

Symantec NetBackup™ Administrator's Guide, Volume I

UNIX, Windows, and Linux

Release 7.7.1



Symantec NetBackup™ Administrator's Guide, Volume I

Documentation version: 7.7.1

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, NetBackup, Veritas, and the Veritas 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.

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 website at the following URL:

www.support.symantec.com

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, DVDs, 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_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Contents

Technical Support	3
Section 1 About NetBackup	31
Chapter 1 Introducing NetBackup interfaces	32
About NetBackup	32
NetBackup documentation	34
About NetBackup administration interfaces	34
About setting up the NetBackup Administration Console on	
UNIX	36
About using the NetBackup Administration Console	37
Administering remote servers of different versions	39
Using the NetBackup Administration Console	41
About security certificates for NetBackup media servers	45
Standard and user toolbars	45
About customizing the NetBackup Administration Console	46
Configuring log files for the NetBackup Administration	
Console	46
NetBackup configuration wizards	49
Running the Troubleshooter	50
Chapter 2 Administering NetBackup licenses	52
About administering NetBackup licenses	52
Section 2 Configuring hosts	57
Chapter 3 Configuring Host Properties	58
About the NetBackup Host Properties	60
Methods to set the NetBackup configuration options	62
Connecting to a host to view the host properties	62
Changing the host properties on multiple hosts at the same	
time	63
Property states for multiple hosts	64

Exporting host properties	65
Standard host property dialog box options	66
Access Control properties	67
Authentication Domain tab of the Access Control properties	68
Authorization Service tab of the Access Control properties	69
Network Attributes tab of the Access Control properties	70
Active Directory host properties	73
Bandwidth properties	75
Bandwidth limit usage considerations and restrictions	76
Add Bandwidth Settings dialog box for Bandwidth properties	77
Busy File Settings properties	77
Activating the Busy File Settings in host properties	79
Clean-up properties	80
Client Name properties	82
Client Attributes properties	84
Add Client dialog box	86
General tab of the Client Attributes properties	86
Connect Options tab of the Client Attributes properties	91
Windows Open File Backup tab of the Client Attributes properties	93
Back-level and upgraded clients that use Windows Open File Backup	98
Client Settings properties for UNIX clients	99
VxFS file change log for incremental backups property	102
Client Settings properties for Windows clients	104
How to determine if change journal support is useful in your NetBackup environment	107
Guidelines for enabling NetBackup change journal support	108
Credential Access properties	108
Data Classification properties	110
Creating a Data Classification	112
Default Job Priorities properties	113
Understanding the Job Priority setting	115
Distributed application restore mapping properties	117
Encryption properties	118
Additional encryption methods for Windows clients	121
Enterprise Vault properties	122
Enterprise Vault Hosts properties	123
Exchange properties	124
About the Exchange credentials in the client host properties	127
Exclude Lists properties	128
About the Add to Exclude List and Add Exceptions to Exclude List dialog boxes	131

Syntax rules for exclude lists	133
About creating an include list on a UNIX client	136
Configuring the exclude list on a Windows client from a UNIX master server	136
Traversing excluded directories	138
Fibre Transport properties	139
About Linux concurrent FT connections	142
Firewall properties	142
Enabling logging for vneta	146
General Server properties	146
Forcing restores to use a specific server	149
Global Attributes properties	150
About constraints on the number of concurrent jobs	153
Setting up email notifications about backups	155
Configuring the nbmail.cmd script on Windows	156
Sending messages to the global administrator about unsuccessful backups only	158
Sending messages to the administrator about successful and unsuccessful backups	158
Installing and testing the email utility on Windows	160
Logging properties	161
Disabling the logging of NetBackup processes	171
Login Banner Configuration properties	173
Removing login banner screen and text	175
Enabling the Auto log off timeout option	176
Lotus Notes properties	177
Media properties	180
Results when media overwrites are not permitted	184
Recommended use for Enable SCSI reserve property	185
NDMP Global Credentials properties	186
Network properties	188
Network Settings properties	190
Reverse Host Name Lookup property	191
IP Address Family Support property	192
Port Ranges properties	192
Registered ports and dynamically-allocated ports	194
Preferred Network properties	195
Add or Change Preferred Network Settings dialog box	197
How NetBackup uses the directives to determine which network to use	199
Configurations to use IPv6 networks	202
Configurations to use IPv4 networks	203

Order of directive processing in the Preferred Network properties	204
Order of directives can affect processing	205
bptestnetconn utility to display Preferred Network information	205
Configuration to prohibit using a specified address	207
Configuration that uses the same specification for both the network and the interface—less constrictive	207
Configuration that uses the same specification for both the network and the interface—more constrictive	208
Configuration that limits the addresses, but allows any interfaces	208
Resilient Network properties	209
Resilient connection resource usage	212
Specifying resilient connections	212
Resource Limit properties	214
Restore Failover properties	216
Assigning an alternate media server as a failover restore server	218
Retention Periods properties	219
Changing a retention period	220
Determining retention periods for volumes	221
Retention Periods with end dates beyond 2038, excluding Infinity	222
Scalable Storage properties	222
Configuring advanced bandwidth throttling settings	224
Advanced bandwidth throttling settings	225
Servers properties	227
Adding a server to a servers list	230
Removing a server from a server list	231
Enabling NetBackup clustered master server inter-node authentication	231
Adding a trusted master server	233
Removing a trusted master server	234
Changing the master server that performs backups and restores for a client	237
SharePoint properties	237
Consistency check options for SharePoint Server	239
SLP Parameters properties	239
About batch creation logic in Storage Lifecycle Manager	244
Deprecated storage lifecycle policy parameters	245
Throttle Bandwidth properties	245

Add Bandwidth Settings dialog box for Throttle Bandwidth properties	246
Timeouts properties	247
Universal Settings properties	249
Logging the status of a redirected restore	252
UNIX Client properties	253
UNIX Server properties	254
VMware Access Hosts properties	255
Windows Client properties	257
Configuration options not found in the Host Properties	257
About using commands to change the configuration options on UNIX clients and servers	259
Configuration options for NetBackup servers	260
About robust logging	260
ALLOW_MEDIA_OVERWRITE option for NetBackup servers	262
ALLOW_NON_RESERVED_PORTS option for NetBackup servers and clients	262
AUTHENTICATION_DOMAIN option for NetBackup servers and clients	263
AUTHORIZATION_SERVICE option for NetBackup servers	265
AUTO_ADD_ALL_ALIASES_FOR_CLIENT option for NetBackup servers	266
BPBRM_VERBOSE option for NetBackup servers	267
BPDBJOBS_COLDEFS options for UNIX master servers	268
BPDBM_VERBOSE option for NetBackup servers	272
BPRD_VERBOSE option for NetBackup servers	273
BPTM_VERBOSE option for NetBackup servers	274
BPEND_TIMEOUT option for NetBackup servers	276
BPSTART_TIMEOUT option for NetBackup servers	276
CHECK_RESTORE_CLIENT option for NetBackup servers	277
CLIENT_CONNECT_TIMEOUT option for NetBackup servers	278
CLIENT_PORT_WINDOW option for NetBackup servers and clients	278
CLIENT_READ_TIMEOUT option for NetBackup servers	279
CLIENT_RESERVED_PORT_WINDOW option for NetBackup servers and clients	281
CONNECT_OPTIONS option for NetBackup servers and clients	281
DEFAULT_CONNECT_OPTIONS option for NetBackup servers and clients	285
DISABLE_JOB_LOGGING option for NetBackup servers	287

DISABLE_STANDALONE_DRIVE_EXTENSIONS option for NetBackup servers	288
DISALLOW_BACKUPS_SPANNING_MEDIA option for NetBackup servers	289
DISALLOW_CLIENT_LIST_RESTORE option for NetBackup servers	289
DISALLOW_CLIENT_RESTORE option for NetBackup servers	290
EMMSERVER option for NetBackup servers	291
FAILOVER_RESTORE_MEDIA_SERVERS option for NetBackup servers	292
FORCE_RESTORE_MEDIA_SERVER option for NetBackup servers	292
GENERATE_ENGLISH_LOGS option for NetBackup servers and clients	293
INCOMPLETE_JOB_CLEAN_INTERVAL option for NetBackup servers and clients	294
INITIAL_BROWSE_SEARCH_LIMIT option for NetBackup servers and clients	295
IP_ADDRESS_FAMILY option for NetBackup servers	296
JOB_PRIORITY option for NetBackup servers and clients	297
LIMIT_BANDWIDTH option for NetBackup servers	299
MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging	300
MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging	301
KEEP_LOGS_SIZE_GB for NetBackup servers and clients	302
ENABLE_CRITICAL_PROCESS_LOGGING	303
MEDIA_UNMOUNT_DELAY option for NetBackup servers	304
MEDIA_REQUEST_DELAY option for NetBackup servers	304
MEDIA_SERVER option for NetBackup servers	305
MPX_RESTORE_DELAY option for NetBackup servers	306
MUST_USE_LOCAL_DRIVE option for NetBackup servers	306
NBRNTD_IDLE_TIMEOUT option for NetBackup servers	307
PREFERRED_NETWORK option for NetBackup servers	308
RANDOM_PORTS option for NetBackup servers and clients	311
RE_READ_INTERVAL option for NetBackup servers	312
REQUIRED_NETWORK option for NetBackup servers	313
RESILIENT_NETWORK option for NetBackup master servers and clients	314
RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers	315

REVERSE_NAME_LOOKUP option for NetBackup servers and clients	315
SERVER option for NetBackup servers	316
SERVER_CONNECT_TIMEOUT option for NetBackup servers	319
SERVER_PORT_WINDOW option for NetBackup servers	319
SERVER_RESERVED_PORT_WINDOW option for NetBackup servers and clients	320
SKIP_RESTORE_TO_SYMLINK_DIR option for NetBackup servers	321
TELEMETRY_UPLOAD option for NetBackup servers	322
THROTTLE_BANDWIDTH option for NetBackup servers	323
TRUSTED_MASTER option for NetBackup servers	327
ULINK_ON_OVERWRITE option for NetBackup servers	327
USE_AUTHENTICATION option for NetBackup servers	329
USE_VXSS option for NetBackup servers and clients	329
VERBOSE option for NetBackup servers and clients	330
VXSS_NETWORK option for NetBackup servers	331
Configuration options for NetBackup clients	333
BPARCHIVE_POLICY option for NetBackup clients	334
BPARCHIVE_SCHED option for NetBackup clients	334
BPBACKUP_POLICY option for NetBackup clients	335
BPBACKUP_SCHED option for NetBackup clients	336
BUSY_FILE_ACTION option for NetBackup clients	337
BUSY_FILE_DIRECTORY option for NetBackup clients	338
BUSY_FILE_NOTIFY_USER option for NetBackup clients	339
BUSY_FILE_PROCESSING option for NetBackup clients	340
CLIENT_NAME option for NetBackup clients	340
COMPRESS_SUFFIX option for NetBackup clients	342
CRYPT_CIPHER option for NetBackup clients	342
CRYPT_KIND option for NetBackup clients	343
CRYPT_OPTION option for NetBackup clients	344
CRYPT_STRENGTH option for NetBackup clients	345
CRYPT_LIBPATH option for NetBackup clients	346
CRYPT_KEYFILE option for NetBackup clients	347
DISALLOW_SERVER_FILE_WRITES option for NetBackup clients	348
DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients	349
IGNORE_XATTR option for NetBackup clients	350
INFORMIX_HOME option for NetBackup clients	353
KEEP_DATABASE_COMM_FILE option for NetBackup clients	354

	KEEP_LOGS_DAYS option for NetBackup clients	354
	LIST_FILES_TIMEOUT option for NetBackup clients	355
	LOCKED_FILE_ACTION option for NetBackup clients	356
	MEDIA_SERVER option for NetBackup clients	356
	MEGABYTES_OF_MEMORY option for NetBackup clients	357
	NFS_ACCESS_TIMEOUT option for NetBackup clients	358
	OLD_VNETD_CALLBACK option for NetBackup clients	359
	REPORT_CLIENT_DISCOVERIES option for NetBackup clients	359
	RESTORE_RETRIES option for NetBackup clients	360
	SERVER option for NetBackup clients	361
	SYBASE_HOME option for NetBackup clients	362
	USE_CTIME_FOR_INCREMENTALS option for NetBackup clients	363
	USE_FILE_CHG_LOG option for NetBackup clients	364
	USEMAIL option for NetBackup clients	365
	VXSS_NETWORK option for NetBackup clients	365
Chapter 4	Configuring server groups	367
	About NetBackup server groups	367
	Configuring a NetBackup server group	367
	NetBackup server group properties	369
	Deleting a server group	369
Chapter 5	Configuring host credentials	370
	About configuring credentials	370
Chapter 6	Managing media servers	371
	Activating or deactivating a media server	371
	Adding a media server	372
	Registering a media server	373
	Generating a security certificate for a NetBackup media server	374
	Deleting all devices from a media server	375
	Removing a device host from the EMM database	378
	About decommissioning a media server	378
	About decommissioning limitations	379
	Before you decommission a media server	380
	Post decommission recommendations	380
	Decommission actions	381
	Previewing references to a media server	384
	Decommissioning a media server	385

Section 3	Configuring storage	388
Chapter 7	Configuring disk storage	389
	About configuring BasicDisk storage	389
	About configuring disk pool storage	389
	Inventorying a NetBackup disk pool	390
	Changing the disk pool state	390
Chapter 8	Configuring robots and tape drives	392
	NetBackup robot types	393
	About the device mapping files	394
	Downloading the device mapping files	394
	About configuring robots and tape drives in NetBackup	395
	About device discovery	396
	About device serialization	397
	About adding devices without discovery	398
	About robot control	398
	About drive name rules	400
	Configuring robots and tape drives by using the wizard	401
	Updating the device configuration by using the wizard	401
	Adding a robot to NetBackup manually	402
	Robot configuration options	404
	Managing robots	407
	Changing robot properties	407
	Configuring a robot to operate in manual mode	408
	Deleting a robot	409
	Moving a robot and its media to a new media server	409
	Adding a tape drive to NetBackup manually	410
	Tape drive configuration options	412
	Configuring drive name rules	416
	Adding a tape drive path	417
	About SCSI reserve on drive paths	417
	Drive path options	418
	About no rewind device files on UNIX	419
	Adding a shared tape drive to a NetBackup environment	419
	Correlating tape drives and SCSI addresses on Windows hosts	419
	Windows device correlation example	420
	Correlating tape drives and device files on UNIX hosts	421
	UNIX device correlation example	422
	Managing tape drives	423
	Changing a drive comment	423

About downed drives	424
Changing a drive operating mode	424
Changing a tape drive path	425
Changing a drive path operating mode	425
Changing tape drive properties	426
Changing a tape drive to a shared drive	426
Cleaning a tape drive from the Device Monitor	427
Deleting a drive	428
Resetting a drive	428
Resetting the mount time of a drive	429
Setting the drive cleaning frequency	429
Viewing drive details	430
Performing device diagnostics	430
Running a robot diagnostic test	430
Running a tape drive diagnostic test	432
Managing a diagnostic test step that requires operator intervention	433
Obtaining detailed information for a diagnostic test step	433
Verifying the device configuration	434
About automatic path correction	434
Enabling automatic path correction	435
Replacing a device	435
Updating device firmware	437
About the NetBackup Device Manager	438
About external access to NetBackup controlled devices on UNIX	438
Stopping and restarting the device manager	439

Chapter 9	Configuring tape media	440
	About NetBackup tape volumes	441
	About NetBackup volume pools	441
	About reserved volume pool name prefixes	442
	About scratch volume pools	443
	About NetBackup volume groups	444
	NetBackup media types	445
	About WORM media	447
	About using volume pools to manage WORM media	448
	About using unique drive and media types to manage WORM media	449
	Disabling WORM volume pool name verification	450
	About WORM media and the Quantum drive	450
	Supported WORM drives	450
	About adding volumes	451

About adding robotic volumes	451
About adding standalone volumes	452
About configuring media name and attribute rules	453
Adding volumes by using the wizard	453
About media settings rules	454
Configuring media settings	454
Media settings options	456
About barcodes	464
About barcode rules	466
Configuring barcode rules	467
Barcode rules settings	470
About media ID generation rules	472
Configuring media ID generation rules	473
Media ID generation options	475
About media type mapping rules	476
Adding volumes by using the Actions menu	477
Volume properties	478
Configuring media type mappings	481
About adding media type mapping entries	483
Default and allowable media types	483
Managing volumes	488
Changing the group of a volume	488
About rules for moving volumes between groups	489
Changing the owner of a volume	489
Changing volume properties	490
About assigning and deassigning volumes	491
Deleting a volume	492
Erasing a volume	493
About exchanging a volume	494
About frozen media	496
About injecting and ejecting volumes	497
About rescanning and updating barcodes	500
About labeling NetBackup volumes	501
About moving volumes	503
About recycling a volume	505
Suspending or unsuspending volumes	506
Managing volume pools	507
Adding a volume pool	507
Changing the properties of a volume pool	507
Deleting a volume pool	508
Managing volume groups on UNIX	509
Moving a volume group	509
Deleting a volume group	511

	About media sharing	511
	Configuring unrestricted media sharing	512
	Configuring media sharing with a server group	512
Chapter 10	Inventorying robots	514
	About robot inventory	514
	When to inventory a robot	515
	About showing a robot's contents	517
	About inventory results for API robots	518
	Showing the media in a robot	520
	About comparing a robot's contents with the volume configuration	521
	Comparing media in a robot with the volume configuration	522
	About previewing volume configuration changes	523
	Previewing volume configuration changes for a robot	524
	About updating the NetBackup volume configuration	526
	Volume update prerequisites	526
	Updating the NetBackup volume configuration with a robot's contents	527
	Robot inventory options	529
	About the vmphyinv physical inventory utility	530
	How vmphyinv performs a physical inventory	532
	Example volume configuration updates	536
Chapter 11	Configuring storage units	546
	About the Storage utility	546
	Using the Storage utility	547
	Creating a storage unit	548
	Creating a storage unit by copying a storage unit	551
	Changing storage unit settings	551
	Deleting storage units	552
	Media Manager storage unit considerations	553
	Disk storage unit considerations	555
	NDMP storage unit considerations	559
	About storage unit settings	561
	Absolute pathname to directory or absolute pathname to volume setting for storage units	561
	Density storage unit setting	562
	Disk pool storage unit setting	562
	Disk type storage unit setting	562
	Enable block sharing storage unit setting	563
	Enable multiplexing storage unit setting	563

High water mark storage unit setting	563
Low water mark storage unit setting	564
Maximum concurrent write drives storage unit setting	564
Maximum concurrent jobs storage unit setting	565
Maximum streams per drive storage unit setting	567
Media server storage unit setting	567
NDMP host storage unit setting	569
On demand only storage unit setting	570
Only use the following media servers storage unit setting	570
Properties option in the Change Storage Units dialog box	571
Reduce fragment size storage unit setting	572
Robot number storage unit setting	573
Robot type storage unit setting	573
Staging schedule option in Change Storage Units dialog	573
Storage device setting for storage units	574
Storage unit name setting	574
Storage unit type setting	574
Enable temporary staging area storage unit setting	574
Use any available media server storage unit setting	575

Chapter 12	Staging backups	576
	About staging backups	576
	About basic disk staging	577
	Creating a basic disk staging storage unit	578
	Creating a schedule for a basic disk staging storage unit	579
	Configuring multiple copies in a relocation schedule	580
	Disk staging storage unit size and capacity	582
	Finding potential free space on a BasicDisk disk staging storage unit	584
	Disk Staging Schedule dialog box	585
	Initiating a relocation schedule manually	588

Chapter 13	Configuring storage unit groups	589
	About storage unit groups	589
	Creating storage unit groups for backups	590
	Creating storage unit groups for snapshots	592
	Deleting a storage unit group	594
	Storage unit selection criteria within a group	594
	Media server load balancing	595
	Other load balancing methods	596
	Exception to the storage unit selection criteria	598
	About disk spanning within storage unit groups	598

Section 4	Configuring storage lifecycle policies (SLPs)	599
Chapter 14	Configuring storage lifecycle policies	600
	About storage lifecycle policies	600
	Creating a storage lifecycle policy	601
	Modifying the hierarchy of operations in a storage lifecycle policy	604
	Deleting a storage lifecycle policy	605
	Lifecycle operation administration using the nbstlutil command	606
Chapter 15	Storage operations	608
	Operation types in a storage lifecycle policy	608
	Backup operation in an SLP	608
	Backup From Snapshot operation in an SLP	610
	Duplication operation in an SLP	612
	Import operation in an SLP	614
	Index From Snapshot operation in an SLP	616
	Determining where and when the Index From Snapshot operation occurs	618
	Replication operation in an SLP	619
	Snapshot operation in an SLP	622
	About configuring storage lifecycle policies for snapshots and snapshot replication	624
	Primary snapshot storage unit	626
	Primary + Replication source snapshot storage unit	626
	Replication source + Replication target snapshot storage unit	627
	Replication target snapshot storage unit	627
	Replication source + Replication target + Mirror snapshot storage unit	628
	Replication target + Mirror snapshot storage unit	628
	Creating a hierarchy of storage operations in a storage lifecycle policy	628
Chapter 16	Retention types for SLP operations	631
	Retention types for storage lifecycle policy operations	631
	Capacity managed retention type for SLP operations	633
	Rules and recommendations for using the Capacity Managed retention type	634

	Capacity managed retention type and disk types that support SIS	634
	Expire after copy retention type for SLP operations	634
	Fixed retention type for SLP operations	635
	Maximum snapshot limit retention type for SLP operations	636
	Mirror retention type for SLP operations	637
	Target retention type for SLP operations	638
Chapter 17	Storage lifecycle policy options	639
	Storage Lifecycle Policy dialog box settings	639
	New or Change Storage Operation dialog box settings	643
	Properties tab of the Storage Operation dialog box	643
	Window tab of the Storage Operation dialog box	647
	Creating a new window for a storage lifecycle policy operation	649
	Excluding days from a window for a storage lifecycle policy operation	650
	Storage lifecycle policy validation dialog box	651
	Storage lifecycle policy Validation Report tab	652
Chapter 18	Using a storage lifecycle policy to create multiple copies	654
	About writing multiple copies using a storage lifecycle policy	654
	How the order of the operations determines the copy order	655
	About ensuring successful copies using lifecycles	655
Chapter 19	Storage lifecycle policy versions	657
	About storage lifecycle policy versions	657
	Storage lifecycle changes and versioning	658
	When changes to storage lifecycle policies become effective	659
	Deleting old storage lifecycle policy versions	660
Section 5	Configuring backups	662
Chapter 20	Creating backup policies	663
	About the Policies utility	664
	Planning for policies	666
	Windows example of one client in multiple policies	668
	Policy attributes that affect how clients are grouped in policies	669

About Microsoft DSFR backups	670
Creating a backup policy	673
Adding or changing schedules in a policy	675
Changing multiple policies at one time	676
Copying or moving policy items to another policy	677
Copying a policy to create a new policy	678
Copying a schedule into the same policy or different policy	679
Deleting schedules, backup selections, or clients from a policy	679
Policy Attributes tab	680
Policy type (policy attribute)	681
Data classifications (policy attribute)	684
Policy storage (policy attribute)	685
Policy volume pool (policy attribute)	688
Take checkpoints every __ minutes (policy attribute)	689
Limit jobs per policy (policy attribute)	693
Job priority (policy attribute)	695
Media Owner (policy attribute)	695
Go into effect at (policy attribute)	695
Follow NFS (policy attribute)	696
Backup Network Drives (policy attribute)	697
Cross mount points (policy attribute)	699
Compression (policy attribute)	703
Encryption (policy attribute)	705
Collect disaster recovery information for Bare Metal Restore (policy attribute)	707
Collect true image restore information (policy attribute) with and without move detection	707
Allow multiple data streams (policy attribute)	711
Disable client-side deduplication (policy attribute)	715
Enable granular recovery (policy attribute)	715
Use Accelerator (policy attribute)	716
Enable optimized backup of Windows deduplicated volumes	737
Keyword phrase (policy attribute)	740
Snapshot Client and Replication Director (policy attributes)	740
Perform block level incremental backups (policy attributes)	741
Use Replication Director (policy attributes)	741
Perform snapshot backups (policy attributes)	743
Microsoft Exchange Attributes (policy attributes)	743
Schedules tab	743
Schedule Attributes tab	744
Name (schedule attribute)	745
Type of backup (schedule attribute)	745
Synthetic backup (schedule attribute)	754

Accelerator forced rescan option (schedule attribute)	755
Calendar (schedule attribute)	757
Frequency (schedule attribute)	757
Instant Recovery (schedule attribute)	759
Multiple copies (schedule attribute)	760
Override policy storage (schedule attribute)	764
Override policy volume pool (schedule attribute)	765
Override media owner (schedule attribute)	765
Retention (schedule attribute)	766
Media multiplexing (schedule attribute)	769
Start Window tab	776
Adding, changing, or deleting a time window in a policy	
schedule	776
Example of schedule duration	779
Excluding days from a schedule	780
Include Dates tab	781
Calendar scheduling with the Include Dates tab	781
How NetBackup determines which schedule to run next	782
About schedule windows that span midnight	785
How open schedules affect calendar-based and frequency-based	
schedules	786
Creating an open schedule in the NetBackup Administration	
Console	789
Runtime considerations that affect backup frequency	790
About the Clients tab	790
Adding or changing clients in a policy	791
Installing client software on secure and trusting UNIX clients	793
Browse for Hyper-V virtual machines	795
Backup Selections tab	796
Adding backup selections to a policy	797
Verifying the Backup Selections list	799
How to reduce backup time	802
Pathname rules for Windows client backups	803
Pathname rules for Windows disk image (raw) backups	806
Pathname rules for Windows registry backups	807
About hard links to files and directories	808
Pathname rules for UNIX client backups	810
Pathname rules for the clients that run extension products	818
About the directives on the Backup Selections list	818
Files that are excluded from backups by default	831
About excluding files from automatic backups	832
Files that are excluded by Microsoft Windows Backup	833
VMware Policy tab	833

Hyper-V Policies tab	835
Disaster Recovery tab	837
Adding policies to the Critical Policies list of a catalog backup policy	840
Creating a Vault policy	840
Performing manual backups	842
Active Directory granular backups and recovery	843
System requirements for Active Directory granular NetBackup backups and recovery	844
Creating a policy that allows Active Directory granular restores	844
Restoring Active Directory objects	846
Troubleshooting granular restore issues	847

Chapter 21	Synthetic backups	850
	About synthetic backups	850
	Recommendations for synthetic backups and restores	851
	Synthetic full backups	854
	Synthetic cumulative incremental backups	855
	Schedules that must appear in a policy for synthetic backups	857
	Adding clients to a policy for synthetic backups	857
	Change journal and synthesized backups	858
	True image restore and synthesized backups	858
	Displaying synthetic backups in the Activity Monitor	858
	Logs produced during synthetic backups	859
	Synthetic backups and directory and file attributes	859
	Using the multiple copy synthetic backups method	860
	Configuring multiple copy synthetic backups	861
	Configuration variables for multiple copy synthetic backups	862
	Multiple copy synthetic backups configuration examples	863
	Optimized synthetic backups	864
	Optimized synthetic backups for deduplication	864

Chapter 22	Protecting the NetBackup catalog	865
	About the NetBackup catalog	865
	Parts of the NetBackup catalog	866
	About the NetBackup image database	868
	About the NetBackup relational database	870
	Protecting the NetBackup catalog	874
	About catalog backups	875
	Using the Catalog Backup Wizard to configure a catalog backup	876

Using the Policy Configuration Wizard to configure a catalog backup	879
Configuring a catalog backup manually	881
Backing up NetBackup catalogs manually	883
Concurrently running catalog backups with other backups	884
About catalog policy schedules	884
How catalog incrementals and standard backups interact on UNIX	885
Determining whether or not a catalog backup succeeded	885
Strategies that ensure successful NetBackup catalog backups	886
Recovering the catalog	886
Disaster recovery emails and the disaster recovery file	886
Archiving the catalog and restoring from the catalog archive	887
Creating a catalog archiving policy	890
Catalog archiving commands	891
Catalog archiving considerations	893
Extracting images from the catalog archives	894
Estimating catalog space requirements	894
NetBackup file size considerations on UNIX systems	896
About the binary catalog format	896
Moving the image catalog	896
About image catalog compression	899

Chapter 23	About the NetBackup relational database	903
	About the NetBackup relational database (NBDB) installation	903
	About NetBackup master server installed directories and files	906
	About the NetBackup configuration entry	915
	Sybase SQL Anywhere server management	916
	Sybase SQL Anywhere and clustered environments	917
	Using the NetBackup Database Administration utility on Windows	917
	About the General tab of the NetBackup Database Administration utility	919
	About the Tools tab of the NetBackup Database Administration utility	925
	Using the NetBackup Database Administration utility on UNIX	931
	About the Select/Restart Database and Change Password menu options	933
	About the Database Space and Memory Management menu options	934

About the Database Validation Check and Rebuild menu options	936
About the Move Database menu options	937
About the Unload Database menu options	938
About the Backup and Restore Database menu options	939
Post-installation tasks	940
Changing the NetBackup database password	940
Moving NBDB database files after installation	941
Adding a mirrored transaction log	943
Creating the NBDB database manually	945
About backup and recovery procedures	948
Database transaction log	948
About catalog recovery	949
Commands for backing up and recovering the relational databases	949
About the catalog backup process	951
Unloading the NetBackup database	952
Terminating database connections	953
Moving the NetBackup database from one host to another	954

Chapter 24 Managing backup images 955

About the Catalog utility	955
About searching for backup images	956
Verifying backup images	959
Viewing job results	959
Promoting a copy to a primary copy	960
Duplicating backup images	962
About multiplexed duplication considerations	967
Jobs that appear while making multiple copies	968
Expiring backup images	969
About importing backup images	970
About importing expired images	970
Importing backup images	971

Section 6 Configuring replication 973

Chapter 25 About NetBackup replication 974

About NetBackup replication	974
About NetBackup Auto Image Replication	975
One-to-many Auto Image Replication model	977
Cascading Auto Image Replication model	977

About the domain relationship for replication	980
About the replication topology for Auto Image Replication	980
Viewing the replication topology for Auto Image Replication	982
About trusted master servers for Auto Image Replication	987
About the storage lifecycle policies required for Auto Image Replication	988
Using NetBackup Access Control (NBAC) with Auto Image Replication	993
Auto Image Replication setup overview	994
How to resolve volume changes for Auto Image Replication	995
Removing or replacing replication relationships in an Auto Image Replication configuration	998
About restoring from a backup at a target master domain	1012
Reporting on Auto Image Replication jobs	1013
About NetBackup Replication Director	1014

Section 7 Monitoring and reporting 1016

Chapter 26 Monitoring NetBackup activity 1017

About the Activity Monitor	1017
Setting Activity Monitor options	1019
About the Jobs tab	1020
Viewing job details in the Activity Monitor	1022
Showing or hiding column heads in the Activity Monitor	1022
Monitoring the detailed status of a selected job in the Activity Monitor	1023
Deleting completed jobs in the Activity Monitor	1023
Canceling a job that has not completed in the Activity Monitor	1023
Restarting a failed (completed) job in the Activity Monitor	1024
Suspending restore or backup jobs in the Activity Monitor	1024
Resuming suspended or incomplete jobs in the Activity Monitor	1024
Printing job list information in the Activity Monitor	1025
Exporting Activity Monitor data to a text file	1025
Printing job detail information	1025
Copying Activity Monitor text to a file	1026
Changing the Job Priority dynamically from the Activity Monitor	1026
About the Daemons tab	1027
Using the nbrbutil utility to configure the NetBackup Resource Broker	1032

Types of NetBackup daemons	1037
Monitoring NetBackup daemons	1037
Starting or stopping a daemon	1038
Displaying all media servers in the Activity Monitor	1038
About the Processes tab	1039
Monitoring NetBackup processes in the Process Details dialog box	1043
About the Drives tab	1044
Monitoring tape drives	1045
Cleaning tape drives from the Activity Monitor	1045
About the jobs database	1045
Changing the default bpdjobs_options values	1046
About the BPDBJOBS_OPTIONS environment variable	1047
bpdjobs command line options	1049
Enabling the bpdjobs debug log	1049
About the Device Monitor	1050
About media mount errors	1051
About pending requests and actions	1052
About pending requests for storage units	1052
Resolving a pending request	1053
Resolving a pending action	1054
Resubmitting a request	1055
Denying a request	1055
 Chapter 27	
Reporting in NetBackup	1056
About the Reports utility	1056
Running a report	1060
Copying report text to another document	1060
Saving or exporting a report	1061
Printing a report	1061
 Chapter 28	
Using the Logging Assistant	1063
About the Logging Assistant	1063
Logging Assistant sequence of operation	1066
Viewing the Logging Assistant records	1067
Viewing the details of a Logging Assistant record	1067
Adding a new Logging Assistant record	1070
Deleting a Logging Assistant record	1072
Setting up debug logging	1073
Set minimum debug logging	1077
Disabling debug logging	1078
Collecting and uploading debug logs	1079

	Collecting and uploading nbsu information	1085
	Uploading evidence	1088
Section 8	Administering NetBackup	1091
Chapter 29	Management topics	1092
	Configuring the NetBackup Client Service	1092
	NetBackup naming conventions	1093
	Wildcard use in NetBackup	1094
	Wildcard characters in backup selections for an NDMP policy	1096
	How to administer devices on other servers	1099
	How to access media and devices on other hosts	1100
Chapter 30	Accessing a remote server	1102
	Prerequisites for accessing a remote server	1102
	Allow access to another server	1103
	Authorize users of one server to access another server	1105
	Using the Change Server operation in the NetBackup-Java Administration Console	1106
	Accessing remote servers	1112
	About running the NetBackup Administration Console on a NetBackup client	1113
	About troubleshooting remote server administration	1113
Chapter 31	Using the NetBackup Remote Administration Console	1115
	About the NetBackup Remote Administration Console	1116
	About authorizing NetBackup users	1119
	Authorization file (auth.conf) characteristics	1120
	About authorizing nonroot users for specific applications	1123
	About authorizing specific tasks in the Backup, Archive, and Restore user interface	1124
	About authorizing users on Windows NetBackup server	1125
	Restricting access to the NetBackup Administration Console on Windows	1125
	Run-time configuration options for the NetBackup Administration Console	1126
	FIREWALL_IN	1126
	FORCE_IPADDR_LOOKUP	1128
	INITIAL_MEMORY, MAX_MEMORY	1129
	MEM_USE_WARNING	1130

NBJAVA_CLIENT_PORT_WINDOW	1130
NBJAVA_CORBA_DEFAULT_TIMEOUT	1131
NBJAVA_CORBA_LONG_TIMEOUT	1131
PBX_PORT	1132
VNETD_PORT	1132
About customizing jnbSA and jbpSA with bp.conf entries	1132
About improving NetBackup-Java performance	1133
About running the Java console locally	1133
About running a console locally and administering a remote server	1134
About enhancing console performance	1134
About determining better performance when console is run locally or uses remote display back	1135
About adjusting time zones in the NetBackup-Java console	1137
Adjusting the time zone in the NetBackup-Java console	1138
Configuring a custom time zone in the NetBackup-Java console	1138

Chapter 32 Alternate server restores 1140

About alternate server restores	1140
About supported configurations for alternate server restores	1141
About performing alternate server restores	1142
About modifying the NetBackup catalogs	1143
Overriding the original server for restores	1144
About enabling automatic failover to an alternate server	1146
Expiring and importing media for alternate server restores	1147

Chapter 33 Managing client backups and restores 1149

About server-directed restores	1149
About client-redirected restores	1150
About restore restrictions	1151
Allowing all clients to perform redirected restores	1151
Allowing a single client to perform redirected restores	1152
Allowing redirected restores of a specific client's files	1152
Examples of redirected restores	1153
About restoring the files that have Access Control Lists (ACLs)	1158
About setting the original atime for files during restores on UNIX	1160
Restoring the System State	1160
About the backup and restore of compressed files on VxFS file systems	1163
About backups and restores on ReFS	1164

Chapter 34	Powering down and rebooting NetBackup servers	1165
	Powering down and rebooting NetBackup servers	1165
	Shutting down and starting up all NetBackup services and daemons	1166
	Rebooting a NetBackup server	1167
	Rebooting a NetBackup media server	1167
	About displaying active processes with bpps on UNIX	1168
	About displaying robotic processes with vmps on UNIX	1169
Chapter 35	About Granular Recovery Technology	1170
	About installing and configuring Network File System (NFS) for Active Directory Granular Recovery	1170
	About configuring Services for Network File System (NFS) on Windows 2012 and Windows 2012 R2	1171
	Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 media server	1172
	Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 client	1176
	About configuring Services for Network File System (NFS) on Windows 2008 and Windows 2008 R2	1179
	Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2	1181
	Disabling the Client for NFS on the media server	1183
	Disabling the Server for NFS	1185
	Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology (GRT)	1187
	Configuring a different network port for NBFSD	1188
Index		1189

About NetBackup

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

Introducing NetBackup interfaces

This chapter includes the following topics:

- [About NetBackup](#)
- [NetBackup documentation](#)
- [About NetBackup administration interfaces](#)
- [Using the NetBackup Administration Console](#)
- [NetBackup configuration wizards](#)
- [Running the Troubleshooter](#)

About NetBackup

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

NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. An administrator can carefully schedule backups to achieve systematic and complete backups over a period of time, and optimize network traffic during off-peak hours. The backups can be full or incremental. Full backups back up all client files. Incremental backups back up only the files that have changed since the last backup.


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

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

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

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

Figure 1-1 NetBackup storage domain example



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

- The master server manages backups, archives, and restores. The master server is responsible for media and device selection for NetBackup. Typically, the master server contains the NetBackup catalog. The catalog contains the internal databases that contain information about NetBackup backups and configuration.
- Media servers provide additional storage by allowing NetBackup to use the storage devices that are attached to them. Media servers can also increase performance by distributing the network load. Media servers can also be referred to by using the following terms:
 - Device hosts (when tape devices are present)
 - Storage servers (when I/O is directly to disk)

- Data movers (when data is sent 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.

NetBackup documentation

For a complete list of NetBackup technical documents for each supported release, see the NetBackup release notes, Administration, Installation, Troubleshooting, Getting Started, and Solutions Guides page at the following URL:

<http://www.symantec.com/docs/DOC5332>

The documents are in Adobe® Portable Document Format (PDF), viewable with the Adobe Acrobat Reader. Download the reader from <http://www.adobe.com>.

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

The [NetBackup landing page](#) on the Symantec Support website contains a list of helpful how-to and product alert topics.

About NetBackup administration interfaces

NetBackup can be administered with several interfaces. The best choice depends on personal preference and the workstation that is available to the administrator.

Table 1-1 NetBackup administration interfaces

Name of interface	Description
NetBackup Administration Console	<p>The NetBackup Administration Console is the recommended interface and is the interface referred to by most procedures and examples in the documentation.</p> <p>Multiple base versions of the NetBackup Administration Console are installed when you install the server software. Select the version of the console that is compatible with the NetBackup server that you want to administer.</p> <p>See “Administering remote servers of different versions” on page 39.</p> <p>To start the NetBackup Administration Console</p> <ul style="list-style-type: none"> ■ On Windows, select NetBackup 7.7.1 Administration Console from the Start menu. ■ On UNIX, run the <code>jnbSA</code> command. <p>Note: To log in to any NetBackup Administration Console, your login credentials must be authenticated from the connecting master or media server. This authentication occurs whether or not NetBackup Access Control (NBAC) is in use.</p>
	<p>Run the <code>tpconfig</code> command to start a character-based, menu interface for device management.</p> <p>Use the <code>tpconfig</code> interface from any terminal (or terminal emulation window) that has a <code>termcap</code> or a <code>terminfo</code> definition.</p>
Command line	<p>NetBackup commands are available on both Windows and UNIX platforms. Enter NetBackup commands at the system prompt or use the commands in scripts.</p> <p>All NetBackup administrator programs and commands require root or administrator user privileges by default.</p> <p>See “About authorizing nonroot users for specific applications” on page 1123.</p> <p>For complete information on all NetBackup commands, see the NetBackup Commands Reference Guide.</p> <p>To view the commands online, use the UNIX <code>man</code> command.</p>

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.

See [“Using the NetBackup Administration Console”](#) on page 41.

See [“Administering remote servers of different versions”](#) on page 39.

About setting up the NetBackup Administration Console on UNIX

NetBackup provides two Java-based administration consoles through which the administrator can manage NetBackup. The consoles can be run on either of the following systems:

- Directly on a supported NetBackup-Java capable UNIX computer by running `/usr/openv/java/jnbSA &`

The `jnbSA` command is described in the [NetBackup Commands Reference Guide](#).

Use the `-r` command to connect to the compatible version of the console for the server that you want to administer.

- On a Windows computer that has the **NetBackup Administration Console** installed.

From the multiple versions of consoles installed, select the version of the console that is compatible with the NetBackup server that you want to administer.

Startup procedures and configuration information is explained in the following topics.

Setting up a CDE for the Java-based NetBackup interfaces on UNIX

Always set the window manager so that a window becomes active only when clicked. The NetBackup Java interfaces do not run properly with auto-focus enabled. Auto-focus causes a window to become active when the pointer moves over the window.

The following are general instructions for correctly setting up the focus on a CDE (Common Desktop Environment) window manager. CDE is the preferred window manager for NetBackup Java applications.

The NetBackup host that is specified on the login dialog box and the computer where you start the **NetBackup Administration Console** must run the same NetBackup version.

The **NetBackup Administration Console** supports remote X Windows display only between same-platform systems. For example, assume that you are on a Solaris system named *system1* and the NetBackup-Java software is on a Solaris system named *system2*. Here, you can display the interface on *system1* by performing `rlogin` to *system2*. Then run `jnbSA -d system1`. However, if *system2* was an HP system, you could display `jnbSA` only directly on *system2*.

In addition, the system where the console appears must run a version of the operating system that the console supports. Refer to the NetBackup release notes for supported versions, including any required release updates.

The `jnbSA` command is described in the [NetBackup Commands Reference Guide](#).

To set up a CDE (Common Desktop Environment) for NetBackup Java interfaces

- 1 On the front panel in the **CDE** window, click the **Style Manager** control icon.
- 2 On the Style Manager toolbar, click the Window control icon.
- 3 In the Style Manager-Window dialog box, click **Click In Window To Make Active**.
- 4 Click **OK**.
- 5 Click **OK** when prompted to Restart the Workspace Manager.

About using the NetBackup Administration Console

Use the **NetBackup Administration Console** to administer Windows or UNIX computers.

To start the NetBackup Administration Console

- 1 On a Windows host where the **NetBackup Administration Console** is installed and configured, select **Start > Programs > Symantec NetBackup > NetBackup 7.7.1 Administration Console**.

On a UNIX computer, use the `jnbSA` command:

```
/usr/opensv/java/jnbSA &
```

- 2 The logon 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 logon 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 logon 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.

Multiple base versions of the **NetBackup Administration Console** are installed when you install NetBackup server software. For example, when you install or upgrade to NetBackup 7.7.1, base versions from 7.0 through 7.7.1 of the console are all installed.

Note: Symantec recommends that after you install or upgrade NetBackup server software, you should uninstall older versions of the NetBackup Remote Administration Console (Windows and Java) present on the host. If the native **NetBackup Administration Console** for Windows is present, it is automatically uninstalled when you install or upgrade the NetBackup server software.

Note: To upgrade any of the multiple versions of consoles to a patch, you must first install the base version of the Remote Administration Console. Use the installer of the base version to install the Remote Administration Console. You must then upgrade to the corresponding required patch of the Remote Administration Console.

Upgrading directly to a patch version of the **NetBackup Administration Console** from the multiple versions of the consoles is not supported.

Note: The NetBackup Administration Console is best viewed at a screen resolution of 1280 X 1024 or higher. The minimum supported screen resolution to use NetBackup Administration Console is 1024 X 768.

Note: The **NetBackup Administration Console** does not support user-defined characters (UDC) and vendor-defined characters (VDC) because of the implementation of Java's encoding converters.

Note: On non-English versions of Windows and UNIX systems, the NetBackup Administration Console may display non-US ASCII characters incorrectly. This issue can lead to functional failures.

This issue results from a character encoding mismatch between the NetBackup server and the **NetBackup Administration Console**. For a detailed description of the configuration, refer to the following tech note on the Symantec Support website:

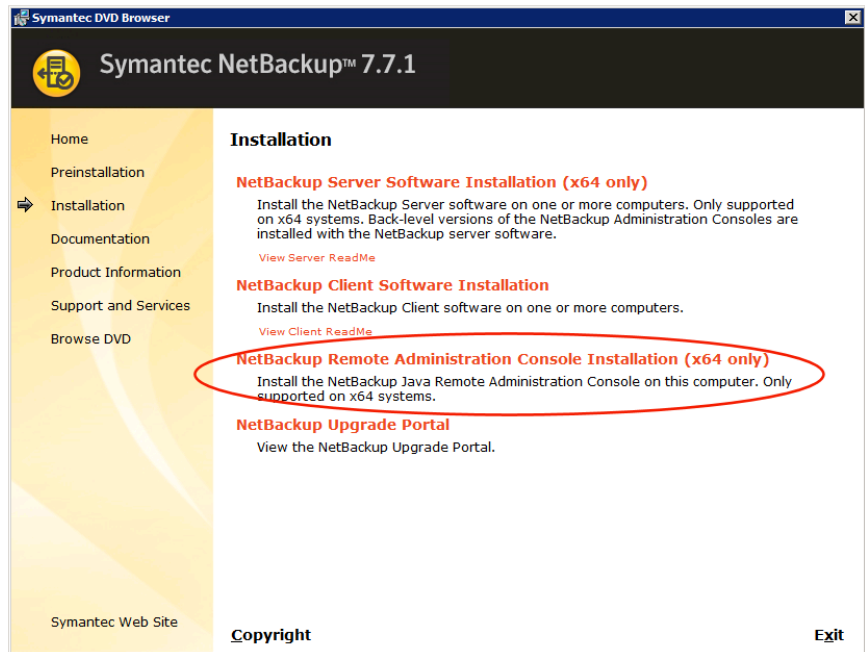
<http://www.symantec.com/docs/TECH75745>

Note: To use the **NetBackup Administration Console** on a Windows computer, the Microsoft Windows UAC (User Access Control) feature must be disabled. See the following link for instructions:

<http://windows.microsoft.com/en-us/windows/tum-user-account-control-on-off#1TC=windows-7>

Figure 1-2 shows the selection to install only the **NetBackup Administration Console**.


Figure 1-2 Selection on the installation panel to install only the NetBackup Administration Console



See [“Accessing remote servers”](#) on page 1112.

Administering remote servers of different versions

The NetBackup server installation provides multiple versions of the NetBackup Administration Console to administer remote servers of different versions. When starting the Administration Console, select the version of the console that is compatible with the NetBackup server that you want to administer.



Alternatively, from the command line, you can use the `jnbSA` command with the `-r` option to launch the console. For example, to connect to a 7.0 master server from a 7.7.1 UNIX master server, enter the following command on the 7.7.1 master server:

```
./jnbSA -r 7.0
```

If no `-r` option is specified, the default is the NetBackup version of the current master server.


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

Backward compatibility with triple-dot versions

The NetBackup Administration Console is backward-compatible between a patch release (x.x.x.x) and a major (x.x) or minor release (x.x.x) that shares the same first and second digits.

For example, the 7.6.0.2 console is compatible with a 7.6 master server. However, a NetBackup 7.6.0.2 console cannot administer a 7.5.0.7 master server. See [Figure 1-3](#) for various examples.

Figure 1-3 Examples of supported and unsupported back-level console configurations



See [“Accessing remote servers”](#) on page 1112.

Using the NetBackup Administration Console

The **NetBackup Administration Console** provides a graphical user interface through which the administrator can manage NetBackup. It runs on both the Windows and the UNIX platforms.

Figure 1-4 NetBackup Administration Console

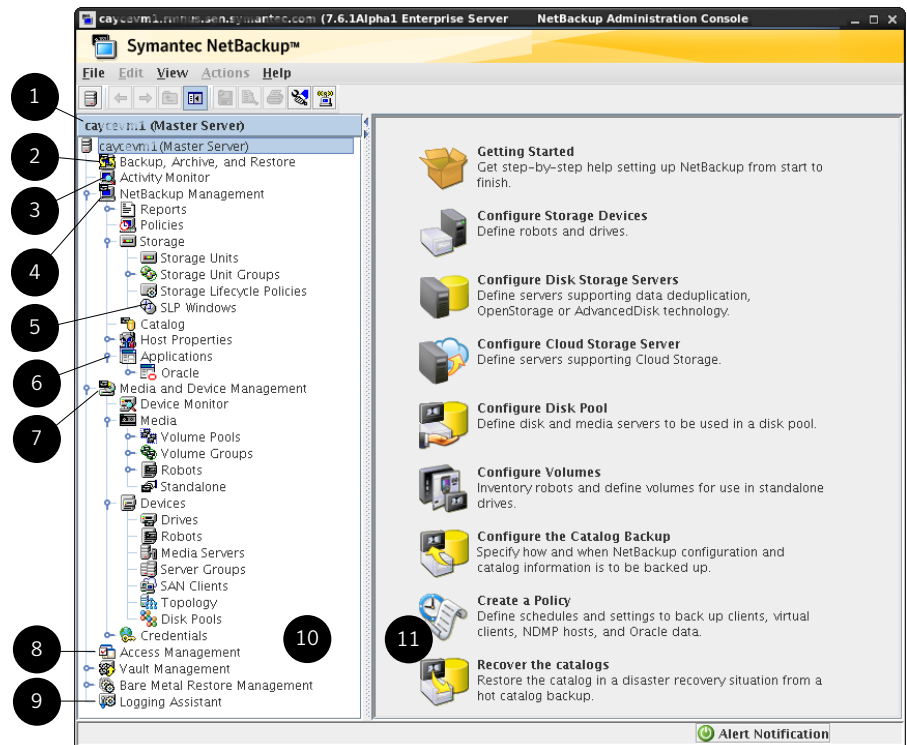


Table 1-2 NetBackup Administration Console description

	Description
1	<p>The name of the currently selected master server.</p> <p>The information in the NetBackup Administration Console applies to this server only.</p>
2	<p>Backup, Archive, and Restore</p> <p>Administrators can use the Backup, Archive, and Restore utility to back up, archive, and restore files, directories, and the formatted raw partitions that reside on the computer that is currently selected. Files can be restored at any time. However, files can be backed up or archived only during the time periods that the administrator defines within a schedule for user backups. Users can view the progress and final status of the operations performed.</p> <p>Note: An archive is a special type of 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 manual, references to backups also apply to the backup portion of archive operations (except where otherwise noted).</p> <p>Documentation for the NetBackup client is available as online Help from the Backup, Archive, and Restore interface.</p>
3	<p>Activity Monitor</p> <p>Displays the NetBackup job information. The Activity Monitor provides the control over the jobs, services, processes, and drives.</p>
4	<p>NetBackup Management</p> <p>The NetBackup Management contains many of the utilities to manage NetBackup:</p> <p>Reports</p> <p>Use to compile information to verify, manage, and troubleshoot NetBackup operations.</p> <p>See “About the Reports utility” on page 1056.</p>

Table 1-2 NetBackup Administration Console description (*continued*)

Description
Policies Use 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 664.
Storage Use to configure storage units, storage unit groups, storage lifecycle policies (SLPs), and windows for secondary operations in SLPs. See “About the Storage utility” on page 546.
Catalog Use 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 the NetBackup catalog, are located on the NetBackup master and media server. The catalog contains 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: <ul style="list-style-type: none">■ 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 955.
Host Properties Use 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 NetBackup Host Properties” on page 60.

Table 1-2 NetBackup Administration Console description (*continued*)

	Description
5	<p>SLP Windows</p> <p>SLP Windows contains any schedules that are defined for secondary operations in a storage lifecycle policy.</p> <p>A window can better define when the job for a secondary operation can run, making sure that the job does not interfere with jobs of a higher priority.</p>
6	<p>Applications</p> <p>Applications contains all discovered instances on SQL and Oracle clients and any manually added instances. These instances are used to build an Oracle or a SQL Intelligent Policy.</p> <p>For more information, see the NetBackup for Oracle Administrator's Guide and the NetBackup for SQL Server Administrator's Guide.</p>
7	<p>Media and Device Management</p> <p>Contains the utilities for managing the media and devices that NetBackup uses to store backups:</p> <hr/> <p>Device Monitor</p> <p>Manages the drives, device paths, and service requests for operators.</p> <hr/> <p>Media</p> <p>Adds and manages removable media.</p> <hr/> <p>Devices</p> <p>Adds, configures, and manages storage devices.</p> <hr/> <p>Credentials</p> <p>Adds, removes, and manages logon credentials for the following:</p> <ul style="list-style-type: none"> ■ NDMP hosts (requires the NetBackup for NDMP license). ■ Storage servers (requires a NetBackup Data Protection Optimization Option license). <p>Credentials appears only if one of these license keys is installed.</p>
8	<p>Access Management</p> <p>NetBackup administrators can protect a NetBackup configuration by defining who may access NetBackup and what functions a user group can perform. Access Management is enabled when NetBackup Product Authentication and Authorization and NetBackup Access Control (NBAC) is configured.</p> <p>For more information about NetBackup Access Control, see the NetBackup Security and Encryption Guide.</p>
9	<p>Logging Assistant</p> <p>The Logging Assistant is a tool that helps administrators configure, gather, and upload debug logs.</p> <p>See “About the Logging Assistant” on page 1063.</p>

Table 1-2 NetBackup Administration Console description (*continued*)

	Description
10	Additional licensed utilities appear under the main NetBackup nodes.
11	The Details pane contains the configuration wizards and details specific to the utility that is selected.

See “[About security certificates for NetBackup media servers](#)” on page 45.

See “[Configuring log files for the NetBackup Administration Console](#)” on page 46.

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

About security certificates for NetBackup media servers

NetBackup requires that media servers have a digital security certificate so that they function correctly in the following use cases:

NetBackup Administration Console authorization for Enhanced Auditing

In the Enhanced Auditing mode, a security certificate is mandatory when a user wants to connect to a media server through the **NetBackup Administration Console**.

See the *NetBackup Security and Encryption Guide*:

<http://www.symantec.com/docs/DOC5332>

For cloud storage

The NetBackup CloudStore Service Container requires that a certificate be installed on the media server. If one is not installed, the Service Container cannot start.

See the *NetBackup Cloud Administrator's Guide*:

<http://www.symantec.com/docs/DOC5332>

The procedure to generate a security certificate is in a different topic.

See “[Generating a security certificate for a NetBackup media server](#)” on page 374.

See “[Using the NetBackup Administration Console](#)” on page 41.

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 NetBackup Administration Console

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

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

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

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

Configuring log files for the NetBackup Administration Console

You can configure the logging options for the **NetBackup Administration Console**.

Caution

Most of the logging options degrade **Administration Console** performance and may write large amounts of data to the log files, possibly filling up the file system. Consider the following items that may help prevent filling up the file system on which NetBackup is installed:

- Use an environment variable named `JNB_LOG` as follows:
 - Set the variable to a log file pathname on a file system that has enough space for the log files.
Delete the environment variable when you stop logging **Administration Console** information.
 - On UNIX or Linux, set the variable to `/dev/tty`. The `/dev/tty` setting writes the log information to the console window.
- On UNIX or Linux: Use the `-l` option to the `jnbSA` or `jbpSA` command to specify one of the following settings:
 - To a log file pathname on a file system that has enough space for the log file.
 - To `/dev/tty`. The `/dev/tty` setting writes the log information to the console window.

The log file directories and names are described in a different topic.

See [“NetBackup Administration Console log files”](#) on page 47.

To configure logging for the NetBackup Administration Console

- 1 In the **NetBackup Administration Console**, click **View > Options** to open the **Options** dialog box.
- 2 Click the **Logging** tab.
- 3 Select the information that you want to write to the log file:

Enable command logs Enable this option to capture the output from the commands that the NetBackup-Java application server (bpjava-susvc) runs.

See [“About the Processes tab”](#) on page 1039.

Note: The output from each command instance is limited to 500 lines in the log file.

Enable command line logs Enable this option to capture command-line activity into the log file. NetBackup writes information about the commands that reside in the following directories:

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

Note: The NetBackup Administration Console often uses APIs and not the NetBackup commands to get data or update data. In those cases, information is not written to the log file.

Add date and timestamp to logs Enable this option to include a timestamp when NetBackup writes information to the log file.

Enable job try logs Enable this option to capture logging information from the NetBackup **Activity Monitor**.

Verbosity Use this option to set the verbosity of the logs. Five (5) is the highest logging level.

4 Click **OK**.

See [“Using the NetBackup Administration Console”](#) on page 41.

NetBackup Administration Console log files

The **NetBackup Administration Console** writes information into log files in the following directories:

UNIX/Linux `/usr/opensv/netbackup/logs/user_ops/nbjlogs`

Windows `install_path\NetBackup\logs\user_ops\nbjlogs`

For the **NetBackup Remote Administration Console**, the log files are in the following directory:

`install_path\Veritas\Java\logs`

The log file names have the following format:

UNIX/Linux `jbp.`whoami`.gui.$$$.log`

The log files always begin with `jbp`. The following items describe the other components of the log file name:

- `gui`: Either `jpbSA` (Backup, Archive, and Restore interface) or `jnbSA` (NetBackup Administration Console).
- `$$$`: The ID of the process for which log information was written when the log file was created.

The following is an example:

`jbp.root.jnbSA.6156.log`

Windows `jbp.yyyymmddthhmtssths.log`

The log files always begin with `jbp`. The following items describe the other components of the log file name:

- `yyyy`: Year
- `mm`: Month
- `dd`: Day
- `thh`: Hour
- `tmm`: Minute
- `tss`: Second
- `ths`: Half-second

The following is an example:

`jbp.201504149323710.log`

NetBackup creates a new log file when the current log file reaches a certain size. The size varies depending on many factors. Those factors include the number of jobs that are active, the amount of activity the Administration Console performs, and others.

Log file retention corresponds to the retention level that you can configure in the master server host properties **Logging** screen.

See [“Logging properties”](#) on page 161.

See [“Configuring log files for the NetBackup Administration Console”](#) on page 46.

NetBackup configuration wizards

The easiest way to configure NetBackup is to use the configuration wizards available in the **NetBackup Administration Console**. The wizards help configure the properties of a NetBackup environment.

The wizard selection varies in the pane on the right, depending on which NetBackup utility is selected in the left of the screen.

Table 1-3 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 Backup Wizard■ Policy Configuration Wizard <p>Configure more advanced properties through the individual utilities.</p> <p>Note: Do not use the Getting Started Wizard to configure policies for Replication Director.</p>
Device Configuration Wizard	<p>Click Configure Storage Devices to configure NetBackup to use robotic tape libraries and tape drives.</p>
Storage Server Configuration Wizard	<p>Click Configure Disk Storage Servers to create the servers that manage disk storage.</p> <p>The wizard appears if a NetBackup Data Protection Optimization Option license is installed.</p>
Cloud Storage Server Configuration Wizard	<p>Click Configure Cloud Storage Servers to add or configure a cloud provider.</p> <p>For more information, see the NetBackup Cloud Administrator's Guide.</p>
Disk Pool Configuration Wizard	<p>Click Configure Disk Pool to create pools of disk volumes for backup by one or more media servers.</p> <p>The wizard appears if a NetBackup Data Protection Optimization Option license is installed.</p>
Volume Configuration Wizard	<p>Click Configure Volumes to configure removable media to use for backups.</p>
Catalog Backup Wizard	<p>Click Configure the Catalog Backup to set up catalog backups. Catalog backups are essential to recover data in the case of a server failure or crash.</p>

Table 1-3 Configuration wizards (*continued*)

Wizard	Description
Policy Configuration Wizard	<p>Click Create a Policy to set up one of the policy types:</p> <ul style="list-style-type: none">■ A policy to back up file systems, databases, or applications.■ A policy to protect VMware or Hyper-V virtual machines.■ A policy to back up NDMP hosts.■ A policy to back up Oracle data. <p>Note: Do not use the Policy Configuration Wizard to configure policies for Replication Director.</p>
Import Images Wizard	<p>This wizard is available only on computers running Windows. Click Import Images to import NetBackup images in a two-part process.</p> <p>See “About importing backup images” on page 970.</p>
Catalog Recovery Wizard	<p>Click Recover the Catalog to recover a catalog backup in a disaster recovery situation.</p> <p>For more information about NetBackup disaster recovery, see the NetBackup Troubleshooting Guide.</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 tech note 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 in the Activity Monitor”](#) on page 1022.

See [“About the Jobs tab”](#) on page 1020.

See [“Copying report text to another document”](#) on page 1060.

See [“Printing a report”](#) on page 1061.

See [“Saving or exporting a report”](#) on page 1061.

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: Restart the **NetBackup Administration Console** after any license updates.

Note: Perform a manual 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 883.

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

- | | |
|-------------------------------------|--|
| Add a new license. | Adding NetBackup license keys |
| Print a license. | Printing a NetBackup license key list |
| Delete a license. | Deleting NetBackup license keys |
| View the properties of one license. | Viewing NetBackup license key properties |

Export the license list. [Exporting NetBackup license keys](#)

On computers running UNIX, the license key utility is a menu-driven interface and can also be used to access licenses.

[Starting the NetBackup license key utility on a UNIX computer](#)

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.
All registered license keys details	<p>Displays the details of the license keys that are registered on this server.</p> <p>The view lists the following:</p> <ul style="list-style-type: none"> ■ Each license key ■ The server where the key is registered ■ When the key was registered, ■ The features that the key provides

Adding NetBackup license keys

Use the following procedure to add NetBackup license keys to the selected server.

