based on possible backup and restore times.

Table 16-25 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 backup 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 16-26 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 schedules that are described in Table 16-26 result in the following series of backups:



The example produces the following results:

- 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 when Windows files are due for backup

On Windows clients, NetBackup performs the incremental backups when the **Perform incrementals based on archive bit** setting is enabled. 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 the 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.

Full Backup

NetBackup backs up files regardless of the state of their archive bit. After a full backup, the archive bit is always cleared.

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 backup back up only the files that changed since the previous full or differential incremental backup.

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. The backup types use the datetime stamp differently.

Full Backup NetBackup backs up files regardless of the datetime stamp.

Differential Incremental Backup NetBackup compares the datetime stamp of the file against the last full or incremental backup.

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.

How NetBackup determines when UNIX files are due for backup

Incremental backups on UNIX clients consider 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*).

The following types of time are associated with UNIX files and directories:

`mtime` The file modification time. The file system updates the `mtime` for a file or directory each time the file is modified. An application can save the `mtime` of the file before it modifies it. The application then resets it with the `utime(2)` system call.

`atime` The file access time. The file system updates the `atime` for a file or directory each time the file is accessed (read or write). An application can save the `atime` of the file before it accesses it. The application then resets it with the `utime(2)` system call.

`ctime` The inode change time. The `ctime` 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 `ctime` for a file or directory cannot be saved before a change, and then reset after a change. The `ctime` of a file or a directory changes when the `mtime` and `atime` (changes with the `utime(2)` 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.

Considerations for user schedules

In order for users to perform backups and archives, an administrator must create a schedule that allows user backups.

User backup schedules and user archive schedules can be included in a policy that contains automatic backup schedules. If you create separate policies for user backups or user 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.

To use a specific policy or schedule for user backups or user archives, perform the tasks that are specified for each client type:

Table 16-27 Tasks for specifying a policy or schedule for user backups or user archives

Client type	Task
Microsoft Windows clients	<ul style="list-style-type: none"> ■ Start the Backup, Archive, and Restore client interface. ■ On the File menu, click NetBackup Client Properties ■ Select the Backups tab, and specify the backup policy and backup schedule.
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).
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.

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 schedules for user backups and user archives

To plan schedules for user backups and user archives, consider the following:

Automatic backups	<p>If possible, do not permit user backups and user 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.)</p> <p>See “Limit jobs per policy (policy attribute)” on page 525.</p> <p>See “Global Attributes properties” on page 131.</p> <p>If the automatic backup continues to run, the user job may miss the backup window depending on how the limiting settings are configured. On the other hand, user jobs can delay automatic backups and can cause backups to miss the backup window.</p>
Storage units	Use a different storage unit to eliminate conflicts with automatic backups.
Volume pools	Use a different volume pool to manage the media separate from the automatic backup media.
Retention periods	<p>Consider setting the retention period for archives to infinite, since the disk copy of the files is deleted.</p> <p>Note: If the retention period expires for a backup, it can be difficult or impossible to restore the archives or backups.</p>

Synthetic backup (schedule 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, to disk storage units, or to a combination of both.

Calendar (schedule attribute)

Calendar-based schedules allow administrators to select specific days to run a policy. Select **Calendar** to display the **Calendar Schedule** tab.

See “Calendar Schedule tab” on page 584.

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.)

Enable **Retries allowed after runday** to have NetBackup attempt to complete the schedule until the backup is successful. With this attribute enabled, the schedule attempts to run, even after a specified runday has passed.

Frequency (schedule attribute)

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.

To set the frequency, select a frequency value from the list. The frequency can be hours, days, or weeks.

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.)

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 tries to run the job at the same time every other day. (NetBackup checks if the frequency is divisible by 24 hours.) If the interval is not divisible by 24, NetBackup does not attempt to run the job at about the same time of day. Instead, NetBackup tries to run the job at the indicated interval after the last successful backup. (For example, 52 hours later.)

Note: **Frequency** does not apply to user schedules because the user can perform a backup or archive whenever the time window is open.

About backup frequency

To determine backup frequency, consider how often data changes. For example, determine if files change several times a day, once a day, 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 important data changes many times during the day and the changes would be difficult to reconstruct.

Daily backups are usually the 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.
- Decrease 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.

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 that has a backup frequency of one month takes priority over a schedule with a backup frequency of two weeks.
- When two schedules are each due to run, the schedule with the schedule name that is first in alphabetical order 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.

NetBackup prioritizes the example schedules in the following order:

Table 16-28 Examples of schedule frequency and priority

Schedule Name	Frequency	Priority
monthly_full	One month	First
weekly_full	One week	Second
daily_incremental	One day	Third

Instant Recovery (schedule attribute)

The **Instant Recovery** attributes 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.

See “Snapshot Client (policy attributes)” on page 548.

This attribute has two options.

Snapshots and copy snapshots to a storage unit The snapshot persists on the client volume with a backup copy made to the storage unit on the media server.

Snapshots only 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.

Transaction logs are not truncated at the end of the backup.

Multiple copies (schedule 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, Inline Tape Copy, or ITC.)

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. **Maximum backup copies** is a **Global Attributes** host property.

See “Global Attributes properties” on page 131.

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.

About configuring for 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 storage unit setting” on page 404.
See “Maximum concurrent write drives storage unit setting” on page 403.
- You can use a storage lifecycle policy to create multiple copies. However, the number of destinations in the lifecycle cannot exceed the **Maximum backup copies** setting in the **Global Attributes** host properties. The lifecycle cannot be saved until the destinations are decreased, or until the **Maximum backup copies** setting is increased.
See “Global Attributes properties” on page 131.

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
- Synthetic backups
- Storage lifecycle policies
Storage lifecycle policies offer their own method to create multiple copies.
See “About writing multiple copies using a storage lifecycle policy” on page 467.

Multiple copies can also be configured for a relocation schedule. A relocation schedule is created as part of basic disk staging storage unit configuration. The **Maximum backup copies** property must be set to include an additional copy beyond the number of copies to be created in the **Configure Multiple Copies** dialog box. For example, to create four copies in the **Configure 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 backup to be one copy. NetBackup automatically counts this copy against the **Maximum backup copies** value.

Configure Multiple Copies dialog box

The **Configure Multiple Copies** dialog box contains the following options:

Table 16-29 Configure Multiple Copies dialog box

Field	Description
Copies	<p>NetBackup can create up to four copies of a backup simultaneously. The storage units must be on the same media server and there must be sufficient resources available for each copy.</p> <p>The maximum is four, or the number of the copies that are specified by the Maximum backup copies Global Attributes host property, whichever is smaller. The Maximum backup copies property specifies the total number of backup copies that can exist in the NetBackup catalog (2 through 10).</p> <p>See “Global Attributes properties” on page 131.</p>
Priority of duplication job	Indicate the priority that the duplication job (based on this schedule) has over other jobs in the queue (0 to 99999).
Primary copy	<p>Copy 1 is the primary copy. If Copy 1 fails for some reason, the first successful copy is the primary copy.</p> <p>See “Promoting a copy to a primary copy” on page 741.</p>
Storage unit	Specify the storage unit where each copy is to be stored. If a Media Manager storage unit has multiple drives, you can use it for both the source and the destination. To let NetBackup decide at runtime, select Any Available .
Volume pool	Indicate where each copy is to be stored.
Retention schedule	<p>Specify how long NetBackup retains the backups.</p> <p>See “Retention (schedule attribute)” on page 569.</p>

Table 16-29 Configure Multiple Copies dialog box (continued)

Field	Description
If this copy fails	<p>In the event that the copy does not complete, select whether you want the entire job to fail (fail all copies), or whether you want the remaining copies to continue.</p> <p>Regardless of how the fail or continue flag is set, all the copy jobs wait in the queue until resources are available for all copies. The first job does not start until the copies have resources.</p> <p>If a copy is configured to allow other copies to continue the job if the copy fails, and if Checkpoint restart for backup jobs is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.</p> <p>See “Take checkpoints every __ minutes (policy attribute)” on page 521.</p>
Media owner	<p>Select who should own the media onto which NetBackup writes the images.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ■ Any Lets NetBackup select the media owner, either a media server or server group. ■ None Specifies that the media server that writes to the media that owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ A server group Specify a media server group to allow 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 that are configured in the NetBackup environment appear in the drop-down list. <p>See “Configuring a server group” on page 210.</p>

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 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.

2 Do one of the following:

To change an existing policy

- Select the policy to change.
- On the **Edit** menu, click **Change**.

To create a new policy

- On the **Actions** menu, click **New > New Policy**.
- Name the policy, and click **OK**.

3 Select the **Schedules** tab.

4 Double-click an existing schedule or click **New** to create a new schedule.

5 In the dialog box that appears, click in the **Attributes** tab, select **Multiple copies**, then click **Configure**.

If the destination for this policy is a storage lifecycle policy, the **Multiple copies** box is unchecked. NetBackup does not allow the two methods for creating multiple copies to be enabled at the same time.

See “Policy storage (policy attribute)” on page 518.

See “About writing multiple copies using a storage lifecycle policy” on page 467.

6 In the **Copies** field, specify the number of copies to be created simultaneously. The number must be between 1 and 4.

The maximum is four, or the number of copies that the **Maximum backup copies** setting specifies, whichever is fewer. You can find this host property in **Global Attributes** properties.

See “Global Attributes properties” on page 131.

Copy 1 is the primary copy. If **Copy 1** fails, the first successful copy is the primary 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.

To create more than four copies, create additional copies at a later time by using duplication.

See “Configure Multiple Copies dialog box” on page 564.

See “About configuring for multiple copies” on page 563.

- 7 In the **Priority of duplication** field, specify the priority of the duplication job in comparison to the 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 (schedule attribute)” on page 569.
- 11 Select one of the following from the **If this copy fails** list:

continue	Continues making the remaining copies. Note: If Take checkpoints every __ minutes is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. See “Take checkpoints every __ minutes (policy attribute)” on page 521.
fail all copies	Fails the entire job.

- 12 For tape media, specify who should own the media onto which NetBackup writes the images:

Any	NetBackup selects 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 that are configured in the NetBackup environment appear in the list.

These settings do not affect images residing on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images.

- 13 Click **OK** until the policy is saved.

Override policy storage (schedule attribute)

The **Override policy storage selection** attribute works as follows:

Disabled	Instructs the schedule to use the Policy storage as specified on the policy Attributes tab.
Enabled	<p>Instructs the schedule to override the Policy storage as specified on the policy Attributes tab.</p> <p>Select the storage from the list of previously configured storage units and storage lifecycle policies. If the list is empty, no storage has been configured.</p>

See “Policy storage (policy attribute)” on page 518.

If a data classification is indicated for the policy, only those storage lifecycles with the same data classification appear in the list.

See “Data classifications (policy attribute)” on page 517.

Note: Storage lifecycle policies cannot be selected within the **Configure Multiple Copies** dialog box.

See “About configuring for multiple copies” on page 563.

Override policy volume pool (schedule attribute)

The **Override policy volume pool** attribute works as follows:

Disabled	Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. 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 .
Enabled	Instructs the schedule to override the volume pool that is specified as the Policy volume pool on the policy Attribute tab. Select the volume pool from the list of previously configured volume pools.

See “Policy volume pool (policy attribute)” on page 519.

Override media owner (schedule attribute)

The **Override media owner** attribute applies only to tape media. It specifies whether to use the policy media owner or another owner for the schedule. The

rules for shared disk media are more flexible so override settings are not needed for disk media.

The attribute works as follows:

Disabled Instructs the schedule to use the media owner that is specified as the **Media Owner** in the policy **Attribute** tab.

Enabled Instructs the schedule to override the media owner that is specified as the **Media Owner** in the policy **Attribute** tab.

Select the new media owner from the list:

- **Any.**
NetBackup selects 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 that are configured in the NetBackup environment appear in the list.

See “Media Owner (policy attribute)” on page 527.

Retention (schedule attribute)

The **Retention** attribute specifies how long NetBackup retains the backups. To set the retention period, select a time period (or level) from the list. When the retention period expires, NetBackup deletes information about the expired backup. After 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 storage lifecycle policy, the **Retention** attribute in the schedule is ignored. The retention period that the lifecycle indicates is followed instead.

See “Adding a storage destination to a storage lifecycle policy” on page 450.

About assigning retention periods

The retention period for data depends on the likelihood of restoring information from media after a certain period of time. Some types of 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.

The following table describes recommended retention periods for different types of backups.

Table 16-30 Recommended retention periods for different types of backups

Type of backup	Description
Full Backup	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 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.
Differential Incremental Backup	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.
Cumulative Incremental Backup	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 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.

The following table suggests several ways that you can prevent backups from expiring earlier than desired.

Table 16-31 Suggestions for preventing prematurely expired backups

Item	Description
Retention period	<p>Assign an adequate retention period. NetBackup does not track backups after the retention period expires. Recovering files is difficult or impossible after the retention period expires.</p> <p>For the backups that must be kept for more than one year, set the retention period to infinite.</p>
Full backups and incremental backups	<p>Assign a longer retention period to full backups than to incremental backups within a policy. A complete restore requires the previous full backup plus all subsequent incremental backups. It may not be possible to restore all the files if the full backup expires before the incremental backups.</p>
Archive schedules	<p>Set the retention period to infinite.</p>
Tape	<p>Set the retention period to infinite. If infinite is unacceptable because of NetBackup database space limitations, set the retention period to match the length of time that the data is to be retained.</p>

Another consideration for data retention is off-site storage of the backup media. Off-site storage protects against the disasters that may occur at the primary site. Consider the following off-site storage methods as precautions for disaster recovery:

- Use the duplication feature to make a second copy for off-site storage.
- Send monthly or weekly full backups to an off-site storage facility.
 To restore the data, request the media from the facility. To restore a total directory or disk with incremental backups requires the last full backup plus all incremental backups.
- Configure an extra set of schedules to create the backups to use as duplicates for off-site storage.

Regardless of the method that is used for off-site storage, ensure that adequate retention periods are configured. Use the NetBackup import feature to retrieve expired backups.

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** in 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.

See “Changing a retention period” on page 187.

Media multiplexing (schedule 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 types and some schedule types do not support media multiplexing. The option cannot be selected in those instances.

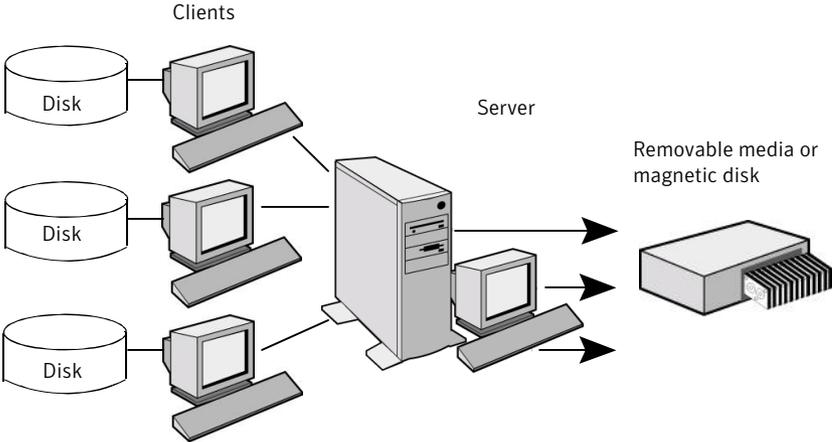
To configure multiplexed backups, multiplexing must be indicated in both the storage unit (**Maximum Streams Per Drive** setting) and the schedule (**Media Multiplexing** setting) configuration. Regardless of the **Media multiplexing** setting, the maximum jobs that NetBackup starts never exceeds the **Maximum Streams Per Drive** value for the storage unit.

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 16-8 shows the multiplexed flow of client data to a server.

Figure 16-8 Multiplexed backups



Multiplexing is generally used to reduce the amount of time that is required to complete backups. The following table describes circumstances where performance improves by using multiplexing:

Table 16-32 Circumstances where performance improves by using multiplexing

Item	Description
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 effect of multiplexing on restore times, set the storage unit maximum fragment size to a value smaller than the largest allowed value.

Consider the following configuration settings when using multiplexing.

Table 16-33 Properties and attributes that affect multiplexing

Item	Description	Where to find it
<p>Limit jobs per policy (policy attribute)</p>	<p>Limits the number of jobs that NetBackup performs concurrently when a policy is run. Set this attribute high enough to support the specified level of multiplexing.</p> <p>See “Limit jobs per policy (policy attribute)” on page 525.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Policies. ■ In the left pane, double-click a policy name. <p>Or, create a new policy and select the Attributes tab.</p>

Table 16-33 Properties and attributes that affect multiplexing (*continued*)

Item	Description	Where to find it
<p>Maximum jobs per client (host property)</p>	<p>Limits the number of backup jobs that can run concurrently on any NetBackup client. This property is part of Global Attributes host properties.</p> <p>See “Global Attributes properties” on page 131.</p> <p>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, NetBackup cannot use the storage unit’s full multiplexing capabilities.</p> <p>Select 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.</p> <p>To estimate the potential load that multiplexing can place on the central processing unit, consider the following limits:</p> <ul style="list-style-type: none"> ■ 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. <p>See “Maximum streams per drive storage unit setting” on page 406.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Host Properties > Master Servers. ■ In the right pane, double-click a master server. ■ In the Master Server Properties dialog box, select Global Attributes from the left pane. ■ The property appears in the right pane.

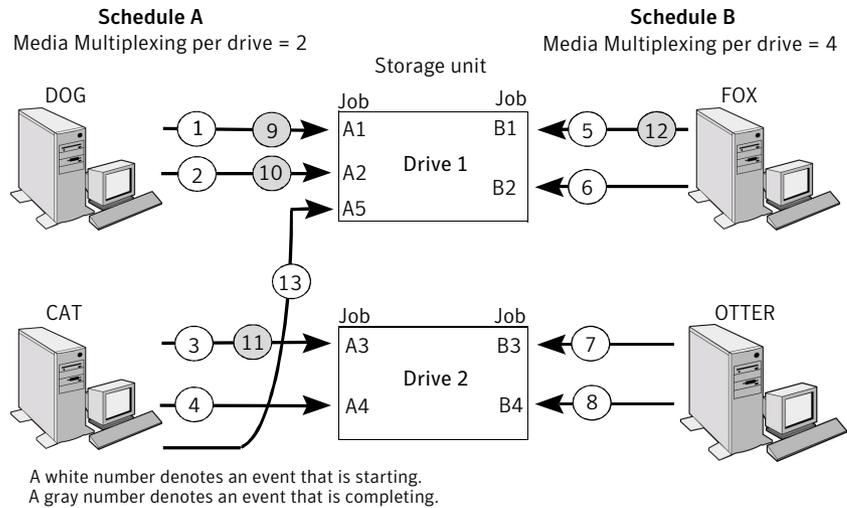
Table 16-33 Properties and attributes that affect multiplexing (continued)

Item	Description	Where to find it
<p>Maximum data streams (host property)</p>	<p>Set the maximum number of jobs that are allowed on a specific client without affecting other clients. This property is part of Client Attributes host properties.</p> <p>See “General tab of the Client Attributes properties” on page 80.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Host Properties > Master Servers. ■ In the right pane, double-click a master server. ■ In the Master Server Properties dialog box, select Client Attributes from the left pane. ■ The property appears in the right pane on the General tab.
<p>Delay on multiplexed restores (host property)</p>	<p>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. This property is part of General Server host properties.</p> <p>See “General tab of the Client Attributes properties” on page 80.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Host Properties > Master Servers. ■ In the right pane, double-click a master server. ■ In the Master Server Properties dialog box, select General Server from the left pane. ■ The property appears in the right pane.
<p>Media Multiplexing (policy schedule attribute)</p>	<p>If the limit is reached for a drive, NetBackup sends jobs to other drives.</p> <p>When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches the Media Multiplexing limit or the Maximum streams per drive limit.</p> <p>See “Media multiplexing (schedule attribute)” on page 572.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Policies. ■ In the left pane, double-click a policy name. Select the Schedules tab. Or, create a new policy and select the Schedules tab. ■ Click New to create a new schedule and configure the Media Multiplexing option.
<p>Maximum streams per drive (storage unit setting)</p>	<p>NetBackup can add jobs from more than one schedule to a drive.</p> <p>When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches the Maximum streams per drive limit or the Media Multiplexing limit</p> <p>See “Maximum streams per drive storage unit setting” on page 406.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Storage. ■ In the left pane, click Storage Units. ■ In the right pane, double-click a storage unit name. Or, create a new storage unit. ■ The setting appears on the dialog box that appears.

Example of using multiplexing with schedules

Figure 16-9 provides an example of how schedules are affected when multiplexing is active.

Figure 16-9 Multiplexing process scenario



Assume the following about Figure 16-9.

- Schedule A begins first.
 Schedules can be in the same or in different policies.
- **Allow Multiple Data Streams** is enabled.
 Consequently, a client can have multiple data streams.
 See “Allow multiple data streams (policy attribute)” on page 542.

Table 16-34 Description of the multiplexing process scenario

Event	Description
1 and 2	<ul style="list-style-type: none"> ■ Jobs A1 and A2 from client <i>DOG</i> start on Drive 1. ■ For Schedule A, the Media Multiplexing limit of 2 is reached for Drive 1.
3 and 4	<ul style="list-style-type: none"> ■ Jobs A3 and A4 from client <i>CAT</i> start on Drive 2. ■ For Schedule A, the Media Multiplexing limit of 2 is reached for Drive 2.

Table 16-34 Description of the multiplexing process scenario (continued)

Event	Description
5 and 6	<ul style="list-style-type: none"> ■ Jobs B1 and B2 for client <i>FOX</i> start on Drive 1. ■ The Maximum streams per drive storage unit setting is reached for Drive 1.
7 and 8	<ul style="list-style-type: none"> ■ Jobs B3 and B4 from client <i>OTTER</i> start on Drive 2. ■ All jobs are now running for Schedule B. ■ The Maximum streams per drive storage unit setting is reached for Drive 2.
9 and 10	<ul style="list-style-type: none"> ■ Jobs A1 and A2 from client <i>DOG</i> finish on Drive 1. ■ However, jobs B1 and B2 for client <i>FOX</i> continue to run. ■ For Schedule A, the Media Multiplexing limit of 2 prevents job A5 from starting on Drive 1
11 and 12	<ul style="list-style-type: none"> ■ Job A3 from client <i>CAT</i> finishes on Drive 2 ■ Job B1 from client <i>FOX</i> finishes on Drive 1. ■ Job B2 is the only job currently running on Drive 1.
13	<ul style="list-style-type: none"> ■ Job A5 from client <i>CAT</i> starts on Drive 1. ■ JobA5 is the last job for Schedule A. ■ For Schedule A, the 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 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 16-9 assumes that the **Activity Monitor** shows jobs A1 through A5 as queued and active.

If only jobs 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. Job A1 in this example.

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 **Configure Multiple Copies** 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 743.

Start Window 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 a 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 different time windows:

- One for the backups that open each day for a specific amount of time
- Another for the backups that keep the window open all week

Adding, changing, or deleting a time window in a schedule

Use the following procedure to add, change, or delete a time window.

To add or change a time window

- 1 In the **Netbackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the left pane, double-click the policy you want to change or add a time window to.
- 3 Select the **Schedules** tab and do one of the following:

To add a time window ■ Click **New**.

- In the **Add New Schedule** dialog box, type the name of a schedule.

To change a time window

Double-click the schedule you want to change. The **Change Schedule** dialog box appears.

9 To indicate the closing of the time window, do one of the following:

- | | |
|---------------------------------------|---|
| Enter the duration of the time window | Enter a length of time in the Duration (days, hours, minutes) fields. |
| Indicate the end of the time window | <ul style="list-style-type: none"> ■ Select a day in the End day list. ■ Select a time in the End time field. |

Time windows show as bars in the schedule display.

To create multiple time windows do the following:

- | | |
|--|---|
| To add time windows on successive days | <ul style="list-style-type: none"> ■ With the cursor over the chosen start time, press and hold the Shift key. ■ Click and drag the cursor to the time when you want to the time window to close. ■ Continue holding the Shift key, and drag the cursor down to the last day of the week you want to include. <p>Duplicates of the time window appear for successive days.</p> |
| To copy a time window | <ul style="list-style-type: none"> ■ Create a time window. ■ Click Duplicate. <p>The time window is duplicated to any days without existing schedules. Duplication stops when it reaches a day that already contains a defined schedule.</p> <ul style="list-style-type: none"> ■ On days that you do not want the time window to be open, select the window and click Delete. |

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.

10 Do any of the following:

- | | |
|--------------------------------------|--|
| To change the start time or end time | <ul style="list-style-type: none"> ■ Adjust the Start time or End time. ■ Click and drag the end of the time window bar to a new position. |
| To move a time window | Click and drag the time window bar to a new position. |
| To delete a time window | Select a time window and click Delete . |

To delete all the time windows Click **Clear**.

To erase the last action Click **Undo**.

11 Click **OK** to save the completed schedule.

Example of schedule duration

Figure 16-10 illustrates the effect of schedule duration on two full backup schedules. The start time for Schedule B begins shortly after the end time for the previous Schedule A. Both schedules have three clients with backups due.

Figure 16-10 Duration example

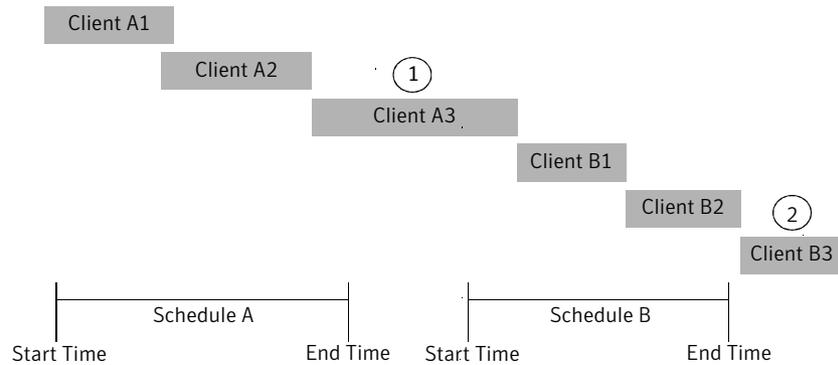


Figure 16-10 illustrates the following points:

- Point 1** Client A3 starts within the Schedule A time window but doesn't complete until after the Schedule B start time. However, Client A3 runs to completion even if the window closes while the backup is running. Client B1, on Schedule B, begins as soon as Client A3 completes.

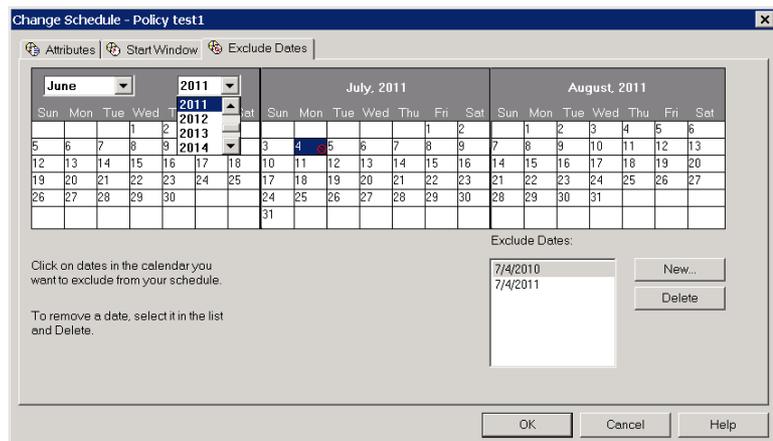
- Point2** Schedule A does not leave enough time for all the clients on Schedule B to be backed up. Consequently, Client B3 is unable to start because the time window has closed. Client B3 must wait until the next time NetBackup runs Schedule B.

Excluding dates from a policy schedule

Use the **Exclude Dates** tab to exclude specific dates from a schedule. If a date is excluded from a schedule, the policy does not run on that day. The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

To exclude a date from the policy schedule

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 On the **Edit** menu, click **Change**, and select the **Schedules** tab on the dialog box that appears.
- 4 Select the schedule you want to modify, and click **Properties**.
- 5 In the dialog box that appears, select the **Exclude Dates** tab.



- 6 Do one of the following:
 - Click the date on the calendar you want to exclude. Use the lists at the top of the calendar to change the first month or year displayed
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box, and click **OK**.

The date appears in the **Exclude Dates** list.

- 7 Add additional dates as necessary, then click **OK** to save the changes.

Calendar Schedule tab

The **Calendar Schedule** tab appears in the **Add New Schedule** or **Change Schedule** dialog box. For the tab to display, you must select the **Calendar** option as the **Schedule type** on the **Attributes** tab. Calendar-based schedules provide several run day options for determining when a task runs.

The tab displays a calendar of three consecutive months. Use the lists at the top of the calendar to change the first month or year displayed.

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

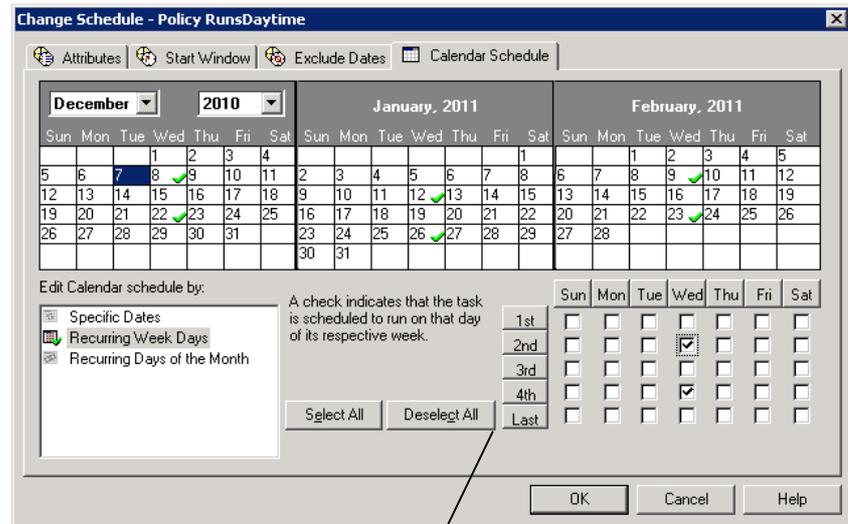
- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
 - 2 Select the policy name in the left pane.
 - 3 On the **Edit** menu, click **Change**, and select the **Schedules** tab on the dialog box that appears.
 - 4 Select the schedule you want to modify, and click **Properties**.
 - 5 In the dialog box that appears, select the **Calendar** schedule type.
 - 6 Select the **Calendar Schedule** tab that appears.
 - 7 In the **Edit Calendar schedule by** list, select **Specific Dates** and do one of the following:
 - Click a 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.
- 8 Add additional dates as necessary, then click **OK** to save the changes.

Scheduling by recurring days of the week

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 the day of that week for each month in the future.

For example, schedule a task to run on the first and the third Thursday of every month. Or, schedule a task to run the last week in every month.

Figure 16-11 Recurring week days setting on the Calendar Schedule tab



Matrix

To schedule a recurring weekly task

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 On the **Edit** menu, click **Change**, and select the **Schedules** tab on the dialog box that appears.
- 4 Select the schedule you want to modify, and click **Properties**.
- 5 In the dialog box that appears, select the **Calendar** schedule type.
- 6 Select the **Calendar Schedule** tab that appears.
- 7 In the **Edit Calendar schedule by** list, select **Recurring Week Days**.
- 8 Do any of the following:
 - Click **Deselect All** to remove existing selections from the matrix.
 - Click **Select All** to select all of the days in every month.
 - Check a box in the matrix to select the day.

- Click the column header with the name of the day to select or clear the corresponding day for each week of the month.
 - Click a row number to select or clear the entire week.
 - Check the 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.
- 9 After the dates are selected, click **OK** to save the changes.

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 (1st through 31st). In addition, a task can be scheduled for the last day of the month, regardless of the actual date.

To schedule a recurring monthly task

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 On the **Edit** menu, click **Change**, and select the **Schedules** tab on the dialog box that appears.
- 4 Select the schedule you want to modify, and click **Properties**.
- 5 In the dialog box that appears, select the **Calendar** schedule type.
- 6 Select the **Calendar Schedule** tab that appears.
- 7 In the **Edit Calendar schedule by** list, select **Recurring Days of the Month**.
- 8 Do any of the following:
 - Click **Deselect All** to remove existing selections from the matrix.
 - Click **Select All** to select all of the days in every month.
 - Click the number for each day to be included in the run schedule. Click the number again to deselect the day.
 - Check **Last Day** to run the schedule on the last day of the month, regardless of the date.
- 9 After the dates are selected, click **OK** to save the changes.

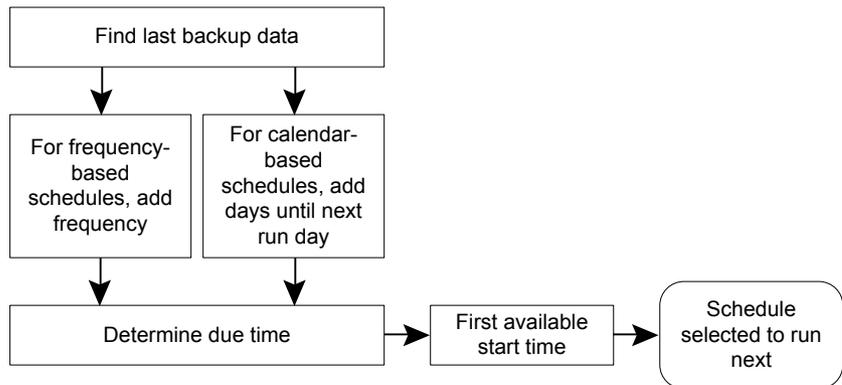
How NetBackup determines which schedule to run next

When a policy contains one schedule, the schedule that is selected to run next is straightforward. But when a policy contains multiple schedules, choosing which schedule to run next can become more complicated.

NetBackup performs the following tasks to determine which schedule to run next:

- 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.
- 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 16-12 Schedule selection overview



When any of the following events occurs, NetBackup recalculates which schedule to run next in a policy:

- A backup job finishes.
- A client backup image expires.
- The Policy Execution Manager (*nbpem*) starts.
- The administrator changes the policy.

NetBackup looks for updated policies every 10 minutes. If the policy has recently been updated, NetBackup waits an additional minute to be sure that changes are not currently underway. You can change the frequency that

NetBackup looks for updates by changing the **Policy Update Interval** in the **Global Attributes** host properties.

See “Global Attributes properties” on page 131.

The due time for each schedule equals the last backup data for the schedule, plus the schedule’s frequency:

$$\text{Due time} = \text{Last backup data} + \text{Frequency}$$

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.

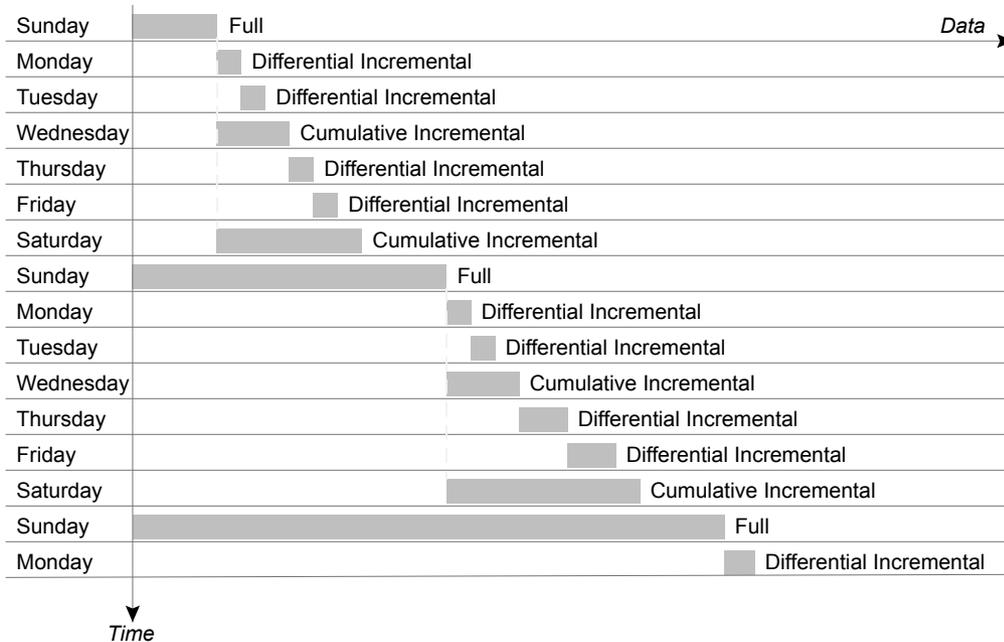
In some cases, the last backup data for a schedule names the schedule itself. In other cases, the last backup data for a schedule is another comparable schedule.

NetBackup makes the following comparisons to identify a comparable schedule:

Full schedules	Compared to other full schedules of the same or longer frequency.
Cumulative incremental schedules	Compared to the following: <ul style="list-style-type: none">■ Full schedules of the same or longer frequency.■ Other cumulative incremental schedules of the same or longer frequency.
Differential incremental schedules	Compared to the following: <ul style="list-style-type: none">■ 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. <p>Note: To have a longer frequency means that the schedule is configured to run less often.</p>

The comparison rules ensure that no schedule is overlooked for consideration, potentially causing a gap in backup coverage.

Figure 16-13 Schedule coverage



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

NetBackup uses the previous synthetic job as the basis for determining when the next synthetic job should run.

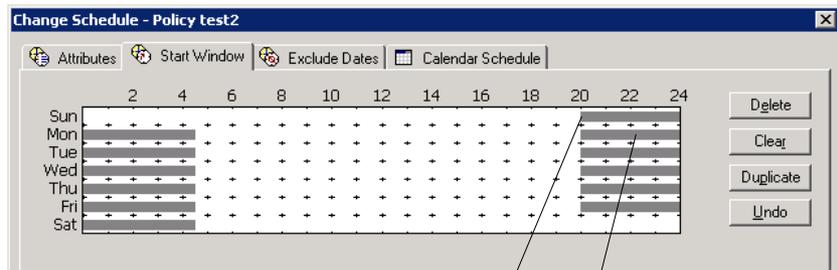
About schedule windows that span midnight

A backup window may begin in one day and end in another. If this kind of 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 16-14 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 16-14 Schedule that spans midnight



The first job begins Sunday.

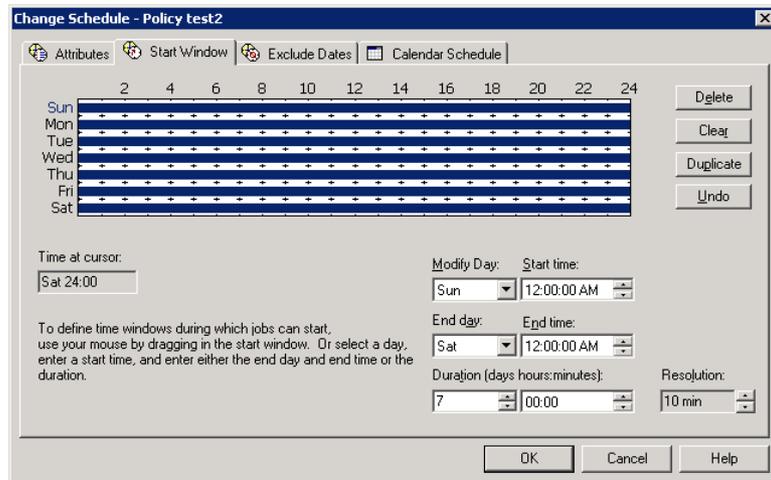
The job is due Monday as well. Instead of running the job again immediately after midnight, NetBackup looks for a window later in the day and runs the job.

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 calendar-based and frequency-based schedules

A single backup window can span the entire week. This kind of schedule is called an open schedule because a job may run at any time of day or night during the week. Open schedules affect calendar-based and frequency-based schedules differently.

Figure 16-15 shows an open schedule.

Figure 16-15 An open schedule

Open schedules affect calendar-based and frequency-based schedules differently:

Calendar-based schedules Calendar-based schedules run whenever the calendar schedule indicates. 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.

Frequency-based schedules Frequency-based schedules run when the frequency setting indicates. For example, with a frequency of one day, NetBackup runs backups at 24-hour intervals based on the start time.

Figure 16-16 shows that the backups on a calendar-based schedule would run Monday through Friday.