To add new license keys to the selected server

- 1 In the toolbar in the **NetBackup Administration Console**, click **Help > License Keys**.
- 2 Click **New**.

- 3 Enter the license key and click **Add** or **OK**.

If you added a license on a computer running UNIX, restart all of the NetBackup utilities. Restart the **NetBackup Administration Console** as well.

- 4 Perform a manual 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 883.

You can add license keys to a remote NetBackup server. In the **NetBackup Administration Console**, select **File > Change Server**, then select another server and click **OK**. Click **Help > License Keys** in the remote server.

Printing a NetBackup license key list

Use the following procedure to print a list of license keys that are installed on a NetBackup server.

To print a list of license keys

- 1 In the toolbar in the **NetBackup Administration Console**, click **Help > License Keys**.
- 2 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

- 3 In the **NetBackup License Keys** dialog box, click the **Print** button.
- 4 Make the print selections and click **OK**.

Deleting NetBackup license keys

Use the following procedure to delete a license key that is installed on a NetBackup server.

To delete a license key

- 1 In the toolbar in the **NetBackup Administration Console**, click **Help > License Keys**.
- 2 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.
- 3 In the **NetBackup License Keys** dialog box, click the **Delete** button.
- 4 Click **OK** or **Yes** to delete the key and all features that are associated with the key. The license key cannot be restored.

If you deleted a license on a computer running UNIX, restart all of the NetBackup utilities. Restart the **NetBackup Administration Console** as well.

Viewing NetBackup license key properties

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

To view the properties of a NetBackup license key

- ◆ In the toolbar in the **NetBackup Administration Console**, click **Help > License Keys**.

Select one license and click the **Properties** button.

Exporting NetBackup license keys

Use the following procedure to export NetBackup license keys.

To export NetBackup license keys

- 1 In the toolbar in the **NetBackup Administration Console**, click **Help > License Keys**.
- 2 In the **NetBackup License Keys** dialog box, click the **Export** button.
- 3 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.
- 4 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

Starting the NetBackup license key utility on a UNIX computer

Use the following procedure to start the NetBackup license key utility on a computer running UNIX.

To start the NetBackup license key utility on a computer running UNIX

- 1 From the command line, run
`/usr/opensv/netbackup/bin/admincmd/get_license_key`
- 2 At the prompt, enter one of the following menu selections, then press `Enter`:
 - Type `A` to add a new license key, then type the license key at the prompt.
 - Type `D` to delete a license from the list, then type the license key at the prompt.
 - Type `F` to list only the licenses that are currently active. Expired licenses do not appear in this listing. Specify a local or a remote host.
 - Type `L` to list all registered licenses—active or inactive. Specify a local or a remote host.
 - Type `H` for help on the License Key Utility.
 - Type `q` to quit the utility.

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:

- [About the NetBackup Host Properties](#)
- [Access Control properties](#)
- [Active Directory host properties](#)
- [Bandwidth properties](#)
- [Busy File Settings properties](#)
- [Clean-up properties](#)
- [Client Name properties](#)
- [Client Attributes properties](#)
- [Client Settings properties for UNIX clients](#)
- [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](#)
- [Network properties](#)
- [Network Settings properties](#)
- [Port Ranges properties](#)
- [Preferred Network properties](#)
- [Resilient Network properties](#)
- [Resource Limit properties](#)
- [Restore Failover properties](#)
- [Retention Periods properties](#)
- [Scalable Storage properties](#)
- [Servers properties](#)
- [SharePoint properties](#)
- [SLP Parameters properties](#)
- [Throttle Bandwidth properties](#)
- [Timeouts properties](#)
- [Universal Settings properties](#)
- [UNIX Client properties](#)
- [UNIX Server properties](#)

- VMware Access Hosts properties
- Windows Client properties
- Configuration options not found in the Host Properties
- About using commands to change the configuration options on UNIX clients and servers
- Configuration options for NetBackup servers
- Configuration options for NetBackup clients

About the NetBackup Host Properties

The configuration options within the **Host Properties** let an administrator customize NetBackup to meet specific site preferences and requirements.

In most NetBackup environments, the **Host Properties** default settings provide satisfactory results and can be left at the default setting.

Figure 3-1 shows the **Host Properties** in the **NetBackup Administration Console**.

Figure 3-1 Host Properties utility in the NetBackup Administration Console

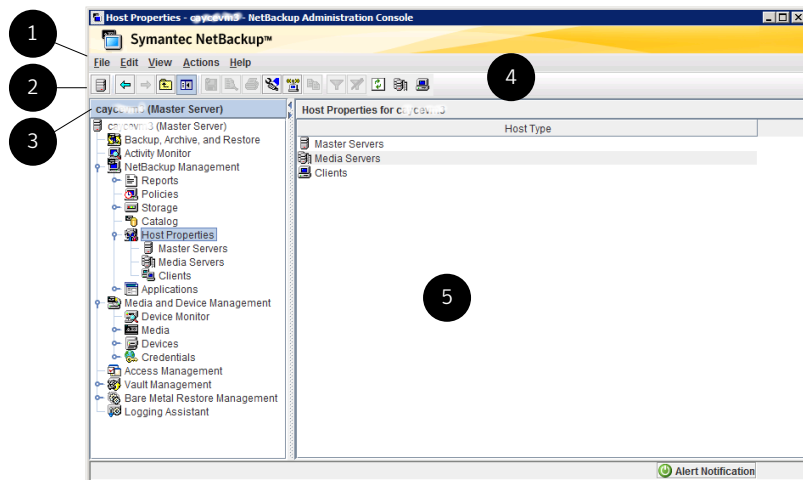


Table 3-1 Host Properties utility

Number	Description
1	The menu toolbar.

Table 3-1 Host Properties utility (*continued*)

Number	Description
2	The standard NetBackup toolbar. See “Standard and user toolbars” on page 45.
3	The name of the currently selected master server.
4	The user toolbar is specific to the Host Properties utility. By default, the Host Properties user toolbar includes the following buttons: <ul style="list-style-type: none"> ■ Select a host in the right pane and click Connect to connect to that host. ■ Select a host in the right pane and click Properties to view the properties of that host. See “Connecting to a host to view the host properties” on page 62. ■ Click Configure media server to name a media server and view its properties. ■ Click Configure Client to name a client and view its properties. You can select which buttons appear on the user toolbar. See “Standard and user toolbars” on page 45.
5	Right-click in the right pane to view the shortcut menu.

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 227.

See [“Allow access to another server”](#) on page 1103.

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 249.

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

See [“Configuration options not found in the Host Properties”](#) on page 257.

Methods to set the NetBackup configuration options

A NetBackup administrator can use one of the following methods to read or set the default configuration options or to change those options that do not appear in the **Host Properties**:

Table 3-2 NetBackup Host Properties configuration methods

Method	Description
NetBackup Administration Console interface	<p>Most properties are listed in the NetBackup Administration Console in the Host Properties. To navigate to the various properties, select NetBackup Management > Host Properties. Depending on the host to be configured, select Master Servers, Media Servers, or Clients.</p> <p>See “About the NetBackup Host Properties” on page 60.</p>
Command line	<p>Use the <code>nbgetconfig</code> command or <code>bpgetconfig</code> command to obtain a list of configuration entries. Then use <code>nbsetconfig</code> or <code>bpsetconfig</code> to change the options as needed.</p> <p>These commands update the appropriate configuration files on both Windows (registry) and UNIX (<code>bp.conf</code> file) master servers and clients.</p> <p>Use the <code>nbemmcmd</code> command to modify some options on hosts.</p> <p>Detailed information on these commands is available in the NetBackup Commands Reference Guide.</p>
<code>vm.conf</code> file	<p>The <code>vm.conf</code> file contains configuration entries for media and device management.</p> <p>See the NetBackup Administrator's Guide, Volume II for more information.</p>
Backup, Archive, and Restore client interface	<p>Administrators can specify configuration options for NetBackup clients.</p> <p>See the NetBackup Backup, Archive, and Restore Getting Started Guide.</p>

Connecting to a host to view the host properties

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

Use the following procedure to connect and view host properties of master server, media server, or client.

To connect and view host properties of master server, media server, or client

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties**.
- 2 Select **Master Servers**, **Media Servers**, or **Clients**.
- 3 In the right pane, perform one of the following:
 - Highlight the host by selecting the row. The host is connected after a delay of approximately 3 seconds. If you select multiple rows, connection to the hosts is not established.
 - Right-click and select **Connect**. If you select multiple rows right-click and select **Connect** by pressing the **Ctrl** key, connection to all the hosts is established.

Note: When you double-click a host that is not connected to view its host properties, an intermittent **Loading Data** dialog box is displayed. This dialog box contains an **Interrupt** button. When you click the **Interrupt** button, the **Loading Data** dialog box closes. The connection to the host continues in the background. However, the **Host Properties** dialog box is not displayed.

To see the properties of a different master server, click **File > Change Server**.

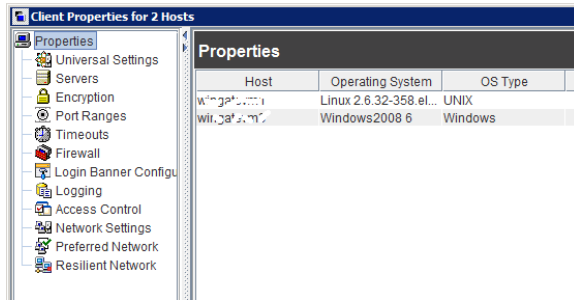
Changing the host properties on multiple hosts at the same time

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

The host property check boxes may appear in one of the following states:

- 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

If the property contains a text field for specifying a value, the field may be in one of the following states:

- 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. In this case, when the cursor is moved 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 focus is on a setting that is set differently between the selected hosts, the following statement appears at the bottom of the dialog box:

This value is different on the selected hosts.

This notice is especially helpful regarding differences in text field settings.

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 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.

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** dialog box, name the file, browse to the directory where you want to save it, and click **Save**.

Standard host property dialog box options

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

Defaults

Click **Defaults** to set all the properties in the current dialog box to the default values.

Note: If you have enabled the Logging for critical processes option on the **Logging** properties dialog box, clicking the **Defaults** button does not modify the Logging for critical processes and Keep logs up to GB options.

While changing the logging properties for multiple hosts, if the Logging for critical processes option on one of the hosts is enabled, clicking the **Defaults** button does not modify the Logging for critical processes and Keep logs up to GB options of any host.

See [“Logging properties”](#) on page 161.

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).

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 of the Access Control properties](#)” on page 68.
- Authorization Service tab
See “[Authorization Service tab of the Access Control properties](#)” on page 69.
- Network Attributes
See “[Network Attributes tab of the Access Control properties](#)” on page 70.

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 3-3 NetBackup Product Authentication and Authorization property options

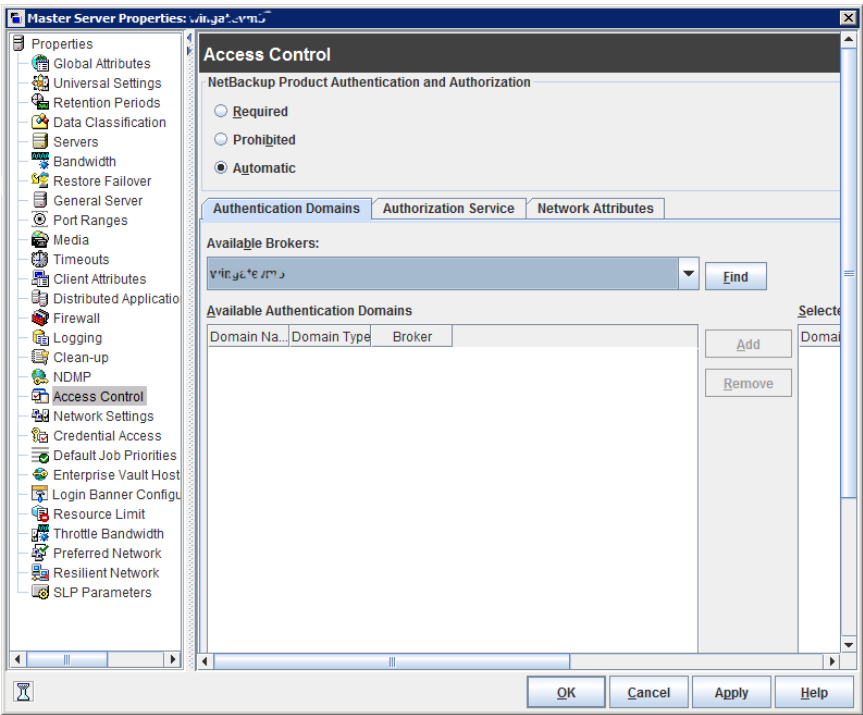
Option	Description
Required	Specifies that the local system should accept requests only from the remote systems that use NetBackup authentication and authorization. Connections from the remote systems that do not use NetBackup authentication and authorization are rejected. Select Required if all systems 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 of the Access Control properties

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 3-4 Authentication Domain tab properties

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

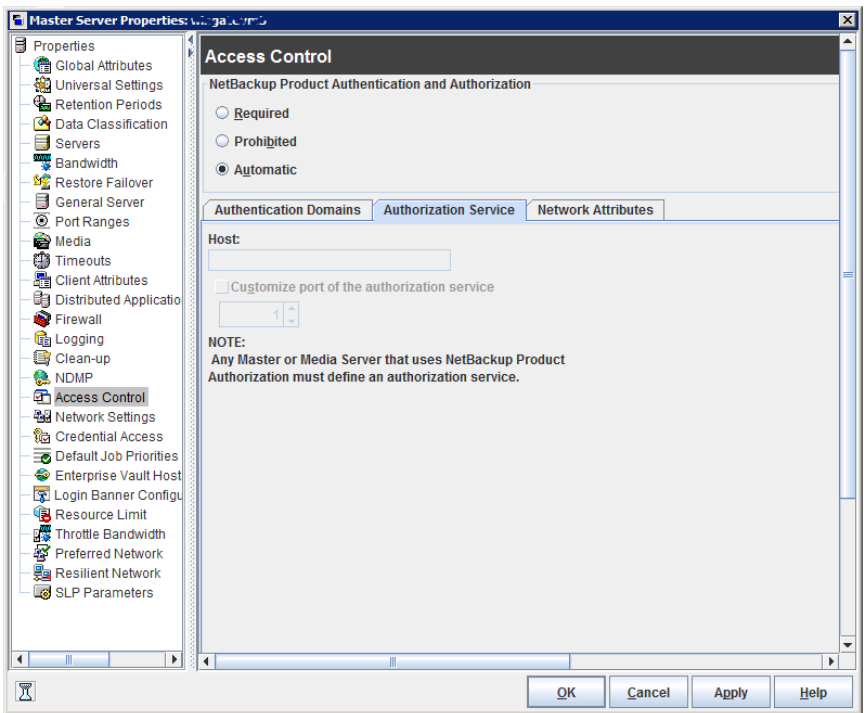
Table 3-4 Authentication Domain tab properties (continued)

Property	Description
Available Authentication Domains list	List of available authentication domains.
Add button	Select the authentication domain(s) that this host can use and click Add .
Selected Authentication Domains list	List of the authentication domains that are 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 of the Access Control properties

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 3-5 Authorization Service property options

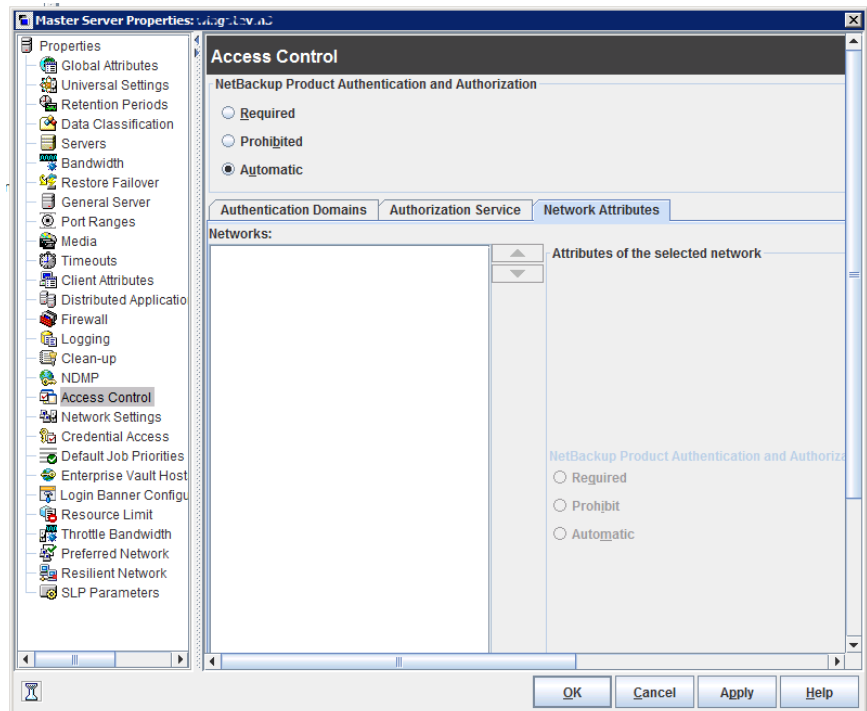
Option	Description
Host name	Specifies the host name or IP address of the authorization service.
Customize the port number of the authorization service	Specifies a nonstandard port number. Select Customize the port number and enter the port number of the authorization service.

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

Network Attributes tab of the Access Control properties

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 67.

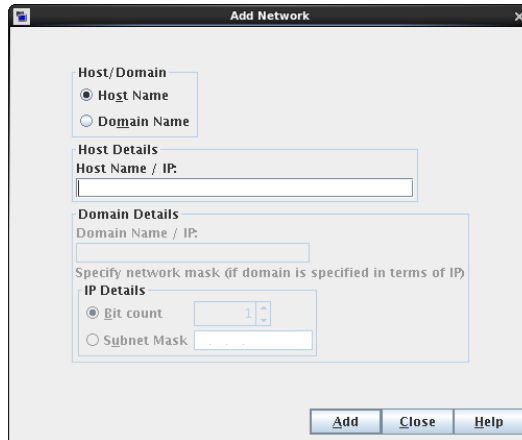
Add Network dialog box

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

Table 3-6 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	<div>Specifies that if the network is a host, one of the following items must be entered:</div> <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 that is 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

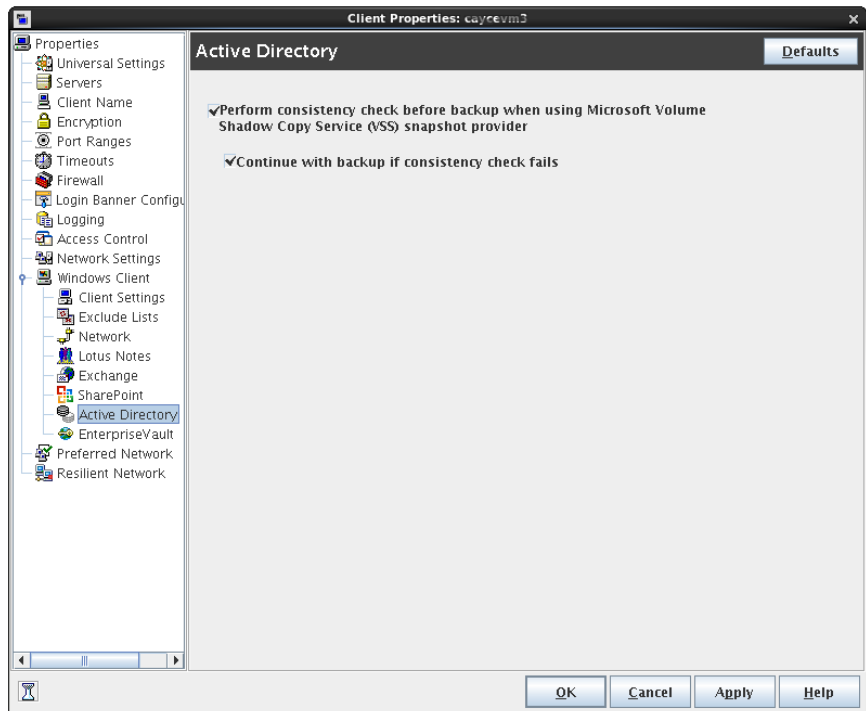
The image shows a Windows-style dialog box titled "Add Network". It contains several sections: "Host/Domain" with radio buttons for "Host Name" (selected) and "Domain Name"; "Host Details" with a text field for "Host Name / IP:"; "Domain Details" with a text field for "Domain Name / IP:"; and "IP Details" with a note "Specify network mask (if domain is specified in terms of IP)" and two options: "Bit count" (selected) with a spin box set to "1", and "Subnet Mask" with a text field. At the bottom right are three buttons: "Add", "Close", and "Help".

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 844.

Figure 3-6 Active Directory dialog box



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

Table 3-7 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 93.
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.

Bandwidth properties

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

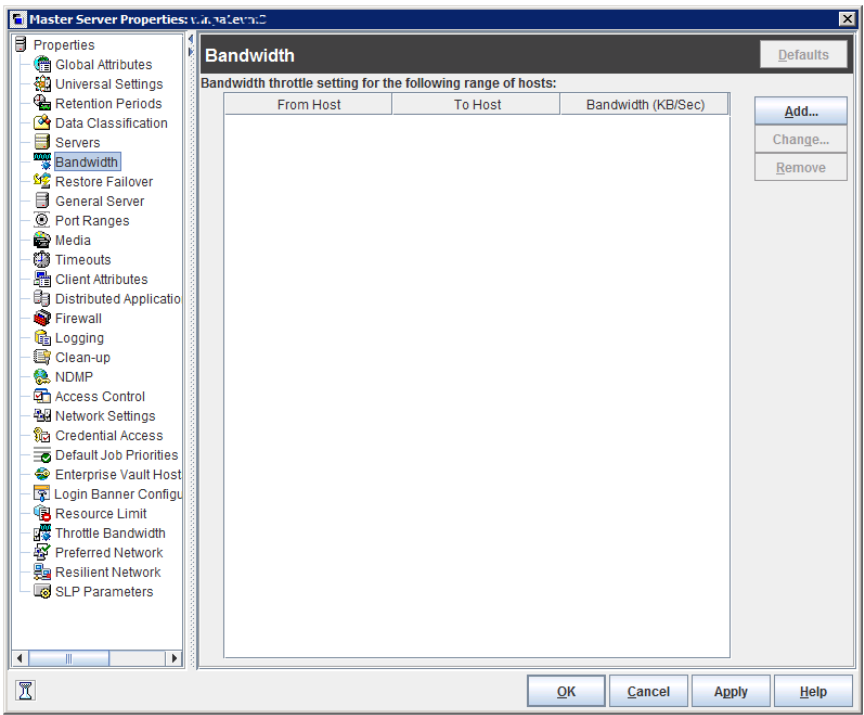
Note: The **Bandwidth** properties apply only to IPv4 networks. Use the **Throttle Bandwidth** properties to limit IPv6 networks.

See “[Throttle Bandwidth properties](#)” on page 245.

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 3-7 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 for each new job based on the number of jobs that are currently running for the IP range. The bandwidth value that is assigned to the job at start time remains the same until the job is completed.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.

Bandwidth limit usage considerations and restrictions

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

Table 3-8 Bandwidth limit usage considerations and restrictions

Client or operation	Bandwidth limit behavior or restrictions
<div><div>■ DataTools SQL-BackTrack</div><div>■ DB2</div><div>■ Enterprise-Vault</div><div>■ FlashBackup</div><div>■ FlashBackup-Windows</div><div>■ Lotus-Notes</div><div>■ MS-SharePoint</div><div>■ MS-SQL-Server</div><div>■ NCR-Teradata</div><div>■ NDMP</div><div>■ SAP</div><div>■ Vault</div></div>	
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.

Table 3-8 Bandwidth limit usage considerations and restrictions (*continued*)

Client or operation	Bandwidth limit behavior or restrictions
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.
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

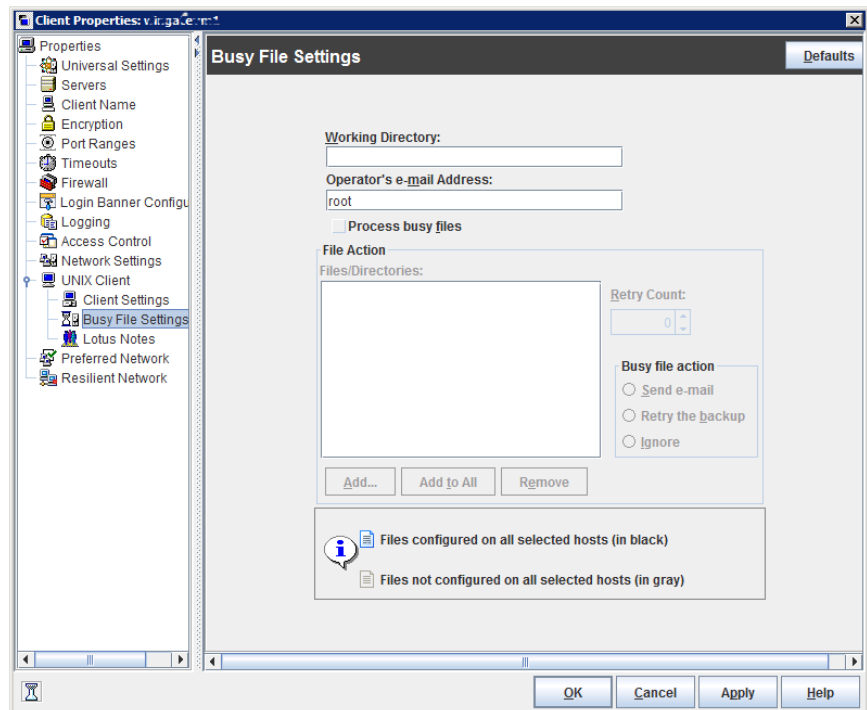
The **Add Bandwidth Settings** and the **Change Bandwidth Settings** dialog boxes contain the following properties.

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 3-8 Busy File Settings dialog box



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

Table 3-9 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 changes during a backup. By default, Process busy files is not enabled and NetBackup does not process the busy files. Additional information about busy file processing is available in the NetBackup Administrator's Guide, Volume II .

Table 3-9 Busy File Settings dialog box properties (*continued*)

Property	Description
File action file list	Specifies the absolute path and file name of the busy file. The metacharacters *, ?, [], [-] can be used for pattern matching of file names or parts of file names.
Add	Adds a new file entry. Enter the file and path directly, or browse to select a file.
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	<p>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.</p> <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
```

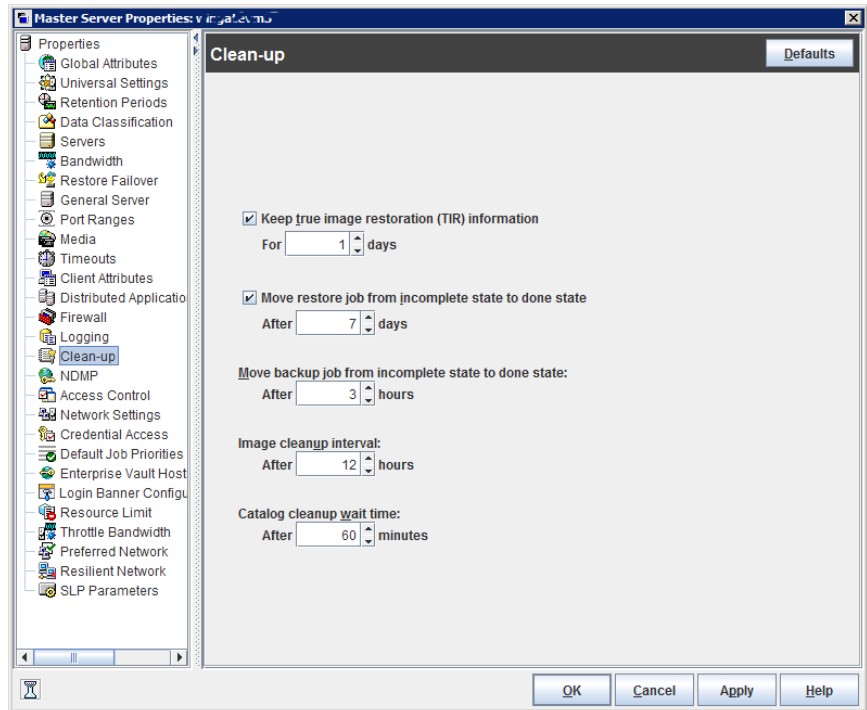
- 2 Set the file access permissions to allow group and others to run `bpend_notify`.
- 3 Configure a policy with a user backup schedule for the busy file backups.

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 3-9 Clean-up dialog box



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

Table 3-10 Clean-up dialog box properties

Property	Description
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 following directories:</p> <p>On Windows:</p> <pre>install_path\NetBackup\db\images</pre> <p>On UNIX:</p> <pre>/usr/opensv/netbackup/db/images</pre> <p>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>
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 249.</p> <p>See “Checkpoint restart for restore jobs” on page 692.</p>

Table 3-10 Clean-up dialog box properties (continued)

Property	Description
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>
Image cleanup interval	<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>

Note: The following properties have been moved from the **Clean-up** host properties to the **Logging** host properties: **Keep logs** and **Keep Vault logs**. On the **Logging** properties screen, these properties are referred to as **Keep logs for days** and **Keep Vault logs for** respectively.

See [“Logging properties”](#) on page 161.

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 host name instead of an IPv6 address.

See “[Client Attributes properties](#)” on page 84.

Figure 3-10 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 3-11 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 .

Table 3-11 Global client attributes group box (*continued*)

Attribute	Description
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 directories:</p> <p>On Windows: <code>Install_path\NetBackup\db\client</code></p> <p>On UNIX: <code>/usr/opensv/NetBackup/db/client</code></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 86.</p> <p>The name that is entered here must match the Client Name property for the specific client. If it does not, the client cannot browse its own backups.</p> <p>See “Client Name properties” on page 82.</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 busy file processing is available in the NetBackup Administrator's Guide, Volume II.</p> <p>On UNIX: You also can create, update, list, and delete client entries by using the <code>bpclient</code> command that is located in the following directory:</p> <p><code>/usr/opensv/netbackup/bin/admincmd</code></p>
General tab	<p>Specifies how to configure the selected Windows master servers (clients).</p> <p>See “General tab of the Client Attributes properties” on page 86.</p>
Connect Options tab	<p>Specifies how to configure the connection between a NetBackup server and a NetBackup client.</p> <p>See “Connect Options tab of the Client Attributes properties” on page 91.</p>
Windows Open File Backup tab	<p>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.</p> <p>See “Windows Open File Backup tab of the Client Attributes properties” on page 93.</p>

Add Client dialog box

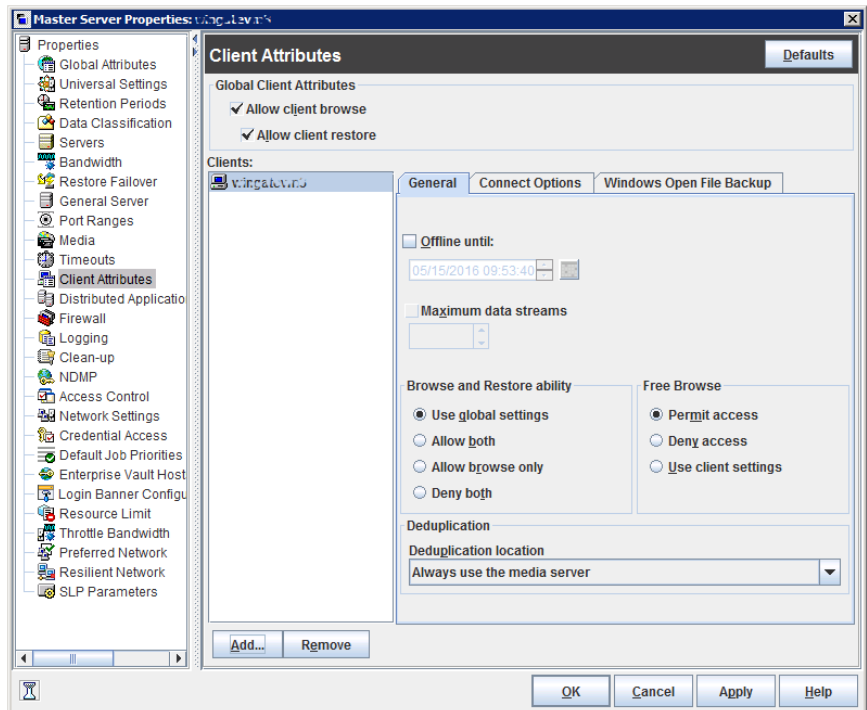
To add a client to the database, enter the name of a client, or browse to find a client. The **Add Client** dialog box contains the following properties.

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 3-11 General tab of Client Attributes dialog box



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

Table 3-12 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 (approximately January 18, 2038), depending on the locale setting.</p> <p>See “Retention Periods with end dates beyond 2038, excluding Infinity” on page 222.</p> <p>Note: Changes to this property do not appear in the audit report.</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 89.</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 150.</p> <p>See “Limit jobs per policy (policy attribute)” on page 693.</p>

Table 3-12 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>Specifies whether the clients can list and restore from scheduled backups. (This setting does not affect user backups and archives.)</p> <p>This property applies to the privileges that are allowed to a non-Windows administrator or non-root user who is logged into the client. This property also applies to the users that do not have backup and restore privileges.</p> <p>Windows administrators and root users 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 the NetBackup Data Protection Optimization Option.</p> <p>For a description of the client-side deduplication options and their actions: See “Where deduplication should occur” on page 90.</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 3-13 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.

Table 3-13 Offline option actions (*continued*)

Type of job or operation	Action or restriction
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 (Client is offline). 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 Data Protection Optimization Option. [Table 3-14](#) describes the client-side deduplication options.

The clients that deduplicate their own data and the master server must use the same name to resolve the storage server. The name must be the host name under which the NetBackup Deduplication Engine credentials were created. If they do not use the same name, backups fail. In some environments, careful configuration may be required to ensure that the client and the master server use the same name for the storage server. Such environments include those that use VLAN tagging and those that use multi-homed hosts.

NetBackup does not support the following for client-side deduplication:

- Multiple copies per job configured in a NetBackup backup policy. For the jobs that specify multiple copies, the backup images are sent to the storage server and may be deduplicated there.
- NDMP hosts. The backup jobs fail if you try to use client-side deduplication for NDMP hosts.

Table 3-14 Client-side 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 is true: <ul style="list-style-type: none">■ The deduplication services on the storage server are inactive.■ The deduplication pool is down.
Prefer to use client-side deduplication	Deduplicates the data on the client and then sends it directly to the storage server. NetBackup first determines if the storage server is active. If it is active, the client deduplicates the backup data and sends it to the storage server to be written to disk. If it is not active, the client sends the backup data to a media server, which deduplicates the data.
Always use client-side deduplication	Always deduplicates the backup data on the client and then sends 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.

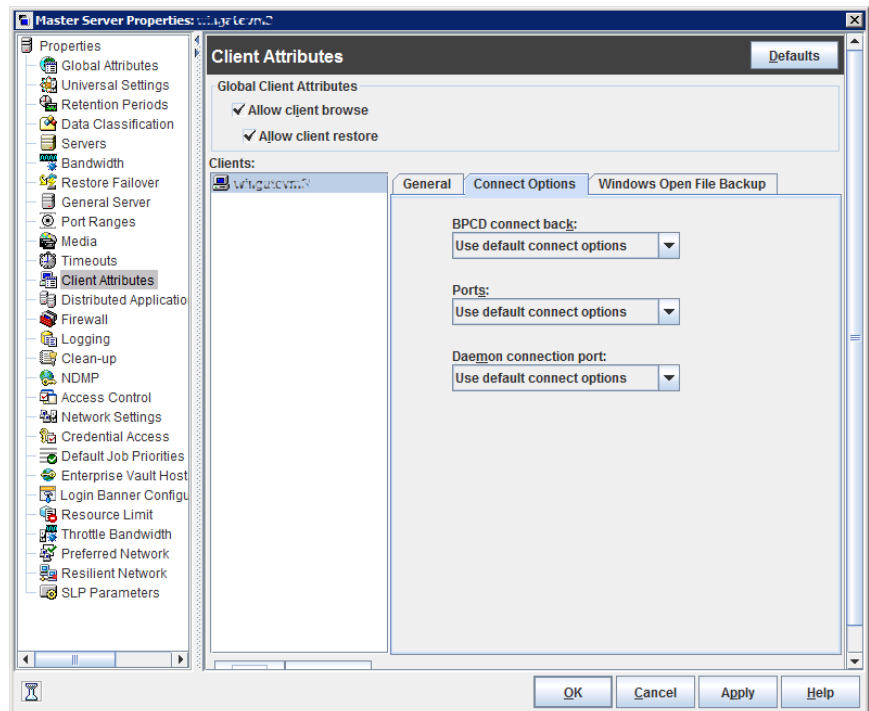
See “[Disable client-side deduplication \(policy attribute\)](#)” on page 715.

More information about client deduplication is available in 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 3-12 Connect Options tab of Client Attributes dialog box



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

Table 3-15 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 142. ■ 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 3-15 Connect Options tab properties (*continued*)

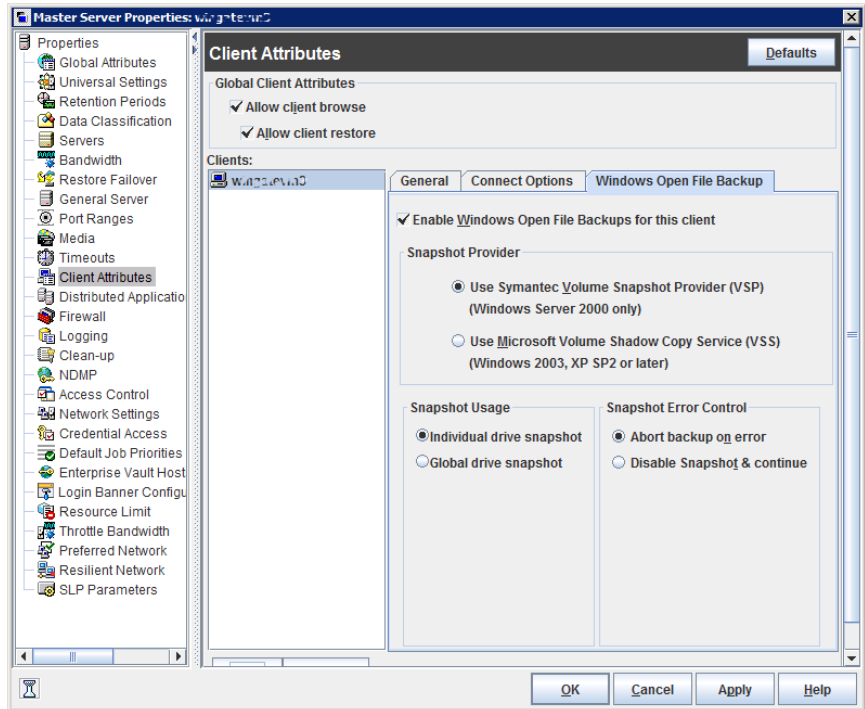
Property	Description
Ports	<p>Specifies the method that the selected clients should use to connect to the server and contains the following options:</p> <ul style="list-style-type: none"> ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See "Firewall properties" on page 142. ■ 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. ■ See "Resilient Network properties" on page 209.

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 3-13 Windows Open File Backup tab of Client Attributes dialog box



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

Table 3-16 Windows Open File Backup tab properties

Property	Description
Add	<p>Adds the 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 “Back-level and upgraded clients that use Windows Open File Backup” on page 98. ■ 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>
Enable Windows Open File Backups	<p>Specifies that Windows Open File Backups be used for the selected clients. Adds the clients to the list only if you want to change the default property settings.</p> <p>For Microsoft Distributed File System Replication (DFS) servers, select Enable Windows Open File Backups. The Snapshot Provider must be Volume Shadow Copy Service.</p> <p>See “About Microsoft DFSR backups” on page 670.</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 NetBackup Snapshot Client Administrator's Guide.</p>

Table 3-16 Windows Open File Backup tab properties (*continued*)

Property	Description
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) This option specifies that Microsoft VSS be used to create volume snapshots of volumes and logical drives for the selected clients. Configure VSS through the Microsoft VSS configuration dialog boxes. For information about how to do Active Directory granular restores when using VSS, see the following topic: See “Active Directory host properties” on page 73. For Microsoft Distributed File System Replication (DFS) servers, select Use Microsoft Volume Shadow Copy Service (VSS). See “About Microsoft DFS backups” on page 670.

Table 3-16 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"> Individual drive snapshot Specifies that the snapshot should be of an individual drive (default). When this property is enabled, snapshot creation and file backup are done sequentially on a per volume basis. For example, assume that drives C and D are to be backed up. If the Individual drive snapshot property is selected, NetBackup takes a snapshot of drive C, backs it up, and discards the snapshot. It then takes a snapshot of drive D, backs it up, and discards the snapshot. Volume snapshots are enabled on only one drive at a time, depending on which drive is to be backed up. This mode is useful when relationships do not have to be maintained between files on the different drives. Use this configuration if snapshot creation fails when all volumes for the backup are snapshot at once when the Global drive snapshot property is enabled. Individual drive snapshot is enabled by default for all non-multistreamed backups by using the Windows Open File Backup option. Global drive snapshot Specifies that the snapshot is of a global drive. All the volumes that require snapshots for the backup job (or stream group for multistreamed backups) are taken at one time. For example, assume that drives C and D are to be backed up. In this situation, NetBackup takes a snapshot of C and D. Then NetBackup backs up C and backs up D. NetBackup then discards the C and D snapshots. This property maintains file consistency between files in different volumes. The backup uses the same snapshot that is taken at a point in time for all volumes in the backup. Note: The Individual drive snapshot property and the Global drive snapshot property only apply to non-multistreamed backups that use Windows Open File Backup. All multistreamed backup jobs share the same volumes snapshots for the volumes in the multistreamed policy. The volume snapshots are taken in a global fashion.

Table 3-16 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"> Abort backup on error Specifies that a backup aborts if it fails for a snapshot-related issue after the snapshot is created and while the backup uses the snapshot to back up open or active files on the file system. The most common reason for a problem after the snapshot is created and is in use by a backup, is that the cache storage is full. If the Abort backup on error property is checked (default), the backup job aborts with a snapshot error status if the backup detects a snapshot issue. This property does not apply to successful snapshot creation. The backup job continues regardless of whether a snapshot was successfully created for the backup job. The Abort backup on error property applies only to the snapshot errors that occur after the snapshot is successfully created and is in use by a backup job. Disable snapshot and continue Specifies that if the snapshot becomes invalid during a backup, the volume snapshots for the backup are destroyed. The backup continues with Windows open file backups disabled. Regarding the file that had a problem during a backup—it may be that the file was not backed up by the backup job. The file may not be able to be restored. Note: Volume snapshots typically become invalid during the course of a backup because insufficient cache storage was allocated for the volume snapshot. Reconfigure the cache storage configuration of the Windows Open File Backup snapshot provider to a configuration that best suits your client's installation.

Back-level 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 3-17 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.

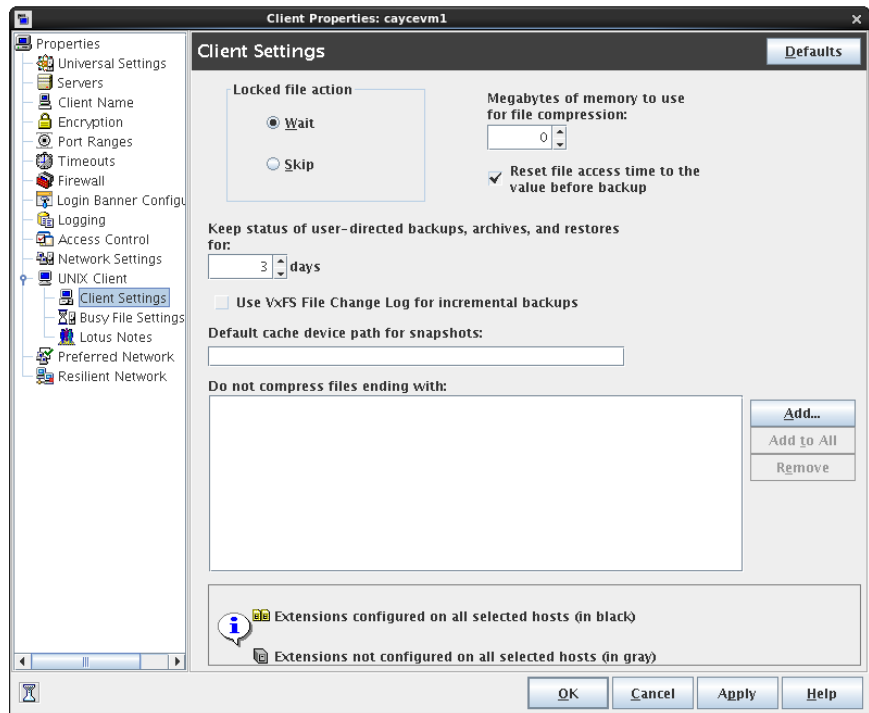
Table 3-17 Snapshot Provider behavior for clients in a 7.x environment
(continued)

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 UNIX clients

The UNIX **Client Settings** properties in the **NetBackup Administration Console** apply to currently selected NetBackup clients running on the UNIX platform.

Figure 3-14 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 247. 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.

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

Property	Description
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 3 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 to the value before backup	<p>Specifies that the access time (<code>atime</code>) 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 computer resources that are used. If other processes also need memory, use a maximum value of half the actual physical memory on a computer 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 102.</p>
Default cache device path for snapshots	<p>For additional information, see the NetBackup Snapshot Client Administrator's Guide.</p>
Do not compress files ending with list	<p>Corresponds to adding a <code>COMPRESS_SUFFIX =.suffix</code> option to the <code>bp.conf</code> file.</p> <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, <code>.A1</code> is allowed, but not <code>.A*</code> or <code>.A[1-9]</code></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>

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

Property	Description
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 later.
- AIX running VxFS 5.0 or later.
- HP 11.23 running VxFS 5.0 or later.
- Linux running VxFS 4.1 or later.

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 796.

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 one of the following ways:
 - Specify ALL_LOCAL_DRIVES.
 - Specifying the actual VxFS mount point.
 - 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 699.

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 the 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>

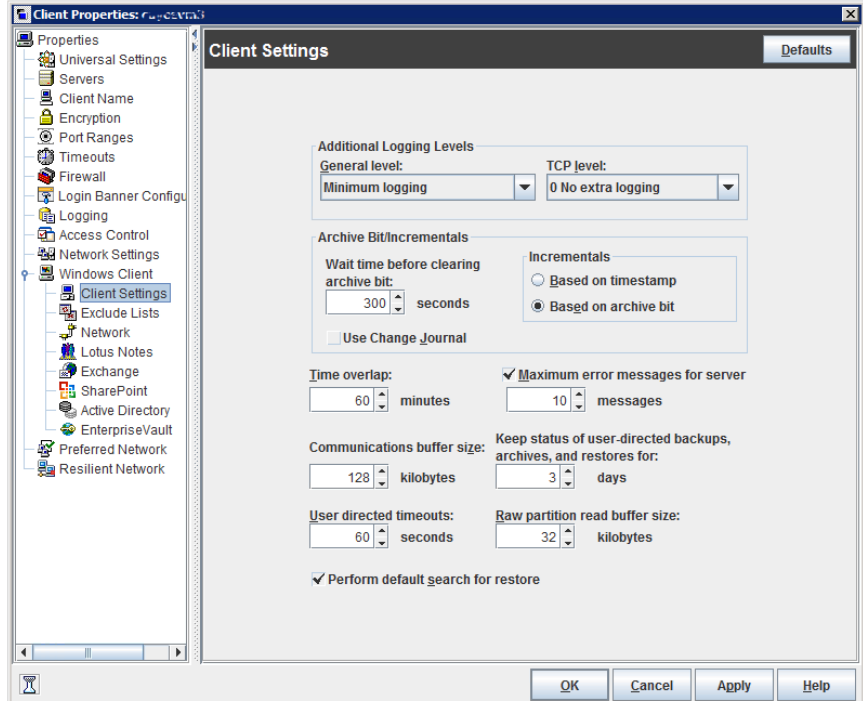
Table 3-19 VxFS file change log feature options (*continued*)

Option	Description
Keeping the data files synchronized with the FCL	The data files must be in sync with the FCL for this property to work. To keep the data files synchronized with the FCL, do not turn the FCL on the VxFS client off and on. Note: If NetBackup encounters any errors as it processes the FCL, it switches to the normal files system scan. If this switch occurs, it appears in the Activity Monitor.
VxFS administration	Additional VxFS commands are available to administrate the FCL in the Veritas File System Administrator's Guide .

Client Settings properties for Windows clients

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

Figure 3-15 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>
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	<p>Note: The Use Change Journal option applies to Windows clients only.</p> <p>This option works together with the Use Accelerator policy attribute and the Accelerator forced rescan schedule attribute.</p> <p>See “Accelerator and the Windows change journal” on page 719.</p> <p>See “Use Accelerator (policy attribute)” on page 716.</p> <p>See “Accelerator forced rescan option (schedule attribute)” on page 755.</p>
Incrementals based on timestamp	Specifies that files are selected for the backups that are based on the date that the file was last modified. When Use change journal is selected, Incrementals based on timestamp is automatically selected.

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 the IP buffers that NetBackup uses 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 128 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	<p>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.</p> <p>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.</p>

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:

- 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 692.
- Support for change journal is not offered with several NetBackup options or Symantec products.

If **Use change journal** is enabled, it has no effect while you use the following options or products:

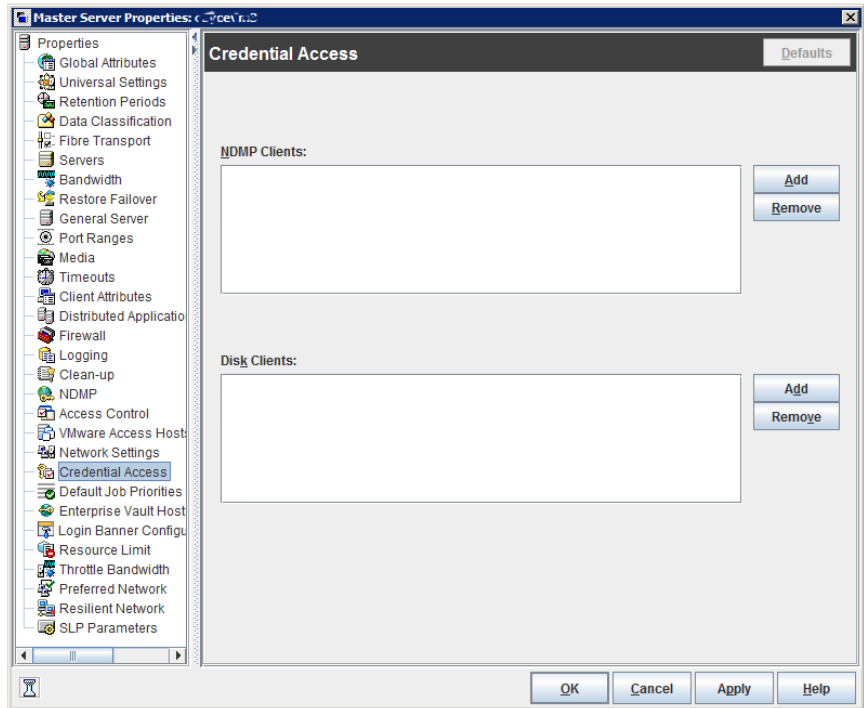
- True image restore (TIR) or True image restore with Move Detection
See [“Collect true image restore information \(policy attribute\) with and without move detection”](#) on page 707.
- Synthetic backups
See [“About synthetic backups”](#) on page 850.
- Bare Metal Restore (BMR)
For more information, 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 107.

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 3-16 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.

Table 3-21 Credential Access dialog box properties (continued)

Property	Description
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 lifecycle policies can be configured.

See “[Data classifications \(policy attribute\)](#)” on page 684.

Figure 3-17 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. This column is not displayed by default.</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 lifecycle policies 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.

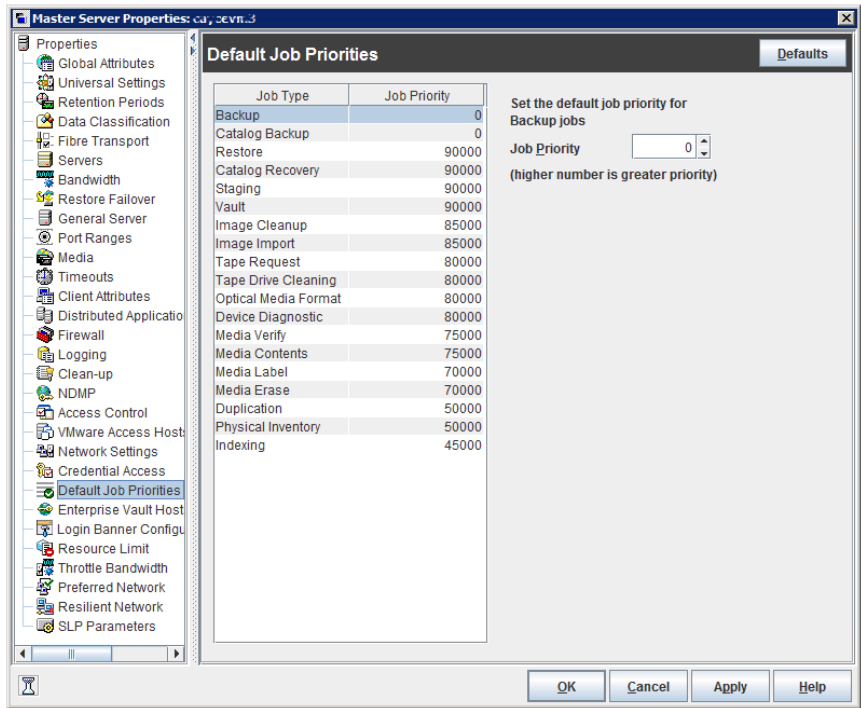
Note: Data classifications cannot be deleted.

- 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 3-18 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 from the Activity Monitor”](#) on page 1026.
- 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.

Table 3-23 Default Job Priorities dialog box properties (*continued*)

Property	Description
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.

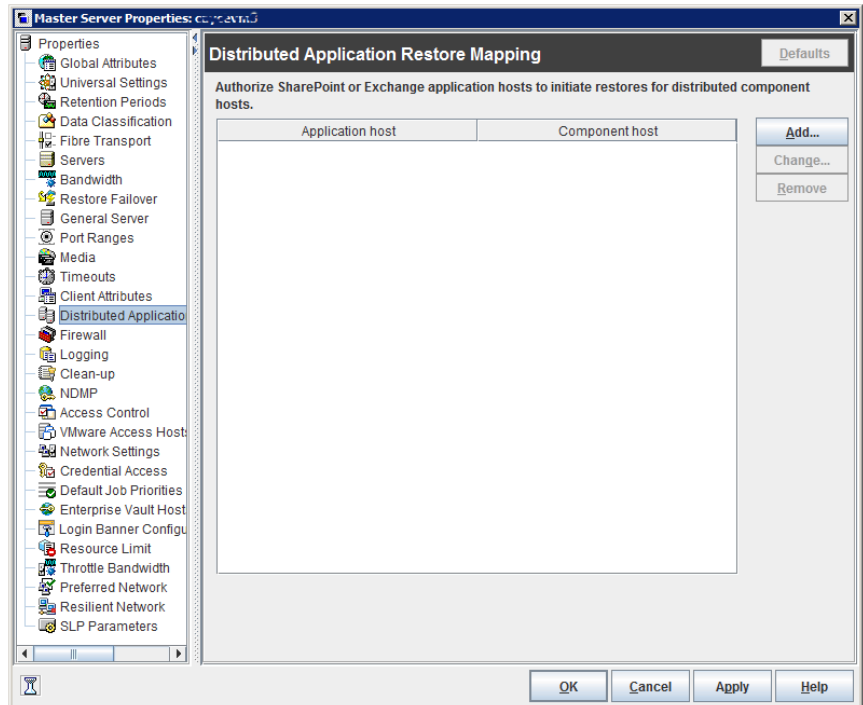
Because a request with a higher priority is listed in the queue before a request with a lower priority, the request with a higher priority is evaluated first. Even though the chances are greater that the higher priority request receives resources first, it is not always definite.

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

Figure 3-19 Distributed Application Restore Mapping dialog box



Some applications, such as SharePoint and Exchange, distribute and replicate data across multiple hosts. You need to provide a mapping of the hosts in the database environment so that NetBackup can successfully restore the databases.

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 component host that is authorized to run restores on an SharePoint or Exchange application host.</p> <p>For SharePoint, NetBackup catalogs backup images under the front-end server name. To allow NetBackup to restore SQL databases to the correct hosts in a farm, you must provide a list of the SharePoint hosts. For Exchange, any operations that use Granular Recovery Technology (GRT) require that you provide a list of the Exchange virtual and the physical host names. Also include the off-host client and the granular proxy host.</p> <p>Note: For VMware backups and restores that protect SharePoint or Exchange, you only need to add the hosts that browse for backups or perform restores.</p> <p>Note: You can use the client's short name or its fully qualified domain name (FQDN). You do not need to provide both names in the list.</p> <p>For more details, see the following:</p> <p>NetBackup for SharePoint Server Administrator's Guide</p> <p>NetBackup for Exchange Server Administrator's Guide</p>
Change	Changes the application host or component host of the currently selected mapping.
Remove	Removes the currently selected mapping.

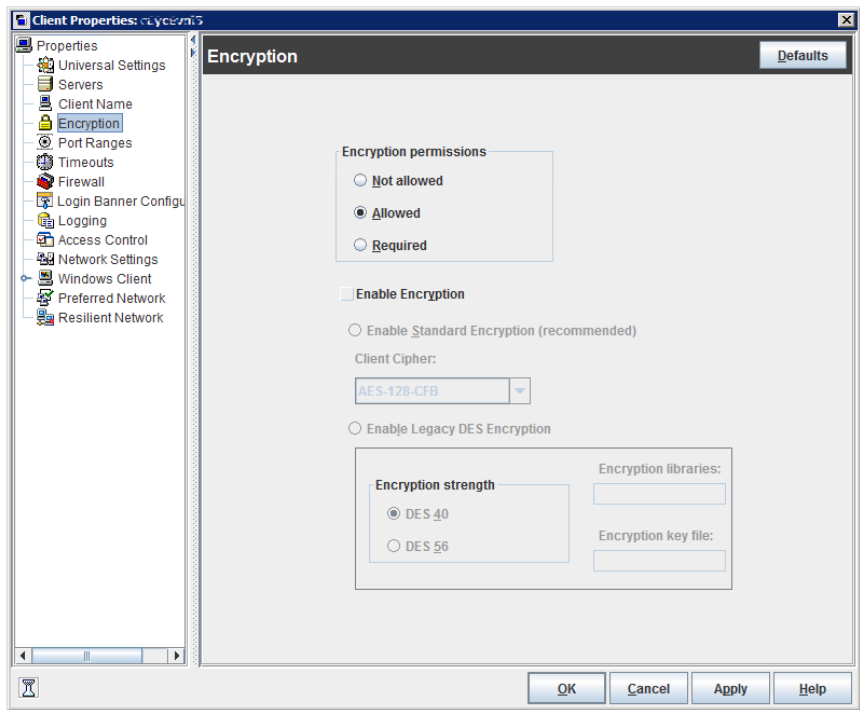
Encryption properties

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.

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 about the ciphers file is available 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: <code>install_path\netbackup\bin\</code> Where <code>install_path</code> is the directory where NetBackup is installed and by default is <code>C:\Program Files\VERITAS</code>.■ On UNIX: <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.</p> <p>The default location is as follows:</p> <ul style="list-style-type: none">■ On Windows: <code>install_path\NetBackup\bin\keyfile.dat</code> Where <code>install_path</code> is the folder where NetBackup is installed and by default is C:\Program Files\VERITAS.■ On UNIX: <code>/usr/opensv/netbackup/keyfile</code> <p>If it is necessary to change the setting, specify the new name.</p>

Additional encryption methods for Windows clients

In addition to NetBackup client and server data encryption, Microsoft Windows clients also have access to methods of encrypting the data on the original disk.

Each of the following methods has its own costs and benefits. NetBackup supports each method for protecting Microsoft Windows clients.

Encrypting File System

The Encrypting File System (EFS) on Microsoft Windows is a feature introduced with Windows 2000 and provides file system-level encryption. EFS is a form of encryption where individual files or directories are encrypted by the file system itself.

The technology enables files to be transparently encrypted to protect confidential data from attackers with physical access to the computer. Users can enable encryption on a per-file, per-directory, or per-drive basis. The Group Policy in a Windows domain environment can also mandate some EFS settings.

No NetBackup settings are involved in protecting these encrypted objects. Any object with an encrypted file system attribute is automatically backed up and restored in its encrypted state.

BitLocker Drive Encryption

BitLocker Drive Encryption is a full disk encryption feature included with certain editions of Microsoft's Windows desktop and server versions, starting with Vista and Windows Server 2008.

Disk encryption is a technology which protects information by converting it into unreadable code that cannot be deciphered easily by unauthorized people. Disk encryption uses disk encryption software or hardware to encrypt every bit of data that goes on a disk or a disk volume.

As with EFS, no NetBackup settings are involved to use BitLocker for encryption. Unlike EFS, the encryption layer is invisible to NetBackup, with the data being automatically decrypted and encrypted by the operating system.

NetBackup does nothing to manage the encryption process and therefore backs up and restores the unencrypted data.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must re-enable BitLocker encryption following the restore.

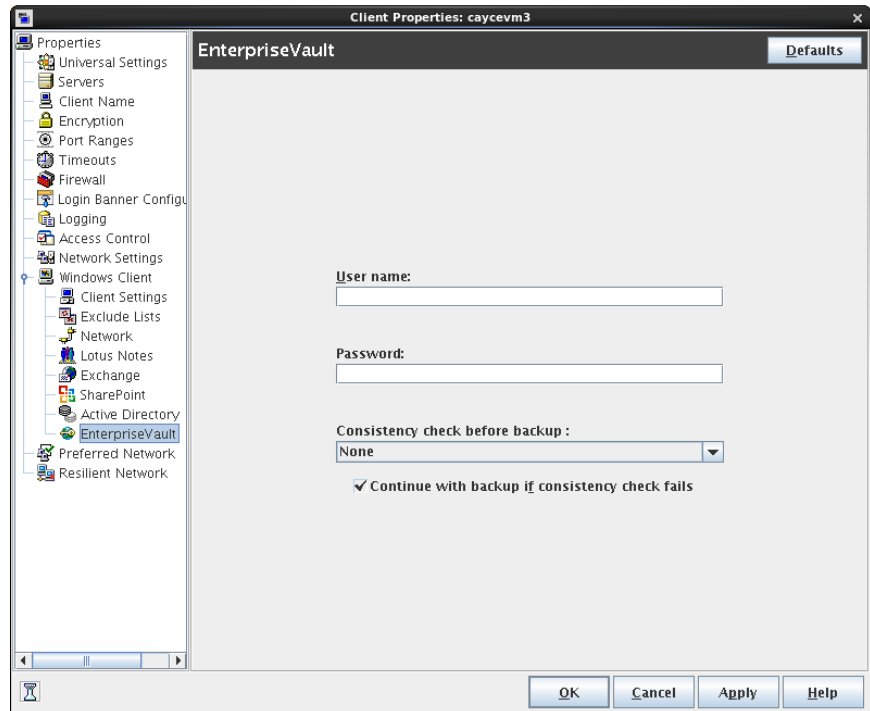
Off-host backup is not supported with volumes that run Windows BitLocker Drive Encryption.

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.

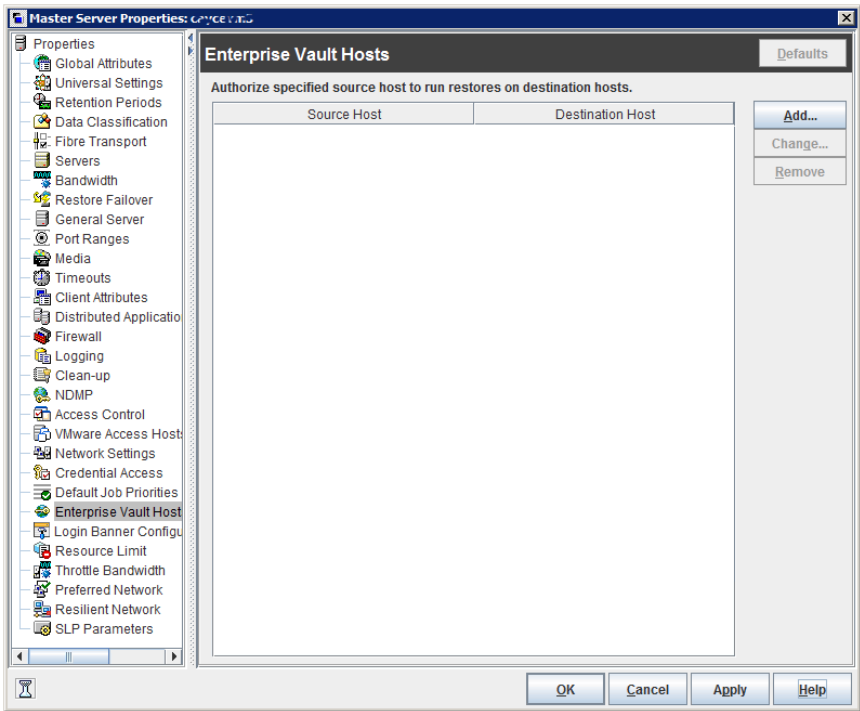
Enterprise Vault Hosts properties

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.

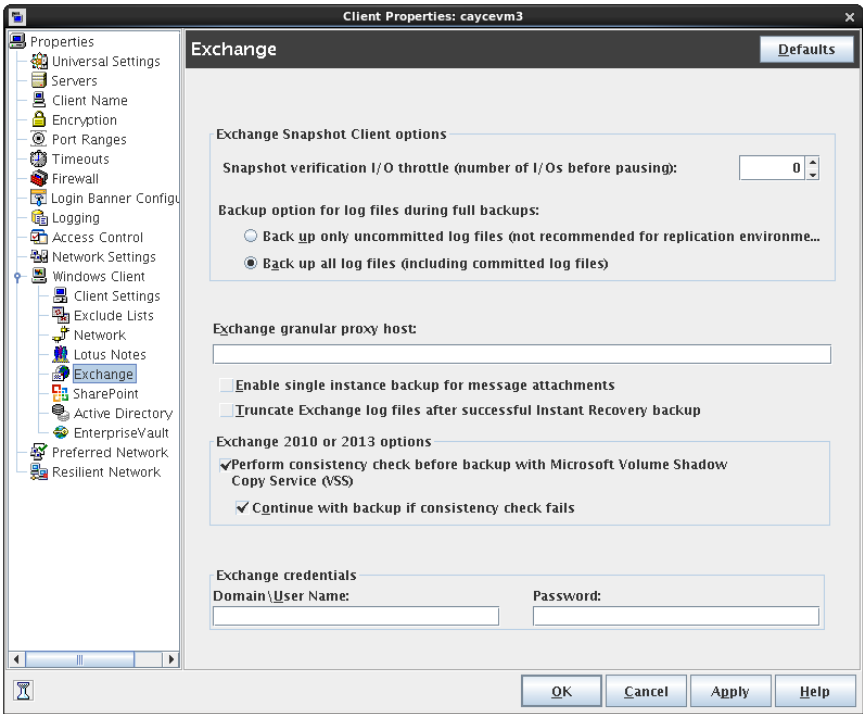
Exchange properties

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 CCR active node or DAG host server is updated.

For complete information on these options, see the [NetBackup for 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	<p>Note: This property only applies to MS-Exchange-Server backup policies with Exchange 2007.</p> <p>This option applies to off-host Exchange 2007 snapshot backups if the Exchange System Management Tools are not installed on the alternate client. Specify the number of I/Os to process for each 1-second pause.</p>

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

Property	Description
Backup option for log files during full backups	<p>Note: This property only applies to MS-Exchange-Server backup policies.</p> <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 DAGs or Exchange 2007 CCR environments.■ Back up all log files (including committed log files)
Truncate log after successful Instant Recovery backup	<p>Note: This property only applies to MS-Exchange-Server backup policies.</p> <p>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.</p>
Exchange granular proxy host	<p>Note: This property applies when you duplicate or browse a backup that uses Granular Recovery Technology (GRT).</p> <p>You can specify a different Windows system to act as a proxy for the source client when you duplicate or browse a backup (with <code>bplist</code>) that uses GRT. Use a proxy if you do not want to affect the source client or if it is not available.</p> <p>For complete information on this option, see the NetBackup for Exchange Server Administrator's Guide.</p>
Enable single instance backup for message attachments	<p>Note: This property only applies to MS-Exchange-Server backup policies with Exchange 2007.</p> <p>Enable this option to back up the data that is stored on a Single Instance Store (SIS) volume. This feature only applies to mailbox and public folder backups with Exchange Server 2007.</p>
Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS)	<p>Note: This property only applies to MS-Exchange-Server backup policies.</p> <p>Disable this option if you do not want to perform a consistency check during a DAG backup. If you check Continue with backup if consistency check fails, NetBackup continues to perform the backup even if the consistency check fails.</p>

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

Property	Description
Exchange credentials	<p>Note the following for this property:</p> <ul style="list-style-type: none">■ This property applies to MS-Exchange-Server and VMware backup policies with Exchange recovery.■ For Exchange 2013 and later, you must configure this property if you want to use GRT.■ This property does not apply to any backup and restore operations that have a policy that uses the <code>Microsoft Exchange Mailboxes:\</code> and the <code>Microsoft Exchange Public Folders:\</code> directives. Instead, you must configure the logon account for the NetBackup Client Service with the credentials for the account for NetBackup Exchange operations. <p>Provide the credentials for the account for NetBackup Exchange operations. This account must have the necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have. The account also needs the right to “Replace a process level token.”</p> <p>For complete information on this property, see the NetBackup for Exchange Server Administrator's Guide.</p>

About the Exchange credentials in the client host properties

The Exchange credentials in the client host properties indicate the account that has necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have.

Note the following:

- The account that you configured for the **Exchange credentials** must also have the right to “Replace a process level token.”
- For database restores from VMware backups, the Exchange credentials that you provide must have permissions to restore VM files.
- If you want to restore from a VMware snapshot copy that was created with Replication Director, do the following:
 - Provide the Exchange credentials in the **Domain\user** and **Password** fields.
 - Configure the NetBackup Client Service with an account that has access to the CIFS shares that are created on the NetApp disk array.

Additional notes for Exchange 2013 and later

To use GRT, configure the Exchange credentials on all granular clients.

Alternatively, you can configure the Exchange credentials only on the granular clients that perform restores. In this case, for the entire domain add “Exchange Servers” to the “View-Only Organization Management” role group. Perform this configuration in the Exchange Administration Center (EAC) or in Active Directory. See the following Microsoft article for more information:

<http://technet.microsoft.com/en-us/library/jj657492>

Additional notes for Exchange 2010 and 2007

These additional items apply to Exchange 2010 and 2007:

- For granular restores from VMware backups, only the Exchange client that performs the granular restore requires configuration of the Exchange credentials. The Exchange credentials are not required for backup or browse operations.
- If you configure the NetBackup Client Service with a logon account and configure the Exchange credentials in the Exchange client host properties, you must configure the “Replace a process level token” for both users.
- In NetBackup 7.5 and earlier, you added the credentials to the logon account for the NetBackup Client Service and the NetBackup Legacy Network Service (Exchange 2010). Customers upgrading from earlier versions of NetBackup can continue to use this configuration (unless you use Replication Director). Symantec recommends that you use the new configuration in the Exchange client host properties.
- (Exchange 2007) If you use the `Microsoft Exchange Mailboxes:\` and the `Microsoft Exchange Public Folders:\` directives in a policy, you must use the old configuration. Configure the NetBackup Client Service to log on as the account for NetBackup Exchange operations.

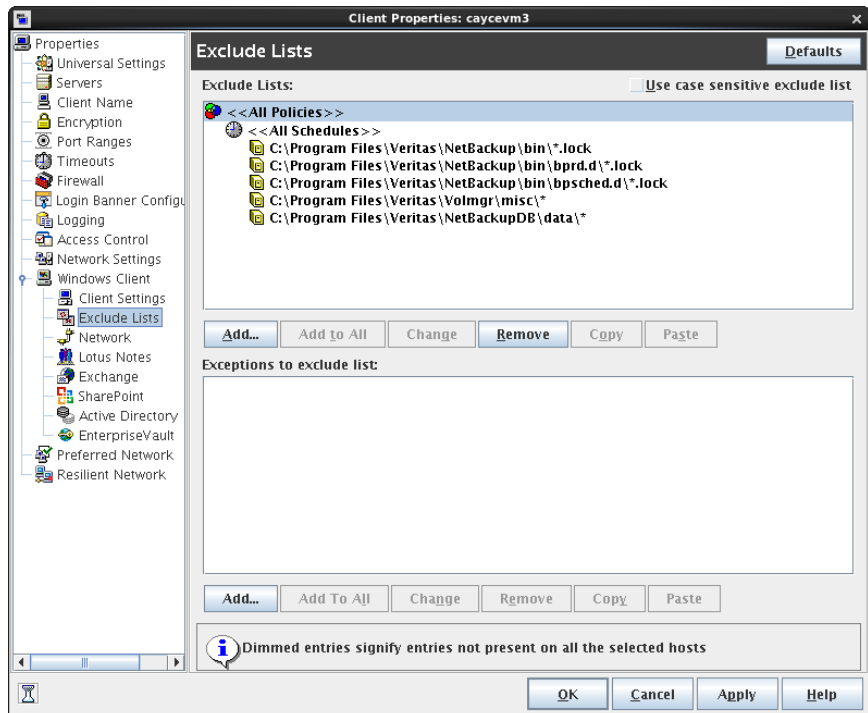
Exclude Lists properties

Use the **Exclude Lists** host properties to create and to modify the exclude lists for Windows clients. An exclude list names the files and directories to be excluded from backups of the selected Windows clients.

For UNIX clients, use the `nbgetconfig` and `nbsetconfig` commands to gather and modify the exclude list files from the `/usr/opensv/netbackup` directory on each client. `nbgetconfig` and `nbsetconfig` can also be used on a UNIX server to create an exclude list on a Windows client.

See “Configuring the exclude list on a Windows client from a UNIX master server” on page 136.

Figure 3-24 Exclude Lists dialog box



Exclude Lists properties apply only to Windows clients.

For information about creating exclude lists for UNIX clients, see the following topic:

See [“About excluding files from automatic backups”](#) on page 832.

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 a policy or for a policy and schedule 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 and a drive or a mount point is excluded, that job will report an error status when it completes. To avoid the situation, do not use the policy or the policy and schedule lists to exclude an entire volume.

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

Table 3-30 Excludes Lists dialog box properties

Property	Description
Use case sensitive exclude list	Indicates that the files and directories to exclude are case-sensitive.
Exclude lists	<div>Displays the policies that contain schedule, file, and directory exclusions as follows:</div> <ul style="list-style-type: none">Add Click to exclude a file from being backed up by a policy. The exclusion is configured in the Add to exclude list dialog box, then added to the Exclude list. When the policies in this list are run, the files and directories that are specified on the list are not backed up.Add to all When multiple Windows clients are selected, click Add to all to add the selected list item to the Exclude list for all selected clients. Add to all is enabled only when multiple Windows clients are selected for configuration.Remove Click to remove the selected policy, schedule, or file from the Exclude list.

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 if they were part of the original file list and were subsequently excluded. The list is useful to exclude all files in a directory but one.</p> <p>For example, if the file list of items to back up contains <code>/foo</code>, and the exclude list contains <code>/foo/bar</code>, adding <code>/fum</code> to the exceptions list does not back up the <code>/fum</code> directory. However, adding <code>fum</code> to the exceptions list will back up any occurrences of <code>fum</code> (file or directory) that occur within <code>/foo/bar</code>.</p> <p>Options include the following:</p> <ul style="list-style-type: none">■ Add Click to create an exception to the Exclude list. The exception is configured in the Add exceptions to exclude list dialog box, then added to the Exceptions to the exclude list. When the policies on the Exceptions to the exclude list run, the items on the exceptions list are backed up (if they were part of the original file list and subsequently excluded). Effectively, you add files back into the backup list of a policy.■ Add to all When multiple Windows clients are selected, click Add to all to add the selected list item to the Exceptions to the exclude list for all selected clients. When the policies on the exclude list run, the items on the exceptions list are backed up on all selected clients.■ Remove Click to remove the selected policy, schedule, or file from the Exceptions list. After the item is removed from the Exclude list, the item is excluded from the backup.


About the Add to Exclude List and Add Exceptions to Exclude 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 .
Files/Directories	Either browse or enter the full path to the files 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, either browse or 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, either browse or 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 1094.
See [“NetBackup naming conventions”](#) on page 1093.
- If all files are excluded in the backup selections list, NetBackup backs up only what is specified by full path names in the include list. Files can be excluded by using / or * or by using both symbols together (/*).
- 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.

Example of a Windows client exclude list

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