Figure 16-16 An open schedule that is calendar-based

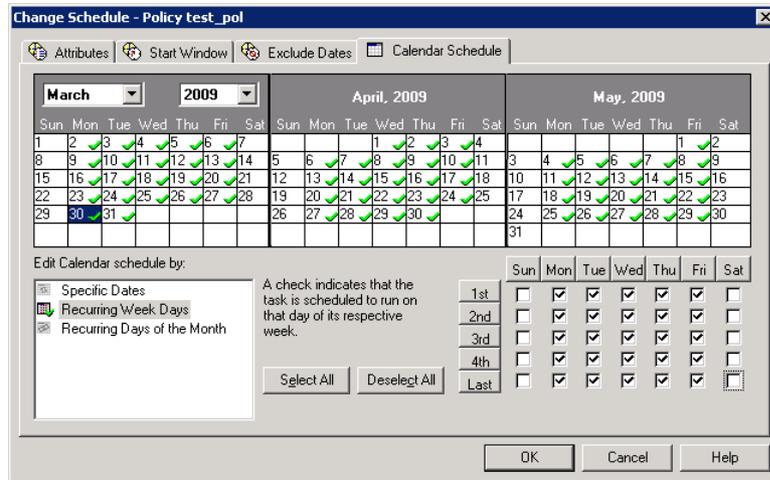
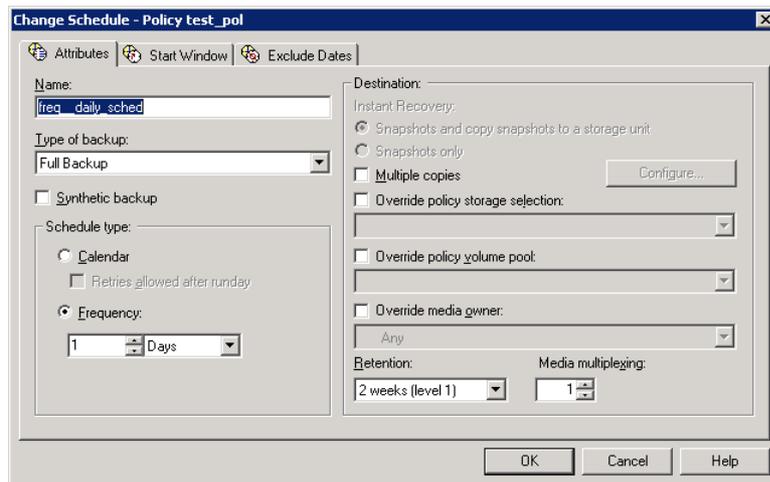
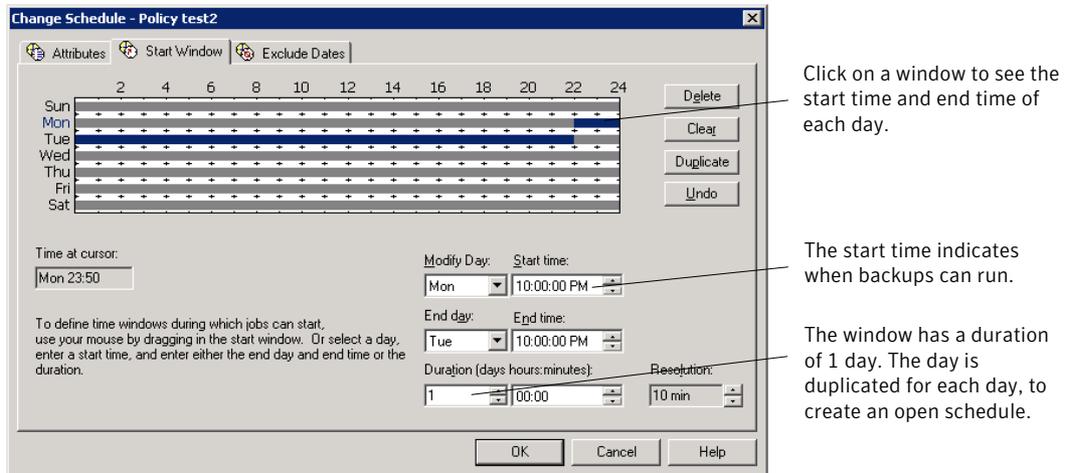


Figure 16-17 and Figure 16-18 show that the backups based on a frequency-based schedule should run every day of the week, including Saturday and Sunday.

Figure 16-17 An open schedule that is frequency-based



In Figure 16-18, backups run at 10:00 P.M. nightly based on the start time.

Figure 16-18 Example of a frequency-based schedule with an open schedule

Creating an open schedule in the NetBackup Administration Console

The following procedure describes how to create an open schedule in an existing policy. In this procedure, the open schedule is configured to begin at 10:00 P.M.

To create an open schedule in the NetBackup Administration Console

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the left pane, double-click on the policy name where you want to create an open schedule.
- 3 Select the **Schedules** tab.
- 4 Click **New** to create a new schedule.
- 5 Complete the information on the **Attributes** tab.
- 6 Select the **Start Window** tab.
- 7 Select Sunday as the **Modify Day** and **10:00:00 PM** as the **Start time**.
- 8 Select Monday as the **End Day** and **10:00:00 PM** as the **End time**. The **Duration** is then automatically set to one day.
- 9 Click **Duplicate** to copy this window to each day of the week.
- 10 Click **OK** to add the schedule to the policy.

Runtime considerations that affect backup frequency

The following items may cause a NetBackup job to run more frequently than expected, or may prevent a job from meeting its backup frequency requirement.

Table 16-35 Items that can affect backup frequency

Item	Description
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	<p>Whether the schedule is calendar-based or frequency-based, a job cannot run if windows are not open on the configured runday.</p> <ul style="list-style-type: none"> ■ 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. <p>Note: A frequency is not configurable for a calendar-based schedule. For this schedule type, NetBackup assumes a daily backup frequency.</p> <ul style="list-style-type: none"> ■ For frequency-based schedules, a daily frequency requires that a window is open each day.
Backup attempt limit	<p>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.</p> <p>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.</p> <p>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.</p> <p>See “Global Attributes properties” on page 131.</p>

Runtime considerations

The following items may cause a NetBackup 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	<p>Whether the schedule is calendar-based or frequency-based, a job cannot run if windows are not open on the configured rundays.</p> <ul style="list-style-type: none">■ 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. <p>Note: A frequency is not configurable for a calendar-based schedule. For this schedule type, NetBackup assumes a daily backup frequency.</p> <ul style="list-style-type: none">■ For frequency-based schedules, a daily frequency requires that a window is open each day.
Backup attempt limit	<p>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.</p> <p>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.</p> <p>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.</p> <p>See “Global Attributes properties” on page 131.</p>

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 or changing clients in a policy

A client must be included in the list of at least one active backup policy to be backed up. Use the following procedure to add, change, or delete clients in an existing NetBackup policy.

- 5 If **Detect operating system when adding or changing a client** is not checked, do the following:
 - Select the appropriate hardware and operating system in the list.
Add only clients with the hardware and the operating system that the policy supports. For example, do not add a Novell NetWare client to an MS-Windows policy. If you add a client to more than one policy, designate the same hardware and operating system in each of the policies.
 - Click **OK** to add the client to the list of clients and close the **Client Hardware and Operating System** dialog box.
- 6 To change an existing client:
 - Select the client name in the list and hover until the name becomes active. Type in the active field to change the client name. Press **Enter** to accept the change.
 - To change the operating system of the client, select one from the list in the **Client Hardware and Operating System** dialog box.
 - Click **OK** to accept the change and close the **Client Hardware and Operating System** dialog box.
- 7 When you are finished in the **Clients** tab, do one of the following:
 - Click **OK** to close and save the policy.
 - Click **Cancel** to close the policy without saving any additions or changes.

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 Machines**
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 virtual 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.

Backup Selections tab

The **Backup Selections** tab lists the paths, directives, scripts, and the templates that specify which files and directories are backed up on each client. NetBackup uses the same backup selection list for all of the clients that are backed up according to the policy.

The policy type determines whether the backup selections list contains paths, directives, scripts, templates, or a combination. Paths identify the location of files. Directives are the predefined sets of instructions that NetBackup uses to perform specific actions. Administrators create scripts to define and control database backups. Scripts include instructions for how the client uses multiple streams. Templates are used exclusively for Oracle and DB2 database backups.

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. (Use the Troubleshooter to find the description of a status code.)

See “Running the Troubleshooter” on page 42.

Note: Windows clients support only the asterisk (*) and the question mark (?) as valid wildcards in the **Backup Selections** tab.

See “Wildcard use in NetBackup” on page 828.

The 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.

Table 16-36 Items allowed in the Backup Selections list for specific policy types

Policy type	Items allowed
Standard	Paths and directives
Exchange	Paths and directives
Lotus Notes	Paths and directives
MS-SQL-Server,	Scripts
Informix-On-BAR	Scripts
SAP	Scripts
Sybase	Scripts
Oracle	Scripts and templates
DB2	Scripts and templates
Vault	Vault commands

See “Policy type (policy attribute)” on page 514.

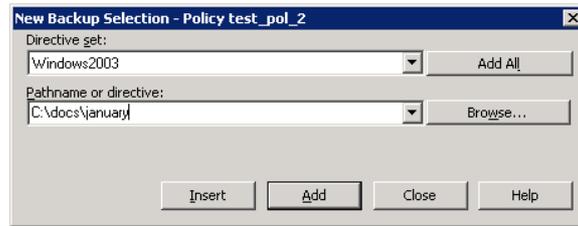
Adding backup selections to a policy

Use the following procedure to add backup selections to a NetBackup policy, without opening up the tab view of the policy.

To add backup selections to a policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane where you'd like to add the backup selection.
- 3 On the **Actions** menu, click **New > New Backup Selection**.
- 4 In the **New Backup Selection** dialog box, indicate the path to the directory that you want to back up.

The policy type determines whether the backup selections list can contain paths, directives, scripts, templates, or a combination.



Entering a path to a directory

Click **Browse** to browse to a remote folder to select a path. Or, enter the pathname directly in the **Pathname or Directive** field.

A path may contain up to 1023 characters.

Selecting a directive set or directive

Select or enter a directive set in the **Directive set** drop-down menu.

Select or enter a directive in the **Pathname or Directive** drop-down menu.

See “About the directives on the Backup Selections list” on page 620.

Selecting a script or template

- Select or enter a script or template in the **Script or template** drop-down menu.
Templates are stored in a known location on the master server and do not need to be installed on each client in the **Clients** list. Enter only the template file name, without a path. For example: `weekly_full_backup.tpl`
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 **Clients** tab.
- For Oracle policies, select a template set based on an operation from the **Template set** list.
For Oracle policies, select **Oracle_RMAN** or **Oracle_XML_Export**.
- For Oracle policies or DB2 policies, select a template from the **Script or template** list, or type the name of a template.

See “Policy type (policy attribute)” on page 514.

See “Verifying the Backup Selections list” on page 601.

5 Click **Add** to add the item to the Backup Selections list for the policy.

6 Click **Close** to close the **New Backup Selection** dialog box.

Verifying the Backup Selections list

Verify the **Backup Selections** list to make sure that the file paths are correct for the clients in the policy.

Table 16-37 Steps to verify the Backup Selections list

Step	Action	Description
Step 1	Check the syntax for the directives and the file path rules.	<p>Do the following:</p> <ul style="list-style-type: none"> ■ If the list includes directives, verify that the syntax for the directives is correct. ■ Check all entries against the file path rules for the clients in the policy. <p>See “Pathname rules for Windows client backups” on page 604. See “Pathname rules for Windows disk image (raw) backups” on page 605. See “Pathname rules for Windows registry backups” on page 607. See “Pathname rules for NetWare NonTarget clients” on page 617. See “Pathname rules for NetWare Target clients” on page 619. See “Pathname rules for UNIX client backups” on page 610.</p> <p>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.)</p>
Step 2	Check for warning messages.	<p>Do the following:</p> <ul style="list-style-type: none"> ■ Run a set of backups. ■ Check the Problems report or the All Log Entries report for warning messages. <p>The backup status code does not always indicate errors on the Backup Selection list. Because NetBackup does not require all paths in the Backup Selections list to be present on all clients, an error may not be especially helpful.</p> <p>See “Problems report” on page 821. See “All Log Entries report” on page 821.</p>

Table 16-37 Steps to verify the Backup Selections list (continued)

Step	Action	Description
Step 3	Create a File System Backup Coverage Report .	<p>Run the <code>check_coverage</code> script to create a File System Backup Coverage Report.</p> <p>The script is located in <code>install_path\NetBackup\bin\goodies</code>. The script can reveal mistakes in the selections list that make it impossible for NetBackup to find the files. Mistakes in the selections list can result in files being skipped in the backup.</p> <p>If a path is not found, NetBackup logs a trivial (TRV) message or a 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 because not all files are expected to be present on every client. However, check the logs or use the <code>check_coverage</code> script to ensure that files are not missed due to bad or missing backup selections list entries.</p> <p>See “Example log messages from the File System Backup Coverage Report (<code>check_coverage</code>)” on page 602.</p>

Example log messages from the File System Backup Coverage Report (`check_coverage`)

The **File System Backup Coverage Report** is created by running the `check_coverage` script. The following example shows the log message that appears when files expected to be 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:

```
9/1/10 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.

How to reduce backup time

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).
- **Maximum jobs per client** (in **Global Attributes** host properties) and the **Limit jobs per policy** policy attribute are set to allow it.
See "Global Attributes properties" on page 131.
See "Limit jobs per policy (policy attribute)" on page 525.

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 **Maximum jobs 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 (policy attribute)" on page 542.

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 disk 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.

Pathname rules for Windows client backups

To back up Windows clients, use the following conventions for entries in the backup selections list.

Table 16-38 Pathname rules for Windows client backups

Item	Description
Paths per line	Enter one path per line.
Colons and backslashes	<p>Begin all paths with the drive letter followed by a colon (:) and a backslash (\).</p> <p>To specify an entire volume, append a backslash (\) to the entry to ensure that all data is protected on that volume:</p> <p>Correct entry: c:\</p> <p>Incorrect entry: c:</p>
Case sensitivity	<p>The drive letter is case-insensitive, but the path is case-sensitive. For example, c:\Worklists\Admin\</p>
Wildcards	<p>Asterisks (*) and question marks (?) are the only wildcard characters allowed in the backup selection list for Windows clients.</p> <p>Square brackets and curly brackets are not valid for Windows clients and can cause backups to fail with a status 71.</p> <p>See “Wildcard use in NetBackup” on page 828.</p>
All local drives	<p>To back up all local drives except for those that use removable media, specify the following:</p> <p>: \</p> <p>Or</p> <p>*: \ or ALL_LOCAL_DRIVES</p> <p>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.</p>

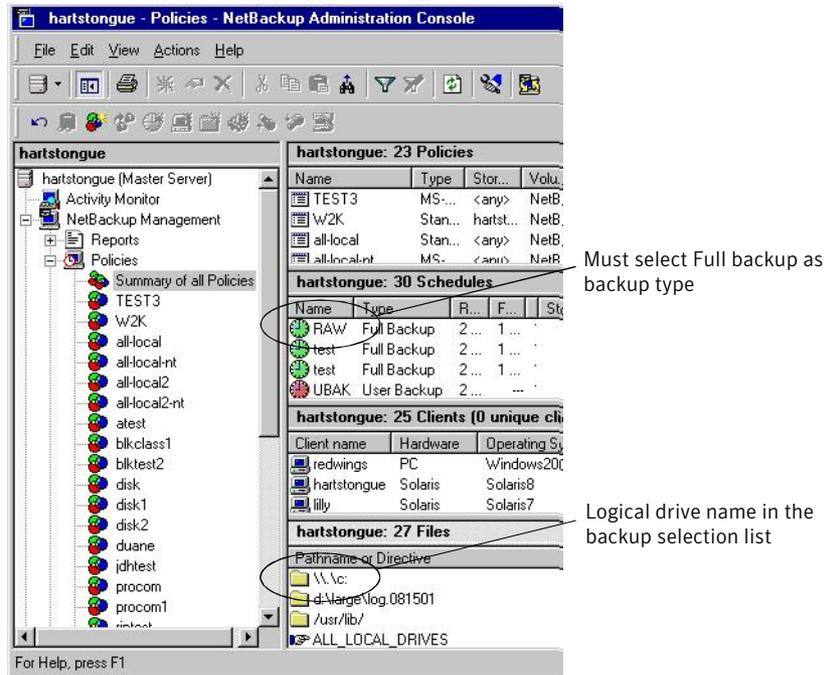
Table 16-38 Pathname rules for Windows client backups (*continued*)

Item	Description
Omitted or excluded files	<p>By default, NetBackup does not back up some files.</p> <p>See “Files that are excluded from backups by default” on page 629.</p> <p>Exclude specific files from backups by creating an exclusion list on the client.</p> <p>See “About excluding files from automatic backups” on page 630.</p> <p>The following backup selection list uses Windows conventions:</p> <pre>c:\ d:\workfiles\ e:\Special\status c:\tests*.exe</pre>

Pathname 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. Use the **Full backup** backup type to perform a disk image backup.

Figure 16-19 Disk image backups



To specify a disk image backup, add the logical name for the drive to the policy backup selection list. Disk images can be included in the same backup selection list with other backups. In the following sample backup selection list, the first entry (\\.\c:) creates a disk image backup of a logical drive C.

```
\\.\c:
d:\workfiles\
e:\Special\status
HKEY_LOCAL_MACHINE:\
```

To restore the backup, the user clicks **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. For the previous example, the file name would show as follows:

```
\\.\c:
```

When you enter the destination to restore the file, use the following format:

\\.*drive*:

Where *drive* is the location where the partition is to be restored.

Consider the following when working with disk image backups:

Windows Open File Backup methods	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.
Open files	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.
Copy On Write snapshots	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.
Raw partitions	NetBackup does not support raw partition backups on unformatted partitions.
Paging file	If the volume is configured to contain a paging file (<code>pagefile.sys</code>), 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.

Pathname rules for Windows registry backups

The Windows registry can be backed up for disaster recover or individual HKEYs can be backed up. Consider the following items when configuring a Windows registry backup.

Disaster recovery	<p>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.</p> <p>On most Windows systems, this directory is located at:</p> <pre>%systemroot%\system32\config</pre> <p>Where <code>%systemroot%</code> is the directory where Windows is installed.</p> <p>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.</p> <p>To restore the registry in the case of a disk failure, see the Disaster Recovery chapter in the <i>NetBackup Troubleshooting Guide</i>.</p>
Individual HKEYs	<p>Do not back up individual HKEYs for disaster recovery. You cannot perform a disaster recovery by restoring HKEYs. 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.</p> <p>The following is an example HKEY entry for a policy backup selection list:</p> <pre>HKEY_LOCAL_MACHINE:\</pre> <p>Backups and restores are slower than if the entire registry was backed up.</p>

About hard links to files and directories

A hard link is a directory entry for a file. Every file can be considered to have at least one hard link. 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.

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.

Hard-linked files and hard-linked directories are different in the following ways:

- 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.

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 indicate the actual, unique file on the volume. Collectively, the VSN and File Index are referred to as the file ID.

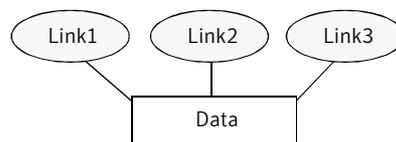
During a backup, if the backup selection list includes hard-linked files, the data is backed up only once. NetBackup uses the first file name reference that is found in the directory structure. If a subsequent file name reference is found, the reference is backed up as a link to the name of the first 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: Restoring Link2 and Link3

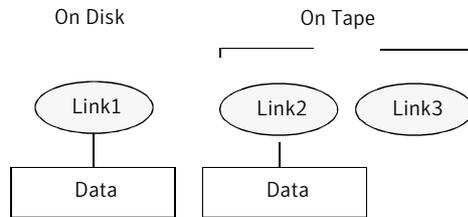
Assume that three hard links point to the same data. During a backup of Link2 and Link3, Link2 is encountered first and backed up. Then Link3 is backed up as a link to Link2. The three files are all hard-linked to the same data.

Figure 16-20 Example of hard links to the same data



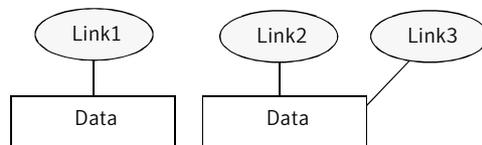
The original copies of Link2 and Link3 are backed up to tape, then deleted. Only Link1 is left on the disk.

Figure 16-21 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 16-22 Example of restored hard links



Example 2: Restoring Link3

Assume that this time you attempt to restore only Link3. However, 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 client backups

To back up UNIX clients, use the following conventions for entries in the backup selections list.

Table 16-39 Pathname rules for UNIX client backups

Item	Description
Pathnames per line	Enter one pathname per line. NetBackup supports a maximum path length of 1023 characters for UNIX clients.
Forward slash	Begin all pathnames with a forward slash (/).

Table 16-39 Pathname rules for UNIX client backups (continued)

Item	Description
Wildcard characters	<p>The following wildcard characters are allowed:</p> <ul style="list-style-type: none"> * ? [] { } <p>See “Wildcard use in NetBackup” on page 828.</p>
Trailing spaces	<p>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:</p> <pre>TRV - cannot process path pathname: No such file or directory. Skipping TRV - Found no matching file system for pathname</pre>
Mount points	<p>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.</p> <p>See “Follow NFS (policy attribute)” on page 528.</p> <p>See “Cross mount points (policy attribute)” on page 533.</p>
Bootable tapes	<p>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.</p>
Omitted or excluded files	<p>By default, NetBackup does not back up all files.</p> <p>See “Files that are excluded from backups by default” on page 629.</p> <p>Exclude specific files from backups by creating an exclusion list on the client.</p> <p>See “About excluding files from automatic backups” on page 630.</p>
Busy File Settings	<p>The Busy File Settings host properties for UNIX clients offers alternatives for handling busy and locked files.</p> <p>See “Busy File Settings properties” on page 72.</p>
Access Control Lists (ACLs)	<p>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).</p>
Sun PC NetLink	<p>NetBackup can back up and restore Sun PC NetLink files.</p>

Table 16-39 Pathname rules for UNIX client backups (*continued*)

Item	Description
Extended attribute files and named data streams	<ul style="list-style-type: none"> ■ By default, NetBackup backs up and restores Solaris 9 and 10 extended attribute files. ■ The FlashBackup single file restore program (<code>sfr</code>) 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 (<code>sfr</code>) does not restore extended attribute files. <p>See “About backing up and restoring extended attribute files and named data streams” on page 615.</p>
VxFS extent attributes	<p>On Hewlett-Packard and Solaris SPARC platforms, NetBackup backs up VxFS extent attributes.</p>
Symbolic links	<p>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.</p> <p>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.</p> <p>See “About hard links to files and directories” on page 608.</p> <p>Note: If NetBackup restores a symbolic link as <code>root</code>, NetBackup changes the owner and group to the original owner and group. When NetBackup restores a symbolic link as a non-root user, the owner and group are set to the owner and the group of the person who performs the restore. Resetting the owner and group does not cause problems. When the UNIX system checks permissions, NetBackup uses the owner and group of the file to which the symbolic link points.</p>
Directory junctions	<p>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.</p> <p>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.</p>

See “About the Reports utility” on page 818.

UNIX raw partitions

Save a copy of the partition table before a raw partition backup is performed. Retain the copy for reference. To restore the raw partition, make sure that a device file exists. Also, the partition where the table is restored must be large enough or the results of the restore are unpredictable.

Consider the following items when creating UNIX raw partition backups.

File changes during the backup	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.
Backup archives	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.
File systems	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 <code>bpstart_notify</code> and the <code>bpend_notify</code> scripts to unmount and remount the backed-up file systems.
Mount points	<p>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.</p> <p>See “Cross mount points (policy attribute)” on page 533.</p> <p>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 computers by using NFS mounts to access the raw partitions. The devices are not accessible on other computers through NFS.</p> <p>See “Follow NFS (policy attribute)” on page 528.</p>
Disk volume managers	Specify the logical partition names for any disks that disk volume managers manage. (For example, Veritas Volume Manager (VxVM).)
FlashBackup policy	For clients in a FlashBackup policy, refer to the <i>NetBackup Snapshot Client Administrator’s Guide</i> for the differences between Standard and FlashBackup policies.

Microsoft Cluster (MSCS) environment The use of FlashBackup in a Microsoft Cluster (MSCS) environment is supported, with the following limitation: Raw partition restores can only be performed when the disk being restored is placed in extended maintenance mode or removed from the MSCS resource group.

Note: Earlier versions of MSCS (such as those versions that were shipped with Windows versions before Windows 2003 SP1) do not allow extended maintenance mode functionality. If the cluster does not support placing disks in extended maintenance mode, it is still possible to perform raw restores to an alternate, non-shared disk.

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 (schedule attribute)” on page 550.

About 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: Access Control Lists (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**. The error message informs 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)
- 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/opensv/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.

Pathname 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
```

The elements of the example path are described as follows:

<i>SMDR</i>	The 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.)
<i>TSA</i>	The 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.
<i>TS</i>	The Target Service is the NetWare entity that contains the data that the selected TSA handles. For example, in the case of the DOS TSA (<i>tsasms.com</i>) 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.
<i>resources</i>	The resources on the target service. For example, it can be NetWare file systems such as BINDERY, SYS, and USER.
<i>directory/file</i>	The directory and file that are in the resource (if it is a path to a specific file).

To back up NetWare NonTarget clients, use the following conventions for entries in the backup selections list.

Table 16-40 Pathname rules for NetWare NonTarget clients

Item	Description
Server access	<p>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.</p> <p>For more information, see the <i>NetBackup for Novell NetWare Client Administrator's Guide</i>.</p>
Paths per line	<p>Enter one path per line.</p>
Forward slash	<ul style="list-style-type: none"> ■ Begin all paths with a forward slash (/). ■ Precede each component in the path with a forward slash. <p>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.</p> <pre data-bbox="352 695 1033 716">/client1/client1.NetWare File System/client1/SYS/DOC/</pre> <p>If the last component is a file, include the file extension and omit the slash from the end of the name.</p> <pre data-bbox="352 825 1139 846">/client1/client1.NetWare File System/client1/SYS/DOC/TEST.TXT</pre>
Case sensitivity	<p>All components in a path must show uppercase and lowercase letters as they appear in the actual path on the client.</p>
Wildcards	<p>Wildcard usage is the same as for Windows clients.</p> <p>See “Wildcard use in NetBackup” on page 828.</p>
All clients	<p>To back up all NetBackup for NetWare clients that are in the policy, enter only one forward slash (/) on a line:</p> <pre data-bbox="319 1151 333 1173">/</pre>
One client	<p>To back up an entire NetBackup for NetWare client, enter a forward slash (/) followed by the client name and another forward slash:</p> <pre data-bbox="319 1303 435 1324">/client1/</pre>

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: 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*.

Pathname 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.

To back up NetWare Target clients, use the following conventions for entries in the backup selections list.

Table 16-41 Conventions for specifying pathnames for NetWare Target clients

Item	Description
Targets per line	Enter one target per line.
Forward slash	Begin all target names with a forward slash (/).
Case sensitivity	All target names must be in uppercase.
Wildcards	Follow the wildcard use as described in the following topic for Windows clients. See "Wildcard use in NetBackup" on page 828.

The following example backs up the targets: `NETWARE`, `SYSTEM`, and `BINDERY`:

```
/NETWARE/  
/SYSTEM/  
/BINDERY/
```

For more information, see the *NetBackup Administrator's Guide for Novell NetWare Clients*.

Pathname rules for the 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.)

About the directives on the Backup Selections list

Directives on the **Backup Selections** list signal NetBackup to perform specific, predefined actions when it processes the files on 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 **MS-Windows** policy type is selected, and **Allow multiple data streams** is enabled.

```
NEW_STREAM
D:\Program Files
NEW_STREAM
C:\Winnt
```

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.

The following table summarizes many of the directives available on the **Backup Selections** list.

Table 16-42 Summary of directives on the Backup Selections list

Directive	Description
ALL_LOCAL_DRIVES	Instructs NetBackup to back up all local drives except for those drives that use removable media. This directive gives different results depending on whether Allow multiple data streams is enabled for the policy. See “ALL_LOCAL_DRIVES directive” on page 621.
System_State:\	Instructs NetBackup to back up critical system-related components. The exact set of system components that are backed up depends on the operating system version and system configuration. See “System_State:\ directive” on page 623.
Shadow Copy Components:\	Instructs NetBackup to back up all writers for the Volume Shadow Copy component. See “Shadow Copy Components:\ directive” on page 624.

Table 16-42 Summary of directives on the Backup Selections list (*continued*)

Directive	Description
Policy-specific directives	Apply only to specific policy types and can appear only in backup selections lists for those policies. See “Directives for specific policy types” on page 625.
UNSET and UNSET_ALL	Interrupt the streaming of policy-specific directives. The Allow multiple data streams policy attribute must be enabled before these directives can be used. See “UNSET and UNSET_ALL directives” on page 628.
NEW_STREAM	When on the first line of the Backup Selections list, this directive determines how a backup is performed in the following modes: <ul style="list-style-type: none"> ■ Administrator-defined streaming ■ Auto-discovery streaming <p>The Allow multiple data streams policy attribute must be enabled before this directive can be used.</p> <p>See “NEW_STREAM directive” on page 625.</p>

ALL_LOCAL_DRIVES directive

Use the `ALL_LOCAL_DRIVES` directive to back up all local drives except for those drives that use removable media. If this directive is used, this directive must be the only entry in the backup selections list for the policy. No other files or directives can be listed. The directive applies only to the following policy types:

- Standard (except for NetWare target clients)
- MS-Windows
- NetWare
 - Only for NonTarget clients
 - Only when **Allow multiple data streams** is disabled

`ALL_LOCAL_DRIVES` gives different results depending on whether **Allow multiple data streams** is enabled for the policy:

- Allow multiple data streams enabled** Applies only to Standard (except for NetWare target clients) and 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.
- Allow multiple data streams disabled** NetBackup backs up the entire client and includes all drives and file systems in the same stream.

See “Allow multiple data streams (policy attribute)” on page 542.

Caution: Do not select **Cross mount points** for policies where you use the `ALL_LOCAL_DRIVES` directive.

See “`ALL_LOCAL_DRIVES` example: Auto-discovery mode” on page 622.

See “`ALL_LOCAL_DRIVES` example: Without multiple data streams” on page 622.

ALL_LOCAL_DRIVES example: Auto-discovery mode

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 the following directive:

```
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 “`ALL_LOCAL_DRIVES` example: Without multiple data streams” on page 622.

See “Allow multiple data streams (policy attribute)” on page 542.

ALL_LOCAL_DRIVES example: Without multiple data streams

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 the following directive:

```
ALL_LOCAL_DRIVES
```

For this backup selections list, 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 (policy attribute)” on page 542.

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. This 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.

Note: 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 | <ul style="list-style-type: none">■ System files■ COM+ Class Registration Database■ SYSVOL■ Active Directory■ Cluster quorum■ Certificate Services■ Registry■ Internet Information Services |
| System Service writers | <ul style="list-style-type: none">■ 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 | Items that the computer does not require to operate. For example, Active Directory Application Mode. |

Other Data A category that is intended for future NetBackup releases.

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. All policy-specific directives that are passed to a client in a stream are passed in all subsequent streams.

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.

NEW_STREAM directive

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.

In this 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” on page 621.
- 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 “Wildcard use in NetBackup” on page 828.

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. Depending on the length of the preprocess interval, preprocessing may not occur before every backup.

About 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 has 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 problems can occur:

Interval 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.
Interval 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.

Use the following form of the `bpconfig` command to change the interval:

```
install_path\NetBackup\bin\admincmd\bpconfig [-prep hours]
```

For more information on the `bpconfig` command, see the *NetBackup Commands Reference Guide*.

UNSET and UNSET_ALL directives

`UNSET`, `UNSET_ALL` The `UNSET` and `UNSET_ALL` directives interrupt the streaming of policy-specific directives.

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 625.

See “Allow multiple data streams (policy attribute)” on page 542.

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 set destpath=/etc/home [or UNSET_ALL]
/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.

Files that are excluded from backups by default

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.

See “About excluding files from automatic backups” on page 630.

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)
- `cachets` (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)

See “Follow NFS (policy attribute)” on page 528.

See “Cross mount points (policy attribute)” on page 533.

About 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 list and the include list depends on the type of client as follows:

Microsoft Windows clients	<p>Specify exclude and include lists in the Backup, Archive, and Restore client interface. Start Backup, Archive, and Restore. On the File menu, click NetBackup Client Properties. Select the Exclude List tab or the Include List tab. For further instructions, see the NetBackup user's guide for the client.</p> <p>The Exclude List or the Include List can also be specified through the NetBackup Administration Console on the master server.</p> <p>See "Exclude Lists properties" on page 115.</p>
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>
UNIX clients	<p>Create the exclude and include lists in the following files on the client:</p> <ul style="list-style-type: none">■ <code>/usr/opensv/netbackup/include_list</code>■ <code>/usr/opensv/netbackup/exclude_list</code>

Files that are excluded by Microsoft Windows Backup

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.

Disaster Recovery tab

The **Disaster Recovery** tab appears when you select the **NBU-Catalog** policy type on the **Attributes** tab. 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 computer. Symantec recommends that the image file be saved to a network share or a removable device.

Table 16-43 describes the options on the **Disaster Recovery** tab.

Table 16-43 Options on the Disaster Recovery tab

Option	Description
Path	<p>Specify the directory where the disaster recovery information is to be saved. Do not save the disaster recovery information to the local computer. Symantec recommends that you save the image file to a network share or a removable device.</p> <p>The share must be established and available before the hot catalog backup runs.</p> <p>Specify an NFS share or a UNC path (CIFS Windows share).</p> <p>When indicating a UNC path, note the following:</p> <ul style="list-style-type: none"> ■ A Windows master server can indicate a UNC path to a Windows computer. ■ A UNIX master server cannot indicate a UNC path to a Windows computer. ■ 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	<p>Specify the logon and password information that is required to access an established Windows or NFS share.</p> <p>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.</p>
Password	<p>Specify the password that is required to log on to the share.</p>

Table 16-43 Options on the Disaster Recovery tab (*continued*)

Option	Description
Send in an email attachment	<p>Specify the email address where the disaster recovery report should be sent. 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:</p> <pre>email1@domain.com,email2@domain.com</pre> <p>See “Setting up email notifications about backups” on page 135.</p> <p>The <code>nbmail.cmd</code> or <code>mail_dr_info.cmd</code> script must be configured (<code>Install_path\NetBackup\bin\goodies\</code>). In addition specify the email addresses in the Disaster Recovery tab.</p> <p>NetBackup performs the notification by passing the email addresses, subject, and message to <code>nbmail.cmd</code> or <code>mail_dr_info.cmd</code>. 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.</p> <p>The following points describe how <code>mail_dr_info.cmd</code> and <code>nbmail.cmd</code> interact:</p> <ul style="list-style-type: none"> ■ If <code>Install_path\NetBackup\bin\mail_dr_info.cmd</code> is configured, the disaster recovery report is sent to the email address of the administrators 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 <code>mail_dr_info.cmd</code> is not configured, and <code>Install_path\NetBackup\bin\goodies\nbmail.cmd</code> is not configured, the disaster recovery report is sent to the administrators that are indicated in the Disaster Recovery tab by <code>nbmail.cmd</code>. ■ If neither file is configured, NetBackup attempts to use Microsoft internal IMAPI services. <p>Note: By default, neither <code>nbmail.cmd</code> nor <code>mail_dr_info.cmd</code> is configured to send email.</p> <p>See “Configuring the <code>nbmail.cmd</code> script” on page 136.</p> <p>For more information on <code>mail_dr_info.cmd</code>, see the <i>NetBackup Administrator’s Guide, Volume II</i>.</p>
Critical policies	<p>Lists the policies that are 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.</p> <p>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.</p>

Note: Vault protects the disaster recovery data by sending the data to the Vault site as an email attachment of the Vault report email.

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**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Do one of the following:
 - Double-click a configured catalog backup policy.
 - Create a catalog backup policy.
See “Configuring a catalog backup manually” on page 673.
- 3 Select the **Disaster Recovery** tab.
- 4 Near the **Critical Policies** list, click **Add**. An active field appears in the list.
- 5 Click the icon at the far right of the active field to display a list of configured policies. Select a policy to add to the **Critical Policies** list.
- 6 Do any of the following:

To add another policy Click **Add**.

To change a policy Select the policy and click **Change**.

To delete a policy Select the policy and click **Delete**.

- 7 Click **OK** to save the **Critical policies** list and the other settings on the **Disaster Recovery** tab.

Creating a Vault policy

A Vault policy differs from other policies in the following respects:

- **Vault** must be specified as the policy type.
- No clients are specified in Vault policies, so the **Clients** tab does not appear.
- In the **Backup Selections** list, a Vault command is specified instead of files.

To create a Vault policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > New Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- 4 On the **Attributes** tab, select **Vault** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule. The type of backup defaults to **Automatic**.

The **Clients** tab does not appear for Vault policy types.

- 6 Complete the schedule.
- 7 On the **Backup Selections** tab, enter one of two Vault commands:

`vltrun` 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
```

`vlteject` 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*.

- 8 Click **OK**.

Performing manual backups

A manual backup is user-initiated and is 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 new software is installed 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 is used only for manual backups. Create a policy for manual backups 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**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Select the policy name in the left pane.
- 3 On the **Actions** menu, click **Manual Backup**. (To perform a manual backup, you must enable the **Active. Go into effect at** attribute.)

See “Go into effect at (policy attribute)” on page 527.

If the **Go into effect at** attribute 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 it does 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 use NetBackup to restore individual objects and attributes in the Active Directory instead of restoring the entire Active Directory.

Administrators can also restore deleted objects (tombstone objects) from the Active Directory.

The following topics describe 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 NetBackup backups and recovery

Active Directory granular NetBackup 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 899.
See “About configuring Services for Network File System (NFS) on the Windows 2003 R2 SP2 NetBackup media server and NetBackup clients” on page 908.
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 900.
- 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 Legacy Client Service (`bpinetd`) must run under the domain administrator account on the Active Directory domain controller or ADAM server. By default, `bpinetd` runs under the Local System account.

See “Configuring the log on account for the NetBackup Client Service for Windows” on page 916.

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 NetBackup 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 Legacy Client Service (`bpinetd`) 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 for Windows” on page 916.

- 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 639.

- 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 623.
 - See “Shadow Copy Components:\ directive” on page 624.
 - See “ALL_LOCAL_DRIVES directive” on page 621.

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 NetBackup 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.

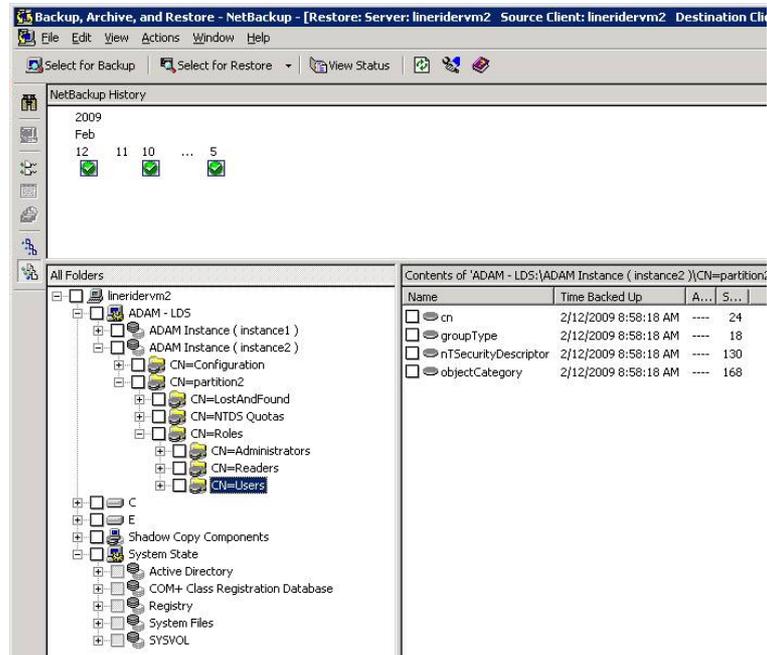
Restoring Active Directory objects

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 “Troubleshooting granular restore issues” on page 641.

Troubleshooting granular restore issues

Some granular restore situations require additional steps to fully restore the objects. In other situations, a granular restore of some part of the Active Directory is not possible.

Table 16-44 describes potential problems for granular restores.

Table 16-44 Troubleshooting restore issues

Situation	Recommendation
Restores that are disabled	<p>When user and computer accounts are restored from a granular Active Directory restore, they are sometimes disabled.</p> <p>The following are possible reasons why the accounts can be disabled:</p> <ul style="list-style-type: none"> ■ 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. ■ When restoring user objects, you must reset the object's user password and enable the object's user account: <ul style="list-style-type: none"> ■ For Active Directory user objects, use the Microsoft Active Directory Users and Computers application. ■ For ADAM/AD LDS user objects, use ADSI Edit. <p>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.</p> <ul style="list-style-type: none"> ■ 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. <p>Use the Microsoft Active Directory Users and Computers application to reset the account of a computer object:</p> <ul style="list-style-type: none"> ■ Remove the computer from the domain. ■ 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.

Table 16-44 Troubleshooting restore issues (*continued*)

Situation	Recommendation
Group and member objects	<p>To restore Active Directory group membership links may require that the restore job be run twice.</p> <p>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.</p> <p>Run the restore again to restore all forward and backward links.</p>
Group policy objects	NetBackup does not support granular restores of Group Policy Objects.

Synthetic backups

This chapter includes the following topics:

- About synthetic backups
- Recommendations for synthetic backups and restores
- Synthetic full backups
- Synthetic cumulative incremental backups
- Schedules that must appear in a policy for synthetic backups
- Adding clients to a policy for synthetic backups
- Change journal and synthesized backups
- True image restore 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
- Optimized synthetic backups for deduplication

About synthetic backups

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.

Because processing takes place on master and media servers instead of the client, synthetic backups help to reduce the network traffic. 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.

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.

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 (policy attribute) with and without move detection” on page 538.
- The schedule that is created for a synthetic backup must have **Synthetic Backup** selected.
See “Synthetic backup (schedule attribute)” on page 559.
- 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 “Recommendations for synthetic backups and restores” on page 646.
 - A combination of disk storage unit(s) and tape library(s).

Recommendations for synthetic backups and restores

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.

Synthetic backups are supported on all media server platforms and tier one master server platforms.

The following items describe recommendations to use synthesized backups to full advantage, and situations under which synthesized backups are not supported:

Recommendations concerning backups:

- Do not multiplex any backups that are to be synthesized because it is inefficient. To synthesize multiplexed client images requires multiple passes over the source media—one per client.
Performance issues can also occur if multiple streams are selected for synthesized backups. The issues are similar to those encountered while multiplexing synthesized backups. 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.
- The option to create multiple copies is allowed for synthetic backups using the multiple copies synthetic backup method.
See “Using the multiple copy synthetic backups method” on page 655.
- 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.

Recommendations concerning restores:

- 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.
- 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.