```
C:\users\doe\john
C:\users\doe\abc\
C:\users\*\test
C:\*\temp
core
```

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.

Example of a UNIX exclude list

In this example of a UNIX exclude list, the list contains the following entries:

```
# this is a comment line
/home/doe/john
/home/doe/abc/
/home/*/test
/*/temp
core
```

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

- The file or directory named `/home/doe/john`.
- The directory `/home/doe/abc` (because the exclude entry ends with `/`).
- All files or directories named `test` that are two levels beneath `home`.
- All files or directories named `temp` that are two levels beneath the root directory.
- All files or directories named `core` at any level.

About creating an include list on a UNIX client

To add a file that is eliminated with the exclude list, create a `/usr/opensv/netbackup/include_list` file. The same syntax rules apply as for the exclude list.

Note: Exclude and include lists do not apply to user backups and archives.

To illustrate the use of an include list, we use the example from the previous discussion. The exclude list in that example causes NetBackup to omit all files or directories named `test` from all directories beneath `/home/*/test`.

In this case, add a file named `/home/jdoe/test` back into the backup by creating an `include_list` file on the client. Add the following to the `include_list` file:

```
# this is a comment line
/home/jdoe/test
```

To create an include list for a specific policy or policy and schedule combination, use a `.policyname` or `.policyname.schedulename` suffix. The following are two examples of include list names for a policy that is named `wkstations` that contains a schedule that is named `fulls`.

```
/usr/opensv/netbackup/include_list.workstations
/usr/opensv/netbackup/include_list.workstations.fulls
```

The first file affects all scheduled backups in the policy that is named `wkstations`. The second file affects backups only when the schedule is named `fulls`.

For a given backup, NetBackup uses only one include list: the list with the most specific name. Given the following two files:

```
include_list.workstations
include_list.workstations.fulls
```

NetBackup uses only `include_list.workstations.fulls` as the include list.

Configuring the exclude list on a Windows client from a UNIX master server

An administrator can change the exclude list on a Windows client from a UNIX master server by using the `nbgetconfig` (or `bpgetconfig`) and the `nbsetconfig` (or `bpsetconfig`) commands.

Configuring the exclude list on a Windows client from a UNIX master server

- 1 From the UNIX server, use the `nbgetconfig` command to access the exclude list on a Windows client.

For example, where:

nbmaster1 is a UNIX master server.

nbmedia1 is a Windows computer and a NetBackup client.

On *nbmaster1*, run the following command:

```
/usr/opensv/netbackup/bin/nbgetconfig -M nbmedia1 EXCLUDE
```

The output lists the exclude list on *nbmedia1*:

```
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bprd.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bpsched.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackupDB\data\*
EXCLUDE = C:\Program Files\Veritas\Volmgr\misc\*
```

- 2 Create an exclude list on the UNIX server.

For example, file `ex_list` on *nbmaster1*:

```
vi /tmp/ex_list
```

The contents of `ex_list` consists of the following files and directories:

```
more /tmp/ex_list
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bprd.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bpsched.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackupDB\data\*
EXCLUDE = C:\Program Files\Veritas\Volmgr\misc\*
EXCLUDE = C:\
EXCLUDE = D:\
```

- 3 From the UNIX server, use the `nbsetconfig` command to set the exclude list on a Windows client.

For example, run the following command on *nbmaster1* to make *ex_list* the exclude list on client *nbmedia1*:

```
/usr/opensv/netbackup/bin/nbsetconfig -h nbmedia1 /tmp/ex_list
```

- 4 On the UNIX server, use the `nbgetconfig` command to view the new exclude list on the Windows client.

For example:

```
/usr/opensv/netbackup/bin/nbgetconfig -M nbmedia1 EXCLUDE
```

The output lists the updated exclude list on *nbmedia1*:

```
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bprd.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackup\bin\bpsched.d\*.lock
EXCLUDE = C:\Program Files\Veritas\NetBackupDB\data\*
EXCLUDE = C:\Program Files\Veritas\Volmgr\misc\*
EXCLUDE = C:\
EXCLUDE = D:\
```

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

NetBackup Fibre Transport properties control how your Fibre Transport media servers and SAN clients use the Fibre Transport service for backups and restores. The **Fibre Transport** properties apply to the host type that you select in the **NetBackup Administration Console**, as follows:

Table 3-32 Host types for Fibre Transport properties

Host type	Description
Master server	Global Fibre Transport properties that apply to all SAN clients.
Media server	The Fibre Transport Maximum concurrent FT connections property applies to the FT media server or servers that you selected in the NetBackup Administration Console .
Client	The Fibre Transport properties apply to the SAN client or clients that you selected in the NetBackup Administration Console . The default values for clients are the global property settings of the master server. Client properties override the global Fibre Transport properties.

Figure 3-26 Fibre Transport host properties for a master server

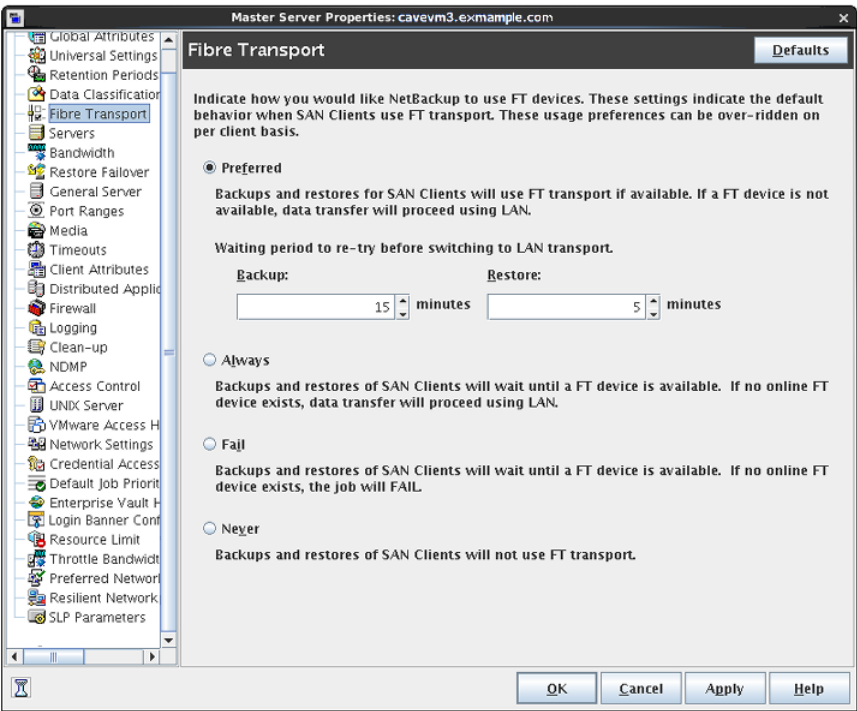


Table 3-33 describes the **Fibre Transport** properties. All properties are not available for all hosts. In this table, FT device is an HBA port on a Fibre Transport media server. The port carries the backup and restore traffic. A media server may have more than one FT device.

Table 3-33 Fibre Transport dialog box properties

Property	Description
Maximum concurrent FT connections	<p>This property appears only when you select an FT media server or servers in the NetBackup Administration Console.</p> <p>This property specifies the number of FT connections to allow to the selected media server or media servers. A connection is equivalent to a job.</p> <p>NetBackup supports 644 buffers per media server for Fibre Transport. To determine the number of buffers that each connection uses, divide 644 by the value you enter. More buffers per connection equal better performance for each connection.</p>

Table 3-33 Fibre Transport dialog box properties (*continued*)

Property	Description
Use defaults from the master server configuration	<p>This property appears only when you select a client or client in the NetBackup Administration Console.</p> <p>This property specifies that the client follow the properties as they are configured on the master server.</p>
Preferred	<p>The Preferred property specifies to use an FT device if one 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>
Always	<p>The Always property specifies that NetBackup should always use an FT device 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 online and up. If not, NetBackup uses the LAN. An FT device may be unavailable because none are active, none have been configured, or the SAN Client license expired.</p>
Fail	<p>The Fail property specifies that NetBackup should fail the job if an FT device is not online and up. If the FT devices are online but busy, NetBackup waits until a device is available and assigns the next job to the device. An FT device may be unavailable because none are 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>

NetBackup provides one finer level of granularity for Fibre Transport. SAN client usage preferences override the FT properties that you configure through **Host Properties**.

For more information about NetBackup Fibre Transport, see the *NetBackup SAN Client and Fibre Transport Guide*:

<http://www.symantec.com/docs/DOC5332>

About Linux concurrent FT connections

NetBackup uses the **Maximum concurrent FT connections Fibre Transport** host property to configure the number of concurrent connections to a Fibre Transport media server, up to the total that is allowed per host.

See “[Fibre Transport properties](#)” on page 139.

```
/usr/opensv/netbackup/db/config/NUMBER_DATA_BUFFERS_FT
```

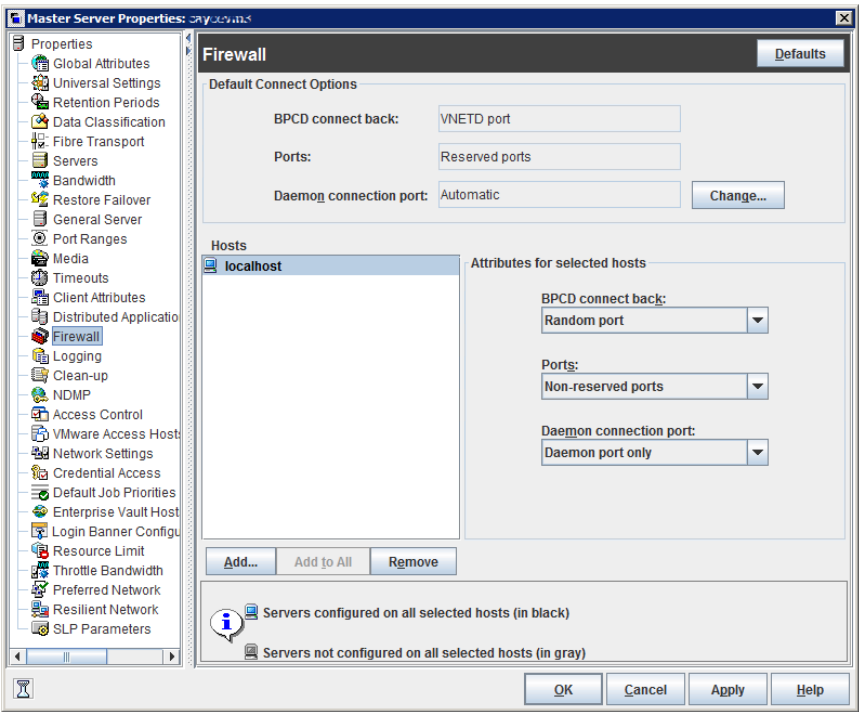
Firewall properties

The **Firewall** properties describe how the selected master servers and media servers that 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 84.

Figure 3-27 Firewall dialog box



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

Table 3-34 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>These properties correspond to the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option.</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>
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	<p>Connect options can be configured for individual servers.</p> <p>These properties correspond to the <code>CONNECT_OPTIONS</code> configuration option.</p>

Table 3-34 Firewall dialog box properties (*continued*)

Property	Description
BPCD connect back	<p>This property specifies how daemons are to connect back to the NetBackup Client daemon (<code>BPCD</code>) 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>bpbrrm</code> on a media server initially connects with <code>bpcd</code> on a client. The situation does not pose a firewall problem because <code>bpbrrm</code> uses the well-known <code>bpcd</code> port.
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 249.

Table 3-34 Firewall dialog box properties (*continued*)

Property	Description
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`.

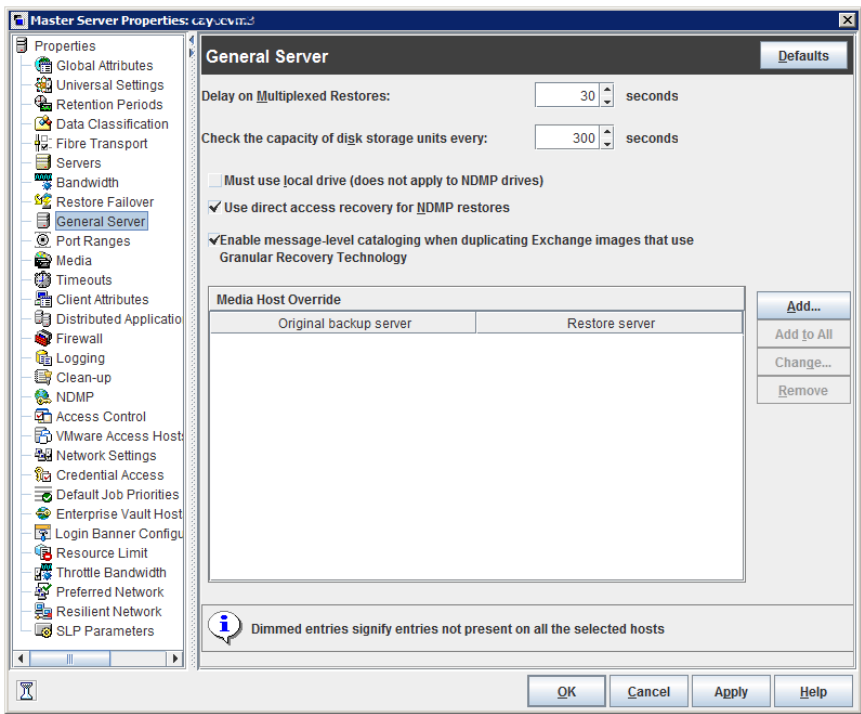
To enable logging for vnetd

- ◆ 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 3-35 General Server dialog box properties

Property	Description
Delay on multiplexed restores	<p>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).</p> <p>The default is a delay of 30 seconds.</p>
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>

Table 3-35 General Server dialog box properties (*continued*)

Property	Description
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 NetBackup for Exchange Administrator's Guide.</p>

Table 3-35 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 (`bprd`) on the master server.

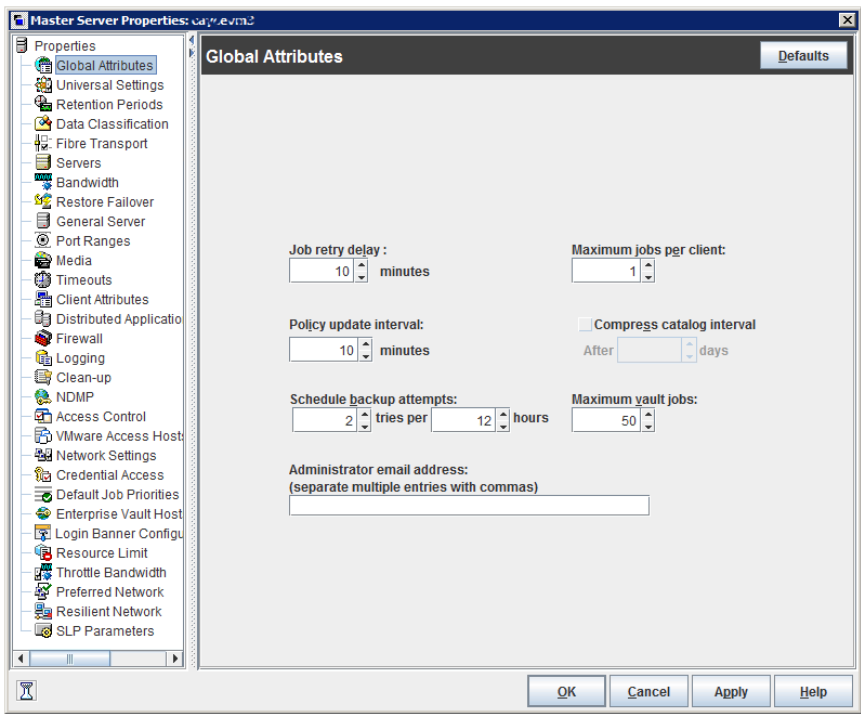
This process applies to all storage units on the original backup server. Restores for any storage unit on the **Original backup server** go to the server that is listed as the **Restore server**.

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

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 3-36 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.

Table 3-36 Global Attributes dialog box properties (*continued*)

Property	Description
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	<p>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.</p>
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 153.</p>
Compress catalog interval	<p>This property specifies how long NetBackup waits after a backup before it compresses the image catalog file.</p>

Table 3-36 Global Attributes dialog box properties (*continued*)

Property	Description
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>If the active vault jobs limit is reached, subsequent vault jobs are queued and their status is shown as Queued in the Activity Monitor.</p> <p>If a duplication job or eject job waits, its status is shown as Active in the Activity Monitor. See “About the Jobs tab” on page 1020.</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><code>useraccount1@company.com,useraccount2@company.com</code></p> <p>Disaster recovery information that is created during catalog backups is not sent to the addresses indicated here. Disaster recovery information is sent to the address that is indicated on the Disaster Recovery tab in the catalog backup policy. See “Disaster Recovery tab” on page 837.</p>

About constraints on the number of concurrent jobs


Specify any number of concurrent jobs within the following constraints.

Table 3-37 Constraints on concurrent jobs

Constraint	Description
Number of storage devices	<p>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.</p>

Table 3-37 Constraints on concurrent jobs (*continued*)

Constraint	Description
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 Servers > Double-click on a specific master server > Client Attributes > General tab.) See “General tab of the Client Attributes properties” on page 86. ■ 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 693.
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: Catalog backups can run concurrently with other backups. To do so, set the **Maximum jobs per client** value to greater than two for the master server. 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 emails can be sent, the computing environment must be configured correctly.

Additional Windows configuration to receive emails

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

The `nbmail.cmd` script must be configured on any Windows host that is to receive email.

Windows systems also 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.

Additional UNIX configuration to receive emails

NetBackup uses the mail transfer agent `sendmail` to send email notifications. Install `sendmail` from the following website and configure the environment accordingly:

<http://www.sendmail.org/releases>

Notification methods

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.
To do so, configure the **Administrator email address** in the **Global Attributes** properties. Since email is sent from the master server or the media server, the `nbmail.cmd` script or `sendmail` must be configured on the master and the media servers.
See [“Sending messages to the global administrator about unsuccessful backups only”](#) on page 158.
- Send emails about successful and failed backups.

Send a message to the local administrator(s) of each client about successful and unsuccessful backups.

To do so, configure either the **Client sends mail** property or the **Server sends mail** property in the **Universal Settings** properties for the client. Email is sent from the client (**Client sends mail**) or the master or the media server (**Server sends mail**). Therefore, `nbmail.cmd` or `sendmail` must be configured on the master and the media servers, or the client if specified.

See [“Sending messages to the administrator about successful and unsuccessful backups”](#) on page 158.

See [“Installing and testing the email utility on Windows”](#) on page 160.

See [“About constraints on the number of concurrent jobs”](#) on page 153.

Configuring the nbmail.cmd script on Windows

For Windows hosts to receive email notifications about backups, the `nbmail.com` script must be configured on the host.

By default, `nbmail.cmd` does not send email. Use the following procedure to configure the `nbmail.cmd` script.

To configure the nbmail.cmd script

- 1 On a NetBackup master server, copy `nbmail.cmd` from the following directory:

```
install_path\VERITAS\NetBackup\bin\goodies\nbmail.cmd
```

- 2 Paste it into the following directory on the server or client:

```
install_path\VERITAS\NetBackup\bin\
```


- 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 mail 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 mail 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 153.

Sending messages to the global administrator about unsuccessful backups only

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 Windows servers, install and configure a mail client.
See [“Installing and testing the email utility on Windows”](#) on page 160.
Edit the `nbmail.cmd` script.
See [“Configuring the nbmail.cmd script on Windows”](#) on page 156.
- 2 On the master server, open the **NetBackup Administration Console**.
- 3 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Server**.
- 4 In the right pane, double-click the master server you want to modify.
- 5 Open the host properties of the master server.
- 6 In the properties dialog box, in the left pane, click **Global Attributes**.
- 7 In the **Administrator’s email address** field, enter the email address of the administrator to receive the notification emails. (Separate multiple addresses with commas.) Click **Apply**.

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

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

For example:

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


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 or through the media server that backed up the client. An email can be sent to the administrator after every backup.

To send email notifications for all backups from a client

- 1 On a Windows client, install and configure a mail client.
See [“Installing and testing the email utility on Windows”](#) on page 160.
On the client or the media server, edit the `nbmail.cmd` script, depending on whether **Client sends mail** is enabled or **Server sends mail** is enabled in step 6.
- 2 On the master server, open the **NetBackup Administration Console**.
- 3 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Clients**.
- 4 In the right pane, double-click the client you want to modify. Multiple clients can also be selected.
- 5 In the properties dialog box, in the left pane, select **Universal Settings**.
- 6 Enable the **Client sends mail** option or if the client cannot send mail, enable the **Server sends mail** option.
- 7 In the **Client administrator's email** field, enter the email address of the administrator(s) to receive the notification emails. (Separate multiple addresses with commas.) Click **Apply**.

See [“Universal Settings properties”](#) on page 249.



Installing and testing the email utility on Windows

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 host name or IP address of the mail 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 `-t i 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 mail server.
- If the mail 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 the 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](#).

[NetBackup Administration Console logging](#) describes the logging for the NetBackup Administration Console.

Legacy logging

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

Note: If you have used the Logging Assistant tool earlier, you do not need to manually create log directories. The Logging Assistant tool automatically creates log directories.

Create the NetBackup legacy log directories in the following locations:

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

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

To create all of the NetBackup debug log directories at one time, run the following files:

- On Windows, run `mklogdir.bat`:
`install_path\NetBackup\logs\mklogdir.bat`
- On UNIX, run the `mklogdir` shell script: `/usr/opensv/netbackup/logs/mklogdir`

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

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

In the case of legacy logging, NetBackup uses the configuration file (the Windows registry or the `bp.conf` file on UNIX) to set the maximum size of a log file. Use the `bpsetconfig` command-line interface to configure the following parameters:

- The `MAX_LOGFILE_SIZE` parameter indicates the maximum size of a log file. When the log file size in NetBackup matches the `MAX_LOGFILE_SIZE` setting, the next logs are stored in a new log file.
See [“MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging”](#) on page 300.
- The `MAX_NUM_LOGFILES` parameter indicates the maximum number of log files that can be created in NetBackup. When the number of log files matches the `MAX_NUM_LOGFILES` setting, the older log files are purged.
See [“MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging”](#) on page 301.

Note: Do not edit the configuration file manually because incorrect syntax or characters in the file may adversely affect NetBackup operations.

See the [NetBackup Commands Reference Guide](#) for information about the `bpsetconfig` command.

Unified logging

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

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

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

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

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.

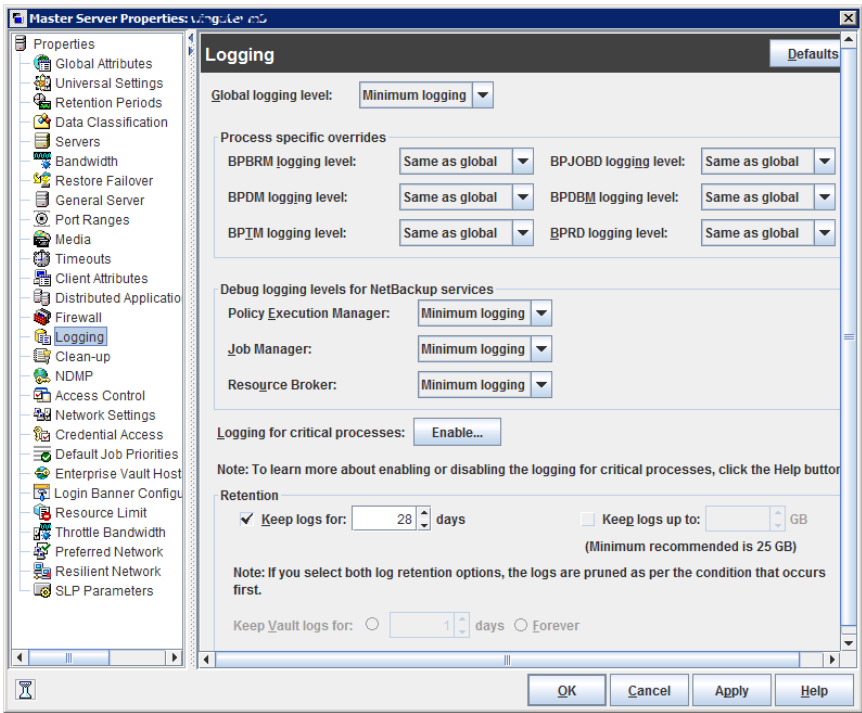
In the case of unified logging, NetBackup uses the `nblog.conf` configuration file to set the maximum size of a log file. Use the `vxlogcfg` command to configure the following log settings in the `nblog.conf` file:

- Configure the `MaxLogFileSizeKB` parameter to set the maximum size of a log file. When the log file size in NetBackup matches the `MaxLogFileSizeKB` setting, the next logs are stored in a new log file.
- Configure the `NumberOfLogFiles` parameter to set the maximum number of log files that you want to create in NetBackup. When the number of log files in NetBackup matches the `NumberOfLogFiles` setting, the older log files are purged. See the [NetBackup Logging Reference Guide](#) for more information about controlling the log file size.

Note: Do not edit the `nblog.conf` file manually. Incorrect syntax or characters in the file may adversely affect NetBackup operations.

See the [NetBackup Commands Reference Guide](#) for information about the `vxlogcfg` command and the `nblog.conf` file.

Figure 3-31 Logging dialog box



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

Table 3-38 Logging dialog box properties

Property	Description
Global logging level	<p>This property controls the logging level for all processes in the Logging dialog box that are set to Same as global. The logging levels control the amount of information that the NetBackup server writes to the logs.</p> <ul style="list-style-type: none"> ■ No logging No log is created for any process that is set to Same as global. ■ Minimum logging (default) A small amount of information is logged for each process that is set to Same as global. Use this setting unless advised otherwise by Symantec Technical Support. Other settings can cause the logs to accumulate large amounts of information. ■ Levels 1 through 4 Progressively more information is logged at each level for each process that is set to Same as global. ■ 5 (Maximum) The maximum amount of information is logged for each process that is set to Same as global. <p>Some NetBackup processes allow individual control over the amount of information the process writes to logs. For those processes you can specify a different logging level other than Same as global.</p>
Process specific overrides	<p>The processes 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"> ■ On Windows: <i>install_path\NetBackup\logs\process_name</i> ■ On UNIX: <i>/usr/opensv/netbackup/logs/process_name</i> <p>The following settings are available for each process:</p> <ul style="list-style-type: none"> ■ Same as global Make this selection to use the same logging level as the Global logging level property. If the Global logging level is set to No logging, no log is created for the individual process. ■ No logging No log is created for the specific process. ■ Minimum logging A small amount of information is logged for the specific process. ■ Levels 1 through 4 Progressively more information is logged at each level for the specific process. ■ 5 (Maximum) The maximum amount of information is logged for the specific process. <p>Table 3-39 lists and describes the processes that use legacy logging.</p>

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

Property	Description
Debug logging levels for NetBackup services	<p>These services use unified logging.</p> <p>Each service creates a log automatically in the following directories:</p> <ul style="list-style-type: none">■ On Windows: <i>install_path</i>\NetBackup\logs■ On UNIX: /usr/openv/logs <p>The following settings are available for each process:</p> <ul style="list-style-type: none">■ Same as global Make this selection to use the same logging level as the Global logging level property. If the Global logging level is set to No logging, no log is created for the individual process.■ No logging No log is created for the specific process.■ Minimum logging A small amount of information is logged for the specific process.■ Levels 1 through 4 Progressively more information is logged at each level for the specific process.■ 5 (Maximum) The maximum amount of information is logged for the specific process. <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 3-38 Logging dialog box properties (*continued*)

Property	Description
	<p>Click Enable to create logs for critical NetBackup processes. In previous NetBackup versions, logging starts once you have created the log directories for the required processes. After you click Enable, the log directories are automatically created for the critical NetBackup processes and logging begins.</p> <p>See “ENABLE_CRITICAL_PROCESS_LOGGING” on page 303.</p> <p>Note: Enabling Logging for critical processes is a one-time activity that you can do using the NetBackup Administration Console. After the logging is enabled, you cannot disable it using the Logging dialog box. To manually disable the logging for a process, disable the logging level of the process.</p> <p>See “Disabling the logging of NetBackup processes” on page 171.</p> <p>If you have enabled Logging for critical processes, clicking Defaults does not modify Logging for critical processes and Keep logs up to GB options. While changing the logging properties for multiple hosts, if the Logging for critical processes option on one of the hosts is enabled, clicking Defaults does not modify the Logging for critical processes and Keep logs up to GB options of any host.</p> <p>See “Standard host property dialog box options” on page 66.</p> <p>After the Logging for critical processes is enabled, NetBackup sets the log retention to the default log size in GB. You can modify it by reconfiguring the Keep logs up to GB option.</p> <p>Note: Clearing the Keep logs up to GB option may adversely affect NetBackup operations. Symantec recommends that you specify the log retention in GB when Logging for critical processes is enabled.</p> <p>The following critical NetBackup processes are logged when you select the Enable option:</p> <p>Master server processes:</p> <ul style="list-style-type: none">■ bprd■ bpdbm <p>Media server processes:</p> <ul style="list-style-type: none">■ bpbm■ bptm■ bpdm <p>Client processes:</p> <ul style="list-style-type: none">■ bpfis <p>Note: You can enable Logging for critical processes for master servers, media servers, and clients. You cannot enable Logging for critical processes for media servers and clients at versions earlier than NetBackup 7.7, even if they are associated with current version master servers.</p>

Table 3-38 Logging dialog box properties (continued)

Property	Description
Keep logs for days	<p>Specifies how many days 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 media servers and clients.</p>

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

Property	Description
Keep logs up to GB	<p>Specifies the size of the NetBackup logs that you want to retain. When the NetBackup log size grows to this value, the older logs are deleted. To set the log size in GB, select the check box and specify the value in GB.</p> <p>See “KEEP_LOGS_SIZE_GB for NetBackup servers and clients” on page 302.</p> <p>Note: In the case of NetBackup servers, the recommended value for the Keep logs up to GB option is 25 GB or greater.</p> <p>In the case of clients, the recommended value for the Keep logs up to GB option is 5 GB or greater.</p> <p>While configuring the Keep logs up to GB option, consider the NetBackup log pruning behavior. All logs are retained until the log size reaches the high watermark, that is, 95% of the Keep logs up to GB value. NetBackup checks for the log size every 10 minutes. Once the log size reaches the high watermark, the deletion of older logs is started. The deletion continues until the log size reaches the low watermark, that is, 85% of the Keep logs up to GB value.</p> <p>Consider the following example:</p> <ul style="list-style-type: none">■ The Keep logs up to GB option is set to 100 GB.■ NetBackup logs are created and retained until the log size reaches the high watermark, that is, 95 GB.■ NetBackup checks the log size every 10 minutes and determines that the log size is 95 GB.■ Log pruning is started and older logs are deleted until the log size reaches the low watermark, that is, 85 GB.■ Log pruning stops when the log size reaches 85 GB and all existing logs are retained until the log size again reaches the high watermark. <p>If you do not select the Keep logs up to GB option, the default value is unlimited. The log retention or log pruning does not happen based on the log size.</p> <p>If both Keep logs up to GB and Keep logs for days are selected, the logs are pruned depending on the condition that occurs first.</p> <p>If no log retention options are selected, the increasing log size may eventually lead to a situation of low disk space, as all logs are retained.</p> <p>Note: You can verify the log pruning behavior in NetBackup using the logs at the following location:</p> <p>On Windows: <code>install_path\NetBackup\logs\nbutil</code></p> <p>On UNIX: <code>/usr/opensv/netbackup/logs/nbutil</code></p> <p>Note: The Keep logs up to GB option is available for the NetBackup master server, media server, and clients. The older versions of NetBackup media servers and clients that are associated with NetBackup 7.7 master servers do not have the Keep logs up to GB option.</p>

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

Property	Description
Keep Vault logs for	<p>If NetBackup Vault is installed, the Keep Vault logs for property is enabled. Select the number of days to keep the Vault session directories, or select Forever.</p> <p>Session directories are found in the following location:</p> <pre>install_path\netbackup\vault\sessions\vaultname\ session_x</pre> <p>where x is the session number. This directory contains vault log files, temporary working files, and report files.</p>

Note: While collecting NetBackup logs using Logging Assistant, the total NetBackup log size may increase. If you have enabled the new **Keep logs up to GB** option and the total NetBackup log size reaches its high water mark, logs are deleted. Logs that you want to retain may also be deleted. To avoid the deletion of logs that you want to retain, disable the **Keep logs up to GB** option while you collect logs using Logging Assistant. Alternatively, you can set the **Keep logs up to GB** option to a value higher than the current value, so that the important logs are not deleted before the log collection is complete.

See [“About the Logging Assistant”](#) on page 1063.

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.
BPJOBDM logging level	The NetBackup Jobs Database Management daemon.
BPDBM logging level	The NetBackup database manager.
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 master servers. NBP creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, NBP is notified and the appropriate Policy/Client tasks are updated.
Job Manager	This property appears for master servers. NBJ accepts the jobs that the Policy Execution Manager submits and acquires the necessary resources.
Resource Broker	NBR makes the allocations for storage units, tape drives, client reservations.

Note: The **Enable robust logging** property has been removed from this release. Robust logging is now enabled by default.

The **Enable robust logging** property appears for earlier versions of master servers, media servers, or clients.

For more information about robust logging, refer to the section about legacy logging in the *NetBackup Logging Reference Guide*.

NetBackup Administration Console logging

Legacy logging and unified logging write information about the NetBackup processes that backup, restore, duplicate, and so on. You can also configure NetBackup so that it writes log information about the **NetBackup Administration Console**. The log information may be helpful to troubleshoot problems.

See “[Configuring log files for the NetBackup Administration Console](#)” on page 46.

The *NetBackup Logging Reference Guide* describes how to troubleshoot the **NetBackup Administration Console**:

<http://www.symantec.com/docs/DOC5332>

Disabling the logging of NetBackup processes

You can disable the logging of a NetBackup process by changing the logging level of the process.

Note: To enable **Logging for critical processes**, in the **NetBackup Administration Console**, select **NetBackup Management > Host Properties**. Double-click on **Master Servers**, **Media Servers**, or **Clients** and select a host. Open the **Logging** properties and click **Enable**.

To disable the logging of a process by changing the logging level

- ◆ Do any of the following:
 - In the **Logging** dialog box, in the **Process specific overrides** pane, set the logging level of the process to **No logging**.
 - Use the Logging Assistant to change the logging level of a process to disable the logging of the process.
See [“Disabling debug logging”](#) on page 1078.

Note: For the processes for which you cannot change or disable logging levels using the **Logging > Process specific overrides** pane or the Logging Assistant, use the configuration file option.

- Change or disable the logging level of a process by modifying using the configuration file (the Windows registry or the `bp.conf` file in UNIX). Use the `bpsetconfig` command to set the `ProcessName_VERBOSE` parameter to `-2`.

Example 1 - To disable logging for the `bprd` process, use the `bpsetconfig` command to set the value of `BPRD_VERBOSE` as follows:

```
bpsetconfig> BPRD_VERBOSE = -2
```

Example 2 - To disable logging for the `bpfis` process that is specific to NetBackup clients, run the `bpsetconfig` command on the associated media server or master server with the `-h host` option to set the value of `BPFIS_VERBOSE` as follows:

```
bpsetconfig> -h client_host_name  
bpsetconfig> BPFIS_VERBOSE = -2
```

See the [NetBackup Commands Reference Guide](#) for information on the `bpsetconfig` command.

Note: Do not edit the configuration file manually. Incorrect syntax or characters in the file may adversely affect NetBackup operations.

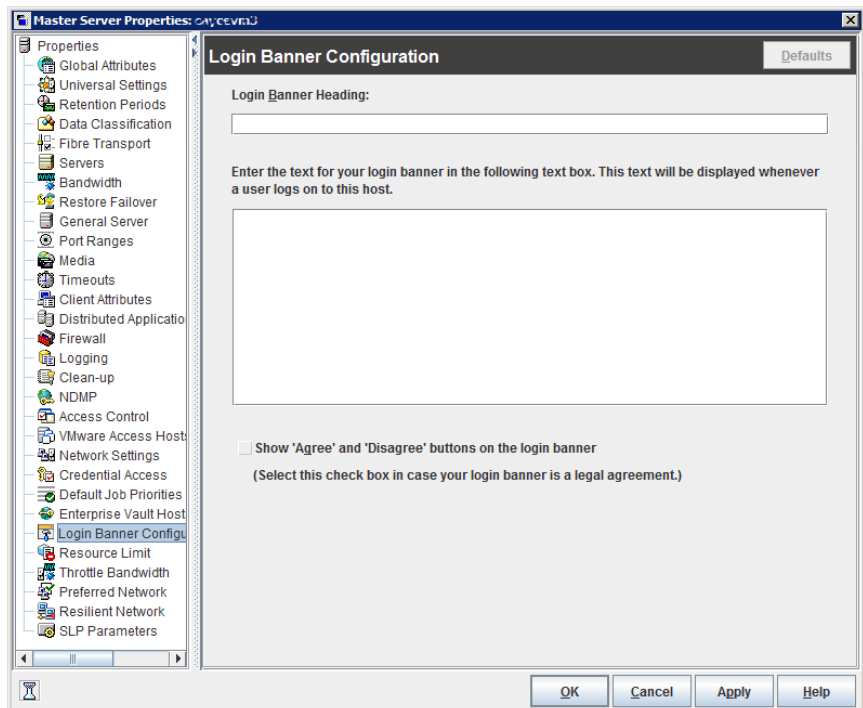
Login Banner Configuration properties

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 before 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.
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

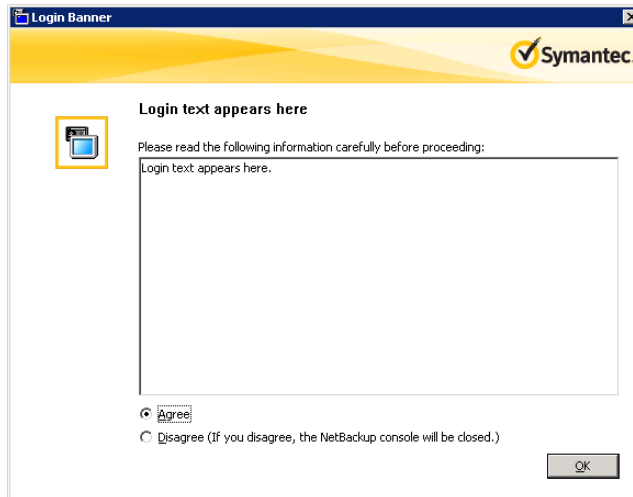



Figure 3-34 Login Banner without agreement option



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 **Administrative Options** tab (UNIX) or the **Administration Console** tab (Windows).
- 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** client 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**

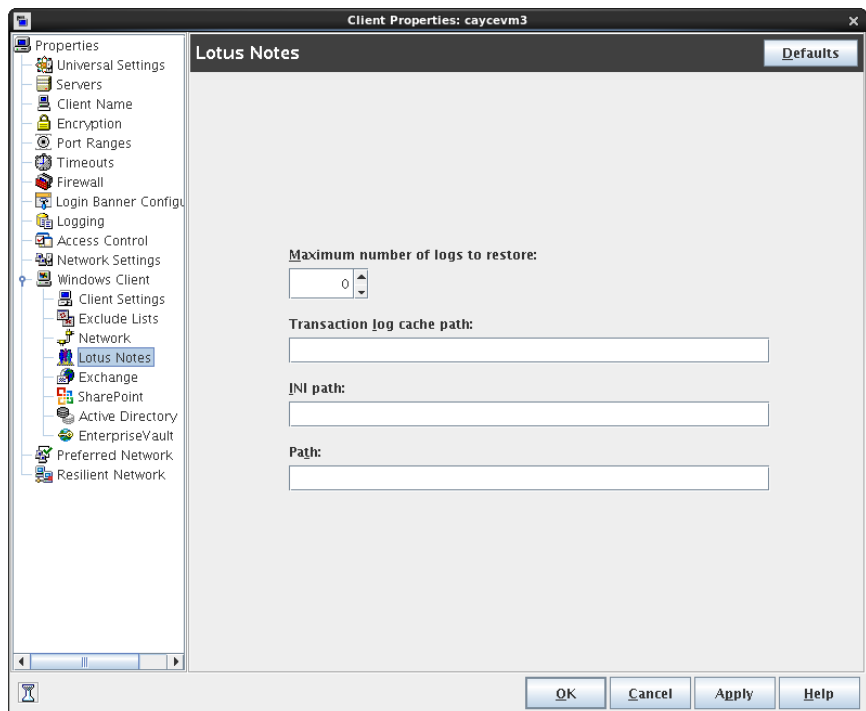
If the user selects this option, the session ends and NetBackup logs off the user immediately.

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



The following properties can be configured in the **Lotus Notes** host properties or by using the `nbgetconfig` and `nbsetconfig` commands.

See [“Methods to set the NetBackup configuration options”](#) on page 62.

For UNIX 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	Configuration options	Description
Maximum number of logs to restore	LOTUS_NOTES_LOGCACHESIZE In the Windows registry, this value is a <code>DWORD</code> value.	The maximum number of logs that can be prefetched in a single restore job during recovery. Specify a value greater than 1. 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. <code>LOTUS_NOTES_LOGCACHESIZE = 3</code>
Transaction log cache path	LOTUS_NOTES_LOGCACHEPATH In the Windows registry, this value is a string value.	Specify a path where NetBackup can temporarily store the prefetched transaction logs during recovery. For example: <ul style="list-style-type: none"> ■ On Windows: <code>D:\LogCache</code> ■ On UNIX: <code>/tmp/logcache</code> If you do not specify a path, during recovery NetBackup restores the logs to the Domino server's transaction log directory. Note the following before specifying the Transaction log cache path : <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.

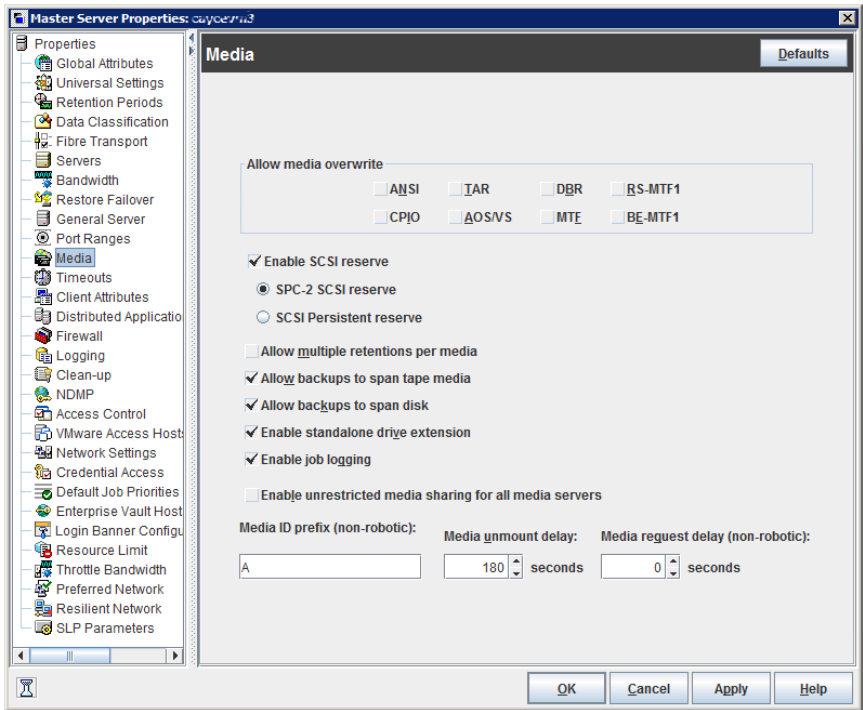
Table 3-42 Lotus Notes client host properties (*continued*)

Client host properties	Configuration options	Description
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"> ■ On Windows: If the notes.ini file is not located in the default directory, indicate its location in the INI path box. For example: <code>D:\Lotus\Domino\notes.ini</code> ■ On UNIX: If the notes.ini is not located in the directory that is specified in the Path, indicate its location here. For example: <code>/db/notesdata/notes.ini</code> Include the directory and the notes.ini file name.
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"> ■ On Windows: Specify the path for Lotus program directory (where nserver.exe resides). For example: <code>D:\Lotus\Domino</code> ■ On UNIX: Specify a path that includes the Domino data directory, the Lotus program directory, and the Lotus resource directory. For example: <code>/export/home/notesdata:/opt/lotus/notes/latest/sunspa:/opt/lotus/notes/latest/sunspa/res/C</code> <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
Allow media overwrite property	<p>This property overrides the NetBackup 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 184.</p>

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

Property	Description
Enable SCSI reserve	<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 185.</p> <p>See “Drive path options” on page 418.</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>
Allow multiple retentions per media	<p>This property lets NetBackup 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>
Allow backups to span tape media	<p>This property, when checked, lets backups span to multiple tape media. This property lets NetBackup 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 lets backups 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 563. ■ 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 lets NetBackup 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 NetBackup server groups” on page 367.</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
 - Sets the volume's state to FROZEN
 - Selects a different volume
 - Logs an error

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

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.

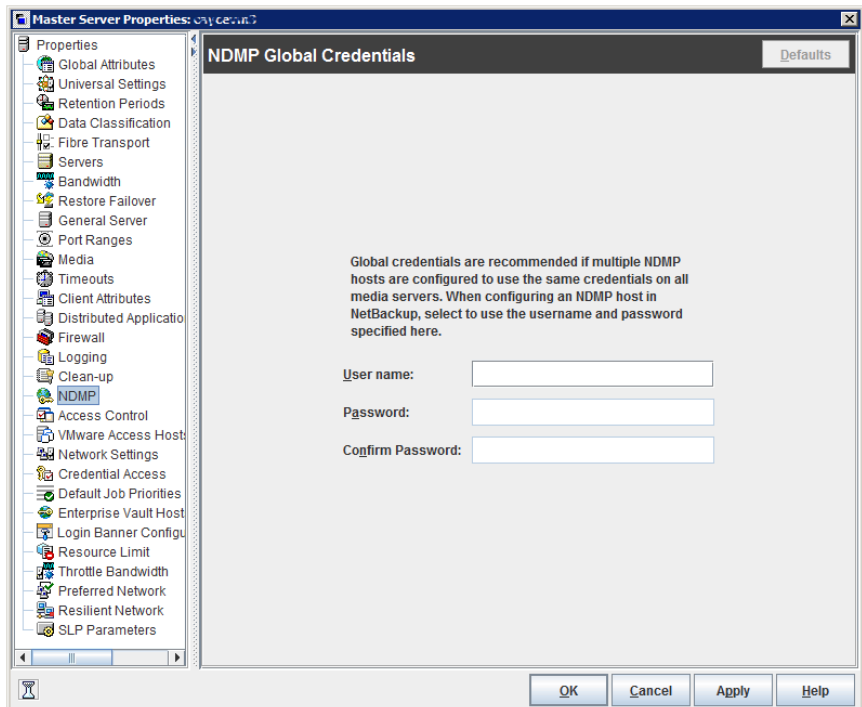
- The operating system tape drivers may require extensive configuration to use SCSI persistent reserve. For example, if the tape drives do not support SPC-3 Compatible Reservation Handling (CRH), ensure that the operating system does not issue SPC-2 reserve and release commands.

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.

Note: Be careful when you add an NDMP host. If you add NDMP host credentials using the fully qualified domain name (FQDN), you must also indicate the fully qualified domain name on the client for lookups. That is, the server list in the Backup, Archive, and Restore administration console must list the NDMP host by the FQDN as well.

If you add NDMP host credentials using its short name, you can use either the short name or the FQDN in the client server list.

Figure 3-38 Add NDMP Host dialog box

For NDMP Global Credentials to apply, this must be selected on the NDMP host as well

New NDMP Host - test

NDMP host: test

NDMP Host Credentials

☒ Use global NDMP credentials for this NDMP host

☐ Use the following credentials for this NDMP host on all media servers

User name:

Password:

Confirm Password:

☐ Use different credentials for this NDMP host on each media server
(Use Advanced Configuration)

To configure individual media server credentials or to override global and NDMP host level credentials, use Advanced Configuration.

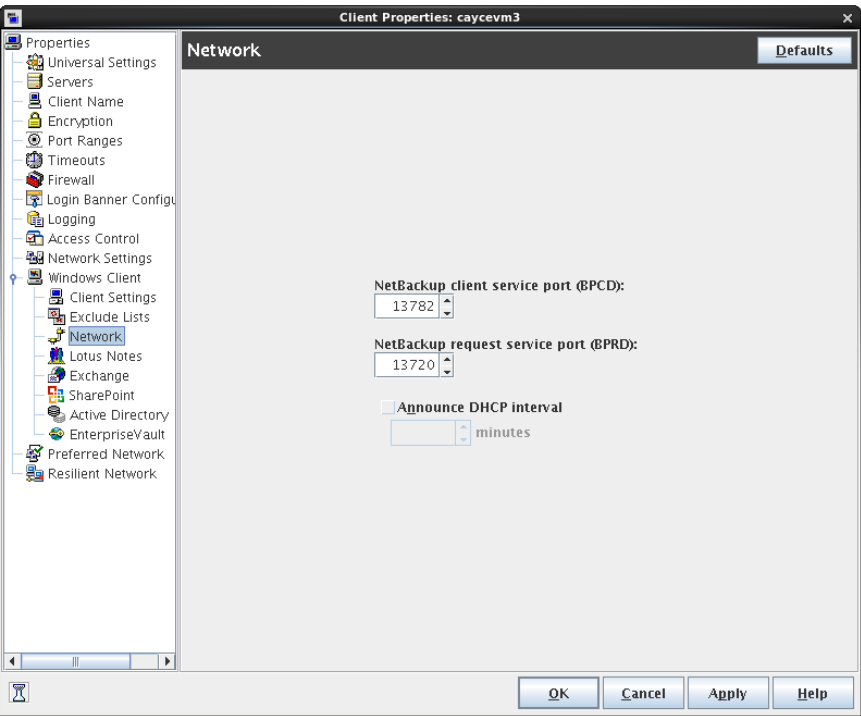
[Advanced Configuration...](#)

OK Cancel Help

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>
NetBackup request service port (BPRD)	<p>This property specifies the port for the client to use when it sends requests to the NetBackup request service (<code>bprd</code> process) on the NetBackup server. The default is 13720.</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>

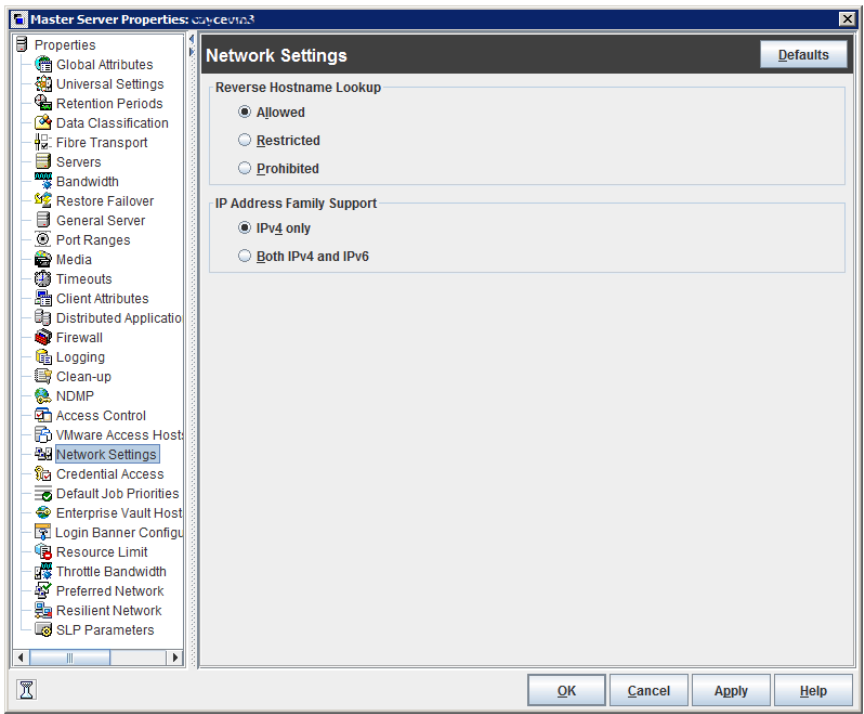
Table 3-45 Network dialog box properties (continued)

Property	Description
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 **NetBackup 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 requires being able to view the **Host Properties**, you'll need to use another method to change it.

Configure the `REVERSE_NAME_LOOKUP` option by using the `nbgetconfig` and `nbsetconfig` commands. The `nbsetconfig` command configures the option on Windows and UNIX master servers and clients.

See [“Methods to set the NetBackup configuration options”](#) on page 62.

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.

IP Address Family Support property

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

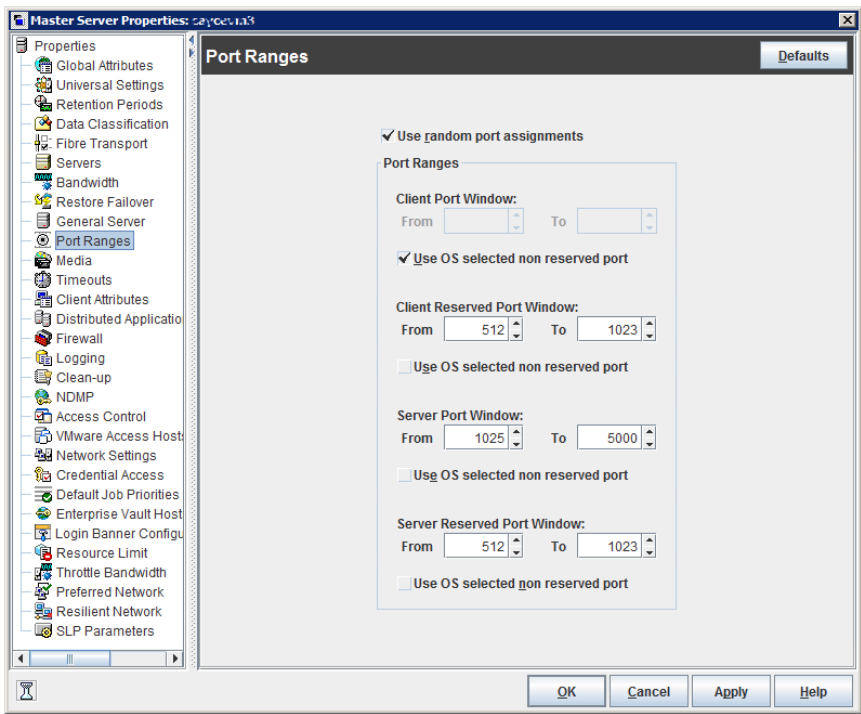
- **IPv4 only** (Default)
- **Both IPv4 and IPv6**

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

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>
Client port window	<p>Lets the administrator define the range of non-reserved ports on the selected computer. NetBackup can use any available port within this range as the source port when communicating with NetBackup on another computer.</p>

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

Property	Description
Use OS selected non reserved port	Lets the operating system determine which non-reserved port to use.
Client reserved port window	<p>This property is not used by NetBackup 7.0.1 and later unless one the following conditions is true:</p> <ul style="list-style-type: none"> ■ The connect options for the remote host specify that the daemon port be used. ■ Attempts to connect to ports 1556 and 13274 on the remote host fail, which causes the daemon port to be used. <p>If either condition is true, then this setting defines the range of reserved ports to use as the source port for the connection.</p> <p>Beginning with NetBackup 7.0.1, connections to other daemon ports use only the Client port window, not the Client reserved port window.</p>
Server port window	<p>This property specifies the range of non-reserved ports on which NetBackup processes on this computer to accept connections from NetBackup on other computers when the inbound connection is not to a well known port. This primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options for the remote NetBackup server or client and that host is configured for non-reserved ports.</p> <p>This property also applies in the situation where a third-party protocol is used, such as NDMP. It specifies the range of non-reserved ports on which this server accepts NetBackup connections from other computers. The default range is 1024 through 5000.</p> <p>Instead of indicating a range of ports, you can enable Use OS selected non reserved port to let the operating system determine which non-reserved port to use.</p> <p>This setting applies to the selected master or media server that connects to a server or client that is configured to accept only non-reserved ports. The Accept connections on non reserved ports property must be enabled in the Universal Settings dialog box for hosts running NetBackup 7.0.1 and prior versions.</p> <p>See “Universal Settings properties” on page 249.</p>
Server reserved port window	This entry specifies the range of local reserved ports on which this computer accepts connections from NetBackup on other computers when the inbound connection is not to a well known port. This primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options for the remote NetBackup server or client.

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 daemon (`bpcd`) is 13782.

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

On Windows: `%systemroot%\system32\drivers\etc\services`

On UNIX: `/etc/services`

Note: Symantec recommends that the port numbers that are associated with PBX (1556 and 1557) not be changed.

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.

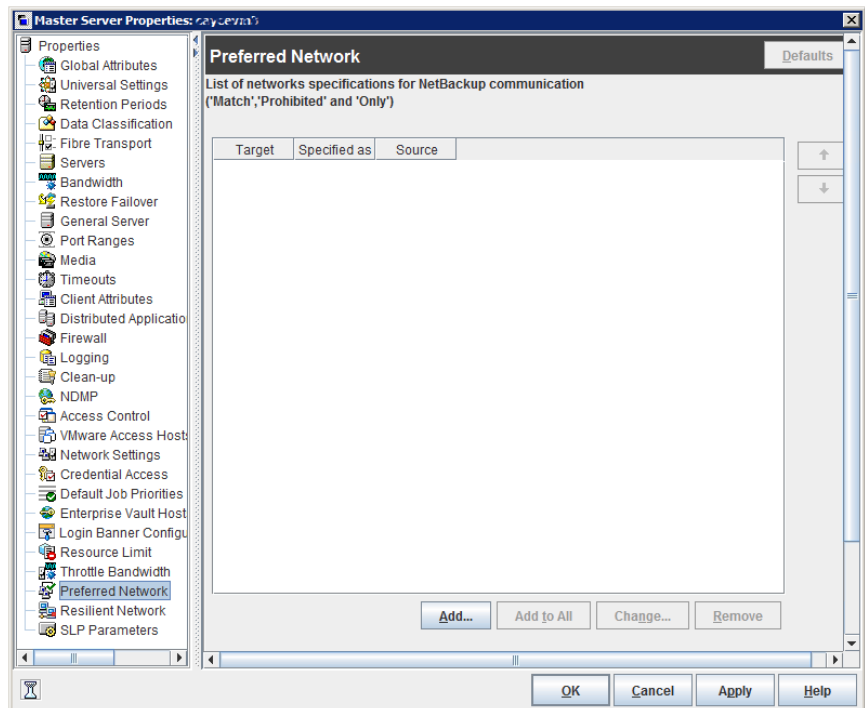
The **Preferred Network** properties are useful in NetBackup environments that include multihomed hosts—the hosts that are connected to two or more networks, or hosts that have two or more network addresses. The properties are especially helpful to administrators who must configure an environment that 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 205.

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 host names 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 204.</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 198.</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 197.</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 host name:</p> <ul style="list-style-type: none"> ■ If an address is specified as the network, it's 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 host name 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 host name is preferred for communication with the selected host.</p> <p>The Match directive does not reject other networks, addresses, or host names 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 host name 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 host name 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> configuration option.</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 host name, 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 configuration option for this property is <code>PREFERRED_NETWORK</code>. This property replaces the <code>REQUIRED_INTERFACE</code> entry.</p> <p>See “PREFERRED_NETWORK option for NetBackup servers” on page 308.</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 NetBackup** 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

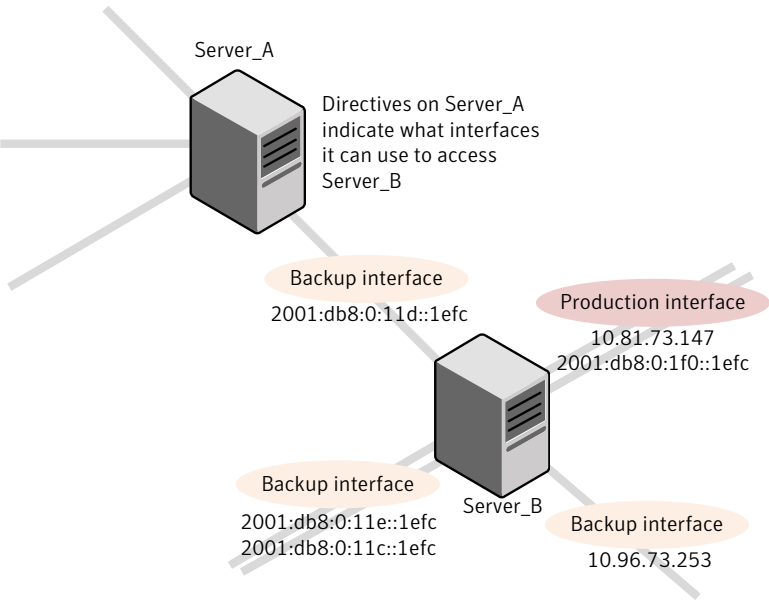


Figure 3-44 shows a table for Server_B. Server_B has multiple network interfaces, some of which have multiple IP addresses. In the table, yes indicates that NetBackup can use the network-IP combination as a source. In this example, no directives have been created for the host. Since no networks are listed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

Note: The following topic shows the `bptestnetconn` output for this example configuration:

See Figure 3-51 on page 206.

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 202.
- See [“Configurations to use IPv4 networks”](#) on page 203.
- See [“Configuration to prohibit using a specified address”](#) on page 207.
- See [“Configuration that uses the same specification for both the network and the interface—less constrictive”](#) on page 207.

- See “Configuration that uses the same specification for both the network and the interface—more constrictive” on page 208.
- See “Configuration that limits the addresses, but allows any interfaces” on page 208.

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 192.

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, 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 host name `host_fred` is used.

See [“How NetBackup uses the directives to determine which network to use”](#) on page 199.

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 host names 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 205.

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.

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 configuration file the directive appears as follows:

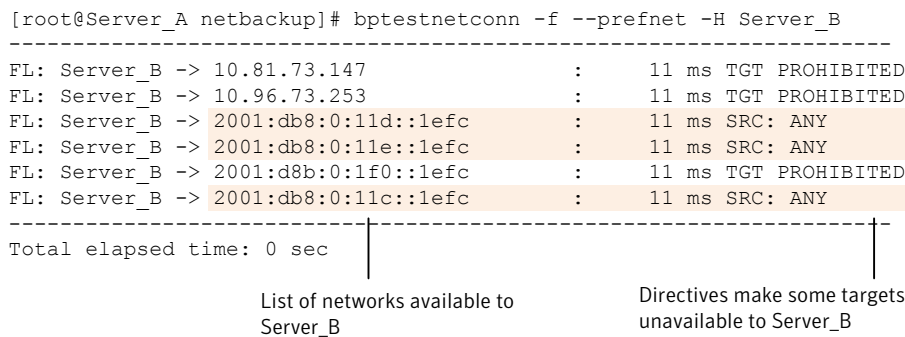
```
PREFERRED_NETWORK = 2001:0db8:0:11c::/62 ONLY
```

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.

See [“PREFERRED_NETWORK option for NetBackup servers”](#) on page 308.

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


Figure 3-52 bptestnetconn for Server_B with directive



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 the 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



Resilient Network properties

The **Resilient Network** properties appear for the master server, for media servers, and for clients. For media servers and clients, the **Resilient Network** properties are read only. When a job runs, the master server updates the media server and the client with the current properties.

The **Resilient Network** properties let you configure NetBackup to use resilient network connections for backups and restores. A resilient connection allows backup and restore traffic between a client and a NetBackup media server to function effectively in high-latency, low-bandwidth networks such as WANs. The data travels across a wide area network (WAN) to media servers in a central datacenter.

NetBackup monitors the socket connections between the remote client and the NetBackup media server. If possible, NetBackup re-establishes dropped connections and resynchronizes the data stream. NetBackup also overcomes latency issues to maintain an unbroken data stream. A resilient connection can survive network interruptions of up to 80 seconds. A resilient connection may survive interruptions longer than 80 seconds.

The NetBackup Remote Network Transport Service manages the connection between the computers. The Remote Network Transport Service runs on the master server, the client, and the media server that processes the backup or restore job. If the connection is interrupted or fails, the services attempt to re-establish a connection and synchronize the data.

NetBackup protects only the network socket connections that the NetBackup Remote Network Transport Service (`nbrntd`) creates. Examples of the connections that are not supported are:

- Clients that back up their own data (deduplication clients and SAN clients)
- Granular Recovery Technology (GRT) for Exchange Server or SharePoint Server
- NetBackup `nbfspd` process.

NetBackup protects connections only after they are established. If NetBackup cannot create a connection because of network problems, there is nothing to protect.

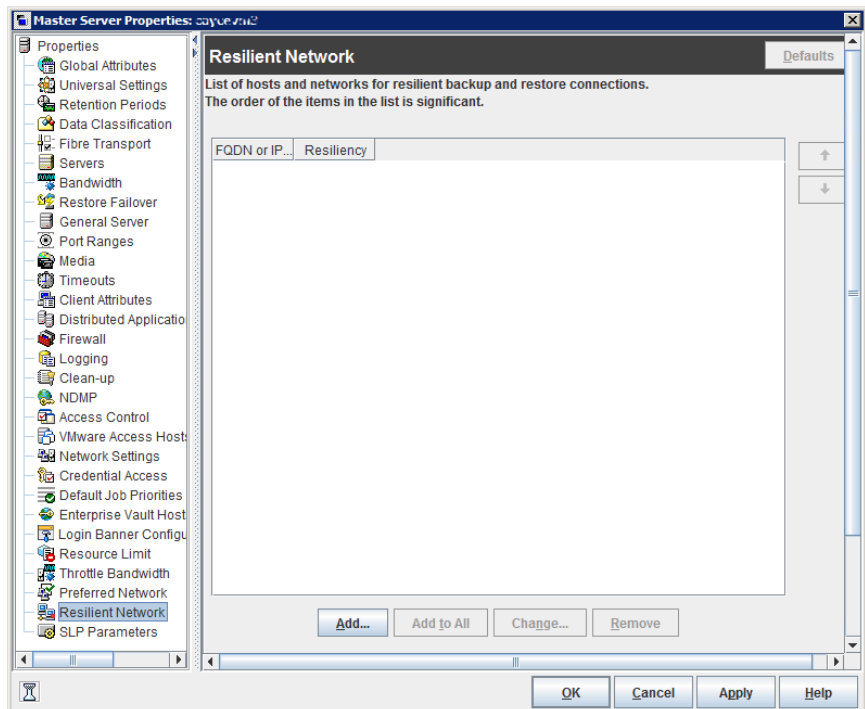
Resilient connections apply between clients and NetBackup media servers, which includes master servers when they function as media servers. Resilient connections do not apply to master servers or media servers if they function as clients and back up data to a media server.

Resilient connections can apply to all of the clients or to a subset of clients.

Note: If a client is in a different subdomain than the server, add the fully qualified domain name of the server to the client's hosts file. For example, **india.symantecs.org** is a different subdomain than **china.symantecs.org**.

When a backup or restore job for a client starts, NetBackup searches the **Resilient Network** list from top to bottom looking for the client. If NetBackup finds the client, NetBackup updates the resilient network setting of the client and the media server that runs the job. NetBackup then uses a resilient connection.

Figure 3-57 Master server Resilient Network host properties



See [“About the Daemons tab”](#) on page 1027.

[Table 3-51](#) describes the **Resilient Network** properties.

Table 3-51 Resilient Network dialog box properties

Property	Description
Host Name or IP Address	<p>The Host Name or IP Address of the host. The address can also be a range of IP addresses so you can configure more than one client at once. You can mix IPv4 addresses and ranges with IPv6 addresses and subnets.</p> <p>If you specify the host by name, Symantec recommends that you use the fully qualified domain name.</p> <p>Use the arrow buttons on the right side of the pane to move up or move down an item in the list of resilient networks.</p>
Resiliency	Resiliency is either ON or OFF .

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

The resilient status of each client also appears as follows:

- In the **NetBackup Administration Console**, select **NetBackup Management > Policies** in the left pane and then select a policy. In the right pane, a **Resiliency** column shows the status for each client in the policy.
- In the **NetBackup Administration Console**, select **NetBackup Management > Host Properties > Clients** in the left pane. In the right pane, a **Resiliency** column shows the status for each client.

Other NetBackup properties control the order in which NetBackup uses network addresses.

See [“Preferred Network properties”](#) on page 195.

The NetBackup resilient connections use the SOCKS protocol version 5.

Resilient connection traffic is not encrypted. Symantec recommends that you encrypt your backups. For deduplication backups, use the deduplication-based encryption. For other backups, use policy-based encryption.

Resilient connections apply to backup connections. Therefore, no additional network ports or firewall ports must be opened.

Note: If multiple backup streams run concurrently, the Remote Network Transport Service writes a large amount of information to the log files. In such a scenario, Symantec recommends that you set the logging level for the Remote Network Transport Service to 2 or less. Instructions to configure unified logs are in a different guide.

See the [NetBackup Logging Reference Guide](#).

Resilient connection resource usage

Resilient connections consume more resources than regular connections, as follows:

- More socket connections are required per data stream. Three socket connections are required to accommodate the Remote Network Transport Service that runs on both the media server and the client. Only one socket connection is required for a non-resilient connection.
- More sockets are open on media servers and clients. Three open sockets are required rather than one for a non-resilient connection. The increased number of open sockets may cause issues on busy media servers.
- More processes run on media servers and clients. Usually, only one more process per host runs even if multiple connections exist.
- The processing that is required to maintain a resilient connection may reduce performance slightly.

Specifying resilient connections

Use the following procedure to specify resilient connections for NetBackup clients.

See “[Resilient Network properties](#)” on page 209.

Alternatively, you can use the `resilient_clients` script to specify resilient connections for clients:

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

To specify resilient connections

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 2 In the right pane, select the master server on which to specify properties.
- 3 On the **Actions** menu, click **Properties**.
- 4 In the properties dialog box left pane, select **Resilient Network**.
- 5 In the **Resilient Network** dialog box, use the following buttons to manage resiliency for clients:

Add

To add resilient settings

- 1 Click **Add**.

The **Add Resilient Network Settings** dialog box appears

- 2 Enter a client host name, an IP address, or an address range.

If you specify the client host by name, Symantec recommends that you use the fully qualified domain name.

For address ranges, use Classless Inter-Domain Routing (CIDR) notation (for example, **192.168.100.0/24** or **fd00::/8**).

- 3 Ensure that the **Resiliency On** option is selected.
- 4 Click **Add**.
- 5 Repeat until you have finished entering clients or address ranges.
- 6 When you finish adding network settings, click **Close**.

Add To All

If you select multiple hosts in the **NetBackup Administration Console**, the entries in the **Resilient Network** list may appear in different colors, as follows:

- The entries that appear in black type are configured on all of the hosts.
- The entries that appear in gray type are configured on some of the hosts only.

For the entries that are configured on some of the hosts only, you can add them to all of the hosts. To do so, select them and click **Add To All**.

Change

To change resilient settings

- 1 Select the client host name, the IP address, or the address range.
- 2 Click **Change**.
The **Change Resilient Network Settings** dialog box appears
- 3 Select the desired **Resiliency** setting.
- 4 Click **OK**.

Remove

Remove the select host or address rang

- 1 Select the client host name, the IP address, or the address range.
- 2 Click **Remove**.
The client is removed immediately; a confirmation dialog box does not appear.



Move an item in the list of items

- 1 Select the client host name, the IP address, or the address range.
- 2 Click the appropriate button to move up the item or move down the item.

The order of the items in the list is significant.

See [“Resilient Network properties”](#) on page 209.

- 6 After specifying resilient connections, click **OK**.

The settings are propagated to the affected hosts through normal NetBackup inter-host communication, which can take up to 15 minutes.

- 7 If you want to begin a backup immediately, restart the NetBackup services on the master server.

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 or Hyper-V resource type. These settings apply to all policies for the currently selected master server.

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

Note: The **Resource Limit** properties apply 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.

The **Resource Limit** dialog box contains the following properties, depending on which application is selected.

For more information about the specific resource types, click **Help**.

Figure 3-58 Resource Limit dialog box with VMware application selected

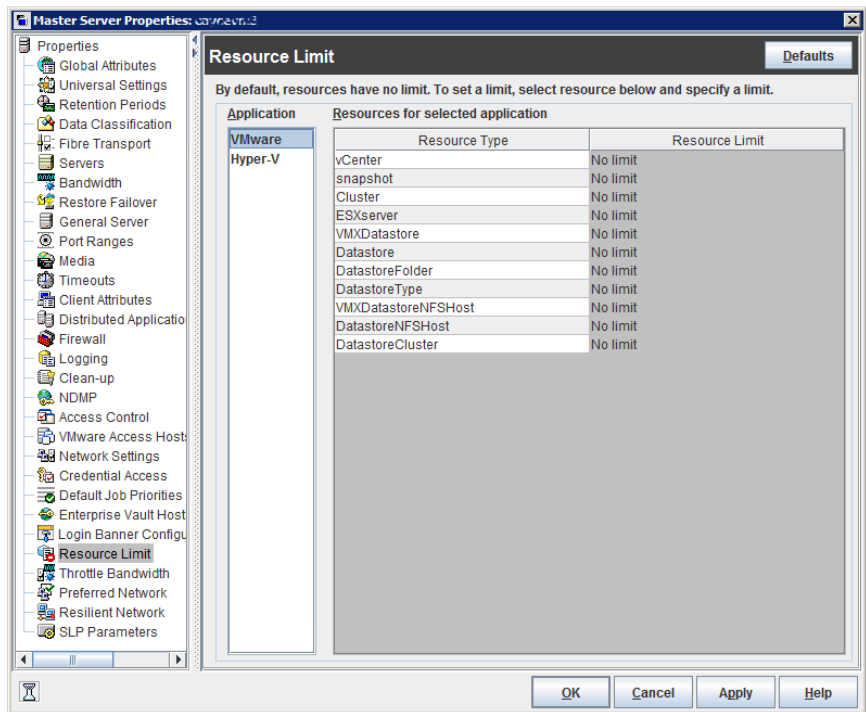
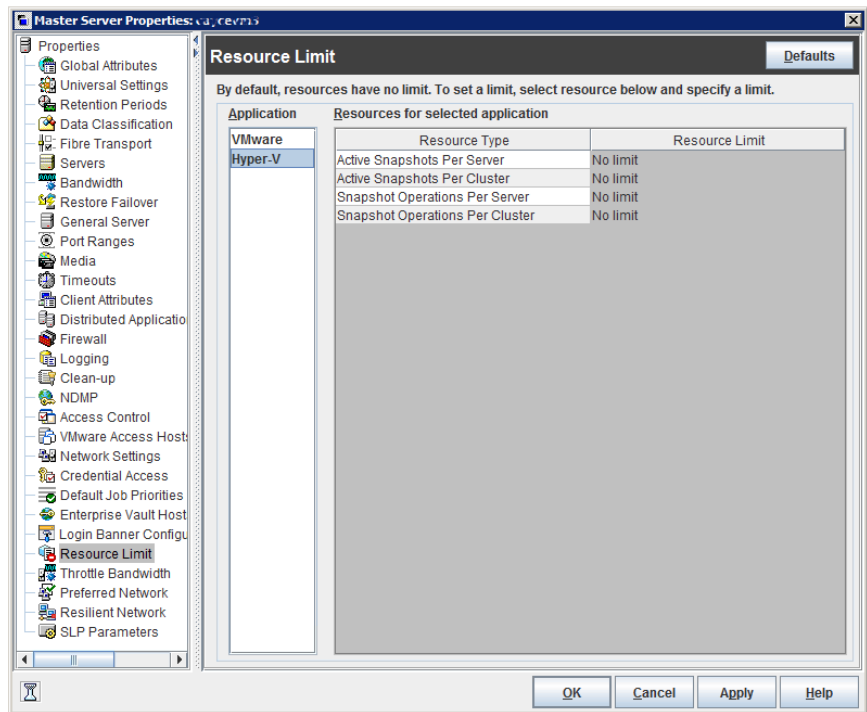



Figure 3-59 Resource Limit dialog box with Hyper-V application selected



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 3-60 Restore Failover dialog box



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

Table 3-52

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.

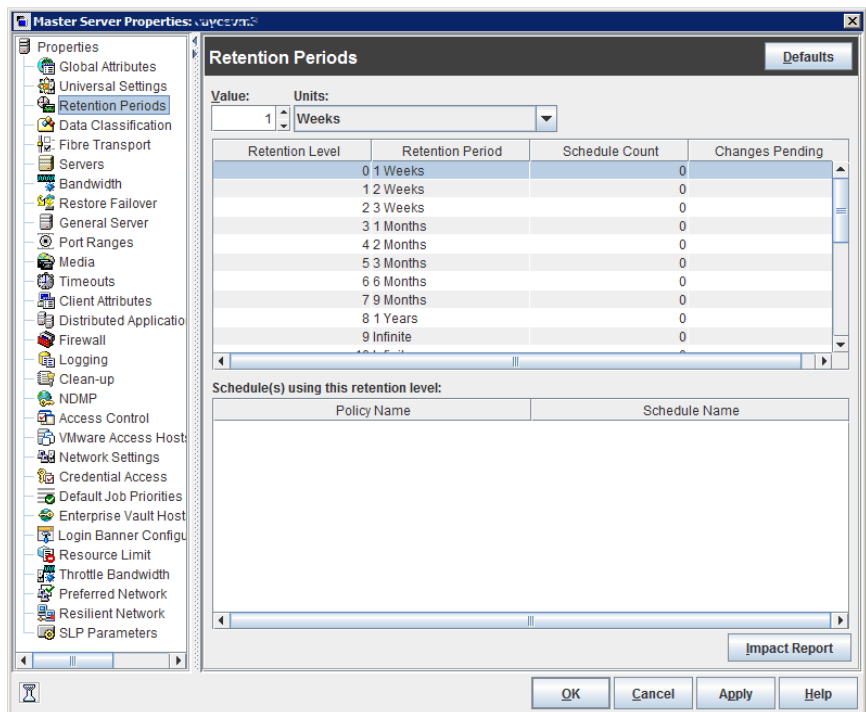
See [“About enabling automatic failover to an alternate server”](#) on page 1146.

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 3-61 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 221.

See [“Suspending or unsuspending volumes”](#) on page 506.

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 Level	The retention level number (0 through 24).
Retention Period	<p>A list of the current definitions for the 25 possible levels of retention. 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>See “Retention Periods with end dates beyond 2038, excluding Infinity” on page 222.</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>
Schedule Count	Lists the number of schedules that use the currently selected retention level.
Schedules using this retention level	Displays a list of the current policy names and schedule names that use the retention level.
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.

See [“Retention Periods with end dates beyond 2038, excluding Infinity”](#) on page 222.

The dialog box displays the names of all schedules that use the selected retention level as well as 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.

Retention Periods with end dates beyond 2038, excluding Infinity

Due to UNIX epoch time and the year 2038 problem, any retention end date that exceeds January 19, 2038 is automatically set to expire on January 19, 2038, regardless of whether the expiration end date is reached

This issue does not apply to retention levels for which the retention period is set to **Infinity**. NetBackup never expires media with a retention set to **Infinity** unless instructed to do so by the NetBackup administrator.

Scalable Storage properties

The **Scalable Storage Cloud Settings** properties contain information about encryption, metering, bandwidth throttling, and network connections between the NetBackup hosts and your cloud storage provider.

The **Scalable Storage** properties appear only if you install a NetBackup Data Protection Optimization Option license key.

The **Scalable Storage** properties apply to currently selected media servers.

Figure 3-62 Scalable Storage Cloud Settings host properties

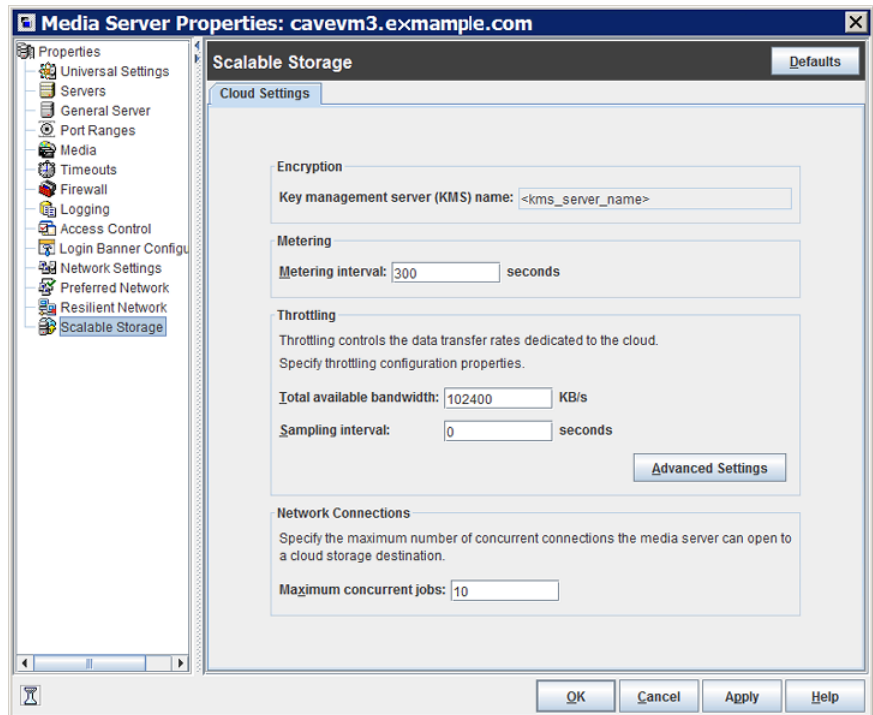


Table 3-54 describes the properties.

Table 3-54 Cloud storage host properties

Property	Description
Key Management Server (KMS) Name	If you configured the NetBackup Key Management Service (KMS), the name of the KMS server.
Metering Interval	Determines how often NetBackup gathers connection information for reporting purposes. NetBackup OpsCenter uses the information that is collected to create reports. The value is set in seconds. The default setting is 300 seconds (5 minutes). If you set this value to zero, metering is disabled.
Total Available Bandwidth	Use this value to specify the speed of your connection to the cloud. The value is specified in kilobytes per second. The default value is 102400 KB/sec.
Sampling interval	The time, in seconds, between measurements of bandwidth usage. The larger this value, the less often NetBackup checks to determine the bandwidth in use.

Table 3-54 Cloud storage host properties (*continued*)

Property	Description
Advanced Settings	<p>Click Advanced Settings to specify additional settings for throttling.</p> <p>See “Configuring advanced bandwidth throttling settings” on page 224.</p> <p>See “Advanced bandwidth throttling settings” on page 225.</p>
Maximum concurrent jobs	<p>The default maximum number of concurrent jobs that the media server can run for the cloud storage server.</p> <p>This value applies to the media server not to the cloud storage server. If you have more than one media server that can connect to the cloud storage server, each media server can have a different value. Therefore, to determine the total number of connections to the cloud storage server, add the values from each media server.</p> <p>If you configure NetBackup to allow more jobs than the number of connections, NetBackup fails any jobs that start after the number of maximum connections is reached. Jobs include both backup and restore jobs.</p> <p>You can configure job limits per backup policy and per storage unit.</p> <p>See “Limit jobs per policy (policy attribute)” on page 693.</p> <p>See “Maximum concurrent jobs storage unit setting” on page 565.</p> <p>Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of connections per media server, the number of media servers, and the job load-balancing logic. Therefore, NetBackup may not fail jobs exactly at the maximum number of connections. NetBackup may fail a job when the connection number is slightly less than the maximum, exactly the maximum, or slightly more than the maximum.</p> <p>In practice, you should not need to set this value higher than 100.</p>

For more information about NetBackup cloud storage, see the *NetBackup Cloud Administrator's Guide*:

<http://www.symantec.com/docs/DOC5332>

Configuring advanced bandwidth throttling settings

Advanced bandwidth throttling settings let you control various aspects of the connection between the NetBackup hosts and your cloud storage provider.

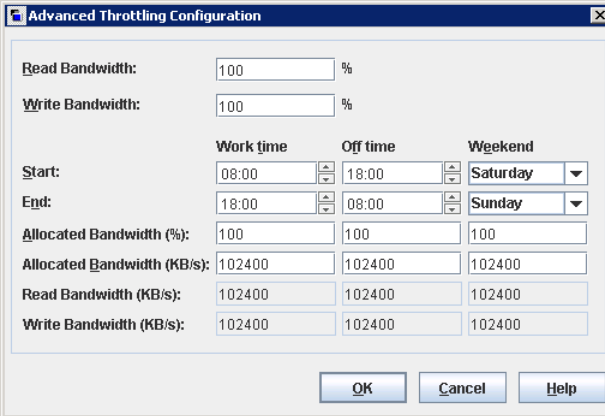
The total bandwidth and the bandwidth sampling interval are configured on the **Cloud Settings** tab of the **Scalable Storage** host properties screen.

See “[Scalable Storage properties](#)” on page 222.

To configure advanced bandwidth throttling settings

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Media Servers** in the left pane.
- 2 In the right pane, select the host on which to specify properties.
- 3 Click **Actions > Properties**.
- 4 In the properties dialog box left pane, select **Scalable Storage**.
- 5 In the right pane, click **Advanced Settings**. The **Advanced Throttling Configuration** dialog box appears.

The following is an example of the dialog box:



The dialog box titled "Advanced Throttling Configuration" contains the following fields and controls:

- Read Bandwidth:** Input field with value "100" and a "%" suffix.
- Write Bandwidth:** Input field with value "100" and a "%" suffix.
- Scheduling Section:**
 - Work time:** Start: 08:00, End: 18:00.
 - Off time:** Start: 18:00, End: 08:00.
 - Weekend:** Saturday (selected), Sunday (dropdown).
- Bandwidth Limits Table:**

	Work time	Off time	Weekend
Allocated Bandwidth (%):	100	100	100
Allocated Bandwidth (KB/s):	102400	102400	102400
Read Bandwidth (KB/s):	102400	102400	102400
Write Bandwidth (KB/s):	102400	102400	102400
- Buttons:** OK, Cancel, Help.

- 6 Configure the settings and then click **OK**.

See [“Advanced bandwidth throttling settings”](#) on page 225.

Advanced bandwidth throttling settings

The following table describes the advanced bandwidth throttling settings.

Table 3-55 Advanced Throttling Configuration settings

Property	Description
Read Bandwidth	<p>Use this field to specify the percentage of total bandwidth that read operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated.</p> <p>If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, restore or replication failures may occur due to timeouts.</p> <p>Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth.</p> <p>Default value: 100</p> <p>Possible values: 0 to 100</p>
Write Bandwidth	<p>Use this field to specify the percentage of total bandwidth that write operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated.</p> <p>If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, backup failures may occur due to timeouts.</p> <p>Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth.</p> <p>Default value: 100</p> <p>Possible values: 0 to 100</p>
Work time	<p>Use this field to specify the time interval that is considered work time for the cloud connection.</p> <p>Specify a start time and end time in 24-hour format. For example, 2:00 P.M. is 14:00.</p> <p>Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.</p>

Table 3-55 Advanced Throttling Configuration settings (*continued*)

Property	Description
Off time	<p>Use this field to specify the time interval that is considered off time for the cloud connection.</p> <p>Specify a start time and end time in 24-hour format. For example, 2:00 P.M. is 14:00.</p> <p>Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.</p>
Weekend	<p>Specify the start and stop time for the weekend.</p> <p>Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.</p>
Read Bandwidth (KB/s)	<p>This field displays how much of the available bandwidth the cloud storage server transmits to a NetBackup media server during each restore job. The value is expressed in kilobytes per second.</p>
Write Bandwidth (KB/s)	<p>This field displays how much of the available bandwidth the NetBackup media server transmits to the cloud storage server during backup jobs. The value is expressed in kilobytes per second.</p>

Servers properties

The **Servers** properties display the NetBackup server lists on selected master servers, media servers, and clients. The server lists display the NetBackup servers that each host recognizes.

The **Master Server** field contains the name of the master server for the selected host. (The name of the selected host appears in the title bar.)


The **Servers** dialog box contains the following tabs:

Table 3-56 Servers dialog box properties

Tab	Description
Additional Servers tab	<p>This tab (Figure 3-63) 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>Note: For a Fibre Transport (FT) media server that has multiple network interfaces for VLANs: Ensure that the FT server's primary host name appears before any other interface names for that FT media server host.</p> <p>See "Accessing remote servers" on page 1112.</p>
Media Servers tab	<p>This tab 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>If a server appears only on the Media servers tab, the server is a media server. Media servers can back up and restore clients, but they have limited administrative privileges.</p> <p>A server that appears on the Media servers tab and the Additional servers tab 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.</p>
OpsCenter Servers tab	<p>This tab lists the Symantec NetBackup OpsCenter servers that can access the currently selected master server.</p> <p>To add a server to the list, ensure that the NetBackup server can reach the OpsCenter server.</p> <p>For more information, see the Symantec NetBackup OpsCenter Administrator's Guide.</p>

Table 3-56 Servers dialog box properties (continued)

Tab	Description
Trusted Master Servers tab	<p>Use this tab to add the remote master servers that you trust and to view the master servers that are already trusted.</p> <p>See “About trusted master servers for Auto Image Replication” on page 987.</p> <p>See “Adding a trusted master server” on page 233.</p> <p>Note: If either the source or remote master server is clustered, you must enable inter-node communication on all of the nodes in the cluster. Do so before you add the trusted master server.</p> <p>See “Enabling NetBackup clustered master server inter-node authentication” on page 231.</p> <p>You cannot remove a trusted master server when an Import SLP is configured in that master server's domain.</p> <p>Information about Auto Image Replication and storage lifecycle policies is available.</p> <p>See “About NetBackup Auto Image Replication” on page 975.</p> <p>See “About storage lifecycle policies” on page 600.</p> <p>If you use NetBackup Access Control, read about access control with Auto Image Replication.</p> <p>See “Using NetBackup Access Control (NBAC) with Auto Image Replication” on page 993.</p>

Figure 3-63 Additional servers tab in the Servers dialog box

Adding a server to a servers list

Depending on the tab that is selected, you can add a master server, media server, client, or OpsCenter server to the server list in the **Additional servers** tab, the **Media servers** tab, or the **OpsCenter servers** tab. All tabs appear in the **Servers** properties dialog box.

To add a server to a 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 that you want to modify.
- 4 In the properties dialog box, in the left pane, click **Servers**.
- 5 Select the tab that contains the server list that you want to modify.

6 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.

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

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

9 Click **Close**.

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

Removing a server from a server 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 remove a server from a server 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 Select a server from the **Additional servers** list, the **Media servers** list, or the **OpsCenter servers** list.
- 6 Click **Remove**.

Enabling NetBackup clustered master server inter-node authentication

NetBackup requires inter-node authentication among the master servers in a cluster. For authentication, you must provision an authentication certificate on all of the nodes of the cluster. The certificates are used to establish SSL connections between the NetBackup hosts. The inter-node authentication allows the following NetBackup functionality:

NetBackup Administration Console	The NetBackup Administration Console in master server clusters requires the NetBackup authentication certificates for correct functionality.
Targeted A.I.R. (Auto Image Replication)	<p>Auto Image Replication in which a master server is in a cluster requires inter-node authentication among the hosts in that cluster. The NetBackup authentication certificates provide the means to establish the proper trust relationships.</p> <p>Provision the certificates on the cluster hosts before you add the trusted master server. This requirement applies regardless of whether the clustered master server is the source of the replication operation or the target.</p> <p>See “About trusted master servers for Auto Image Replication” on page 987.</p>

See [“Adding a trusted master server”](#) on page 233.

To enable clustered master server inter-node authentication

- ◆ On the active node of the NetBackup master server cluster, run the following NetBackup command:
 - **Windows:** `install_path\Veritas\NetBackup\bin\admincmd\bpnbaz -setupat`
 - **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpnbaz -setupat`

NetBackup creates the certificates on every node in the master server cluster.

The following is example output:

```
# bpnbaz -setupat
You will have to restart NetBackup services on this machine after
the command completes successfully.
Do you want to continue(y/n)y
Gathering configuration information.
Please be patient as we wait for 10 sec for the security services
to start their operation.
Generating identity for host 'bitl.remote.example.com'
Setting up security on target host: bitl.remote.example.com
nbatd is successfully configured on NetBackup Master Server.
Operation completed successfully.
```

Adding a trusted master server

You can configure a trust relationship between multiple NetBackup domains. To do so, in a source domain you specify the remote master servers with which you want to add a trust relationship. Use the following procedure in the source domain to add a remote master server as a trusted master server.

A trust relationship between domains helps with replication operations.

See [“About trusted master servers for Auto Image Replication”](#) on page 987.

Note: If either the source or remote master server is clustered, you must enable inter-node communication on all of the nodes in the cluster. Do so before you add the trusted master server.

See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 231.


To add a trusted master server

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 2 In the right pane, select the master server.
- 3 On the **Actions** menu, click **Properties**.
- 4 In the properties dialog box left pane, select **Servers**.
- 5 In the **Servers** dialog box, select the **Trusted Master Servers** tab.

- 6 On the **Trusted Master Servers** tab, click **Add**.

The **Add a New Trusted Master Server** dialog box appears.

The following is an example of the dialog box:



- 7 In the **Add a New Trusted Master Server** dialog box, enter the following and then click **OK**:
- The fully-qualified host name of the remote master server.
 - The logon account **User name** of the remote master server host.
 - The **Password** for the logon account of the remote master server host.
- 8 Repeat step 6 and step 7 for each master server with which you want to add a trust relationship.
- 9 When you finish adding trusted master servers, click **OK**.

Removing a trusted master server

You can remove a remote master server with which you have configured a trust relationship.

A trust relationship between domains helps with replication operations.

See [“About trusted master servers for Auto Image Replication”](#) on page 987.

Two procedures exist, as follows:

The remote master server is online	<p>If the remote master server is online and active, use the NetBackup Administration Console to remove the trusted master server. This method also removes the reciprocal trust relationship.</p> <p>See “To remove a trusted master server for inter-domain operations” on page 235.</p>
The remote master server is offline	<p>If the remote master server is offline, use the NetBackup <code>nbseccmd</code> command to remove the trusted master server. This method does not remove the reciprocal trust relationship.</p> <p>See “To remove a trusted master server that is offline” on page 236.</p>

To remove a trusted master server for inter-domain operations

- 1 Ensure that all replication jobs to the trusted master are complete. You can use `nbstlutil stlilist` to list the state of all storage lifecycle policy-managed operations. To cancel jobs use `nbstlutil cancel`.

See the [NetBackup Commands Reference Guide](#) for information about the `nbstlutil` command.

- 2 Delete all storage lifecycle policies (SLPs) that use the trusted master as a destination.


Note: Before deleting a storage lifecycle policy, ensure that there are no backup policies that indicate the SLP for the **Policy storage**.

- 3 In the **NetBackup Administration Console**, expand **NetBackup Management > Host Properties > Master Servers** in the left pane.
- 4 In the right pane, select the master server.
- 5 On the **Actions** menu, click **Properties**.
- 6 In the properties dialog box left pane, select **Servers**.
- 7 In the **Servers** dialog box, select the **Trusted Master Servers** tab.

- 8 On the **Trusted Master Servers** tab, select the trusted master server that you want to remove and click **Remove**.

The **Remove Trusted Master Server** dialog box appears.

The following is an example of the dialog box:



In the **Remove Trusted Master Server** dialog box, enter the credentials for that master server host and then click **Remove**.

- 9 Repeat step 8 for each master server from which you want to remove a trust relationship.
- 10 When you finish removing trusted master servers, click **OK**.

To remove a trusted master server that is offline

- ◆ On the source master server, enter the following NetBackup `nbseccmd` command:

```
UNIX: /usr/opensv/netbackup/bin/admincmd/nbseccmd
-setuptrustedmaster -remove -masterserver master_server_name
-remotemasterserver remote_master -localonly
```

```
Windows: install_path\NetBackup\bin\admincmd\nbseccmd
-setuptrustedmaster -remove -masterserver master_server_name
-remotemasterserver remote_master -localonly
```


Changing the master server that performs backups and restores for a client

Use the **Make Master** option to change the master server that performs backups and restores for a client. This option does not change a host into a master server.

This option is useful in a disaster recovery situation or in a NetBackup environment where Auto Image Replication is configured. For example, select a client in the source domain, then use **Make Master** to temporarily point the client to the master server of the target domain. After changing the master server, restores from the target domain can be initiated.

To change the master server that a client uses for backups and restores

- 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**.

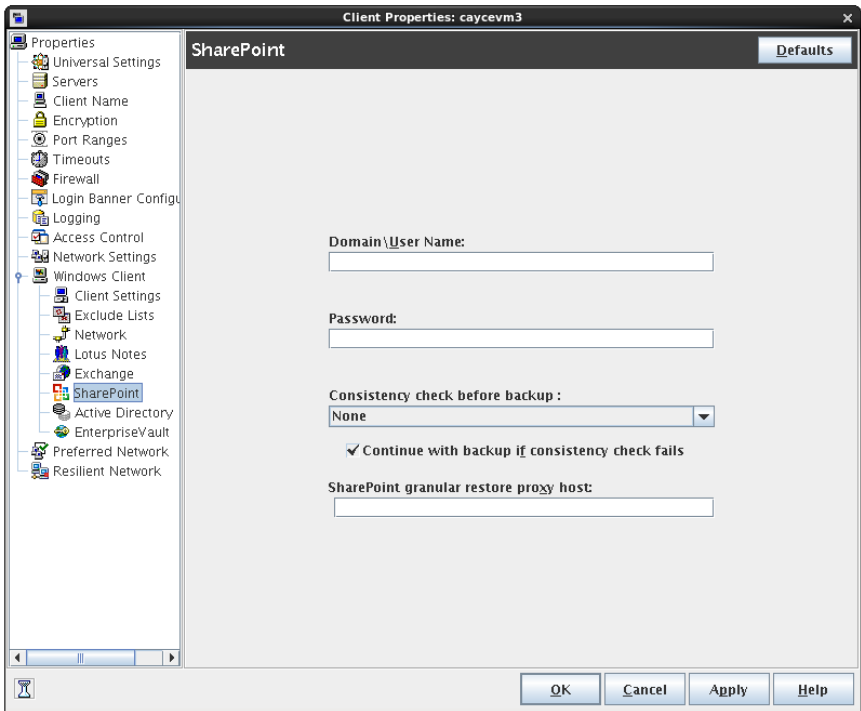
In the configuration file, the new master server appears as the first server entry in the list.

Changing the master server does not prevent the former master server from initiating backups for the client. As long as that server continues to be listed on the client's server list, the master server can perform backups.

SharePoint properties

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

Figure 3-64 SharePoint dialog box



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

Table 3-57 SharePoint dialog box properties

Property	Description
Domain\User name	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	<p>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.</p> <p>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.</p>

Table 3-57 SharePoint dialog box properties (*continued*)

Property	Description
SharePoint granular restore proxy host	For VMware backups that protect Federated SharePoint configurations, provide the name of the back-end SQL server. This server acts as the granular restore proxy host for the catalog hosts (front-end servers in the farm).

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.

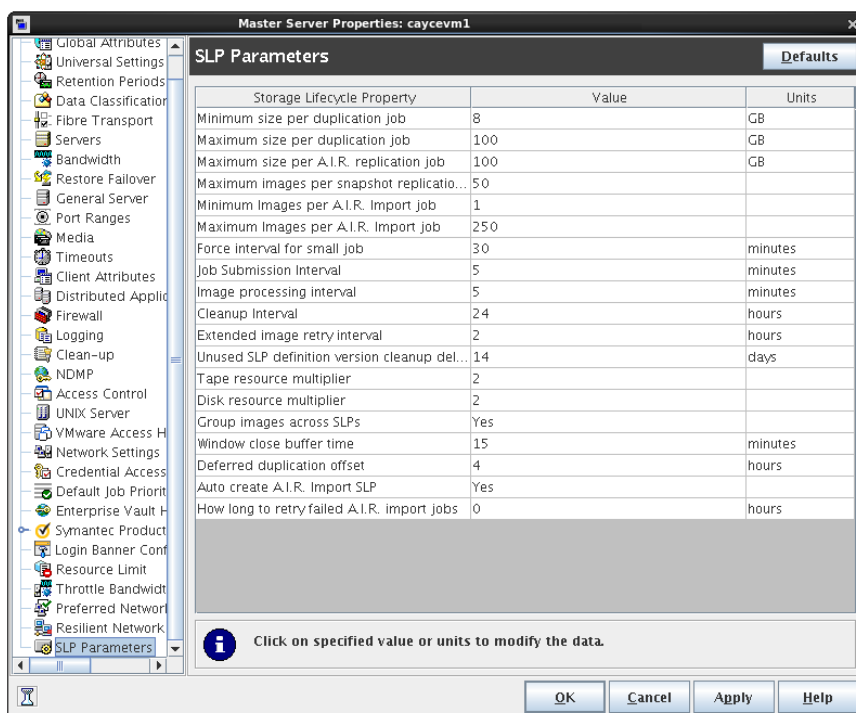
Table 3-58 Consistency check options

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.

SLP Parameters properties

The **SLP Parameters** properties in the **NetBackup Administration Console** allow administrators to customize how storage lifecycle policies are maintained and how SLP jobs run. These properties apply to the SLPs of the currently selected master servers.

Figure 3-65 SLP Parameters dialog box



In NetBackup 7.6 and earlier, the parameters that appear in the **SLP Parameters** were configured in the `LIFECYCLE_PARAMETERS` configuration file. The host properties replace that file. An upgrade to 7.6 migrates the settings in the `LIFECYCLE_PARAMETERS` file to the **SLP Parameters**.

To change the parameters, either change them in the host properties or use the command line. If either method is used to change a storage lifecycle policy parameter default, the change creates the `nbcl.conf` configuration file.

The `nbcl.conf` file is found in the following locations, but is not present if the default of any parameter has not been changed.

- On Windows:


```
install_path\veritas\NetBackup\var\global\nbcl.conf
```
- On UNIX:


```
/usr/opensv/var/global/nbcl.conf
```

To use the command-line method, use the `nbgetconfig` and the `nbsetconfig` commands to change the defaults. For information about these commands, see the [NetBackup Commands Reference Guide](#).

Table 3-59 describes the **SLP Parameters** and the default of each. It also lists the configuration options and the syntax to use with the command-line method.

Use the drop-down menu in the **Units** column to change the units of measurement for size or time.

Table 3-59 SLP Parameters

Property	Description
Minimum size per duplication job	<p>The smallest batch size that can run as a single duplication job. The job does not run until enough images accumulate to reach this minimum batch size or until the Force interval for small jobs time is reached. Minimum: 1 kilobyte; no maximum size. Default: 8 gigabytes.</p> <p>Configuration option default: <code>SLP.MIN_SIZE_PER_DUPLICATION_JOB = 8 GB</code></p>
Maximum size per duplication job	<p>The largest batch size that can run as a single duplication job. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes.</p> <p>Configuration entry default: <code>SLP.MAX_SIZE_PER_DUPLICATION_JOB = 100 GB</code></p>
Maximum size per A.I.R. replication job	<p>The largest batch size that can run as a single job for Auto Image Replication. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes.</p> <p>Configuration entry default: <code>SLP.MAX_SIZE_PER_BACKUP_REPLICATION_JOB = 100 GB</code></p>
Maximum images per snapshot replication job	<p>The largest number of images in a single batch that can run as a single job. Default: 50 images, with no minimum number or maximum number.</p> <p>Use this parameter with the Limit I/O streams disk pool option which limits the number of jobs that can run concurrently to each volume in the disk pool.</p> <p>Configuration entry default: <code>SLP.MAX_IMAGES_PER_SNAPSHOT_REPLICATION_JOB = 50</code></p>
Minimum images per A.I.R. Import job	<p>The fewest number of images in a single batch that can run as an Auto Image Replication import job. The job does not run until either the minimum size is reached or the Force interval for small jobs time is reached. Minimum: 1 image; no maximum number of images. Default: 1 image.</p> <p>Configuration entry default: <code>SLP.MIN_IMAGES_PER_IMPORT_JOB = 1</code></p>
Maximum images per A.I.R. Import job	<p>The largest number of images in a single batch that can run as an Auto Image Replication import job. Minimum: 1 job; no maximum number of images. Default: 250 images.</p> <p>Configuration entry default: <code>SLP.MAX_IMAGES_PER_IMPORT_JOB = 250</code></p>

Table 3-59 SLP Parameters (*continued*)

Property	Description
Force interval for small jobs	<p>The age that the oldest image in a batch must reach after which the batch is submitted as a duplication job. This value prevents many small duplication jobs from running at one time or running too frequently. It also prevents NetBackup from waiting too long before it submits a small job. Default: 30 minutes, with no minimum number or maximum number.</p> <p>Configuration entry default: <code>SLP.MAX_TIME_TIL_FORCE_SMALL_DUPLICATION_JOB = 30 MINUTES</code></p>
Job submission interval	<p>Indicates the frequency of the job submission for all operations. No minimum interval or maximum interval. Default: 5 minutes.</p> <p>By default, all jobs are processed before more jobs are submitted. Increase this interval to allow NetBackup to submit more jobs before all jobs are processed. Set the interval when the list of available images is scanned for those that can be batched together and jobs submitted. A shorter interval allows for a better response to changing system workloads at the cost of increased processing.</p> <p>Configuration entry default: <code>SLP.JOB_SUBMISSION_INTERVAL = 5 MINUTES</code></p>
Image processing interval	<p>The number of minutes between image-processing sessions. Set the interval when newly created images are recognized and set up for SLP processing. Default: 5 minutes.</p> <p>Configuration entry default: <code>SLP.IMAGE_PROCESSING_INTERVAL = 5 MINUTES</code></p>
Cleanup interval	<p>The time between when a job finishes and before NetBackup removes the job artifacts for the completed job. No minimum interval or maximum interval. Default: 24 hours.</p> <p>Configuration entry default: <code>SLP.CLEANUP_SESSION_INTERVAL = 24 HOURS</code></p>
Extended image retry interval	<p>The amount of time to wait before an unsuccessful operation is added to the first job that runs after the delay. (This applies to all SLP jobs.) The extra time gives the administrator additional time to solve a problem that prevents job completion. No minimum interval or maximum interval. Default: 2 hours.</p> <p>Configuration entry default: <code>SLP.IMAGE_EXTENDED_RETRY_PERIOD = 2 HOURS</code></p>
Unused SLP definition version cleanup delay	<p>Concerns the deletion of SLP versions where a more recent version exists. The setting controls how long a version must be inactive before NetBackup deletes it. Default: 14 days.</p> <p>Configuration entry default: <code>SLP.VERSION_CLEANUP_DELAY = 14 DAYS</code></p> <p>See "Deleting old storage lifecycle policy versions" on page 660.</p>

Table 3-59 SLP Parameters (*continued*)

Property	Description
Tape resource multiplier	<p>Limits the number of concurrently active duplication jobs that can access a single tape media storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two).</p> <p>Configuration entry default: <code>SLP.TAPE_RESOURCE_MULTIPLIER = 2</code></p>
Disk resource multiplier	<p>Limits the number of concurrently active duplication jobs that can access a single disk storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two).</p> <p>Configuration entry default: <code>SLP.DISK_RESOURCE_MULTIPLIER = 2</code></p>
Group images across SLPs	<p>If this parameter is set to Yes (default), multiple SLPs of the same priority can be processed in the same job. If No, batching can occur only within a single SLP.</p> <p>Configuration entry default: <code>SLP.DUPLICATION_GROUP_CRITERIA = 1</code></p> <p>Configuration entry for no, do not allow batching: <code>SLP.DUPLICATION_GROUP_CRITERIA = 0</code></p>
Window close buffer time	<p>Sets the amount of time before a window closes when NetBackup does not submit new jobs using that window. Minimum 2 minute; maximum: 60 minutes. Default: 15 minutes.</p> <p>Configuration entry default: <code>SLP.WINDOW_CLOSE_BUFFER_TIME = 15 MINUTES</code></p>
Deferred duplication offset time	<p>For deferred operations, jobs are submitted x time before the source copy is due to expire. Default: 4 hours.</p> <p>Configuration entry default: <code>SLP.DEFERRED_DUPLICATION_OFFSET_TIME = 4 HOURS</code></p>
Auto create A.I.R. Import SLP	<p>Used for Auto Image Replication, indicates whether an SLP (that contains an Import operation) is created automatically in the target domain if no SLP is configured there. Default: Yes, an SLP is created in the target domain.</p> <p>Configuration entry default: <code>SLP.AUTO_CREATE_IMPORT_SLP = 1</code></p>
How long to retry failed A.I.R. import jobs	<p>How long NetBackup retries an Import job before it stops and deletes the record. After the initial four attempts, the retries become less frequent. Default: 0 (do not retry after the initial four attempts).</p> <p>Configuration entry default: <code>SLP.REPLICA_METADATA_CLEANUP_TIMER = 0 HOURS</code></p>

Command-line units of measurement for the SLP parameters

The abbreviations are case-insensitive for units of measurement.

The following abbreviations can be used where sizes are indicated:

bytes	kb	kilobyte	kilobyte(s)	kilobytes	mb	megabyte
megabyte(s)	megabytes	gb	gigabyte	gigabyte(s)	gigabytes	tb
terabyte	terabyte(s)	terabytes	pb	petabyte	petabyte(s)	petabytes

The following abbreviations can be used where units of time are indicated:

sec	second	second(s)	seconds	min	minute	minute(s)	minutes
hour	hour(s)	hours	day	day(s)	days	mon	month
month(s)	months	week	week(s)	weeks	year	year(s)	years

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.

Note: Restart `nbstserv` after making changes to the underlying storage for any operation in an SLP.

One objective of the batching logic is to prevent media contention for tape operations, including virtual tape libraries (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 **Group images across SLPs** parameter in the **SLP Parameters** host properties. By default, multiple storage lifecycle policies with the same priority can be batched together.

See “[SLP Parameters properties](#)” on page 239.

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, set the **Group images across SLPs** parameter to **No** in the **SLP Parameters** host properties.

Deprecated storage lifecycle policy parameters

Three storage lifecycle policy parameters were deprecated in NetBackup 7.6 due to changes in SLP processing.

`DUPLICATION_SESSION_INTERVAL_MINUTES` is no longer effective in controlling the timing of SLP processing sessions. Two parameters now control session processing:

- `SLP.IMAGE_PROCESSING_INTERVAL` controls how often new SLP-managed images are recognized and set up for SLP processing.
- `SLP.JOB_SUBMISSION_INTERVAL` controls how often SLP processing jobs are submitted for execution.

Import operations are controlled by the same session that performs other SLP operations. As a result, `IMPORT_SESSION_TIMER` and `IMPORT_EXTENDED_RETRY_SESSION_TIMER` are no longer used.

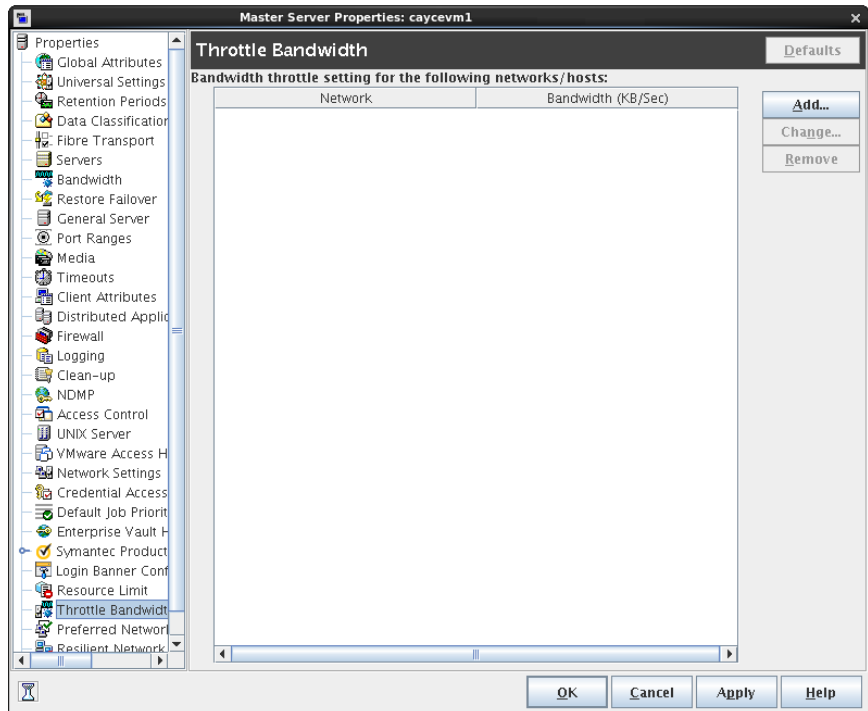
As with other SLP operations, `SLP.IMAGE_PROCESSING_INTERVAL` and `SLP.JOB_SUBMISSION_INTERVAL` now control the frequency of SLP session operations. `SLP.IMAGE_EXTENDED_RETRY_PERIOD` controls the retry behavior for all SLP operations including import operations.

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 3-66 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.

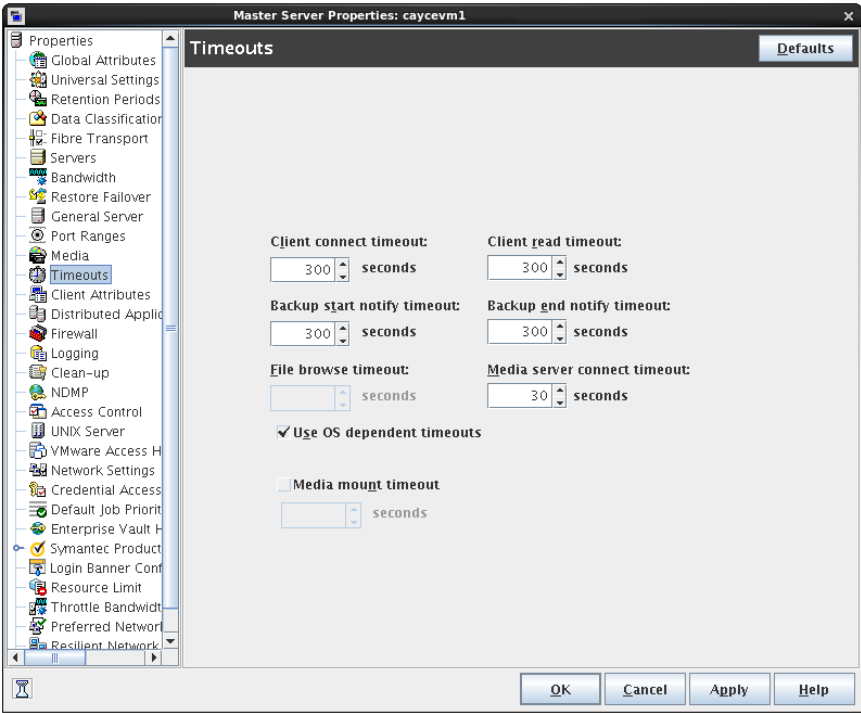
Bandwidth (KB/Sec)

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 3-67 Timeouts dialog box



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

Table 3-60 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.

Table 3-60 Timeouts dialog box properties (*continued*)

Property	Description
Backup start notify timeout	<p>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.</p> <p>Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running.</p>
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>

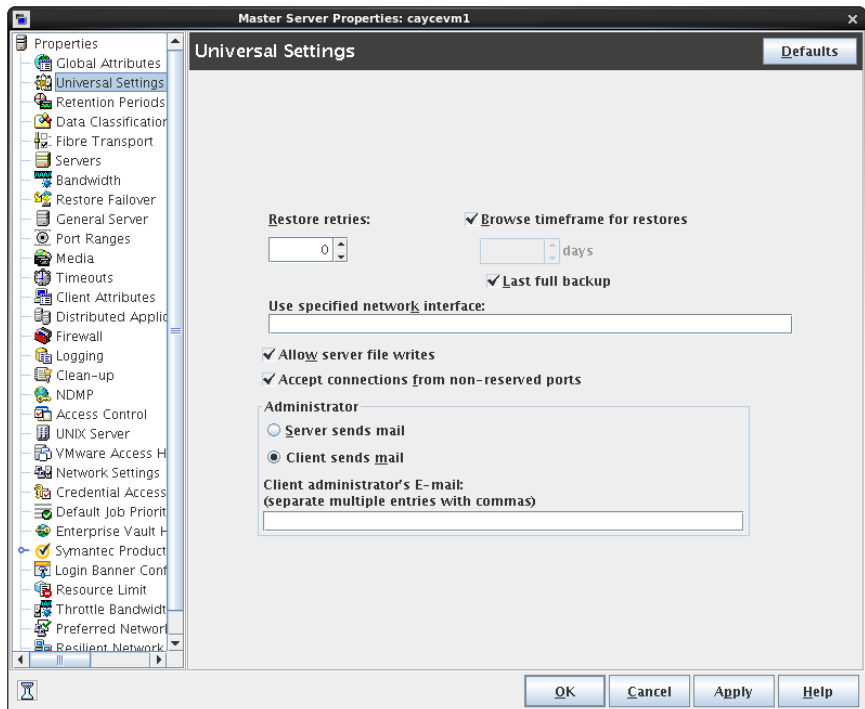
Table 3-60 Timeouts dialog box properties (*continued*)

Property	Description
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>If the server does not get a response from a client within the Client read timeout period, the backup or restore operation, for example, fails.</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 properties for UNIX clients" on page 99.</p> <p>Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running.</p>
Backup end notify timeout	<p>Specifies the number of seconds that the server waits for the <code>bpend_notify</code> script on a client to complete. The default is 300 seconds.</p> <p>Note: If this timeout is changed, verify that Client read timeout is set to the same or higher value.</p>
Media server connect timeout	<p>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.</p>

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 3-68 Universal Settings dialog box



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

Table 3-61 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 80.</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 689.</p>

Table 3-61 Universal Settings dialog box properties (*continued*)

Property	Description
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>
Use specified network interface	<p>Specifies the network interface that NetBackup uses to connect to another NetBackup client or server. A NetBackup client or server can have more than one network interface. To force NetBackup connections to be made on a specific network interface, use this entry to specify the network host name of that interface. By default, the operating system determines the one to use.</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>

Table 3-61 Universal Settings dialog box properties (*continued*)

Property	Description
Logging the status of a redirected restore	<p>UNIX only:</p> <p>No progress log is produced for a redirected restore if the requesting server is not included in the server list for the server that performs the restore. (A progress log is an entry in the Task Progress tab of the Backup, Archive, and Restore client interface.)</p> <p>Without the entry in the server list, the restoring server has no access to write the log files to the requesting server.</p> <p>Consider the following solutions:</p> <ul style="list-style-type: none"> ■ To produce a progress log, add the requesting server to the server list and log into the requesting server. <ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Host Properties > Master Servers. ■ In the details pane, double-click on master server you want to modify, and click Servers from the options pane. ■ 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.
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>
Client sends mail	<p>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.</p>
Server sends mail	<p>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.</p> <p>See “Global Attributes properties” on page 150.</p>
Client administrator’s email	<p>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.</p>

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:
 - **Additional servers**
From the **Media servers** list, click **Add**.
 - **Media servers**
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.

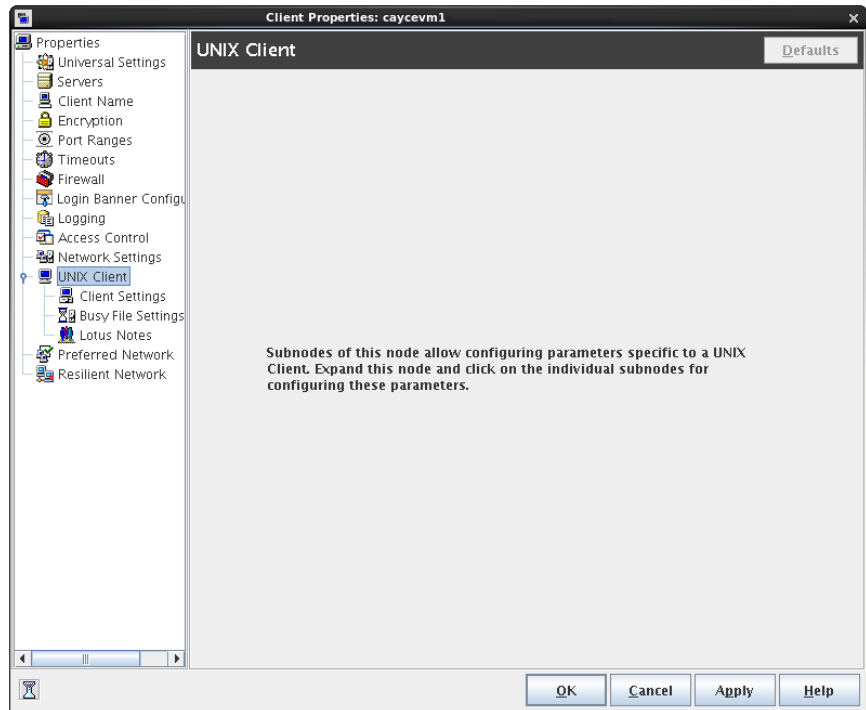
Check the **Activity Monitor** to determine the success of the restore operation.

See [“About the Jobs tab”](#) on page 1020.

UNIX Client properties

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

Figure 3-69 UNIX Client dialog box



See [“Client Settings properties for UNIX clients”](#) on page 99.

See [“Busy File Settings properties”](#) on page 77.


See [“Lotus Notes properties”](#) on page 177.

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 3-70 UNIX Server dialog box

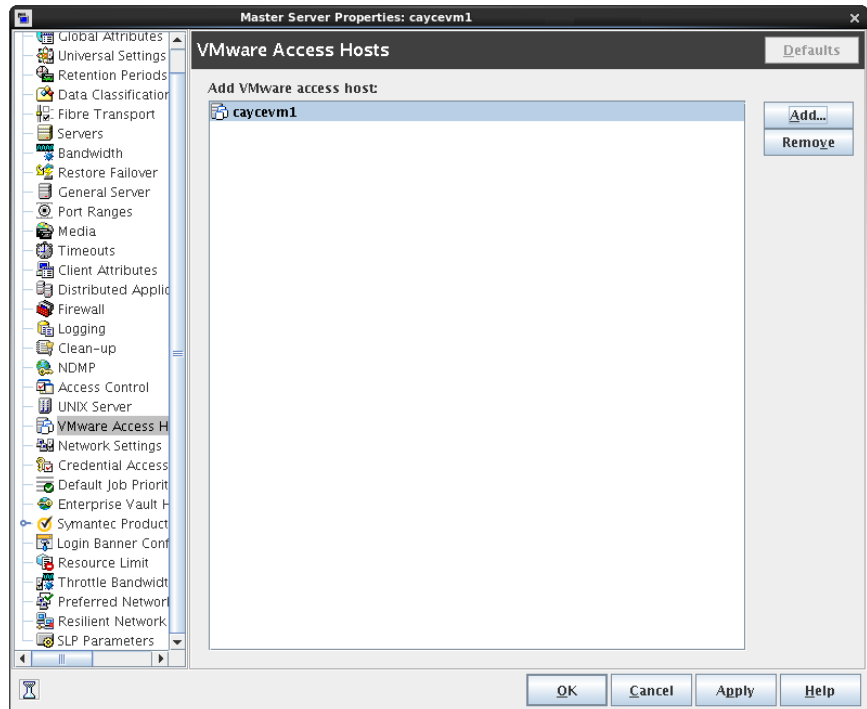


See [“Follow NFS \(policy attribute\)”](#) on page 696.

VMware Access Hosts properties

Use the **VMware Access Hosts properties** in the **NetBackup Administration Console** to add or remove VMware backup hosts. These properties apply to currently selected master servers.

These properties appear when the NetBackup Enterprise Client license is installed.

Figure 3-71 Vmware Access Hosts dialog box

The backup host is a NetBackup client that performs backups on behalf of the virtual machines. (This host was formerly known as the VMware backup proxy server.) The backup host is the only host on which NetBackup client software is installed. As an option, the backup host can also be configured as a NetBackup master server or media server.

The backup host is referred to as the recovery host when it performs a restore. You can add servers to and remove servers from the access hosts list:

Add Click **Add** and enter the fully qualified domain name of the backup host.


Remove Select the backup host in the list and click **Remove**.

For more information, see the [NetBackup for VMware Administrator's Guide](#).

Windows Client properties

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

Figure 3-72 Windows Client dialog box



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

Configuration options not found in the Host Properties

Most NetBackup configuration options can be found in the **Host Properties** of the **NetBackup Administration Console**. However, some options cannot be accessed in the **Host Properties**.

To change the default value for an option that is not found in the **Host Properties**, first use the `nbgetconfig` command to obtain a list of configuration options. Then use `nbsetconfig` to change the options as needed.

For information about these commands, see the [NetBackup Commands Reference Guide](#).

The following NetBackup administration options cannot be configured by using the **NetBackup Administration Console**.

AUTO_ADD_ALL_ALIASES_FOR_CLIENT	See “ AUTO_ADD_ALL_ALIASES_FOR_CLIENT option for NetBackup servers” on page 266.
BPARCHIVE_POLICY	See “ BPARCHIVE_POLICY option for NetBackup clients” on page 334.
BPARCHIVE_SCHED	See “ BPARCHIVE_SCHED option for NetBackup clients” on page 334.
BPBACKUP_POLICY	See “ BPBACKUP_POLICY option for NetBackup clients” on page 335.
BPBACKUP_SCHED	See “ BPBACKUP_SCHED option for NetBackup clients” on page 336.
BPDBJOBS_COLDEFS	See “ BPDBJOBS_COLDEFS options for UNIX master servers” on page 268.
CHECK_RESTORE_CLIENT	See “ CHECK_RESTORE_CLIENT option for NetBackup servers” on page 277.
EMMSERVER	See “ EMMSERVER option for NetBackup servers” on page 291.
GENERATE_ENGLISH_LOGS	See “ GENERATE_ENGLISH_LOGS option for NetBackup servers and clients” on page 293.
IGNORE_XATTR	See “ IGNORE_XATTR option for NetBackup clients” on page 350.
INFORMIX_HOME	See “ INFORMIX_HOME option for NetBackup clients” on page 353.
KEEP_DATABASE_COMM_FILE	See “ KEEP_DATABASE_COMM_FILE option for NetBackup clients” on page 354.
REPORT_CLIENT_DISCOVERIES	See “ REPORT_CLIENT_DISCOVERIES option for NetBackup clients” on page 359.
RESUME_ORIG_DUP_ON_OPT_DUP_FAIL	See “ RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers” on page 315.
SYBASE_HOME	See “ SYBASE_HOME option for NetBackup clients” on page 362.

ULINK_ON_OVERWRITE

See “[ULINK_ON_OVERWRITE option for NetBackup servers](#)” on page 327.

vm.conf file (media and device configuration options) See the [NetBackup Administrator's Guide, Volume II](#).

About using commands to change the configuration options on UNIX clients and servers

When commands (`nbsetconfig` or `bpsetconfig`) are used to change the configuration options on UNIX NetBackup servers or clients, the commands change the appropriate configuration files.

Most options are found in the following configuration file:

```
/usr/opensv/netbackup/bp.conf
```

If a single UNIX system is running as both a client and a server, the `bp.conf` file contains options for both the client and the server.

The `bp.conf` file observes the following syntax:

- Use the `#` symbol to comment out lines.
- Any number of spaces or tabs are allowed on either side of `=` signs.
- Blank lines are allowed.
- Any number of blanks or tabs are allowed at the start of a line.

Each nonroot user on a UNIX client can also have a personal `bp.conf` file in their home directory:

```
$HOME/bp.conf
```

The options in personal `bp.conf` files apply only to user operations. During a user operation, NetBackup checks the `$HOME/bp.conf` file before

```
/usr/opensv/netbackup/bp.conf.
```

Root users do not have personal `bp.conf` files. NetBackup uses the `/usr/opensv/netbackup/bp.conf` file for root users.

Stop and restart all NetBackup daemons and utilities on the server after you make a change to the `bp.conf` file on a UNIX master server. This action ensures that all of the NetBackup processes use the new `bp.conf` values. This action is not required for changes to `bp.conf` files on a client or to a `$HOME/bp.conf` file on the master server.

The `SERVER` option must be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX clients and servers. During installation, NetBackup sets the `SERVER` option to the name of the master server where the software is installed. It is the only required option in the `bp.conf` files. NetBackup uses internal software defaults for all options in the `bp.conf` file, except `SERVER`.

The `SERVER` entries must be the same on all servers in a master and a media server cluster. Symantec recommends (but does not require) that all other entries also match on all servers. (The `CLIENT_NAME` option is an exception.)

Configuration options for NetBackup servers

The following topics are about configuration options for NetBackup servers. Nearly all of these options can also be set in the Host Properties in the **NetBackup Administration Console**.

See [“Configuration options not found in the Host Properties”](#) on page 257.

About robust logging

The Robust logging option indicates to NetBackup to limit the amount of disk space that a log directory consumes. When a log file grows to the specified maximum size, the log file is closed and 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.

In NetBackup 7.6.1 or earlier versions, you can enable or disable robust logging in the **NetBackup Administration Console** in the Logging host properties. The **Enable Robust Logging** option was disabled by default.

See [“Logging properties”](#) on page 161.

Starting from NetBackup 7.7, the **Enable Robust Logging** option is removed from the Logging host properties because the option is always enabled. The error logs are always created based on the robust logging configuration parameters.

However, the following configuration parameter still appears when you run the `bpgetconfig` command: `ENABLE_ROBUST_LOGGING = NO`

Regardless of the value of the `ENABLE_ROBUST_LOGGING` configuration parameter, robust logging is always enabled.

Before NetBackup 7.7, robust logging was applicable only for unified logging. Robust logging is also enabled for legacy logging for NetBackup servers and clients.

Known issue: Robust logging is enabled only for Informix and SAP database agents for NetBackup. For the remaining NetBackup database agents, robust logging is not enabled.

Robust logging functionality comprises the following configuration parameters:

<code>MAX_LOGFILE_SIZE</code>	<p>This configuration parameter specifies the maximum size that you want to set for a log file in MB.</p> <p>The default value for <code>MAX_LOGFILE_SIZE</code> is 50 MB.</p> <p>See “MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging” on page 300.</p>
<code>MAX_NUM_LOGFILES</code>	<p>This configuration parameter specifies the maximum number of log files that you want to be created in NetBackup.</p> <p>The default value for <code>MAX_NUM_LOGFILES</code> is infinite.</p> <p>See “MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging” on page 301.</p>

Note: Initially, `MAX_LOGFILE_SIZE` and `MAX_NUM_LOGFILES` parameter entries are not present in the output for `bpgetconfig`. In this case, they are set to their default values. In UNIX, you can update the `bp.conf` configuration file using the `bpsetconfig` command to set these parameters to different values. In Windows, you can add the parameters using the `bpsetconfig` command in the registry.

About robust logging functionality after upgrade

After an upgrade to NetBackup 7.7 or later, **Enable Robust Logging** is enabled. This is true even if the setting had been disabled before the upgrade.

After an upgrade to NetBackup 7.7 or later, the `MAX_LOGFILE_SIZE` is set to a new default.

The following points concern various upgrade scenarios to NetBackup 7.7 or later with respect to `MAX_LOGFILE_SIZE`:

- If robust logging was enabled before the upgrade:
 - The old unified logging default value for `MaxLogFileSizeKB` was 51200 KB. After the upgrade, the legacy logging default for `MAX_LOGFILE_SIZE` is 50 MB.
 - The old unified logging default value for `NumberOfLogFiles` was 3. After the upgrade, the legacy logging value for `MAX_NUM_LOGFILES` is 3.
- If robust logging was disabled before the upgrade:

- After the upgrade, robust logging is enabled with the following configuration settings:

```
MAX_LOGFILE_SIZE = 50 MB
MAX_NUM_LOGFILES = INFINITE
```

ALLOW_MEDIA_OVERWRITE option for NetBackup servers

This option overrides the NetBackup overwrite protection for various media formats on removable media.

Table 3-62 ALLOW_MEDIA_OVERWRITE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>ALLOW_MEDIA_OVERWRITE = media_format</pre> <p>This option should appear only once in the configuration file.</p>
Example	<p>On the master server (and media servers if applicable), add the following entry to permit overwriting the <code>cpio</code> format:</p> <pre>ALLOW_MEDIA_OVERWRITE = CPIO</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers or Media Servers > Double-click on server > Media > Allow media overwrite.</p> <p>See “Media properties” on page 180.</p>

ALLOW_NON_RESERVED_PORTS option for NetBackup servers and clients

This option is not applicable to NetBackup 7.0.1 or later.

This option specifies that the NetBackup Client daemon (`bpcd`) can accept remote connections from unauthorized ports (port numbers 1024 or greater). If this option is not present, `bpcd` requires remote connections to come from privileged ports (port numbers 1024 or smaller). This option is useful when NetBackup clients and servers are on opposite sides of a firewall.

Table 3-63 ALLOW_NON_RESERVED_PORTS information

Usage	Description
Where to use	On NetBackup master servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>ALLOW_NON_RESERVED_PORTS = YES NO</pre> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers, Media Servers, Clients > Double-click on server > Universal Settings > Accept Connections on Non-reserved Ports.</p> <p>See “Universal Settings properties” on page 249.</p>

AUTHENTICATION_DOMAIN option for NetBackup servers and clients

This option defines a set of NetBackup product authentication principles.

Table 3-64 AUTHENTICATION_DOMAIN information

Usage	Description
Where to use	<p>A master server, media server, or client that uses NetBackup product authentication and authorization must have at least one <code>AUTHENTICATION_DOMAIN</code> entry listed for this option. More than one domain can be specified.</p> <p>If a media server or client does not define an authentication domain, it uses the authentication domains of its master server.</p>

Table 3-64 AUTHENTICATION_DOMAIN information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>AUTHENTICATION_DOMAIN = domain "comment" mechanism broker [port]</pre> <p>This option may appear more than once in the configuration file.</p> <p>The following variables are defined:</p> <ul style="list-style-type: none"> ■ <i>domain</i> is an Internet domain name or a Windows domain name. ■ <i>"comment"</i> is a quoted comment that describes the authentication domain. ■ <i>mechanism</i> is the authentication mechanism. <p>The keywords for various mechanisms are as follows:</p> <ul style="list-style-type: none"> ■ NIS: Network Information Service version 1 ■ NIS+: Network Information Service version 2 ■ PASSWD: Local UNIX password file on the specified broker ■ VXPDP: The NetBackup product authentication and authorization private database ■ WINDOWS: Windows Active Directory or primary domain controller ■ <i>broker</i> is the host name or IP address of the authentication broker. ■ <i>port</i> is the port number of the authentication broker. The default is the standard port number for authentication brokers.

Table 3-64 AUTHENTICATION_DOMAIN information (*continued*)

Usage	Description
Example	<p>AUTHENTICATION_DOMAIN = mycompany.com "Typical UNIX logins" NIS broker1.mycompany.com AUTHENTICATION_DOMAIN = OurEnterprise "Typical Windows logins" WINDOWS broker2.mycompany.com 5544 AUTHENTICATION_DOMAIN = mycompany.com "VxSS-Only Identities" VXPD broker1.mycompany.com AUTHENTICATION_DOMAIN = broker3.mycompany.com "Local UNIX Logins on host broker3" PASSWD broker3.mycompany.com</p> <p>In this example:</p> <ul style="list-style-type: none"> ■ mycompany.com is the Internet domain name and OurEnterprise is the Windows domain name. ■ The broker on host name broker1 handles both NIS and private authentication for NetBackup product authentication. ■ The broker on host name broker2 handles Windows authentication for NetBackup product authentication. ■ broker2 uses the non-standard port number 5544. ■ The broker on host name broker3 uses its local /etc/passwd file for NetBackup product authentication.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers, Media Servers, Clients > Double-click on server > Access Control > Authentication Domain.</p> <p>See "Authentication Domain tab of the Access Control properties" on page 68.</p>

AUTHORIZATION_SERVICE option for NetBackup servers

This option defines the NetBackup Product Authorization Service that the local NetBackup server uses.

Table 3-65 AUTHORIZATION_SERVICE information

Usage	Description
Where to use	<p>A master server that uses NetBackup product authorization must define an authorization service. If a media server does not define an authorization service, it uses the authorization service of the master server.</p>

Table 3-65 AUTHORIZATION_SERVICE information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>AUTHORIZATION_SERVICE = host [port]</pre> <p>The following variables are defined:</p> <ul style="list-style-type: none"> ■ <code>host</code> is the host name or IP address of the authorization service. ■ <code>port</code> is the port number of the authorization service. The default is the standard port number for the authorization service. <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers or Media Servers > Double-click on server > Access Control > Authorization Service.</p> <p>See “Authorization Service tab of the Access Control properties” on page 69.</p>

AUTO_ADD_ALL_ALIASES_FOR_CLIENT option for NetBackup servers

This option allows client aliases to be automatically added to the NetBackup database when `bpdbm` detects a new client in a backup policy.

Table 3-66 AUTO_ADD_ALL_ALIASES_FOR_CLIENT information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-66 AUTO_ADD_ALL_ALIASES_FOR_CLIENT information (*continued*)

Usage	Description
How to use	<p>By default, <code>AUTO_ADD_ALL_ALIASES_FOR_CLIENT</code> is not present in the configuration file. When <code>AUTO_ADD_ALL_ALIASES_FOR_CLIENT</code> is not present, the option is enabled. That is, <code>bpdbm</code> is allowed to add client aliases automatically.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>AUTO_ADD_ALL_ALIASES_FOR_CLIENT = YES NO</pre> <p>This entry should appear only once in the configuration file.</p>
Example	<p>The following entry prohibits <code>bpdbm</code> from adding a client alias automatically:</p> <pre>AUTO_ADD_ALL_ALIASES_FOR_CLIENT = NO</pre>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

BPBRM_VERBOSE option for NetBackup servers

The `BPBRM_VERBOSE` option is used for debugging purposes. It controls the amount of information that NetBackup includes in the `bpbrm` debug log.

Table 3-67 BPBRM_VERBOSE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>The default is that <code>BPBRM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPBRM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPBRM_VERBOSE = -1 0 1 2 3 4 5</pre> <p>This entry should appear only once in the configuration file.</p>

Table 3-67 BPBRM_VERBOSE information (*continued*)

Usage	Description
Example	<ul style="list-style-type: none"> To use the same value as the <code>VERBOSE</code> option, enter: <code>BPBRM_VERBOSE = 0</code> This option is the same as setting the BPBRM logging level in the Logging host properties to Same as Global. To log the minimum amount of information, enter: <code>BPBRM_VERBOSE = -1</code> This option is the same as setting the BPBRM logging level to 0 in the Logging host properties. To log additional information, enter a value of 1 through 5: <code>BPBRM_VERBOSE = 1</code> This option is the same as setting the BPBRM logging level to 1 in the Logging host properties. To log the maximum amount of information, enter: <code>BPBRM_VERBOSE = 5</code> This option is the same as setting the BPBRM logging level to 5 in the Logging host properties.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Logging > BPBRM logging level.</p> <p>See the NetBackup Logging Reference Guide for more information about the debug log.</p> <p>See “Logging properties” on page 161.</p>

BPDBJOBS_COLDEFS options for UNIX master servers

Use `BPDBJOBS_COLDEFS` entries to customize the output of the `bpdbjobs` process. Add a `BPDBJOBS_COLDEFS` option for every column you want to include in the output.

Add `BPDBJOBS_COLDEFS` entries to the `bp.conf` file to customize the output of the `bpdbjobs` process.

Table 3-68 BPDBJOBS_COLDEFS information

Usage	Description
Where to use	On a UNIX NetBackup master server.

Table 3-68 BPDBJOBS_COLDEFS information (*continued*)

Usage	Description
How to use	<p>Add BPDBJOBS_COLDEFS to the <code>/usr/opensv/netbackup/bp.conf</code> file.</p> <p>Add an entry for every column to include in the output by using the following format:</p> <pre>BPDBJOBS_COLDEFS = COLDEFS_ENTRY [minimum_size [true false]]</pre> <p>The following variables are defined:</p> <ul style="list-style-type: none"> ■ COLDEFS_ENTRY is the name of the column to include in the output. ■ minimum_size is the minimum column width. If not specified, the default is a width of 5. ■ true indicates that the column should expand as needed. If not specified, true is the default. ■ false indicates that the column should not expand beyond the <i>minimum_size</i>.
Example	<p>The order of the entries determines the order in which the column headings appear.</p> <pre>BPDBJOBS_COLDEFS = JOBID 5 true BPDBJOBS_COLDEFS = TYPE 4 true BPDBJOBS_COLDEFS = STATE 5 true BPDBJOBS_COLDEFS = STATUS 6 true BPDBJOBS_COLDEFS = POLICY 6 true BPDBJOBS_COLDEFS = SCHEDULE 8 true BPDBJOBS_COLDEFS = CLIENT 6 true BPDBJOBS_COLDEFS = DSTMEDIA_SERVER 12 true BPDBJOBS_COLDEFS = ACTPID 10 true</pre> <p>The appearance of BPDBJOBS_COLDEFS entries in the <code>bp.conf</code> file has the following ramifications:</p> <ul style="list-style-type: none"> ■ The addition of any BPDBJOBS_COLDEFS option overrides all default columns. ■ All users on the local system see only those columns that are specified in the <code>bp.conf</code> file.
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

[Table 3-69](#) shows possible COLDEFS entries and the column which is created by each.

Table 3-69 COLDEFS entries

COLDEFS entry	Column Name
ACTIVEELAPSED	Active Elapsed (elapsed active time)
ACTPID	Active PID (PID of job)
ATTEMPT	Attempt
BACKUPTYPE	Backup Type
CLIENT	Client
COMPLETION	Completion (percent complete)
COMPRESSION	Compression (yes or no)
DEDUPRATIO	Dedup (shows deduplication rate in <code>bpd jobs</code> command output)
DSTMEDIA_SERVER	Dest Media Svr (writing media server)
DSTMEDIAID	Dest Media ID (writing media ID)
DSTSTORAGE_UNIT	Dest StUnit (writing storage unit)
ELAPSED	Elapsed (elapsed time)
ENDED	Ended
ESTFILE	Est File (estimated number of files)
ESTKB	Est KB (estimated number of kilobytes)
FILES	Files
GROUP	Group
JOBID	JobID
KBPERSEC	KB Per Sec
KILOBYTES	Kilobytes
LASTBACKUP	Last Backup (date and time)
MAINPID	Main PID (PID that spawns job, if applicable)
NUMTAPESEJECT	Media to Eject (number of tapes to eject; Vault only)

Table 3-69 COLDEFS entries (*continued*)

COLDEFS entry	Column Name
OPERATION	Operation (current operation)
OWNER	Owner
PATHNAME	Pathname
PARENTJOBID	Parent JobID
POLICY	Policy
POLICYTYPE	Policy Type
PRIORITY	Priority
PROFILE	Profile (Vault only)
RETENTION	Retention (retention period)
RESUMABLE	Resumable
ROBOT	Robot (Vault only)
RQSTPID	Request PID (PID requesting job, if applicable)
SCHEDULE	Schedule
SCHEDULETYPE	Schedule Type
SESSIONID	Session ID (Vault only)
SRCMEDIA_SERVER	Src Media Svr
SRCMEDIAID	Src Media ID
SRCSTORAGE_UNIT	Src StUnit
STARTED	Started
STATE	State
STATUS	Status
STREAMNUMBER	Stream Number
SUSPENDABLE	Suspendable
TYPE	Type (job type)

Table 3-69 COLDEFS entries (*continued*)

COLDEFS entry	Column Name
VAULT	Vault (Vault only)

BPDBM_VERBOSE option for NetBackup servers

The `BPDBM_VERBOSE` option is used for debugging purposes. It controls the amount of information NetBackup includes in the `bpdbm` debug log.

Table 3-70 BPDBM_VERBOSE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>The default is that <code>BPDBM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPDBM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPDBM_VERBOSE = -1 0 1 2 3 4 5</pre> <p>This entry should appear only once in the configuration file.</p>

Table 3-70 BPDBM_VERBOSE information (*continued*)

Usage	Description
Example	<ul style="list-style-type: none"> To use the same value as the <code>VERBOSE</code> option for, enter: <code>BPDBM_VERBOSE = 0</code> This option is the same as setting the BPDBM logging level to Same as Global in the Logging host properties. To log the minimum amount of information, enter: <code>BPDBM_VERBOSE = -1</code> This option is the same as setting the BPDBM logging level to 0 in the Logging host properties. To log additional information, enter a value of 1 through 5: <code>BPDBM_VERBOSE = 1</code> This option is the same as setting the BPDBM logging level to 1 in the Logging host properties. To log the maximum amount of information, enter: <code>BPDBM_VERBOSE = 5</code> This option is the same as setting the BPDBM logging level to 5 in the Logging host properties. <p>The following examples show two entries which enable logging, while they minimize the growth rate of the <code>bpdbm</code> debug file:</p> <pre>VERBOSE = 5 BPDBM_VERBOSE = -1</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Logging > Global logging level > BPDBM logging level.</p> <p>See the NetBackup Logging Reference Guide for more information about logging.</p> <p>See “Logging properties” on page 161.</p>

BPRD_VERBOSE option for NetBackup servers

Used for debugging purposes, the `BPRD_VERBOSE` option controls the amount of information that NetBackup includes in the `bprd` debug logs.

Table 3-71 BPRD_VERBOSE information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-71 BPRD_VERBOSE information (*continued*)

Usage	Description
How to use	<p>The default is that the value is the same as the <code>VERBOSE</code> option (Global logging level). The <code>BPRD_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPRD_VERBOSE = -1 0 1 2 3 4 5</pre> <p>This entry should appear only once in the configuration file.</p>
Example	<ul style="list-style-type: none"> ■ To use the same value as the <code>VERBOSE</code> option, enter: <pre>BPRD_VERBOSE = 0</pre> <p>This option is the same as setting the BPRD logging level in the Logging host properties to Same as Global.</p> ■ To log the minimum amount of information, enter: <pre>BPRD_VERBOSE = -1</pre> <p>This option is the same as setting the BPRD logging level to 0 in the Logging host properties.</p> ■ To log additional information, enter a value of 1 through 5: <pre>BPRD_VERBOSE = 1</pre> <p>This option is the same as setting the BPRD logging level to 1 in the Logging host properties.</p> ■ To log the maximum amount of information, enter: <pre>BPRD_VERBOSE = 5</pre> <p>This option is the same as setting the BPRD logging level to 5 in the Logging host properties.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Logging > BPRD logging level.</p> <p>See the NetBackup Logging Reference Guide for more information about the debug log.</p> <p>See “Logging properties” on page 161.</p>

BPTM_VERBOSE option for NetBackup servers

The `BPTM_VERBOSE` option is used for debugging purposes. It controls the amount of information that NetBackup includes in the `bptm` debug logs.

Table 3-72 BPTM_VERBOSE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>The default is that <code>BPTM_VERBOSE</code> is the same value as the <code>VERBOSE</code> option (Global logging level). The <code>BPTM_VERBOSE</code> option overrides the <code>VERBOSE</code> option in the configuration file.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPTM_VERBOSE = -1 0 1 2 3 4 5</pre> <p>This entry should appear only once in the configuration file.</p>
Example	<ul style="list-style-type: none"> ■ To use the same value as the <code>VERBOSE</code> option, enter: <code>BPTM_VERBOSE = 0</code> This option is the same as setting the BPTM logging level in the Logging host properties to Same as Global. ■ To log the minimum amount of information, enter: <code>BPTM_VERBOSE = -1</code> This option is the same as setting the BPTM logging level to 0 in the Logging host properties. ■ To log additional information, enter a value of 1 through 5: <code>BPTM_VERBOSE = 1</code> This option is the same as setting the BPTM logging level to 1 in the Logging host properties. ■ To log the maximum amount of information, enter: <code>BPTM_VERBOSE = 5</code> This option is the same as setting the BPTM logging level to 5 in the Logging host properties.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Logging > BPTM logging level.</p> <p>See the NetBackup Logging Reference Guide for more information about the debug log.</p> <p>See “Logging properties” on page 161.</p>

BPEND_TIMEOUT option for NetBackup servers

The `BPEND_TIMEOUT` option specifies the number of seconds to wait for the `bpend_notify` script on a client to complete.

Table 3-73 BPEND_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPEND_TIMEOUT = seconds</pre> <p>The default timeout is 300 seconds (five minutes).</p> <p>Note: If this option is changed, verify that the <code>CLIENT_READ_TIMEOUT</code> option is set to the same value or higher.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Timeouts > Backup end notify timeout.</p> <p>See "Timeouts properties" on page 247.</p>

BPSTART_TIMEOUT option for NetBackup servers

The `BPSTART_TIMEOUT` option specifies the number of seconds to wait for the `bpstart_notify` script on a client to complete.

Table 3-74 BPSTART_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-74 BPSTART_TIMEOUT information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPSTART_TIMEOUT = seconds</pre> <p>The default timeout is 300 seconds (five minutes).</p> <p>Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Timeouts > Backup start notify timeout.</p> <p>See “Timeouts properties” on page 247.</p>

CHECK_RESTORE_CLIENT option for NetBackup servers

The `CHECK_RESTORE_CLIENT` option specifies that the client to be restored to is checked before the restore starts. An unresponsive client can slow restores for other clients that have data on the same tapes.

Table 3-75 CHECK_RESTORE_CLIENT information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

CLIENT_CONNECT_TIMEOUT option for NetBackup servers

This option specifies the number of seconds that the server waits when it connects to a client. If the server needs to wait longer than the time specified, it times out.

Table 3-76 CLIENT_CONNECT_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CLIENT_CONNECT_TIMEOUT = seconds</pre> <p>The default timeout is 300 seconds (five minutes).</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Timeouts > Client connect timeout.</p> <p>See "Timeouts properties" on page 247.</p>

CLIENT_PORT_WINDOW option for NetBackup servers and clients

This option specifies the range of non-reserved ports on this computer that are used as source ports when connecting to NetBackup on other computers. This setting applies to daemon or service socket connections to the server and to the client hosts as well as call-back from `bpcd`.

Table 3-77 CLIENT_PORT_WINDOW information

Usage	Description
Where to use	On NetBackup servers and clients.

Table 3-77 CLIENT_PORT_WINDOW information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CLIENT_PORT_WINDOW = start_port_range end_port_range</pre> <p>If 0 is specified for the first number (default), the operating system determines the non-reserved port to use.</p> <p>This option should appear only once in the configuration file.</p>
Example	<p>The following example permits ports from 4800 through 5000:</p> <pre>CLIENT_PORT_WINDOW = 4800 5000</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Port Ranges > Client port window.</p> <p>See “Port Ranges properties” on page 192.</p>

CLIENT_READ_TIMEOUT option for NetBackup servers

The `CLIENT_READ_TIMEOUT` option specifies the number of seconds to use for the client-read timeout.

If the master server does not get a response from a client within the `CLIENT_READ_TIMEOUT` period, the backup or restore operation, for example, fails.

Table 3-78 CLIENT_READ_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-78 CLIENT_READ_TIMEOUT information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CLIENT_READ_TIMEOUT = seconds</pre> <p>By default, <code>CLIENT_READ_TIMEOUT</code> is not present on the server or the database agent and the client-read timeout is 300 seconds (five minutes). This time is a reasonable default. Change only in the event of problems.</p> <p><code>CLIENT_READ_TIMEOUT</code> on a database agent is a special case because these types of clients can initially require more time to get ready than other clients. Database backup utilities frequently start several backup jobs at the same time, which can slow the CPU.</p> <p>The sequence on a database agent is as follows:</p> <ul style="list-style-type: none"> ■ NetBackup on the database agent reads the client's <code>CLIENT_READ_TIMEOUT</code> to find the value to use initially. If the option is not set, the standard default of five minutes is used. ■ When the database agent API receives the server's value, it uses it as the <code>CLIENT_READ_TIMEOUT</code>. <p>Note: Symantec suggests setting <code>CLIENT_READ_TIMEOUT</code> on the database agent to a value greater than 5 minutes. A setting of 15 minutes is adequate for many installations.</p> <p>Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can timeout while the <code>bpstart_notify</code> script is running.</p>
Example	<p>The following example configures a client read timeout of 15 minutes.</p> <pre>CLIENT_READ_TIMEOUT = 900</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Timeouts > Client read timeout.</p> <p>See "Timeouts properties" on page 247.</p>

CLIENT_RESERVED_PORT_WINDOW option for NetBackup servers and clients

The `CLIENT_RESERVED_PORT_WINDOW` option specifies the range of reserved ports on this computer that are used for connecting to NetBackup on other computers. This setting applies when the computer connects to a client (`bpcd`) configured to accept only reserved ports.

Beginning with NetBackup 7.0.1, connections to other daemon ports use only the `CLIENT_PORT_WINDOW`, **not the** `CLIENT_RESERVED_PORT_WINDOW`.

Table 3-79 `CLIENT_RESERVED_PORT_WINDOW` information

Usage	Description
Where to use	On NetBackup master servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CLIENT_RESERVED_PORT_WINDOW = start_port_range end_port_range</pre> <p>The default is from 512 through 1023. Note that if 0 is specified for the first number, a non-reserved port is used instead. The operating system chooses the nonreserved port.</p> <p>This option should appear only once in the configuration file.</p>
Example	<p>The following example permits ports from 900 through 1023:</p> <pre>CLIENT_RESERVED_PORT_WINDOW = 900 1023</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Port Ranges > Client reserved port window.</p> <p>See “Port Ranges properties” on page 192.</p>

CONNECT_OPTIONS option for NetBackup servers and clients

The `CONNECT_OPTIONS` option specifies the following options that are designed to enhance firewall efficiency with NetBackup:

- Whether processes on *host* are reached:
 - By using the Veritas Network daemon (`vneta`) port only.
 - By using the daemon ports only.

- By first using the `vnetd` port and then the daemon ports if needed.
- Whether subsequent call-back connections with *host* use the traditional call-back method or use `vnetd`.
- Whether connections to *host* use reserved or a non-reserved port number.

Table 3-80 CONNECT_OPTIONS information

Usage	Description
Where to use	On NetBackup master servers, media servers, or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CONNECT_OPTIONS = host [0 1 2] [0 1 2] [0 1 2 3]</pre> <p>The following variables are defined:</p> <p><i>Host</i> is a NetBackup system to which this computer connects. You may have multiple <code>CONNECT_OPTIONS</code> entries in the configuration. If a host is not specified in any <code>CONNECT_OPTIONS</code> entries, the values from the <code>DEFAULT_CONNECT_OPTIONS</code> option are used.</p> <p>See “DEFAULT_CONNECT_OPTIONS option for NetBackup servers and clients” on page 285.</p> <p>The first setting indicates the type of port to use as the source port for connections to service daemons on <i>host</i>:</p> <p>0 = Connections from this computer should be from a reserved source port number. It is selected from the <code>CLIENT_RESERVED_PORT_WINDOW</code> range.</p> <p>1 = Connections from this computer should be from a non-reserved source port number that is selected from the <code>CLIENT_PORT_WINDOW</code> range. (The default is 1.) If this option is selected for use with the daemon port and <i>host</i> is running a version of NetBackup older than 7.1, enable Allow non reserved ports for the selected <i>host</i>.</p> <p>In the NetBackup Administration Console, see the Universal Settings dialog box under Host Properties > Media Servers.</p> <p>2 = Use the method that the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option defines.</p> <p>See “Universal Settings properties” on page 249.</p>

Table 3-80 CONNECT_OPTIONS information (*continued*)

Usage	Description
	<p>The second setting indicates the call-back method to use with <i>host</i>. (This applies if configured to use the daemon port or if <i>bpcd</i> cannot be reached using ports 1556 or 13724.)</p> <p>0 = Use the traditional call-back method. <i>Host</i> connects back to a random port number that this computer has selected from the <code>SERVER_RESERVED_PORT_WINDOW</code> range, or <code>SERVER_PORT_WINDOW</code> range as determined by the first setting.</p> <p>1 = Use the <i>vnetd</i> no call-back method. Connect to <i>vnetd</i> instead of a random port. If this computer is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <i>vnetd</i>. The connection is outbound to <i>host</i> if the third setting is 0 or 1. The connection is inbound from <i>host</i> if the third setting is 2.</p> <p>2 = Use the method that the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option defines (default).</p>
	<p>The third setting affects which destination port number to use for connections to daemon services on <i>host</i>.</p> <p>0 = Use <i>vnetd</i> if possible, to connect to a daemon on this computer. If it is not possible, connect by using the traditional port number of the daemon. If this host is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <i>vnetd</i>. If <i>bpcd</i> is reachable through port 1556 or <i>vnetd</i>, then the first two settings are ignored and all connections to <i>host</i> are outbound and use non-reserved ports.</p> <p>1 = Use <i>vnetd</i> to connect to a daemon on <i>host</i>. If this computer is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <i>vnetd</i>. The traditional port number of the daemon is not used if the connection cannot be established using port 1556 or <i>vnetd</i>. If <i>bpcd</i> is reachable through port 1556 or <i>vnetd</i>, then the first two settings are ignored and all connections to <i>host</i> are outbound and use non-reserved ports.</p> <p>2 = Use only the traditional port number of the daemon to connect to a daemon on <i>host</i>. If the daemon is <i>bpcd</i>, the second setting is then used to determine the call-back method.</p> <p>3 = Use the method that the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option defines (default).</p>

Table 3-80 CONNECT_OPTIONS information (*continued*)

Usage	Description
Example 1	<p>The configuration file can contain <code>CONNECT_OPTIONS</code> settings for multiple hosts.</p> <pre>CONNECT_OPTIONS = shark 0 0 0</pre> <p>In this example, connections to daemons on <code>shark</code> attempt to use port 1556 if this computer is running NetBackup 7.0.1 or later. If this computer is running a version of NetBackup that is earlier than 7.0.1, or if the previous attempt was unsuccessful, then the connections try <code>vnetd</code>. If that attempt is unsuccessful, the daemon port is used. If the connections are successful using 1556 or <code>vnetd</code>, then the first two settings are ignored.</p> <pre>\$ bptestbpcd -host shark 0 0 1 10.82.105.11:40402 -> 10.82.104.249:1556 10.82.105.11:40404 -> 10.82.104.249:1556</pre>
Example 2	<pre>CONNECT_OPTIONS = dolphin 1 0 1</pre> <p>In this example, connections to daemons on <code>dolphin</code> attempt to use port 1556 if this computer is running NetBackup 7.0.1 or later. If this computer is running a version of NetBackup that is earlier than 7.0.1, or if the previous attempt was unsuccessful, then the connections try <code>vnetd</code>. The first two settings are ignored.</p> <pre>\$ bptestbpcd -host dolphin 1 0 1 10.82.105.11:40414 -> 10.82.104.249:13724 10.82.105.11:40416 -> 10.82.104.249:13724</pre>
Example 3	<pre>CONNECT_OPTIONS = perch 0 1 2</pre> <p>In this example:</p> <ul style="list-style-type: none"> ■ Connections to daemons on <code>perch</code> must use the daemon port numbers. ■ Call-back connections are inbound to <code>vnetd</code> on this computer. ■ The source ports for all connections are bound from the reserved port number range.

Table 3-80 CONNECT_OPTIONS information (*continued*)

Usage	Description
Example 4	<p><code>CONNECT_OPTIONS = trout 1 1 2</code></p> <p>In this example:</p> <ul style="list-style-type: none">■ Connections to daemons on <code>trout</code> must use the daemon port numbers.■ Call-back connections are inbound to <code>vnetd</code> on this computer.■ The source ports for all connections are bound from the non-reserved port number range.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Click Master Servers, Media Servers, Clients > Double-click on server > Firewall.</p> <p>See “Firewall properties” on page 142.</p>

DEFAULT_CONNECT_OPTIONS option for NetBackup servers and clients

The `DEFAULT_CONNECT_OPTIONS` option specifies the default values for the `CONNECT_OPTIONS` configuration option. If a host name is not specified in any `CONNECT_OPTIONS` option, the value from the `DEFAULT_CONNECT_OPTIONS` option is used.

See [“CONNECT_OPTIONS option for NetBackup servers and clients”](#) on page 281.

Table 3-81 DEFAULT_CONNECT_OPTIONS information

Usage	Description
Where to use	On NetBackup master servers, media servers, or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DEFAULT_CONNECT_OPTIONS = [0 1][0 1][0 1 2]</pre> <p>The default value is <code>1 1 0</code>.</p> <p>This option should appear only once in the configuration file.</p>

Table 3-81 DEFAULT_CONNECT_OPTIONS information (*continued*)

Usage	Description
	<p>The first setting indicates the type of port to use as the source port when connecting to service daemons on a destination host. It also indicates the type of server port if using the traditional call-back method.</p> <p>0 = Connections on this computer should use a reserved port number. They are selected from the <code>CLIENT_RESERVED_PORT_WINDOW</code> range for source ports and from the <code>SERVER_RESERVED_PORT_WINDOW</code> range if using the traditional call-back method.</p> <p>1 = Connections on this computer should use a non-reserved port number. (1 is the default.) Connections are selected from the <code>CLIENT_PORT_WINDOW</code> range for source ports and from the <code>SERVER_PORT_WINDOW</code> range if using the traditional call-back method. If this option is selected and the destination host is earlier than NetBackup 7.1, enable Allow non reserved ports for the selected <i>host</i>. See the Universal Settings dialog box under Host Properties > Media Servers.</p> <p>See "Universal Settings properties" on page 249.</p>
	<p>The second setting indicates the call-back method to use. (This setting applies if configured to use the daemon port or if <code>bpcd</code> cannot be reached using ports 1556 or 13724.)</p> <p>0 = Use the traditional call-back method. The destination host connects back to a random port number that this computer has selected from the <code>SERVER_RESERVED_PORT_WINDOW</code> range, or the <code>SERVER_PORT_WINDOW</code> range as determined by the first setting.</p> <p>1 = Use the <code>vnetd</code> no call-back method. Connect to <code>vnetd</code> instead of a random port. If this computer is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <code>vnetd</code>. The connection is outbound to the destination host if the third setting is 0 or 1. The connection is inbound if the third setting is 2.</p>

Table 3-81 DEFAULT_CONNECT_OPTIONS information (*continued*)

Usage	Description
	<p>The third setting affects which destination port number to use when attempting connections to daemon services on the destination host.</p> <p>0 = Use <code>vnetd</code> to connect to a daemon on the destination host if possible, otherwise connect by using the traditional port number of the daemon. If this computer is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <code>vnetd</code>. If <code>bpcd</code> is reachable through port 1556 or <code>vnetd</code>, then the first two settings are ignored and all connections are outbound and use non-reserved ports.</p> <p>1 = Connect to a daemon on the destination host by using <code>vnetd</code>. If this computer is NetBackup 7.0.1 or later, attempt to connect to port 1556 before attempting to connect to <code>vnetd</code>. The traditional port number of the daemon will not be used if the connection cannot be established using port 1556 or <code>vnetd</code>. If <code>bpcd</code> is reachable through port 1556 or <code>vnetd</code>, then the first two settings are ignored and all connections are outbound and use non-reserved ports.</p> <p>2 = Connect to a daemon on the destination host by using the traditional port number of the daemon only. The first setting is used to determine the port range and the second setting is used to determine the call-back method.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers, Media Servers, Clients > Double-click on server > Firewall.</p> <p>See "Firewall properties" on page 142.</p>

DISABLE_JOB_LOGGING option for NetBackup servers

This option disables the logging of the job information that the NetBackup Activity Monitor requires.

Table 3-82 DISABLE_JOB_LOGGING information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-82 DISABLE_JOB_LOGGING information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DISABLE_JOB_LOGGING</pre> <p>The default is that this option is not present in the configuration file and that job logging occurs.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Media > Enable job logging.</p> <p>See “Media properties” on page 180.</p>

DISABLE_STANDALONE_DRIVE_EXTENSIONS option for NetBackup servers

This option disables the nonrobotic drive operations. During a backup, NetBackup automatically attempts to use standalone volumes in nonrobotic drives.

Table 3-83 DISABLE_STANDALONE_DRIVE_EXTENSIONS information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbemmcmd</code> command to change the option. For example:</p> <pre>nbemmcmd -changesetting -DISABLE_STANDALONE_DRIVE_EXTENSIONS no</pre> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>This option should appear only once in the configuration file.</p>
Example	<p>The following command enables nonrobotic drive operations.</p> <pre>nbemmcmd -changesetting -DISABLE_STANDALONE_DRIVE_EXTENSIONS no</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Media > Enable standalone drive extension. The default is that this option is enabled.</p> <p>See “Media properties” on page 180.</p>

DISALLOW_BACKUPS_SPANNING_MEDIA option for NetBackup servers

This option prevents backups from spanning media.

Table 3-84 DISALLOW_BACKUPS_SPANNING_MEDIA information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DISALLOW_BACKUPS_SPANNING_MEDIA</pre> <p>The default is that the entry is not present in the configuration file and backups are allowed to span media.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Media > Allow backups to span tape media.</p> <p>See "Media properties" on page 180.</p>

DISALLOW_CLIENT_LIST_RESTORE option for NetBackup servers

This option denies the list and restore requests for all clients. When this option is present, clients cannot list or restore any files that they have backed up through this master server.

Table 3-85 DISALLOW_CLIENT_LIST_RESTORE information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-85 DISALLOW_CLIENT_LIST_RESTORE information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DISALLOW_CLIENT_LIST_RESTORE</pre> <p>The default is that the entry is not present in the configuration file and clients can list and restore their files.</p> <p>Note: Override the <code>DISALLOW_CLIENT_LIST_RESTORE</code> option for individual clients by changing their <code>list_restore</code> setting.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Client Attributes > Allow Client Restore.</p> <p>See "Client Attributes properties" on page 84.</p>

DISALLOW_CLIENT_RESTORE option for NetBackup servers

This option denies the restore requests for all clients. When this option is present, clients cannot restore the files that they have backed up through this master server.

Table 3-86 DISALLOW_CLIENT_RESTORE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DISALLOW_CLIENT_RESTORE</pre> <p>The default is that the entry is not present in the configuration file and clients can restore their files.</p> <p>Note: To override the <code>DISALLOW_CLIENT_RESTORE</code> option for individual clients, change their <code>list_restore</code> setting.</p>

Table 3-86 DISALLOW_CLIENT_RESTORE information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Master Servers > Double-click on server > Client Attributes > Allow Client Browse . See “Client Attributes properties” on page 84.

EMMSERVER option for NetBackup servers

In releases before NetBackup 7.6, `EMMSERVER` indicated the master or the media server that acted as the Enterprise Media Manager server for one or more master servers.

Support for remote-EMM and shared-EMM server configurations was withdrawn in 7.6. In a remote- or shared-EMM server configuration, the NetBackup relational database (NBDB), the Enterprise Media Manager (EMM), and the Resource Broker (RB) have been moved to a server that is not the master server. Customers who currently use this option should contact Symantec Support. Support engages Symantec Engineering to review the options to disengage this configuration.

Table 3-87 EMMSERVER information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>EMMSERVER = server_name</pre>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

See [“Moving the NetBackup database from one host to another”](#) on page 954.

FAILOVER_RESTORE_MEDIA_SERVERS option for NetBackup servers

This option specifies that an automatic failover media server be used if a server is temporarily inaccessible for a restore. This failover does not require administrator intervention.

Table 3-88 FAILOVER_RESTORE_MEDIA_SERVERS information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>FAILOVER_RESTORE_MEDIA_SERVERS = failed_host host1 host2 ... hostN</pre> <ul style="list-style-type: none"> ■ <i>failed_host</i> is the server that is not operational. ■ <i>host1 ... hostN</i> are the servers that provide failover capabilities. <p>The default is that NetBackup does not perform automatic failover.</p> <p>When automatic failover is necessary for a server, NetBackup searches from left to right through the associated <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> list. It stops when it finds one that is eligible to perform the restore.</p> <p>Note: The configuration file can contain multiple <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> entries and each entry can list multiple servers. However, a NetBackup server can be a <i>failed_host</i> in only one option.</p> <p>After a <code>FAILOVER_RESTORE_MEDIA_SERVERS</code> option is added, stop and restart the NetBackup Request daemon on the master server where you plan to change the configuration.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Restore Failover.</p> <p>See "Restore Failover properties" on page 216.</p>

FORCE_RESTORE_MEDIA_SERVER option for NetBackup servers

This option forces the restore to go to a specific server, regardless of where the files were backed up.

Table 3-89 FORCE_RESTORE_MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>FORCE_RESTORE_MEDIA_SERVER = <i>from_host to_host</i></pre> <p>Where <i>from_host</i> is the server that performed the original backup and <i>to_host</i> is the server to use for the restore.</p> <p>Stop and restart the NetBackup Request daemon on the master server after adding the <code>FORCE_RESTORE_MEDIA_SERVER</code> option. Physically move the media to <i>to_host</i> before attempting a restore. Update the Media Manager volume database to reflect the move.</p> <p>This setting applies to all storage units on the original server. Restores for any storage unit on <i>from_host</i> go to <i>to_host</i>. To revert to the original configuration for future restores, delete the option.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > General Server > Media Host Override.</p> <p>See “General Server properties” on page 146.</p>

GENERATE_ENGLISH_LOGS option for NetBackup servers and clients

This option enables the generation of an English error log, and English trace logs for the `bptarchive`, `bpbbackup`, `bpduplicate`, `bpimport`, and `bprestore` commands. This option is useful to support personnel to assist in distributed environments where different locales result in logs that contain various languages.

An English text error log (indicated by the suffix `_en`) is created in the following directory:

- On Windows: `Install_path\Veritas\NetBackup\db\error`
- On UNIX: `/usr/opensv/netbackup/db/error`

Table 3-90 GENERATE_ENGLISH_LOGS information

Usage	Description
Where to use	On NetBackup master servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>GENERATE_ENGLISH_LOGS</pre> <p>This entry should appear only once in the configuration file.</p> <p>Setting the <code>GENERATE_ENGLISH_LOGS</code> option also forces the <code>-en</code> argument on the execution of the following commands when the progress log is specified (<code>-L</code>): <code>bparchive</code>, <code>bpbackup</code>, <code>bpduplicate</code>, <code>bpimport</code>, and <code>bprestore</code>.</p> <p>The suffix <code>_en</code> indicates the English text progress log.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

INCOMPLETE_JOB_CLEAN_INTERVAL option for NetBackup servers and clients

This option indicates the number of days a failed restore job can remain in the incomplete state before it is moved to the done state.

Table 3-91 INCOMPLETE_JOB_CLEAN_INTERVAL information

Usage	Description
Where to use	On NetBackup master servers or clients.

Table 3-91 INCOMPLETE_JOB_CLEAN_INTERVAL information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>INCOMPLETE_JOB_CLEAN_INTERVAL = number_of_days</pre> <p>The default is 7 days.</p> <p>Where <i>x</i> is a value between 0 and 365. A value of 0 indicates that failed, incomplete jobs are never automatically moved to the done state.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Clients > Double-click on server > Clean-up > Move restore job from incomplete state to done state.</p> <p>See “Clean-up properties” on page 80.</p>

INITIAL_BROWSE_SEARCH_LIMIT option for NetBackup servers and clients

This option specifies the number of days back that NetBackup searches for files to restore.

Table 3-92 INITIAL_BROWSE_SEARCH_LIMIT information

Usage	Description
Where to use	On NetBackup master servers or clients.

Table 3-92 INITIAL_BROWSE_SEARCH_LIMIT information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>INITIAL_BROWSE_SEARCH_LIMIT = number_of_days</pre> <p>The default is that 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 the browse starts with the earliest of the set of last full backups.</p> <p>This entry should appear only once in the configuration file.</p> <p>When this option is specified on a UNIX client, it applies only to that client. The option can reduce the size of the Search window from what is specified on the server (the client setting cannot make the window larger).</p>
Example	<p>The following example limits the browse range to the seven days before the current date.</p> <pre>INITIAL_BROWSE_SEARCH_LIMIT = 7</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Universal Settings > Browse timeframe for restores.</p> <p>See “Universal Settings properties” on page 249.</p>

IP_ADDRESS_FAMILY option for NetBackup servers

The `IP_ADDRESS_FAMILY` option indicates whether NetBackup on this host supports only IPv4 or both IPv4 and IPv6.

If any of the `SERVER` entries do not support IPv4, NetBackup uses the setting that indicates both IPv4 and IPv6.

Table 3-93 IP_ADDRESS_FAMILY information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 3-93 IP_ADDRESS_FAMILY information (*continued*)

Usage	Description
How to use	<p>On the hosts that use both IPv4 and IPv6, use this option to indicate which address family to use.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>IP_ADDRESS_FAMILY = AF_INET AF_UNSPEC</pre> <p><code>AF_INET</code> indicates that the host supports only IPv4.</p> <p><code>AF_UNSPEC</code> indicates that the host supports both IPv4 and IPv6.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Network Settings.</p> <p>See “Network Settings properties” on page 190.</p>

JOB_PRIORITY option for NetBackup servers and clients

Use this option to set the priority for a job type.

Table 3-94 JOB_PRIORITY information

Usage	Description
Where to use	On NetBackup master servers or clients.

Table 3-94 JOB_PRIORITY information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>JOB_PRIORITY = P1 P2 P3 P4 P5 P6 P7 P8 P9 P10 P11 P12 P13 P14 P15 P16 P17 P18 P19 P20 P21 P22 P23 P24</pre> <p>Variables <i>P1</i>, <i>P2</i>, and so on indicate the priority for each backup type.</p> <p>Table 3-95 shows the default priority values.</p> <p>The actual default values for the option appear as follows:</p> <pre>JOB_PRIORITY = 0 0 90000 90000 90000 90000 85000 85000 80000 80000 80000 80000 75000 75000 70000 70000 50000 50000 0 0 0 0 0 0</pre> <p>This entry should appear only once in the configuration file.</p>
Example	<p>To give backup jobs a priority of 50000 and duplication jobs a priority of 30000, change the <code>JOB_PRIORITY</code> entry as follows:</p> <pre>JOB_PRIORITY = 50000 0 90000 90000 90000 90000 85000 85000 80000 80000 80000 80000 75000 75000 70000 70000 30000 50000 0 0 0 0 0 0</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Clients > Double-click on server > Default Job Priorities > Job Priority.</p> <p>See “Default Job Priorities properties” on page 113.</p>

[Table 3-95](#) lists the order of the job types and the various job type defaults.

Table 3-95 Default job type priorities

Field	Represents this action	Default
P1	Performing a backup	0
P2	Performing a database backup (a catalog backup)	0
P3	Performing a restore	90000
P4	Recovering a catalog	90000

Table 3-95 Default job type priorities (*continued*)

Field	Represents this action	Default
P5	Performing a staging operation	90000
P6	Performing the duplication jobs that Vault starts	90000
P7	Cleaning up images	85000
P8	Importing images	85000
P9	Requesting tapes	80000
P10	Cleaning a tape	80000
P11	Tape formatting	80000
P12	Performing device diagnostics	80000
P13	Verifying an image	75000
P14	Running a media contents report	75000
P15	Labeling tape media	70000
P16	Erasing media	70000
P17	Running a duplication job	50000
P18	Performing an inventory	50000
P19	This field is not currently in use	0
P20	This field is not currently in use	0
P21	This field is not currently in use	0
P22	This field is not currently in use	0
P23	This field is not currently in use	0
P24	This field is not currently in use	0

LIMIT_BANDWIDTH option for NetBackup servers

This option specifies a limit for the network bandwidth that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. This option limits only backups. Restores are unaffected. The default is that the bandwidth is not limited.

Note: `LIMIT_BANDWIDTH` applies only to IPv4 networks. Use the `THROTTLE_BANDWIDTH` option to limit IPv6 networks.

See [“Throttle Bandwidth properties”](#) on page 245.

MAX_LOGFILE_SIZE option for NetBackup servers and clients for legacy logging

The `MAX_LOGFILE_SIZE` option specifies the maximum size that you want to set for a log file. When the log file size in NetBackup matches the `MAX_LOGFILE_SIZE` setting, the next logs are stored in a new log file.

See [“About robust logging”](#) on page 260.

Table 3-96 `MAX_LOGFILE_SIZE` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> (or <code>bpgetconfig</code>) and the <code>nbsetconfig</code> (or <code>bpsetconfig</code>) commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Initially, the <code>bp.conf</code> file does not contain <code>MAX_LOGFILE_SIZE</code> entry. In this case, the parameter is set to its default value, that is 50 MB.</p> <p>To set it to a different value, update the <code>bp.conf</code> file using the <code>nbsetconfig</code> command.</p> <p>Use the following format to set <code>MAX_LOGFILE_SIZE</code> to a new value in the <code>bp.conf</code> file:</p> <pre>MAX_LOGFILE_SIZE = X</pre> <p>'X' indicates maximum size of a NetBackup log file in MB.</p> <p>Note: <code>MAX_LOGFILE_SIZE</code> should be set to an integer number, which should be greater than '0'. If you have set <code>MAX_LOGFILE_SIZE</code> to an invalid value such as 0 or -100, it is automatically set to the default value.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

MAX_NUM_LOGFILES option for NetBackup servers and clients for legacy logging

The `MAX_NUM_LOGFILES` option specifies the maximum number of log files that you want to be retained in a NetBackup log directory. When the number of log files in the directory matches the `MAX_NUM_LOGFILES` setting, the oldest log file is deleted.

See “[About robust logging](#)” on page 260.

Table 3-97 MAX_NUM_LOGFILES information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> (or <code>bpgetconfig</code>) and the <code>nbsetconfig</code> (or <code>bpsetconfig</code>) commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Initially, the <code>bp.conf</code> file does not contain <code>MAX_NUM_LOGFILES</code> entry. In this case, the parameter is set to its default value, which is infinite.</p> <p>To set it to a different value, update the <code>bp.conf</code> file using the <code>nbsetconfig</code> command.</p> <p>Use the following format to set <code>MAX_NUM_LOGFILES</code> to a new value in the <code>bp.conf</code> file:</p> <pre>MAX_NUM_LOGFILES = X</pre> <p>'X' indicates maximum number of NetBackup log files that are created in a log directory.</p> <p><code>MAX_NUM_LOGFILES</code> should be set to a number that should be greater than one (1). If you have set <code>MAX_NUM_LOGFILES</code> to an invalid value such as 0 or 1, it is automatically set to the default value, which is infinite. However, in the <code>bp.conf</code> file, the <code>MAX_NUM_LOGFILES</code> entry appears as follows:</p> <pre>MAX_NUM_LOGFILES = 0</pre> <p>Zero (0) indicates an infinite value.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

KEEP_LOGS_SIZE_GB for NetBackup servers and clients

The `KEEP_LOGS_SIZE_GB` option specifies the size of the NetBackup logs that you want to retain. When the NetBackup log size grows up to this configuration value, the older logs are deleted.

Table 3-98 `KEEP_LOGS_SIZE_GB` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Initially, the <code>bp.conf</code> file does not contain <code>KEEP_LOGS_SIZE_GB</code> entry.</p> <p>Enable the Keep logs up to GB option in the Logging dialog box on the NetBackup Administration Console to set the log retention in GB. A corresponding entry is added in the <code>bp.conf</code> file as follows:</p> <pre>KEEP_LOGS_SIZE_GB = 25</pre> <p>If you disable the Keep logs up to GB option, the <code>bp.conf</code> file shows the corresponding entry as follows:</p> <pre>KEEP_LOGS_SIZE_GB = 0</pre> <p>To set it to a different value, update the <code>bp.conf</code> file using the <code>nbsetconfig</code> command.</p> <p>Use the following format to set <code>KEEP_LOGS_SIZE_GB</code> to a new value in the <code>bp.conf</code> file:</p> <pre>KEEP_LOGS_SIZE_GB = X</pre> <p>'X' indicates the log size in GB.</p> <p>Note: For NetBackup servers, the recommended value for the <code>KEEP_LOGS_SIZE_GB</code> option is 25 GB or greater.</p> <p>For NetBackup clients, the recommended value for the <code>KEEP_LOGS_SIZE_GB</code> option is 5 GB or greater.</p> <p>This entry should appear only once in the <code>bp.conf</code> configuration file.</p>

Table 3-98 KEEP_LOGS_SIZE_GB information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Master Servers or Clients > Double-click on server > Logging > Keep logs up to GB. See “Logging properties” on page 161.

ENABLE_CRITICAL_PROCESS_LOGGING

The `ENABLE_CRITICAL_PROCESS_LOGGING` option lets you automatically log critical NetBackup processes. Log directories for the critical processes are created and logging begins when this option is enabled in the **Logging** host properties.

Table 3-99 ENABLE_CRITICAL_PROCESS_LOGGING information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Initially, the <code>bp.conf</code> file does not contain <code>ENABLE_CRITICAL_PROCESS_LOGGING</code> entry.</p> <p>After logging for critical processes is enabled, a corresponding entry is added in the <code>bp.conf</code> file as follows:</p> <pre>ENABLE_CRITICAL_PROCESS_LOGGING = YES</pre> <p>Note: You should not modify the <code>ENABLE_CRITICAL_PROCESS_LOGGING</code> parameter. To disable the logging for critical processes, you need to modify their logging levels.</p> <p>See “Disabling the logging of NetBackup processes” on page 171.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Click Master Servers, Media Servers, Clients > Double-click on server > Logging > Logging for critical processes.</p> <p>See “Logging properties” on page 161.</p>

MEDIA_UNMOUNT_DELAY option for NetBackup servers

When `MEDIA_UNMOUNT_DELAY` is specified, the media unload is delayed for the specified number of seconds after the requested operation has completed. (Applies only to user operations.)

Table 3-100 MEDIA_UNMOUNT_DELAY information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>MEDIA_UNMOUNT_DELAY = seconds</pre> <p>The default is a media mount delay time of 180 seconds.</p> <p>This entry should appear only once in the configuration file.</p>
Example	<p>The delay is set to 120 seconds in the following example:</p> <pre>MEDIA_UNMOUNT_DELAY = 120</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Media > Media unmount delay.</p> <p>See "Media properties" on page 180.</p>

MEDIA_REQUEST_DELAY option for NetBackup servers

This option specifies the number of seconds that NetBackup waits for a non-robotic drive to become ready.

Table 3-101 MEDIA_REQUEST_DELAY information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-101 MEDIA_REQUEST_DELAY information (*continued*)

Usage	Description
How to use	<p>Change the <code>MEDIA_REQUEST_DELAY</code> by using the <code>nbemmcmd</code> command.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>The default is that NetBackup does not wait for non-robotic drives to become ready. This option should appear only once in the configuration file.</p>
Example	<p>The following command indicates to NetBackup to wait 150 seconds for a non-robotic drive to become ready for use.</p> <pre>nbemmcmd -changesetting -MEDIA_REQUEST_DELAY 150</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Media > Media request delay.</p> <p>See "Media properties" on page 180.</p>

MEDIA_SERVER option for NetBackup servers

This option is similar to the `SERVER` option.

A host that is listed as a `MEDIA_SERVER` can back up and restore clients. However, if the host is not specified as a `SERVER`, the host has limited administrative capabilities.

Table 3-102 MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>MEDIA_SERVER = media_server_name</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Servers > Media Servers.</p> <p>See "Servers properties" on page 227.</p>

MPX_RESTORE_DELAY option for NetBackup servers

This option applies to multiplexed restores. The `MPX_RESTORE_DELAY` specifies how long the server waits for restore requests of files and raw partitions. In this case, the option applies to the files and raw partitions in a set 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).

Table 3-103 MPX_RESTORE_DELAY information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>MPX_RESTORE_DELAY = seconds</pre> <p>The default is 30 seconds.</p> <p>This option should appear only once in the configuration file.</p>
Example	<p>The following example configures a server to wait 15 minutes.</p> <pre>MPX_RESTORE_DELAY = 900</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > General Server > Delay on multiplexed restores.</p> <p>See “General Server properties” on page 146.</p>

MUST_USE_LOCAL_DRIVE option for NetBackup servers

This option instructs NetBackup that if the client is also a media server and this option is present, backups for this client must occur on a local drive. If all drives are down, another may be used. If the client is not a media server, this option has no effect.

Table 3-104 MUST_USE_LOCAL_DRIVE information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-104 MUST_USE_LOCAL_DRIVE information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbemmcmd</code> command to change the option. For example:</p> <pre>nbemmcmd -changesetting -MUST_USE_LOCAL_DRIVE yes</pre> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > General Server > Must use local drive.</p> <p>See “General Server properties” on page 146.</p>

NBRNTD_IDLE_TIMEOUT option for NetBackup servers

The `NBRNTD_IDLE_TIMEOUT` option specifies the number of seconds to wait before the Remote Network Transport Service (`nbrntd`) shuts itself down. The Remote Network Transport Service provides resilient network connections. After it is shut down, NetBackup must restart the service if a new resilient connection for backups or restores is required.

Table 3-105 NBRNTD_IDLE_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>NBRNTD_IDLE_TIMEOUT = seconds</pre> <p>The default timeout is 300 seconds (five minutes).</p> <p>By default, this entry is not present in the configuration file.</p>
Example	<p>In the following example, the Remote Network Transport Service shuts off after 15 minutes.</p> <pre>NBRNTD_IDLE_TIMEOUT = 900</pre>

See [“RESILIENT_NETWORK option for NetBackup master servers and clients”](#) on page 314.

PREFERRED_NETWORK option for NetBackup servers

The `PREFERRED_NETWORK` option replaces the `REQUIRED_INTERFACE` and `REQUIRED_NETWORK` entries.

The `PREFERRED_NETWORK` option is not needed in an environment if NetBackup is configured with appropriate host names. The operating system must resolve to the correct IP addresses and then route the addresses correctly.

When external constraints prevent the environment from being corrected, `PREFERRED_NETWORK` entries can be useful in the following situations:

- To prevent NetBackup from connecting to specific destination addresses.
- To cause NetBackup to connect only to specific destination addresses.
- To request a subset of local interfaces for source binding when outbound connections are made.

Table 3-106 `PREFERRED_NETWORK` information

Usage	Description
Where to use	On NetBackup servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>The option uses the following syntax:</p> <pre>PREFERRED_NETWORK = target[/subnet] directive source[/subnet]</pre>

Table 3-106 PREFERRED_NETWORK information (*continued*)

Usage	Description
	<p>The <i>target[/subnet]</i> is a host name or range of addresses to be compared to the prospective source or destination addresses being evaluated. The following are examples of how to indicate a target or a subnet:</p> <ul style="list-style-type: none"> ■ A host name: <code>myserver.domain</code> ■ An IP address: <code>10.82.105.11</code> ■ A network with subnet: <code>10.82.105.0/21</code> ■ Any IPv4 address: <code>0.0.0.0</code> ■ Any IPv6 address: <code>0::0</code> ■ Any address: <code>0/0</code> <p>Note: A host or a network name that cannot resolve causes the <i>target</i> to be ignored. However, any associated <i>source</i> is added to the source binding list.</p>
	<p>The <i>directive</i> (MATCH, ONLY, or PROHIBITED) determines how the <i>target</i> is compared to the source or the destination address that is under evaluation.</p> <ul style="list-style-type: none"> ■ MATCH This directive applies to destination addresses. If the address being evaluated matches the <i>target</i>, then the address is immediately selected to be used and evaluation stops. If a <i>target</i> is not matched, evaluation continues with the next entry. ■ ONLY This directive applies to destination addresses. If the address under evaluation does not match a <i>target</i>, it is not used and evaluation stops for this address. If the destination is the only potential destination, the connection is not attempted. If there is an additional potential destination, it is evaluated starting over with the first entry. ■ PROHIBITED The <i>target</i> applies to both the source and the destination addresses. If a <i>source</i> is specified, it is ignored and the <i>target</i> is still prohibited. If the matched address is a destination address, evaluation stops. If the matched address was the only potential destination, no connection is attempted. If there are additional potential destinations, they are evaluated starting over with the first entry. If the matched address is a source address, it is removed from the binding list. If <i>source</i> entries exist, the shortened binding list may not be used. On some platforms, prohibiting a local interface may cause unexpected results when NetBackup connects to remote hosts. Prohibiting a local interface does not affect the connections that are internal to the host.

Table 3-106 PREFERRED_NETWORK information (*continued*)

Usage	Description
	<p>Indicating a <i>source</i> [/subnet] is optional. This is a host name or IP address range that is requested to be used as the local interface for outbound connection to addresses in <i>target</i>. The <i>source</i> applies to the directives ONLY and MATCH, but not to the directive PROHIBITED.</p> <p>Notes:</p> <ul style="list-style-type: none"> ■ The operating system may not honor the <i>source</i> request. ■ NetBackup does not request a <i>source</i> which has already been PROHIBITED. ■ A host or a network name that cannot be resolved is ignored, but the evaluation status of the <i>target</i> is still used.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Preferred Network.</p> <p>See “Preferred Network properties” on page 195.</p>

The following Support document contains more information about using the PREFERRED_NETWORK option:

<http://www.symantec.com/docs/HOWTO60779>

Examples of the PROHIBITED directive

- An example that prevents the host from using IPv4 addresses:

```
PREFERRED_NETWORK = 0.0.0.0 PROHIBITED
```

- An example that prevents the host from using IPv6 addresses:

```
PREFERRED_NETWORK = 0::0 PROHIBITED
```

- This example prevents the host from using the address of the *production_server*:

```
PREFERRED_NETWORK = production_server PROHIBITED
```

- This example prevents NetBackup from using any address between 12.34.0.0 and 12.34.255.255 as the source or destination for a connection. If it matches a local interface, NetBackup provides an ordered list of the remaining interfaces as the source binding list for the outbound interface when other entries do not specify a *source*. The operating system may always use the first one in the list instead of using the network routing tables to select the best one.

```
PREFERRED_NETWORK = 12.34.56.78/16 PROHIBITED
```

Examples of the ONLY directive

- This example instructs the host to use only IPv6 addresses in that subnet range:

```
PREFERRED_NETWORK = 2001:1234:1efc::/48 ONLY
```

- This example prevents NetBackup from connecting to any destination address outside the range of 12.34.56.0 to 12.34.56.255. The entry does not affect outbound interface selection.

```
PREFERRED_NETWORK = 12.34.56.78/24 ONLY
```

- This example prevents NetBackup from connecting to any destination address outside the range of 12.34.56.0 to 12.34.56.255. The entry requests that the operating system use 98.76.54.32 as the outbound interface.

```
PREFERRED_NETWORK = 12.34.56.78/24 ONLY 98.76.54.32
```

Examples of the MATCH directive

- This example instructs the host to use the interface IPs of *Host_A* for all IPv4 and IPv6 addresses. The entry can appear multiple times in the configuration file.

```
PREFERRED_NETWORK = 0/0 MATCH Host_A
```

- This example allows connectivity to the host names that resolve to 12.34.0.0 to 12.23.255.255. The entry does not affect outbound interface selection:

```
PREFERRED_NETWORK = 12.34.0.0/16 MATCH
```

- This example allows connectivity to the host name that resolves to 12.34.56.78, and requests that the operating system use 98.76.54.32 as the outbound interface:

```
PREFERRED_NETWORK = 12.34.56.78 MATCH 98.76.54.32
```

RANDOM_PORTS option for NetBackup servers and clients

This option specifies whether NetBackup chooses port numbers randomly or sequentially when it requires one for communication with NetBackup on other computers.

Table 3-107 RANDOM_PORTS information

Usage	Description
Where to use	On NetBackup master servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <ul style="list-style-type: none"> ■ If <code>RANDOM_PORTS = YES</code> (default), NetBackup chooses port numbers randomly from those that are free in the allowed range. For example, if the range is from 1024 through 5000, it chooses randomly from the numbers in this range. ■ If <code>RANDOM_PORTS = NO</code>, NetBackup chooses numbers sequentially, starting with the highest number available in the allowed range. For example, if the range is from 1024 through 5000, NetBackup chooses 5000 (if the number is available). If 5000 is not available, port 4999 is chosen. <p>By default, this option is not present in the configuration file and NetBackup uses the random method for selecting port numbers.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Port Ranges > Use random port assignments.</p> <p>See "Port Ranges properties" on page 192.</p>

RE_READ_INTERVAL option for NetBackup servers

The `RE_READ_INTERVAL` option determines how often NetBackup checks disk storage units for available capacity.

Table 3-108 RE_READ_INTERVAL information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>RE_READ_INTERVAL = seconds</pre> <p>The default is 300 seconds (5 minutes).</p> <p>This entry should appear only once in the configuration file.</p>

Table 3-108 RE_READ_INTERVAL information (*continued*)

Usage	Description
Example	The reread interval is changed to 15 minutes in the following example: <code>RE_READ_INTERVAL = 900</code>
Equivalent Administration Console property	NetBackup Management > Host Properties > Master Servers > Double-click on server > General Server > Check the capacity of disk storage units. See “General Server properties” on page 146.

REQUIRED_NETWORK option for NetBackup servers

The `REQUIRED_NETWORK` option specifies the required route for backup traffic in an environment where the network traffic is segregated.

For example, an environment can contain a production network at `145.21.14.0` and a backup network at `192.132.28.0`.

Table 3-109 REQUIRED_NETWORK information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>REQUIRED_NETWORK = IP_address</pre> <p>This entry should appear only once in the configuration file.</p> <p>Note: If the variable is set and the network is not available, all connections fail and no backups are performed.</p>
Example	The required network is set to <code>192.132.28.0</code> in the following example: <code>REQUIRED_NETWORK = 192.132.28.0</code>
Equivalent Administration Console property	NetBackup Management > Host Properties > Master Servers > Double-click on server > Preferred Network > Only. See “Preferred Network properties” on page 195.

RESILIENT_NETWORK option for NetBackup master servers and clients

The `RESILIENT_NETWORK` option specifies the computers that should use a resilient connection for backups and restores.

Table 3-110 `RESILIENT_NETWORK` information

Usage	Description
Where to use	On NetBackup master servers or clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use any of the following formats:</p> <pre>RESILIENT_NETWORK = hostname ON OFF RESILIENT_NETWORK = ip_address ON OFF RESILIENT_NETWORK = network address/network_mask ON OFF</pre> <p>You can mix IPv4 addresses and ranges with IPv6 addresses and subnets.</p> <p>By default, <code>RESILIENT_NETWORK</code> is not present in the configuration file.</p>
Examples	<p>The following are examples of valid forms for this entry:</p> <pre>RESILIENT_NETWORK = client.symantecs.org ON RESILIENT_NETWORK = 192.0.2.0 ON RESILIENT_NETWORK = 192.0.2.0/26 OFF RESILIENT_NETWORK = 2001:db8:0:0:0:0:0:0 ON</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Clients > Double-click on server > Resilient Network.</p> <p>See “Specifying resilient connections” on page 212.</p>

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

See “[NBRNTD_IDLE_TIMEOUT option for NetBackup servers](#)” on page 307.

RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers

The `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL` option specifies that NetBackup should perform normal duplication if an OpenStorage optimized duplication fails.

Table 3-111 `RESUME_ORIG_DUP_ON_OPT_DUP_FAIL` information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>RESUME_ORIG_DUP_ON_OPT_DUP_FAIL = TRUE FALSE</pre> <p>By default, this entry is not present and NetBackup does not perform normal duplication when an optimized duplication fails.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

REVERSE_NAME_LOOKUP option for NetBackup servers and clients

This option lets administrators allow, restrict, or prohibit reverse host name lookup.

Table 3-112 `REVERSE_NAME_LOOKUP` information

Usage	Description
Where to use	On NetBackup servers or clients.

Table 3-112 REVERSE_NAME_LOOKUP information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>REVERSE_NAME_LOOKUP = ALLOWED RESTRICTED PROHIBITED</pre> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Clients > Double-click on server > Network Settings > Reverse Host Name Lookup.</p> <p>See “Network Settings properties” on page 190.</p>

SERVER option for NetBackup servers

The first `SERVER` option in the `bp.conf` file must point to the master server where the `bp.conf` option resides. During installation, `SERVER` is automatically set to the name of the system where the NetBackup master server software is installed.

Table 3-113 SERVER information

Usage	Description
Where to use	<p>On NetBackup master servers and media servers.</p> <p>Note: For a Fibre Transport (FT) media server that has multiple network interfaces for VLANs, ensure that the FT server's primary host name appears before any other interface names for that FT media server host.</p> <p>For more information, see the NetBackup SAN Client and Fibre Transport Guide.</p>

Table 3-113 SERVER information (*continued*)

Usage	Description
How to use	<p>An entry for the <code>SERVER</code> option must be present in the configuration file on all NetBackup servers and clients. It is the only required NetBackup option. This option is not used in <code>\$HOME/bp.conf</code> files on a client.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Note: This topic discusses the <code>bp.conf</code> entries on the server. However, note that every <code>SERVER</code> option in a client <code>bp.conf</code> file must be a NetBackup master or media server. That is, each system that is listed as a <code>SERVER</code> must have either NetBackup master or media server software installed. The client service on some clients cannot be started if the client name is incorrectly listed as a server.</p> <p>If you configure NetBackup media servers for a master server, the <code>bp.conf</code> file on the master server must have a <code>SERVER</code> entry or <code>MEDIA_SERVER</code> entry for each. As previously mentioned, the first <code>SERVER</code> entry in the list designates the master server itself. The <code>SERVER</code> entry or the <code>MEDIA_SERVER</code> entries should be added after the first, self-referencing option.</p> <p>A NetBackup master server can be backed up as a NetBackup client by the servers that belong to another cluster. In that case the <code>bp.conf</code> file on the master server should have <code>SERVER</code> entries for those servers as well.</p>

Table 3-113 SERVER information (*continued*)

Usage	Description
Example	<p>The following is an example entry on a master server:</p> <pre>SERVER = Master_server (this master server itself) SERVER = NB_server (master server of another cluster) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> <p>The first <code>SERVER</code> entry on all the media servers must point to the master server for those media servers. A media server can have only one master server. However, a media server can be backed up as a NetBackup client by the servers that belong to another cluster, in which case the configuration file on the media server should have <code>SERVER</code> entries for those servers as well.</p> <p>The following is an example entry on a media server:</p> <pre>SERVER = Master_server (for this media server) SERVER = NB_server (master server of another cluster) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> <p>The <code>SERVER</code> entries must be the same on all servers in a master and a media server cluster.</p> <p>If a <code>SERVER</code> entry is added or modified in the <code>bp.conf</code> file on the master server, stop and restart <code>bprd</code> and <code>bpdcm</code> so that NetBackup recognizes the change. (The NetBackup request daemon and NetBackup database manager.)</p> <p>Note: If the first <code>SERVER</code> entry (the master server) is modified on a media server, the Enterprise Media Manager (EMM) also needs to be updated. To update EMM, run <code>nbemmcmd -updatehost</code> to change the master server for a media server.</p>

Table 3-113 SERVER information (*continued*)

Usage	Description
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Media Servers > Double-click on server > Servers.</p> <p>See “Servers properties” on page 227.</p> <p>See “FAILOVER_RESTORE_MEDIA_SERVERS option for NetBackup servers” on page 292.</p>

SERVER_CONNECT_TIMEOUT option for NetBackup servers

`SERVER_CONNECT_TIMEOUT` specifies the number of seconds that the master server waits before it times out when it connects to a media server.

Table 3-114 SERVER_CONNECT_TIMEOUT information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>SERVER_CONNECT_TIMEOUT = seconds</pre> <p>The default timeout period is 30 seconds.</p> <p>This entry should appear only once in the configuration file.</p>
Example	<p>The example permits a timeout of 60 seconds:</p> <pre>SERVER_CONNECT_TIMEOUT = 60</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Timeouts > Server connect timeout.</p> <p>See “Timeouts properties” on page 247.</p>

SERVER_PORT_WINDOW option for NetBackup servers

The `SERVER_PORT_WINDOW` option specifies the range of non-reserved ports on which NetBackup processes on this computer accept connections from NetBackup on other computers when the inbound connection is not to a well known port. This

primarily applies to `bpcd` call-back when `vnetd` is disabled in the connect options for the remote NetBackup server or client and that host is configured for non-reserved ports. This also applies to NDMP call-back to the media server during remote NDMP backups.

Table 3-115 SERVER_PORT_WINDOW information

Usage	Description
Where to use	On NetBackup master servers or media servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>SERVER_PORT_WINDOW = start_port end_port</pre> <p>The default range is from 1024 through 5000.</p> <p>This entry should appear only once in the configuration file.</p>
Example	<p>The following example permits ports from 4900 through 5000:</p> <pre>SERVER_PORT_WINDOW = 4900 5000</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Media Servers > Double-click on server > Port Ranges > Server port window.</p> <p>See "Port Ranges properties" on page 192.</p>

SERVER_RESERVED_PORT_WINDOW option for NetBackup servers and clients

This option specifies the range of local reserved ports on which this computer accepts connections from NetBackup on other computers when the inbound connection is not to a well known port. This primarily applies to `bpcd` call-back when `vnetd` is disabled in the connect options for the remote NetBackup server or client.

The `SERVER_RESERVED_PORT_WINDOW` option applies when a server connects to a client that is configured to accept only reserved ports. This option is generally not useful on clients.

Table 3-116 SERVER_RESERVED_PORT_WINDOW information

Usage	Description
Where to use	On NetBackup master servers or media servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>SERVER_RESERVED_PORT_WINDOW = start_port end_port</pre> <p>The default range is from 512 through 1023.</p> <p>This entry should appear only once in the configuration file.</p>
Example	<p>The following example permits ports from 900 through 1023:</p> <pre>SERVER_RESERVED_PORT_WINDOW = 900 1023</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers or Media Servers > Double-click on server > Port Ranges > Server reserved port window.</p> <p>See "Port Ranges properties" on page 192.</p>

SKIP_RESTORE_TO_SYMLINK_DIR option for NetBackup servers

The `SKIP_RESTORE_TO_SYMLINK_DIR` option forces NetBackup to check all directories on a UNIX client into which files are restored. If the file to be restored is under a symbolically linked directory, NetBackup does not restore the file.

Table 3-117 SKIP_RESTORE_TO_SYMLINK_DIR information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-117 SKIP_RESTORE_TO_SYMLINK_DIR information (*continued*)

Usage	Description
How to use	<p>This option should appear only once in the configuration file.</p> <p>For example, if the UNIX client requests a restore for <code>/home/user/.cshrc</code> and <code>/home/user</code> is a symbolic link, NetBackup does not restore <code>.cshrc</code>.</p> <p>The addition of <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> helps minimize potential security and data loss problems if the restore is performed with root permissions. Without <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> in the <code>bp.conf</code> file, NetBackup follows any symbolically linked directories and restores files to that location.</p> <p>Note: Restore job performance is reduced by using this option.</p> <p><code>SKIP_RESTORE_TO_SYMLINK_DIR</code> and <code>UNLINK_ON_OVERWRITE</code> do not affect each other if both are specified, with one exception:</p> <p>When the following options are enabled:</p> <ul style="list-style-type: none"> ■ Overwrite existing files option ■ <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> ■ <code>UNLINK_ON_OVERWRITE</code> <p>Then, when a restore job comes across a symbolic link, the link is unlinked before the job checks, and the files and directory are restored.</p> <p>For example, <code>/home/user/</code> is backed up as a directory and, when restored, it is a symbolic link to a directory.</p> <p>These settings have the following outcomes:</p> <ul style="list-style-type: none"> ■ With <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> set (and Overwrite existing files indicated), no files are restored into the directory the symbolic link points to, and the symbolic link remains. ■ With both <code>UNLINK_ON_OVERWRITE</code> and <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> (and Overwrite existing files indicated), the symbolic link directory is unlinked, the original directory is restored, and all files within the directory are also restored. ■ With neither option set (and Overwrite existing files indicated), NetBackup follows the symbolic link and restore all files into the directory to which the symbolic link points.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Port Ranges > Server reserved port window.</p> <p>See “Port Ranges properties” on page 192.</p>

TELEMETRY_UPLOAD option for NetBackup servers

`TELEMETRY_UPLOAD` allows NetBackup to collect data about how the product is used in a NetBackup environment. The information that Symantec receives becomes

part of a continuous quality improvement program (NetBackup Product Improvement Program) that helps Symantec understand how customers configure, deploy, and use the NetBackup product. The data is restricted to NetBackup and OpsCenter development at Symantec.

The telemetry collection agent runs on every NetBackup server and every OpsCenter server. The telemetry collection agent does not run on NetBackup clients.

Table 3-118 TELEMETRY_UPLOAD information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>TELEMETRY_UPLOAD = YES NO</pre> <p>The default is YES.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	During the NetBackup or OpsCenter installation process, a prompt or check box appears that asks for the users' agreement to participate in the Product Improvement Program.

THROTTLE_BANDWIDTH option for NetBackup servers

This option specifies 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. This option limits only backups. Restores are unaffected. The default is that the bandwidth is not limited.

While `LIMIT_BANDWIDTH` associates a bandwidth or transfer rate with all client IP addresses in a range between two IP addresses, `THROTTLE_BANDWIDTH` is more useful in an IPv6 environment. `THROTTLE_BANDWIDTH` associates a bandwidth setting with a subnet description.

For example, the following subnet will get 400kbs bandwidth:

```
2001:db8:cb30:120::/64 400
```

Table 3-119 THROTTLE_BANDWIDTH information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format in a single line:</p> <pre>THROTTLE_BANDWIDTH = xxxx:xxxx:xxxx:xxxx: yyyy:yyyy:yyyy:yyyy::/nnn zzz</pre> <p>Each <code>THROTTLE_BANDWIDTH</code> option specifies the bandwidth value and the IP address of the clients and networks to which it applies.</p> <p>The following variables are defined:</p> <ul style="list-style-type: none"> ■ <code>xxxx.xxxx.xxxx.xxxx</code> is the subnet portion (64-bits) of the IPv6 address range. (For example, <code>2001:db8:1:110</code>.) ■ <code>yyyy.yyyy.yyyy.yyyy</code> is the host portion (64-bits) of the IPv6 address range. (For example, <code>0:0:0:8b72</code>.) ■ <code>nnn</code> is the number of mask bits that, when applied to the IPv6 address, identifies the range of addresses that are considered for throttling. The valid range is 0 to 128. Mask bits are applied left to right across the address range. ■ <code>zzz</code> is the bandwidth limitation in kilobytes per second. (For example, 200.) A value of 0 disables throttling IPv6 addresses covered by this option.
Bandwidth examples	<p>The following are <code>LIMIT_BANDWIDTH</code> examples:</p> <ul style="list-style-type: none"> ■ Configure a bandwidth limit of 500 kilobytes per second for all computers on the subnet <code>2001:db8:1:110</code> as follows: <pre>LIMIT_BANDWIDTH = 2001:db8:1:110::/64 500</pre> ■ Configure a bandwidth limit of 700 kilobytes per second for a particular client (<code>2001:db8:1:110:0:0:0:8b72</code>) as follows: <pre>LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 700</pre> ■ To disable bandwidth limiting for a client in a subnet that has a bandwidth limit, specify 0 for the kilobytes per second: <pre>LIMIT_BANDWIDTH = 2001:db8:1:110::/64 500 LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 0</pre> <p>In this case, no limiting occurs for the client with IPv6 address <code>2001:db8:1:110:0:0:0:8b72</code></p>

Table 3-119 THROTTLE_BANDWIDTH information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Master Servers > Double-click on server > Bandwidth . See "Bandwidth properties" on page 75.

IPv6 address rules for NetBackup clients

The IPv6 address ranges can specify individual clients or entire subnets.

- An IPv6 address can take the following forms:
 - a.b.c.d.
Where a, b, c, and d are hexadecimal integers in the range 0-ffff.
 - a
A 32-bit integer that represents the full IP address in network byte order.
(The big endian, the most significant byte is first on the wire.)
- Enter IPv6 addresses as hexadecimal numbers.
- Neither the subnet nor the host part of an IPv6 address can be zero.
- Only ordinary IPv6 addresses are accepted.
- Do not create multiple entries that specify the same range of IPv6 addresses.
If multiple entries are created, NetBackup uses the last one found.
In the following example, NetBackup uses the second entry:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 200
```

This rule also applies to multiple entries that specify an exact client address, as follows:

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 200
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 100
```

- Do not specify IPv6 address ranges that overlap one another.
Consider the following:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
```

The ranges overlap, and bandwidth limiting results are unpredictable.

- Specify a range of addresses in one entry and an address for a specific client in other entries.

If a client is covered by an entry that specifies its exact IPv6 address and by another entry that specifies a range of IPv6 addresses, NetBackup uses the bandwidth value in the entry with the exact IP address.

The following sets the bandwidth for a range of IPv6 addresses:

```
LIMIT_BANDWIDTH = 2001:db8:1:110::/48 500
```

The following sets the bandwidth for a specific address that is within the range:

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 200
```

In this case, NetBackup uses the specific entry (bandwidth of 200) for the client whose address is 2001:db8:1:110:0:0:0:8b72. This capability can also be used to exclude specific clients from bandwidth limiting. The order of the range and specific address entries in the `bp.conf` file is not significant.

Rules for setting bandwidth values for NetBackup clients

Set bandwidths for individual clients to one of the following values:

- 0 (no bandwidth limiting), or
- Less than or equal to any value that is set for the IPv6 address range that contains the IP address for the client.

For example, the following is valid:

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/64 500
```

```
LIMIT_BANDWIDTH = 2001:db8:1:110:0:0:0:8b72::/128 300
```

If the bandwidth is set higher for a client than is set for the range, NetBackup ignores the individual setting. NetBackup uses the value for the range instead. In this case, the client receives a share of the bandwidth that is specified for the network.

If the bandwidth limit for a client is equal to or lower than the value for the range, the client uses the lower of the following settings:

- Its share of the network bandwidth value.
- Its individual bandwidth value.

The bandwidth value that NetBackup uses for a client is always at least one kilobyte per second.

TRUSTED_MASTER option for NetBackup servers

The `TRUSTED_MASTER` option lets administrators indicate a specific storage lifecycle policy in a target master server domain to configure Auto Image Replication. The ability to replicate to a specific target domain SLP is supported between MSDP storage servers and PDDO storage servers.

Table 3-120 `TRUSTED_MASTER` information

Usage	Description
Where to use	On NetBackup master servers or media servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Add <code>TRUSTED_MASTER</code> on the master server in the source domain and the master server in the target domain.</p>
Example	<p>On the source domain master server, enter the name of the target master server:</p> <pre>TRUSTED_MASTER = Target_Server_Name</pre> <p>On the target domain master server, enter the name of the source master server:</p> <pre>TRUSTED_MASTER = Source_Server_Name</pre>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Servers > Inter-domain Master Servers tab.</p> <p>See “Servers properties” on page 227.</p>

ULINK_ON_OVERWRITE option for NetBackup servers

When a UNIX client indicates **Overwrite existing files** as a restore option, the `UNLINK_ON_OVERWRITE` option forces NetBackup to perform the following actions:

- Check for the existence of a file to be restored.
- Unlink the file if it exists.
- Restore the file.

The file can be any normal file, symbolic link, hard link, or empty directory.

The addition of `UNLINK_ON_OVERWRITE` helps minimize potential security and data loss problems from following existing symbolic links. It also guarantees that files are restored exactly as they were backed up.

Table 3-121 UNLINK_ON_OVERWRITE information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>This option should appear only once in the configuration file.</p> <p>Note: Restore job performance is reduced by using this option.</p> <p>If the <code>UNLINK_ON_OVERWRITE</code> option is not indicated in the <code>bp.conf</code> file but the Overwrite existing files option is specified, the behavior of NetBackup is different when it restores symbolic links. (Or, if the option is set to NO.) NetBackup unlinks existing files or empty directories when it restores symbolic links, hard links, or special files (CHR, BLK, and FIFO).</p> <p>However, NetBackup does not unlink when it restores normal files or directories, which can be problematic. NetBackup follows the symbolic link and creates or replaces the files that the symbolic link points to. Or, it replaces the directory that the symbolic link points to.</p> <p><code>SKIP_RESTORE_TO_SYMLINK_DIR</code> and <code>UNLINK_ON_OVERWRITE</code> do not affect each other if both are specified, with one exception:</p> <p>When the following options are enabled:</p> <ul style="list-style-type: none"> ■ Overwrite existing files option ■ <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> ■ <code>UNLINK_ON_OVERWRITE</code> <p>Then, when a restore job comes across a symbolic link, the link is unlinked before the job checks, and the files and directories are restored.</p>
Example	<p>For example, if the <code>/home/user/</code> path was backed up as a directory and, when restored, it is a symbolic link to a directory:</p> <ul style="list-style-type: none"> ■ With <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> set (and Overwrite existing files indicated), no files are restored into the directory the symbolic link points to, and the symbolic link remains. ■ With both <code>UNLINK_ON_OVERWRITE</code> and <code>SKIP_RESTORE_TO_SYMLINK_DIR</code> (and Overwrite existing files indicated), the symbolically linked directory is unlinked, the original directory is restored, and all files within the directory are also restored. ■ With neither set (and Overwrite existing files indicated), NetBackup follows the symbolic link and restore all files into the directory the symbolic link points to.
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

USE_AUTHENTICATION option for NetBackup servers

The `USE_AUTHENTICATION` option can be used to audit user information without enabling NetBackup Access Control (NBAC).

Table 3-122 USE_AUTHENTICATION information

Usage	Description
Where to use	On NetBackup master servers.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>USE_AUTHENTICATION = ON OFF</pre> <p>The option uses the following parameters:</p> <ul style="list-style-type: none"> ■ ON The user is authenticated and is audited. Authentication is achieved without enabling NBAC. ■ OFF The user is not authenticated and is not allowed to perform any operation. Only a root user is allowed to make changes and is audited accordingly. <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

USE_VXSS option for NetBackup servers and clients

The `USE_VXSS` option specifies whether the local system uses NetBackup product authentication and authorization.

Table 3-123 USE_VXSS information

Usage	Description
Where to use	On NetBackup master servers or clients.

Table 3-123 USE_VXSS information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>USE_VXSS = REQUIRED PROHIBITED AUTOMATIC</pre> <p>The option uses the following parameters:</p> <ul style="list-style-type: none"> ■ REQUIRED Indicates that the local system always uses NetBackup product authentication and authorization. Connections from the systems that do not use NetBackup product authentication and authorization are rejected. ■ PROHIBITED (Default.) Indicates that the local system never uses NetBackup product authentication and authorization. Connections from the systems that use NetBackup product authentication and authorization are rejected. ■ AUTOMATIC Indicates that the local system negotiates with the remote system whether to use NetBackup product authentication and authorization. If the <code>USE_VXSS = AUTOMATIC</code> option is specified, <code>VXSS_NETWORK</code> entries can be used to require or prohibit NetBackup product authentication and authorization connections with specified remote systems. See “VXSS_NETWORK option for NetBackup servers” on page 331. <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Access Control > NetBackup product authentication and authorization. Change this option to Required.</p> <p>See “Network Attributes tab of the Access Control properties” on page 70.</p>

VERBOSE option for NetBackup servers and clients

Used for debugging purposes, the `VERBOSE` option controls the amount of information NetBackup includes in its legacy logs.

Table 3-124 VERBOSE information

Usage	Description
Where to use	On NetBackup master servers or clients.

Table 3-124 VERBOSE information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>VERBOSE = [0 1 2 3 4 5]</pre> <p>By default, this option is disabled.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Logging > Global logging level.</p> <p>See “Logging properties” on page 161.</p>

VXSS_NETWORK option for NetBackup servers

This option identifies whether a specific network or remote system must or must not use **NetBackup product authentication and authorization** with the local system.

If a media server or client does not define a **NetBackup product authentication and authorization** network, it uses the **NetBackup product authentication and authorization** networks of its master server.

Table 3-125 VXSS_NETWORK information

Usage	Description
Where to use	On NetBackup master servers.

Table 3-125 VXSS_NETWORK information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>VXSS_NETWORK = hostname IP_address .domain network. [AUTOMATIC REQUIRED PROHIBITED]</pre> <p>VXSS_NETWORK is relevant only if USE_VXSS is set to AUTOMATIC (USE_VXSS = AUTOMATIC). More than one VXSS_NETWORK option can be specified.</p> <p>See “USE_VXSS option for NetBackup servers and clients” on page 329.</p> <p>The option uses the following parameters:</p> <ul style="list-style-type: none"> ■ <code>hostname</code> The host name of the remote system. ■ <code>IP_address</code> The IP address of the remote system. ■ <code>.domain</code> A dot followed by the Internet domain name of the remote systems. ■ <code>network.</code> The network for the remote systems, followed by a dot. <p>The optional second value can be one of the following keywords:</p> <ul style="list-style-type: none"> ■ AUTOMATIC ■ REQUIRED ■ PROHIBITED <p>Note: If multiple VXSS_NETWORK entries specify one particular system, the first occurrence takes precedence.</p>

Table 3-125 VXSS_NETWORK information (*continued*)

Usage	Description
Example	<pre>USE_VXSS = AUTOMATIC VXSS_NETWORK = fred.mycompany.com VXSS_NETWORK = 10.0.0.37 REQUIRED VXSS_NETWORK = 10.0.0. PROHIBITED VXSS_NETWORK = .theircompany.com VXSS_NETWORK = wilma.theircompany.com PROHIBITED VXSS_NETWORK = barney.mycompany.com PROHIBITED</pre> <p>In the example, NetBackup product authentication and authorization is required for connections between the local system and the system with host fred.mycompany.com.</p> <p>NetBackup product authentication and authorization is required for connections between the local system and the system with IP address 10.0.0.37.</p> <p>NetBackup product authentication and authorization is prohibited for connections between the local system and systems in the 10.0.0 network except for 10.0.0.37.</p> <p>NetBackup product authentication and authorization is required for connections between the local system and systems within the theircompany.com Internet domain.</p> <p>NetBackup product authentication and authorization is required for connections between the local system and the system with host name wilma.theircompany.com despite the PROHIBITED option for wilma.theircompany.com. The REQUIRED option for .theircompany.com takes precedence.</p> <p>NetBackup product authentication and authorization is prohibited for connections between the local system and the system with host name barney.mycompany.com</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Master Servers > Double-click on server > Access Control. Change the networks list property.</p> <p>See “Network Attributes tab of the Access Control properties” on page 70.</p>

Configuration options for NetBackup clients

The following topics are about configuration options for NetBackup clients. Nearly all of these options can also be set in the Host Properties in the **NetBackup Administration Console**.

See [“Configuration options not found in the Host Properties”](#) on page 257.

BPARCHIVE_POLICY option for NetBackup clients

The `BPARCHIVE_POLICY` entry specifies the name of the policy to use for user archives.

Table 3-126 BPARCHIVE_POLICY information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPARCHIVE_POLICY = policy_name</pre> <p>By default, <code>BPARCHIVE_POLICY</code> is not present in the configuration file. By default, NetBackup uses the first policy that it finds that has the client and a user archive schedule.</p> <p>If it is used, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists.</p> <p>See "Type of backup (schedule attribute)" on page 745.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

BPARCHIVE_SCHED option for NetBackup clients

This entry specifies the name of the schedule for user archives.

Table 3-127 BPARCHIVE_SCHED information

Usage	Description
Where to use	On NetBackup clients.

Table 3-127 BPARCHIVE_SCHED information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPARCHIVE_SCHED = schedule_name</pre> <p>By default, <code>BPARCHIVE_SCHED</code> is not present in the configuration file. By default, NetBackup uses the first archive schedule in the first policy that it finds that contains this client.</p> <p>If it is used, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists.</p> <p>See "Considerations for user schedules" on page 753.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

BPBACKUP_POLICY option for NetBackup clients

This entry specifies the name of the policy to use for user backups.

Table 3-128 BPBACKUP_POLICY information

Usage	Description
Where to use	On NetBackup clients.

Table 3-128 BPBACKUP_POLICY information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPBACKUP_POLICY = policy_name</pre> <p>By default, <code>BPBACKUP_POLICY</code> is not present in the configuration file. By default, NetBackup uses the first policy it finds that has both the client and a user backup schedule.</p> <p>If present, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists.</p> <p>See "Type of backup (schedule attribute)" on page 745.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

BPBACKUP_SCHED option for NetBackup clients

This entry specifies the name of the schedule to use for user backups.

Table 3-129 BPBACKUP_SCHED information

Usage	Description
Where to use	On NetBackup clients.

Table 3-129 BPBACKUP_SCHED information (continued)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BPBACKUP_SCHED = schedule_name</pre> <p>By default, <code>BPBACKUP_SCHED</code> is not present in the configuration file. By default, NetBackup uses the first policy it finds that contains both the client and a user backup schedule.</p> <p>If present, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists.</p> <p>See "Considerations for user schedules" on page 753.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

BUSY_FILE_ACTION option for NetBackup clients

The `BUSY_FILE_ACTION` entry directs the action that NetBackup performs on busy files when busy-file processing is enabled.

Table 3-130 BUSY_FILE_ACTION information

Usage	Description
Where to use	On NetBackup clients.

Table 3-130 BUSY_FILE_ACTION information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p><code>BUSY_FILE_ACTION = filename_template action_template</code></p> <ul style="list-style-type: none"> ■ filename_template The absolute pathname and file name of the busy file. The shell language metacharacters <code>*</code>, <code>?</code>, <code>[]</code>, <code>[-]</code> can be used for matching patterns of file names or parts of file names. ■ action_template Use one of the following parameters: <ul style="list-style-type: none"> ■ <code>MAIL mail</code> Directs NetBackup to email a busy file notification message to the user that the <code>BUSY_FILE_NOTIFY_USER</code> option specifies. ■ <code>REPEAT repeat [repeat_count]</code> Directs NetBackup to retry the backup on the specified busy file. A repeat count can be specified to control the number of backup attempts. The default repeat count is 1. ■ <code>IGNORE ignore</code> Directs NetBackup to exclude the busy file from processing. <p>Multiple <code>BUSY_FILE_ACTION</code> entries are allowed.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Busy File Settings.</p> <p>See “Busy File Settings properties” on page 77.</p>

BUSY_FILE_DIRECTORY option for NetBackup clients

The `BUSY_FILE_DIRECTORY` entry specifies the path to the busy-files working directory when busy-file processing is enabled.

Table 3-131 BUSY_FILE_DIRECTORY information

Usage	Description
Where to use	On NetBackup clients.

Table 3-131 BUSY_FILE_DIRECTORY information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BUSY_FILE_DIRECTORY = pathname</pre> <p>By default, <code>BUSY_FILE_DIRECTORY</code> is not present in the configuration file. By default, NetBackup creates the <code>busy_files</code> directory in <code>/usr/opensv/netbackup</code>.</p> <p>If present, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Busy File Settings.</p> <p>See “Busy File Settings properties” on page 77.</p>

BUSY_FILE_NOTIFY_USER option for NetBackup clients

The `BUSY_FILE_NOTIFY_USER` entry specifies who receives a notification when the `BUSY_FILE_ACTION` entry is set to `MAIL`.

Table 3-132 BUSY_FILE_NOTIFY_USER information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BUSY_FILE_NOTIFY_USER = email</pre> <p>By default, <code>BUSY_FILE_NOTIFY_USER</code> is not present in the configuration file. By default, the email recipient is <code>root</code>.</p> <p>If present, this option should appear only once in the configuration file.</p> <p>The value in the user's <code>\$HOME/bp.conf</code> file takes precedence, if it exists.</p>

Table 3-132 BUSY_FILE_NOTIFY_USER information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Busy File Settings . See “Busy File Settings properties” on page 77.

BUSY_FILE_PROCESSING option for NetBackup clients

The `BUSY_FILE_PROCESSING` entry lets the administrator control what NetBackup does when a file changes while it is in the process of being backed up.

Table 3-133 BUSY_FILE_PROCESSING information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>BUSY_FILE_PROCESSING = yes</pre> <p>By default, <code>BUSY_FILE_PROCESSING</code> is not present in the configuration file and busy-file processing does not occur.</p> <p>If present, this option should appear only once in the configuration file.</p>
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Busy File Settings . See “Busy File Settings properties” on page 77.

CLIENT_NAME option for NetBackup clients

The `CLIENT_NAME` entry specifies the name of the client as it's known to NetBackup.

Table 3-134 CLIENT_NAME information

Usage	Description
Where to use	On NetBackup clients.

Table 3-134 CLIENT_NAME information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CLIENT_NAME = clientname</pre> <p>This option should appear only once in the configuration file.</p> <p>If more than one entry appears, NetBackup observes only the last <code>CLIENT_NAME</code> entry that is listed for the option. The client name in a policy that backs up the client should match the client name that is specified with <code>CLIENT_NAME</code>.</p> <p>Note: Do not use an IP address as a client name in a policy or the backup may fail. Specify a host name instead.</p> <p>The <code>bp.conf</code> of the master server does not require the addition of other clients, other than the master server as <code>CLIENT_NAME = master server name</code>. The name is added by default.</p> <p>During a restore, the default is to restore to the client that is named in the policy that was used for the backup. For an alternate client restore, indicate the name of the alternate client in the Backup, Archive, and Restore user interface. (Within the user interface, the Destination client for restores field is located in the Specify NetBackup Machines and Policy Type dialog box.)</p> <p>To use the <code>bprestore</code> command, a parameter can be given to <code>bprestore</code> to indicate the destination client.</p> <p>See “About client-redirected restores” on page 1150.</p> <p>The client installation procedures automatically set <code>CLIENT_NAME</code> to the value that is specified in <code>ftp_to_client</code> command or <code>install_client</code> command in the installation scripts.</p> <p>If the value is not in any <code>bp.conf</code> file, NetBackup uses the value that the <code>gethostname()</code> library function returns.</p> <p>See “Client Name properties” on page 82.</p>
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Client Name.

COMPRESS_SUFFIX option for NetBackup clients

The `COMPRESS_SUFFIX` entry 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.

Table 3-135 `COMPRESS_SUFFIX` information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>COMPRESS_SUFFIX = .suffix</pre> <p>By default, <code>COMPRESS_SUFFIX</code> is not present in the <code>bp.conf</code> file. This option has a reasonable default. Change only if problems result.</p> <p>Multiple <code>COMPRESS_SUFFIX</code> entries are allowed.</p> <p>Do not use wildcards to specify these extensions. Do not specify <code>.A*</code> or <code>.A [1-9]</code> (For example, specify <code>.A1</code>)</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Clients Double-click on client > Client Settings.</p> <p>See “Client Settings properties for UNIX clients” on page 99.</p>

CRYPT_CIPHER option for NetBackup clients

The `CRYPT_CIPHER` entry applies to clients with the NetBackup Encryption option installed.

Table 3-136 `CRYPT_CIPHER` information

Usage	Description
Where to use	On NetBackup clients.

Table 3-136 CRYPT_CIPHER information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_CIPHER = value</pre> <p>Where <i>value</i> is one of the following:</p> <ul style="list-style-type: none">■ AES-128-CFB (used when no method is specified; default)■ AES-256-CFB■ BF-CFB■ DES-EDE-CFB <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p> <p>See the NetBackup Security and Encryption Guide.</p>

CRYPT_KIND option for NetBackup clients

The `CRYPT_KIND` entry on the client determines whether the standard encryption or legacy encryption is used in the backup. Normally, `CRYPT_KIND` is set automatically.

The `CRYPT_KIND` entry applies to clients with the NetBackup Encryption option installed.

See the [NetBackup Security and Encryption Guide](#).

Table 3-137 CRYPT_KIND information

Usage	Description
Where to use	On NetBackup clients.

Table 3-137 CRYPT_KIND information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_KIND = type</pre> <p>The following values can be entered:</p> <ul style="list-style-type: none">■ <code>NONE</code> No encryption is used on the client (default)■ <code>LEGACY</code> Legacy pertains to 40-bit and 56-bit data encryption standard (DES). Legacy encryption is not recommended.■ <code>STANDARD</code> Standard pertains to 128-bit and 256-bit encryption (AES, 3DES, Blowfish cipher). Standard encryption is recommended. <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p>

CRYPT_OPTION option for NetBackup clients

The `CRYPT_OPTION` entry specifies the encryption options on NetBackup clients. NetBackup creates this entry automatically in the `/usr/openv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup master server.

The `CRYPT_OPTION` entry applies to clients with the NetBackup Encryption option installed.

See the [NetBackup Security and Encryption Guide](#).

Do not alter the entry or create this file manually unless it was accidentally deleted.

Table 3-138 CRYPT_OPTION information

Usage	Description
Where to use	On NetBackup clients.

Table 3-138 CRYPT_OPTION information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_OPTION = type</pre> <p>The following values can be entered:</p> <ul style="list-style-type: none">■ <code>DENIED denied</code> Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, it is considered an error. This option is the default for a client that has not been configured for encryption.■ <code>ALLOWED allowed</code> Specifies that the client allows either encrypted or unencrypted backups.■ <code>REQUIRED required</code> Specifies that the client requires encrypted backups. If this value is specified and the server requests an unencrypted backup, it is considered an error. <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p>

CRYPT_STRENGTH option for NetBackup clients

The `CRYPT_STRENGTH` entry specifies the encryption strength on NetBackup clients. It applies to clients with the NetBackup Encryption option installed.

See the [NetBackup Security and Encryption Guide](#).

Table 3-139 CRYPT_STRENGTH information

Usage	Description
Where to use	On NetBackup clients.

Table 3-139 CRYPT_STRENGTH information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_STRENGTH = value</pre> <p>The allowable values are as follows:</p> <ul style="list-style-type: none">■ <code>DES_40 des_40</code> Specifies 40-bit DES encryption. 40-bit is the default value for a client that has not been configured for encryption.■ <code>DES_56 des_56</code> Specifies 56-bit DES encryption. <p>NetBackup creates this entry automatically on a UNIX client when the <code>bpinst_crypt</code> command is run on the NetBackup master server. This entry should appear only once in the configuration file.</p> <p>Do not alter the entry or create it manually unless it was accidentally deleted.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p>

CRYPT_LIBPATH option for NetBackup clients

The `CRYPT_LIBPATH` entry specifies the directory that contains the encryption libraries for NetBackup clients. NetBackup creates this entry automatically in the `/usr/opensv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup master server.

The `CRYPT_LIBPATH` entry applies to clients with the NetBackup Encryption option installed. Do not alter the entry or create it manually unless it was accidentally deleted. This entry should appear only once in the configuration file.

See the [NetBackup Security and Encryption Guide](#).

Table 3-140 CRYPT_LIBPATH information

Usage	Description
Where to use	On NetBackup clients.