Table 17-1 Recommendations when using disk storage or tape storage for synthetic backups

Storage unit type	Recommendations
Disk storage units	<p>Disk-based images are more efficient for synthesizing. 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 inefficient compared to disk-based images.</p> <p>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.</p>
Tape storage units	<p>If tape is used instead of disk, the tape for the synthetic image must be different from the tape where the component images reside.</p> <p>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.</p> <p>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.</p>

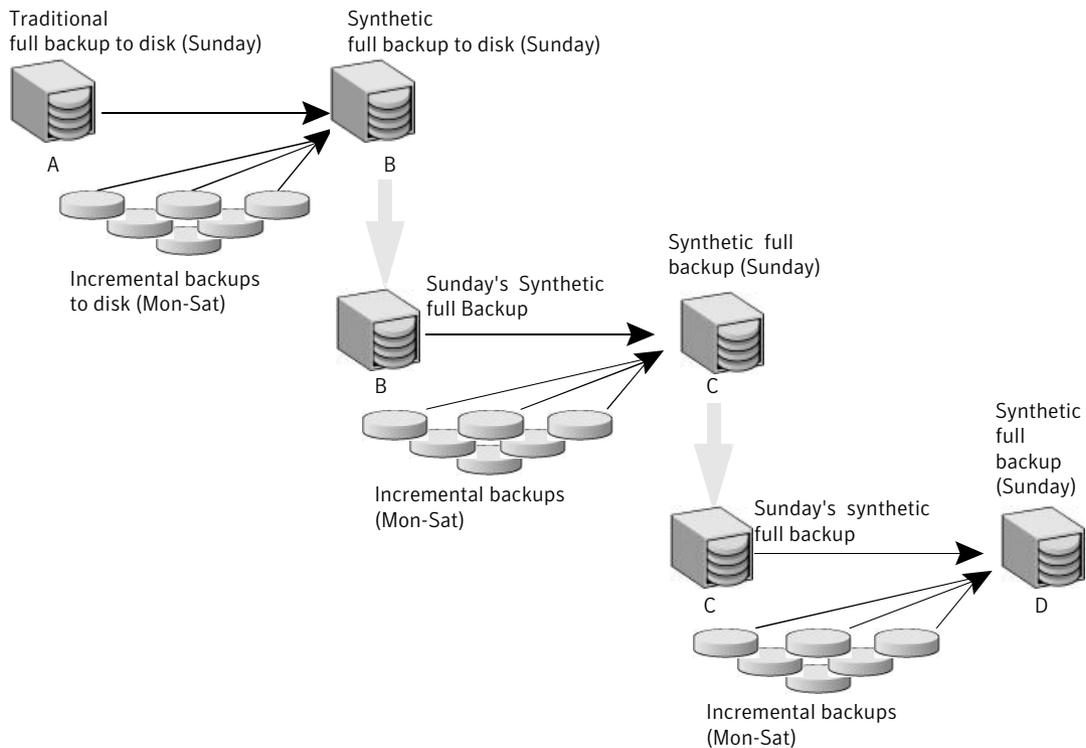
Synthetic full backups

A synthetic backup can be a synthetic full or a synthetic cumulative backup.

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.

Figure 17-1 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 17-1 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 650.

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

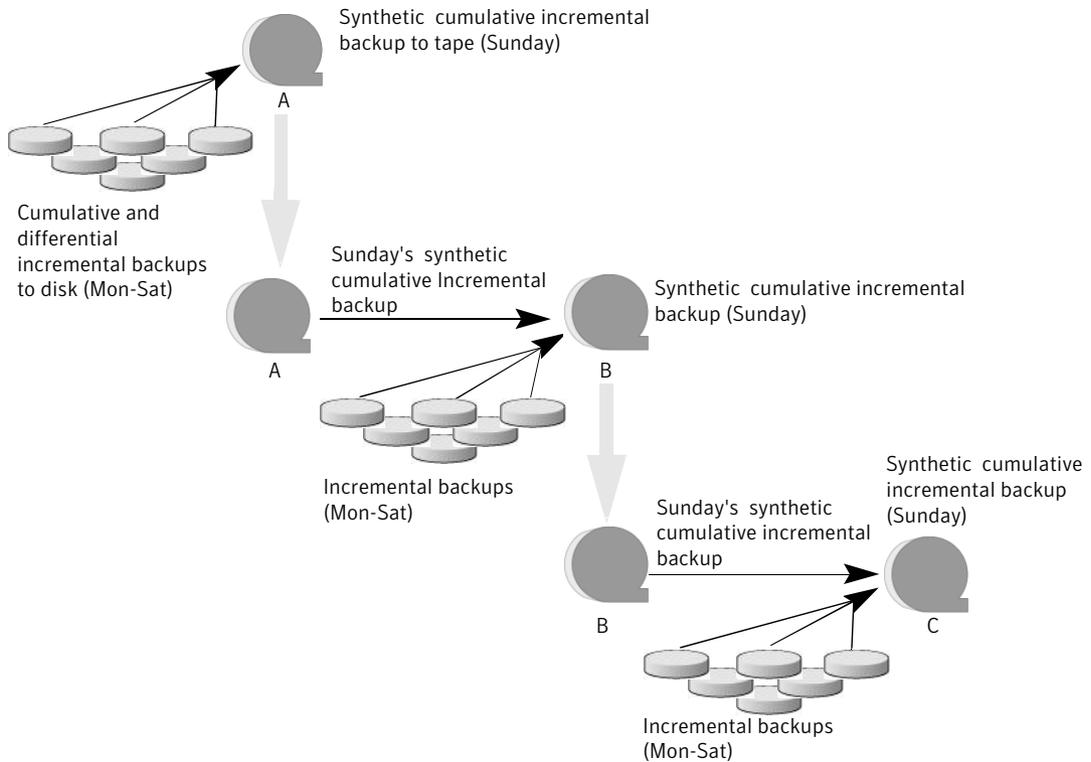
The scenario to create a synthetic, cumulative incremental backup is similar to the scenario to create a synthetic full backup. Remember, a cumulative incremental backup includes all changes since the last full backup.

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 17-2 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 17-2 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.

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 (schedule attribute)” on page 559.

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 (policy attribute) with and without move detection” on page 538.

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 “Client Settings properties for Windows clients” on page 96.

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 (policy attribute) with and without move detection” on page 538.

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 “Clean-up properties” on page 75.

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.

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 “Client Settings properties for Windows clients” on page 96.

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 17-2 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 17-2 Comparing synthetic copy process with and without method enabled

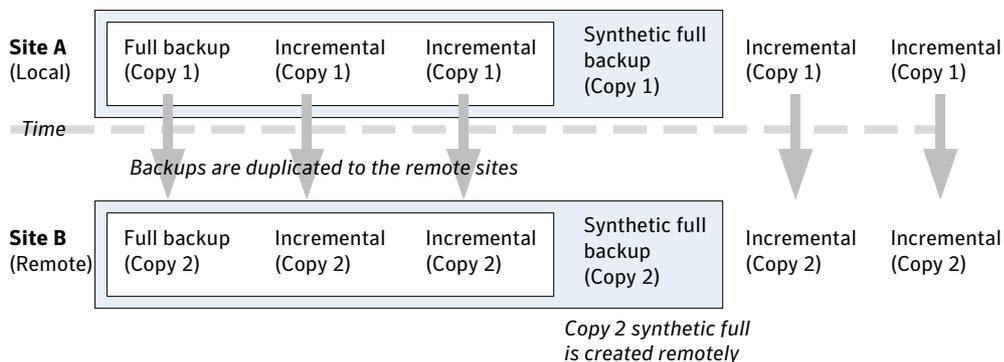
Step	Without using the multiple copy synthetic backups method:	Using the multiple copy synthetic backups method:
1	A full backup is performed at the local site (Site A).	Step 1 remains the same.
2	The full backup is duplicated to the remote site (Site B).	Step 2 remains the same.
3	An incremental backup 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.

Table 17-2 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:
7	The full backup is duplicated to Site B.	A full synthetic backup is produced at Site B from images at Site B. 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 17-3 shows how no extra bandwidth is used to copy the synthetic full backup from Site A to Site B.

Figure 17-3 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 17-3 describes the configuration entries that can be included in the configuration file.

Table 17-3 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. If COPY is specified and either SRC_COPY and TARGET_COPY is not specified, the value for COPY is used.
TARGET_STU	Specifies the storage unit name or storage unit group name where the second copy synthetic backup is to be written. Use the special identifier __ANY__ 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 ANY: TARGET_STU=__ANY__
FAIL_MODE	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 FAIL_MODE value specifies the fate of the first copy job and image. Specify one of the following: <ul style="list-style-type: none"> ■ FAIL_MODE=ALL ALL means that if the second copy fails, the first copy and its job also fail. (Default.) ■ FAIL_MODE=ONE ONE means that if the second copy fails, the failure does not affect the first copy job.

Table 17-3 Configuration entries (*continued*)

Entry	Purpose
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"> ■ ENABLED=YES Production of the second copy is enabled. (Default.) ■ ENABLED=NO Production of the second copy is disabled.
SRC_COPY_FALLBACK	<p>Specifies that if a copy by the number given in SRC_COPY or COPY does not exist, the synthetic backup should use the primary backup.</p> <p>The only valid value is the following:</p> <p>SRC_COPY_FALLBACK=PRIMARY</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>

Configuration examples

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.

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.

Optimized synthetic backups for deduplication

NetBackup environments that use the NetBackup Deduplication Option license key environment can benefit from the optimized synthetic backup method.

This method constructs the synthetic image by using calls from the backup server to the storage server. The backup 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 Deduplication Guide* for more information.

Protecting the NetBackup catalog

This chapter includes the following topics:

- About NetBackup catalogs
- Parts of the NetBackup catalog
- Protecting the NetBackup 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 clustered environments, see the *NetBackup Clustered Master Server Administrator's Guide*.

Parts of the NetBackup catalog

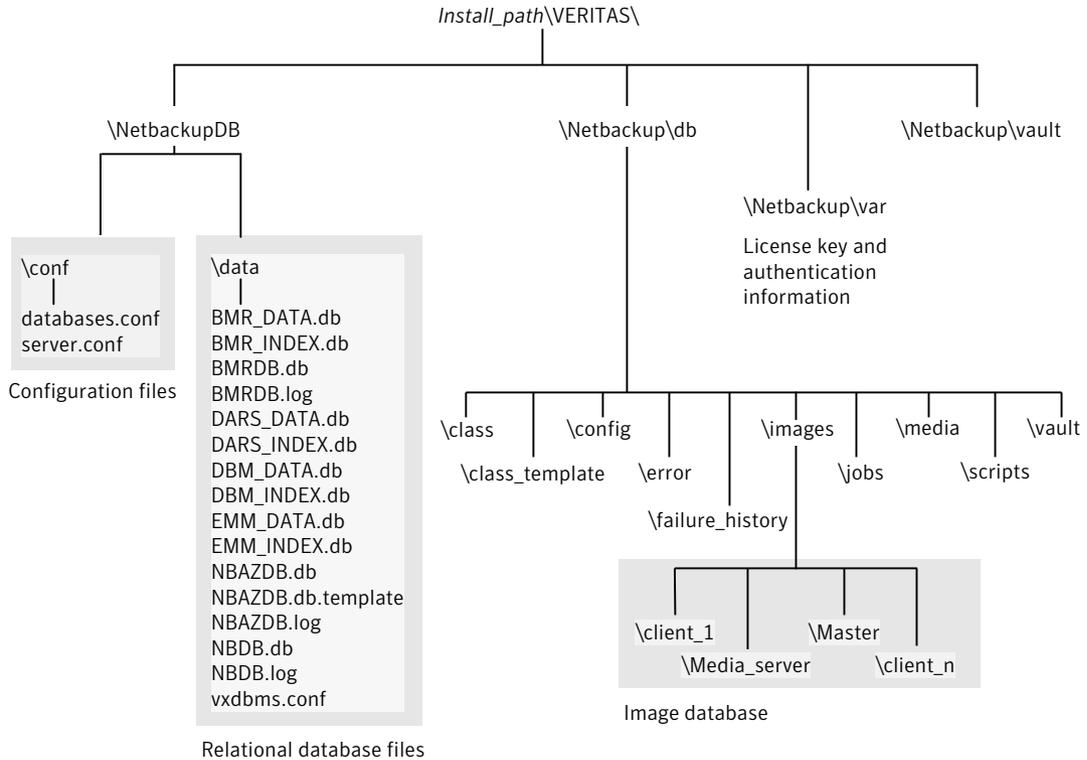
The NetBackup catalog resides on the NetBackup master server.

Figure 18-1 shows the default files and directories in a NetBackup catalog.

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 NetBackup image database” on page 663.
- 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 665.
- 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 NetBackup server.conf file” on page 696.
See “About the databases.conf file” on page 697.

Figure 18-1 Catalog configuration



About the NetBackup 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 .*ε* 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 684.

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 687.

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 “Estimating catalog space requirements” on page 684.

About NetBackup 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 NetBackup 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. The NetBackup database (NBDB) is also known as the Enterprise Media Manager (EMM) database. It 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 database (including the NetBackup Authorization database, NBAZDB) 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*
- *Install_path\VERITAS\NetBackupDB\data\NBAZDB.db*
- *Install_path\VERITAS\NetBackupDB\data\NBAZDB.db.template*
- *Install_path\VERITAS\NetBackupDB\data\NBAZDB.log*

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) 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.

See “About the NetBackup image database” on page 663.

See “About the NetBackup relational database (NBDB) installation” on page 693.

See “Post-installation tasks” on page 721.

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 NetBackup relational database” on page 665.

See “Moving the NetBackup database from one host to another” on page 731.

Protecting the NetBackup 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 679.

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

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. For Sybase SQL Anywhere, an incremental backup means a backup of the transaction log only. Transaction logs are managed automatically and truncated after each successful backup.

The online, hot catalog lets you recover either the entire catalog or pieces of the catalog. (For example, the databases separately from the image catalog.)

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 669.
 - The Backup Policy Configuration Wizard.
See “Using the Backup Policy Wizard to configure a catalog backup” on page 672.
Either wizard automatically includes all the necessary catalog files to include the database files (NBDB, NBAZDB, and BMRDB) and any catalog configuration files (`vxdbms.conf`, `server.conf`, `databases.conf`).
- By creating a backup policy manually and indicating the **NBU-Catalog** policy type.
See “Configuring a catalog backup manually” on page 673.

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.

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. 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 “Concurrently running online, hot catalog backups with other backups” on page 676.

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 computer.

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 678.

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.

The NetBackup **Disaster Recovery** report lists the media that is used for backups of critical policies. The report lists the media for only incremental and full backup schedules, so critical policies should use only incremental or full backup schedules.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

See “Determining whether or not a catalog backup succeeded” on page 677.

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.

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 **Next**.

- 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 “Concurrently running online, hot catalog backups with other backups” on page 676.

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 computer.

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 678.

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.

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 report 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. 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.

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:
- 5 Select the **Schedules** tab to set up a schedule for an online catalog backup. See “Concurrently running online, hot catalog backups with other backups” on page 676.
See “About catalog policy schedules” on page 676.

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:

- 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 computer.
- 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 report lists the media for only incremental and full backup schedules, so critical policies should use only incremental or full backup schedules.
- Click **OK** to save the policy.

Backing up NetBackup 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 scheduled 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 636.

You can also run the `bpbackup` command from the command line to perform an online, hot catalog backup.

More information is available in the *NetBackup Commands Reference Guide*.

See “About online, hot catalog backups” on page 668.

See “Configuring a catalog backup manually” on page 673.

Concurrently running online, hot catalog backups 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.
- Increase the **Maximum concurrent jobs** setting on the storage unit where the backups are sent.

See “Determining whether or not a catalog backup succeeded” on page 677.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

About catalog policy schedules

When you work with catalog policy schedules, consider the following:

- The schedules that are supported in the online, hot catalog backup policy type are as follows:
 - 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 incremental schedules that are session-based:
 - 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.

See “About online, hot catalog backups” on page 668.

Determining whether or not 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 the *Administrator's Guide, Volume II* for more information on setting up this script.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

Strategies that ensure successful NetBackup 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 sensitive to the location of the catalog. Storing the catalog on a remote file system may have critical performance issues for catalog backups. 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.

See “About NetBackup catalogs” on page 661.

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.

Archiving the catalog

The catalog archiving feature helps administrators solve the kinds of 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 680.

- 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 2011 | 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 2011 | bpcatres
```

This command restores all the catalog archive files before Jan 1, 2011.

See “Catalog archiving commands” on page 681.

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**.
If Vault is used, the files are duplicated and vaulted similarly to other backups.
- 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 **Backup Selections** 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.

Catalog archiving uses the following commands.

Table 18-1 Catalog archiving commands

Command	Description
bpcatlist	<p>The <code>bpcatlist</code> command queries the catalog data. Then, <code>bpcatlist</code> 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. <code>bpcatlist</code> outputs the formatted image summary information of matched images to standard output.</p> <p>The other catalog archiving commands, <code>bpcatarc</code>, <code>bpcatrm</code>, and <code>bpcatres</code>, all depend on input from <code>bpcatlist</code> by a piped command.</p> <p>For example, to archive (backup and delete) all of the <code>.f</code> files that were created before January 1, 2010, the following would be entered:</p> <pre>Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist -client all -before Jan 1 2011 bpcatarc bpcatrm</pre> <p><code>bpcatlist</code> is also used to provide status information.</p> <p>For each catalog, it lists the following information:</p> <ul style="list-style-type: none"> ■ Backup ID (Backupid) ■ Backup date (Backup Date) ■ Catalog archive ID (catarcid). After one <code>.f</code> 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) <p>The following is an example of the <code>bpcatlist</code> output, showing all of the backups for client alpha since October 23:</p> <pre># bpcatlist -client alpha -since Oct 23 Backupid Backup Date ...Catarcid S C Files file alpha_0972380832 Oct 24 10:47:12 2010 ... 973187218 1 0 alpha_0972380832_UBAK.f alpha_0972336776 Oct 23 22:32:56 2010 ... 973187218 1 0 alpha_0972336776_FULLL.f alpha_0972327197 Oct 23 19:53:17 2010 ... 973187218 1 0 alpha_0972327197_UBAK.f</pre> <p>More information is available in the <i>NetBackup Commands Reference Guide</i>.</p>
bpcatarc	<p>The <code>bpcatarc</code> command reads the output from <code>bpcatlist</code> and backs up the selected list of <code>.f</code> files. After one <code>.f</code> file is successfully backed up, a catalog archive ID is entered into the catarcid field in the image file. For archiving of the <code>.f</code> 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 <code>bpcatarc</code> command is run.</p> <p>See “Creating a catalog archiving policy” on page 680.</p>

Table 18-1 Catalog archiving commands (*continued*)

Command	Description
bpcatrm	The <code>bpcatrm</code> command reads the output from <code>bpcatlist</code> or <code>bpcatarc</code> . If the image file has valid catarcid entries, <code>bpcatrm</code> deletes selected image <code>.f</code> files from the online catalog. <code>bpcatrm</code> does not remove one <code>.f</code> file unless the file has been previously backed up using the catarc policy.
bpcatres	Use the <code>bpcatres</code> command to restore the catalog. The <code>bpcatres</code> command reads the output from <code>bpcatlist</code> and restores selected archived <code>.f</code> files to the catalog. For example: <pre>Install_path\VERITAS\NetBackup\bin\admincmd\bpcatlist -client all -before Jan 1 2011 bpcatres</pre>

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 186.

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 \times 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. 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 number of bytes per file record.

If you are unsure of the average number of bytes per file record, use 132. The results from the examples in step 3 yield:

$$(8400 \times 132) + (12,000 \times 132) = 2692800 \text{ bytes (or about 2630 kilobytes)}$$

- 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.

NetBackup 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.

See “Estimating catalog space requirements” on page 684.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

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 improve 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

See “About NetBackup image .f files” on page 664.

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.

See “Backing up NetBackup catalogs manually” on page 675.

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 Daemon and NetBackup Database Manager service by using the **Services** tab in the Activity Monitor.

Backups and restores can now resume for the client.

See “NetBackup file size considerations” on page 686.

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.

To index image header files, run the following command:

```
bpimage -create_image_list [-client name]
```

Run this command to create the following index files in each client image directory:

```
IMAGE_FILES
```

```
IMAGE_INFO
```

```
IMAGE_LIST
```

To stop image header indexing for a client, remove these files.

Additional information is available about `bpimage` in the *NetBackup Commands Reference Guide*.

See “About NetBackup catalogs” on page 661.

About image catalog compression

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 “Global Attributes properties” on page 131.

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:

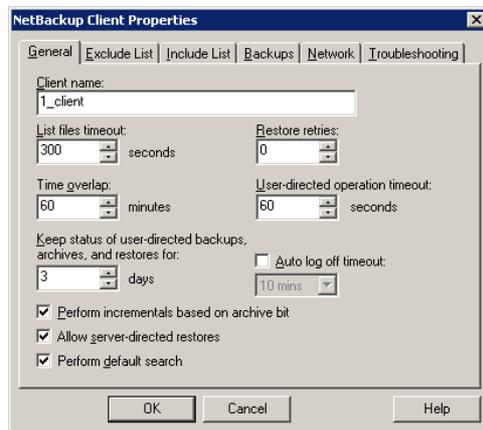
- Reduces greatly the disk space that is consumed.
- Reduces the media that is required to back up the catalog.

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 18-2 List Files Timeout General property on the client



See “Uncompressing the NetBackup catalog” on page 690.

See “Indexing the catalog for faster access to backups” on page 688.

Uncompressing the NetBackup 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.

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 684.
- 2 Stop the NetBackup Request Daemon service, `bprd`. Use the Activity Monitor or the Services application in the Windows Control Panel.
- 3 Verify that the NetBackup Database Manager, `bpdbm`, 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 Daemon `bprd`.
See “About image catalog compression” on page 689.

About the NetBackup relational database

This chapter includes the following topics:

- About the NetBackup relational database (NBDB) installation
- Using the NetBackup Database Administration utility
- Post-installation tasks
- About backup and recovery procedures
- Unloading the NetBackup database
- Terminating database connections
- Moving the NetBackup database from one host to another

About the NetBackup relational database (NBDB) installation

The following information can help you to install and operate 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). NBDB contains the NetBackup Authorization database, 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 computer as the Enterprise Media Manager.

See “About the Enterprise Media Manager” on page 832.

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 722.

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) to a remote file system such as NFS or CIFS.

The following procedure is performed automatically during installation in the order presented. You can also use the same procedure to manually install the database independently.

Installing the NetBackup database

- 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 NetBackup server.conf file” on page 696.

- 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 701.

- 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 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 authorization database files for NetBackup Access Control:

```
Install_path\VERITAS\NetBackupDB\data\NBAZDB.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
```

See “About online, hot catalog backups” on page 668.

About NetBackup master server installed directories and files

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.

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 722.

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) 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.

See “Moving NBDB database files after installation” on page 722.

See “Moving the NetBackup database from one host to another” on page 731.

See “Moving the NetBackup database files” on page 713.

About the NetBackup 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"
```

In this example, `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.

Table 19-1 Commands used in the server.conf file

Command	Description
<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 that use port 13785.

Table 19-1 Commands used in the server.conf file (continued)

Command	Description
-gp 4096	Indicates the maximum page size (in bytes) for the database. This parameter is given during database creation.
-ct+	Indicates that character set translation is used. UTF8 encoding is used.
-gd DBA -gk DBA -gl DBA	Indicates that the DBA user is the account used to start, stop, load, and unload data.
-ti 0	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.
-c 25M	Indicates the initial memory that is reserved for caching database pages and other server information. The value may be changed for performance reasons.
-ch 500M	Indicates the maximum cache size, as a limit to automatic cache growth. The value may be changed for performance reasons.
-cl 25M	Indicates the minimum cache size, as a limit to automatic cache resizing. The value may be changed for performance reasons.
-gn 10	Indicates the number of requests the database server can handle at one time. This parameter limits the number of threads upon startup. The value may be changed for performance reasons.
-o <i>Install_path</i> \VERITAS\ NetBackupDB\log\server.log	Indicates the location of server output messages. The messages include start and stop events, checkpoints, error conditions, and cache change size. This log is not managed, but growth is slow.
-ud	Indicates that the server should run as a daemon.
-ec SIMPLE	Indicates the encryption method. Default: SIMPLE. NONE SIMPLE TLS (TLS_TYPE=cipher; [FIPS={Y N}]) CERTIFICATE=server-identity-filename; CERTIFICATE=PASSWORD=password)

See “About NetBackup master server installed directories and files” on page 695.

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\NBAZDB.db" -n NBAZDB
```

```
"C:\Program Files\VERITAS\NetBackupDB\data\BMRDB.db" -n BMRDB
```

See “About NetBackup master server installed directories and files” on page 695.

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
```

See “About NetBackup master server installed directories and files” on page 695.

About the bin directory

`NetBackup\bin` contains NetBackup-specific binaries and commands for administrating NBDB and BMRDB:

- `NbDbAdmin.exe`
This file launches the NetBackup Database Administration utility, which provides administrators with a way to more easily perform the tasks based on the `nbdb` commands.
See “Using the NetBackup Database Administration utility” on page 703.
- `create_nbdb.exe`
Used during installation and upgrades to create and upgrade the NetBackup database, NBDB.
- `nbdb_admin.exe`
Among other things, use `nbdb_admin.exe` to change the DBA and NetBackup account passwords, or to start and stop individual databases.
- `nbdb_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 NetBackup Database Administration utility) 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 NetBackup Database Administration utility) to restore the NetBackup catalog only as a precautionary measure.

- `nbdb_move.exe`
Use to change the location of the SQL Anywhere database files from the default location.
- `nbdb_ping.exe`
Displays the status of the SQL Anywhere database.
- `nbdb_restore.exe`
Use to recover from an online backup in a file system directory that was created using `nbdb_backup`.
- `nbdb_unload.exe`
Use to create a dump of all or part of the NBDB database or the BMRDB database schema and data.
- `nbdbms_start_server.exe`
Use to start and stop the SQL Anywhere service.
- `nbdb_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 the *NetBackup Commands Reference Guide* and the online Help.

See “Using the NetBackup Database Administration utility” on page 703.

See “About NetBackup master server installed directories and files” on page 695.

About the content of the NetBackup directories

The following table describes the contents of the NetBackup directories.

Table 19-2 NetBackup directory contents

Directory	Description
Charsets	The directory <i>Install_path\VERITAS\NetBackupDB\Charsets</i> contains SQL Anywhere-specific information.
log	The directory <i>Install_path\VERITAS\NetBackupDB\log</i> contains the SQL Anywhere server log file <i>server.log</i> that contains only Sybase logs.
scripts	The directory <i>Install_path\VERITAS\NetBackupDB\scripts</i> 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.
staging	The directory <i>Install_path\VERITAS\NetBackupDB\staging</i> is used as a temporary staging area during online, hot catalog backup, and recovery.
WIN32	The directory <i>Install_path\VERITAS\NetBackupDB\WIN32</i> contains SQL Anywhere commands and .dll files.
java	Symantec OpsCenter uses the directory <i>Install_path\VERITAS\NetBackupDB\java</i> .
shared	The directory <i>Install_path\VERITAS\NetBackupDB\shared</i> is a directory used by Symantec OpsCenter.

See “About the data directory” on page 700.

See “About NetBackup master server installed directories and files” on page 695.

About the data directory

Install_path\VERITAS\NetBackupDB\data is the default location of the NetBackup database, NBDB, and includes the following files:

- **NBDB.db**
The main NetBackup database file; considered a **dbspace**.
- **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.
- **NBAZDB.db**
The NetBackup Authorization database is present whether or not NetBackup Access Control (NBAC) is configured and used.
- **EMM_DATA.db**

An additional **dbspace** that contains EMM data.

- `EMM_INDEX.db`

File that enhances the EMM database performance.

- `vxdbms.conf`

File that 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 the DBA accounts for NBDB, NBAZDB, 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 721.

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`

See “About NetBackup master server installed directories and files” on page 695.

See “About the content of the NetBackup directories” on page 699.

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`, `NBAZDB.db`, `BMRDB.db`, and the `vxdbms.conf` files are located.

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\
Config\VXDBMS_NB_DATA
```

See “About the data directory” on page 700.

See “About the content of the NetBackup directories” on page 699.

See “About NetBackup master server installed directories and files” on page 695.

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\bpdwn -e SQLANYs_VERITAS_NB`
- `Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB`

Individual databases can be started or stopped, while the SQL Anywhere service continues. To do so, use the NetBackup Database Administration utility 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 “Using the NetBackup Database Administration utility” on page 703.

See “Commands for backing up and recovering the relational databases” on page 727.

Sybase SQL Anywhere and 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 computers 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`.

See “About NetBackup master server installed directories and files” on page 695.

See “About the NetBackup relational database (NBDB) installation” on page 693.

Using the NetBackup Database Administration utility

The NetBackup Database Administration utility is a stand-alone application (`NbDbAdmin.exe`) and is located in the following directory:

`InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe`

To use the utility, you must be an administrator with administrator privileges.

When you start the NetBackup Database Administration utility, you must enter the DBA password. If you use the default password that was used during the NetBackup installation (**nbusql**), you are encouraged to change the password.

The NetBackup Database Administration utility displays the following information:

Table 19-3 NetBackup Database Administration properties

Property	Description
Database name and status	<p>Select either the NBDB or the BMRDB database to administer.</p> <p>The list of possible databases is derived from the <code>vxdbsm.conf</code> file. The <code>vxdbsm.conf</code> file is located in the directory that is specified in the <code>bp.conf</code> file or in the Windows registry parameter <code>VXDBMS_NB_DATA</code>.</p> <p>The database must reside on the same computer where the NetBackup Database Administration console runs.</p> <p>One of the following status reports display for the selected database:</p> <ul style="list-style-type: none"> ■ If the database is available, the screen displays Alive and well. ■ If the database is unavailable, the screen displays Not available.
Stop	Shuts down the selected database.
Start	Starts the selected database.
General tab	<p>Contains information about database utilization.</p> <p>See “About the General tab of the NetBackup Database Administration utility” on page 704.</p>
Tools tab	<p>Contains a variety of tools to administer the selected database.</p> <p>See “About the Tools tab of the NetBackup Database Administration utility” on page 711.</p>

Table 19-3 NetBackup Database Administration properties (*continued*)

Property	Description
Drive Space	<p>Displays 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.</p> <p>The Drive Space dialog displays the following information:</p> <ul style="list-style-type: none">■ Drive■ Capacity■ Used space■ Free space■ % Utilized■ Space
Close	Closes the Database Administration utility.
Help	Provides additional assistance in the console.

About the General tab of the NetBackup Database Administration utility

The **General** tab contains information about database space utilization. The tab contains tools to let the administrator reorganize fragmented database objects, add free space to the database files, and validate and rebuild the database.

Figure 19-1 shows the **General** tab of the Database Administration utility after a user logs on.

Figure 19-1 General tab of the NetBackup Database Administration utility

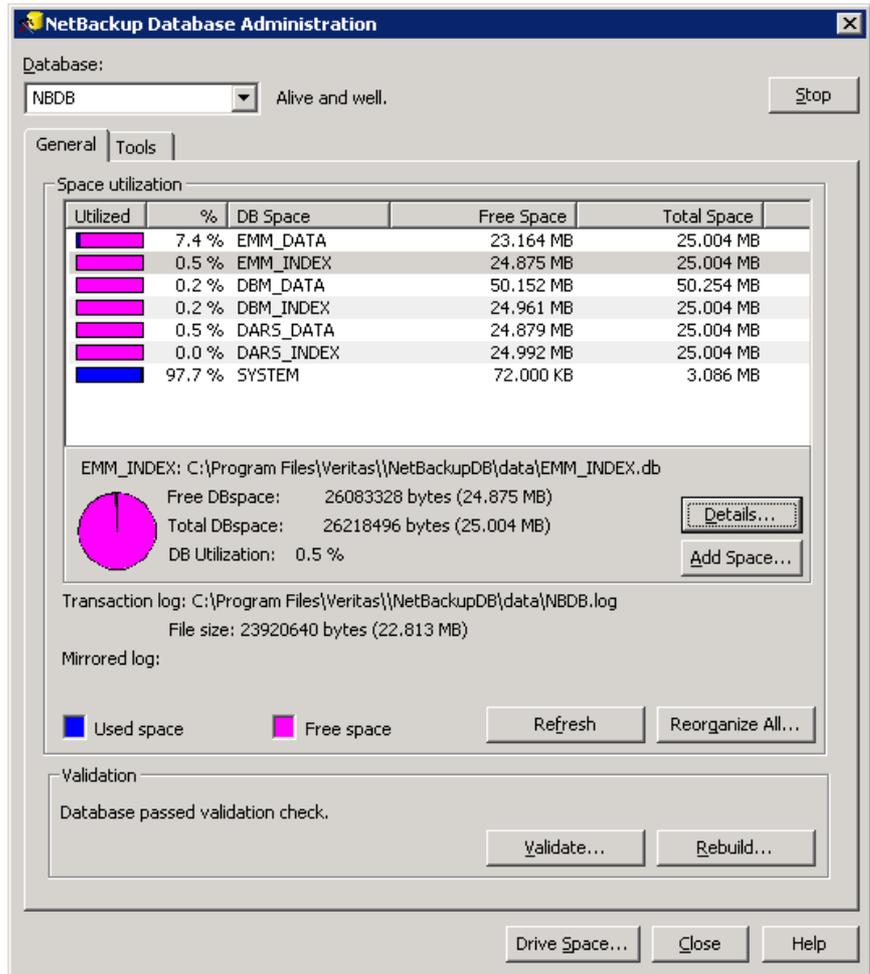


Table 19-4 General tab options

Option	Description
Space Utilization list	<p>Displays the information about used space and free space in pie chart format for the following database system, data, and index files:</p> <ul style="list-style-type: none"> ■ Free DBspace The amount of free space available. ■ Total DBspace The total amount of space that is allocated for the dbspace. ■ DB Utilization The utilization percentage is the percentage of the Total DBSpace used for data. When the NBDB or the BMRDB databases are created, extra space is allocated so that contiguous space is available as needed. As a result, a low space utilization percentage is a positive indication unless the database is very large and disk space is in short supply.
Details	<p>Use to display information about the selected database table or index file and elect to defragment (reorganize) fragmented files.</p> <p>Select a database table or an index file in the Space Utilization list and click Details.</p> <ul style="list-style-type: none"> ■ Database tables Select one or more tables in the Data (Table) Details dialog box and click Defragment. The database table is evaluated for fragmentation and if any fragmentation is detected, it is reorganized. Select one or more database tables to reorganize and then click Defragment. Each selected database table is evaluated for fragmentation and if any fragmentation is detected, it is reorganized. Table 19-5 describes the columns in the Data (Table) Details dialog box. ■ Index files Select one or more indexes in the Index Details dialog box and click Compress. Each selected index is evaluated for fragmentation and if any fragmentation is detected, it is reorganized. Table 19-6 describes the columns in the Index Details dialog box. <p>Click Close after defragmenting the files.</p>

Table 19-4 General tab options (*continued*)

Option	Description
Add Space	<p>Use to add additional free space to individual dbspaces. Additional free space helps to reduce future fragmentation of the database objects that are stored in the database.</p> <p>When the relational database is initially created or rebuilt, 25MB of free space is automatically added to the data and the index dbspaces.</p> <p>Click Add Space, then select one of the following amounts to add:</p> <ul style="list-style-type: none"> ■ A small amount of space to add is 25MB. ■ A medium amount of space to add is 50MB. ■ A large amount of space to add is 100MB. <p>Click OK to add the space or Cancel to close the dialog box.</p>
Transaction log information	The location and the file size of the transaction log.
Mirrored log information	The location and the file size of the mirrored log, if one exists.
Refresh	Displays the most current information.
Reorganize All	<p>This option automatically determines the database tables and indexes that are fragmented. The option then uses the SQL Anywhere REORGANIZE command to defragment the tables and compress the indexes.</p> <p>To click Reorganize All is equivalent to running the following command:</p> <pre>nbdb_admin.exe -reorganize</pre>
Validation status	<p>This option informs you whether or not the selected database has passed the utility's validation check.</p> <p>See Table 19-7 on page 710.</p>

Table 19-4 General tab options (*continued*)

Option	Description
Validate	<p>This option performs a database validation on all of the database tables and indexes in the selected database.</p> <p>Choose one of the following validation checks in the Validate Database dialog box:</p> <ul style="list-style-type: none"> ■ Standard The Standard validation option lets you validate the indexes and keys on all of the tables in the database. Each table is scanned, and for each row, a check is made that it exists in the appropriate indexes. The number of rows in the table must match the number of entries in the index. The equivalent command is <code>nbdb_admin.exe -validate</code> ■ Full In addition to the Standard validation checks, a Full validation ensures that every row that is referenced in each index exists in the corresponding table. For foreign key indexes, it also ensures that the corresponding row exists in the primary table. The equivalent command is <code>nbdb_admin.exe -validate -full</code> <p>Note: To perform a full database validation, shut down NetBackup and start only the database service.</p> <p>After a validation check runs, the Results screen lists each database object. Each error is listed next to the database object where it was found. The total number of errors are listed at the end of the list of database objects. If no errors were found, that is indicated.</p> <p>If any validation errors are reported, perform the following tasks:</p> <ul style="list-style-type: none"> ■ Shut down NetBackup (all daemons and services). ■ Start only the SQL Anywhere database server (<code>SQLANYs_VERITAS_DB</code>, the NetBackup Relational Database Manager). ■ Click Validate to repeat the validation check or use the <code>nbdb_admin.exe</code> command line utility. <p>If validation errors persist, contact Symantec customer support. The administrator may be asked to rebuild the database using the Rebuild option or the <code>nbdb_unload.exe</code> command line utility.</p>

Table 19-4 General tab options (*continued*)

Option	Description
Rebuild	<p>This option unloads and reloads the database. A new database with all of the same options is built in its place.</p> <p>A Database Rebuild may be required if validation errors are reported using the Standard or Full validation options using the Validate option.</p> <p>Note: Before you rebuild the database, Symantec suggests that you create a copy of the database files by performing a backup from the Tools tab.</p> <p>To rebuild the database temporarily suspends NetBackup operations and can take a long time depending on the database size.</p> <p>The equivalent command is <code>nbdb_unload -rebuild</code></p>

Table 19-5 Data (Table) Details dialog box

Column	Description
Table Name	The name of the table. The tables most in need of reorganizing are listed first.
Rows	The number of rows in the table.
Row Segments	The total number of row segments for a table. A row segment is all or part of one row that is contained on one page. A row may have one or more row segments.
Segments Per Row	The average number of segments per row. A Segments Per Row value of 1 is ideal. Any value above 1 indicates a high degree of fragmentation. For example, a value of 1.5 means that half of the rows are partitioned.
State	The state of the table. Upon opening the Data Details dialog box, the state may show as being OK (does not need defragmentation) or Fragmented (requires defragmentation). After it is reorganized, the state shows as Defragmented.

Table 19-6 Index Details dialog box

Column	Description
Table Name	The name of the table.
Index Name	The name of the index. The indexes most in need of reorganizing are listed first.

Table 19-6 Index Details dialog box (*continued*)

Column	Description
Index Type	<p>The type of the index.</p> <p>The Index Type can be one of the following values:</p> <ul style="list-style-type: none"> ■ PKEY (primary key) ■ FKEY (foreign key) ■ UI (unique index) ■ UC (unique constraint) ■ NUI (non-unique index)
Index Level	<p>The number of index levels in the index tree.</p> <p>The index level and index density indicate whether or not an index needs to be reorganized. The number of levels in the index tree determines the number of input and output operations that are needed to access a row using the index.</p> <p>Indexes with fewer levels are more efficient than indexes with greater numbers of levels. The density is a fraction between 0 and 1 providing an indication of how full each index page is on average.</p> <p>An Index Level value of 1 is ideal. An index with a value of 4 or above or with a value of 2 or 3 and an Index Density greater than 0.5 is a good candidate for reorganization.</p>
Index Density	<p>The index density and the index level indicate whether or not an index needs to be reorganized. (See the Index Level description.)</p>
State	<p>The state of the index. Upon opening the Index Details dialog box, the state may show as being OK (does not need defragmentation) or Fragmented (requires defragmentation).</p> <p>After it is reorganized, the state shows as Defragmented.</p>

Table 19-7 Validation status messages

Message	Description
Database passed validation check.	The database does not require further validation.
Not available.	No statistics on the database can be gathered because the database is not available.
Database is corrupt.	<p>Validate, then rebuild the database. Before you rebuild the database, Symantec suggests that you create a copy of the database files by doing a Backup from the Tools tab.</p> <p>Table 19-4 describes how to use the Validate option.</p>

See “About fragmentation” on page 711.

About fragmentation

Table fragmentation can impede performance. When rows are not stored contiguously, or if rows are split into more than one page, performance decreases because these rows require additional page accesses.

When an update to a row causes it to grow beyond the originally allocated space, the row is split. The initial row location contains a pointer to another page where the entire row is stored. As more rows are stored on separate pages, more time is required to access the additional pages.

Use the **Defragment** option to defragment rows in a table or the **Compress** option to defragment the indexes which have become sparse due to deletions.

Reorganizing may also reduce the total number of pages that are used to store the table and its indexes. It may reduce the number of levels in an index tree.

Note that the reorganization does not result in a reduction of the total size of the database file. The **Rebuild** option on the **General** tab completely rebuilds the database, eliminating any fragmentation, and free space. This option may result in a reduction of the total size of the database files.

See “Estimating catalog space requirements” on page 684.

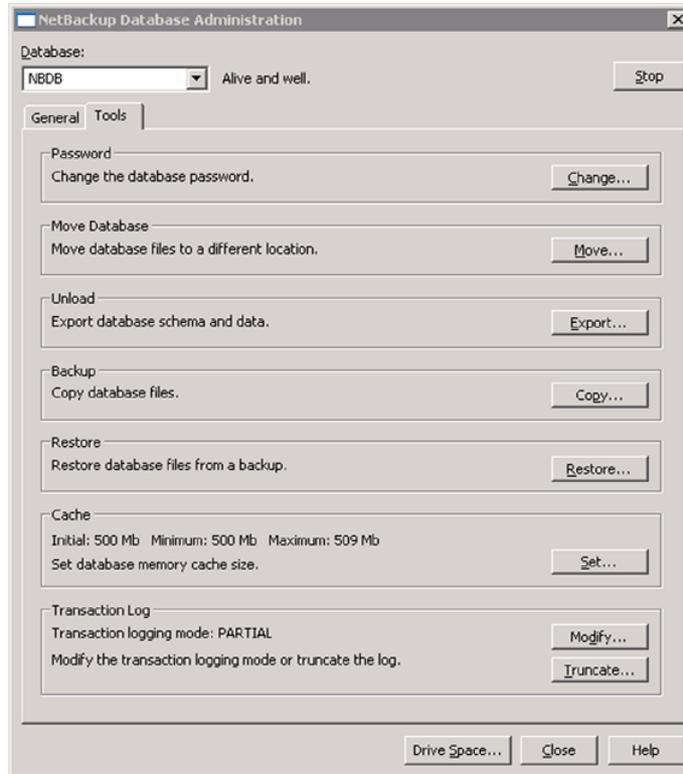
About the Tools tab of the NetBackup Database Administration utility

The **Tools** tab of the NetBackup Database Administration utility contains a variety of tools to administer the selected database:

Password section	See “Changing the DBA password” on page 712.
Move Database section	See “Moving the NetBackup database files” on page 713.
Unload section	See “Exporting database schema and data” on page 714.
Backup section	See “Copying or backing up the database files” on page 715.
Restore section	See “Restoring database files from a backup” on page 716.
Cache section	See “Changing NetBackup database cache memory settings” on page 717.
Transaction Log section	See “Setting the transaction mode for NBDB.log and BMRDB.log” on page 719. See “Truncating the transaction log” on page 720.

Figure 19-2 displays the **Tools** tab of the utility.

Figure 19-2 Tools tab of the Database Administration utility



Changing the DBA password

When you start the Database Administration utility, you must enter the DBA password. If the DBA password is the default password that is used when NetBackup is installed, you are encouraged to change the password. You are not required to change the password, however.

To change the DBA password

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe

- 2 Enter the database logon password and click **OK**.

See “Using the NetBackup Database Administration utility” on page 703.

- 3 Select the **Tools** tab.

- 4 In the **Password** section, click **Change**.
- 5 In the **Change password** dialog box, enter the new password and confirm the new password. Changing the password changes it for both NBDB and BMRDB, if a BMR database is present.
- 6 Enable **Create a backup file of your new DBA password** to keep track of the password. Then, browse to a directory to store the file that contains the new password.
- 7 Click **OK**.

The **Change password** dialog box warns you that it is important to remember the password. Symantec may not be able to recover information within the EMM database if the password is unavailable.
- 8 Restart the database for the password change to take effect.

Moving the NetBackup database files

Use the NetBackup Database Administration utility to change the location of the database files or to split the database files into multiple directories. Changing the location of the database files can improve performance when the database is very large.

Note: Due to performance issues, NetBackup supports database files only on locally attached drives.

The database files are moved for both NBDB and BMRDB, if present.

To move the NetBackup database files

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

```
InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe
```
- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Move Database** section, click **Move**.

- 5 In the **Move database files** dialog box, select one or more of the following options:

Move data to	Use to change the location of the data dbspaces. Browse to the new location.
Move index to	Use to change the location of the index dbspaces. Browse to the new location.
Move transaction log to	Use to change the location of the transaction log. The transaction logs, NBDB.log, and BMRDB.log, are critical files that are used to recover the relational databases. Browse to the new location.
Create mirrored transaction log at	Use to create a mirrored transaction log. Create the mirrored log in a different directory from the original log. Symantec also recommends the mirrored transaction log be placed on a different drive. A mirrored transaction log offers extra protection.
Move mirrored transaction log to	This option is displayed if a mirrored transaction log exists. Use to change the location of the mirrored transaction log. Browse to the new location. Create the mirrored log in a different directory from the original log.
Stop mirroring	This option is displayed if mirroring is used. Use this option to stop mirroring of the transaction log. This option removes any existing mirrored transaction log from the directory.

- 6 Click **OK**. The NetBackup operations are temporarily suspended.

Exporting database schema and data

Use the NetBackup Database Administration utility to unload either the schema or the schema and data from the relational database.

To export database schema and data

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe

- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.

- 4 In the **Unload** section, click **Export**.
- 5 In the **Export database** dialog box, browse to a destination directory.
- 6 Select one or more of the following options:

Schema	Unload only the database schema. The schema is unloaded as a file that is named <code>reload.sql</code> in the named directory.
Schema and data	Unload both the database schema and the data. The data is unloaded as a set of files in comma-delimited format. One file is created for each database table.

- 7 Click **OK**.

Copying or backing up the database files

Use the NetBackup Database Administration utility to back up the relational database to a specified directory.

Symantec recommends creating a backup copy of the database files in the following situations:

Before you move the database.	See “Moving the NetBackup database files” on page 713.
Before you rebuild the database.	See “About the General tab of the NetBackup Database Administration utility” on page 704.
Before you add data space.	See “About the General tab of the NetBackup Database Administration utility” on page 704.
Before you add index space.	See “About the General tab of the NetBackup Database Administration utility” on page 704.
Before you modify the transaction logging mode from FULL to PARTIAL.	See “Setting the transaction mode for NBDB.log and BMRDB.log” on page 719.
Before you truncate the transaction log.	See “Truncating the transaction log” on page 720.

Note: Using the NetBackup Database Administration utility to back up and 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 the tool to back up and restore the NetBackup catalog only as a precautionary measure.

To copy or back up the database files

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe

- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Backup** section, click **Copy**.
- 5 In the **Copy database files** dialog box, browse to a destination directory.
The destination directory contains the files that are created by the backup. A copy of all of the database files is made in this directory. This directory is also the location of the database files that are used by the **Restore** option.
See “Restoring database files from a backup” on page 716.
- 6 Select one of the following options:

Online	Makes a copy of the database files while the database is active. Other NetBackup activity is not suspended during this time.
Offline	Makes a copy of the database files with all other NetBackup activity suspended. The database is shut down before the copy is made, and restarted after the copy has completed. Since NetBackup activity is suspended, do not perform an offline backup while active backups or restores run.

Note: Neither option is a catalog backup, performed as part of regular NetBackup operations.

- 7 Click **OK**.

Restoring database files from a backup

Use the NetBackup Database Administration utility to restore a database from a backup copy. The backup copy may be either online or offline.

The restore overwrites the current database files. The database is shut down and restarted after the restore is completed.

A database restore causes NetBackup activity to be suspended, so do not perform a database restore while active backups or other restores run.

Note: Using the Database Administration utility to back up and 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 the tool to back up and restore the NetBackup database only as a precautionary measure.

To restore database files from a backup

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

```
InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe
```
- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Restore** section, click **Restore**.
- 5 In the **Restore database** dialog box, browse to the directory that contains the backup database.
See “Copying or backing up the database files” on page 715.
- 6 Click **OK**.

Changing NetBackup database cache memory settings

You can use the NetBackup Database Administration utility to view and change the SQL Anywhere memory cache settings of the relational database server.

Changes to these settings affect all of the relational databases that the database server manages. The changes do not take effect until the NetBackup services are restarted.

The database cache is an area of the memory that is used by the database server to store database pages for repeated fast access. The more pages that are accessible in the cache, the fewer times the database server needs to read data from disk.

Reading data from disk is a slow operation and the amount of cache available is often a key factor in determining performance. The database cache is automatically resized as needed.

The cache grows when the database server can usefully use more, as long as memory is available. The cache shrinks when other applications require cache

memory so that the database server does not unduly affect other applications on the system.

To change the NetBackup database cache memory settings

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

```
InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe
```

- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Cache** section, click **Set**.
- 5 In the **NetBackup Database Cache settings** dialog box, choose from the pre-set cache settings as described in Table 19-8.

The selection determines the following memory cache settings to control the size of the database cache:

Minimum cache size	<p>Sets the minimum cache size as a lower limit to automatic cache resizing.</p> <p>This setting represents the <code>-cl</code> option in the <code>server.conf</code> file.</p>
Initial cache size	<p>Sets the initial memory that is reserved for caching database pages and other server information.</p> <p>This setting represents the <code>-c</code> option in the <code>server.conf</code> file.</p>
Maximum cache size	<p>Sets the maximum cache size as an upper limit to automatic cache growth.</p> <p>This setting represents the <code>-ch</code> option in the <code>server.conf</code> file.</p> <p>If the settings are too large, the database server may not start.</p>

- 6 Click **OK**.

Table 19-8 Database cache settings

Option	Minimum cache size	Initial cache size	Maximum cache size
Current	As configured	As configured	As configured

Table 19-8 Database cache settings (continued)

Option	Minimum cache size	Initial cache size	Maximum cache size
Small	25MB (50MB with BMR)	25MB (50MB with BMR)	500MB (750MB with BMR)
Medium	200MB (400MB with BMR)	200MB (400MB with BMR)	750MB (850MB with BMR)
Large	500MB (750MB with BMR)	500MB (750MB with BMR)	1000MB (1000MB with BMR)
Custom	Configurable	Configurable	Configurable

The database cache settings can be configured in the NetBackup Database Administration utility or in the `server.conf` file. The database server reads the file when it is started.

The `server.conf` file is found in the following location:

`installpath\VERITAS\NetBackupDB\conf`

Setting the transaction mode for NBDB.log and BMRDB.log

You can use the NetBackup Database Administration utility to set the transaction log mode for `NBDB.log` and `BMRDB.log`. The transaction mode determines when the transaction log is automatically truncated outside of the catalog backup process.

To set the transaction log mode for NBDB.log and BMRDB.log

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

`InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe`

- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Transaction Log** section, click **Modify**.

5 In the **Transaction Log Mode** dialog box, select from one of the following log modes:

- | | |
|----------------|---|
| Full (default) | <p>With Full selected, the transaction log is truncated automatically after a successful online (hot) or offline (cold) catalog backup. All of the database files are included in the backup. (NBDB, NBDB.db, EMM_DATA.db, EMM_INDEX.db, NBDB.log.)</p> <p>The differential incremental schedule for the online (hot) catalog backup includes only the transaction log file.</p> <p>To recover using a full and an incremental backup, all of the database files are restored. The transaction logs are applied one at a time in order.</p> |
| Partial | <p>With Partial selected, all of the schedules used for the hot catalog backup policies include backups of all of the relational database files.</p> <p>Partial mode forces a deletion of the transaction log whenever a database checkpoint occurs.</p> <p>In Partial mode, the online (hot) catalog backup must always be a full backup. All incremental schedules are automatically converted to full schedules by NetBackup. The cold catalog backup is always a full backup.</p> |

6 Click **OK**. The new log settings go into effect after the NetBackup services are restarted.

Truncating the transaction log

You can use the NetBackup Database Administration utility to truncate the transaction log.

Truncating the transaction log forces a full, hot catalog backup the next time any schedule of the catalog backup policy is due. Without running a full, hot catalog backup, a gap would exist in the transaction logs due to the truncation. This results in an error during catalog recovery.

If the next scheduled hot catalog backup is a differential incremental, a backup of all of the relational database files is included.

To truncate the transaction log (NBDB.log and BMRDB.log)

- 1 Start the NetBackup Database Administration utility that is located in the following directory:

```
InstallPath\VERITAS\NetBackup\bin\NbDbAdmin.exe
```
- 2 Enter the database logon password and click **OK**.
See “Using the NetBackup Database Administration utility” on page 703.
- 3 Select the **Tools** tab.
- 4 In the **Transaction Log** section, click **Truncate**.
- 5 In the **Truncate Transaction Log** dialog box, browse to a **Temporary Directory For Truncation** where the log is copied. Make sure that enough space is available for a copy of the existing transaction log before it is truncated.
- 6 Click **OK**. After the transaction log is successfully copied and truncated, the temporary copy is deleted.

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 721.
- Move NBDB and BMRDB database files (possibly to tune performance).
See “Moving NBDB database files after installation” on page 722.
- Add a mirrored transaction log.
See “Adding a mirrored transaction log” on page 723.
- Recreate NBDB.
See “Creating the NBDB database manually” on page 724.

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 NetBackup Database Administration utility.
See “Using the NetBackup Database Administration utility” on page 703.
 - Run the following command to update the `vxdbs.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, you can change the location of the database files or split 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 NetBackup Database Administration utility.
See “Moving the NetBackup database files” on page 713.
 - 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.

See “About NetBackup master server installed directories and files” on page 695.

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, 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 NetBackup Database Administration utility.
See “Setting the transaction mode for NBDB.log and BMRDB.log” on page 719.
See “Truncating the transaction log” on page 720.

- Type the following command:

```
Install_path\NetBackup\bin\nbdb_move.exe  
-mlog log_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  
-data data_directory -index index_directory -tlog  
log_directory -mlog log_mirror_directory
```

- 5 Start all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

- 6 Perform a catalog backup.

See “About online, hot catalog backups” on page 668.

See “Moving NBDB database files after installation” on page 722.

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: Recreating 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 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.

See “Sybase SQL Anywhere server management” on page 702.

See “About the NetBackup relational database (NBDB) installation” on page 693.

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 following actions. In each command, `NB_server_name` matches the name in `server.conf`.

See “About the NetBackup `server.conf` file” on page 696.

- 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.

See “Relocating the NetBackup database” on page 695.

See “Moving the NetBackup database from one host to another” on page 731.

See “Moving NBDB database files after installation” on page 722.

About backup and recovery procedures

The online, hot catalog method can be performed while regular backup activity takes place.

It runs according to 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, NBAZDB, 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.

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\NetBackupDE\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.

See “Adding a mirrored transaction log” on page 723.

See “About NetBackup master server installed directories and files” on page 695.

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*.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

See “Commands for backing up and recovering the relational databases” on page 727.

Commands for 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 19-9 Commands used to back up and recover relational databases

Command	Description
nbdb_backup.exe	<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\NetBackup\bin\nbdb_backup.exe [-dbn database_name] [-online -offline] destination_directory</i></pre> <p><code>-dbn database_name</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>Note: Using this command (or the NetBackup Database Administration utility) 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 NetBackup Database Administration utility) to back up the NetBackup catalog only as a precautionary measure.</p> <p>Note: The transaction logs are not truncated by using <code>nbdb_backup</code>. A catalog backup must be run to truncate the logs.</p>
nbdb_restore.exe	<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\NetBackup\bin\nbdb_restore.exe -recover source_directory</i></pre> <p>Logs are recorded in the <code>\admin</code> directory.</p> <p>Note: Using this command (or the NetBackup Database Administration utility) 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 NetBackup Database Administration utility) to restore the NetBackup catalog only as a precautionary measure.</p>

See “About the Enterprise Media Manager (EMM) database” on page 666.

See “Configuring a catalog backup manually” on page 673.

See “Strategies that ensure successful NetBackup catalog backups” on page 677.

About the online, hot catalog backup process

Normally, a hot, online catalog backup consists of one parent job and two or more child jobs. Events for these jobs appear in the `dbm` log.

An overview of the hot catalog backup process consists of the following process:

- Make a temporary copy of database files to a staging directory by typing the following command:

```
Install_path\NetBackupDB\staging
```

Once the copy is made, NetBackup can back up the catalog files.

- A child job backs up files in a single stream as follows:
 - Configuration files (*server.conf, database.conf, vxdbms.conf*)

- Database files

```
NBDB.db
NBDB.log
NBAZDB.db
NBAZDB.db.template
NBAZDB.log
EMM_DATA.db
EMM_INDEX.db
```

If BMR was installed

```
BMRDB.db
BMRDB.log
BMR_DATA.db
BMR_INDEX.db
```

- A second child job begins the image catalog backup.

If BMR is installed and a remote EMM server is in use, the backup of the EMM server appears as a separate job.

- Transaction logs are truncated after a successful full or incremental backup.

If the transaction logs are manually changed or deleted, a hole could exist in the recovery.

The child job for the relational database backup is normally run on the master server. The master server is the default location for NBDB and the required location for BMRDB.

If NBDB was moved to a media server, the child job runs on the media server. In this case, additional logging for the job appears in the admin log on the media server.

If NBDB was moved to a media server and BMRDB is installed on the master server, two child jobs exist for the relational database backup portion of the online, hot catalog backup. One on the media server for NBDB and one on the master server for BMRDB.

Unloading the NetBackup database

Use the NetBackup Database Administration utility or the `nbdb_unload` command line utility to dump the entire NetBackup or Bare Metal Restore databases. These utilities can also be used to dump 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.

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.

See “Exporting database schema and data” on page 714.

See “Terminating database connections” on page 730.

Terminating database connections

Before you run `nbdb_unload`, shut down NetBackup to terminate all active connections to the database. Shutting down NetBackup eliminates any possible concurrency problems.

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 NetBackup Database Administration utility.
See “Using the NetBackup Database Administration utility” on page 703.
- 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 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 NetBackup Database Administration utility or `nbdb_backup` to make a physical copy or use `nbdb_move` to relocate the database files.

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, and NetBackup Authorization Database, NBAZDB, must also reside on the master server. So, if NBDB and EMM server are moved to a media server from a master server, BMRDB and NBAZDB 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 735.

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
B_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 NetBackup Database Administration utility 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`

`NBDB.log`

`NBDB.m.log` (optional)

`EMM_DATA.db`

`EMM_INDEX.db`

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/bpdown -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`

`NBDB.log`

`NBDB.m.log` (optional)

`EMM_DATA.db`

`EMM_INDEX.db`

- 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  
B_emmservername
```

This command is not internationalized.

- 12** On A, delete the following database files and configuration files:

```
NBDB.db  
  
NBDB.log  
  
NBDB.m.log (optional)  
  
NBAZAD.db  
  
NBAZAD.db.template  
  
NBAZAD.log  
  
EMM_DATA.db  
  
EMM_INDEX.db
```

- 13** On A do the following:

- If BMRDB and NBAZDB do not exist on A, delete the following configuration files:

```
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.

See “Commands for backing up and recovering the relational databases” on page 727.

See “Changing the database password” on page 721.

See “About NetBackup master server installed directories and files” on page 695.

See “Removing the EMM server from a Windows cluster” on page 736.

See “Terminating database connections” on page 730.

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 735.

See “Removing the EMM server from a Windows cluster” on page 736.

See “Sybase SQL Anywhere and clustered environments” on page 702.

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).
- See “Removing the EMM server from a Windows cluster” on page 736.
- See “Moving the NetBackup database from one host to another” on page 731.
- See “Cluster considerations with the EMM server” on page 735.

Removing 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.
- 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).
See “Moving the EMM server to a Windows cluster” on page 735.
See “Moving the NetBackup database from one host to another” on page 731.
See “Cluster considerations with the EMM server” on page 735.

Managing backup images

This chapter includes the following topics:

- About the Catalog utility
- About searching for backup images
- Verifying backup images
- Viewing job results
- Promoting a copy to a primary copy
- Duplicating backup images
- Expiring backup images
- About importing backup images

About the Catalog utility

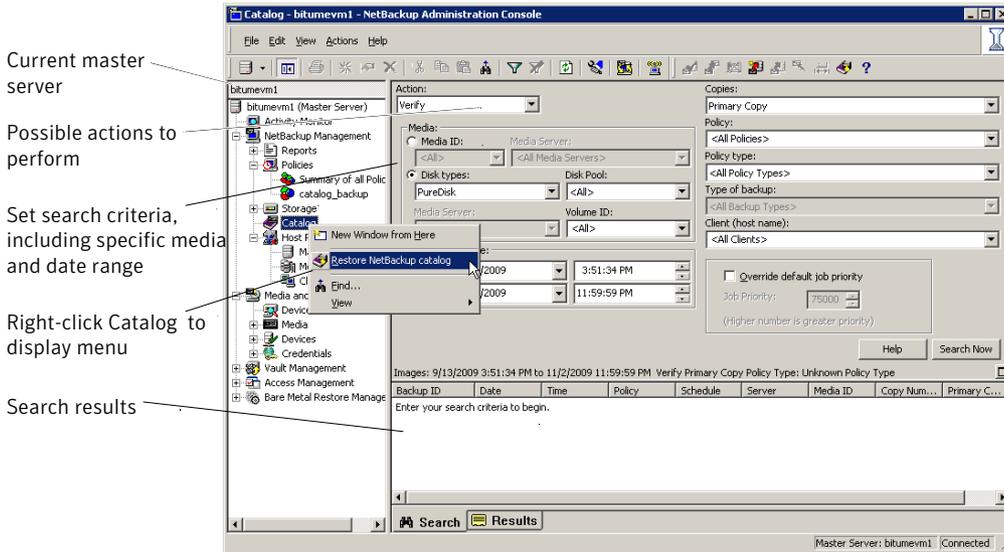
Use the **Catalog** utility in the **NetBackup Administration Console** to create and configure catalog backups. Catalog backups are required for NetBackup to protect NetBackup internal databases. The catalogs contain setup information as well as critical information about client backups. The catalog 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 20-1 Catalog utility options



About 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 “About importing backup images” on page 750.

Table 20-1 lists the search criteria for backup images.

Table 20-1 Catalog utility search properties

Property	Description
Action	Specifies the action that was used to create the image: Verify, Duplicate, Import . See “Verifying backup images” on page 740. See “Duplicating backup images” on page 743. See “Expiring backup images” on page 750.
Media ID	Specifies 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 <All> .
Media Server	Specifies 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 All Media Servers .
Disk type	Specifies the type of the disk storage unit on which to search for backup images.
Disk pool	Specifies the name of the disk pool on which to search for backup images.
Volume ID	Specifies the ID of the disk volume in the disk pool on which to search for backup images.
NearStore Server	Specifies 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 All NearStore Servers .
Path	Searches for an image on a disk storage unit, if the path is entered. Or, searches all of the disk storage on the specified server, if All was selected. Appears if the disk type is BasicDisk or NearStore.
Date/time range	Specifies the range of dates and times that includes all the backups for which you want to search. The Global Attributes property Policy Update Interval determines the default range.
Copies	Specifies the source you want to search. From the scroll-down list, select either Primary or the copy number.
Policy	Specifies 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 All Policies .

Table 20-1 Catalog utility search properties (*continued*)

Property	Description
Client (host name)	Specifies 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 All Clients .
Type of backup	Specifies 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 All Backup Types .
Override default job priority	<p>Selects the job priority for verify, duplicate, and import actions.</p> <p>To change the default for the selected action, enable Override default job priority. Then, select a value in the Job Priority field.</p> <p>Changes in the catalog dialog box affect the priority for the selected job only.</p> <p>If this option is not enabled, the job runs using the default priority as specified in the Default Job Priorities host properties.</p> <p>See “Default Job Priorities properties” on page 105.</p>

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**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, on the **Search** tab, set up the search criteria for the image you want to verify. Click **Search Now**.
- 3 In the right pane, click the **Results** tab, then select the verification job to view the job results.

See “Viewing job results” on page 741.

Viewing job results

The results of verify, duplicate, or import jobs appear in the **Results** tab for the Catalog options. 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**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, click the **Results** tab.
- 3 Select a log file.
- 4 On the **View** menu, click **View > Full View** to display the entire log file in a screen editor.

On the **Edit** menu, 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.

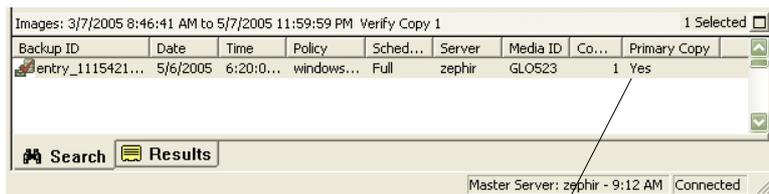
There are three different ways to promote a copy to a primary copy:

Promote a backup copy to a primary copy using search criteria	See “To promote a backup copy to a primary copy” on page 742.
Promote a copy to a primary copy for many backups using the <code>bpchangeprimary</code> command	See “To promote a copy to a primary copy for many backups” on page 742.

Promote a backup copy to a primary copy using the `bpduplicate` command

See “To use `bpduplicate` to promote a backup copy to a primary copy” on page 743.

Figure 20-2 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**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, 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 “About searching for backup images” on page 738.

- 3 Select the image you want to promote.
- 4 On the **Actions** menu, 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
```

In the next example, 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 the *NetBackup Commands Reference Guide*.

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 20-2 lists the scenarios in which duplication is possible and scenarios in which duplication is not possible:

Table 20-2 Backup duplication scenarios

Duplication possible	Duplication not possible
<ul style="list-style-type: none"> ■ From one storage unit to another. ■ From one media density to another. ■ From one server to another. ■ From multiplex to nonmultiplex format. ■ 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.) 	<ul style="list-style-type: none"> ■ While the backup is created (unless making multiple copies concurrently). ■ When the backup has expired. ■ By using NetBackup to schedule duplications automatically (unless you use a Vault policy to schedule duplication) ■ When it is a multiplexed duplicate of the following type: <ul style="list-style-type: none"> ■ FlashBackup ■ NDMP backup ■ Backups from disk type storage units ■ Backups to disk type storage units ■ Nonmultiplexed backups

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.

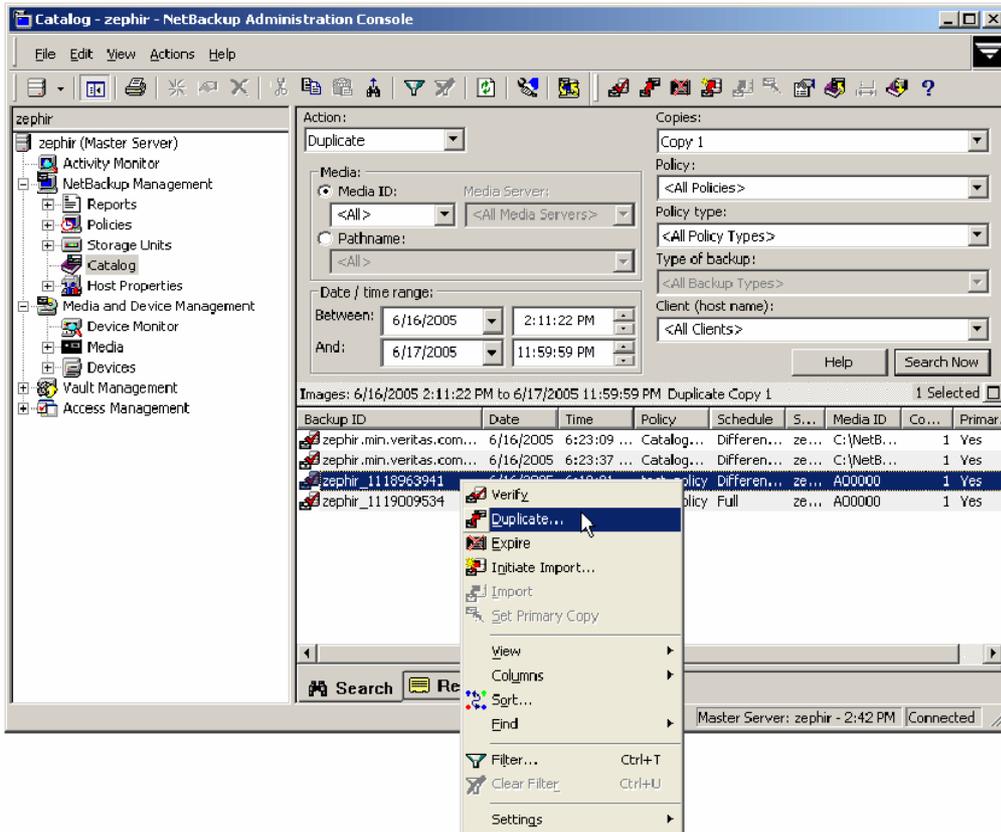
See “About writing multiple copies using a storage lifecycle policy” on page 467.

To duplicate backup images

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, 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.

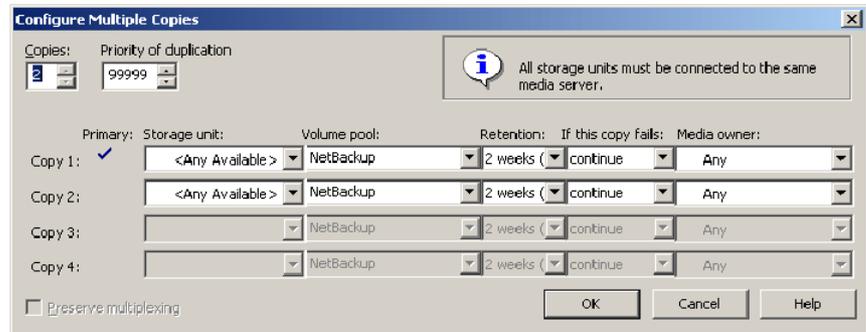
If you duplicate an online, hot catalog backup, select all child jobs that were used to create the catalog backup. All jobs must be duplicated to duplicate the catalog backup.



- 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 “Global Attributes properties” on page 131.



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 “About configuring for multiple copies” on page 563.

7 Specify the volume pool where each copy is stored.

The following volume pool selections are based on the policy type setting that was used for the query.

Policy type set to All Policy Types (default)	Specifies that all volume pools are included in the drop-down list. Both catalog and non-catalog volume pools are included.
Policy type set to NBU-Catalog	Specifies that only catalog volume pools are included in the drop-down list.
Policy type set to a policy type other than NBU-Catalog or All Policy Types	Specifies that 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 `bexpdate` 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, 2010 and its retention period is one week, the new copy's expiration date is November 21, 2010.

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	Specifies that NetBackup chooses the media owner, either a media server or server group.
None	Specifies 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 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 741.

See “About multiplexed duplication considerations” on page 748.

About multiplexed duplication considerations

Consider the following items about multiplexed duplication.

Table 20-3 Multiplexed duplication considerations

Consideration	Description
Multiplex settings are ignored	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.
Backups in a multiplexed group are duplicated and duplicated group is identical	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: <ul style="list-style-type: none"> ■ 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.

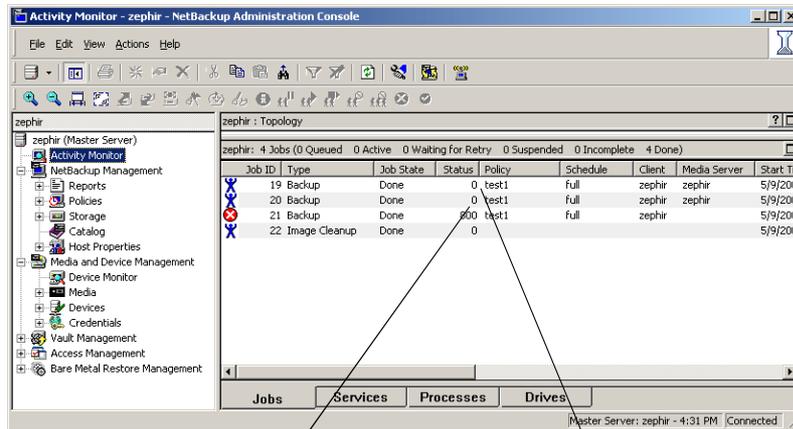
Jobs that appear while making multiple copies

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.

Figure 20-3 Backup that contains two copies



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**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, set up the search criteria for the image you want to expire, then click **Search Now**.
See “About searching for backup images” on page 738.
- 3 Select the image you want to expire and on the **Actions** menu, 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.

About importing backup images

NetBackup can import the backups that have expired, the backups from another NetBackup server, or the backups written by Backup Exec for Windows.

See “About importing Backup Exec media” on page 755.

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:

Table 20-4 Phases to import an image

Phase	Description
Phase I	NetBackup creates a list of expired catalog entries for the backups on the imported volume. No actual import occurs in Phase I. See “Importing backup images, Phase I” on page 751.
Phase II	Images are selected for importing from the list of expired images that was created in Phase I. See “Importing backup images, Phase II” on page 753.

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 `vmphyinv` physical inventory utility” on page 361.

- 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 In the **Media Host** field, type the name of the host that contains the volume to import. Click **Next**.

This media server becomes the media owner.

- 6 In the **Image Type** field, 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 741.

- 9 Complete the import.

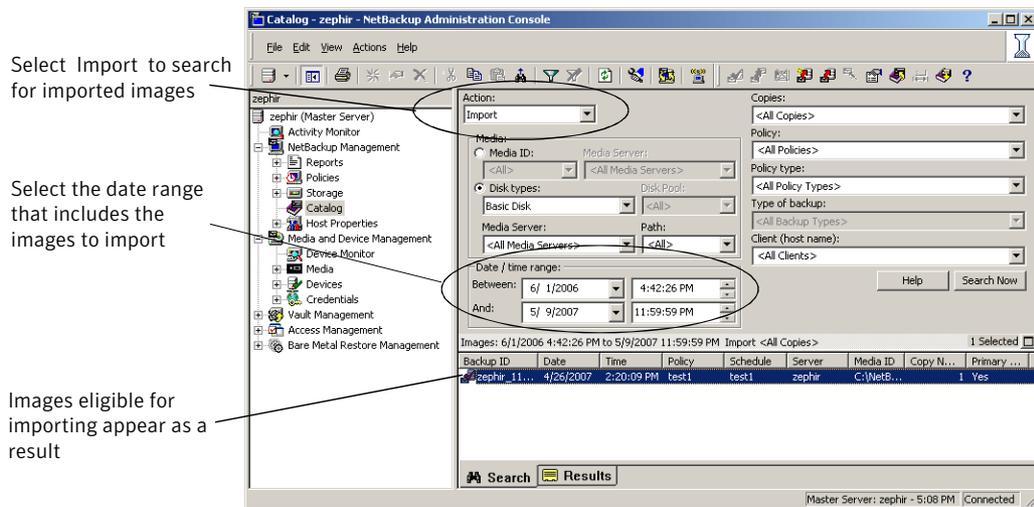
See “Importing backup images, Phase II” on page 753.

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**, in the left pane, expand **NetBackup Management > Catalog**.
- 2 In the right pane, 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.



- 3 Select the image(s) you want to import and on the **Actions** menu, select **Actions > Import**.
- 4 To view the log, click the **Results** tab, then select the import job log.

About 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, 2010, and its retention period is one week, the new expiration date is November 21, 2010.

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\bpexpdate  
-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 750.

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 `vmphyinv` physical inventory utility” on page 361.

- 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**, in the left pane, expand **NetBackup Management > Catalog**.

- 4 On the **Actions** menu, 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 **Initialize Import** window, 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.

The following table shows the actions to take depending on the location of the image.

If images are on tape	<p>In the Media ID field, type the Media ID of the volume that contains the backups to import.</p> <p>Check whether or not the images to import are password-protected Backup Exec images.</p> <p>Validate the Backup Exec password by retyping the password in the field provided.</p>
If images are on disk	<p>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.</p> <p>If the disk type references a disk pool, enter or select the disk pool and the disk volume ID.</p> <p>If the disk type references a disk pool, enter or select the disk pool and the disk volume ID.</p> <p>For a BasicDisk type, enter or browse to the path to the images in the field provided.</p> <p>For a NearStore disk type, select or enter the name of the NearStore server and the NearStore volume.</p>

Then, 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.

About importing Backup Exec media

Consider the following situations and results when importing Backup Exec media:

Table 20-5 Importing Backup Exec media results

Situation	Result
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.
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 <code>bpimport</code> command.
Importing from Backup Exec media and conversion/migration of job information	Does not convert or migrate Backup Exec job history, job schedules, or job descriptions to NetBackup.
Importing from Backup Exec media and conversion of application setup or configuration information	Does not convert Backup Exec application setup or configuration information to NetBackup.
Backup Exec backups created with the Intelligent Image Option	Cannot be restored.
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 . ϵ 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 68.

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

The Backup Exec Tape Reader provides support for the following versions of Windows images, Exchange Server images, and SQL images.

Table 20-6 Backup Exec Tape Reader supported images and versions

Image	Versions supported
Windows images	<p>The Backup Exec Tape Reader provides support for all Windows versions that NetBackup currently supports.</p> <p>The support includes the following:</p> <ul style="list-style-type: none"> ■ 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.

Table 20-6 Backup Exec Tape Reader supported images and versions (*continued*)

Image	Versions supported
Exchange Server images	<p>The Backup Exec Tape Reader provides support for the following:</p> <ul style="list-style-type: none"> ■ 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. <p>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.</p> <p>The following functionality is not available for Backup Exec images of Exchange 2003 and 2007:</p> <ul style="list-style-type: none"> ■ 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.
SQL images	<p>The Backup Exec Tape Reader provides support for the following:</p> <ul style="list-style-type: none"> ■ 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 table describes the differences between Backup Exec and NetBackup when importing, browsing, and restoring images.

Table 20-7 Differences between Backup Exec and NetBackup when importing, browsing, and restoring images

Topic	Differences
Run <code>vmphyinv</code> for Backup Exec media	<p>To import Backup Exec media requires <code>vmphyinv</code> 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.</p> <p>See “About the <code>vmphyinv</code> physical inventory utility” on page 361.</p>
To import and restore QIC media	<p>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.</p>

Table 20-7 Differences between Backup Exec and NetBackup when importing, browsing, and restoring images (*continued*)

Topic	Differences
Spanned media: Importing differences	<p>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.</p> <p>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.</p>
SQL: Browsing and restoring differences	<p>Backup Exec SQL images are browsed, then restored using the NetBackup Backup, Archive, and Restore client interface.</p> <p>NetBackup SQL images are browsed, then restored using the NetBackup SQL interface.</p>
File level objects: Browsing and restoring differences	<p>When a user selects a Backup Exec file to restore, the directory where that file is located is restored.</p> <p>When a user selects a NetBackup file to restore, only the single file is restored.</p>
NetWare: Restoring differences	<p>NetBackup does not support restoring Backup Exec NetWare non-SMS backups that were created using the NetWare redirector.</p> <p>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.</p>

Table 20-7 Differences between Backup Exec and NetBackup when importing, browsing, and restoring images (*continued*)

Topic	Differences
Restoring NTFS hard links, NTFS SIS files, and Exchange SIS mail messages	<ul style="list-style-type: none"> ■ 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: <pre style="margin-left: 20px;">C:\hard_links\one.txt C:\hard_links\two.txt C:\hard_links\three.txt</pre> 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: <pre style="margin-left: 20px;">C:\hard_links\one.txt to a location on C:\ C:\hard_links\two.txt to a location on D:\</pre> 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 <pre style="margin-left: 20px;">C:\redir_hard_links.</pre>

Monitoring and reporting

- Chapter 21. Monitoring NetBackup activity
- Chapter 22. Auditing NetBackup operations
- Chapter 23. Reporting in NetBackup

Monitoring NetBackup activity

This chapter includes the following topics:

- About 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

About the Activity Monitor

Use the Activity Monitor in the **NetBackup Administration Console** to monitor and control NetBackup jobs, services, processes, and drives.

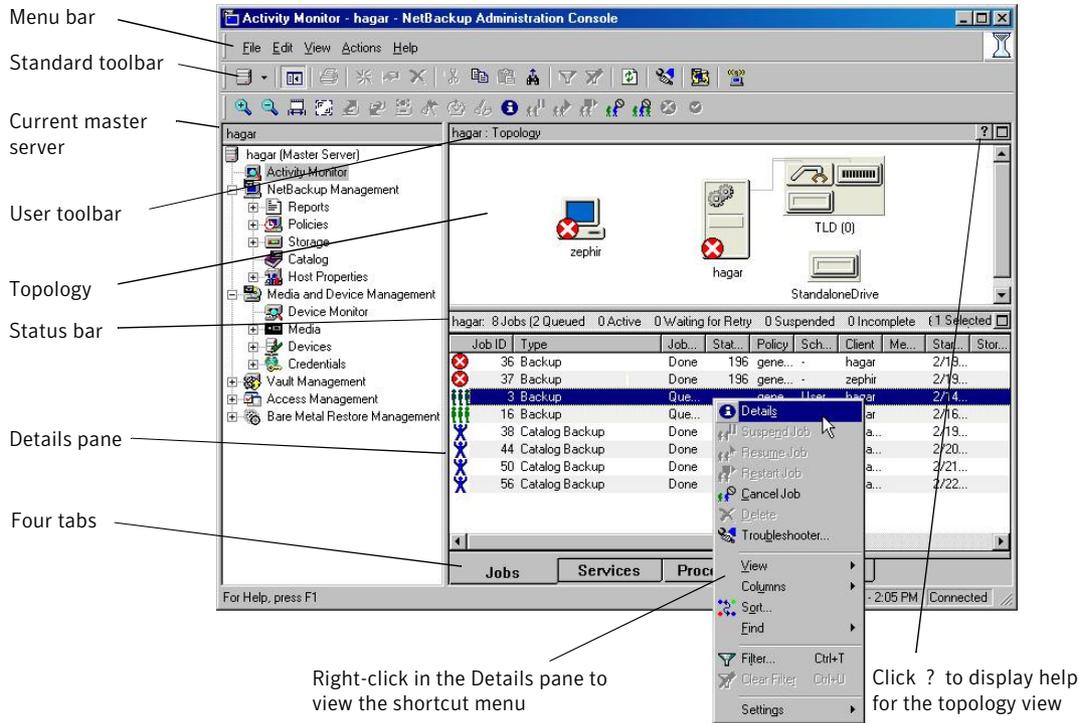
Note: The **Filter** option on the **View** menu is useful for displaying in Activity Monitor only those jobs with specified characteristics. For example, the jobs that were started before a specific date; jobs in the queued state; jobs with status completion codes within a specified range.

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 21-1 Activity Monitor



Activity Monitor topology

The Activity Monitor topology view displays icons that represent the state and configuration of the entire NetBackup system that is 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 topology view displays the following physical and logical connections between the devices:

- Relationships between master servers and media servers.
- Relationships between media servers and robots and drives.
- Back up and restore activity, job failures, down services, and drive state status.

- Robots and drives are represented as connected to a media server. Drives that are physically located in a robotic library appear directly beneath the robotic library. Stand-alone drives are represented as individual drive objects.
- Drive-to-device host connections are shown only if the drive is used for a request and the drive is in use by NetBackup. Robot-to-server connections and robotic library-to-volume database connections are always shown.
- 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.

Figure 21-2 shows some of the icons you may see in the Activity Monitor.

Figure 21-2 Example of Activity Monitor icons



About 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

In the **NetBackup Administration Console**, the **Jobs** tab in the Activity Monitor displays all of the 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.

Table 21-1 Tasks initiated by parent jobs

Task	Description
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. (Snapshots and copy snapshots to a storage unit is selected in the policy Schedule Attributes 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. (Snapshots only is selected in the policy Schedule Attributes 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"> ■ A Sybase backup ■ A file system backup of the master server ■ A backup of the BMR database, if necessary
Multiple copies	<p>See “Multiple copies (schedule attribute)” on page 562.</p> <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>

Table 21-1 Tasks initiated by parent jobs (*continued*)

Task	Description
Multiple data streams	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.
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

- ◆ In the **NetBackup Administration Console**, click **Activity Monitor**. To view the details for a specific job, double-click on the job displayed in the **Jobs** tab pane. 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.

Showing or hiding column heads

The following procedure describes how to show or hide column heads.

To show or hide column heads

- 1 In the **NetBackup Administration Console**, 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 a selected job

- 1 In the **NetBackup Administration Console**, 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**.

Deleting completed jobs

The following procedure describes how to delete a completed job.

To delete completed jobs

- 1 In the **NetBackup Administration Console**, 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 In the **NetBackup Administration Console**, 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 completed job

The following procedure describes how to restart a completed job.