Table 3-140 CRYPT_LIBPATH information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_LIBPATH = <i>directory</i></pre> <p>If necessary, create the entry in the following locations:</p> <ul style="list-style-type: none">■ The default value on UNIX systems is <code>/usr/opensv/lib/</code>■ The default value on Windows systems is <code>install_path\bin\</code> Where <i>install_path</i> is the directory where NetBackup is installed and by default is <code>C:\Program Files\VERITAS</code>.
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p>

CRYPT_KEYFILE option for NetBackup clients

The `CRYPT_KEYFILE` entry specifies the file that contains the encryption keys on NetBackup clients. NetBackup creates this entry automatically in the `/usr/opensv/netbackup/bp.conf` file on a UNIX client when the `bpinst_crypt` command is run on the NetBackup master server.

See the [NetBackup Security and Encryption Guide](#).

Table 3-141 CRYPT_KEYFILE information

Usage	Description
Where to use	On NetBackup clients.

Table 3-141 CRYPT_KEYFILE information (*continued*)

Usage	Description
How to use	<p>The <code>CRYPT_KEYFILE</code> entry applies to clients with the NetBackup Encryption option installed. Do not alter the entry or create it manually unless it was accidentally deleted.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>CRYPT_KEYFILE = directory</pre> <p>The default values follow:</p> <ul style="list-style-type: none"> On Windows systems <code>install_path\bin\keyfile.dat</code> Where <code>install_path</code> is the directory where NetBackup is installed and by default is <code>C:\Program Files\VERITAS</code>. On UNIX systems <code>/usr/opensv/netbackup/keyfile</code> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Encryption.</p> <p>See “Encryption properties” on page 118.</p>

DISALLOW_SERVER_FILE_WRITES option for NetBackup clients

The `DISALLOW_SERVER_FILE_WRITES` entry prevents the NetBackup server from creating files on the NetBackup client. For example, this entry prevents server-directed restores or server-directed updates of the configuration file on the client.

Table 3-142 DISALLOW_SERVER_FILE_WRITES information

Usage	Description
Where to use	On NetBackup clients.

Table 3-142 DISALLOW_SERVER_FILE_WRITES information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DISALLOW_SERVER_FILE_WRITES</pre> <p>By default, server writes are allowed.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Universal Settings.</p> <p>See "Universal Settings properties" on page 249.</p>

DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients

The `DO_NOT_RESET_FILE_ACCESS_TIME` entry specifies that if a file is backed up, its access time (`atime`) displays the time of the backup. The default is that NetBackup preserves the access time by resetting it to the value it had before the backup.

Note: The `DO_NOT_RESET_FILE_ACCESS_TIME` entry affects the software and the administration scripts that examine a file's access time. Do not use this option or `USE_CTIME_FOR_INCREMENTALS` if Storage Migrator is on the system. Setting these options causes the `atime` for files to be updated every time they are backed up. Each time the `atime` updates, it appears as if the files are frequently used and prevents Storage Migrator from selecting the files for migration.

Table 3-143 DO_NOT_RESET_FILE_ACCESS_TIME information

Usage	Description
Where to use	On NetBackup clients.

Table 3-143 DO_NOT_RESET_FILE_ACCESS_TIME information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>DO_NOT_RESET_FILE_ACCESS_TIME</pre> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Client Settings.</p> <p>See “Client Settings properties for UNIX clients” on page 99.</p>

IGNORE_XATTR option for NetBackup clients

By default, extended attribute files (Solaris 9 or later) and named data streams are backed up. Use `IGNORE_XATTR` to exclude extended attributes and named data streams from backups. (`IGNORE_XATTR` was formerly `IGNORE_XATTR_SOLARIS`.)

Table 3-144 IGNORE_XATTR information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>IGNORE_XATTR</pre> <p>NetBackup does not check for extended attributes or named data streams if the <code>IGNORE_XATTR</code> entry is present in the <code>bp.conf</code> file.</p> <p>This option should appear only once in the configuration file.</p> <p>See “About backing up and restoring extended attribute files and named data streams” on page 815.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

Backing up and restoring the ACLs, extended attributes, and metadata attributes for GPFS volumes

NetBackup recognizes the access control lists, extended attributes, and metadata attributes of General Parallel File System files and folders. By default, NetBackup backs up and restores all of these for GPFS volumes. No additional configuration is necessary.

Restoring files with GPFS Extended Attributes to operating systems other than AIX or RHEL or file systems other than GPFS may generate errors or failed jobs. The errors occur because the restore target does not recognize the metadata.

Preventing the backup or restore of the ACL and extended attributes for a GPFS volume

- To prevent backups of GPFS extended attributes:
Add the `IGNORE_XATTR` entry to the `bp.conf` file on the client:

```
/usr/opensv/netbackup/bp.conf
```


The entry does not require a value setting; the entry is only `IGNORE_XATTR`.
- To prevent backups of the GPFS ACL:
Add the `IGNORE_ACL` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_ACL
```

Note: If extended attributes are allowed to be backed up, the ACL is also backed up, regardless of whether or not the `IGNORE_ACL` touch file is present.

- To prevent restores of GPFS extended attributes (provided that they were backed up in the first place):
Add the `IGNORE_XATTR` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_XATTR
```
- To prevent restores of the GPFS ACL (provided that it was backed up in the first place):
Add the `IGNORE_ACL` touch file to the client:

```
/usr/opensv/netbackup/IGNORE_ACL
```

Note: If extended attributes are allowed for restore, the ACL is also restored, regardless of whether or not the `IGNORE_ACL` touch file is present.

Considerations when backing up and restoring GPFS extended attributes

- Upon restore of a file or folder that has an ACL or extended attributes, the ACL and extended attributes are also restored if they were originally backed up. However, if the restore job was interrupted, the last file that was in-process may be skipped when the restore job resumes. In that case, the ACL and the extended attributes of that file and all subsequent files will not be restored.
- The `IGNORE_XATTR` option instructs NetBackup to ignore extended attributes when backing up files and folders that have extended attributes. In case of a NetBackup Accelerator initial backup, if `IGNORE_XATTR` is added to the `bp.conf` file or the registry and then is later removed for subsequent backups, perform a backup with the **Accelerator forced rescan** option enabled.
See [“Accelerator forced rescan option \(schedule attribute\)”](#) on page 755.
- If the ACL interface is used to change ACL permissions after a backup is performed, the restore may not preserve the ACL upon restore.

Support for GPFS metadata attributes: storage pools, metadata replication, and data replication

NetBackup has added support for the backup and restore of the following GPFS metadata attributes: storage pools, metadata replication, and data replication. By default, NetBackup backs up and restores the ACLs, the extended attributes, and these additional metadata attributes for GPFS volumes. No additional configuration is necessary.

The following information is pertinent if, after restoring a file, the `illplaced` or the `illreplicated` flags display in the status of the file. (View the status of the file by using the GPFS `mmfsattr` command.)

The `illplaced` flag displays if the file was restored to a different GPFS storage pool than where it was when it was backed up. The change could be due to a change in creation rules or the result of migration rules in the GPFS file placement policy. Or the GPFS administrator may have used the `mmchattr` command to manually move the file to a different storage pool.

The `illreplicated` flag displays when the storage pool containing the restored file has fewer failure groups than required by the data replication setting or the system pool has fewer failure groups than the metadata replication setting.

The GPFS administrator can rebalance the replication factor of the file and resolve the `illplaced` or `illreplicated` flags by running one of two GPFS commands, depending on the number of the files that are involved and the network traffic:

- For a single file: `mmrestripefile`

- For the complete file system: `mmrestripefs`

Note: Restoring files with GPFS attributes and ACLs to an alternate platform does not restore the metadata attributes or the ACLs. For example, backing up files from an RHEL system and restoring to an AIX system. The restore may generate an error such as “Invalid system call.” (Extended attributes can be restored to the alternate platform, however.)

Notes about support for GPFS metadata attributes

- Restoring files with GPFS attributes and ACLs to an alternate platform does not restore the metadata attributes or the ACLs. For example, backing up files from an RHEL system and restoring to an AIX system. The restore may generate an error such as “Invalid system call.” (Extended attributes can be restored to the alternate platform, however.)
- When NetBackup 7.6.1 is used to restore a backup that was created with a NetBackup 7.6.1.1 client, any GPFS Extended Attributes or newly supported GPFS metadata attributes that were backed up are not restored. (ACLs are restored, however.) Depending on the verbose logging level used, the progress log may display a message similar to the following:

```
Unknown extended attribute (3 - META_GPFS) - information will be
lost
```

INFORMIX_HOME option for NetBackup clients

The `INFORMIX_HOME` entry specifies the path to the Informix home directory and is required when the client uses NetBackup for Informix.

Table 3-145 INFORMIX_HOME information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>INFORMIX_HOME</pre> <p>This option should appear only once in the configuration file.</p>

Table 3-145 INFORMIX_HOME information (*continued*)

Usage	Description
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

KEEP_DATABASE_COMM_FILE option for NetBackup clients

The `KEEP_DATABASE_COMM_FILE` entry causes NetBackup to keep database agent logs for seven days. The default is that NetBackup keeps database agent logs for only one day.

Table 3-146 KEEP_DATABASE_COMM_FILE information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>KEEP_DATABASE_COMM_FILE</pre> <p>For example, add it to a client that runs NetBackup for Informix.</p> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

KEEP_LOGS_DAYS option for NetBackup clients

The `KEEP_LOGS_DAYS` entry specifies how long to keep job and progress logs of the NetBackup-Java program, **Backup**, **Archive**, and **Restore**. The default is 3 days.

NetBackup writes these files in 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 **Backup**, **Archive**, and **Restore** console. This entry also controls how long to keep the log files from the NetBackup-Java console.

The log files are located in `/usr/opensv/netbackup/logs/user_ops/nbjlogs`.

Table 3-147 KEEP_LOGS_DAYS information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option file.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>KEEP_LOGS_DAYS</pre> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Client Settings.</p> <p>See “Client Settings properties for UNIX clients” on page 99.</p> <p>See “Client Settings properties for Windows clients” on page 104.</p>

LIST_FILES_TIMEOUT option for NetBackup clients

The `LIST_FILES_TIMEOUT` entry specifies how long to wait for a response from the server when it lists files by using the client-user interface or `bplist`. If this time is exceeded, the user receives a `socket read failed` error even if the server continues to process the user’s request. The default is that `LIST_FILES_TIMEOUT` is not in any `bp.conf` file and NetBackup uses a value of 30 minutes.

Table 3-148 LIST_FILES_TIMEOUT information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>LIST_FILES_TIMEOUT</pre> <p>This option should appear only once in the configuration file.</p> <p>The value in the <code>\$HOME/bp.conf</code> file takes precedence if it exists.</p>

Table 3-148 LIST_FILES_TIMEOUT information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Timeouts . See “Timeouts properties” on page 247.

LOCKED_FILE_ACTION option for NetBackup clients

The `LOCKED_FILE_ACTION` entry specifies the behavior of NetBackup when it backs up a file that has mandatory file locking enabled in its file mode. (See `chmod(1)`). If this entry is set to `SKIP`, NetBackup skips the files that currently have mandatory locking set by another process. NetBackup logs a message to this effect.

Table 3-149 LOCKED_FILE_ACTION information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>LOCKED_FILE_ACTION = SKIP</pre> <p>This option should appear only once in the configuration file.</p> <p>Note: <code>SKIP</code> is the only legal value for this entry. The default is that NetBackup waits for files to become unlocked.</p>
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Client Settings . See “Client Settings properties for UNIX clients” on page 99.

MEDIA_SERVER option for NetBackup clients

The `MEDIA_SERVER` entry specifies that the listed computer is a media server only. Computers that are listed as media servers can back up and restore clients, but have limited administrative privileges.

Table 3-150 MEDIA_SERVER information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>MEDIA_SERVER = media_server_name</pre> <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Servers.</p> <p>See “SERVER option for NetBackup servers” on page 316.</p>

MEGABYTES_OF_MEMORY option for NetBackup clients

The `MEGABYTES_OF_MEMORY` entry specifies how much memory is available on the client to use to compress files during backup. If compression is selected, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to the compress code, the greater the compression. The percentage of computer resources that are used is also greater. If other processes also need memory, use a maximum value of one half the actual physical memory on a computer to avoid excessive swapping.

Table 3-151 MEGABYTES_OF_MEMORY information

Usage	Description
Where to use	On NetBackup clients.

Table 3-151 MEGABYTES_OF_MEMORY information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>MEGABYTES_OF_MEMORY = memory_value</pre> <p>The default is that NetBackup assumes a value of zero megabytes.</p> <p>This option should appear only once in the configuration file.</p> <p>Note: The <code>MEGABYTES_OF_MEMORY</code> entry has a reasonable default. Change it only if problems are encountered.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Client Settings.</p> <p>See “Client Settings properties for UNIX clients” on page 99.</p>

NFS_ACCESS_TIMEOUT option for NetBackup clients

The `NFS_ACCESS_TIMEOUT` entry specifies the number of seconds that the backup process waits to process an NFS mount table. After the time is exceeded, the process considers an NFS file system to be unavailable.

Table 3-152 NFS_ACCESS_TIMEOUT information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>NFS_ACCESS_TIMEOUT = seconds</pre> <p>By default, <code>NFS_ACCESS_TIMEOUT</code> is set to five seconds.</p> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > UNIX Server.</p> <p>See “UNIX Server properties” on page 254.</p>

OLD_VNETD_CALLBACK option for NetBackup clients

The `OLD_VNETD_CALLBACK` entry specifies that the client should use the client-direct restore path rather than use the `bptm` process on a NetBackup media server. The client-direct restore path bypasses the NetBackup media server for the restore jobs. Configure this entry on the NetBackup clients that you want to receive data directly from a **Media Server Deduplication Pool** storage server.

Before you decommission a media server that hosts a **Media Server Deduplication Pool**, deactivate MSDP on that media server. See the [NetBackup Deduplication Guide](#) for your release.

Table 3-153 OLD_VNETD_CALLBACK information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>OLD_VNETD_CALLBACK = YES</pre> <p>This entry should appear only once in the configuration file.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

REPORT_CLIENT_DISCOVERIES option for NetBackup clients

By default, the NetBackup Discovery Service (`nbdisco`) runs on all clients in the NetBackup environment. The service helps you build Intelligent Policies by reporting to the master server when it finds instances of applications (such as Oracle).

Table 3-154 REPORT_CLIENT_DISCOVERIES information

Usage	Description
Where to use	On NetBackup clients.

Table 3-154 REPORT_CLIENT_DISCOVERIES information (*continued*)

Usage	Description
How to use	<p>By default, <code>REPORT_CLIENT_DISCOVERIES</code> is not present in the configuration file. When <code>REPORT_CLIENT_DISCOVERIES</code> is not present, the NetBackup Discovery Service is enabled.</p> <p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>REPORT_CLIENT_DISCOVERIES = TRUE FALSE</pre> <p>This option should appear only once in the configuration file.</p> <ul style="list-style-type: none">■ If the entry is set to <code>FALSE</code>, the <code>nbdisco</code> process on that client stops reporting discoveries to the master server. The service shuts down within 10 minutes after being set to <code>FALSE</code> and remains down.■ To turn on the Discovery Service again, either change the entry on that client to <code>REPORT_CLIENT_DISCOVERIES = TRUE</code> or remove the entire option. Then, run <code>bp.start_all</code> on the client to restart the service.■ To set this value on a client remotely, run the following command from the master server: <pre>nbsetconfig -h clientname</pre>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

RESTORE_RETRIES option for NetBackup clients

The `RESTORE_RETRIES` entry specifies the number of times to retry a restore after a failure.

Table 3-155 RESTORE_RETRIES information

Usage	Description
Where to use	On NetBackup clients.

Table 3-155 RESTORE_RETRIES information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>RESTORE_RETRIES = number_of_retries</pre> <p>The default is 0 (no retries).</p> <p>This option should appear only once in the configuration file.</p> <p>Note: The <code>RESTORE_RETRIES</code> entry has a reasonable default. Change it only if problems are encountered.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Universal Settings > Restore retries.</p> <p>See “Universal Settings properties” on page 249.</p>

SERVER option for NetBackup clients

The `SERVER` entry defines the list of NetBackup master servers and media servers that can access the NetBackup client. During client installation, `SERVER` is automatically set to the name of the primary master server for this client.

Other `SERVER` entries can be added for any other master servers for this client, and for media servers for this client. Media servers for this NetBackup client can also be added by using the `MEDIA_SERVER` option.

Table 3-156 SERVER information

Usage	Description
Where to use	On NetBackup clients.

Table 3-156 SERVER information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Note: Every <code>SERVER</code> entry in a client <code>bp.conf</code> file must be a NetBackup master or media server. That is, each system that is listed as a <code>SERVER</code> must have either NetBackup master or media server software installed. The client service on some clients cannot be started if the client name is incorrectly listed as a server.</p> <p>If you configure media servers, each media server must have a <code>SERVER</code> or a <code>MEDIA_SERVER</code> entry in the <code>bp.conf</code> file of the client.</p>
Example	<p>The following is an example <code>bp.conf</code> file on a client:</p> <pre>SERVER = Master_server (default master server) SERVER = NB_server (other master server) SERVER = Media_server_#1 MEDIA_SERVER = Media_server_#2 . . .</pre> <p>The first <code>SERVER</code> entry denotes the master server to which the client connects to by default for any requests. (For example, to back up, to list, or to restore files). The <code>SERVER</code> entry must be present in the <code>/usr/opensv/netbackup/bp.conf</code> file on all UNIX clients. The <code>SERVER</code> entry is the only required entry in the <code>bp.conf</code> file for clients. The <code>SERVER</code> entry is not used in a <code>\$HOME/bp.conf</code> file. On NetBackup UNIX servers, the <code>SERVER</code> entry applies to both client and the server.</p> <p>See “RESUME_ORIG_DUP_ON_OPT_DUP_FAIL option for NetBackup servers” on page 315.</p> <p>See “MEDIA_SERVER option for NetBackup clients” on page 356.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Servers.</p> <p>See “SERVER option for NetBackup servers” on page 316.</p>

SYBASE_HOME option for NetBackup clients

The `SYBASE_HOME` entry specifies the path to the Sybase home directory. The entry is required for NetBackup to use Sybase to back up Sybase databases.

Table 3-157 SYBASE_HOME information

Usage	Description
Where to use	On a NetBackup for Sybase client.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>SYBASE_HOME = path_to_Sybase_home_directory</pre> <p>The default is that <code>SYBASE_HOME</code> is not in the configuration file.</p> <p>This option should appear only once in the configuration file.</p> <p>Note: This entry is not required to back up the Sybase SQL Anywhere database that NetBackup uses as part of the NetBackup catalog.</p> <p>For additional information, see the NetBackup for Sybase Administrator's Guide.</p>
Equivalent Administration Console property	No equivalent exists in the NetBackup Administration Console host properties.

USE_CTIME_FOR_INCREMENTALS option for NetBackup clients

The `USE_CTIME_FOR_INCREMENTALS` entry changes how NetBackup determines whether or not a file has changed. This entry causes the client software to use both modification time and inode change time during incremental backups to determine if a file has changed. (`mtime` and `ctime`.)

Table 3-158 USE_CTIME_FOR_INCREMENTALS information

Usage	Description
Where to use	On NetBackup clients.

Table 3-158 USE_CTIME_FOR_INCREMENTALS information (*continued*)

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>USE_CTIME_FOR_INCREMENTALS</pre> <p>This option should appear only once in the configuration file.</p> <p>By default, NetBackup uses only <code>mtime</code>.</p> <p>Note: If you specify <code>USE_CTIME_FOR_INCREMENTALS</code>, you must also specify <code>DO_NOT_RESET_FILE_ACCESS_TIME</code>. Do not use these options if Storage Migrator is on the system. Setting these options causes the file <code>atime</code> to be updated every time they are backed up. When the <code>atime</code> updates, it appears as if the files are frequently used and prevents Storage Migrator from selecting the files for migration.</p> <p>See “DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients” on page 349.</p>
Equivalent Administration Console property	<p>NetBackup Management > Host Properties > Double-click on client > Client Settings > Reset file access time to the value before backup.</p>

USE_FILE_CHG_LOG option for NetBackup clients

The `USE_FILE_CHG_LOG` entry specifies whether NetBackup uses the file change log on VxFS clients. The default is off.

Table 3-159 USE_FILE_CHG_LOG information

Usage	Description
Where to use	On NetBackup clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>USE_FILE_CHG_LOG</pre> <p>This option should appear only once in the configuration file.</p>

Table 3-159 USE_FILE_CHG_LOG information (*continued*)

Usage	Description
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Client Settings > Use VxFS File Change Log for Incremental Backups. See “VxFS file change log for incremental backups property” on page 102.

USEMAIL option for NetBackup clients

The `USEMAIL` entry specifies the email address where NetBackup sends status on the outcome of operations for a UNIX client.

Table 3-160 USEMAIL information

Usage	Description
Where to use	On NetBackup UNIX clients.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>USEMAIL = name1@xxx.com,name2@xxx.com,name3@xxx.com</pre> <p>Note: Separate multiple email addresses using a comma, with no spaces.</p> <p>Add <code>USEMAIL</code> as follows:</p> <ul style="list-style-type: none"> ■ If the <code>USEMAIL</code> specifies an address, NetBackup sends automatic backup and manual backup status to that address. ■ If the <code>\$HOME/bp.conf</code> file specifies an address, NetBackup also sends status on the success or failure of user operations to that address. <p>This option should appear only once in the configuration file.</p>
Equivalent Administration Console property	NetBackup Management > Host Properties > Double-click on client > Universal Settings. See “Universal Settings properties” on page 249.

VXSS_NETWORK option for NetBackup clients

This entry identifies whether a specific network or remote system must or must not use **NetBackup product authentication and authorization** with the local system. It can be configured on NetBackup servers and clients.

See [“VXSS_NETWORK option for NetBackup servers”](#) on page 331.

See [“Network Attributes tab of the Access Control properties”](#) on page 70.

UNIX client examples

The following is an example of a `bp.conf` file of a UNIX client:

```
SERVER = server1
CLIENT_NAME = client1
USEMAIL = abc@bdev.com
COMPRESS_SUFFIX = .Addr
COMPRESS_SUFFIX = .Counts
VERBOSE = 1
RESTORE_RETRIES = 1
BPBACKUP_POLICY = Uuserdir
BPBACKUP_SCHEDULE = userbackups
BPARCHIVE_POLICY = Uuserdir
BPARCHIVE_SCHEDULE = userarchives
LOCKED_FILE_ACTION = SKIP
```

Nonroot users on UNIX clients can have a personal `bp.conf` file in their home directory: `$HOME/bp.conf`.

Note: A root user cannot have a personal `bp.conf` file. For root users, NetBackup uses the `/usr/opensv/netbackup/bp.conf` file.

A personal `bp.conf` file can have any of the following options:

```
USEMAIL = abc@bdev.com
BPBACKUP_POLICY = user1
BPBACKUP_SCHEDULE = userback
BPARCHIVE_POLICY = user1
BPARCHIVE_SCHEDULE = userarch
LIST_FILES_TIMEOUT = 10
CLIENT_NAME = alternate_client_name
```

Specify `CLIENT_NAME` only when doing restores to an alternate client.

See [“Configuration options for NetBackup clients”](#) on page 333.

Configuring server groups

This chapter includes the following topics:

- [About NetBackup server groups](#)
- [Configuring a NetBackup server group](#)
- [Deleting a server group](#)

About NetBackup server groups

A server group is a group of NetBackup servers that are used for a common purpose. A NetBackup **Media Sharing** group is a server group that shares tape media for write purposes (backups).

A **Media Sharing** group can contain the following:

- NetBackup master server
- NetBackup media servers
- NDMP tape servers

Servers can be in more than one group. All members of a **Media Sharing** server group must have the same NetBackup master server.

See [“About media sharing”](#) on page 511.

See [“Configuring media sharing with a server group”](#) on page 512.

See [“Configuring a NetBackup server group”](#) on page 367.

Configuring a NetBackup server group

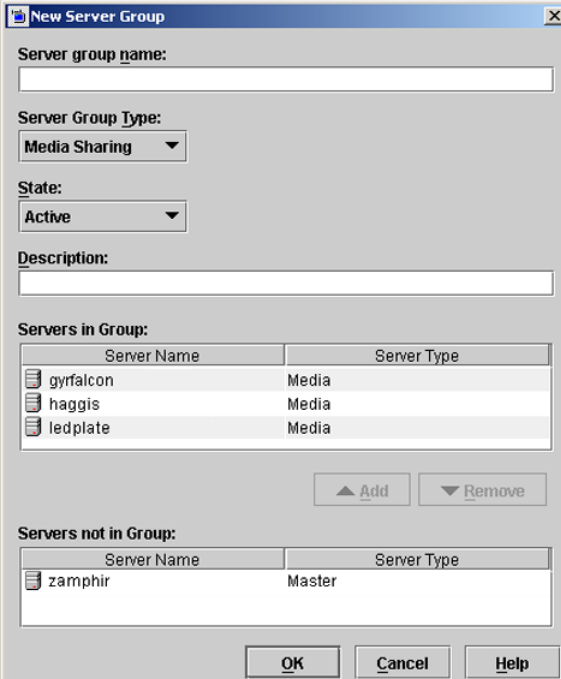
Use the following procedure to configure a server group.

Caution: NetBackup allows a server group name to be the same as the name of a media server. To avoid confusion, Symantec recommends that you do not use the same name for a server group and a media server.

See [“About NetBackup server groups”](#) on page 367.

To configure a server group

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Server Groups**.
- 2 In the **Actions** menu, select **New > Server Group**.



The "New Server Group" dialog box contains the following fields and sections:

- Server group name:** A text input field.
- Server Group Type:** A dropdown menu with "Media Sharing" selected.
- State:** A dropdown menu with "Active" selected.
- Description:** A text input field.
- Servers in Group:** A table with two columns: "Server Name" and "Server Type". It lists three servers: gyrfalcon (Media), haggis (Media), and ledplate (Media).
- Servers not in Group:** A table with two columns: "Server Name" and "Server Type". It lists one server: zamphir (Master).

Buttons at the bottom include "Add" (with an up arrow), "Remove" (with a down arrow), "OK", "Cancel", and "Help".

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

See [“NetBackup server group properties”](#) on page 369.

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**.

NetBackup 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>
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 server group.</p>
Servers in group	<p>Specifies the servers that belong to the group.</p>
Servers not in group	<p>Specifies the servers 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 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:

- Cloud storage credentials.
Configure the credentials when you configure the storage server.
See the *NetBackup Cloud Administrator's Guide*:
<http://www.symantec.com/docs/DOC5332>
- NetBackup Deduplication Engine credentials.
Create the credentials when you configure the storage server.
See the *NetBackup Deduplication Guide*:
<http://www.symantec.com/docs/DOC5332>
- NDMP hosts.
See the *NetBackup for NDMP Administrator's Guide*:
<http://www.symantec.com/docs/DOC5332>
- OpenStorage storage servers.
Configure the credentials when you configure the storage server.
See the *NetBackup OpenStorage Solutions Guide for Disk*.
See the *NetBackup Replication Director Solutions Guide*.
<http://www.symantec.com/docs/DOC5332>

Managing media servers

This chapter includes the following topics:

- [Activating or deactivating a media server](#)
- [Adding a media server](#)
- [Registering a media server](#)
- [Generating a security certificate for a NetBackup media server](#)
- [Deleting all devices from a media server](#)
- [Removing a device host from the EMM database](#)
- [About decommissioning a media server](#)
- [Previewing references to a media server](#)
- [Decommissioning a media server](#)

Activating or deactivating a media server

When you activate a media server, NetBackup can use it for backup and restore jobs. You can deactivate a media server. A common reason to do so is 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.
- 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 **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> , available at the following URL: http://www.symantec.com/docs/DOC5332
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 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 227.

Table 6-1 Adding a media server (*continued*)

Procedure	Section
Restart the NetBackup services on the master server and the media servers where a new server name was added.	See “Starting or stopping a daemon” on page 1038.
Install the NetBackup media server software.	See the <i>NetBackup Installation Guide</i> , available at the following URL: http://www.symantec.com/docs/DOC5332
On the master server, configure the robots and drives that are attached to the media server.	See “Configuring robots and tape drives by using the wizard” on page 401.
On the master server, configure the volumes.	See “About adding volumes” on page 451.
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” on page 548.
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 664.
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 842.

Registering a media server

If the master 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 master server.

To register a media server

- 1 Start the EMM service on the master server.
- 2 On the master server, run the following command. (For *hostname*, use the host name of the media server.)

On Windows:

```
install_path\Veritas\NetBackup\bin\admincmd\nbemmcmd -addhost  
-machinename hostname -machinetype media -masterserver server_name  
-operatingsystem os_type -netbackupversion  
level.major_level.minor_level
```

On UNIX:

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

```
/usr/opensv/netbackup/bin/admincmd/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*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

Generating a security certificate for a NetBackup media server

Under certain conditions, NetBackup media servers require security certificates so that NetBackup can function properly.

See “[About security certificates for NetBackup media servers](#)” on page 45.

Use the following procedure to generate a security certificate for a media server.

To generate a certificate for a media server

- 1 Run the following command on the master server, depending on your environment:

UNIX/Linux: `/usr/opensv/netbackup/bin/admincmd/bpnbaz -ProvisionCert Media_server_name`

Windows: `install_path\NetBackup\bin\admincmd\bpnbaz -ProvisionCert Media_server_name`

NetBackup appliance (as a NetBackupCLI user): `bpnbaz -ProvisionCert Media_server_name`

- 2 Restart the services on the media server.

Note: Generating a security certificate is a one-time activity.

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:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbemcmd -deletealldevices
-machinename server_name -machinetype media
```

On UNIX:

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

On Windows:

```
install_path\NetBackup\bin\admincmd\bpsturep -label  
storage_unit_label -delhost host_name
```

On UNIX:

```
/usr/opensv/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:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -changestate  
-stype server_type -dp disk_pool_name -state DOWN
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -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:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -changedp -dp  
-disk_pool_name -stype server_type -del_storage_servers  
storage_server
```

On UNIX:

```
/usr/opensv/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:

On Windows:

```
install_path\NetBackup\bin\admincmd\nbdevconfig -deletedp -dp  
disk_pool_name -stype server_type
```

On UNIX:

```
/usr/opensv/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:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpexpdate -dp  
disk_pool_name -stype server_type -nodelete
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpexpdate -dp  
disk_pool_name -stype server_type -nodelete
```

- Determine the media IDs in the disk pool as follows:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpimmedia -dp  
disk_pool_name -stype server_type -nodelete
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpimmedia -dp  
disk_pool_name -stype server_type
```

- Delete each media ID in the disk pool as follows:

On Windows:

```
install_path\NetBackup\bin\nbdelete -dt disk_type -media_id  
name
```

On UNIX:

```
/usr/opensv/netbackup/bin/nbdelete -dt disk_type -media_id  
name
```

Removing a device host from the EMM database

Use the following procedure to remove a device host from the Enterprise Media Manager (EMM) in the NetBackup database.

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.

About decommissioning a media server

You can use the NetBackup `nbdecommission` to decommission a media server. 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 retire a server that is no longer required. After all of the backup images on the old server expire, you can use the `nbdecommission` to remove the old server from the NetBackup environment.
- 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 tried 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 understand completely 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 384.

The following topics provide more information and provide guidance for the entire process.

See [“About decommissioning limitations”](#) on page 379.

See [“Before you decommission a media server”](#) on page 380.

See [“Decommission actions”](#) on page 381.

See [“Post decommission recommendations”](#) on page 380.

See [“Previewing references to a media server”](#) on page 384.

See [“Decommissioning a media server”](#) on page 385.

About decommissioning limitations

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

- Does not decommission clustered media servers. Those include NetBackup failover media servers or application clusters.
- Does not process the Media Server Deduplication Pools. Before you decommission a media server that hosts a **Media Server Deduplication Pool**, you must deactivate MSDP on that media server. See the *NetBackup Deduplication Guide*, available at the following URL: <http://www.symantec.com/docs/DOC5332>
- 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 server list of the new client.
- 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 384.
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

Symantec recommends the following actions after you run the `nbdecommission` command:

- 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	<p>Deletes the following tape storage units:</p> <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 standalone 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. <p>The BasicDisk storage must be shared, and the same disk path must be available on the new server.</p>
Storage unit - AdvancedDisk	<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.)

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

Component	Action
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.)
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 operation of the storage lifecycle policy.
Storage lifecycle policy	<p>Specifies that for each storage lifecycle policy in which one or more operations uses a storage unit that 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 was used by a Backup or Snapshot operation, 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
bp.conf file	<p>On UNIX NetBackup servers, removes the old server from the following bp.conf file entries:</p> <ul style="list-style-type: none"> ■ SERVER ■ MEDIA_SERVER ■ CLIENT_NAME ■ BROWSER <p>On UNIX master servers, also removes the old server from the FORCE_RESTORE_MEDIA_SERVER and FAILOVER_RESTORE_MEDIA_SERVERS 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"> ■ SERVER ■ MEDIA_SERVER ■ CLIENT_NAME ■ BROWSER <p>On Windows master servers, also removes the old server from the FORCE_RESTORE_MEDIA_SERVER and FAILOVER_RESTORE_MEDIA_SERVERS 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 378.

See [“Decommissioning a media server”](#) on page 385.

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:

On Windows:

```
install_path\Veritas\NetBackup\bin\admincmd\nbdecommission  
-list_ref -oldserver OldServer > file.txt
```

On UNIX:

```
/usr/openv/netbackup/bin/admincmd/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

This topic is part of a group of topics that provide information about decommissioning a NetBackup media server.

See [“About decommissioning a media server”](#) on page 378.

Your path through the `nbdecommission` 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 [“Previewing references to a media server”](#) on page 384.

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

Two procedures exist, as follows:

Replace an old media server with a new media server	See “To replace an old media server with a new media server” on page 386.
Decommission a media server	See “To decommission a media server” on page 387.

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:

On Windows:

```
install_path\Veritas\NetBackup\bin\admincmd\nbdecommission  
-oldserver OldServer [-newserver NewServer] [-file decom_ops.txt]
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/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.

On Windows:

```
install_path\Veritas\NetBackup\bin\admincmd\nbdecommission  
-oldserver OldServer
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbdecommission -oldserver  
OldServer
```

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

- 2 Follow the prompts and perform the requested actions.

Configuring storage

- [Chapter 7. Configuring disk storage](#)
- [Chapter 8. Configuring robots and tape drives](#)
- [Chapter 9. Configuring tape media](#)
- [Chapter 10. Inventorying robots](#)
- [Chapter 11. Configuring storage units](#)
- [Chapter 12. Staging backups](#)
- [Chapter 13. Configuring storage unit groups](#)

Configuring disk storage

This chapter includes the following topics:

- [About configuring BasicDisk storage](#)
- [About configuring disk pool storage](#)
- [Inventorying a NetBackup disk pool](#)
- [Changing the disk pool state](#)

About configuring BasicDisk storage

A **BasicDisk** type storage unit consists of a directory on locally-attached disk or network-attached disk. The disk storage is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory.

No special configuration is required for **BasicDisk** storage. You specify the directory for the storage when you configure the storage unit.

See [“Creating a storage unit”](#) on page 548.

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 AdvancedDisk Storage Solutions Guide.*
- *The NetBackup Cloud Administrator's Guide.*
- *The NetBackup Deduplication Guide.*
- *The NetBackup OpenStorage Solutions Guide for Disk.*

- The *NetBackup Replication Director Solutions Guide*.

The Guides are available through the following URL:

<http://www.symantec.com/docs/DOC5332>

Inventorying a NetBackup disk pool

An inventory of a NetBackup disk pool reads the capacity of the disk volumes in the pool. An inventory operation lets you update NetBackup with the new capacity values if you do the following:

- Increase or decrease the size of the disk volumes in a disk pool.
- Add volumes to or remove volumes from a disk pool.

How you increase or decrease the underlying storage capacity depends on your storage implementation. You must complete that process before you inventory the disk pool.

To inventory a NetBackup disk pool

- 1 in the **NetBackup Administration Console**, select **Media and Device Management > Devices > Disk Pools**.
- 2 On the **Actions** menu, select **Inventory Disk Pools**.
- 3 In the **Inventory Disk Pool** dialog box, select the disk pool to inventory and then click **Start Inventory**.
- 4 To update the NetBackup catalog with the capacity values returned by the inventory, click **Update Configuration**.
- 5 To inventory another disk pool, go to step 3.
- 6 To exit, click **Close**.

Changing the disk pool state

You can change the state of a NetBackup disk pool. Pool states are UP or DOWN.

To change the state to DOWN, the disk pool must not be busy. If backup jobs are assigned to the disk pool, the state change fails. Cancel the backup jobs or wait until the jobs complete.

To change the disk pool state

- 1 In the **NetBackup Administration Console**, in the left pane, select **Media and Device Management > Device Monitor**.
- 2 Select the **Disk Pools** tab.

- 3 Select the disk pool.
- 4 Select either **Actions > Up** or **Actions > Down**.

Configuring robots and tape drives

This chapter includes the following topics:

- [NetBackup robot types](#)
- [About the device mapping files](#)
- [Downloading the device mapping files](#)
- [About configuring robots and tape drives in NetBackup](#)
- [Configuring robots and tape drives by using the wizard](#)
- [Updating the device configuration by using the wizard](#)
- [Adding a robot to NetBackup manually](#)
- [Managing robots](#)
- [Adding a tape drive to NetBackup manually](#)
- [Configuring drive name rules](#)
- [Adding a tape drive path](#)
- [Adding a shared tape drive to a NetBackup environment](#)
- [Correlating tape drives and SCSI addresses on Windows hosts](#)
- [Correlating tape drives and device files on UNIX hosts](#)
- [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
- About external access to NetBackup controlled devices on UNIX
- Stopping and restarting the device manager

NetBackup robot types

A robot is a peripheral device that moves tape volumes into and out of tape drives. NetBackup uses robotic control software to communicate with the robot firmware.

NetBackup classifies robots according to one or more of the following characteristics:

- The communication method the robotic control software uses; SCSI and API are the two main methods.
- The physical characteristics of the robot. Library usually refers to a larger robot, in terms of slot capacity or number of drives. Stacker usually refers to a robot with one drive and low media capacity (6 - 12 media slots).
- The media type commonly used by that class of robots. HCART (1/2-inch cartridge tape) and 8 mm are examples of media types.

Table 8-1 lists the NetBackup robot types, with drive and slot limits for each type.

To determine which robot type applies to the model of robot that you use, see the Hardware Compatibility List for your NetBackup version:

<http://www.netbackup.com/compatibility>

Table 8-1 NetBackup robot types

Robot type	Description	Drive limits	Slot limits	Note
ACS	Automated Cartridge System	1680	No limit	API control. The ACS library software host determines the drive limit.
TL4	Tape library 4mm	2	15	SCSI control.
TL8	Tape library 8mm	No limit	16000	SCSI control.

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

Robot type	Description	Drive limits	Slot limits	Note
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.

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 following webpage:

<http://www.netbackup.com/compatibility>

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 website 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 394.

See “[About configuring robots and tape drives in NetBackup](#)” on page 395.

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 394.

To download the current device mapping files

- 1 Open the NetBackup compatibility lists knowledge base article at the following URL:

<http://www.netbackup.com/compatibility>

- 2 In the NetBackup Device Mapping Files row in the table in the **Solutions** section, select the link for your operating system.

A Symantec Knowledge Base article appears that contains installation instructions and an archive file of the device mappings.

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

About configuring robots and tape drives in NetBackup

Before you configure robots and tape drives in NetBackup, they 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. See the Guide at the following URL:

<http://www.symantec.com/docs/DOC5332>

Configure robots and tape drives in NetBackup as follows:

Device Configuration Wizard

Symantec recommends using 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)

See “[Configuring robots and tape drives by using the wizard](#)” on page 401.

The wizard discovers the devices that are attached to the media servers and helps to configure them.

See “[About configuring robots and tape drives in NetBackup](#)” on page 395.

Manually

Alternatively, add robots and drives manually as follows:

- Use menu options in the **NetBackup Administration Console**.
 See “[Adding a robot to NetBackup manually](#)” on page 402.
 See “[Adding a tape drive to NetBackup manually](#)” on page 410.
- Use NetBackup commands.
 See the *NetBackup Commands Reference Guide*, available at the following URL:
<http://www.symantec.com/docs/DOC5332>

Manual methods do not use device discovery.

To add a robot and drives, first add the robot and then add the drives that are in the robot.

Before configuring robots and drives, read the following topics to understand the process.

See “[About device serialization](#)” on page 397.

See “[About adding devices without discovery](#)” on page 398.

See “[About robot control](#)” on page 398.

See “[About drive name rules](#)” on page 400.

See “[Correlating tape drives and device files on UNIX hosts](#)” on page 421.

See “[Correlating tape drives and SCSI addresses on Windows hosts](#)” on page 419.

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. See the *NetBackup Device Configuration Guide*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

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 435.

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 419.

See [“Correlating tape drives and device files on UNIX hosts”](#) on page 421.

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 405.

The following table 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 8-2 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
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 8-1 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 8-1 Robot control host example



TLD robot (HP EML E-Series)

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 416.

Configuring robots and tape drives by using the wizard

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

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.

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.

Adding a robot to NetBackup manually

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

See [“NetBackup robot types”](#) on page 393.

See [“About robot control”](#) on page 398.

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

See [“Adding a tape drive to NetBackup manually”](#) on page 410.

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 > 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 404.

- 4 After you specify properties, click **OK**.

After you click **OK**, the **Stop/Restart Media Manager Device Daemon** dialog box appears.

- 5 If you intend to make other changes, click **Cancel** in the **Stop/Restart Media Manager Device Daemon** dialog box. You can restart the Device Manager or the device daemon after you make the final change.

If the device changes are complete, restart the device daemon by clicking **OK** in the **Stop/Restart Media Manager Device Daemon** dialog box.

If you restart the device daemon, 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 compatibility list website:

<http://www.netbackup.com/compatibility>

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.

- 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 in the *NetBackup Device Configuration Guide*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

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 8-3 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).

Table 8-3 Robot configuration properties (*continued*)

Property	Description
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.</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 in the <i>NetBackup Device Configuration Guide</i>, available at the following URL: http://www.symantec.com/docs/DOC5332 ■ 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.</p> <p>For compatibility information, see the Hardware Compatibility List: http://www.netbackup.com/compatibility</p> <p>An overview of ACS robots is available in the <i>NetBackup Device Configuration Guide</i>, available at the following URL: http://www.symantec.com/docs/DOC5332</p>
Library name	<p>The following applies only to a TLH robot:</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\ibmatl.conf</code> file. For example, in the following example entry in that file, 3494AH is the library name: <code>3494AH 176.123.154.141 ibmpcl</code> ■ Enter the library name. <p>An overview of TLH robots is available in the <i>NetBackup Device Configuration Guide</i>. See the <i>NetBackup Device Configuration Guide</i>, available at the following URL: http://www.symantec.com/docs/DOC5332</p>
LMCP device file	<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>

Table 8-3 Robot configuration properties (*continued*)

Property	Description
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.</p> <p>If the browse operation fails for any other reason, a dialog box appears that lets you enter the information.</p> <p>Use the Windows management tools to find the Port, Bus, Target, and LUN numbers.</p> <p>If the browse operation does not find attached robots, an error dialog box appears.</p>
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 in the <i>NetBackup Device Configuration Guide</i>. See the <i>NetBackup Device Configuration Guide</i>, available at the following URL: http://www.symantec.com/docs/DOC5332</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 hosts 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>

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 404.
- 5 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box or the **Media and Device Management** dialog box to restart the Device Manager or the device daemon.

If you intend to make other changes, click **No**; you can restart the Device Manager or the device daemon after you make the final change.

If you restart the Device Manager or the device daemon, 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 1052.

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 on the host to which the robot is attached:

On Windows:

```
install_path\Volmgr\bin\tpconfig -update -robot robot_number -pird  
yes
```

On UNIX:

```
/usr/opensv/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 492.

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 375.

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 8-4 Move a robot and media to a new server overview

Task	Procedure
Determine which tapes on the <i>old_server</i> contain NetBackup images that have not expired.	Run the following <code>bpmedialist</code> command: <code>bpmedialist -mlist -l -h old_server</code> The <code>-l</code> option produces one line of output per tape.
Move the tapes in the robot that is attached to the <i>old_server</i> to non-robotic status (standalone).	See “Moving volumes by using the Actions menu” on page 504.

Table 8-4 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>bpmedia -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>bpmedia -movedb -ev media_ID -oldserver old_server -newserver new_server</pre> <p>For the media that has active images, see the <code>bpmedialist</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 149.
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.
Create the appropriate NetBackup storage units.	See “Creating a storage unit” on page 548.
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 NetBackup volume configuration with a robot’s contents” on page 527.

Adding a tape drive to NetBackup manually

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 > 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 412.
 - Select **Use drive name rules**. This option uses rules to name the drive automatically.
See [“About drive name rules”](#) on page 400.
See [“Configuring drive name rules”](#) on page 416.
- 4 To configure the host and the 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 412.

- 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 413.
- 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 or the **Media and Device Management** dialog box to restart the Device Manager or the device daemon.

If you intend to make other changes, click **No**; you can restart the Device Manager or the device daemon after you make the final change.

If you restart the Device Manager or the device daemon, any backups, archives, or restores that are in progress also may be stopped.

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 rules to create a unique drive name.

Use drive name rules (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 400.

See [“Configuring drive name rules”](#) on page 416.

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 417.

See [“Drive path options”](#) on page 418.

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.

The following table describes the tape drive configuration options.

Table 8-5 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 standalone 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 8-5 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 the 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, Volume II</i>: http://www.symantec.com/docs/DOC5332</p>
Drive Status	<p>On UNIX only.</p> <p>Specifies the availability of the drive.</p> <p>When you add a drive, the default drive status is UP, which means the drive is available. When a drive is UP, the default mode is AVR (Automatic Volume Recognition).</p> <p>To change the drive status, click UP or DOWN.</p> <p>You can also change the drive status by using the commands on the Actions menu in Device Monitor.</p>
Serial Number	A read-only field that shows the serial number of the drive.
Robotic library	Specifies a robot that controls the drive. You can select any configured robot that can control the drive.

Table 8-5 Tape drive configuration options (*continued*)

Option	Description
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, as follows:</p> <ul style="list-style-type: none"> ■ On Windows: You must determine which physical drive in the robot is identified by the logical device name. See “Correlating tape drives and SCSI addresses on Windows hosts” on page 419. ■ On UNIX: You must determine which physical drive in the robot is identified by the device file name. See “Correlating tape drives and device files on UNIX hosts” on page 421. <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>
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>

Configuring drive name rules

Use the following procedure to configure the rules that NetBackup uses to name tape drives. The procedure describes how to configure the rules in the **NetBackup Administration Console**.

Alternatively, if you use the **NetBackup Device Configuration Wizard**, click **Configure Drive Name Rules** in the **Device Hosts** screen. The same **Configure Drive Name Rules** dialog box that is described in the following procedure.

To configure drive name rules

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Drives**.

See [“Adding a tape drive to NetBackup manually”](#) on page 410.

- 2 On the **Actions** menu, select **New > Tape Drive**.
- 3 In the **New Drive** dialog box, click **Configure**.

Configure Drive Name Rules

Use this dialog to automatically create drive names based on rules you specify. The default drive name rule creates names in the format VendorID.ProductID.INDEX.

Host selection
Select the hosts for which to configure the rule.

☐ Global Rule
☐ cave1.abc-domain.com
☐ cave2.abc-domain.com

Fields and order
Select the fields that will be part of the drive name.

host name
drive type
drive position
serial number
vendor ID
robot number
product ID
index
robot type

Choose the order in which the fields appear.

Custom text:

Configured drive name rules

Host Name	Rule
GLOBAL	<vendor ID> . <product ID> . <index>

OK Cancel Help

- 4 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.

See “[About drive name rules](#)” on page 400.

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**.
- 2 Select **Add a New Drive**. In the dialog box, click **Add**.
- 3 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 417.

See “[Drive path options](#)” on page 418.

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 180.

More information about how NetBackup reserves drives is available in the *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

Drive path options

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

Table 8-6 Add drive path options

Option	Description
Host name (Windows)	Specifies the device host for the drive.
Device host (UNIX)	
Enable host path	Specifies that the path is active and that NetBackup can use it for backups and restores.
NDMP host	<p>Specifies the NDMP host for the device (if an NDMP host is configured in your NetBackup environment).</p> <p>See the <i>NetBackup for NDMP Administrator's Guide</i> for additional information about NDMP drives:</p> <p>http://www.symantec.com/docs/DOC5332</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 media server default 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 the SCSI devices that conform to the reserve and the release management method. That method 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 the 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 180.</p>
Path	<p>On UNIX hosts.</p> <p>Specifies the path of the character-mode, no rewind device file on the specified host.</p> <p>See “About no rewind device files on UNIX” on page 419.</p> <p>Device files are in the <code>/dev</code> directory on the UNIX host. If the entries do not exist, see the <i>NetBackup Device Configuration Guide</i> for information about how to create them:</p> <p>http://www.symantec.com/docs/DOC5332</p>
Port, Bus, Target, and LUN	<p>On Windows hosts.</p> <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>

Table 8-6 Add drive path options (*continued*)

Option	Description
This path is for a Network Attached storage device	Specifies that the path is for a Network Attached Storage (NAS) device.

About no rewind device files on UNIX

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, you must create them. See the *NetBackup Device Configuration Guide*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

Adding a shared tape drive to a NetBackup environment

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.

For more information about the Shared Storage Option, see the *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

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 *NetBackup Troubleshooting Guide*:

<http://www.symantec.com/docs/DOC5332>

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.

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

For Windows hosts only.

Assume that 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 that the tape is 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.

If the data path to the drive in which the tape is mounted is not on the robot control host, you may have to unload the drive. To do so, use a command on another host or use the drive's front panel.

Correlating tape drives and device files on UNIX hosts

If your tape drives do not support device serialization, you may have to determine which device file or SCSI address matches the physical drive. You also may have to do so if you add the tape drives manually.

Correlate device files to physical drives when you create the device files for each drive.

To correlate tape drives and device files on UNIX

- 1 Determine the physical location of each drive within the robotic library. The location usually is shown on the connectors to the drives or in the vendor's documentation.
- 2 Physically connect the drives to SCSI adapters in the host.
- 3 Record the adapter and SCSI addresses to which you connected each drive.

- 4 Create device files for each drive by using the SCSI addresses of the drives and adapters.

Add the device file by using the notes from a previous step to complete the correlation between device files and physical drive location.

- 5 Configure the robot in NetBackup and then add the drives.

When you add the drives, verify that you assign the correct drive address (for example, robot drive number) to each device path.

Optionally, use the appropriate NetBackup robotic test utility to verify the configuration.

Information about the robotic test utilities is available.

See *NetBackup Troubleshooting Guide*:

<http://www.symantec.com/docs/DOC5332>

To verify the device correlation on UNIX

- 1 Stop the NetBackup device daemon (`ltid`).
- 2 Start `ltid`, which starts the Automatic Volume Recognition daemon (`avrd`). Stop and restart `ltid` to ensure that the current device configuration is activated.

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 Administration Console Device Monitor** to verify that the tape was mounted on the correct robot drive.

UNIX device correlation example

On UNIX only.

Assume a TLD robot includes three drives and the operating system includes the following device paths:

Drive 1 `/dev/rmt/0cbn`

Drive 2 `/dev/rmt/1cbn`

Drive 3 `/dev/rmt/3cbn`

Also assume that you requested that the tape be mounted on drive 1.

If the device path for the drive is configured correctly, the **NetBackup 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 device path for that drive is not configured correctly. For example, if the Device Monitor shows that the tape is mounted on Drive 2, the device path for drive 1 is incorrect. Replace the drive 1 device path (`/dev/rmt/0cbrn`) with the correct device path (`/dev/rmt/1cbrn`) for drive 2. You may need to use a temporary device path while you make these changes. You also know that the device path for drive 2 is incorrect. Possibly, the device paths were swapped during configuration.

Use the robotic test utility to unload and unmount the tape from drive 1. Repeat the test for each drive.

If the path to the drive where the tape is 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.

Managing tape drives

You can perform various tasks to manage tape drives.

Changing a drive comment

You can change the comment that is 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 you licensed a disk pool feature, 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 1093.
- 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 `nbemmcmd` command with the `-drive_error_threshold` and `-time_window` options to change the default values.

Additional information about `nbemmcmd` is available.

See the *NetBackup Commands Reference Guide*:

<http://www.symantec.com/docs/DOC5332>

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 424.

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 you licensed a disk pool feature, 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 425.

To change a drive path

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices > Drives**. Double-click the drive that you want to change.
- 2 In the **Change Tape Drive** dialog box, select the drive path in the **Host and Path information** list. Click **Change**.
- 3 In the **Change Path** dialog box, configure the properties for the drive path.
The properties you can change depend on drive type, server platform, or NetBackup server type.
See [“About SCSI reserve on drive paths”](#) on page 417.
See [“Drive path options”](#) on page 418.
- 4 Click **OK** to save the changes.

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 you licensed a disk pool feature, 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** or the **Change 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 412.

- 5 If the device changes are complete, select **Yes** on the **Restart Device Manager** dialog box or the **Media and Device Management** dialog box to restart the Device Manager or the device daemon.

If you intend to make other changes, click **No**; you can restart the Device Manager or the device daemon after you make the final change.

If you restart the Device Manager or the device daemon, 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 daemon.

- 6 After you change the properties, click **OK**.

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 can configure the automatic, frequency-based cleaning interval.

You can also 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 of a drive”](#) on page 429.

See the *NetBackup Administrator's Guide, Volume I*:

<http://www.symantec.com/docs/DOC5332>

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

See [“Cleaning tape drives from the Activity Monitor”](#) on page 1045.

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** 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 standalone 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 375.

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 you licensed a disk pool feature, 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 of a drive

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 you licensed a disk pool feature, 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 the 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 you licensed a disk pool feature, 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 you licensed a disk pool feature, select the **Drives** tab.
- 3 In the **Drive Status** pane, select a drive.
- 4 Select **Actions > Drive Details**.
- 5 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.

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 standalone drive, insert the prelabeled test tape that is shown in the **Step Information** column of the **Results** window.

The **Results** window shows results of each step in the test.

- 6 If operator intervention is required, the State column of the Results window displays Waiting. For example, a test step may require that you to load a new tape into a drive before the test can continue.

Complete the intervention and then click **Continue**.

Select the test step in the **Results** window and click **Details** to determine what you must do. Complete the requested operation task and then click **Continue** in the **Test Details** dialog box to resume the test

To stop a test and change the device

- 1 Click **Stop**.

The test ends after it performs any necessary clean-up work and updates the test records to reflect that the test run has been stopped.

- 2 In the **Device Host** and the **Drive** boxes, select the host and the drive that you want to test.
- 3 Click **Start** to restart the diagnostic test.

Managing a diagnostic test step that requires operator intervention

Operator intervention is required if the **Status** column of the **Results** display contains **Waiting**. For example, a test step may prompt for a new tape to be loaded into a drive before the test continues.

To manage a diagnostic step

- 1 Complete the requested operations task.
- 2 Click **Continue** to resume the test.

If you clicked **Details** for a test step that requires operator intervention, you can click **Continue** from the **Test Details** dialog box.

Obtaining detailed information for a diagnostic test step

You can get information for a test step at any time during the test.

To obtain detailed information for a diagnostic test step

- 1 Select a test step in the **Results** display.
- 2 Click **Details**. A dialog box appears that displays information for the step.

The information includes a brief explanation of the checks that are performed by a specific step and the instructions that are associated with any step that requires manual intervention. For example, a step may prompt for a new tape to be loaded into a tape drive before the diagnostic session continues.
- 3 Click **Close** to return to the **Device Diagnostics** dialog box.

Verifying the device configuration

Verify the device configuration by running the Device Configuration Wizard. However, some details of a device configuration cannot be validated without attempting tape mounts. Use the NetBackup `robtest` utility to mount tapes and validate the configuration.

To verify robots and drives by using the wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

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 NetBackup Device Manager (`ltid`) performs automatic path correction. See [“About the NetBackup Device Manager”](#) on page 438.
- 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 435.

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**.

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*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

- 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 434.

To configure automatic path correction

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

On Windows:

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

On UNIX:

```
/usr/opensv/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

Two processes exist for replacing a device, as follows:

Replace a device on a single host See [Table 8-7](#) on page 436.

Replace a shared device See [Table 8-8](#) on page 436.

Table 8-7 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 424.
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 424.
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 401.

Table 8-8 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 424.
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> <p>On Windows:</p> <pre>install_path\Veritas\Volmgr\bin\tpautoconf -report_disc</pre> <p>On UNIX:</p> <pre>/usr/opensv/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 daemon” on page 1038.

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

Task	Instructions
Read the serial number from the new device and update the EMM database.	<p>If the device is a robot, run the following command:</p> <p>On Windows:</p> <pre>install_path\Veritas\Volmgr\bin\tpautoconf -replace_robot robot_number -path robot_path</pre> <p>On UNIX:</p> <pre>/usr/opensv/volmgr/bin/tpautoconf -replace_robot robot_number -path robot_path</pre> <p>If the device is a drive, run the following commands:</p> <p>On Windows:</p> <pre>install_path\Veritas\Volmgr\bin\tpautoconf -replace_drive drive_name -path path_name</pre> <p>On UNIX:</p> <pre>/usr/opensv/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 401.
If the device is a drive, change the drive state to UP.	See “Changing a drive operating mode” on page 424.

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 8-9 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 424.
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 424.

About the NetBackup Device Manager

The NetBackup Device Manager (`ltid`) manages robot and tape processes for NetBackup. The Device Manager processes requests to mount and unmount tapes in robotically controlled devices through the robotic control processes. NetBackup starts `ltid` on the hosts that have storage devices configured. The Device Manager starts the Volume Manager (`vmd`), the automatic volume recognition process (`avrd`), and any robotic processes as needed.

In the **NetBackup Administration Console**, the NetBackup Device Manager is exposed as follows:

In the Activity Monitor	For Windows hosts, as the NetBackup Device Manager . For UNIX hosts, as <code>ltid</code> .
On the Actions menu	As the Media Manager Device Daemon .

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 439.

About external access to NetBackup controlled devices on UNIX

On UNIX hosts, the NetBackup Device Manager restricts access to drives that are in an `UP` state by changing the permissions of the device files for those drives. The Device Manager changes the permissions to 0600 when it starts and back to their

original settings when it is terminated. The permissions also are returned to their original settings when a drive's state is changed to `DOWN`.

See “About the NetBackup Device Manager” on page 438.

Do not modify the permissions of these device files when the Device Manager is active. The automatic volume recognition process (`avrd`) periodically tries to rewind and read data from media in the drives that are `UP` and are not currently assigned in NetBackup.

To ensure reliable operation, do not use UNIX tape and drive commands on the drives that are `UP` and controlled by the Device Manager. Users can use the NetBackup `tpreq` and `tpunmount` commands and the `drive_mount_notify` and `drive_unmount_notify` scripts on those drives.

See “NetBackup notify scripts” in the *NetBackup Administrator's Guide, Volume II* <http://www.symantec.com/docs/DOC5332>

Stopping and restarting the device manager

Use the following procedure to stop and restart the NetBackup Device Manager.

See “About the NetBackup Device Manager” on page 438.

To start or stop the Device Manager

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Devices**.
- 2 On the **Actions** menu, select **Stop/Restart Media Manager Device Daemon**.
- 3 In the **Stop/Restart Media Manager Device Daemon** dialog box, do the following:
 - 1 In the **Device host** dropdown list, select the wanted media server.
 - 2 Select the **Action: Start, Stop, or Stop/Restart**.
The actions that are available depend on the state of the device manager / daemon.
 - 3 Select the wanted **Options: Eject media from standalone drive(s) or Enable verbose logging**.
 - 4 Click **Apply** or **OK**, depending on the following results:
 - **Apply** does not close the dialog box so you can select device hosts and actions for more than another device host.
 - **OK** closes the dialog box.

Configuring tape media

This chapter includes the following topics:

- [About NetBackup tape volumes](#)
- [About NetBackup volume pools](#)
- [About NetBackup volume groups](#)
- [NetBackup media types](#)
- [About WORM media](#)
- [About adding volumes](#)
- [About configuring media name and attribute rules](#)
- [Adding volumes by using the wizard](#)
- [About media settings rules](#)
- [Configuring media settings](#)
- [About barcodes](#)
- [Configuring barcode rules](#)
- [About media ID generation rules](#)
- [Configuring media ID generation rules](#)
- [About media type mapping rules](#)
- [Adding volumes by using the Actions menu](#)
- [Configuring media type mappings](#)
- [Managing volumes](#)

- [Managing volume pools](#)
- [Managing volume groups on UNIX](#)
- [About media sharing](#)
- [Configuring unrestricted media sharing](#)
- [Configuring media sharing with a server group](#)

About NetBackup 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.

NetBackup uses two volume types, as follows:

Robotic volumes	Volumes that are located in a robot. The robotic library moves the volumes into and out from the robotic drives as necessary.
Standalone volumes	Volumes that are allocated for the drives that are not in a robot. Operator intervention is required to load volumes into and eject volumes from standalone drives.

NetBackup uses volume pools to organized volumes by usage.

See [“About NetBackup volume pools”](#) on page 441.

You can use WORM media with NetBackup.

See [“About WORM media”](#) on page 447.

Volume information is stored in the EMM database.

See [“About the Enterprise Media Manager \(EMM\)”](#) on page 873.