To restart a completed job

- 1 In the **NetBackup Administration Console**, 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 In the **NetBackup Administration Console**, 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 or incomplete jobs

The following procedure describes how to resume suspended or incomplete jobs.

To resume a suspended or an incomplete job

- 1 In the **NetBackup Administration Console**, 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 from a list of jobs.

To print job list information from a list of jobs

- 1 In the **NetBackup Administration Console**, 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 from a single job.

To print job detail information from a single job

- 1 In the **NetBackup Administration Console**, 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
Attempt 1
Started: 6/7/2007 6:55:10 PM
Elapsed: 00:02:49
Ended: 6/7/2007 6:57:59 PM
Status: the requested operation was successfully completed(0)
-----
File list:
C:\Documents and Settings

```

Copying Activity Monitor text to a file

The following procedure describes how to copy Activity Monitor text to a file.

To copy Activity Monitor text to a file

- 1 In the **NetBackup Administration Console**, 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 job priority.

Table 21-2 Change Job Priority options

Option	Description
Set Job Priority to	Enters the specific job priority for the selected jobs.
Increment the Job Priority by	Raises the priority of the job by the selected internal.
Decrement the Job Priority by	Lowers the priority of the job by the selected internal.

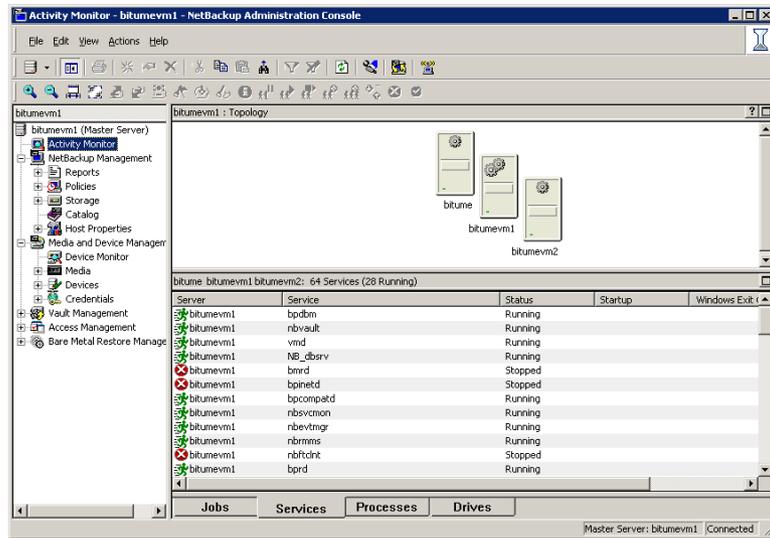
Changes in the **Change job priority** dialog box affect the priority for the selected job only, and not all other jobs of that type.

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 in the Activity Monitor in the **NetBackup Administration Console** displays the status of NetBackup services on the master server and all media servers that the selected master server uses.

Figure 21-3 Services tab in the Activity Monitor



Note: To see any services or processes on another computer, the other computer 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 21-3 NetBackup services

Service	Description
NetBackup Authentication (nbatd.exe)	NetBackup Product Authentication validates identities and forms the basis for authorization and access control in NetBackup.
NetBackup Authorization (nbazd.exe)	NetBackup Product Authorization provides access control in NetBackup applications.
NetBackup Bare Metal Restore Master Server (bmrtd.exe)	Appears if Bare Metal Restore is installed.

Table 21-3 NetBackup services (continued)

Service	Description
<p>NetBackup Client Service (bpinetd.exe)</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>Note: 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>To configure a BasicDisk storage unit that uses CIFS, nbrmms must share the same logon credentials as bpinetd on the media server.</p> <p>See “Configuring credentials for CIFS and disk storage units” on page 395.</p>
<p>NetBackup Compatibility Service (bpcompatd.exe)</p>	<p>Communicates with legacy NetBackup services.</p>
<p>NetBackup Database Manager (bpdbm.exe)</p>	<p>Manages the NetBackup internal databases and catalogs. BPDBM must be running on the NetBackup master server during all normal NetBackup operations.</p>
<p>NetBackup Deduplication Engine (spoold.exe)</p>	<p>Runs on the NetBackup deduplication storage server host to store and manage deduplicated client data. The file name spoold.exe 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 storage server.</p>
<p>NetBackup Deduplication Manager (spad.exe)</p>	<p>Runs on the NetBackup deduplication storage server host to maintain the NetBackup deduplication configuration, control deduplication internal processes, control replication, control security, and control event escalation.</p> <p>Active only if the NetBackup Deduplication Option is licensed and the media server is configured as a deduplication storage server.</p>
<p>NetBackup Device Manager (ltid.exe)</p>	<p>Starts the Volume Manager (vmd), the automatic volume recognition process (avrd), 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.</p>

Table 21-3 NetBackup services (continued)

Service	Description
NetBackup Enterprise Media Manager (nbemm.exe)	<p>Accesses and manages the database where media and device configuration information is stored (EMM_DATA.db). nbemm.exe must be running in order for jobs to run.</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>
NetBackup Event Manager Service (nbevtmgr.exe)	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 (nbjm.exe)	Accepts the jobs that the Policy Execution Manager (nbpem.exe) submits and acquires the necessary resources. The Job Manager then starts the job and informs nbpem.exe that the job is completed.
NetBackup Key Management Service (nbkms.exe)	A master server-based symmetric Key Management Service that provides encryption keys to media server BPTM processes.
NetBackup Policy Execution Manager (nbpem.exe)	Creates Policy or Client tasks and determinate when jobs are due to run. If a policy is modified or if an image expires, nbpem is notified and the Policy/Client task objects are updated.
NetBackup Relational Database Manager (dbsrv11.exe)	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 SQLANYs_VERITAS_NB.
NetBackup Remote Manager and Monitor Service (nbrmms.exe)	<p>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.</p> <p>To configure a BasicDisk storage unit that uses CIFS, nbrmms must share the same logon credentials as bpinetd on the media server.</p> <p>See “Configuring credentials for CIFS and disk storage units” on page 395.</p>
NetBackup Request Daemon (bprd.exe)	Processes the requests from NetBackup clients and servers. bprd also prompts NetBackup to perform automatically scheduled backups. bprd must be running on the NetBackup master server to perform any backups or restores.
NetBackup Resource Broker (nbrb.exe)	<p>Allocates the storage units, tape drives, and client reservations for jobs. nbrb works with the Enterprise Media Manager (NBEMM).</p> <p>The nbrbutil utility can be used to add or change the Resource Broker settings.</p> <p>See “Using the nbrbutil utility to configure the NetBackup Resource Broker” on page 777.</p>

Table 21-3 NetBackup services (continued)

Service	Description
NetBackup Service Layer (nbsl.exe)	<p>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.</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>
NetBackup Service Monitor (nbsvcmon.exe)	<p>Monitors the NetBackup services that run on the local computer. If a service unexpectedly terminates, the service tries to restart the terminated service. If <code>nbsvcmon</code> determines that NetBackup is configured for a cluster, the service shuts down, and the monitoring is taken over by the cluster.</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>
NetBackup Storage Lifecycle Manager (nbstserv.exe)	<p>Manages storage lifecycle operations and schedules duplication jobs. Monitors disk capacity on capacity-managed volumes and removes older images when required.</p> <p>The Duplication Manager and the Import Manager run within <code>nbstserv</code>:</p> <ul style="list-style-type: none"> ■ The Duplication Manager creates batches of the images to be imported based on SLP name and storage device (disk media id or robot number). ■ The Import Manager monitors a worklist in EMM for images to be imported and initiates <code>bpimport</code> jobs for those images.
Symantec Private Branch Exchange (pbx_exchange.exe)	<p>Note: This service does not appear in the Activity Monitor but is represented in the Windows Services utility.</p> <p>Provides single-port access to clients outside the firewall that connect to Symantec product services. Service name: <code>VRTSpbx</code>.</p>
NetBackup Vault Manager (nbvault.exe)	<p>Manages NetBackup Vault. <code>NEVAULT</code> must be running on the NetBackup Vault server during all NetBackup Vault operations.</p>
NetBackup Volume Manager (vmd.exe)	<p>Manages the volumes (tapes) needed for backup or restore and starts local device management daemons and processes.</p>

Types of services

The following table describes additional information about NetBackup services.

Table 21-4 Additional information about NetBackup services

NetBackup service	Description
Stand-alone services	Always run and listen to accept connections. Examples include <code>bpdbm</code> , <code>bprd</code> , <code>bpjobd</code> , and, <code>vmd</code> .
Multiprocess stand-alone services	"Fork" a child process to handle requests. Examples include <code>bpdbm</code> and <code>bprd</code> .
Single-process stand-alone services	Accept connections and handle requests in the same process.
<code>inetd</code> services	<code>inetd</code> (<code>lm</code>) or <code>bpinetd</code> usually launch these NetBackup services. Examples include <code>bpcd</code> , <code>bpjava-msvc</code> , and <code>vnetd</code> .

Using the `nbrbutil` utility to configure the NetBackup Resource Broker

The NetBackup Resource Broker (`nbrb`) allocates resources and maintains resource requests for jobs in the job queue. Use the `nbrbutil` utility to configure the Resource Broker.

The `nbrbutil` utility is located in the following directory:

- On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbrbutil
```

- On Windows:

```
Install_path\VERITAS\NetBackup\bin\admincmd\nbrbutil
```

For a complete description of `nbrbutil`, see the *NetBackup Commands Reference Guide*.

Table 21-5 describes the options available to `nbrbutil` command.

Table 21-5 `nbrbutil` options

Option	Description
<code>-cancel requestID</code>	Cancels the allocation request within the given identifier.
<code>-changePriority requestID</code>	Changes the request priority.
<code>-changePriorityClass requestID</code> <code>-priorityClass priorityClass</code>	Changes the request priority class.
<code>-changeSettings</code> <code>parameterparameter_value</code>	Adds or changes the <code>nbrb</code> configuration settings. Table 21-6 describes the configuration settings in detail.

Table 21-5 nbrbutil options (*continued*)

Option	Description
-deleteSetting <i>settingname</i>	Deletes a Resource Broker configuration setting identified by <i>settingname</i> .
-dump	Dumps all Resource Broker allocation and request lists.
-dumptables [-f <i>filename</i>]	Enables the Resource Broker to log its internal state in the specified file name.
-disablePerfMon	Disables performance monitoring.
-enablePerfMon	Enables performance monitoring.
-help	Lists the help for this command.
-listActiveDriveJobs [<i>driveName</i>]	Lists all the active jobs for a drive.
-listActiveJobs	Lists all the active jobs.
-listActiveMediaJobs <i>mediaId</i>	Lists all the active jobs for a media ID (disk or tape).
-listActivePoolJobs <i>poolName</i>	Lists all the active jobs for a volume pool.
-listActiveStuJobs <i>stuName</i> <i>stugroup</i>	Lists all the active jobs for a storage unit or a storage unit group.
-listOrphanedDrives	Lists the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
-listOrphanedMedia	Lists the media that is reserved in EMM but has no corresponding allocation in the Resource Broker.
-listOrphanedPipes	Lists the orphaned fibre transport pipes.
-listOrphanedStus	Lists the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
-listSettings	Lists the configuration settings of the Resource Broker.
-priority <i>priority</i>	Changes the request priority.
-release <i>allocationID</i>	Release the allocation with the given identifier.
-releaseAllocHolds	Releases the allocation holds caused by allocation errors for drives and media.
-releaseDrive <i>drivename</i>	Releases all allocations for the named drive.

Table 21-5 nrbutil options (continued)

Option	Description
<code>-releaseMDS mdsAllocationKey</code>	Releases the EMM and the MDS allocations that are allocated by the MDS with the specified identifier.
<code>-releaseMedia mediaid</code>	Releases all allocations for the specified volume.
<code>-releaseOrphanedDrive drivekey</code>	Releases the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-releaseOrphanedMedia mediakey</code>	Releases the media that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-releaseOrphanedPipes</code>	Releases the orphaned fibre transport pipes.
<code>-releaseOrphanedStu stuName</code>	Releases the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-reportInconsistentAllocations</code>	Reports inconsistent allocations between the Resource Broker and MDS.
<code>-resetAll</code>	Resets all Resource Broker allocations, requests, and persisted states.
<code>-resetMediaServer mediaserver</code>	Resets all Resource Broker EMM and MDS allocations that are related to <code>toltid</code> on the media server.
<code>-resume</code>	Resumes the Resource Broker processing.
<code>-setDriveGroupUnjoinable</code>	Disables the future job from joining the group for this drive.
<code>-setMediaGroupUnjoinable</code>	Disables the future job from joining the group for this media.
<code>-suspend</code>	Suspends the Resource Broker processing.
<code>-syncAllocations</code>	Syncs up any allocation difference between the Resource Broker and MDS.

Table 21-6 lists the parameters for the `nrbutil -changesettings` option, and describes the use of each.

Use the `nrbutil` command with the `-changesettings` option to add or change Resource Broker configuration settings.

Table 21-6 nbrbutil -changesettings parameters

Parameter	Description
RB_DO_INTERMITTENT_UNLOADS	<p>When the RB_DO_INTERMITTENT_UNLOADS parameter is set to <i>true</i> (default), nbrb initiates unloads of the drives that have exceeded the media unload delay. Drives become available more quickly to jobs that require different media servers or different media than the job that last used the drive. However, the loaded media or drive pair may not be available for jobs with less priority in the prioritized evaluation queue that can use the drive or media without unload.</p> <p>RB_DO_INTERMITTENT_UNLOADS=true</p>
RB_ENABLE_OPTIMIZATION	<p>When the RB_ENABLE_OPTIMIZATION parameter is set to <i>true</i> (default), this entry instructs nbrb to cache states of resource requests.</p> <p>RB_ENABLE_OPTIMIZATION=true</p>
RB_RESPECT_REQUEST_PRIORITY	<p>When the RB_RESPECT_REQUEST_PRIORITY parameter is set to <i>false</i> (default), nbrb continues to evaluate jobs in the prioritized job queue. As a result, a job is likely to reuse a drive more quickly after the drive has been released. However, some lower priority jobs may receive drives before higher priority jobs do.</p> <p>When the RB_RESPECT_REQUEST_PRIORITY parameter is set to <i>true</i>, nbrb restarts its evaluation queue at the top of the prioritized job queue after resources have been released.</p> <p>RB_RESPECT_REQUEST_PRIORITY=false</p>
RB_BREAK_EVAL_ON_DEMAND	<p>When a high priority request appears (for example, a tape span request, or a request for a synthetic or a duplication job), nbrb immediately interrupts the evaluation cycle. nbrb releases and unloads drives, if required before the evaluation cycle begins again.</p> <p>If the RB_BREAK_EVAL_ON_DEMAND parameter is set to <i>true</i> (default), interruptions of high priority jobs are not allowed and the evaluation cycle continues.</p> <p>RB_BREAK_EVAL_ON_DEMAND=true</p>
RB_MAX_HIGH_PRIORITY_QUEUE_SIZE	<p>Spanning requests and additional resources for an active duplication job are put in a special queue for priority processing. The RB_MAX_HIGH_PRIORITY_QUEUE_SIZE parameter sets the maximum number of requests that NetBackup allows in that queue. (Default: 100 requests.)</p> <p>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE=100</p>

Table 21-6 nbrbutil -changesettings parameters (continued)

Parameter	Description
RB_RELEASE_PERIOD	<p>The RB_RELEASE_PERIOD parameter indicates the interval that NetBackup waits before it releases a resource. (Default: 180 seconds.)</p> <p>RB_RELEASE_PERIOD=180</p>
RB_CLEANUP_OBSOLETE_DBINFO	<p>The RB_CLEANUP_OBSOLETE_DBINFO parameter indicates the number of seconds that can elapse between the cleanup of obsolete information in the nbrb database. (Default: 60 seconds.)</p> <p>RB_CLEANUP_OBSOLETE_DBINFO=60</p>
RB_MPX_GROUP_UNLOAD_DELAY	<p>The RB_MPX_GROUP_UNLOAD_DELAY parameter indicates the number of seconds that nbrb waits for a new job to appear before a tape is unloaded. (Default: 10 seconds.)</p> <p>RB_MPX_GROUP_UNLOAD_DELAY=10</p> <p>This setting can help avoid unnecessary reloading of tapes and applies to all backup jobs. During user backups, nbrb uses the maximum value of RB_MPX_GROUP_UNLOAD_DELAY and the Media mount timeout host property setting when nbrb unmounts the tape.</p> <p>During restores, Media mount timeout is used, not RB_MPX_GROUP_UNLOAD_DELAY.</p> <p>See “Timeouts properties” on page 199.</p>
RB_RETRY_DELAY_AFTER_EMM_ERR	<p>The RB_RETRY_DELAY_AFTER_EMM_ERR parameter 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. (Default: 60 seconds.)</p> <p>RB_RETRY_DELAY_AFTER_EMM_ERR=60</p>
RB_REEVAL_PENDING	<p>The RB_REEVAL_PENDING parameter 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. (Default: 60 seconds.)</p> <p>RB_REEVAL_PENDING=60</p>
RB_REEVAL_PERIOD	<p>The RB_REEVAL_PERIOD parameter indicates the time between evaluations if an outstanding request is not satisfied, and if no other requests or resources have been released. (Default: Five minutes must pass before the initial request is reevaluated.)</p> <p>RB_REEVAL_PERIOD=300</p>

For additional information about the `nbrbutil` utility, see the *Commands Reference Guide*.

Starting or stopping a service

The following procedure describes how to start or stop a NetBackup service.

To start or stop a service

- 1 In the **NetBackup Administration Console**, select **Activity Monitor** and then 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 In the **NetBackup Administration Console**, select **Activity Monitor** and then 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

In the **NetBackup Administration Console**, the Activity Monitor **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.

Table 21-7 lists and describes the NetBackup processes.

Table 21-7 NetBackup processes

Process	Port	Description
acsd	13702	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	None	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	13782	The NetBackup Client daemon, this process issues requests to and from the master server and the media server to start programs on remote hosts. On UNIX clients, <code>bpcd</code> can only be run in stand-alone mode. On Windows, <code>bpcd</code> always runs under the supervision of <code>bpinetd.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.
bpcompatd	None	The process for the NetBackup Compatibility service.
bpdbm	13721	The process for the NetBackup Database Manager service. The process that responds to queries that are related to the NetBackup catalog.
bpinetd	None	The process for the NetBackup Client service. The process that provides a listening service for connection requests. Note: To configure a BasicDisk storage unit that uses CIFS, <code>bpinetd</code> , <code>nbrmms</code> , and <code>vnetd</code> must share the same logon credentials as on the media server.
bpjava-msvc	13722	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	None	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	13723	The NetBackup Jobs Database Management daemon. This process queries and updates the jobs database.

Table 21-7 NetBackup processes (*continued*)

Process	Port	Description
bprd	13720	<p>The process for the NetBackup Request Daemon.</p> <p>The process that starts the automatic backup of clients and responds to client requests for file restores and user backups and archives.</p> <p>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.</p>
ltid	None	The process for the NetBackup Device Manager service.
NBConsole	None	The NetBackup Administration Console on the Windows platform.
nbemm	None	<p>The process for the NetBackup Enterprise Media Manager service.</p> <p>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.</p>
nbEvtMgr	None	<p>The process for the NetBackup Event Manager service.</p> <p>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.</p>
nbfdrv64	None	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	None	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	None	<p>The process for the NetBackup Job Manager service.</p> <p>The process that accepts the jobs that the Policy Execution Manager (<code>NBPEM</code>) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.</p>
nbpem	None	<p>The process for the NetBackup Policy Execution Manager service.</p> <p>It creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, <code>NBPEM</code> is notified and the appropriate Policy/Client tasks are updated.</p>
nbproxy	None	The process that safely allows multi-threaded NetBackup processes to use existing multi-threaded unsafe libraries.

Table 21-7 NetBackup processes (continued)

Process	Port	Description
nbrb	None	This process allocates storage units, tape drives, and client reservations for jobs. nbrb works with the Enterprise Media Manager (NBEMM).
nbrmms	None	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). Note: To configure a BasicDisk storage unit that uses CIFS, bpinetd, nbrmms, and vnetd must share the same logon credentials as on the media server. See “Configuring credentials for CIFS and disk storage units” on page 395.
nbsl	None	The process for the NetBackup Service Layer service. nbsl facilitates the communication between the graphical user interface and NetBackup logic.
nbstserv	None	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	None	The process for the NetBackup Service Monitor. Monitors the NetBackup services. When a service unexpectedly terminates, nbsvcmon attempts to restart the terminated service.
nbvault	None	If Vault is installed, the process for the NetBackup Vault Manager service.
ndmp	10000	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.
opr	None	The NetBackup Volume Manager (vmd) starts the opr operator request daemon. This process receives requests to mount and unmount volumes and communicates the requests to the NetBackup Device Manager ltid. The NetBackup Device Manager communicates the requests to the robotics through SCSI interfaces.
postgres	10085	The process for the NetBackup deduplication database. It runs on the deduplication storage server. Active only if the NetBackup Media Server Deduplication option is licensed.
spoold	None	The process for the NetBackup Deduplication Engine service. It runs on the deduplication storage server. Active only if the NetBackup Media Server Deduplication option is licensed.
t14d	13713	The t14d 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.

Table 21-7 NetBackup processes (*continued*)

Process	Port	Description
t18d t18cd	13705	<p>The <code>t18d</code> 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 <code>t18cd</code>.</p> <p>The <code>t18cd</code> process communicates with the TL8 robotics through SCSI interfaces.</p> <p>To share the tape library, <code>t18cd</code> runs on the NetBackup server that provides the robotic control.</p>
t1dd t1dcd	13711	<p>The <code>t1dd</code> 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 <code>t1dcd</code>.</p> <p>The <code>t1dcd</code> process communicates with the Tape Library DLT robotics through SCSI interfaces.</p> <p>To share the tape library, <code>t1dcd</code> runs on the NetBackup server that provides the robotic control.</p>
t1hd t1hcd	13717	<p>The <code>t1hd</code> 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 <code>t1hcd</code>.</p> <p>The <code>t1hcd</code> process runs on the NetBackup server that provides the robotic control and communicates with the TLH robotics through SCSI interfaces.</p>
t1md	13716	<p>The <code>t1md</code> 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.</p>
vmd	13701	<p>The process for the NetBackup Volume Manager service.</p>
vnetd	13724	<p>This process is preserved for backward compatiability. For example, when the 7.0.1 Java interface communicates with a 7.0 NetBackup server.</p> <p>The 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 <code>vnetd</code> port number.</p> <p>Note: To configure a BasicDisk storage unit that uses CIFS, <code>bpinetd</code>, <code>nbrmms</code>, and <code>vnetd</code> must share the same logon credentials as on the media server.</p> <p>See “Configuring credentials for CIFS and disk storage units” on page 395.</p>
vrts-auth-port	4032	<p>The Veritas Authorization Service verifies that an identity has permission to perform a specific task.</p>

Table 21-7 NetBackup processes (*continued*)

Process	Port	Description
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. Connections to NetBackup 7.0.1 and later use the <code>veritas_pbx</code> port.

Monitoring NetBackup processes in the Process Details dialog box

The following procedure describes how to view the details for a process.

To view the details for a process

- 1 In the **NetBackup Administration Console**, click **Activity Monitor**.

- 2 To view the details for a specific process, double-click on the process you want to display in the **Processes** tab. The **Process Details** dialog box appears that contains detailed information about your selected process.

Elapsed time	Specifies the total time (in seconds) since the process was created.
Handle count	Specifies the number of handles that a process currently uses.
Page faults per second	Specifies the rate of virtual memory Page Faults by the threads that run in this process.
Page file bytes	Specifies the current number of bytes that the process has used in the paging file(s).
Pool non-paged bytes	Specifies the number of bytes in the Non-paged Pool. The non-paged pool is a system memory area that acquires space from operating system components as they accomplish their tasks.
Peak page file bytes	Specifies the maximum number of bytes that the process has used in the paging file(s).
Peak virtual bytes	Specifies the maximum number of bytes of virtual address space that the process has used at any one time.
Peak working set	Specifies the maximum number of bytes in the set of memory pages that the process has used at any point in time.
Pool paged bytes	Specifies the number of bytes in the Paged Pool. The paged pool is a system memory area that acquires space from operating system components as they accomplish their tasks.
Privileged time	Specifies the percent of processor time that the process has spent in privileged mode.
Priority base	Specifies the current base priority of this process.
Private bytes	Specifies the current number of bytes this process allocated that cannot be shared with other processes.
Process ID (PID)	Specifies the unique identifier of this process. The ID numbers are reused, so they only identify a process for the lifetime of that process.
Process name	Specifies the name of the process.

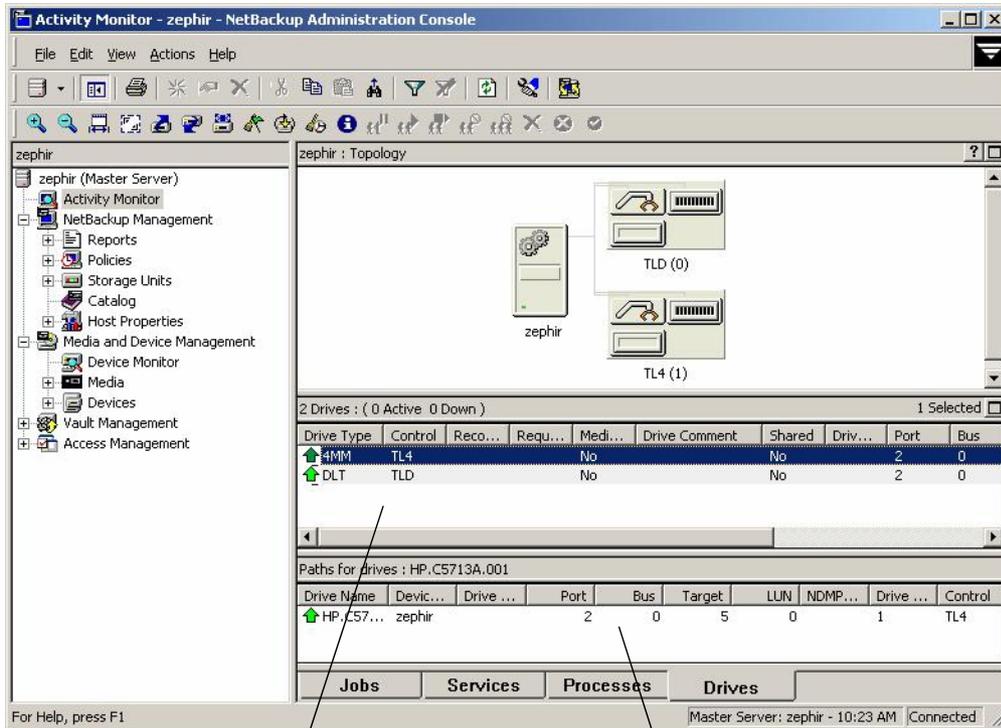
Processor time	Specifies the percentage of processor time (since the last refresh) that the process threads have used.
Server	Specifies the server currently selected.
Thread count	Specifies the number of threads currently active in this process.
Total privileged time	Specifies the Amount of processor time (in seconds) that the process has spent in privileged mode
Total processor time	Specifies the amount of process time (in seconds) that this process spent.
Total user time	Specifies the amount of processor time (in seconds) that the process has spent in user mode.
User time	Specifies the percentage of processor time that the process's threads have spent in user mode.
Virtual bytes	Specifies the current size in use of the virtual address space for a process.
Working set	Specifies the current number of bytes in use in the set of memory pages for a process.

- 3 In the **Process Details** dialog box, click the up or down arrow to see the details of the next process in the list.

About the Drives tab

In the **NetBackup Administration Console**, the **Drives** tab in the Activity Monitor displays the status of NetBackup drives on the master server being monitored.

Figure 21-4 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 tape drives.

To monitor NetBackup tape drives

- 1 In the **NetBackup Administration Console**, click the **Activity Monitor**.
- 2 In the right pane, 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.

Action	Description
Clean Now	Starts 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	Resets 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	Sets the number of mount hours between drive cleanings.

About the jobs database

NetBackup uses the `install_path\NetBackup\bin\admincmd\bpdbjobs -clean` command to delete done jobs periodically.

By default, the `bpdbjobs` process deletes all completed jobs that are more than three days old. By default, the `bpdbjobs` process retains more recent done jobs until the three-day retention period expires.

You may want to keep jobs in the jobs database longer than the default of three days. To do this, you must change the default value.

If the `bprd` NetBackup request daemon is active, `bprd` starts the `bpdbjobs` 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 `bpdbjobs` at other times by using `cron` or alternate methods.

About changing the default values

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 | bpssetconfig
```

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 | bpssetconfig
```

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 values:

- 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.

Table 21-8 BPDBJOBS_OPTIONS environment variable options

Option	Description
<code>-keep_hours <i>hours</i></code>	Use with the <code>-clean</code> option to specify how many hours <code>bpdbjobs</code> keeps unsuccessfully completed jobs. Default: 78 hours. To keep both successful and both failed jobs longer than the default of 78 hours, <code>keep_successful_hours</code> must be used with <code>keep_hours</code> .
<code>-keep_successful_hours <i>hours</i></code>	Use with the <code>-clean</code> option to specify how many hours <code>bpdbjobs</code> keeps successfully completed jobs. The number of hours must be less than or equal to <code>keep_hours</code> . Values outside the range are ignored. Default: 78 hours.
<code>-keep_days <i>days</i></code>	Use with the <code>-clean</code> option to specify how many days <code>bpdbjobs</code> keeps completed jobs. Default: 3 days.
<code>-keep_successful_days <i>days</i></code>	This value must be less than the <code>-keep_days</code> value. Use with the <code>-clean</code> option to specify how many days <code>bpdbjobs</code> keeps successfully completed jobs. Default: 3 days.

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:

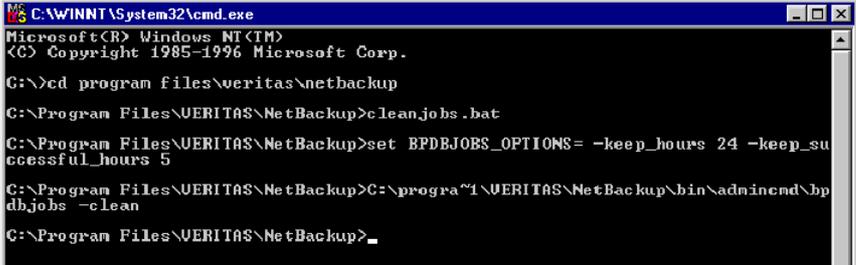
```
set BPDBJOBS_OPTIONS= -keep_hours 24 -keep_successful_hours 5
C:\progra~1\VERITAS\NetBackup\bin\admincmd\bpdbjobs -clean
```

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 21-5 is a screen capture of `cleanjobs.bat` being run:

Figure 21-5 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 the *NetBackup Commands Reference Guide*.

Enabling the bpdbjobs debug log

If you need detailed information on `bpdbjobs` activities, use the following procedure:

Enabling the bpdbjobs debug log

- ◆ 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*.

About the Device Monitor

Use the **NetBackup Administration Console 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 **NetBackup Administration Console 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

In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**. If requests await action or if NetBackup acts on a request, the **Pending Requests** pane appears. For example, if a tape mount requires a specific volume, the request appears in the **Pending Requests** pane. If NetBackup requires a specific volume for a restore operation, NetBackup loads or requests the volume. 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.

Table 21-9 Pending states

Pending state	Description
Pending request	<p>Specifies that 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 Pending Requests pane.</p> <p>NetBackup assigns pending status to a mount request when it cannot determine the following:</p> <ul style="list-style-type: none"> ■ Which stand-alone drive to use for a job. ■ Which drive in a robot is in Automatic Volume Recognition (AVR) mode.
Pending action	<p>Specifies that 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 Pending Requests pane. Pending actions usually occur with drives in robotic libraries.</p>

About pending requests for storage units

In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**. 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

These 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 **NetBackup Administration Console 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 258.

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 800.
See “Denying a request” on page 800.

Resubmitting a request

After you correct a 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 336.

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**.

Auditing NetBackup operations

This chapter includes the following topics:

- About NetBackup auditing
- Viewing the current audit settings
- Configuring auditing on a NetBackup master server
- User identity in the audit report
- Auditing host property changes
- Using the command line `-reason` or `-r` option
- Viewing the audit report
- `nbaudit` log behavior
- Retaining and backing up audit trail records

About NetBackup auditing

An audit trail is a record of user-initiated actions in a NetBackup environment. Essentially, auditing gathers the information to help answer who changed what and when they changed it.

Auditing NetBackup operations can help provide information in the following areas:

General tracking	Customers can gain insight from audit trails while they investigate unexpected changes in a NetBackup environment. For example, it might be found that the addition of a client or a backup path has caused a significant increase in backup times. The audit report can indicate that an adjustment to a schedule or to a storage unit configuration might be necessary to accommodate the policy change.
Regulatory compliance	Auditing creates a record of who changed what and when it was changed. The record complies with guidelines such as those required by the Sarbanes-Oxley Act (SOX).
Corporate change management	For customers who must adhere to internal change management policies, NetBackup auditing offers a method to adhere to such policies.
Troubleshooting	The information from NetBackup auditing helps NetBackup Support to troubleshoot problems for customers.

The NetBackup Audit Manager (`nbaudit`) runs on the master server and audit records are maintained in the EMM database. If the master server is not the EMM server in the environment, `nbaudit` uses Remote ODBC to access Sybase ASA on the EMM server.

The Audit Manager provides the mechanism to query and report on auditing information. For example, an administrator can search specifically for information based on when an action occurred, actions performed by a specific user, actions performed in a specific content area, or changes to the audit configuration.

When auditing is configured (by default auditing is turned on), the following NetBackup user-initiated actions are recorded and available to view by using the `nbauditreport` command or by using Symantec OpsCenter:

- The following items in the NetBackup Administration Console are audited:
 - **Policies**
Adding, deleting, or updating policy attributes, clients, schedules, and backup selections lists.
 - **Activity Monitor**
Canceling, suspending, resuming, or deleting any type of job creates an audit record.
 - **Storage units**
Adding, deleting, or updating storage units.

Note: Storage Lifecycle Policies related actions are not audited.

- **Storage servers**
Adding, deleting, or updating storage servers.
- **Disk pools and Volume pools**
Adding, deleting, or updating disk or volume pools.
- **Host properties**
Updating host properties. (NetBackup Access Control (NBAC) must be enabled for host property auditing to occur.)
- Initiating a restore job.
A restore job is the only job type for which the initiation is audited. For example, when a backup job begins, no audit record is created.
- Changes to the audit configuration.
- Starting and stopping the NetBackup Audit Manager (`nbaudit`).

Note: By default, audit configuration changes or starting and stopping `nbaudit` is audited, even if auditing is disabled.

- Changes to the `bp.conf` file (UNIX) or the registry (Windows).
For NetBackup to audit changes to the `bp.conf` file or the registry, NetBackup Access Control (NBAC) must be enabled. These changes must be made by using either `bpsetconfig` or the **Host Properties** utility in the NetBackup Administration Console. Changes that are made by manually editing the `bp.conf` file or the registry are not audited.
See “Auditing host property changes” on page 808.
For more information about configuring NetBackup Access Control, see the *NetBackup Security and Encryption Guide*.

The following actions are not audited and do not display in the audit report:

Any failed actions.

Failed actions are logged in NetBackup error logs. Failed actions do not display in audit reports because a failed attempt does not bring about a change in the NetBackup system state.

The ramifications of a configuration change.

The results of a change to the NetBackup configuration are not audited. For example, the creation of a policy is audited, but the jobs that result from its creation are not.

The completion status of a manually initiated restore job. While the act of initiating a restore job is audited, the completion status of the job is not audited. Nor is the completion status of any other job type, whether initiated manually or not. The completion is displayed in the Activity Monitor.

Internally initiated actions. NetBackup-initiated internal actions are not audited. For example, the scheduled deletion of expired images, scheduled backups, or periodic image database cleanup is not audited.

Viewing the current audit settings

To view the current audit configuration, use either the `nbemmcmd` command on a NetBackup master server or view the settings using Symantec OpsCenter.

For directions about how to use Symantec OpsCenter to configure auditing, see the *OpsCenter Administrator's Guide*.

To view the current audit settings

- 1 From a command prompt, locate the `nbemmcmd` command on the master server in the following directory:

- On UNIX:

```
/usr/opensv/netbackup/bin/admincmd
```

- On Windows:

```
Install_path\Veritas\NetBackup\bin\admincmd
```

- 2 Enter the `nbemmcmd` command using the following syntax:

```
nbemmcmd -listsettings -machinename masterserver
```

Where *masterserver* is the master server in question.

Note: The options are case-sensitive.

- 3 The output lists many configuration settings. Among them are the following:

- `AUDIT="ENABLED"`

Indicates that auditing is turned on.

- `AUDIT="DISABLED"`

Indicates that auditing is turned off.

- `AUDIT_RETENTION_PERIOD="90"`

Indicates that if auditing is enabled, the records are retained for this length of time (in days) and then deleted. The default audit retention period is 90 days. A value of 0 (zero) indicates that the records are never deleted.

Configuring auditing on a NetBackup master server

Auditing is enabled by default in new installations. However, the default may be enabled or disabled after an upgrade to NetBackup 7.1, depending on the version before the upgrade. For information on auditing configuration after an upgrade, see the following topic:

See “Auditing configuration after upgrading to NetBackup 7.1” on page 807.

NetBackup auditing can be configured directly on a NetBackup master server or by using Symantec OpsCenter.

The master server settings for enabling or disabling audit logging and setting the retention period are configured in the **Manage > Hosts** section of OpsCenter. Within OpsCenter, the expiration setting for Audit logs is configured under **Settings > Purge**. See the *OpsCenter Administrator's Guide* for more detail.

To configure auditing on a master server, use the `nbevmcmd` command with the `-changesetting` option.

To configure NetBackup auditing on a master server

- 1 From a command prompt, locate the `nbevmcmd` command on the master server in the following directory:
 - On UNIX:
`/usr/opensv/netbackup/bin/admincmd`
 - On Windows:

Install_path\Veritas\NetBackup\bin\admincmd

2 Enter the `nbemmcmd` command using the following syntax:

```
nbemmcmd -changesetting -AUDIT DISABLED -machinename masterserver
```

Where `-AUDIT DISABLED` turns off auditing on the master server that is indicated.

Note: The options are case-sensitive.

In the following example, auditing has been turned off for `server1`.

For example:

```
nbemmcmd -changesetting -AUDIT DISABLED -machinename server1
```

3 Configure the audit retention period using the following syntax:

```
nbemmcmd -changesetting -AUDIT_RETENTION_PERIOD  
number_of_days -machinename masterserver
```

Where *number_of_days* indicates (in days) how long audit records are to be retained for the audit report. If no retention period is indicated, the default audit retention period is 90 days.

Note: An audit retention period value of 0 (zero) indicates that the records are never deleted.

Symantec OpsCenter downloads the audit records periodically and retains them for a period of time that is configurable in OpsCenter. Therefore, retaining the audit records on the master server is only necessary if you want to view audit reports using the command line on the master server.

See the following topic for more information.

See “Retaining and backing up audit trail records” on page 814.

In the following example, the records of user actions are to be retained for 30 days and then deleted.

```
nbemmcmd -changesetting -AUDIT_RETENTION_PERIOD 30  
-machinename server1
```

The two options can be combined in one command line, as in the following example:

```
nbemmcmd -changesetting -AUDIT ENABLED -machinename server1  
-AUDIT_RETENTION_PERIOD 30
```

4 Run `nbauditreport` to display a report of the audited information.

See “Viewing the audit report” on page 810.

Auditing configuration after upgrading to NetBackup 7.1

The auditing configuration after upgrading NetBackup to 7.1 varies, depending on the NetBackup version before the upgrade.

Table 22-1 Auditing configuration after upgrade

Version before upgrade	Audit configuration after upgrade to NetBackup 7.1
7.0.1	The upgraded configuration is the same as it was before the upgrade.
7.0 or 6.5.x	<p>The upgraded configuration is that same as the default configuration for 7.0.1:</p> <p>Auditing is disabled and the default retention is 365 days.</p> <p>After upgrading both a master server and a remote EMM configuration, be sure to restart the audit service <code>nbaudit</code> on the master server.</p>

User identity in the audit report

The audit report lists the identity of the user who performed a specific action. The identity includes the user name, the domain, and the domain type of the authenticated user.

If NetBackup Access Control (NBAC) is not used in an environment, administrators must have administrator (or root) privileges to configure and run NetBackup. In large environments, multiple administrators may share the same root logon.

To differentiate between administrators in the audit report, NBAC must be configured. When NBAC is enabled, the audit report displays the actual user identities that are associated with audited actions. Information about NBAC installation and configuration is available in the *NetBackup Security and Encryption Guide*.

Auditing host property changes

NetBackup audits host property changes if the administrator uses either the `bpsetconfig` command or the equivalent property in the **Host Properties** utility.

The following criteria must be met for auditing to take place:

- The environment must be configured for NetBackup Access Control (NBAC).
- The host on which the `bp.conf` file or the registry changes are made must be at NetBackup 7.1.
- The administrator must use either the `bpsetconfig` command or the equivalent property in the **Host Properties** utility for auditing to occur. Changes made

directly to the `bp.conf` file or to the registry (that is, without using `bpsetconfig`), are not audited.

For example, taking a client off-line is not performed using the `bpsetconfig` command, so this operation would not show up in the audit log.

Using the command line `-reason` or `-r` option

Many commands offer the `-reason` option for administrators to use to indicate why the action was performed. The reason displays in the audit report.

The `-reason` string must be no more than 512 characters. Command lines that accept the `-reason` option display an error if the string is over 512 characters.

Keep in mind that the audit reason cannot begin with a dash character (`-`). The reason also cannot contain a single quotation mark (`\'`).

The following commands accept the `-reason` option (or `-r` option in the case of `bpsetconfig`):

- `bpdbjobs`
- `bpplcatdrinfo`
- `bpplclients`
- `bppldelete`
- `bpplinclude`
- `bpplinfo`
- `bpplsched`
- `bpplschedrep`
- `bpolicynew`
- `bpsetconfig`

Note: The `bpsetconfig` command accepts the `-r` option instead of the `-reason` option.

- `bpstuadd`
- `bpstudel`
- `bpsturep`
- `nbdecommission`

- `nbdevconfig`
- `vmpool`

For more information on using the commands, see the *NetBackup Commands Reference Guide*.

Viewing the audit report

To view the audit report, use either the `nbauditreport` command on a NetBackup master server or view the settings using Symantec OpsCenter.

Within OpsCenter, the **Monitor > Audit Trails** section provides the details of the Audit logs and allows you to export that information to Excel or save as a .pdf file. See the *OpsCenter Administrator's Guide* for more detail.

If auditing is enabled but a user action fails to create an audit record, the audit failure is captured in the `nbaudit` log.

The failure to create an audit record has no effect on the user action that was performed.

If the user action succeeds, an exit code is returned that reflects the successful action. If auditing of the action fails, NetBackup status code 108 is returned (`Action succeeded but auditing failed`).

Note: The NetBackup Administration Console (Windows and UNIX (`jnbSA`)) does not return an exit status code 108 when auditing fails.

To view the NetBackup audit report

- 1 From a command prompt, locate the `nbauditreport` command on the master server in the following directory:

- On UNIX:

```
/usr/opensv/netbackup/bin/admincmd
```

- On Windows:

```
Install_path\Veritas\NetBackup\bin\admincmd
```

- 2 In its simplest form, enter the `nbauditreport` command using the following syntax:

```
nbauditreport
```

The `nbauditreport` can also be used with a number of options.

Note: The options are case-sensitive.

<code>-help</code>	Use for assistance with the command at the command prompt.
<code>-sdate</code> <"MM/DD/YY [HH:[MM[:SS]]]">	Use to indicate the start date and time of the report data you want to view.
<code>-edate</code> <"MM/DD/YY [HH:[MM[:SS]]]">	Use to indicate the end date and time of the report data you want to view.
<code>-ctgy POLICY</code>	Use <code>-ctgy POLICY</code> to display information pertaining to policy changes.
<code>-ctgy JOB</code>	Use <code>-ctgy JOB</code> to display information pertaining to jobs.
<code>-ctgy STU</code>	Use <code>-ctgy STU</code> to display information pertaining to storage units.
<code>-ctgy STORAGESRV</code>	Use <code>-ctgy STORAGESRV</code> to display information pertaining to storage servers.
<code>-ctgy POOL</code>	Use <code>-ctgy POOL</code> to display information pertaining to storage pools.
<code>-ctgy AUDITCFG</code>	Use <code>-ctgy AUDITCFG</code> to display information pertaining to audit configuration changes.
<code>-ctgy AUDITSVC</code>	Use <code>-ctgy AUDITSVC</code> to display information pertaining to the starting and stopping of the NetBackup Audit service (nbaudit).
<code>-ctgy BPCONF</code>	Use <code>-ctgy BPCONF</code> to display information pertaining to changes in the <code>bp.conf</code> file.
<code>-user</code> <username[:domainname]>	Use to indicate the name of the user for whom you'd like to display audit information.
<code>-fmt SUMMARY</code>	If no report output format option (<code>-fmt</code>) is specified, the <code>SUMMARY</code> option is used by default.

<code>-fmt DETAIL</code>	The <code>-fmt DETAIL</code> option displays a comprehensive list of audit information. For example, when a policy is changed, this view lists the name of the attribute, the old value, and the new value.
<code>-fmt PARSABLE</code>	The <code>-fmt PARSABLE</code> option displays the same set of information as the <code>DETAIL</code> report but in a parsable format. The report uses the pipe character () as the parsing token between the audit report data.
<code>[-nottruncate]</code>	Use the <code>-nottruncate</code> option to display the old and new values of a changed attribute on separate lines in the details section of the report. Note: <code>-nottruncate</code> is valid only with the <code>-fmt DETAIL</code> option.
<code>[-pagewidth <NNN>]</code>	Use the <code>-pagewidth</code> option to set the page width for the details section of the report. Note: <code>-pagewidth</code> is valid only with the <code>-fmt DETAIL</code> option.
<code>[-order <DTU DUT TDU TUD UDT UTD>]</code>	The <code>-order</code> option is valid only with <code>-fmt PARSABLE</code> . Use it to indicate the order in which the information appears. Use the following parameters: <ul style="list-style-type: none">■ D (Description)■ T (Timestamp)■ U (User)

3 The audit report contains the following details:

DESCRIPTION	The details of the action that was performed. The details include the new values that are given to a modified object and the new values of all attributes for a newly created object. The details also include the identification of any deleted objects.
USER	The identity of the user who performed the action. The identity includes the user name, the domain, and the domain type of the authenticated user. See “User identity in the audit report” on page 808.

TIMESTAMP	The time that the action was performed. The time is given in Coordinated Universal Time (UTC) and indicated in seconds. (For example, 12/06/10 10:32:48.)
CATEGORY	The category of user action that was performed. The CATEGORY displays only with the <code>-fmt DETAIL PARSABLE</code> options. Examples include the following: <ul style="list-style-type: none"> ■ AUDITSVC START, AUDITSVC STOP ■ POLICY CREATE, POLICY MODIFY, POLICY DELETE
ACTION	The action that was performed. The ACTION displays only with the <code>-fmt DETAIL PARSABLE</code> options. Examples include the following: <ul style="list-style-type: none"> ■ START, STOP ■ CREATE, MODIFY, DELETE
REASON	The reason that the action was performed. A reason displays if a reason was specified in the command that created the change. The <code>bpsetconfig</code> command accepts the <code>-r</code> option. See “Using the command line <code>-reason</code> or <code>-r</code> option” on page 809. The reason displays only with the <code>-fmt DETAIL PARSABLE</code> options.
DETAILS	An account of all of the changes, listing the old values and the new values. Displays only with the <code>-fmt DETAIL PARSABLE</code> options.

If an exit status appears in the output, look up the code in the NetBackup Administration Console (Troubleshooter), the online Help, or the *Status Codes Reference Guide*.

Figure 22-1 shows the default contents of an audit report that was run on `server1`.

Figure 22-1 Summary audit report example

```
[root@server1 admincmd]# ./nbauditreport
TIMESTAMP      USER           DESCRIPTION
09/23/2010 14:40:54  root@server1  Policy 'test_pol_1' was created
09/23/2010 14:40:54  root@server1  Schedule 'full' was added to Policy
'test_pol_1'
09/22/2010 17:10:23  root@server1  Audit setting(s) of master server 'server1'
were modified

Audit records fetched: 3
```

nbaudit log behavior

The `nbaudit` log is found in the following location:

- On UNIX:

`/usr/opensv/logs/nbaudit`

- On Windows:

`Install_path\Veritas\NetBackup\logs\nbaudit`

If auditing is enabled but a user action fails to create an audit record, the audit failure is captured in the `nbaudit` log.

The `nbaudit` service behaves in the following manner when it creates audit records:

- The audit record limits the details of an entry to a maximum of 4096 characters. (For example, the Policy name.) The remaining characters are truncated while stored in the audit database.

- The audit record limits the restore image IDs to a maximum of 1024 characters. The remaining characters are truncated while stored in the audit database.

- Rollback operations are not audited.

Some operations are carried out as multiple steps. For example, creating an MSDP-based storage server consists of multiple steps. Every successful step is audited. Failure in any of the steps results in a rollback, or rather, the successful steps may need to be undone. The audit record does not contain details about rollback operations.

Retaining and backing up audit trail records

By default, audit records are kept for 90 days. To change the default, use the `nbeemcmd -changesetting` command with the `-AUDIT_RETENTION_PERIOD` option.

See “Configuring auditing on a NetBackup master server” on page 805.

Based on the configured retention setting, the NetBackup Audit Service (`nbaudit`) deletes expired audit records once every 24 hours at 12:00 A.M. (local time).

The audit records are kept in audit tables that are part of the NetBackup database. The tables are retained for as long as the `-AUDIT_RETENTION_PERIOD` indicates and are backed up as part of the NetBackup catalog backup.

To make sure that audit records are not missed from a catalog backup, configure the catalog backup frequency to be less frequent or equal to the

`-AUDIT_RETENTION_PERIOD`.

Symantec OpsCenter downloads the audit records periodically from the EMM database. OpsCenter retains the records for a period of time that is configured

within OpsCenter. Therefore, retaining the audit records on the NetBackup master server is only necessary if you want to view audit reports using the command line on the master server. Audit records can also be exported from OpsCenter.

Reporting in NetBackup

This chapter includes the following topics:

- About the Reports utility
- Running a report
- Copying report text to another document
- Saving or exporting a report
- Printing a report
- Status of Backups report
- Client Backups report
- Problems report
- All Log Entries report
- Images on Media report
- Media Logs report
- Images on Disk report
- Disk Logs report
- Disk Storage Unit Status report
- Disk Pool Status report
- Images on Tape report
- Tape Logs report
- Tape Contents report

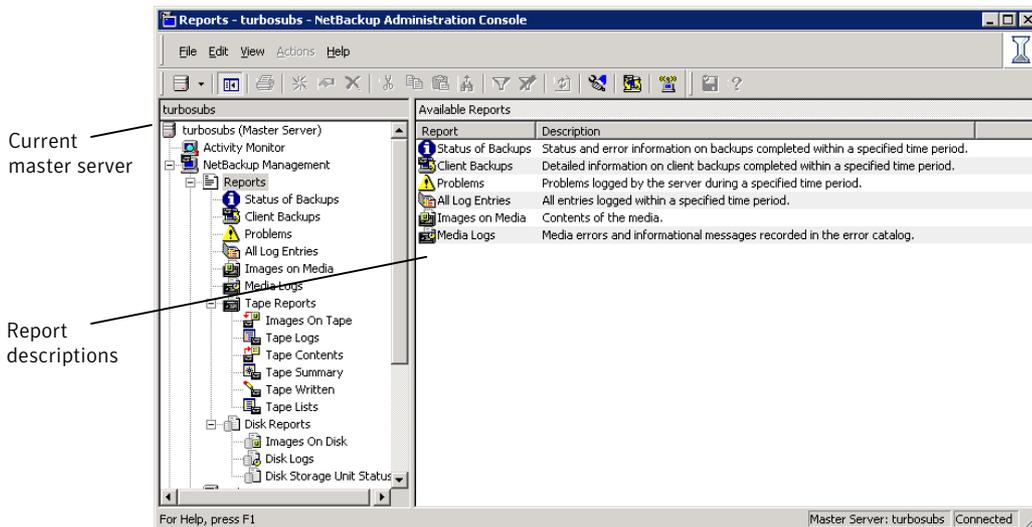
- Tape Summary report
- Tape Written report
- Tape Lists report

About the Reports utility

Use the **Reports** utility to generate reports to verify, manage, and troubleshoot NetBackup operations. NetBackup reports display information according to job status, client backups, and media contents. Use the **Troubleshooter** to analyze the cause of the errors that appear in a NetBackup report.

In the **Reports** window, in the right pane, you can select a report to run or manage report data.

Figure 23-1 NetBackup Reports utility



NetBackup offers many different reports to view the information you need.

For information about Vault reports, see the *NetBackup Vault Administrator's Guide*.

See “Running a report” on page 819.

See “Running the Troubleshooter” on page 42.

See “Copying report text to another document” on page 819.

See “Saving or exporting a report” on page 820.

See “Printing a report” on page 820.

NetBackup also offers auditing and audit reports.

See “Viewing the audit report” on page 810.

Running a report

The following procedure describes how to run a NetBackup report from the **Reports** utility.

To run a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.

NetBackup runs the report for the master server that is currently selected. To run a report on a different master server, on the **File** menu, click **Change Server**.

See “Accessing remote servers” on page 835.

- 2 In the left pane, click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report. For example, select the media servers and clients on which to run the report, and select the time period that the report should span.
- 4 Click **Run Report**.

Copying report text to another document

The following procedure describes how to copy the text from a NetBackup report and paste it into a spreadsheet or other document.

To copy report text to another document

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report, and click **Run Report**.

- 4 Select the rows of the report you want to copy by holding down the **Shift** or **Ctrl** key.
- 5 On the **Edit** menu, click **Copy**.
- 6 Paste the selected rows into a spreadsheet or other document.

Saving or exporting a report

The following procedure describes how to save or export a NetBackup report.

To save or export a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report and click **Run Report**.
- 4 On the **File** menu, click **Export**.
- 5 In the **Save As** dialog box, select the location where you want to save the report, and specify the file name.
- 6 Click **Save**.

Printing a report

The following procedure describes how to print a NetBackup report.

To print a report

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Reports**.
- 2 In the left pane, click the name of the report you want to run.
For some reports, you must first expand a report group, and then click the name of the report.
- 3 Select the criteria for what to include or exclude in the report and click **Run Report**.
- 4 On the **File** menu, click **Print**.

Status of Backups report

The **Status of Backups** report shows status and error information about the jobs that completed within the specified time period. If an error occurred, a short explanation of the error is included in the report.

Client Backups report

The **Client Backups** report shows detailed information about the backups that completed within the specified time period.

Problems report

The Problems report generates a list of the problems that the server has logged during the specified time period. The information in this report is a subset of the information that is obtained from the All Log Entries report.

All Log Entries report

The **All Log Entries** report generates a list of 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 rates and predict backup times for future backups. (The transfer rate does not appear for multiplexed backups.)

Images on Media report

The **Images on Media** report generates a list of the media contents 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 report

The **Media Logs** report shows the media errors or the informational messages that are recorded in the NetBackup error catalog.

Images on Disk report

The **Images on Disk** report generates the image list present on the disk storage units that are connected to the media server. The **Images on Disk** report is a subset of the **Images on Media** report, showing only disk-specific columns.

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.

Disk Logs report

The **Disk Logs** report displays all error logs related to disk-based backup and recovery. The **Disk Logs** report is a subset of the **Media Logs** report.

Disk Storage Unit Status report

The **Disk Storage Unit Status** report displays the state of the disk storage units in the current NetBackup configuration. (For example, the total capacity and the used capacity of the disk storage unit.)

Multiple storage units can point to the same disk pool. When the report query searches by storage unit, the report counts the capacity of disk pool storage multiple times.

Storage units that reference disk groups do not display capacity values.

Disk Pool Status report

The **Disk Pool Status** report generates the details of one or more disk pools.

This report displays only when an Enterprise Disk Option is installed.

Images on Tape report

The **Images on Tape** report generates the contents of the tape-based media as recorded in the NetBackup image catalog. The **Images on Tape** is a subset of the **Images on Media** report.

Tape Logs report

The **Tape Logs** report displays all error logs related to tape-based backup and recovery. The **Tape Logs** report is a subset of the **Media Logs** report.

Tape Contents report

The **Tape Contents** report (formerly known as the Media Contents report) generates a list of 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.

Before running this report, you can choose to override the default job priority for the job. The default priority is specified in the **Default Job Priorities** host properties.

See “Default Job Priorities properties” on page 105.

Tape Summary report

The **Tape Summary** report summarizes active and nonactive volumes for the specified media owner according to expiration date. It also shows how many volumes are at each retention level. In verbose mode, the report shows each media ID and the expiration date.

Nonactive media are those with a status of FULL, FROZEN, SUSPENDED, or IMPORTED. Other volumes are considered active.

Expired volumes with a status of FULL, SUSPENDED, or IMPORTED do not appear in the report. However, expired volumes with a FROZEN status do appear in the report. NetBackup deletes other expired volumes from the media catalog when it runs backups. Also, an expired volume of a different status can display if the report is run between the time the volume expires and the time that the next backup is done.

Tape Written report

The **Tape Written** report 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.

Tape Lists report

The **Tape Lists** report generates information about the volumes that are allocated for backups for the selected media owner or media ID.

This report does not show media for disk type storage units. For the backups that are saved to disk storage units, use the **Images on Media** report or the **Images on Disk** report.

See “Images on Media report” on page 821.

See “Images on Disk report” on page 822.

Administering NetBackup

- Chapter 24. Management topics
- Chapter 25. Accessing a remote server
- Chapter 26. Using the NetBackup-Java administration console
- Chapter 27. Alternate server restores
- Chapter 28. Managing client restores
- Chapter 29. Powering down and rebooting NetBackup servers
- Chapter 30. About Granular Recovery Technology

Management topics

This chapter includes the following topics:

- NetBackup naming conventions
- Wildcard use 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 also used for foreign languages. Spaces are only allowed in a drive comment.

Wildcard use in NetBackup

NetBackup recognizes the following wildcard characters in areas where wildcards can be used. (For example, in the paths of include and exclude file lists.)

The following table shows the wildcards that can be used in various NetBackup dialog boxes and lists.

Table 24-1 Wildcard use in NetBackup

Wildcard	Use
*	<p>An asterisk serves 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 and UNIX 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> <pre>/etc/*.conf</pre>
?	<p>A question mark 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 and UNIX 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> <pre>c:\system\log??_03</pre>

Table 24-1 Wildcard use in NetBackup (*continued*)

Wildcard	Use
[]	<p>A pair of square brackets indicates any single character or range of characters that are separated with a dash.</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> <p>Brackets are not valid wildcards under all circumstances for all clients:</p> <ul style="list-style-type: none"> ■ Brackets used as wildcards in include and exclude lists: UNIX clients: Allowed Windows clients: Allowed ■ Brackets used as wildcards in policy backup selections lists: UNIX clients: Allowed Windows clients: Not allowed; the use of brackets in policy backup selections lists causes backups to fail with a status 71.
{ }	<p>Curly brackets can be used in the backup selection list, the include list, and the exclude list for UNIX clients only.</p> <p>A pair of curly brackets (or braces) indicates 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><code>{*1.doc, *.pdf}</code> refers to <code>file1.doc</code>, <code>file1.pdf</code>, <code>file2.pdf</code></p> <p>Note: 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 examples 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`

See “Backup Selections tab” on page 598.

How to administer devices on other servers

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 a **NetBackup Administration Console** that is connected to the master server.

Alternatively, you can administer the devices on a specific media server from a **NetBackup Administration Console** connected to that media server. To perform this task, change to or log in to the media server by using one of the following methods:

- In an existing instance of the **NetBackup Administration Console**, expand **File > Change Server** and change to the media server.
- Start the **NetBackup Administration Console** on the media server.
- See “About choosing a remote server to administer” on page 839.

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 computer 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 the 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.

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 scenario 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` is 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 666.

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.

About 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.

About 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
- About adding a NetBackup server to a server list
- About choosing a remote server to administer
- About using the Remote Administration Console
- About using the Java Windows Administration Console
- About running the NetBackup Administration Console on a NetBackup client
- About 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.

Use the following procedure to access a remote server.

To access a remote server

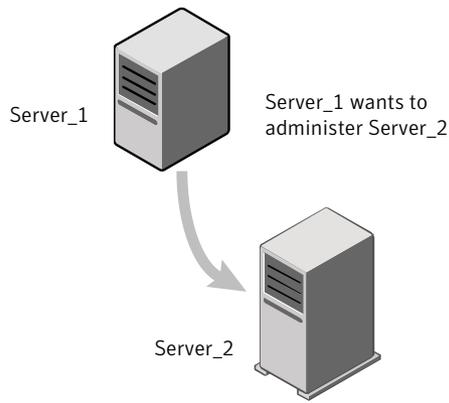
- 1 Ensure that the remote server is accessible to the local server.
See “About adding a NetBackup server to a server list” on page 836.
- 2 Indicate the remote server that you want to administer.
See “About choosing a remote server to administer” on page 839.

About 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.

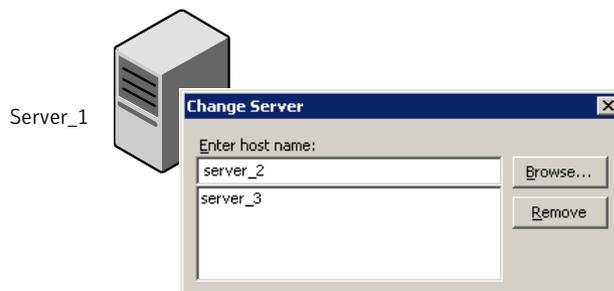
Figure 25-1 assumes that server_1 wants to administer server_2.

Figure 25-1 Server accessing a remote server



On server_1, in the **NetBackup Administration Console**, in the menu bar, select **File > Change Server** and type **server_2** as the host name in the **Change Server** window. Click **OK**.

Figure 25-2 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 837.

Other reasons may exist why a remote server is inaccessible:

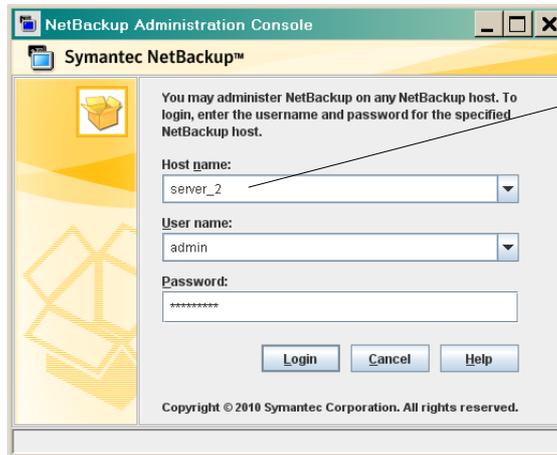
See “About troubleshooting remote server administration” on page 845.

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

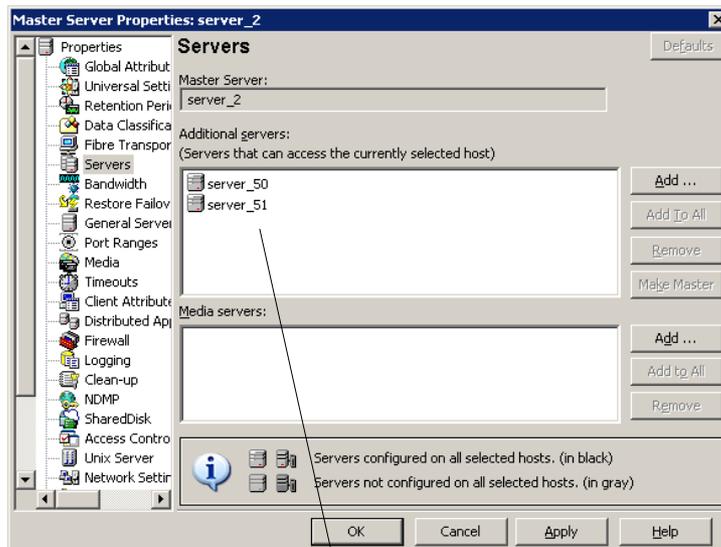
- 1 Access the server properties of the remote server in one of the following ways:
 - Physically go to the Windows destination host (server_2) and start the **NetBackup Administration Console**.
 - If it is installed, start the **Java Windows Administration Console**, on the local Windows host. Indicate 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.

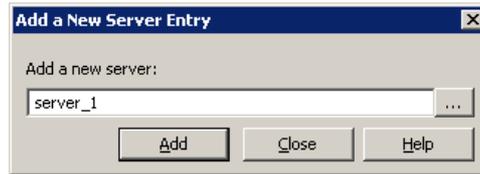
- 2 Expand **Host Properties > Master Server**.
- 3 Double-click the server name (server_2) to view the properties.
- 4 Select the **Servers** tab to display the server list.

Since the server list does not include server_1, server_2 considers server_1 to be an invalid server.

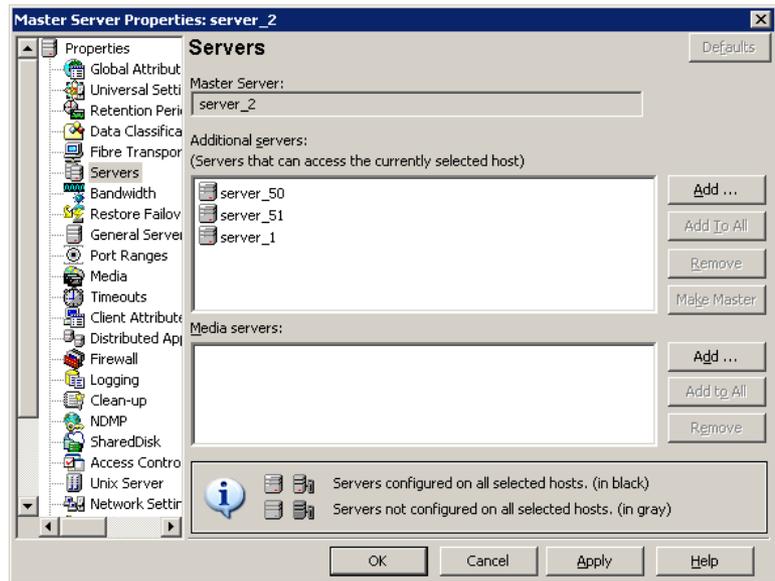


Currently, server_2 allows remote access to two servers: server_50 and server_51

- 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.

About 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 840.

- Specify the remote server in the host name field to start the NetBackup-Java console.

See “Indicating a remote system upon login” on page 841.

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 837.

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 “About using the Remote Administration Console” on page 842.
 - To start the console on the Windows system where the **Java Windows Administration Console** is installed, select **Start > Programs > Symantec NetBackup > NetBackup-Java Version 7.1**.

- 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 837.

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 on as a different user.

Indicating a remote system upon login

Use the following procedure to indicate a remote system upon logging on 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 computer.

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 7.1**.
 - To start the **NetBackup Administration Console** on a NetBackup-Java capable computer, run `j_nbSA` as follows:

```
/usr/opensv/java/jnbSA
```

- 2 In the **NetBackup Administration Console** login 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 “About the NetBackup-Java Administration Console” on page 847.

See “Restricting access to NetBackup-Java applications on Windows” on page 856.

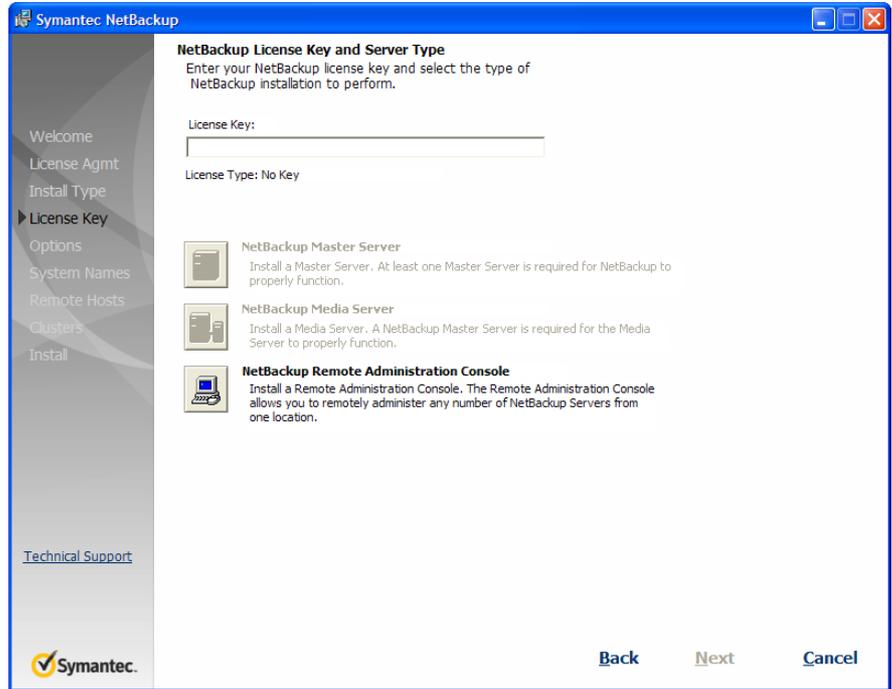
About using the Remote Administration Console

Install the **NetBackup Remote Administration Console** on a Windows computer 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 computer to be backed up like any other client. No master server software or media server software is installed.

Figure 25-3 shows how to install the Remote Administration Console.

Figure 25-3 Remote Administration Console selection on the installation screen



Start the **NetBackup Remote Administration Console** from the menu toolbar. Select **File > Change Server** to change to another NetBackup server.

See “Adding a server to a remote server list” on page 837.

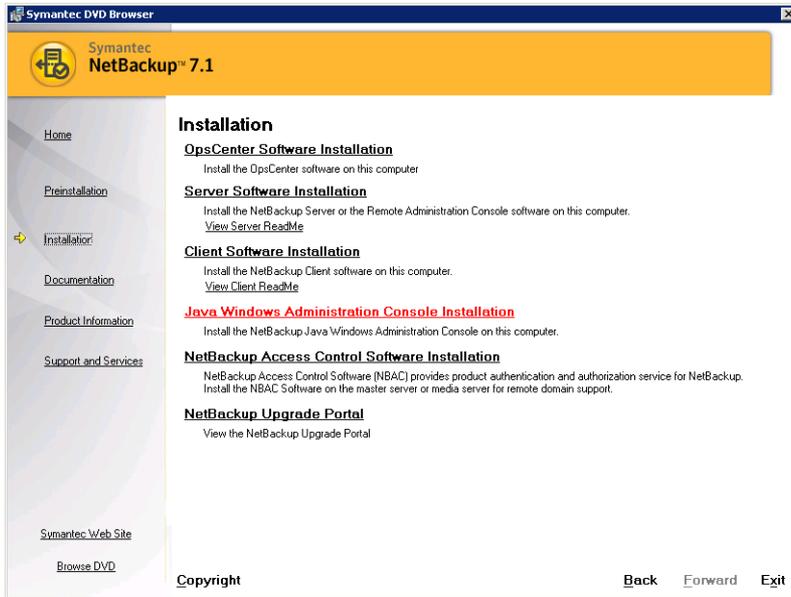
See “About choosing a remote server to administer” on page 839.

About 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 25-4 shows how to install the Java Windows Administration Console.

Figure 25-4 Java Windows Administration Console selection on the installation screen



After it is installed, select **Start > Symantec NetBackup > NetBackup-Java Version 7.1** to start the **Java Windows Administration Console**.

See “About the NetBackup-Java Administration Console” on page 847.

About running the NetBackup 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.

About 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 “About adding a NetBackup server to a server list” on page 836.
 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 (`bpdbm`) and NetBackup Request Daemon (`bpird`) on the remote server if it is Windows.
- The NetBackup Database Manager and NetBackup Request Daemon 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:

- About the NetBackup-Java Administration Console
- About authorizing NetBackup-Java users
- Authorization file (auth.conf) characteristics
- About authorizing nonroot users for specific applications
- About authorizing specific tasks in jbpSA
- About authorizing NetBackup-Java users on Windows
- Restricting access to NetBackup-Java applications on Windows
- Runtime configuration options
- About logging the command lines that the NetBackup interfaces use
- About customizing jnbSA and jbpSA with bp.conf entries
- About improving NetBackup-Java performance
- About adjusting time zones in the NetBackup-Java console

About the NetBackup-Java 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 “About using the Java Windows Administration Console” on page 843.
- 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 7.1** on a Windows system on which the **Java Windows Administration Console** is installed

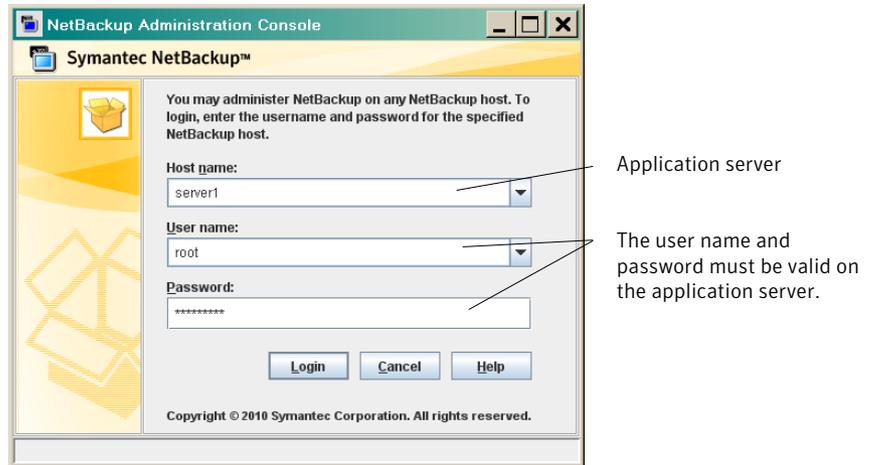
Then the administrator logs on to the application server on the host that is specified in the logon dialog box.

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 is the host that is specified in the **NetBackup Administration Console** logon dialog box and authenticates the logon credentials of the user. The credentials are authenticated by using standard UNIX user account data and associated APIs.

Note: To log in to any **NetBackup Administration Console**, your login credentials must be authenticated from the connecting master or media server. This is true whether or not NetBackup Access Control (NBAC) is in use.

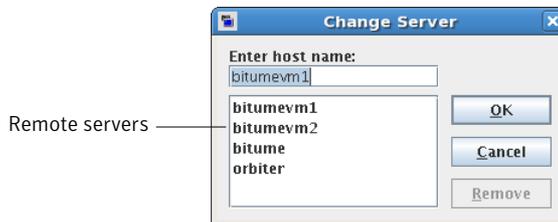
Figure 26-1 NetBackup logon dialog box



The server that is usually the object of all administrative tasks is the host that is 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 26-2 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 “About adding a NetBackup server to a server list” on page 836.

See “Indicating a remote system upon login” on page 841.

About 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 can 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 lets root users 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.

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 “About authorizing nonroot users for specific applications” on page 853.
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 “About authorizing specific tasks in jbpSA” on page 854.
asterisk (*)	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 let anyone with any valid user name 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.

About 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:

<code>ALL</code>	Indicates that the user has administrative privileges for all of the applications that are listed in this table.
<code>AM</code>	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`.

About authorizing specific tasks in jbpSA

The Backup, Archive, and Restore interface can be configured to let only a user 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 “About server-directed restores” on page 881.

See “About client-redirected restores” on page 882.

- 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:

Table 26-1 Identifiers for the `JBP` keyword in the `auth.conf` file

Identifier	Description
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
```

About 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.

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 who 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 “About the NetBackup-Java Administration Console” on page 847.

Restricting access to NetBackup-Java applications on Windows

Use the following procedure to restrict access to one or more of the NetBackup-Java applications.

To restrict access to one or more of the NetBackup-Java applications

- 1 Create the following file on the Windows system:

```
nbservice_install_path\java\auth.conf
```

- 2 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/opensv/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.

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 “VNETD_PORT” on page 862.

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 **Backup, 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_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.

PBX_PORT

The `PBX_PORT` configuration entry specifies the pbx port.

This option is present by default and uses the following format:

```
PBX_PORT=1556
```

VNETD_PORT

The `VNETD_PORT` is the configured port for the `vnetd` daemon process and is registered with the Internet Assigned Number Authority (IANA).

This option uses the following format:

```
VNETD_PORT=13724
```

Symantec recommends that this port not be changed. If changes are necessary, make the change on all NetBackup hosts in the relevant NetBackup cluster.

This option is preserved for backward compatibility when the 7.0.1 JAVA interface is used to communicate with a 7.0 NetBackup server.

See the *NetBackup Installation Guide*.

The value must be set in the corresponding `nbj.conf` (UNIX) or `setconf.bat` (Windows) configuration option.

About logging 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 **Backup, 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.

About customizing 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 let the administrator and users 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`.

About improving 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.

About running the Java console locally

On Windows platforms, select **Start > Symantec NetBackup > NetBackup-Java Version 7.1** 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.

On supported UNIX platforms, the console is run locally if `jnbSA` or `jbpsA` 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 `jnbSA` or `jbpsA` 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.

About running a console locally and administering 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.

About enhancing 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.

About determining better performance when console is run locally or uses 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.

NetBackup-Java performance 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.

NetBackup-Java performance 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 the Backup, Archive, and Restore client console.

Table 26-2 shows the files that contain configuration entries.

Table 26-2 Files containing configuration entries

File	Description
<code>/usr/opensv/java/auth.conf</code>	Authorization options.
<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).

About 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 Administration Console** for administration

of remote NetBackup hosts. (In this context, a remote NetBackup host may either be the host that is specified in the administration 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 Administration 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 Administration Console** when servers in different time zones are administered.

For example, open a **NetBackup-Java Administration 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 Administration Console**.

Do not open a new window in the first **NetBackup-Java Administration 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 in the NetBackup-Java console

Use the following procedure in the NetBackup-Java console to adjust the time zone or to use daylight savings time.

To adjust the time zone in the NetBackup-Java console

- 1 In the **NetBackup Administration Console**, or in the **Backup, Archive, and Restore** client dialog box, 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 in the NetBackup-Java console

Use the following procedure to configure a custom time zone in the NetBackup-Java console.

To configure a custom time zone in the NetBackup-Java console

- 1 In the **NetBackup Administration Console**, or in the **Backup, Archive, and Restore** client dialog box, 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, see the following table to set the DST start time:

Begin DST on a specific date	Select Absolute date and indicate the month and day To begin DST on April 5, set as follows:
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:
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:
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:
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
- About supported configurations for alternate server restores
- About 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.

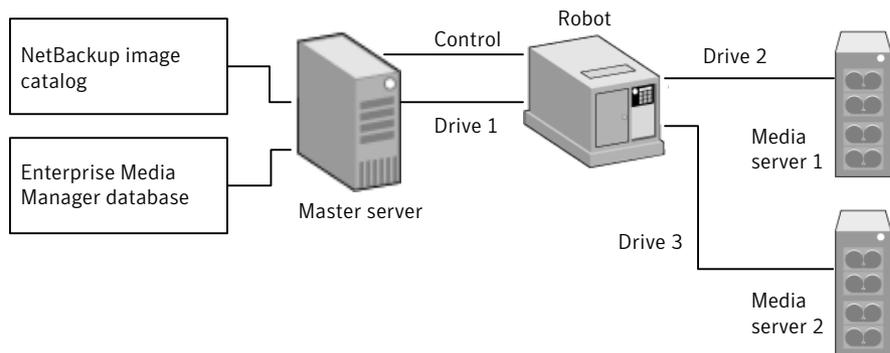
More information is available in the *NetBackup Commands Reference Guide*.
See “Expiring and importing media for alternate server restores” on page 878.

About 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 27-1 and Figure 27-2 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.

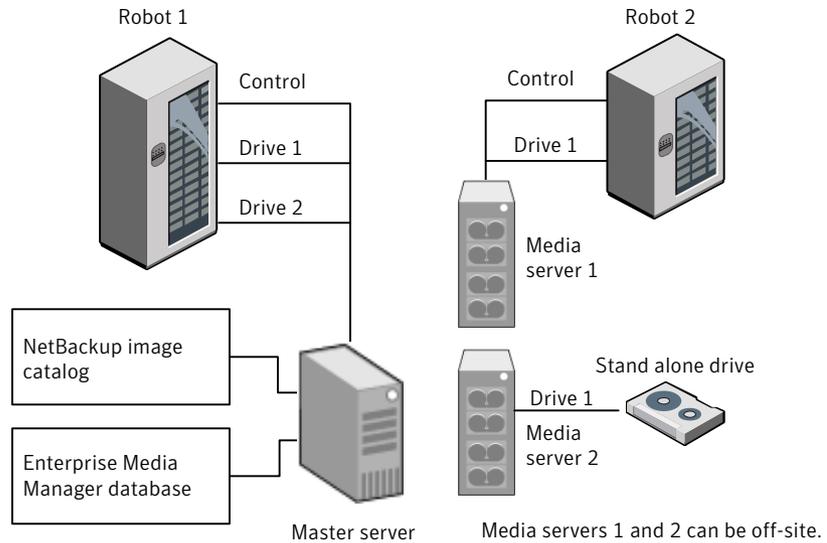
Figure 27-1 NetBackup servers that share robotic peripherals



Assume the following in Figure 27-1:

- 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 27-2 NetBackup servers with separate non-shared peripherals



Assume the following in Figure 27-2:

- 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.

About 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:

- Modify the NetBackup catalogs.
See “About modifying the NetBackup catalogs” on page 874.
- Override the original server for restores.
See “Overriding the original server for restores” on page 875.
- Enable automatic failover to an alternate server

See “About enabling automatic failover to an alternate server” on page 877.

About modifying the NetBackup catalogs

To perform alternate server restores by modifying the NetBackup catalogs, change the contents of the NetBackup catalogs. 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**, in the left pane, expand **NetBackup Management > Host Properties**. Depending on the type of server to override, click on either **Master Servers** or **Media Servers**.
See “General Server properties” on page 128.
- 2 In the right pane, click on the selected server to open the **General Server** host properties dialog box.
- 3 In the **General Server** host properties dialog box, click on the **Add** button to open the **Add Media Override settings** window. Add entries for the original backup server and the restore server and click the **Add** button in the **Add Media Override settings** window.
- 4 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**:
In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**. Click on **Master Servers**. In the right pane, click on the selected server to open the **General Server** host properties dialog box of the master server.
In the **General Server** host properties dialog box, click on the **Add** button to open the **Add Media Override settings** window. Add entries for the original backup server and the restore server and click the **Add** button in the **Add Media Override settings** window.
 - By modifying the `bp.conf` file on a UNIX NetBackup server:
As `root` add the following entry to the

```
/usr/opensv/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*.

About 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 184.

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 are as follows:

■ By using the **NetBackup Administration Console**:

In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Host Properties**. Click on **Master Servers** to open the **Master Server Properties** dialog box. In the left pane, click on **Restore Failover** to open the **Restore Failover** host properties dialog box. In the **Restore Failover** host properties dialog box, click on the **Add** button to open the **Add Failover Servers** window. Add entries for the media server and the failover restore server(s) and click the **Add** button in the **Add Failover Servers** window. Click **OK**.

■ By modifying the `bp.conf` file on a UNIX NetBackup server:

As `root`, add the following entry to the `/usr/opensv/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

It may be necessary to expire media and then import it, even with the alternate server restore capabilities.

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:

- About server-directed restores
- About client-redirected restores
- About restoring files and access control lists
- How to improve search times by creating an image list
- About restoring the System State

About 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
In the **NetBackup Administration Console**, in the toolbar, click **File > Backup, Archive, and Restore**.
Select **File > NetBackup Client Properties > General**, then 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 computer.

About 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.

About 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):
In the **NetBackup Administration Console**, in the toolbar, select **File > Backup, Archive, and Restore**. In the **Backup, Archive, and Restore** dialog box, in the toolbar, 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:
In the **Backup, Archive, and Restore** dialog box, select **File > Specify NetBackup Machines and Policy Type**. In the **Specify NetBackup Machines and Policy Type** dialog box, select the client name as the **Source client for restores**.

About allowing 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.

About allowing 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.

About allowing 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 of a 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 of a redirected client restore using the altnames file

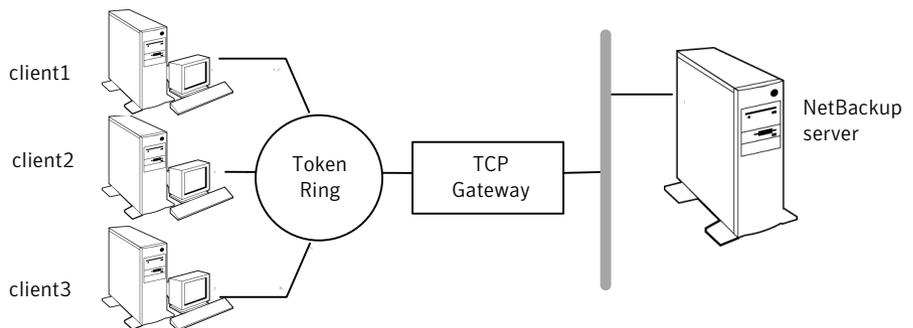
This example explains how `altnames` provides restore capabilities to clients that do not use their own host name when they connect to the NetBackup server.

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 28-1.

Figure 28-1 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 “About allowing a single client to perform redirected restores” on page 884.

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 “About allowing redirected restores of a client’s files” on page 884.

Note that this example requires no changes on the clients.

Restore the files.

See “About allowing redirected restores of a client’s files” on page 884.

See “About allowing a single client to perform redirected restores” on page 884.

Example of how to troubleshoot a 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 Servers** > Double-click on the master server. In the **Master Server Properties** dialog box, in the left pane, click on **Universal Settings**. Select the **Enable Performance Data Collection** property check box.

- Create the debug log directory for the NetBackup Request Daemon:

```
Install_path\NetBackup\logs\bprd
```

- On the master server, stop and restart the NetBackup Request Daemon. 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 Daemon to identify the failing name combination:

```
Install_path\NetBackup\logs\bprd\mmddy.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 Servers** > Double-click on master server. In the **Master Server Properties** dialog box, in the left pane, click on **Clean-up**. Clear the **Keep Logs** property check box.
- 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.

About 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.

About restoring the files that have ACLs

By default, the NetBackup-modified GNU tar (`/usr/openv/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.

About restoring the 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:

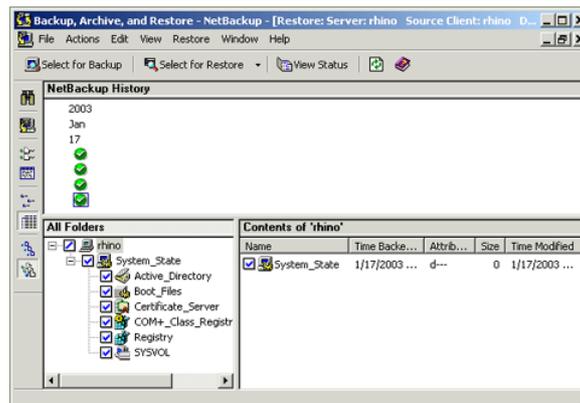
- The System State should be restored in its entirety. Do not restore selected files.
- Although incremental backups of the System State can be configured, NetBackup always performs a full backup. Therefore, only the most recent backup of the System State must be restored.
- Do not redirect a System 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.

See 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
- Shutting down all NetBackup services on Windows
- Starting up all NetBackup services on Windows
- Rebooting a NetBackup server
- Rebooting a NetBackup media server

Powering down and rebooting NetBackup servers

To close and restart NetBackup servers, use the following recommended procedure.

To power down a server

- 1 In the **NetBackup Administration Console**, in the left pane, 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 In the **NetBackup Administration Console**, in the left pane, 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\bpdown.exe
```

- 5 From the command line, enter:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 6 Power down the server.

Shutting down all NetBackup services on Windows

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 on Windows

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 also starts automatically.

- Start the NetBackup Request Daemon 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`) also starts.

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 for Windows

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.

NBFS is the NetBackup File System (NBFS) service that runs on the media server. NBFS 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

Table 30-1 Configuring NFS in a Windows 2008 or Windows 2008 R2 environment,

Step	Action	Description
Step 1	Stop and disable the Portmapper service.	Before you install NFS on the media server or client(s), look for the ONC Portmapper service. If it exists, stop it and disable it. Otherwise, the installation of NFS Services for Windows fails.
Step 2	Enable NFS.	Enable NFS on the following: <ul style="list-style-type: none"> ■ 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 901.
Step 3	Disable Server for NFS.	You can disable the Server for NFS on the following: <ul style="list-style-type: none"> ■ The NetBackup media server ■ All Active Directory domain controllers or ADAM/LDS hosts. See “Disabling the Server for NFS” on page 906.

Table 30-1 Configuring NFS in a Windows 2008 or Windows 2008 R2 environment, *(continued)*

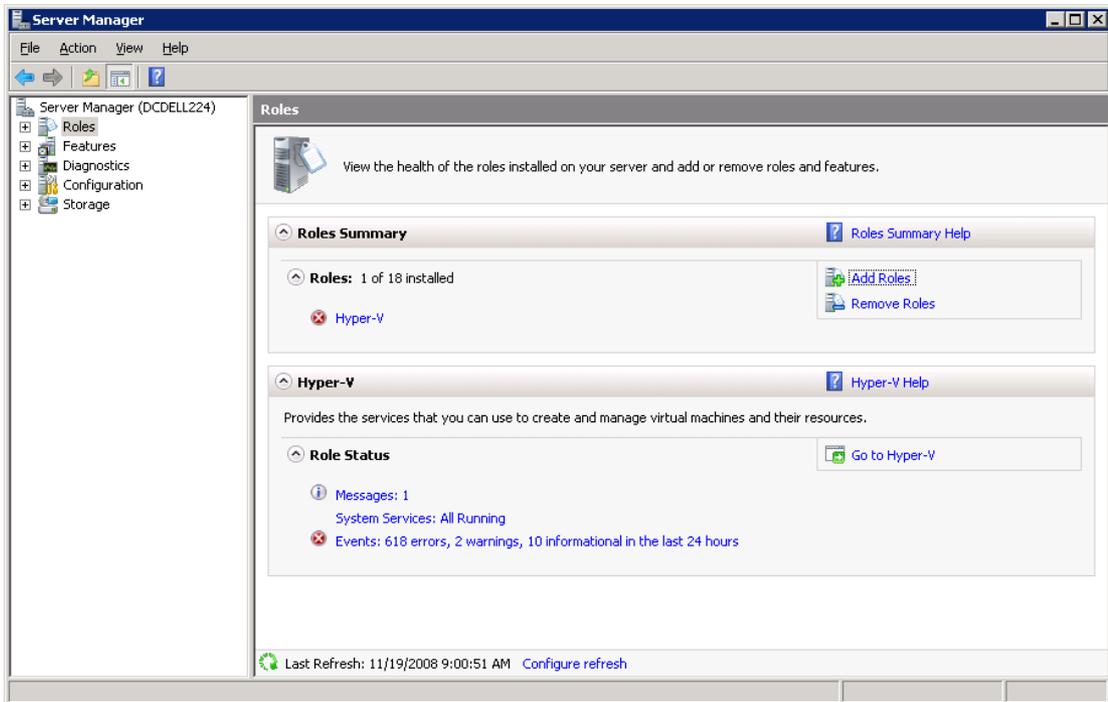
Step	Action	Description
Step 4	Disable Client for NFS.	<p>You can disable the Client for NFS on the NetBackup media server. See “Disabling the Client for NFS on the media server” on page 905.</p> <p>If the Active Directory domain controller or ADAM/LDS host resides on the media server, do not disable the Client for NFS.</p>

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.

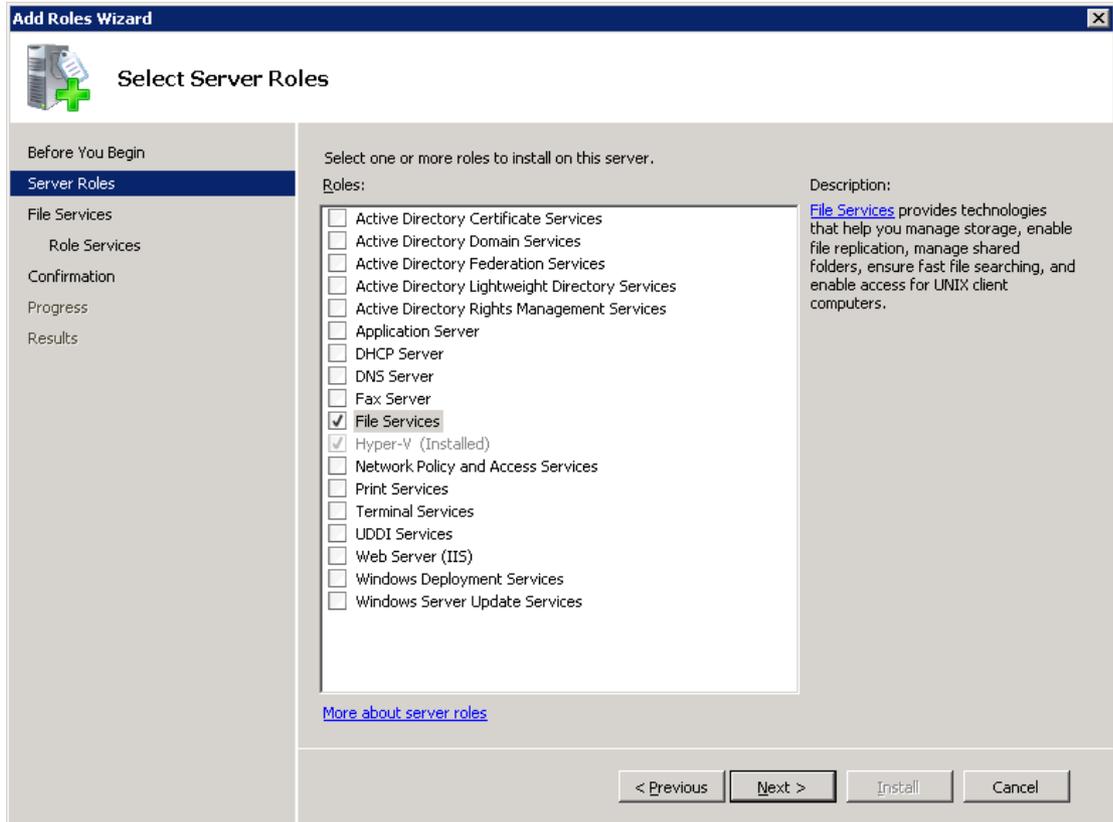
To enable Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2

- 1 Open the Server Manager.
- 2 In the left pane, click **Roles** and, in the right pane, click **Add Roles**.



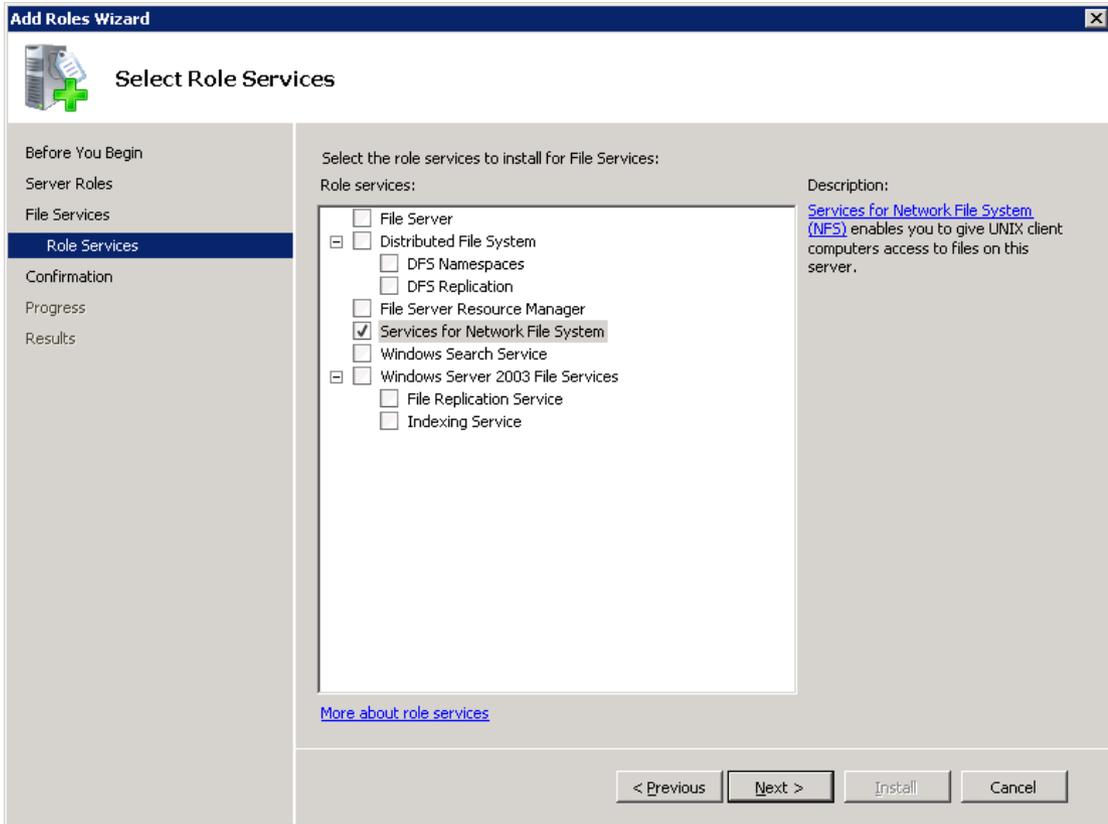
- 3 In the Add Roles Wizard, on the **Before You Begin** page, click **Next**.

- 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.

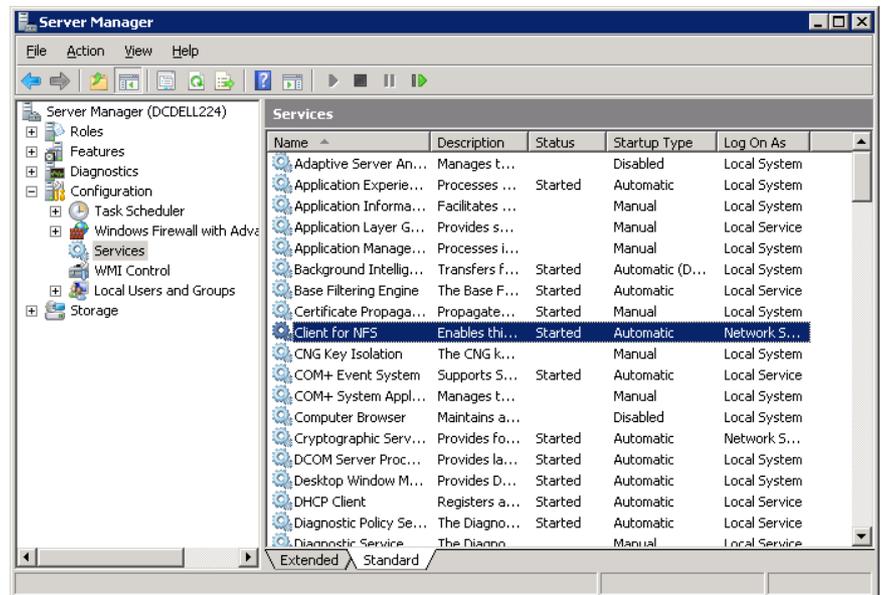
- For a host that is only the NetBackup media server, you can disable the Server for NFS and the Client for NFS.
- For a host that is only an Active Directory domain controllers or ADAM/LDS host, you can disable the Server for NFS.

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.

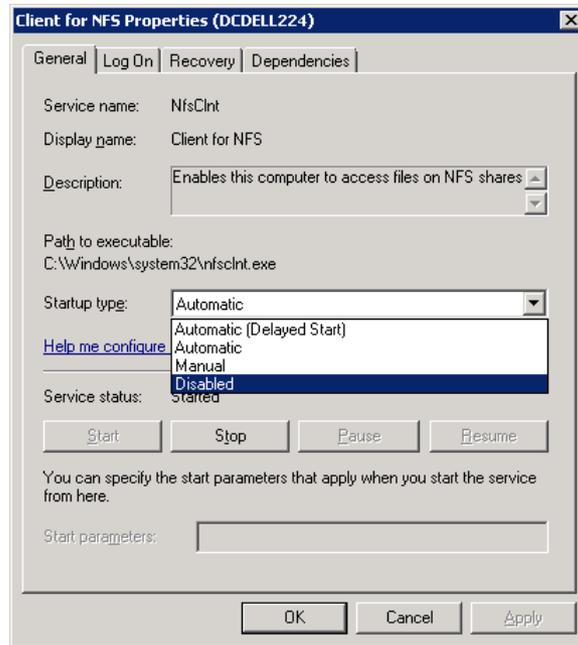
To disable the Client for NFS on the NetBackup media server

- 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**.

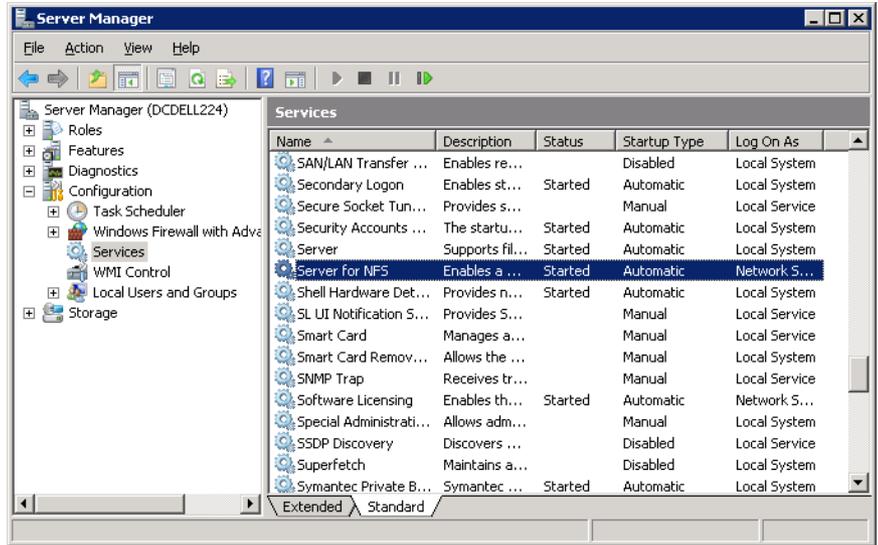
Disabling the Server for NFS

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**.

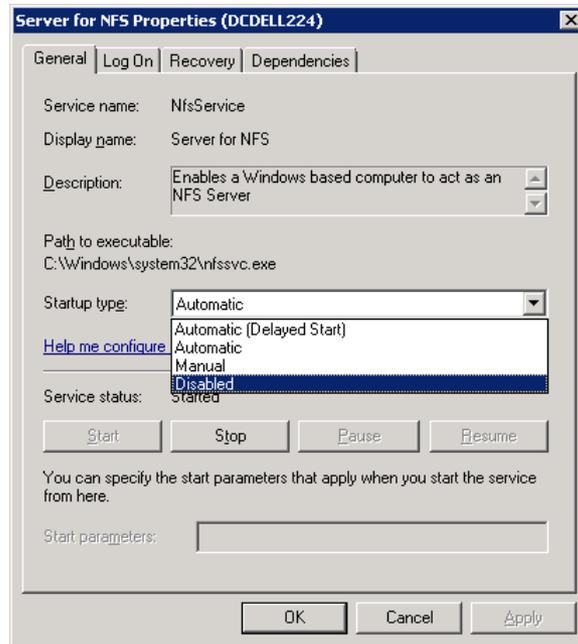
3 Click Services.



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.

Table 30-2

Step	Action	Description
Step 1	Install the necessary NFS components on the NetBackup media server.	See Table 30-3 on page 909. See “Installing Services for NFS on the Windows 2003 R2 SP2 media server” on page 909.
Step 2	Install the necessary NFS components on all Active Directory domain controllers or ADAM/LDS hosts.	See Table 30-3 on page 909. See “Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2” on page 912. Note: If the Active Directory domain controllers or ADAM/LDS host resides on the media server, install all the components on the media server.

Table 30-3 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

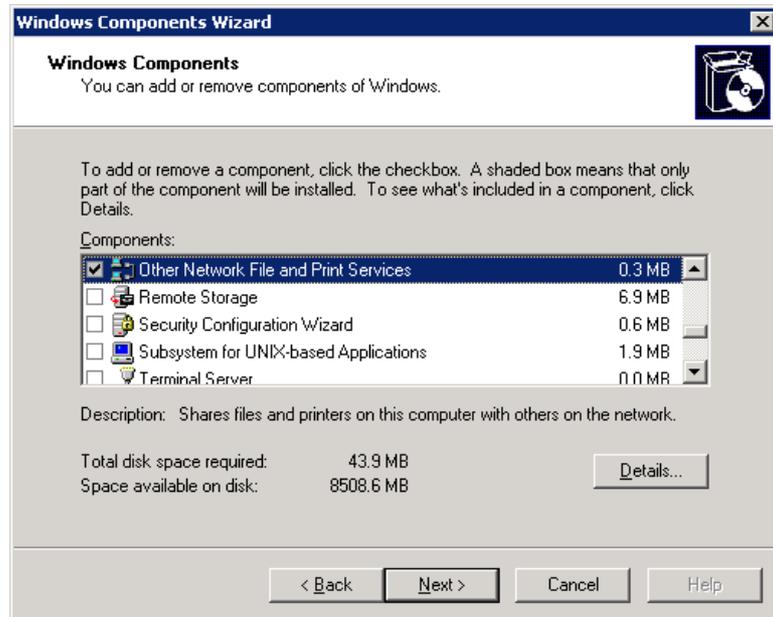
Installing Services for NFS on the Windows 2003 R2 SP2 media server

This topic describes how to install Services for NFS on a Windows 2003 R2 SP2 media server.

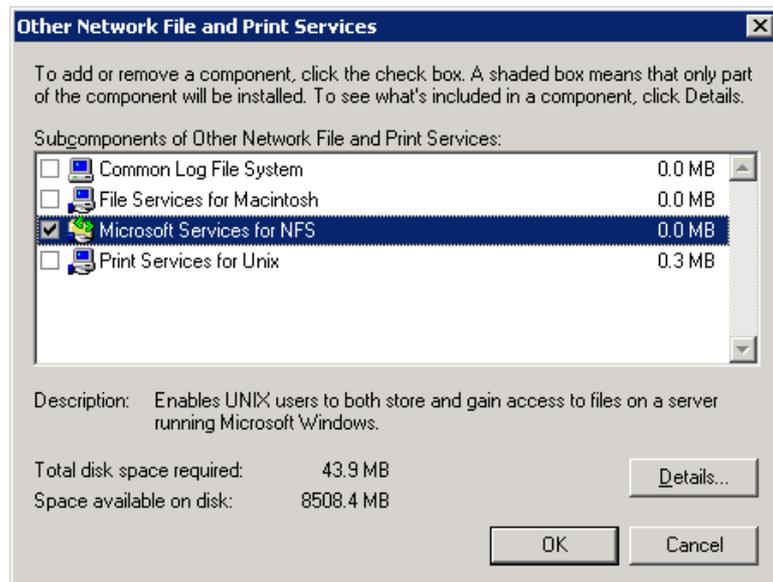
To install Services for NFS on the Windows 2003 R2 SP2 media server

- 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.



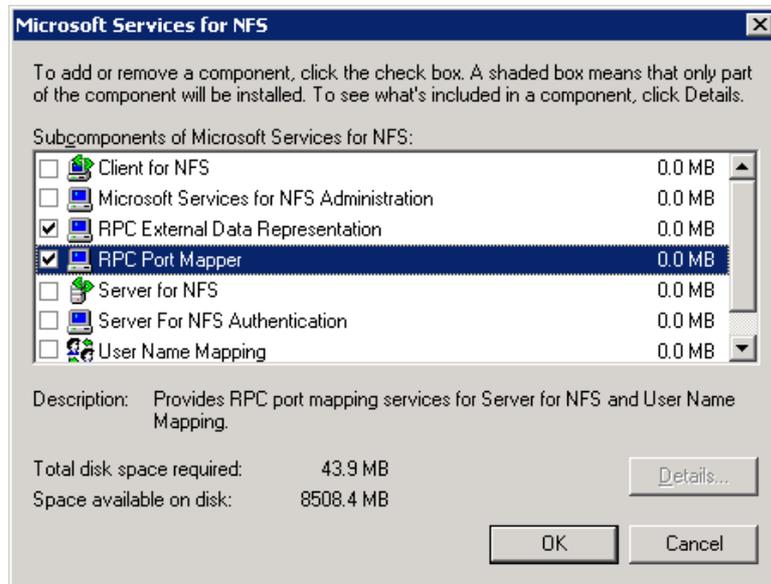
4 Check Microsoft Service for NFS and click Details.



- 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

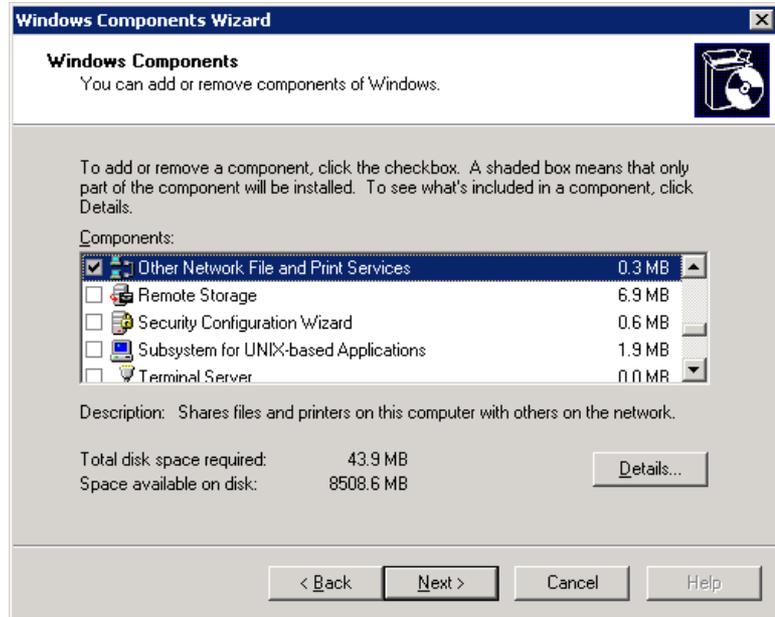
This topic describes how to install NFS on the NetBackup clients with Windows 2003 R2 SP2. Only the clients that are Active Directory domain controllers or ADAM/LDS hosts require NFS. If an Active Directory domain controllers or ADAM/LDS host is also a media server, you must follow a different procedure.

See “Installing Services for NFS on the Windows 2003 R2 SP2 media server” on page 909.

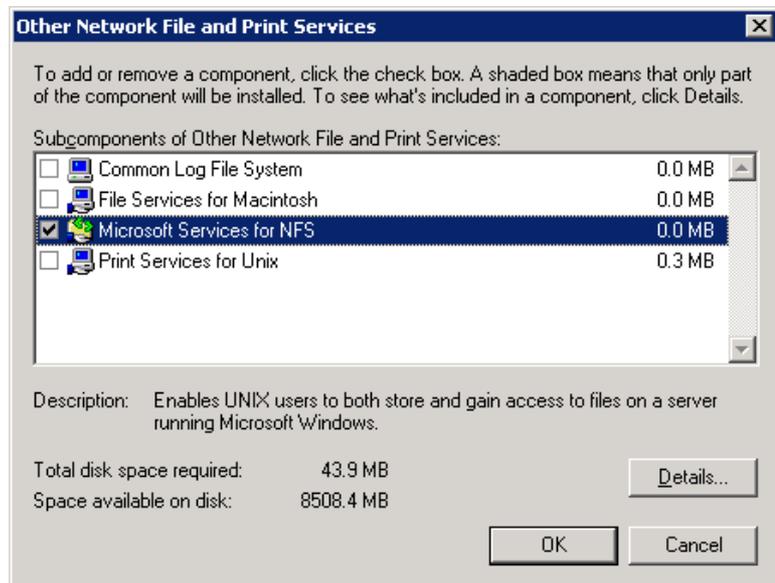
To install Services for NFS on the NetBackup clients with Windows 2003 R2 SP2

- 1 Click **Start > Control Panel > Add or Remove Programs**.
- 2 Click **Add/Remove Windows Components**.

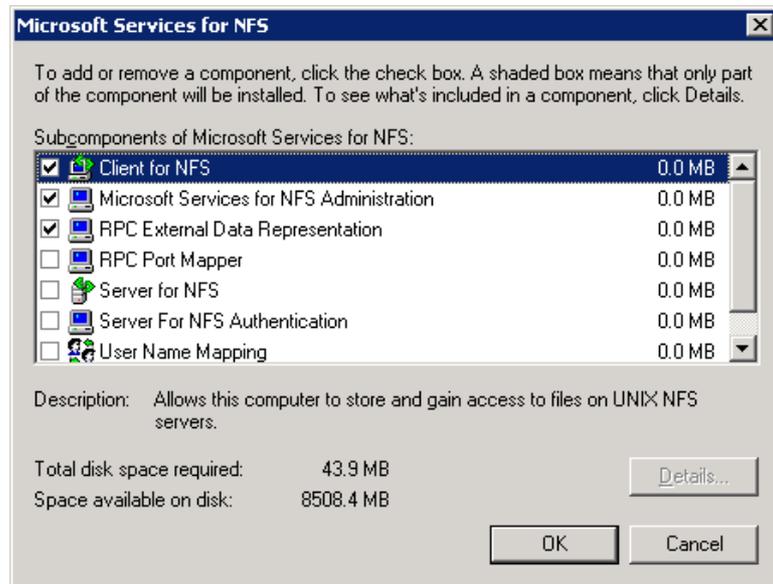
3 Check Other Network File and Print Services and click Details.



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 901.
See “Installing Services for NFS on Active Directory domain controllers or ADAM/LDS hosts with Windows 2003 R2 SP2” on page 912.
- You can configure a different network port for NBFSD.
See “Configuring a different network port for NBFSD” on page 915.

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 for Windows

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 Click **OK**.

7 Stop and start the NetBackup Client Service.

8 Close the Services control panel application.

Index

Symbols

- 259, 275. *See* moving a robot and its media
- .ds files 398
- .f files in catalog 664
- .SeCuRiTy.nnnn files 889

A

- Absolute pathname
 - to directory/volume storage unit setting 400
- Accept connections on non reserved ports
 - property 202
- Access Control
 - authorizing users 856
 - host properties 61–65
 - utility in Administration Console 43
 - within the NetBackup-Java Administration Console 850
- access control lists (ACLs) 611, 889–890
- Access Management utility 43
- ACS robot type 235
- ACS, LSM, Panel, Drive tape drive configuration
 - option 251
- ACSLs host robot configuration option 244
- Active Directory
 - granular recovery 637–639
 - host properties 66–67
 - restoring objects 639
- Active Directory ApplicationMode (ADAM) 637–639, 642
- Activity Monitor
 - bpdbjobs command 795
 - BPDBJOBS_OPTIONS environmental variable 793
 - canceling uncompleted jobs 769
 - copy job information 772
 - deleting completed jobs 769
 - detailed job status 769, 782
 - disabling job logging 157
 - filtering column heads 768
 - killing jobs 769
 - monitoring jobs 769
- Activity Monitor (*continued*)
 - restarting a completed job 770
 - resuming suspended jobs 770
 - Services tab 772
 - starting and stopping services 782
 - suspending a job 770
 - topology 765
 - using the Troubleshooter 42
 - window description 763
- ADAM (Active Directory ApplicationMode) 637–639, 642
- adjust time zone 868
- administering remote systems 841
- administration console 34
- Administration Console options 150
- administrator
 - email address property 133
 - nonroot 853
- AdvancedDisk disk storage units 393
- AdvancedDisk storage units
 - credentials for CIFS 395
- AFS policy type 515
- All log entries report 75, 821
- Allow backups to span tape media property 156
- Allow client browse property 79
- Allow client restore property 79
- Allow media overwrite property 155
- Allow multiple data streams
 - set policy attribute 542
 - when to use 543
- Allow multiple retentions per media property 156, 572
- Allow server file writes property 56, 202
- alternate media types 279
- Alternate read server for storage destinations 454–455
- Alternate restore failover machines host
 - properties 184
- Always property in Fibre Transport host
 - properties 124
- Announce DHCP interval property 163
- ANSI format 155

- AOS/VS format 155
- API robots 309, 345, 354
- application backups 551
- archive bit 98–99, 550, 555, 652
- archive jobs, keeping progress reports 93
- asterisk as wildcard 828
- atime 93, 556
- Audit Manager 802, 814
- audit service (nbaudit) 802–803, 808, 814
- auditing
 - configuration after upgrade 805, 807
 - enabling 805
 - overview 802
 - report 810, 812
 - viewing current settings 804
- auth.conf file
 - capabilities identifiers 855
 - description 851
 - entries for specific applications 853
 - overview 850
- Authentication Domain tab 61–62
- Authorization Service tab 61, 63
- Auto log off timeout option 38, 150
- auto-discovery streaming mode 627
- AUTO_CREATE_IMPORT_SLP 474
- Automated Cartridge System
 - drive information 251
- automatic
 - backups 551
 - failover to an alternate server 877
 - path correction 272
 - Volume Recognition service (avrd) 783
- automounted directories 529
- avrd (Automatic Volume Recognition process) 783

B

- Back up all log files 114
- Back up only uncommitted log files 114
- Backup destination type 458
- Backup end notify timeout property 201
- Backup Exec QIC media 758
- Backup Exec Tape Reader
 - Exchange Server support 758
 - host properties 68
 - limitations 756
 - SQL support 758
 - Windows 2003 support 757
 - Windows 2008 support 757
- Backup migrated files property 91
- Backup network drives policy attribute 529
- Backup option for log files during full backups
 - property 114
- Backup Policy Wizard 672
- Backup start notify timeout property 199
- Backup status report 75
- backups
 - activating policy 527
 - archive 551
 - automatic 550–551
 - Client backups report 821
 - creating copies 467, 565
 - deactivating policy 527
 - duplicating 743
 - expiring 750
 - frequency setting 560
 - full 550
 - how NetBackup determines when files are
 - due 555–556
 - importing 750
 - NetBackup database files 715
 - network drives 529
 - NFS mounted files 513, 528
 - off-site storage 571
 - raw partitions on Windows 523, 605
 - selections list
 - pathname rules 604–605, 607, 610, 617, 619
 - verifying 601
 - send email notification about 135, 137–139
 - Status of Backups report 821
 - types of 550
 - user directed
 - schedules 557
 - type of backup 551
 - verifying 740
 - weekly schedule 584
 - windows 579
 - specifying 580
- Bandwidth host properties 70–72
- bar codes 303, 344, 346, 349–350
- Bare Metal Restore (BMR) 101, 538, 665, 694, 767, 773
- Bare Metal Restore service 783
- basic disk staging
 - creating a storage unit 422
 - Final destination media owner 431
 - Final destination storage unit 431
 - Final destination volume pool 431

- basic disk staging (*continued*)
 - priority of duplication jobs 425
 - relocation schedule 415, 430, 549–550
 - storage units
 - checkpoint restart 522
 - size recommendations 426
 - Use alternate read server attribute 425, 432
 - BasicDisk storage units 393, 454
 - credentials for CIFS 395, 774–775, 783, 785–786
 - spanning within storage unit groups 442
 - BE-MTF1 format 155
 - BLAT mail utility 139
 - Block level incremental backups 523
 - BMRD (NetBackup Bare Metal Restore Master Server) 773, 783
 - BMRDB.db
 - configuration entry 701
 - in catalog 665
 - relocating 696, 722
 - transaction log 719–720
 - bp.conf file
 - auditing changes in 803, 808
 - configuring to use ctime 557
 - customizing jnbSA and jbpSA 863
 - NetBackup-Java Administration Console
 - configuration entries 856
 - BPARCHIVE_POLICY 558
 - BPARCHIVE_SCHED 558
 - bpbackup command 675
 - BPBACKUP_POLICY 558
 - BPBACKUP_SCHED 558
 - BPBERM logging property 145
 - bpcatarc command 682
 - bpcatlist command 682
 - bpcatres command 683
 - bpcatrm command 683
 - bpcd client process 395
 - BPCD connect-back property 85, 126
 - bpcd daemon 783
 - BPCD port setting on client 162
 - bpchangeprimary command 742
 - BPCOMPATD (NetBackup Compatibility Service) 774, 783
 - bpconfig command 628
 - bpdbjobs
 - adding registry key 793
 - command 795
 - debug log 795
 - BPDBJOBS_OPTIONS environmental variable 793
 - BPDBM
 - NetBackup Database Manager, description 774
 - BPDBM (NetBackup Database Manager) 783
 - BPDBM logging property 145
 - BPDM logging property 145
 - bpend 201
 - bpexdate command 747
 - bpgetconfig 56
 - BPINETD (NetBackup Client Service) 774
 - bpinetd client process 395, 783
 - bpjava-msvc service 783
 - bpjava-susvc service 783
 - bpjobd process 783
 - BPRD
 - logging property 145
 - NetBackup Request Daemon, description 775
 - port setting on client 163
 - process 784
 - bpsetconfig 56
 - bpstart 199
 - bpsynth log 654
 - bptestnetconn utility 178
 - BPTM logging level property 145
 - bpvault 145
 - Browse and restore ability property 82
 - buffer size 99
 - Busy action property 74
 - Busy file host properties 73–74
 - BUSY_FILE_NOTIFY_USER 73
- ## C
- cachefs file systems, excluding from backup 630
 - Calendar schedule type 559
 - calendar scheduling, using 584
 - canceling uncompleted jobs 769
 - capacity-based licenses 45–46
 - catalog archiving 527, 681
 - Catalog Backup Wizard 669
 - catalog backups
 - adding critical policies to 634
 - archiving 679–681, 683
 - compressing image catalog 689
 - determining success of 677
 - image files 664
 - manual backup 675
 - master server offline 83
 - Maximum concurrent jobs setting 404
 - moving client images 687
 - multiple file layout 664

- catalog backups (*continued*)
 - online, hot method 668
 - overview 661
 - parent and child jobs 767
 - policy type 631
 - retaining audit records 814
 - running concurrently with other backups 135
 - schedules for 676
 - single file layout 664
 - space required 684
 - strategies 677
 - uncompressing 691
- Catalog cleanup wait time property 76
- catalog indexing 688
- catalog recovery 668, 678
- Catalogbackup volume pool 278, 520
- cdrom file system, excluding from backup 630
- change journal 101
 - and synthetic backups 653
 - determining if enabling is useful 100
 - using in incremental backups 98
- Change server option 840
- changing to another server 840
- Check the capacity of disk storage units
 - property 129, 398
- checkpoint restart
 - and disk staging 422, 432
 - and synthetic backups 648
 - backup jobs 522
 - for restore jobs 524
 - Move job from incomplete state to done state
 - property 77
- CIFS
 - credentials for BasicDisk storage units 395, 774–775
- CIFS credentials for BasicDisk storage units 783, 786
- cipher types for NetBackup encryption 110
- Clean-up host properties 75, 524
- cleaning
 - drives 264, 792
 - frequency 250
 - tape, change cleanings allowed 293
- Cleaning Frequency tape drive configuration
 - option 250
- CLEANUP_SESSION_INTERVAL_HOURS 474
- Client administrator's email property 203
- Client Attributes host properties 78, 81, 84
- Client backups report 821
- Client cipher property 110
- Client connect timeout property 199
- Client name property 77
- Client port window property 168
- Client read timeout property 93, 200
- Client sends mail setting 203
- clients
 - adding a client to the client database 79
 - adding and removing clients 86
 - adding to a policy 595
 - BPCD port 162
 - BPRD port 163
 - choosing a policy type 514
 - deleting from policy 513
 - DHCP interval property 163
 - exclude and include lists 122
 - exclude file list 115–116, 119–120
 - exclude files list 630
 - maximum jobs 132
 - moving image catalog 687
 - name 77, 883
 - peername 882
 - setting host names 596
 - taking offline 81–82, 809
- clustering 59, 209, 661, 695, 702, 872
- Collect disaster recovery information for Bare Metal
 - Restore policy attribute 538
- Collect true image restore information (TIR) policy
 - attribute 540
- Collect true image restore information (TIR) with
 - move detection policy attribute 541
- Collect true image restore information (TIR) with
 - move detection property 76, 539, 652
- column heads, displaying
 - 768
- Communications buffer size property 99
- Compress catalog interval property 133, 689
- Compression policy attribute 536
- concurrent jobs
 - on client 132
 - per policy 525
- Consistency check before backup host property 196
- copies
 - creating using Catalog duplicating option 743
 - creating using storage lifecycle policies 467, 565
 - option in the Configure Multiple Copies dialog
 - box 564
 - primary copy 564
 - third party 563

- copy
 - NetBackup database files 715
 - policy backup selections, clients, schedules 511
 - primary 747
 - report text to another document 819
- Copy on write snapshots 607
- copy window 581
- cpio format 155
- create media ID generation rules 334
- Credential Access host properties 102
- credentials 213
 - about NDMP 213
- Critical Policies list 633–634, 671, 673, 675
- cross mount points
 - effect with UNIX raw partitions 534
 - examples 534
 - interaction with Follow NFS policy attribute 534
 - policy attribute 613
 - policy setting 533
 - separate policies for 533
- ctime 616
- cumulative incremental backups 550, 553
- curly brackets as wildcards 829
- D**
- Daemon connection port property 86, 127
- Daemon port only property 127
- daemons
 - tlmd 786
- DAS drive name tape drive configuration option 251
- DAS server robot configuration option 244
- data
 - deduplication 537–538
 - movers 394–395
- Data Classification setting 446
- Data Classifications
 - creating 104
 - in storage lifecycle policies 444, 447–448
 - policy attribute 517–518, 568
- Database Administration utility 698, 703
- database cache memory settings 717
- Database manager process (bpdbm) 783
- database schema, exporting 714
- database-extension clients, adding file paths for 619
- DataStore
 - policy type 515
 - volume pool 289, 520
- DataTools-SQL-BackTrack policy type 515
- datetime stamp 556
- Daylight savings time 868
- DB2 policy type 515
- DBA password, changing 712, 722
- DBR format 155
- debug logging levels 145
- decommissioning a media server 225
- deduplication disk pool, configuring 382
- Deduplication Option license key 659
- Deduplication property 82–83
- deduplication storage server
 - credentials for 213
 - defining target for remote master server
 - duplication 487
- Default cache device path for Snapshots property 94
- Default Job Priorities host properties 105, 526
- defragment NetBackup database files 707
- Delay on multiplexed restores property 128
- delete all devices for a media server 227
- deleting
 - a device host 229
 - drive 265
 - license keys 48
 - schedules 581
 - schedules, backup selections, or clients from a
 - policy 512
 - storage unit groups 437
 - storage units 390
 - volume pools 316
- Density storage unit setting 401
- denying requests 800
- destination types for storage lifecycle policies 458
- detailed job status 769, 782
- device
 - configuration wizard 257
 - discovery 237–238
 - file 246
 - mapping file 235
- Device Configuration Wizard 388
- device host
 - for move volume 310
 - removing 229
- Device host robot configuration option 243
- device management
 - remote 830
- Device Monitor
 - add drive comment 261
 - assigning requests 798
 - display pending requests 797
 - resubmit request 800

- devpts file system, excluding from backup 630
 - DHCP setting on client 163
 - differential incremental backups 551–552
 - Direct Access Recovery (DAR) 129
 - Directory can exist on the root file system or system
 - disk setting 400
 - directory junctions on UNIX 612
 - Disable client-side deduplication policy attribute 546
 - DISABLE_STANDALONE_DRIVE_EXTENSIONS 284
 - disaster recovery
 - information 133
 - sending e-mails 633
 - tab 631–632, 634
 - Disaster recovery file 671, 678
 - disk
 - array, credentials for 102
 - logs report 822
 - pool status report 822
 - pools 395
 - spanning 157, 442, 519
 - staging storage units
 - selection within a storage unit group 439
 - storage unit status report 822
 - Disk image backups 605
 - Disk pool storage unit setting 401
 - disk pools 382
 - disk staging 419
 - Disk type storage unit setting 401
 - disk-image backups
 - checkpoint restart 523
 - Distributed Application Restore Mapping
 - host properties 107
 - Do not compress files ending with property 94
 - down a device 261
 - drive
 - add comment 261
 - cleaning 250, 264, 267
 - name rules, configuring 252
 - running diagnostics 270
 - servicing requests 797
 - type 249
 - Drive is in a robotic library tape drive configuration
 - option 249
 - Drive name tape drive configuration option 248
 - drives
 - cleaning 792
 - initial state 248
 - initial status 264
 - monitoring 791
 - drives (*continued*)
 - name rules 251
 - replacing 272–273
 - updating firmware 274
 - duplicate backups
 - becoming a primary copy 746
 - creating 743
 - restoring from 741
 - duplicate window 581
 - Duplication destination type 458
 - Duplication job priority 445
 - Duplication Manager 466, 473, 776
 - See also* Storage Lifecycle Manager service (nbstserv)
 - duplication to remote master 451
 - AUTO_CREATE_IMPORT_SLP lifecycle parameter 474
 - DUPLICATION_GROUP_CRITERIA 475, 480
 - DUPLICATION_SESSION_INTERVAL_MINUTES 475, 481
 - duration of backup window
 - examples 582
 - dynamically-allocated ports 168
- ## E
- EFI System partitions 624
 - ejecting volumes 300
 - Email
 - address for administrator of this client 203
 - disaster recovery 633
 - notifications 133, 135–139
 - send from client 203
 - send from server 203
 - EMM database 667
 - containing audit records 802, 807, 815
 - removing a device host from 229
 - shared 193
 - empty media access port prior to update 336
 - Enable block sharing storage unit setting 402
 - Enable encryption property 109
 - Enable granular recovery policy attribute 547, 638
 - Enable job logging property 157
 - Enable multiplexing storage unit setting 402
 - Enable performance data collection property 203, 888
 - Enable robust logging property 144
 - Enable SCSI reserve property 156
 - Enable single instance backup for message attachments property 115

- Enable standalone drive extension property 157
 - Enable standard encryption property 110
 - Encryption host properties 109, 647
 - encryption method for SQL Anywhere 697, 701
 - Encryption policy attribute 538
 - Enterprise Disk license key 659
 - Enterprise Disk Options 395
 - Enterprise Media Manager (EMM) 191–193, 391, 665, 667, 693, 775, 832, 872–876
 - Enterprise Media Manager server 832–833
 - sharing 695
 - Enterprise Vault Hosts properties 112
 - Enterprise Vault properties 111
 - erasing media 295
 - error codes 83
 - See also* status codes
 - escape character
 - on UNIX 829
 - ESX server 183
 - Exchange granular restore proxy host property 115
 - Exchange Server
 - in IPv6-enabled environments 547
 - Exchange Server images, importing with BETR 758
 - exclude
 - dates from schedule 583
 - files and directories from backup 115–116, 119
 - list syntax 120
 - exclude file list 631
 - exclude files list 121–122, 630
 - Exclude list host properties 117–118
 - exclude_list 631
 - expiring backups 750
 - export
 - database schema and data 714
 - host properties 60
 - license keys 49
 - reports 820
 - extended attribute files
 - disabling the restore of 617
 - Solaris 9 612
- F**
- fail all copies when creating multiple copies 426, 567
 - failover
 - media server to alternate media server(s) 184
 - servers, adding or changing 185
 - storage unit selection in group 438
 - failover to an alternate server 877
 - Fibre Transport host properties 123
 - File browse timeout property 200
 - File Change Log (FCL) 93–94
 - file lists
 - extension clients 619
 - links on UNIX 609
 - NetWare clients
 - nontarget 617
 - target 619
 - NetWare NonTarget clients 617
 - NetWare Target clients 619
 - raw partitions 612
 - UNIX clients 610
 - UNIX files not backed up 611, 629
 - Windows clients 604
 - File system backup coverage report 602
 - File System Export option 539
 - files
 - .SeCuRiT.y.nnnn 889
 - catalog space requirements 684
 - excluding from backup 115–116, 119–120
 - linked, UNIX 612
 - NFS mounted 513, 528
 - No.restrictions 883
 - NOTES.INI 152
 - peername 884
 - redirected restores 885
 - restrictions on restores 882
 - FilesNotToBackup list 631
 - Final destination
 - media owner 431
 - storage unit 431
 - volume pool 431
 - Firewall host properties 124
 - FlashBackup 612–614
 - policy type 515
 - Windows policy type 515
 - Flexible Disk Option 393, 401–402, 443
 - Follow NFS mounts
 - cross mount points 529
 - raw partitions 529, 613
 - Follow NFS policy setting 534
 - Follow NFS setting policy attribute 528
 - FORCE_IPADDR_LOOKUP 858
 - Free browse property 82
 - Frequency schedule attribute 560
 - frozen media 299
 - full backups 550, 552, 649

G

General level logging property 97
 General server host properties 128, 130
 Global attributes
 host properties
 Schedule backup attempts 595
 Global attributes host properties 131, 133
 Global logging level property 144
 Go into effect at policy attribute 527
 granular recovery 547
 granular recovery of Active Directory objects 637
 Granular Recovery Technology (GRT) 129, 547
 Group Policy Objects 643

H

hard links
 NTFS volumes 609
 UNIX directories 609
 High water mark storage unit setting 402
 HKEYS, backing up 608
 host
 device 33
 properties
 changing in a clustered environment 59
 exporting 60
 permission to change 56
 host credentials. *See* credentials

I

IBM device number tape drive configuration
 option 251
 If this copy fails option 565
 image catalog file, compressing 133
 Image cleanup property 76
 IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS 475, 481
 images
 changing primary copy 741
 duplicating 743
 moving client catalog 687
 on disk report 822
 on media report 821
 restoring from duplicate 741
 verifying 740
 Import destination type 459
 Import Manager 473, 776
 See also Storage Lifecycle Manager service (nbstserv)

IMPORT_EXTENDED_RETRY_SESSION_TIMER 475
 IMPORT_SESSION_TIMER 476
 importing backups 750
 inactive media 823
 include
 files list 630
 list, on client 122
 include_list 631
 Incrementals based on
 archive bit property 99
 timestamp property 98
 Informix policy type 515
 INI file, for Lotus Notes 152
 Initial browse search limit property 202
 INITIAL_BROWSE_SEARCH_LIMIT 863
 INITIAL_MEMORY 860, 865
 inject volume into robot
 multiple volumes 336
 robot inventory 300
 Inline copy option 562, 744, 749
 installing and configuring Network File System (NFS) 899
 Instant Recovery
 Advanced Backup method 523
 Backups to disk only setting 562
 Internet Assigned Numbers Authority (IANA) 168, 862
 inventory and compare robot contents 325
 IP Address Family Support host property 166
 IP_ADDRESS_FAMILY 166
 IPv4
 addresses 175–176
 IP_ADDRESS_FAMILY entry 166
 networks, limiting bandwidth 70
 IPv6
 addresses 175–177
 and client names 77, 596
 and granular recovery 547
 IP_ADDRESS_FAMILY entry 166
 IPv6 networks, limiting bandwidth 197

J

Java
 auth.conf file 851
 authorizing users 850
 directory 853
 interface 34
 jbpSA configuration options 862
 performance improvement hints 865

- Java (*continued*)
 - Virtual Machine (JVM) 860
 - Java Administration Console 850
 - Java Windows Administration Console 34, 837, 840–841, 844, 848, 855–856
 - improving performance 863, 865
 - installing 843
 - jbpSA 862–863
 - jnbSA 34, 847, 862–863
 - Job Manager logging property 146
 - Job retry delay property 132
 - jobs
 - Concurrent per disk storage unit 404
 - maximum per client 132
 - maximum per policy 525
 - priority for policy 526
 - setting default priority 105
 - SLP_MultipleLifecycles 479
 - specifying filters 768
 - viewing in the Activity Monitor 766
 - JVM (Java Virtual Machine) 860
- K**
- Keep logs property 75
 - Keep status of user-directed backups 92–93, 100
 - Keep true image restoration information property 76
 - Keep vault logs property 76
 - KEEP_LOGS_DAYS 863
 - KeysNotToRestore list 631
 - Keyword phrase policy attribute 548
 - killing jobs 769
- L**
- labeling media 305
 - legacy logging 145
 - Library name robot configuration option 245
 - library sharing 240
 - license keys 45–49
 - LIFECYCLE_PARAMETERS file 472
 - Limit jobs per policy setting 525, 545–546, 559, 574
 - limiting bandwidth 71
 - links
 - UNIX hard-linked directories 609
 - UNIX symbolic 612
 - LiveUpdate 83
 - LMCP device file robot configuration option 245
 - load balancing methods 440
 - Locked file action property 93
 - logging
 - bpsynth 654
 - deleting logs after a set time 75
 - jbpSA 862
 - jnbSA 862
 - legacy 142
 - off of NetBackup automatically 150
 - unified 141
 - Logging host properties 141
 - Login Banner Configuration host properties 146
 - login banner text, removing 149
 - long erase 296
 - Lotus Notes host properties 150
 - Lotus Notes policy type 515
 - Low water mark storage unit setting 400
 - ltid (NetBackup Device Manager) 774, 784
- M**
- Mac OS X 517
 - mail notifications
 - administrator email address 203
 - Disaster Recovery attachment
 - sending 633
 - mail_dr_info.cmd 679
 - Mailbox for message level backup and restore property 115
 - manual backups
 - NetBackup catalogs 675
 - policy for 636
 - master servers
 - rebooting 897
 - sharing EMM database 193
 - switching to another 192
 - Match directive for Preferred Network host
 - properties 171, 176, 178, 180
 - MAX_GB_SIZE_PER_DUPLICATION_JOB 476
 - MAX_MEMORY 860, 865
 - MAX_MINUTES_TIL_FORCE_SMALL_DUPLICATION_JOB 477
 - maximum
 - concurrent FT connections property 124
 - concurrent jobs storage unit setting 404
 - concurrent write drives storage unit setting 563
 - data streams property 81, 546
 - error messages for server property 100
 - jobs per client 132, 545–546
 - jobs per policy 525
 - vault jobs property 133
 - Maximum backup copies property 133
 - maximum bar code lengths 345

- Maximum concurrent write drives setting 403
- Maximum number of logs to restore property 151
- Maximum streams per drive storage unit setting 406, 572
- media
 - active 823
 - ejection timeout period 302
 - formats 279
 - freeze 299
 - frozen 299
 - host override property 130
 - ID generation rules 348
 - ID prefix (non-robotic) property 158
 - inactive 823
 - log entries report 75, 821
 - mount timeout property 200
 - pools (see volume pools) 312
 - request delay property 158
 - server connect timeout property 201
 - server register 226
 - suspend 312
 - type when not an API robot 341
 - unfreeze 299
 - unmount delay property 158
 - unsuspend 312
- Media host properties 153
- media ID
- media ID, prefix for update robot 340
- media mount errors 796
- Media owner policy attribute 527
- media server
 - delete all devices from 227
- Media server copy advanced backup method 523
- Media server load balancing storage unit selection in group 438-440
- Media server storage unit setting 406
- media servers
 - activate or deactivate 215
 - adding a media server to the Alternate restore failover machine list 185
 - decommissioning 217-220, 225
 - moving a robot and its media 259
 - previewing references to 224
 - rebooting 897
 - registering with the EMM server 830
 - Restore failover host properties 184
- media sharing
 - about 317
 - configuring unrestricted 318
- media sharing (*continued*)
 - configuring with a server group 318
- media types 278
- Megabytes of memory property 93
- MEM_USE_WARNING 860
- Microsoft Cluster (MSCS) 614
- Microsoft Exchange policy attributes 548
- Microsoft Volume Shadow Copy Service (VSS) 67, 88
- Microsoft Windows Backup 631
- MIN_GB_SIZE_PER_DUPLICATION_JOB 476
- mirrored transaction log, creating 723
- mixing retention levels on tape volumes 572
- mklogdir.bat 142
- mntfs file system, excluding from backup 630
- monitoring
 - NetBackup drives 791
- monitoring NetBackup processes 787
- monthly backups, scheduling 586
- mount
 - points 533
 - requests, pending 797
- move
 - job from incomplete state to done state
 - property 77
 - restore job from incomplete state to done state 524
 - restore job from incomplete state to done state
 - property 77
 - volumes
 - update volume configuration 308
- Move backup job from incomplete state to done state
 - property 522
- moving NBDB database files 722
- moving NetBackup database files 713
- MS-Exchange-Server policy type 515
- MS-SharePoint policy type 515
- MS-SQL-Server policy type 515
- MS-Windows policy type 516
- MTF format 155
- mtime 616
- multihomed server example 172
- multiple copies
 - checkpoint restart 523
 - creating using a policy schedule 565
 - creating using storage lifecycle policies 467, 564
 - criteria for creating 563
 - dialog box 564
 - fail all copies 426, 567
 - parent and child jobs 767

- multiple copies (*continued*)
 - setting 562
 - multiple copy synthetic backups method 655–658
 - multiple data streams 543, 768
 - multiple file layout for NetBackup catalogs 664
 - multiple installations 151
 - multiplexing (MPX)
 - and synthetic backups 647
 - demultiplexing 578
 - Maximum jobs per client property 575
 - preserving 455
 - set for schedule 572
 - use with Enable block sharing 402
 - multistreamed backups 589
 - multistreaming and synthetic backups 647
 - Must use local drive property 129
- N**
- named data streams
 - disabling the restore of 617
 - naming conventions 827
 - nb_updatedssu script 398
 - NBAC (NetBackup Access Control (NBAC)) 803
 - See also* NetBackup Access Control (NBAC)
 - nbatd (NetBackup Product Authentication) 773
 - nbaudit (NetBackup Audit service) 802–803, 808, 811, 814
 - nbaudit log 814
 - nbauditreport 807, 810
 - nbazd (NetBackup Product Authorization) 773
 - NBAZDB 695, 701
 - See also* NetBackup Authorization database
 - NBDB.db
 - configuration entry 701
 - creating manually 724
 - in catalog 665
 - installation overview 694
 - moving from one host to another 731
 - relocating 695, 722
 - transaction log 719–720
 - NbDbAdmin.exe (Database Administration utility) 703
 - nbdecommission command 225
 - NBEMM (NetBackup Enterprise Media Manager) 775, 784
 - nbemmcmd command 194
 - nbEvtMgr process 784
 - nbfsd port 915
 - nbftsvr process 784
 - nbj.conf 856
 - NBJAVA_CLIENT_PORT_WINDOW 860
 - NBJAVA_CORBA_DEFAULT_TIMEOUT 861
 - NBJAVA_CORBA_LONG_TIMEOUT 861
 - NBJM (NetBackup Job Manager) 146, 784
 - nbmail.cmd script 133, 679
 - NBPEM (NetBackup Policy Execution Manager) 146, 775, 784
 - nbproxy process 784
 - NBRB (NetBackup Resource Broker) 146, 775
 - nbrb process 785
 - nbrbutil configuration utility 779
 - NBRMMS (NetBackup Remote Management and Monitor Service) 396, 775, 785
 - NBSL (NetBackup Service Layer) 776
 - nbsl process 785
 - nbstlutil (lifecycle utility) command 480
 - nbstserv process 785
 - nbsvcmon process 785
 - NBU-Catalog policy type 516, 520, 631
 - NBVAULT (NetBackup Vault Manager) 776, 785
 - NCR-Teradata policy type 516
 - NDMP
 - credentials for 42, 102, 213
 - Direct Access Recovery for restores 129
 - drives 129
 - global credentials 160
 - host storage unit setting 409
 - hosts 160, 237
 - storage units 398, 416, 563
 - NDMP host name robot configuration option 245
 - NDMP policy type 516
 - NearStore storage units 381, 393, 409, 437, 457, 539
 - checkpoint restart 523
 - NetApp 393
 - NetBackup
 - administration
 - console 34
 - client service 162
 - request service port (BPRD) 163
 - NetBackup Access Control (NBAC) 61, 695, 700–701, 803, 850
 - NetBackup Audit Manager 802, 814
 - NetBackup Authorization database 693, 695, 700–701
 - NetBackup Client Service (BPINETD) 774
 - NetBackup Client Service log on account, configuring 916
 - NetBackup Compatibility Service (BPCOMPATD) 774

- NetBackup database files
 - adding space 707
 - backing up 715
 - changing DBA password 712
 - defragmenting space 706
 - exporting database schema 714
 - free and used space 706
 - memory cache settings 717
 - moving 713
 - rebuilding 709
 - restoring 716
 - validating 708
 - NetBackup Database Manager (BPDBM) 774
 - NetBackup Deduplication Engine 213
 - NetBackup Device Manager 275, 774
 - NetBackup for MS-Exchange 619
 - NetBackup Job Manager (NBJM) 146, 775
 - NetBackup Key Management Service (NBKMS) 775
 - NetBackup Legacy Network Service (vnetd) 786
 - NetBackup media kit 33
 - NetBackup Monitor Service 776
 - NetBackup Policy Execution Manager (NBPEM) 146, 775
 - NetBackup Remote Administration Console 840, 842
 - NetBackup Remote Management and Monitor Service (NBRMMS) 775
 - NetBackup Request Daemon (BPRD) 775
 - NetBackup Request Service Port (BPRD) property 163
 - NetBackup Resource Broker (NBRB) 106, 146, 775
 - NetBackup Resource Broker (nbrb) 779
 - NetBackup Service Layer (NBSL) 776, 784
 - NetBackup Storage Lifecycle Manager 472, 776
 - NetBackup support Web site 236
 - NetBackup Vault Manager (NBVAULT) 776
 - NetBackup Volume Manager (VMD) 776
 - NetBackup volume pool 520
 - NetBackup-Java Administration Console 850
 - improving performance 863
 - NetBackup-Java Version 7.1 840–841
 - NetWare client
 - target and nontarget 116
 - NetWare client host properties 91
 - NetWare clients support for checkpoint restart 523
 - NetWare policy type 516
 - network
 - addresses, prohibiting 180
 - drives, backing up 529
 - Network Attached Storage (NAS) 394, 399
 - Network Attributes tab 61, 64–65
 - Network File System (NFS), described 899
 - Network host properties 162
 - Network Settings host properties 163, 165–166
 - Never property in Fibre Transport host
 - properties 124
 - NEW_STREAM
 - file list directive 626
 - NFS (Network File System)
 - Follow NFS policy attribute 528, 534
 - NFS access timeout property 205
 - no disk spanning 157
 - Nirvanix cloud storage 394
 - non reserved ports 202
 - None volume pool 520
 - nonroot administration for specific applications 853
- ## O
- ODBC, remote 802
 - offline
 - master server and catalog backups 83
 - taking clients 81–82, 809
 - On demand only storage unit setting 409, 442
 - Only directive for Preferred Network host
 - properties 171, 178, 181–182
 - open schedules 590
 - OpenStorage
 - storage server. *See* NetBackup Shared Storage Guide
 - OpenStorage Disk Option 393–394, 401, 443
 - OpenStorage disk storage units 394, 437
 - OpenStorage optimized synthetic backup method 659
 - operating mode of tape drive
 - changing 262
 - Operator's email address property 73
 - OpsCenter 776, 802, 804–805, 807, 810, 815
 - optical devices, support in NetBackup 7.0 234
 - Oracle policy type 516–517
 - Oracle_RMAN 600
 - Oracle_XML_Export 600
 - OS/2 policy type 516
 - Override default job priority
 - for Catalog jobs 106, 740
 - for Media Contents report 106
 - for Media contents report 823
 - for queued or active jobs 772
 - Override policy
 - storage selection setting 568
 - volume pool setting 568
 - Overwrite existing files property 616

P

- pagefile.sys 607
- parent jobs 542, 767
 - in Activity Monitor Jobs tab 767
 - Limit jobs per policy setting 526
 - parent_end_notify script 767
 - parent_start_notify script 767
- parent_end_notify script 767
- parent_start_notify script 767
- password, changing 712, 722
- path
 - correction, enabling automatic 272
 - separators 400
- PBX (Symantec Private Branch Exchange) 787
- PBX_PORT 862
- PC NetLink files 611
- peername
 - files 884
 - of client 882
- pending actions
 - resolving 799
- pending requests
 - resolving 798
 - resubmitting 800
- Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS) property 115
- Perform default search for restore property 100
- Perform incrementals based on archive bit 555
- permissions
 - to change NetBackup properties 56
- physical inventory utility 361
- policies
 - activating 527
 - changing multiple policies at one time 510
 - changing properties 508–509, 511, 513, 594–595
 - creating 508–509
 - creating policy for Vault 635
 - deactivating 527
 - for Active Directory granular restores 638
 - media owner attribute 527
 - overview 502
 - planning 504
 - setting priority 105, 526
 - types 514
 - user schedules 557
 - utility, using 502
 - volume pool policy setting 519, 521
- Policy Execution Manager
 - Logging property 146
- Policy storage policy attribute 518, 674
- policy type
 - Vault Catalog Backup 552
- Policy update interval property 132, 588, 739
- Port Ranges host properties 167–168
- Port, Bus, Target, LUN configuration option 246
- ports
 - allow operating system to select non reserved port 168
 - dynamically-allocated 168
 - non reserved 202
- power down NetBackup servers 895
- Preferred Network host properties 169–170, 172, 177–178
- Preferred property in Fibre Transport host properties 123
- prelabel media 305
- preprocess interval 628
- Preserve multiplexing 455
- preview volume configuration update 333
- previewing a media server's references 224
- primary copy
 - becoming a 746
 - changing 741
 - definition 747
 - promoting to 742
 - setting in the Configure Multiple Copies dialog box 564
- print
 - job detail information 771
 - job list information 771
 - license key 48
 - reports 820
- Prioritized storage unit selection in group 438
- priority
 - of a job 105, 526
 - of duplication jobs 425
 - of relocation jobs started from this schedule setting 430
- Priority of duplication job option 564
- Private Branch Exchange 776, 787
- Problems report 75, 821
- proc file system
 - excluding from backups 630
- processes
 - monitoring 787

Prohibited directive for Preferred Network host properties 171, 176, 178, 180

properties

- changing on multiple hosts 58
- exporting 60
- viewing 57

PureDisk

- PureDisk-Export policy type 516
- Storage Option 394, 402
- Storage Pool Authority (SPA) 401
- storage units 457

Q

question mark as wildcard 828
 quick erase 296
 quotas on file systems 395

R

random ports, setting on server 167
 raw partitions

- backing up 523, 550, 605
- backups on UNIX 612, 614
- Follow NFS policy attribute 529
- restoring 606

rebooting NetBackup servers 897
 recommended method of configuring devices 237
 redirected restores 614, 882
 Reduce fragment size storage unit setting 413
 register a media server 226
 registered ports 168
 registry

- auditing changes in 808
- backup/restore 608

reload NetBackup database 709
 reload.sql 730–731
 relocation schedule 424, 431, 549–550, 560

- initiating manually 432

remote

- access, allowing 836–837
- device management 830
- systems
- administering 841

Remote Administration Console 34, 840, 842
 Remote master storage destination 453
 Remote ODBC 802
 removing a device host 229
 REORGANIZE command to defragment NetBackup database 707

replacing a drive 272–273

REPLICA_METADATA_CLEANUP_TIMER 477
 reports

- All log entries report 821
- Client backups report 821
- copying to another document 819
- description of utility in Administration Console 818
- Disk logs report 822
- Disk pool status report 822
- Disk storage unit status report 822
- Images on Disk report 822
- Images on media report 821
- Media log entries report 821
- printing 820
- Problems report 821
- running a report 819
- saving 820
- Status of backups report 821
- Tape contents report 106, 823
- Tape lists report 823
- Tape logs report 822
- Tape summary report 823
- Tape written report 823
- using the Troubleshooter 42

requests

- assigning 798
- denying 800

reset

- file access time property 93
- mount time 267

residence, updating volume configuration 324

Resource Broker (nrb) 779

Resource Broker logging property 146

Resource Limit host properties 182

restarting jobs 770

Restore Failover host properties 184

Restore jobs

- move restore job from incomplete state to done state 524

Restore retries

- property 202

restore retries

- checkpoint restart 524

restores

- adjust time zone for 868
- alternate server 871
- directed from the server 881
- from a specific backup copy 425, 566, 871

- restores *(continued)*
 - keeping progress reports 93
 - NetBackup database files 716
 - raw partition 606
 - redirected 184, 882, 884
 - reducing search time 688
 - registry on Windows clients 608
 - server independent 871
 - symbolic links on UNIX 612
 - System State 891
 - using a specific server 130
 - resuming suspended jobs 770
 - retention levels
 - for archiving catalogs 681
 - retention periods
 - caution for setting 559
 - changing 187
 - expiration 559
 - guidelines for setting 570
 - lifecycle and policy-based 451
 - mixing on tape volumes 156, 572
 - precautions for setting 571
 - redefining 186
 - setting 569
 - user schedule 559
 - volumes 188
 - Retention types for storage lifecycle policies
 - Staged capacity managed 455
 - retire a media server. *See* decommissioning a media server
 - retiring a media server 217
 - Retries allowed after runday policy setting 560
 - Retry count property 74
 - retry restores, setting 202
 - Reverse Host Name Lookup host property 164–165
 - REVERSE_NAME_LOOKUP entry 165
 - robot
 - about configuring 237
 - adding 241
 - compare contents 325
 - destination for move volume 310
 - device file 246
 - device host configuration option 243
 - inventory 322, 325
 - moving to new media server 259
 - number storage unit setting 414
 - robot number configuration option 243
 - robot type configuration option 243
 - running diagnostics 268
 - robot *(continued)*
 - type storage unit setting 415
 - robot configuration
 - about 237
 - changing 258
 - device host option 243
 - robot number option 243
 - robot type option 243
 - Robot control host robot configuration option 245
 - Robot control is attached to an NDMP host robot configuration option 244
 - Robot control is handled by a remote host robot configuration option 244
 - Robot control options 243
 - Robot device path robot configuration option 246
 - Robot device robot configuration option 245
 - Robot drive number tape drive configuration option 250
 - Robot is controlled locally by this device host robot configuration option 244
 - Robot number robot configuration option 243
 - Robot type robot configuration option 243
 - robot types 234
 - Robotic device file robot configuration option 246
 - Robotic library tape drive configuration option 250
 - Round robin storage unit selection in group 438
 - RS-MTF1 format 155
- ## S
- SAP policy type 516
 - Sarbanes–Oxley Act (SOX) 802
 - save a report 820
 - Schedule backup attempts property 132, 522, 543, 594–595
 - schedules
 - adding to a policy 509
 - backups on specific dates 584
 - changing a time window 581
 - copying a time window 581
 - creating a time window 580
 - creating time windows on successive days 581
 - deleting a time window 581
 - determining due time 588
 - duplicating a time window 581
 - excluding dates 583
 - frequency setting 560
 - how NetBackup determines which schedule to run 587
 - monthly backups 586

- schedules *(continued)*
 - moving a time window 581
 - naming 549
 - overview 549
 - priority 561
 - recalculating 587
 - retention periods
 - guidelines 570
 - setting 569
 - specify multiplexing 572
 - Start Windows tab 579
 - storage unit/storage lifecycle policy 568
 - type of backup 550
 - user backup or archive 557
 - volume pool 568
 - windows that span midnight 589, 594–595
- schedules, creating weekly backups 584
- scratch
 - pool and WORM media 282
 - pool, adding 313, 315
 - volume pool 520
- scripts 767
 - bpdbjobs example 794
- SCSI
 - Long Erase 296
 - pass-through command 238
 - persistent reserve
 - drive path override 255
 - Quick Erase 296
 - reserve, configuring 156
 - reserve/release
 - drive path override 255
- SeCuRiT_y.nnnn files 889
- Serial Number tape drive configuration option 250
- SERVER
 - vm.conf entry 831
- server
 - directed restores 855
 - adding to Additional servers list 190
 - adding to Media servers list 191
 - allowing access 836–837
 - alternate server restores 871
 - directed restore 881
 - EMM server 667
 - host properties 189
 - media servers 190
 - using 836–837
 - independent restores 184, 871
 - list definition 189
 - (continued)*
 - list, adding a server 836–837
 - power down 895
 - rebooting 895
 - remove from Additional servers list 192
 - removing from Media list 192
 - sends mail property 203
 - server group
 - configuring 210
 - deleting 212
 - Service Manager 702
 - services
 - description of those for NetBackup 773
 - starting and stopping 782
 - tab in the Activity Monitor 772
 - types of 776
 - Services for NFS
 - installing on Windows 2003 R2 SP2 912
 - Services tab in the Activity Monitor 772
 - setconf.bat file 856
 - Shadow Copy Components 757
 - Shadow Copy Components directive 624
 - Shadow Copy Service 67, 88
 - shared drives
 - configuration wizards 237
 - drive operating mode 263, 265
 - shared tape drives
 - operating mode 262
 - SharedDisk
 - properties 195
 - SharedDisk storage units 382
 - SharePoint 2003 768
 - SharePoint policy type 515
 - SharePoint Server
 - consistency checks options 197
 - in IPv6-enabled environments 547
 - properties 195
 - show robot contents 325
 - shut down NetBackup services 896
 - single file
 - layout for NetBackup catalogs 664
 - restore program
 - FlashBackup 612
 - Single-Instance Storage (SIS) 115, 457, 537–538
 - checkpoint restart 523
 - slot number
 - for move volumes 309
 - for volume 286
 - SLP_MultipleLifecycles job 479

- Snapshot Client 88, 205, 394, 402, 515, 562, 619, 767
 - checkpoint restart 523
 - policy attributes 548
- Snapshot destination type 459
- Snapshot verification I/O throttle property 114
- SnapVault storage units 394, 409, 416, 437, 454
- Solaris 9 extended attributes 612
- source binding 170, 172, 176, 178, 180–182
- SPC-2 SCSI reserve 156
- SQL Anywhere
 - encryption method 697, 701
 - in NetBackup installation 665
- SQL images, importing with BETR 758
- SQL-Server policy type 515
- SQLANYs_VERITAS_NB 702, 775
- square brackets as wildcards 829
- Staged capacity managed retention type 455
- staging
 - backups 419
 - schedule storage unit setting 415
 - using BasicDisk storage unit 400
 - using Storage Lifecycle Policies 443
- Standard policy type 517
- start up NetBackup services 896
- Start Window tab 579
- startup text, removing 149
- status codes
 - NetBackup
 - 1000 83
 - 71 598
- Status of backups report 821
- Storage device storage unit setting 415
- Storage Lifecycle Manager service (nbstserv) 472, 479, 776
- Storage Lifecycle Policies
 - Alternate read server for destination 454–455
 - and the Multiple copies configuration dialog 568
 - copy number 467
 - Data classification setting 446
 - data classifications 447–448
 - deleting 448
 - duplicating to a remote master 451
 - Duplication job priority setting 445
 - hierarchy 460, 462–463, 465
 - Import destination 451
 - Media owner for destination 454–455
 - Preserve multiplexing for destination 455
 - Remote master option 453
 - retention type 451, 569
- Storage Lifecycle Policies (*continued*)
 - retention types 453
 - Staged capacity managed retention type 455
 - storage destination list requirements 452
 - storage destinations 450, 452
 - Storage lifecycle policy name 445
 - Storage unit for destination 454
 - using nbstlutil to administrate lifecycle operations 480
 - utility 443
 - versions of 469, 471–472
 - volume pool for destination 454
 - writing multiple copies 467
- storage lifecycle policies 472, 776
 - See also* Storage Lifecycle Manager service (nbstserv)
 - optional LIFECYCLE_PARAMETERS configuration 472
- storage server
 - AdvancedDisk. *See* NetBackup Shared Storage Guide
 - credentials for deduplication 213
 - define target for remote master server duplication 487
 - OpenStorage. *See* NetBackup Shared Storage Guide
- storage servers 395
- storage unit
 - groups 435, 437–438
 - name setting 415
 - selection within a storage unit group 437–438, 442
 - type setting 415
 - utility 386
- storage units
 - AdvancedDisk disk type 393
 - available storage property of volume 411
 - BasicDisk type 393
 - capacity property of volume 411
 - changing server to manage 835
 - creating 388–389
 - creating a basic disk staging unit 422
 - creation overview 386
 - deleting 390
 - disk pool comment property 411
 - disk storage units 392
 - for policy 518
 - for schedule 568
 - high water mark property of volume 412

- storage units (*continued*)
 - low water mark property of volume 412
 - Media Manager type 390
 - name property 412
 - naming conventions 827
 - NDMP disk type 398
 - NearStore disk type 393, 409, 437
 - number of volumes property 412
 - OpenStorage disk type 394, 437
 - percent full property on volume 412
 - PureDisk disk type 394, 437
 - QIC drive type 563
 - raw size property on volume 412
 - SnapVault disk type 394, 409, 437
 - storage lifecycle policies 396
 - usable size property of volume 412
 - vendor-specific 394
- subnets 71
- Sun PC NetLink 611
- suspended jobs 77, 770
- Sybase policy type 517
- Sybase SQL Anywhere
 - dbsrv11.exe 775
 - default password 701
 - management of 702
 - starting/stopping the service 702
 - use in NetBackup 693
- Symantec OpsCenter 698, 700, 776, 804–805, 810
- Symantec Private Branch Exchange 776, 787
- Symantec products properties 197
- Symantec support Web site 236
- symbolic links
 - included in backup selection list 601
 - UNIX 612
- synthetic backups
 - and encryption 647
 - checkpoint restart 523
 - component images 649
 - deduplication 659
 - full 648
 - logs produced during 654
 - multiple copy backups method 655
 - no multiple copy support 563
 - no NetBackup change journal support 101
 - OpenStorage optimized method 659
 - recommendations for using 646
 - schedules 559, 589
- System State
 - directive 623
- System State (*continued*)
 - restoring 891
- System State backups
 - checkpoint restart 523
- T**
- Take checkpoints every __ minutes (policy attribute) 522
- tape
 - assigning requests 798
 - contents report 823
 - lists report 823
 - logs report 822
 - Media contents report 823
 - summary report 823
 - written report 823
- tape drive
 - ACS, LSM, Panel, Drive configuration option 251
 - changing operating mode 262
 - Cleaning Frequency configuration option 250
 - configuration option 248–249
 - configuration options 248–249
 - DAS drive name configuration option 251
 - Drive is in a robotic library configuration option 249
 - Drive name configuration option 248
 - IBM device number configuration option 251
 - Robot drive number configuration option 250
 - Robotic library configuration option 250
 - Serial Number configuration option 250
 - TapeAlert 250
- tape drive configuration
 - about 237
 - ACS, LSM, Panel, Drive option 251
 - Cleaning Frequency option 250
 - DAS drive name option 251
 - Drive is in a robotic library option 249
 - Drive name option 248
 - IBM device number option 251
 - Robot drive number option 250
 - Robotic library option 250
 - Serial Number option 250
 - tape drive configuration option 248–249
 - tape drive configuration options 248–249
- tape drives
 - about configuring 237
 - adding 246, 248
 - adding a path 254
 - configuring by using the wizard 241

- TAPE_RESOURCE_MULTIPLIER 478
 - TapeAlert 250
 - tar format 155
 - TCDebug_TCPP level logging property 97
 - temporary staging area 403, 414, 416, 424
 - third-party copies 563
 - Third-Party Copy Device Advanced Backup method 523
 - Throttle Bandwidth host properties 197–198
 - THROTTLE_BANDWIDTH 197
 - Time overlap property 99
 - time zones
 - adjustment for restores 868
 - setting Daylight savings time 868
 - Timeouts host properties 199
 - tlmd daemon 786
 - tmpfs file system, excluding from backup 630
 - tpext utility 725
 - transaction log
 - setting full or partial mode 719
 - truncating 720
 - Transaction log cache path property 152
 - transaction log, creating 723
 - Transfer throttle storage unit setting 416
 - traversing directories to back up a file 122
 - Troubleshooter 42
 - True Image Restoration (TIR)
 - Error code 136 653
 - pruning information 653
 - with Move Detection 653
 - True Image Restore (TIR) with Move Detection 101
 - Truncate log after successful Instant Recovery backup property 115
 - truncating the NetBackup transaction log 720
- ## U
- UNC path
 - checkpoint restart 523
 - with CIFS and AdvancedDisk storage units 396
 - with CIFS and BasicDisk storage units 396
 - uncompress
 - NetBackup catalogs 691
 - Uncompress files before backing up property 91
 - unified logging 141, 145
 - Universal Settings host properties 201
 - UNIX Client host properties 204
 - UNIX Client Settings host properties 92, 94
 - UNIX clients
 - checkpoint restart 523
 - UNIX server properties 205
 - unload NetBackup database 709
 - unloading
 - the NetBackup database 730
 - unloading database schema 714
 - UNSET file list directive 629
 - UNSET_ALL file list directive 629
 - unsupported characters 332
 - update
 - robot procedure 332
 - volume configuration 323, 325
 - updating drive firmware 274
 - upgrading and the auditing configuration 805
 - usbdevfs file system, excluding from backup 630
 - Use alternate read server attribute 425, 432
 - Use case sensitive exclude list host property 117
 - Use change journal in incrementals property 98
 - Use defaults from the master server configuration property 124
 - Use Direct Access Recovery for NDMP restores property 129
 - Use legacy DES encryption property 110
 - Use non reserved ports property 127
 - Use OS dependent timeouts property 200
 - Use random port assignments properties 167
 - Use reserved ports property 127
 - Use VxFS file change log for Incremental backups property 93
 - user
 - archive backups 551
 - backups 551
 - schedules, planning 557
 - User directed timeouts property 99
- ## V
- validate NetBackup database 708
 - Vault
 - backup type 552
 - catalog archiving 681
 - designating duplicate as the primary 741
 - Logging property 145
 - Maximum vault host property 133
 - parent and child jobs 768
 - policy
 - creating 635
 - vlteject command 635
 - vltrun command 635
 - Vault policy type 517
 - vCenter server 183

- vendor-specific storage units 394
- verifying backup
 - images 740
 - selections list 601
- Veritas Volume Manager (VxVM) 613
- Veritas Volume Snapshot Provider 88, 206
- veritas_pbx (Symantec Private Branch Exchange) 787
- veritas_pbx port 86, 127
- VERSION_CLEANUP_DELAY_HOURS 478
- view properties of a license key 49
- vlteject Vault command 635
- vltrun Vault command 635
- vm.conf file, adding SERVER entries 831
- VMD (NetBackup Volume Manager) 776
- vmd process 786
- vmphyinv physical inventory utility 361
- VMware backup hosts host properties 205
- VMware cluster 183
- VMX datastore 183
- vnetd
 - enabling logging for 127
 - NetBackup Legacy Network Service 786
 - Only property (for selection of ports) 127
 - Veritas Network Daemon 126
 - with CIFS BasicDisk storage units 783, 785–786
- VNETD_PORT 862
- volume expiration date, changing 293
- volume groups
 - about 316
 - changing name 290–291, 299
 - for move volume 310
 - rules for assigning 316
- volume pools
 - about 312
 - adding 314
 - and WORM media 281
 - changing attributes 315
 - changing for a volume 291
 - DataStore pool 293, 520
 - deleting 315–316
 - for schedule 568
 - indicating one for use by a policy 519
 - overview 312
 - properties 314
 - scratch 313
- Volume Shadow Copy Service (VSS) 67, 88, 623–624
- Volume Snapshot Provider (VSP) 88, 206
- volumes
 - adding 283, 285

- volumes (*continued*)
 - assignments 520
 - changing properties 291
 - cleaning count 293
 - description for new volume 292
 - determining retention period 188
 - ejecting 300
 - exchanging 297–298, 311
 - maximum mounts allowed 287
 - moving 290, 306, 308
 - properties 286
- VRTSpbx (Symantec Private Branch Exchange) 776
- VXDBMS_NB_DATA registry entry 701
- VxFS
 - file change log 93
 - named data streams 615
- vxlogcfg command 141, 145
- vxlogmgr command 141

W

- Wait time before clearing archive bit property 98–99
- weekly backups scheduling 584
- wildcard characters 604, 611, 828
- windows
 - creating schedules on successive days 581
 - Windows Client host properties 207
 - Windows Client Settings host properties 96, 100–101
 - Windows Disk-Image (raw) backups 523, 605
 - Windows Display Console 35
 - Windows Open File Backup host properties 86
 - Windows policy type 516
 - Windows Service Manager 702
- wizards
 - backup policy 508
 - Device Configuration 257, 388
- Working directory property 73
- WORM media 280–283