About NetBackup 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	<p>For NetBackup catalog backups.</p> <p>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.</p> <p>See “About the NetBackup catalog” on page 865.</p>
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.

See [“About scratch volume pools”](#) on page 443.

The volume pool concept is relevant only for tape storage units and does not apply to disk storage units.

You can use any of the approved characters for volume pools names.

See [“NetBackup naming conventions”](#) on page 1093.

NetBackup uses several special prefixes for volume pool names.

See [“About reserved volume pool name prefixes”](#) on page 442.

Examples of volume pool usage are available in the *NetBackup Administrator’s Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

About reserved volume pool name prefixes

NetBackup reserves the following prefixes (case sensitive) for the names of the volume pools that contain media for specific purposes:

ENCR	<p>For volumes on which NetBackup encrypts the data. The volumes in a volume pool that uses this name prefix must be targeted to encrypting tape drives.</p> <p>See “About NetBackup encryption options” on page 705.</p> <p>See the NetBackup Security and Encryption Guide.</p>
WENCR	<p>For WORM volumes on which NetBackup encrypts the data. The volumes in a volume pool that uses this name prefix must be targeted to encrypting tape drives.</p> <p>See “About NetBackup encryption options” on page 705.</p> <p>See the NetBackup Security and Encryption Guide.</p> <p>See the <i>NetBackup Commands Reference Guide</i>, available at the following URL:</p> <p>http://www.symantec.com/docs/DOC5332</p>
WORM	<p>For WORM volumes. NetBackup does not encrypt the data.</p> <p>See “About using volume pools to manage WORM media” on page 448.</p>

NetBackup examines the volume pool names to determine if they are special purpose volume pools. If a volume pool name begins with one of the special prefixes, NetBackup processes the media in that pool according to the requirements for that pool. If not, NetBackup does not use special processing for that volume pool's media.

When you create a volume pool for any of these purposes, you must use upper case characters. For readability, it may be beneficial to use an underscore character after the prefix, for example `WORM_` or `ENCR_`.

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.

See [“About NetBackup volume pools”](#) on page 441.

See [“Configuring media settings”](#) on page 454.

See [“About media settings rules”](#) on page 454.

About NetBackup 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*:

<http://www.symantec.com/docs/DOC5332>

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.

The following table describes the NetBackup media types.

Table 9-1 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
8MM3	8MM cartridge tape 3
8MM3_CLN	8MM cleaning tape 3
DLT	DLT cartridge tape
DLT_CLN	DLT cleaning tape

Table 9-1 NetBackup media types (*continued*)

Media type	Description
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.

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.

See "Media formats" in the *NetBackup Administrator's Guide, Volume II*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

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](#), available at the following URL:

<http://www.symantec.com/docs/DOC5332>

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.

NetBackup provides two methods to manage WORM media, as follows:

- Assign a reserved prefix to WORM volume pool names.
See [“About using volume pools to manage WORM media”](#) on page 448.
- Assign a specific drive type to all WORM drives and a specific media type to all WORM media.
See [“About using unique drive and media types to manage WORM media”](#) on page 449.

About using volume pools to manage WORM media

You can dedicate volume pools for WORM media. This method lets a WORM-capable tape drive back up and restore standard and WORM media. NetBackup uses two reserved volume pool prefixes to indicate that the volumes in a pool are for WORM drives, as follows:

- WORM (uppercase letters) denotes WORM media.
- WENCR (uppercase letters) denotes WORM media on which NetBackup should encrypt the data.

See [“About reserved volume pool name prefixes”](#) on page 442.

For more information about encrypting data on your media, see the *NetBackup Security and Encryption Guide*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

When you create a volume pool for WORM media, specify one of the reserved prefixes as the first characters of the pool name. NetBackup examines the volume pool names to determine if they begin with a reserved prefix. For readability, it may be beneficial to use an underscore character after the prefix, for example **WORM_**.

See [“Adding a volume pool”](#) on page 507.

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 the Quantum media that has never been used or all of its NetBackup images have expired, NetBackup uses the media.

See [“About using a WORM scratch pool”](#) on page 449.

See [“About WORM media”](#) on page 447.

See [“About using unique drive and media types to manage WORM media”](#) on page 449.

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 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.

Because Quantum drives use only a single media type, this method for managing the WORM media is unnecessary.

See [“About WORM media and the Quantum drive”](#) on page 450.

If you use unique drive and media types to manage WORM media, disable the WORM volume pool name verification.

See [“Disabling WORM volume pool name verification”](#) on page 450.

Disabling WORM volume pool name verification

If you use unique drive and media types to manage WORM media, disable NetBackup volume pool name verification. WORM volume pool name verification is used only for the WORM volume pool method of managing WORM media.

See [“About using unique drive and media types to manage WORM media”](#) on page 449.

See [“About using volume pools to manage WORM media”](#) on page 448.

To disable the volume pool name verification

- ◆ Create the following touch file on the media server of the WORM drive:

On Windows:

```
install_path\netbackup\db\config\DISABLE_WORM_POOLCHECK
```

On UNIX:

```
/usr/opensv/netbackup/db/config/DISABLE_WORM_POOLCHECK
```

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.

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*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

For information about the drives that NetBackup supports for WORM media, see the Hardware Compatibility List at the following URL

<http://www.netbackup.com/compatibility>

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*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

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 standalone.

Volumes in a robotic library See “[About adding robotic volumes](#)” on page 451.

Standalone volumes See “[About adding standalone volumes](#)” on page 452.

NetBackup uses the rules to assign names and attributes to volumes.

See “[About configuring media name and attribute rules](#)” on page 453.

About adding robotic volumes

The robotic volumes are the volumes that are located in a robotic tape library. The following table describes the methods for adding robotic volumes.

Table 9-2 Methods for adding robotic volumes

Method	Description
The Volume Configuration Wizard	See “ Adding volumes by using the wizard ” on page 453.

Table 9-2 Methods for adding robotic volumes (*continued*)

Method	Description
Robot inventory	See “ About robot inventory ” on page 514. See “ Updating the NetBackup volume configuration with a robot's contents ” on page 527.
The Actions menu	See “ Adding volumes by using the Actions menu ” on page 477.
NetBackup commands	See the <i>NetBackup Commands Reference Guide</i> , available at the following URL: http://www.symantec.com/docs/DOC5332

About adding standalone volumes

Standalone volumes are the volumes that reside in the drives that are not in a robot or are allocated for standalone 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 standalone drive is full or unusable because of errors, NetBackup ejects (logically) the volume. If you add other standalone volumes, NetBackup requests that volume; NetBackup does not generate an `out of media` error.

The easiest way to add standalone volumes is to use the Volume Configuration Wizard. Then, when NetBackup requests one of the volumes, insert it into the standalone drive and NetBackup labels it.

The `DISABLE_STANDALONE_DRIVE_EXTENSIONS` option of the `nbemmcmd` command can turn off the automatic use of standalone volumes.

Table 9-3 Methods for adding standalone volumes

Method	Description
The Volume Configuration Wizard	See “ Adding volumes by using the wizard ” on page 453.
The Actions menu	See “ Adding volumes by using the Actions menu ” on page 477.
NetBackup commands	See the <i>NetBackup Commands Reference Guide</i> , available at the following URL: http://www.symantec.com/docs/DOC5332

About configuring media name and attribute rules

NetBackup uses the default settings and rules to name and assign attributes to new removeable media. NetBackup uses these rules when you do the following:

- Use the **Volume Configuration Wizard** to add new media.
- Use the **Robot Inventory** dialog box to inventory a robot. If NetBackup discovers new media in the robot, it adds that media to NetBackup.

For most configurations, the default settings work well. However, you can change the default settings and rules that NetBackup uses. Change the settings only if you have special hardware or usage requirements. You can change the settings from the **Volume Configuration Wizard** or from the **Robot Inventory** dialog box.

The following table shows the rules that you can configure:

Table 9-4 Attributes for media

What	Where
Media settings	See “About media settings rules” on page 454. See “Configuring media settings” on page 454.
Barcode rules	See “About barcodes” on page 464. See “Configuring barcode rules” on page 467.
Media ID generation rules	See “About media ID generation rules” on page 472. See “Configuring media ID generation rules” on page 473.
Map media for API robots	See “About media type mapping rules” on page 476. See “Configuring media type mappings” on page 481.

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.

You can change the default settings and rules that NetBackup uses to name and assign attributes to new removable media.

See [“About configuring media name and attribute rules”](#) on page 453.

About media settings rules

The NetBackup media settings rules depend on the following:

- For existing media, the volume group to which the volumes belong.
- For new media, the media ID prefix, the media type, and the pool to which the volume should be assigned.

You can change the default rules.

See [“Configuring media settings”](#) on page 454.

See [“Media settings options”](#) on page 456.

Configuring media settings

Use the **Media Settings** tab of the NetBackup **Advanced Robot Inventory Options** dialog box to configure the attributes for existing and new media.

See [“About media settings rules”](#) on page 454.

To configure media settings

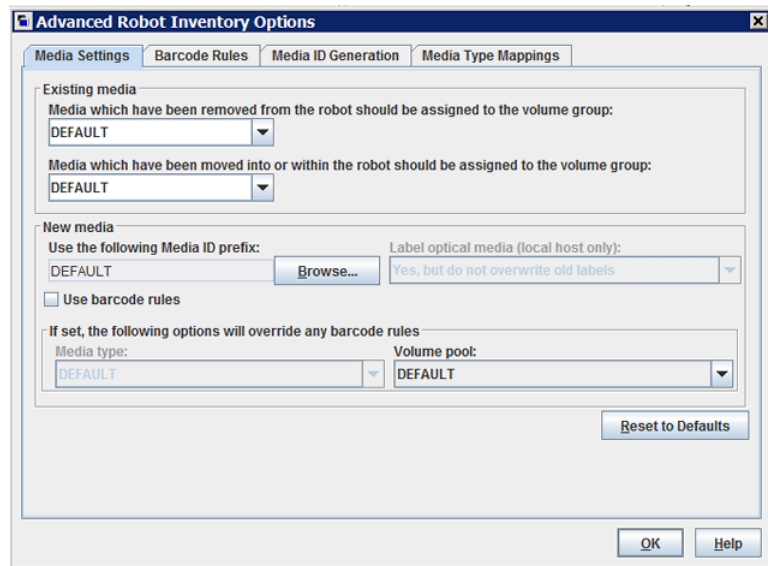
- 1 Open the **Advanced Robot Inventory Options** dialog box, as follows:

From the **Robot Inventory** dialog box

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.
- 2 Select the robot that you want to inventory.
- 3 On the **Actions** menu, select **Inventory Robot**.
- 4 Click either **Preview volume configuration changes** or **Update volume configuration**.
- 5 Click **Advanced Options**.

- 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.
- 3 On the **Robot Inventory** panel of the **Volume Configuration Wizard**, click **Advanced Options**.

- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media Settings** tab.



3 Configure the settings, as follows:

- a. In the **Media which have been removed from the robot should be assigned to the volume group** list, select a volume group for the media that are removed from the robot.

See [“Media which have been removed from the robot... \(existing media setting\)”](#) on page 457.

- b. In the **Media which have been moved into or within the robot should be assigned to the volume group** list, select a volume group for the media that are in or are added to the robot.

See [“Media which have been moved into or within the robot... \(existing media setting\)”](#) on page 457.

- c. If the robotic library supports barcodes and the volume has readable barcodes, NetBackup creates media IDs automatically from the barcodes. You do not need to configure a prefix.

However, if the media in the robotic library has unreadable barcodes or if the robot does not support barcodes, NetBackup assigns a default media ID prefix.

To use a media ID prefix other than the **DEFAULT**, click **Browse** in the **Use the following Media ID prefix** field. Then, specify or choose a media ID prefix in the **Media ID Prefix** dialog box.

See [“Use the following Media ID prefix \(new media setting\)”](#) on page 458.

- d. To use your barcode rules to assign attributes to new volumes, select **Use barcode rules**.

See [“Use barcode rules \(new media setting\)”](#) on page 460.

- e. To override your barcode rules for the new media in the robotic library, select a **Media type** from the list.

See [“Media type \(new media setting\)”](#) on page 460.

- f. To override the default volume pool for the new media in the robotic library, select a **Volume pool** from the list.

See [“Volume pool \(new media setting\)”](#) on page 463.

4 Click **OK**.

Media settings options

The following are the settings for the new media in a robot that you add to your NetBackup volume configuration.

Media which have been removed from the robot... (existing media setting)

For the media that already exist in your volume configuration, you can specify the volume group if the media are removed from the robot. **Media which have been removed from the robot should be assigned to the volume group.**

The **Media which have been removed from the robot should be assigned to the volume group** drop-down box 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 of the **New media** section of the dialog box, as follows:

- If the **Media type** field is **DEFAULT**

The **Media which have been removed from the robot should be assigned to the volume group** dropdown box includes the volume groups that are valid for the robot's default media type.
- If the **Media type** field is other than **DEFAULT**

The **Media which have been removed from the robot should be assigned to the volume group** dropdown box includes the volume groups that are valid for the specified media type.

To specify a volume group other than **DEFAULT**, enter a volume group name or select one from the list.

See [“Media type \(new media setting\)”](#) on page 460.

Media which have been moved into or within the robot... (existing media setting)

You can specify the volume group for the existing media that have been moved into or within a 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 **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box 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.

The following other selections may be available depending on the setting of the **Media type** field of the **New media** section of the dialog box:

If the **Media type** field is **DEFAULT** The **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box includes the volume groups that are valid for the robot's default media type.

If the **Media type** field is other than **DEFAULT** The **Media which have been moved into or within the robot should be assigned to the volume group** drop-down box includes the volume groups that are valid for the specified media type.

To specify a volume group other than **DEFAULT**, enter a volume group name or select one from the list.

If the robotic library contains multiple media types, Symantec recommends a **DEFAULT** setting. If you specify a volume group and volumes of different media types were moved into or within the robot, the new update fails. Volumes of different media types cannot have the same volume group.

See [“Media type \(new media setting\)”](#) on page 460.

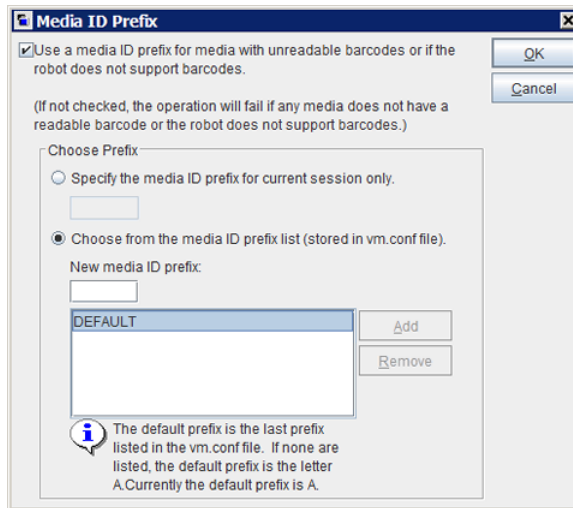
Use the following Media ID prefix (new media setting)

if the media has unreadable barcodes or if the robot does not support barcodes, by default NetBackup assigns media ID prefixes as follows:

- If **MEDIA_ID_PREFIX** entries are defined in the `vm.conf` file, NetBackup assigns the last **MEDIA_ID_PREFIX** entry as the media ID prefix.
- If no **MEDIA_ID_PREFIX** entries are defined in the `vm.conf` file, NetBackup uses the letter A as the media ID prefix.

To configure NetBackup to use a media ID prefix other than the default, select **Use the following Media ID prefix** field in the **Advanced Robot Inventory Options** dialog box and then click **Browse**. The **Media ID Prefix** dialog box appears.

Figure 9-1 Media ID Prefix dialog box



The following are the alternative NetBackup media ID assignment behaviors that you can configure in the dialog box:

- | | |
|---|--|
| To not use a media ID prefix | Deselect the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option. |
| To use a media ID prefix | Select the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option. |
| To use a specific media ID prefix for the current operation only | Select the Specify the media ID prefix for the current session only option then enter the media ID prefix. You can specify a prefix of one to five alphanumeric characters. NetBackup assigns the remaining numeric characters to create a six character media ID.

NetBackup uses the prefix only for the current operation. |
| To configure a media ID prefix to use for the current session and future sessions | Select the Choose from the Media ID prefix list (stored in vm.conf file) option and then select the prefix from the list. |
| To add a new media ID prefix to the <code>vm.conf</code> file | Select the Choose from the Media ID prefix list (stored in vm.conf file) option and then enter the prefix in the New media ID prefix field. Click Add . |

To remove a media ID prefix from the `vm.conf` file Select the **Choose from the Media ID prefix list (stored in `vm.conf` file)** option, select the prefix from the list, and then click **Remove**.

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

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

<http://www.symantec.com/docs/DOC5332>

Use barcode rules (new media setting)

Specifies whether or not to use barcode rules to assign attributes for new media.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

See “[About barcodes](#)” on page 464.

See “[Configuring barcode rules](#)” on page 467.

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

See the *NetBackup Administrator's Guide, Volume II*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

Media type (new media setting)

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.

How NetBackup determines the new media type depends on the **Use barcode rules** setting, as follows:

- **Use barcode rules** is selected.
See “[Media type when using barcode rules](#)” on page 460.
- **Use barcode rules** is not selected.
See “[Media type when not using barcode rules](#)” on page 462.

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 481.

Media type when using barcode rules

If you use barcode rules in NetBackup, choose one of the following:

DEFAULT

NetBackup uses the barcode rules to determine the media type that is assigned.

Each media type to be added should have a barcode 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 barcode rules for DLT and half-inch cartridges and then select the specific media types when you create the barcode 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 barcode rule, NetBackup assigns the default media type for the robot.

A specific media type from the list.

You can use a single barcode 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 barcode rule media type when you create the barcode rule. You can perform one update for DLT and another for half-inch cartridge, and the barcode rule assigns the correct media type.

If you specify a value other than **DEFAULT**, the barcode rule media type must be the same as the media or be **DEFAULT**. If not, the barcode rule does not match the media (except for cleaning media).

[Table 9-5](#) shows some combinations of media types on the **Media Settings** tab and barcode rule media types for a TLD (non-API) robot. It also shows the results when the media are added to the volume configuration.

Table 9-5 Example media type and barcode rule combinations

Media type on Media Settings tab	Barcode 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

Table 9-5 Example media type and barcode rule combinations (*continued*)

Media type on Media Settings tab	Barcode rule media type	Rule matches?	Media type added to volume configuration
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 barcode matches a barcode tag.
- The media type for the barcode rule is cleaning media (DLT_CLN).

Another example is available:

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 barcode rule. In the seventh, specify the cleaning media on the **Media Settings** tab and specify default when you configure the barcode rule.

See [“Configuring barcode rules”](#) on page 467.

Media type when not using barcode rules

Choose one of the following if the barcode rules in NetBackup are not used:

DEFAULT

NetBackup uses the media type that is configured for the drives if:

- The drives in the robot are configured on the robot control host
- All drives the same type
- At least one drive is configured on the robot control host

If the drives are not the same type, NetBackup uses the default media type for the robot.

A specific media type

If the robot supports multiple media types and you do not want to use the default media type, select a specific type.

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.

The following table shows the default media types for robots when drives are not configured on the robot control host:

Table 9-6 Default media types for non-API robots

Robot type	Default media type
Tape Library 4 MM (TL4)	4 MM cartridge tape.
Tape Library 8 MM (TL8)	8 MM cartridge tape. Also supports 8 MM cartridge tape 2 and 8 MM cartridge tape 3.
Tape Library DLT (TLD)	DLT cartridge tape. Also supports the following: <ul style="list-style-type: none">DLT cartridge tape 2 and 3, 1/2-inch cartridge tape1/2-inch cartridge tape 2, 1/2-inch cartridge tape 38 MM cartridge tape, 8 MM cartridge tape 2, 8 MM cartridge tape 3DTF cartridge tape1/4-inch cartridge tape

Volume pool (new media setting)

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

DEFAULT

DEFAULT. If you select DEFAULT and:

- Use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned
- Do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool

A specific volume pool.

If you use barcode rules, this volume pool setting always overrides the rule.

About barcodes

When a robotic library has a barcode reader, it scans the media for barcodes and saves the results. The results associate the slot number and the barcode with the media in that slot. NetBackup obtains the barcode and slot information from the robotic library.

In the robots that have barcode readers, NetBackup verifies the barcode to ensure that the robot loads the correct volume.

If the barcode on the volume does not match the barcode 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.

Barcode advantages

NetBackup functions well whether or not barcodes are used. However, Symantec suggests using media with barcodes in the robots that can read barcodes.

Barcodes 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 barcodes can adversely affect performance for some robots. A robot that reads barcodes performs a scan each time it moves a tape. The robot stores the correct barcode in memory or verifies a previously saved barcode. However,

if a tape does not have a barcode, the robot retries the scan multiple times, degrading performance.

Barcode best practices

Consider the following practices when you select barcodes for volumes:

- Barcodes usually appear on the labels that are attached to the outside of tape volumes.
- The maximum barcode length that NetBackup supports depends on the type of robot.
See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.symantec.com/docs/DOC5332>
- Always follow the robotic library vendor's recommendations when purchasing barcode labels for use with NetBackup.
Ensure that the barcodes have the correct number of characters.
- Barcodes can represent any combination of alpha and numeric characters, but different robots support different lengths of barcodes.
See the robot vendor's documentation to determine the requirements for a specific robot type.
- Use barcodes 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 barcode.
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 barcode for a volume must be identical to the NetBackup media ID.
Match barcodes 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 barcode labels in that series.
- When a robotic library can contain more than one media type, assign specific characters in the barcode to different media types. Do so by using media ID generation rules.
Also, use barcodes to differentiate between data tapes and cleaning tapes or to differentiate between volume pools.

About barcode rules

A barcode rule specifies criteria for assigning attributes to new robotic volumes. NetBackup assigns these attributes by using the barcode for the volume that the robotic library provides and your barcode rules.

In NetBackup, you choose whether to use barcode rules when you set up the robot inventory update operation. The barcode rules are stored on the master server.

Note: NetBackup does not use barcode rules if a volume already uses a barcode.

About NetBackup actions for barcodes

When a robot inventory update operation uses NetBackup barcode rules and a new barcode 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 barcode.
- If the barcode matches a rule, NetBackup verifies that the media type in the rule is compatible with the media type specified for the update.
- If the media types match, NetBackup 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.

Example barcode rules

The following table shows some example barcode rules. Rules are sorted first according to the number of characters in the barcode 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 9-7 Example barcode rules

Barcode 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

Table 9-7 Example barcode rules (*continued*)

Barcode tag	Media type	Volume pool	Max mounts and cleanings	Description
TL	8MM	None	0	8-mm no pool
<NONE>	DEFAULT	None	0	No barcode
<DEFAULT>	DEFAULT	NetBackup	0	Other barcodes

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 barcode rules = YES

Volume pool = DEFAULT

If a new volume in this robotic library has a barcode of TL800001, NetBackup uses the rule with the barcode tag of TL8. NetBackup assigns the following attributes to the volume:

- Media ID = 800001 (last six characters of barcode)
- Volume group = 00_000_TL8
- Volume pool = t_pool
- Maximum mounts = 0 (no maximum)

If a new volume has a barcode of TL000001, NetBackup uses the rule with the barcode tag of TL. NetBackup assigns the following attributes to the volume:

- Media ID = 000001 (last six characters of barcode)
- Volume group = 00_000_TL8
- Volume pool = None
- Maximum mounts = 0 (no maximum)

Configuring barcode 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 barcodes when you select **Use barcode rules** on the **Media Settings** tab.

To enable barcode rule support for API robots, add an `API_BARCODE_RULES` entry to the `vm.conf` file.

Robot types are described in a different topic.

See “[NetBackup robot types](#)” on page 393.

Information about the `vm.conf` file is in the *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

See “[About barcodes](#)” on page 464.

To configure barcode rules

- 1 Open the **Advanced Robot Inventory Options** dialog box, as follows:

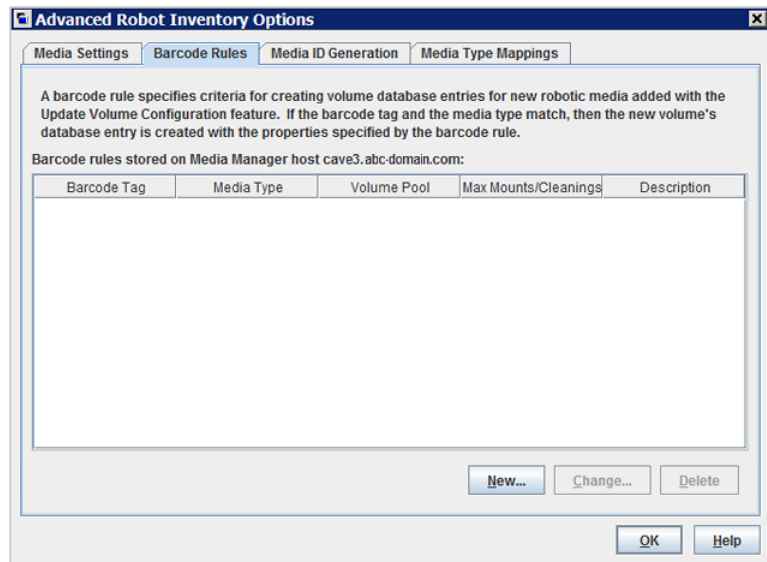
From the **Robot Inventory** dialog box

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.
- 2 Select the robot that you want to inventory.
- 3 On the **Actions** menu, select **Inventory Robot**.
- 4 Click either **Preview volume configuration changes** or **Update volume configuration**.
- 5 Click **Advanced Options**.

From the **Volume Configuration 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.
- 3 On the **Robot Inventory** panel of the **Volume Configuration Wizard**, click **Advanced Options**.

- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Barcode Rules** tab.



3 To configure the rules, do any of the following:

- | | |
|---------------|--|
| Add a rule | Click New and then configure the rule in the dialog box. |
| Change a rule | <p>Select the rule, click Change, and then change the rule in the Change Barcode Rule dialog box.</p> <p>You cannot change the barcode tag of a barcode rule. You first must delete the old rule and then add a rule with a new barcode tag.</p> |
| Delete a rule | Select the rule, click Delete , and click OK in the Delete Barcode Rules dialog box. You can select and delete multiple rules with one operation. |

See [“Barcode rules settings”](#) on page 470.

4 When you are finished configuring rules, click **OK**.

Barcode rules settings

The following table describes the settings you can configure for barcode rules. NetBackup uses these rules to assign barcodes to new media.

Table 9-8 Barcode rule settings

Barcode rule setting	Description
Barcode tag	<p>A unique string of barcode characters that identifies the type of media.</p> <p>For example, use DLT as the barcode tag for a barcode rule if the following is true:</p> <ul style="list-style-type: none"> ■ You use DLT on the barcodes to identify DLT tapes ■ DLT is not used on any other barcodes in the robot <p>Similarly, if you use CLND for DLT cleaning media, use CLND as the barcode tag for the rule for DLT cleaning media.</p> <p>The barcode tag can have from 1 to 16 characters but cannot contain spaces.</p> <p>The following are the special barcode rules that can match special characters in the barcode tags:</p> <ul style="list-style-type: none"> ■ NONE Matches when rules are used and the volume has an unreadable barcode or the robot does not support barcodes. ■ DEFAULT For volumes with barcodes, this tag matches when none of the other barcode tags match. However, the following must be compatible: the media type in the DEFAULT rule and the media type on the Media Settings tab. <p>You cannot change the barcode tag of a barcode rule. Instead, first delete the old rule, then add a rule with a new barcode tag.</p> <p>Use the Media Settings tab to set up the criteria for a robot update.</p> <p>See "Configuring media settings" on page 454.</p>
Description	A description of the barcode 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 barcodes for the cleaning media that cannot be confused with barcodes for data media. Doing so can avoid a value of 0 for cleaning tapes.</p>

Table 9-8 Barcode rule settings (*continued*)

Barcode rule setting	Description
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 barcode rule. If you specify a value other than <code>DEFAULT</code> on the Media Settings tab, the barcode rule media type must be the same as the media or be <code>DEFAULT</code>. If not, the barcode rule does not match the media (except for cleaning media).</p> <p>See “Media type when using barcode rules” on page 460.</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 445.</p>
Volume pool	<p>The volume pool for the new media. The actions depend on whether you use barcode 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 barcode rules, the barcode rules determine the volume pool to which new volumes are assigned.■ If you do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool. <ul style="list-style-type: none">■ A specific volume pool <p>This volume pool setting always overrides any barcode rules.</p>

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 barcode the robot provides to generate the media ID.

Note: To use media ID generation rules, the robot must support barcodes and the robot cannot be an API robot. Media ID generation rules are saved in the Media Manager configuration file (`vm.conf`). Information about the `vm.conf` file is in the *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

For example, two eight-character barcodes are `S000006L1` and `0000006L1`. Without any media ID generation rules NetBackup uses the last six characters of the barcode to generate media IDs. In this example, the same media ID for the two barcodes is created (`0006L1`).

Use a rule to control how NetBackup creates media IDs by specifying which characters of a barcode 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 barcode lengths. Define rules to specific robots and for each barcode format that has different numbers or characters in the barcode. Multiple rules allow flexibility for the robots that support multiple media types.

Configuring media ID generation rules

For non-API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 393.

Use the **Media ID Generation** tab of the NetBackup **Advanced Robot Inventory Options** dialog box to configure the rules that override the default naming method. To use media ID generation rules, the robot must support barcodes and the robot cannot be an API robot.

See [“About media ID generation rules”](#) on page 472.

To configure media ID generation rules

- 1
- Open the **Advanced Robot Inventory Options** dialog box, as follows:

From the **Robot Inventory** dialog box

- 1
- In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.
- 2
- Select the robot that you want to inventory.
- 3
- On the **Actions** menu, select **Inventory Robot**.
- 4
- Click either **Preview volume configuration changes** or **Update volume configuration**.
- 5
- Click **Advanced Options**.

- From the **Volume Configuration 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.

3

On the **Robot Inventory** panel of the **Volume Configuration Wizard**, click **Advanced Options**.
- 2 In the **Advanced Robot Inventory Options** dialog box, click the **Media ID Generation** tab.



3 To configure the rules, do any of the following:

Add a rule	Click New and then configure the rule in the dialog box.
Change a rule	Select the rule, click Change , and then change the rule in the dialog box. You cannot change the robot number or barcode length of a rule. To change those properties, first delete the old rule and then add a rule.
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.

See [“Media ID generation options”](#) on page 475.

4 When you are finished configuring rules, click **OK**.

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 barcode 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 barcode 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 barcode for tapes in the robot.

You cannot change the barcode length of a rule. Rather, first delete the rule and then add a new rule.

■ **Media ID generation rule**

A **Media ID generation rule** consists of a maximum of six colon-separated fields. Numbers define the positions of the characters in the barcode that are to be extracted. For example, the number 2 in a field extracts the second character (from the left) of the barcode. 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.

Barcode 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.

About media type mapping rules

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 393.

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 483.

See [“Default and allowable media types”](#) on page 483.

See [“NetBackup media types”](#) on page 445.

Note: You can write a barcode rule that contains the media types that are incompatible with vendor media types. However, the robot inventory update may assign NetBackup media types that are inconsistent with the vendor media types. Avoid this problem by grouping barcode rules by media type.

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 > Volumes**.

New Volumes

Media Manager host:

Media type:

☒ Volume is in a robotic library

Select robot

Device host:

Robot:

Number of volumes: Media ID naming style:

Media ID: Partner ID:

First slot number: Maximum mounts:

Volume group: Volume pool:

Media description: Label optical media:

- 4 In the **New Volumes** dialog box, specify the properties for the volumes.

The properties that appear in the dialog box vary.

See ["Volume properties"](#) on page 478.

- 5 Click **Apply** or **OK**.

If the robot has a barcode reader, NetBackup performs the following actions:

- Adds the volume to the EMM database using the specified media ID.

- Reads the barcode of each new volume.
 - Adds the barcodes as attributes in the EMM database.
- The **Apply** option adds the volume without closing the dialog box or refreshing the display. You can then add more volumes.

Volume properties

[Table 9-9](#) describes the properties for volumes in NetBackup. The properties depend on whether you add, change, or move volumes.

The properties are arranged alphabetically.

Table 9-9 Volume properties

Property	Description	Operation
Device host	The name of the NetBackup media server to which the robot is attached.	Add, move
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 494.</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>	Change
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 1093.</p>	Add
First slot number	<p>The number of the first slot in the robot in which the range of volumes resides. If you add or move more than one media, 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>	Add, move

Table 9-9 Volume properties (*continued*)

Property	Description	Operation
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>	Add
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>	Add, change
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 1093.</p>	Add, change
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 barcode on the media (for API robots, NetBackup supports barcodes from 1 to 6 characters). Therefore, obtain a list of the barcodes 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 1093.</p>	Add, change

Table 9-9 Volume properties (*continued*)

Property	Description	Operation
Media ID naming style	<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 barcode on the media. For API robots, NetBackup supports barcodes from 1 to 6 characters. Therefore, obtain a list of the barcodes 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 1093.</p>	Add
Media type	<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 445.</p>	Add
Number of volumes	<p>The number of volumes to add. For a robotic library, enough slots must exist for the volumes.</p>	Add
Robot	<p>The robotic library to add or move 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>	Add, move
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"> ■ Standalone 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 NetBackup volume groups" on page 444.</p> <p>See "About rules for moving volumes between groups" on page 489.</p>	Add, move

Table 9-9 Volume properties (*continued*)

Property	Description	Operation
Volume is in a robotic library	<p>When you add a volume:</p> <ul style="list-style-type: none">■ If the volume is in a robot, select Volume is in a robotic library.■ If the volume is a standalone volume, do not select Volume is in a robotic library. <p>When you move a volume:</p> <ul style="list-style-type: none">■ To inject a volume into a robotic library, select Volume is in a robotic library. Then, select a robot and the slot number (First slot number) for the volume.■ To eject a volume from a robot, clear Volume is in a robotic library.	Add, move
Volume pool	<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 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 NetBackup volume pools” on page 441.</p>	Add, change
Volumes to move	The Volumes to move section of the dialog box shows the media IDs of the volumes that you selected to move.	Move

Configuring media type mappings

Applies to API robots only. Robot types are described in a different topic.

See [“NetBackup robot types”](#) on page 393.

Use the **Media Type Mappings** tab of the NetBackup **Advanced Robot Inventory Options** dialog box to configure the attributes for existing and new media.

See [“About media type mapping rules”](#) on page 476.

To configure media type mappings

- 1 Open the **Advanced Robot Inventory Options** dialog box, as follows:

- From the **Robot Inventory** dialog box

1

In the **NetBackup Administration Console**, expand **Media and Device Management > Media > Robots** in the left pane.

2

Select the robot that you want to inventory.

3

On the **Actions** menu, select **Inventory Robot**.

4

Click either **Preview volume configuration changes** or **Update volume configuration**.

5

Click **Advanced Options**.
- From the **Volume Configuration Wizard**

1

In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Devices**.

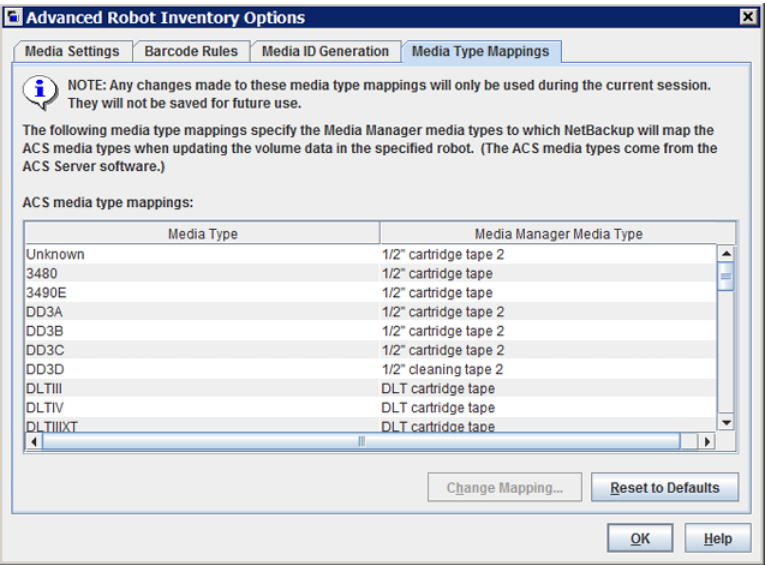
2

From the list of wizards in the right pane, click **Configure Volumes** and follow the wizard instructions.

3

On the **Robot Inventory** panel of the **Volume Configuration Wizard**, 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 default mappings and any mappings you added or changed appear.

- 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 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. Robot types are described in a different topic.

See “[NetBackup robot types](#)” on page 393.

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 run the **NetBackup Administration Console**.

Information about the `vm.conf` file is in the *NetBackup Administrator's Guide, Volume II*:

<http://www.symantec.com/docs/DOC5332>

Table 9-10 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. Robot types are described in a different topic.

See “[NetBackup robot types](#)” on page 393.

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:

- NetBackup ACS type robots: [Table 9-11](#)

- NetBackup TLH type robots: [Table 9-12](#)
- NetBackup TLM type robots: [Table 9-13](#)

The following items provide information to help you understand the tables:

- The first column of each table shows the vendor's media type.
- The second column of each table shows the default media type in NetBackup.
- 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
```

The following table shows the default media types and the allowable media types for ACS robots.

Table 9-11 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

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
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
LTO_800W	1/2-inch cartridge tape (HCART)	HCART, HCART2, HCART3
LTO_1_5_T	1/2-inch cartridge tape 2 (HCART2)	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

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
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
T10000T1	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3
T10000TS	1/2-inch cartridge tape 3 (HCART3)	HCART, HCART2, HCART3

Table 9-11 Default and allowable media types for ACS robots (*continued*)

ACS media type	Default media type	Allowable media types through mappings
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

The following table shows the default and allowable media types for TLH robots.

Table 9-12 Default and allowable media types for TLH robots

TLH media type	Default Media Manager media type	Allowable media types through mappings
3480	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3490E	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3590J	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3590K	1/2-inch cartridge (HCART)	HCART, HCART2, HCART3
3592JA	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JB	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JX	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JJ	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
3592JR	1/2-inch cartridge tape 2 (HCART2)	HCART, HCART2, HCART3
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

The following table shows the default and allowable media types for TLM robots.

Table 9-13 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.

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 489.

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 standalone volume group.
- All members of a group must be in the same robotic library or be standalone. That is, if volume group already exists in another robotic library, you cannot add it (or part of it) to a robotic library.

See [“About NetBackup volume groups”](#) on page 444.

See [“About moving volumes”](#) on page 503.

Changing the owner of a volume

You can change the media server or server group that owns the volume.

See [“About NetBackup server groups”](#) on page 367.

See [“About media sharing”](#) on page 511.

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**.

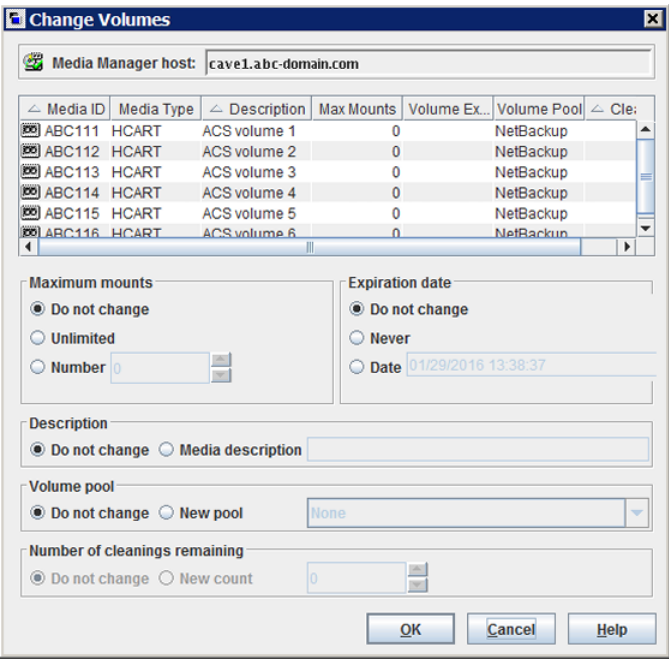
Changing volume properties

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 [“Volume properties”](#) on page 478.
- 5
- Click **OK**.

About assigning and deassigning volumes

An assigned volume is one that is reserved for exclusive use by NetBackup. A volume is set to the assigned state when either application writes data on it for the first time. The time of the assignment appears in the **Time Assigned** column for the volume in the **NetBackup Administration Console Volumes** pane. When a volume is assigned, you cannot delete it or change its volume pool.

A volume remains assigned until NetBackup deassigns it.

To determine which application currently uses a volume, see the **Application** column of the right pane, labeled **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 969.

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 970.

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.

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 assigning and deassigning volumes”](#) on page 491.

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 assigning and deassigning volumes”](#) on page 491.

The following table describes the types of erase.

Table 9-14 Types of erase

Type of erase	Description
SCSI long erase	<p>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.</p> <p>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</p>

Table 9-14 Types of erase (continued)

Type of erase	Description
SCSI quick erase	<p>Rewinds the media and an erase gap is recorded on the media. The format of this gap is drive dependent. It can be an end-of-data (EOD) mark or a recorded pattern that the drive does not recognize as data.</p> <p>Some drives do not support a quick erase (such as QUANTUM DLT7000). For the drives that do not support a quick erase, the new tape header that is written acts as an application-specific quick erase.</p>

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

To erase a volume

- 1 In the **NetBackup Administration Console**, 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.

The following table describes the procedure that used to exchange a volume and use a new media ID.

Table 9-15 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 standalone group.	See “About moving volumes” on page 503.
Step 2	Add a new volume or move an existing volume in as a replacement for the volume you removed. 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).	See “About adding volumes” on page 451.
Step 3	Physically replace the old volume. Do not delete the old volume in case you need to retrieve the data on the volume.	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.

The following table describes the procedure to exchange a volume and use the old media ID.

Table 9-16 Exchange a volume and use the old media ID

Step	Task	Instructions
Step 1	Delete the volume.	See “Deleting a volume” on page 492.
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 497.
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 451.
Step 4	Set a new expiration date for the volume.	See “Changing volume properties” on page 490.
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 502.

About frozen media

Frozen media is the media that NetBackup does not use for backups. NetBackup stops directing the backups and the archives to frozen media. NetBackup never deletes a frozen media ID from the NetBackup media catalog, even after the retention period ends for all backups on the media. NetBackup does not unassign a frozen volume from the NetBackup volume pool when its backup images expire. All unexpired backup images on frozen media continue to be available for restores. NetBackup freezes the tape volumes for a variety of reasons, as follows:

- NetBackup freezes a volume 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). 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 `nbenmcmd` command with the `-media_error_threshold` and `-time_window` options to change the default values.

Additional information about the `nbenmcmd` is available in the *NetBackup Commands Reference Guide*:

<http://www.symantec.com/docs/DOC5332>

- NetBackup freezes a volume if a write failure makes future attempts at positioning the tape unreliable.
- NetBackup freezes the catalog volumes during catalog recovery.
- NetBackup freezes the volumes in some circumstances with write once read many (WORM) media or WORM-capable drives.
See “[About using volume pools to manage WORM media](#)” on page 448.

You can unfreeze the frozen volumes manually.

Freezing or unfreezing a volume

NetBackup freezes volumes under circumstances.

Use the following procedure to manually freeze or unfreeze a volume.

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 to use 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 NetBackup volume configuration with a robot's contents”](#) on page 527.
- 3 Select **Empty media access port prior to update** on the **Robot Inventory** dialog box.

Ejecting volumes

You can eject single or multiple volumes.

You cannot eject multiple volumes with one operation if they reside in multiple robots.

Operator intervention is only required if the robotic library does not contain a media access port large enough to eject all of the selected volumes. For these robot types, NetBackup prompts an operator to remove the media from the media access port so the eject operation can continue.

See [“Media ejection timeout periods”](#) on page 499.

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**.
- 4 In the **Eject Volumes** dialog box, do one of the following actions:

- ACS and TLM robots only Select the media access port to use for the ejection, then click **Eject**.
- All other robots Click **Eject**.

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.

See “[NetBackup robot types](#)” on page 393.

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.

The following table shows the ejection timeout periods for robots.

Table 9-17 Media ejection timeout periods

Robot types	Timeout period
Automated Cartridge System (ACS) Tape Library Multimedia (TLM)	One week
Tape Library 8MM (TL8) Tape Library DLT (TLD)	30 minutes.
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 barcodes

You can rescan the media in a robot and then update NetBackup with the barcodes of that media.

You should rescan and update only in certain circumstances.

Note: Rescan and update barcodes does not apply to volumes in API robot types.

See [“NetBackup robot types”](#) on page 393.

When to rescan and to update barcodes

Rescan and update barcodes only to add the barcodes 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 barcode to the database. Use this command to add the barcode after you insert the tape into the robotic library.

When not to rescan and to update barcodes

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 NetBackup volume configuration with a robot's contents”](#) on page 527.

- Physically move the volume into the correct slot.

To obtain an inventory of the robot without updating the barcode information in the database, inventory the robot and use the show contents option.

See [“Showing the media in a robot”](#) on page 520.

See [“About barcodes”](#) on page 464.

Rescanning and updating barcodes

Use the following procedure to rescan the media in a robot and to update NetBackup with the barcodes.

Note: Rescan and update barcodes does not apply to volumes in API robot types.

See [“NetBackup robot types”](#) on page 393.

See “[About rescanning and updating barcodes](#)” on page 500.

To rescan barcodes 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**.
- 5 Click **Start**.

The results of the update appear in the output section of the dialog box.

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 barcode, as follows:

- If the robot supports barcodes and the media has barcodes, NetBackup uses the last six characters of the barcode 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 473.
- For volumes without barcodes, 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. Information about the `vm.conf` file is available in the *NetBackup Administrator's Guide, Volume II*:
<http://www.symantec.com/docs/DOC5332>

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 assigning and deassigning volumes”](#) on page 491.

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 501.

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.

See [“About NetBackup volume groups”](#) on page 444.

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 barcode reader and the volumes must contain readable barcodes.
To remove volumes from a robot.	Use this procedure even if the volumes do not contain barcodes 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 514.

Moving volumes by using the Actions menu

If you move a volume to a robotic library that has a barcode reader, NetBackup updates the EMM database with the correct barcode.

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**.
- 5 In the **Move Volumes** dialog box, specify the properties for the move.

If you move a single volume, the dialog box entries show the current location of the volume.

See [“Volume properties”](#) on page 478.

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 assigning and deassigning volumes”](#) on page 491.

Recycling a volume and using a new media ID

Recycle a volume if it is a duplicate of another volume with the same media ID. Also recycle a volume if you change how you name volumes and you want to match the barcodes on the volume.

The following table describes the procedure to recycle a volume and use a new media ID.

Table 9-18 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 498.
Step 2	If the volume is in a robotic library, move it to stand alone.	See “About moving volumes” on page 503.
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 NetBackup Administration Console).
Step 4	Delete the volume entry.	See “Deleting a volume” on page 492.

Table 9-18 Recycling a volume and using a new media ID (*continued*)

Step	Action	Description
Step 5	Add a new volume entry.	<p>See “Adding volumes by using the Actions menu” on page 477.</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 498.
Step 7	Configure the number of mounts	<p>Set the number of mounts to the value you recorded earlier by using the following command:</p> <p>On Windows hosts:</p> <pre>install_path\Volmgr\bin\vmchange -m media_id -n number_of_mounts</pre> <p>On UNIX hosts:</p> <pre>/usr/openv/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 490.

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**.

Managing volume pools

The following sections describe the operations you can perform to manage volume pools.

Adding a volume pool

Use this procedure to add a new volume pool.

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 > Volume Pool**.
- 3 In the **New Volume Pool** dialog box, specify the properties for the volume pool.
See [“Volume pool properties”](#) on page 508.
- 4 Add volumes to the pool by adding new volumes to NetBackup or by changing the pool of existing volumes.
See [“Adding volumes by using the Actions menu”](#) on page 477.
See [“Adding volumes by using the wizard”](#) on page 453.
See [“Changing volume properties”](#) on page 490.

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 508.

Volume pool properties

You can specify various properties for a volume pool.

The following table describes the properties that you can configure for volume pools, either when you add a new pool or change an existing one.

Table 9-19 Volume pool properties

Property	Description
Catalog backup pool	Select this option to use this volume pool for catalog backups. 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.
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. See “About reserved volume pool name prefixes” on page 442.
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.

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.

Managing volume groups on UNIX

These topics describe operations you can perform to manage volume groups.

Moving a volume group

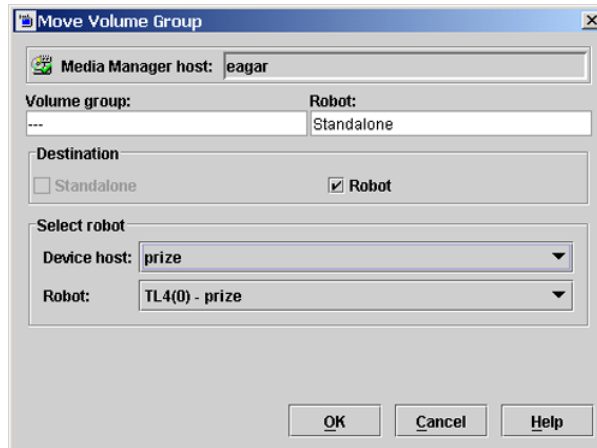
You can move a volume group from a robotic library to standalone storage or from standalone storage to a robotic library.

Moving a volume group changes only the residence information in the EMM database. You must move the volumes physically to their new locations.

To move a volume group

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the right pane, select the volume group that you want to move.
- 3 On the **Actions** menu, select **Move**.

- 4 In the **Move Volume Group** dialog box, specify the properties for the move.



You can only specify the properties that apply for the move type.

The dialog box shows a move from standalone to the robot TL4(0) on host prize.

See [“Move volume group properties”](#) on page 510.

- 5 After you move the volume group logically, physically move the volumes to their new locations.

Move volume group properties

The following table describes the properties you can set when you move a volume group.

Table 9-20 Move volume group properties

Property	Description
Destination	<p>The destination for the move, as follows:</p> <ul style="list-style-type: none"> ■ If you move the volume group from a robotic library, Standalone is selected as the destination. ■ If you move the volume group from standalone, Robot is selected as the destination.
Device host	The host that controls the robotic library.
Robot	The destination robotic library.

Table 9-20 Move volume group properties (continued)

Property	Description
Volume group	The volume group to move. Displays "---" when you move standalone volumes.

Deleting a volume group

Use the following procedure to delete a volume group.

To delete a volume group

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Media**.
- 2 In the volumes list, verify that all of the volumes in the group are unassigned. You cannot delete the group until the application unassigns the volumes. If the **Time Assigned** column contains a value, the volume is assigned.

See [“About assigning and deassigning volumes”](#) on page 491.
- 3 Select a volume group in the right pane.
- 4 On the **Edit** menu, select **Delete**.
- 5 In the confirmation dialog box, confirm the action.
- 6 Remove the deleted volumes from the storage device.

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 512.
- Media media sharing with server groups.
See [“Configuring media sharing with a server group”](#) on page 512.

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 both 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 NetBackup server groups”](#) on page 367.

[Table 9-21](#) outlines the process for configuring media sharing with a server group.

Note: Do not use both unrestricted media sharing and media sharing server groups. If you use both, NetBackup behavior is undefined.

Table 9-21 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.
Step 2	Configure the media sharing server group.	See “Configuring a NetBackup server group” on page 367.
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 507. See “Changing the properties of a volume pool” on page 507.
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 backup policy” on page 673.

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 previewing volume configuration changes](#)
- [Previewing volume configuration changes for a robot](#)
- [About updating the NetBackup volume configuration](#)
- [Updating the NetBackup volume configuration with a robot's contents](#)
- [Robot inventory options](#)
- [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.

The following table describes the **NetBackup Administration Console** robot inventory options for the robotic libraries that contain barcode readers and contain barcoded media.

Table 10-1 Robot inventory options

Inventory option	Description
Show contents	<p>Queries the robot for its contents and 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 517.</p> <p>For the robotic libraries without barcode readers (or that contain media without barcodes), 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 vmphyinv physical inventory utility” on page 530.</p>
Compare contents with volume configuration	<p>Queries the robot for its contents and compares the contents with the contents of the EMM database. Does not change the database.</p> <p>See “About comparing a robot's contents with the volume configuration” on page 521.</p>
Preview volume configuration changes	<p>Queries the robot for its contents and compares the contents 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 523.</p>
Update volume configuration	<p>Queries the robot for its contents; if necessary, 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 NetBackup volume configuration” on page 526.</p>

When to inventory a robot

The following table describes the criteria to use to determine when to inventory a robot and which options to use for the inventory.

Table 10-2 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 barcode numbers.</p> <p>See “Showing the media in a robot” on page 520.</p>
To determine if volumes were moved physically within a robot	<p>For the robots with barcode readers and the robots that contain media with barcodes, use the Compare contents with volume configuration option.</p> <p>See “Comparing media in a robot with the volume configuration” on page 522.</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. The update creates media IDs (based on barcodes or a prefix that you specify).</p> <p>See “Updating the NetBackup volume configuration with a robot's contents” on page 527.</p>
To determine whether new media have barcodes before you add them to NetBackup	<p>Use the Preview volume configuration 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 NetBackup volume configuration with a robot's contents” on page 527.</p>
To insert existing volumes into a robot (an existing volume is one that already has a NetBackup media ID)	<p>If the robot supports barcodes and the volumes have readable barcodes, 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 NetBackup volume configuration with a robot's contents” on page 527.</p> <p>If the robot does not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 503.</p> <p>See “About the vmphyinv physical inventory utility” on page 530.</p>
To move existing volumes between robotic and standalone (an existing volume is one that already has a NetBackup media ID)	<p>If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic or standalone location.</p> <p>See “Updating the NetBackup volume configuration with a robot's contents” on page 527.</p>

Table 10-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
To move existing volumes within a robot (an existing volume is one that already has a NetBackup media ID)	<p>If the robot supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new slot location.</p> <p>See “Updating the NetBackup volume configuration with a robot's contents” on page 527.</p> <p>If the robot does not support barcodes or if the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 503.</p> <p>See “About the vmphyinv physical inventory utility” on page 530.</p>
To move existing volumes from one robot to another (an existing volume is one that already has a NetBackup media ID)	<p>If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the NetBackup volume configuration information.</p> <p>See “Updating the NetBackup volume configuration with a robot's contents” on page 527.</p> <p>If the robots do not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility.</p> <p>See “About moving volumes” on page 503.</p> <p>See “About the vmphyinv physical inventory utility” on page 530.</p> <p>For either operation, perform the following updates:</p> <ul style="list-style-type: none"> ■ First move the volumes to standalone ■ 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>
To remove existing volumes from a robot (an existing volume is one that already has a NetBackup media ID)	<p>For any robot NetBackup supports, use the Update volume configuration option to update the NetBackup volume configuration information.</p> <p>See “Updating the NetBackup volume configuration with a robot's contents” on page 527.</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.

The following table describes the report contents.

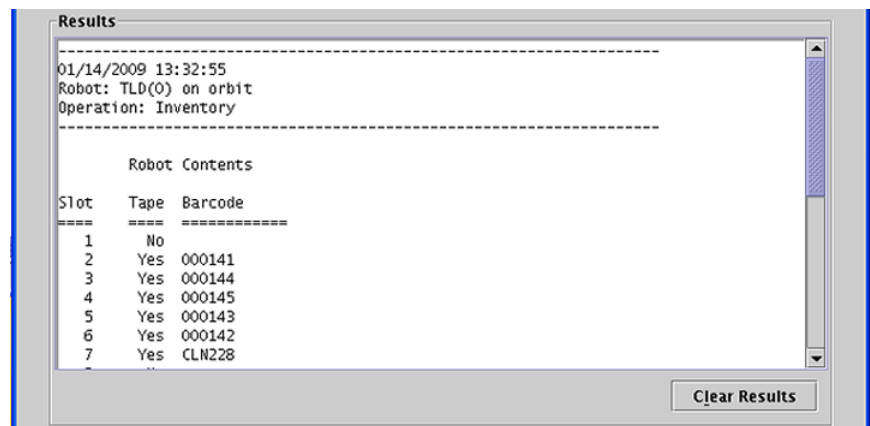
Note: On UNIX: If a volume is mounted in a drive, the inventory report lists the slot from which the volume was moved to the drive.

Table 10-3 Show contents description

Robot and media	Report contents
The robot has a barcode reader and the robot contains media with barcodes.	Shows if each slot has media and lists the barcode for the media.
The robot does not have a barcode reader or the robot contains media without barcodes.	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 518.

The following figure shows an example of the report.

Figure 10-1 Show contents report



See [“Showing the media in a robot”](#) on page 520.

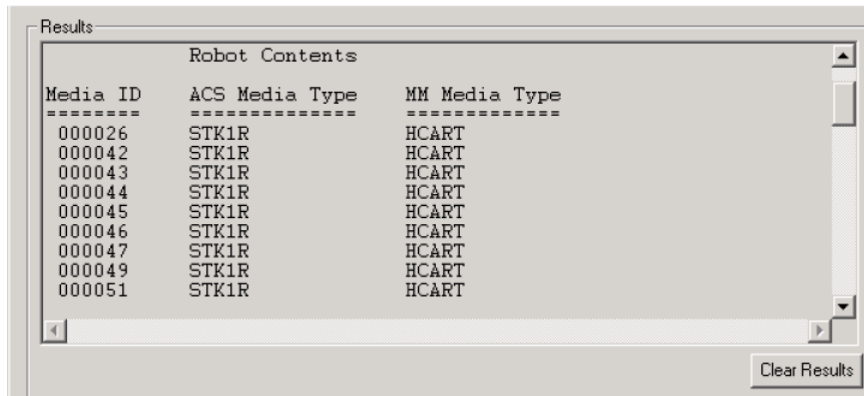
About inventory results for API robots

The following table describes the contents of the robot inventory for the API robots.

Table 10-4 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 barcode 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 barcode 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 barcode rules).

The following figure shows the results for an ACS robot; the results for other API robots are similar.

Figure 10-2 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 514.

See [“Robot inventory options”](#) on page 529.

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.

- On the **Actions** menu, select **Inventory Robot**.



- In the **Robot Inventory** dialog box, select **Show contents**.
- 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 10-5 Compare contents description

Robot and media	Report contents
The robot can read barcodes	The report shows the differences between the robot and the EMM database

Table 10-5 Compare contents description (continued)

Robot and media	Report contents
The robot cannot read barcodes	<p>The report shows only whether a slot contains a volume</p> <p>If the media have barcodes, this operation is useful for determining if volumes have been physically moved within a robot.</p>
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.

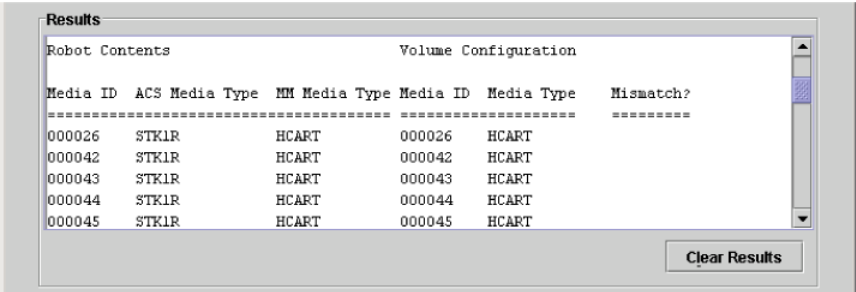
If the results show that the EMM database does not match the contents of the robotic library, perform the following actions:

- Physically move the volume.
- Update the EMM database. Use **Actions > Move** or use the **Update volume configuration** option.

See [“About updating the NetBackup volume configuration”](#) on page 526.

The following figure shows a sample compare report.

Figure 10-3 Compare contents report (API robot)



See [“Comparing media in a robot with the volume configuration”](#) on page 522.

Comparing media in a robot with the volume configuration

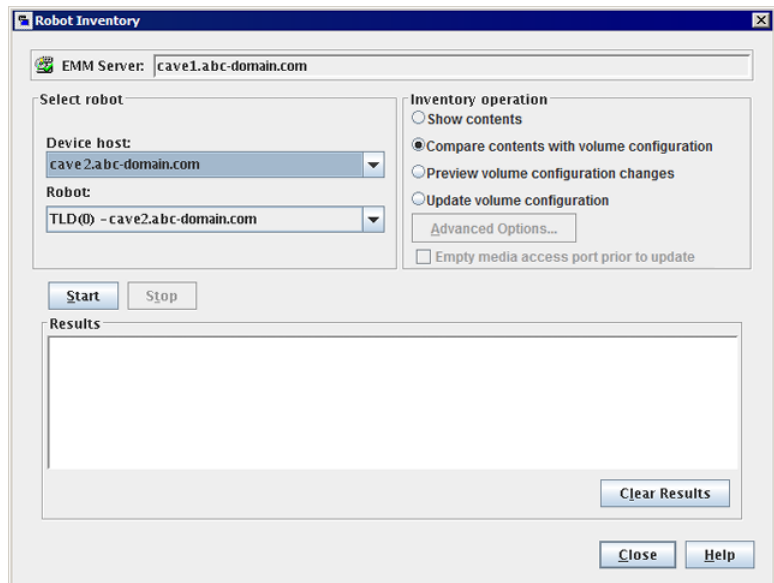
Use the following procedure to compare the media in a robot with the EMM database.

See [“About robot inventory”](#) on page 514.

See [“Robot inventory options”](#) on page 529.

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 that 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 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 barcodes before you add them to the EMM database.

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 barcode rules, and so on.

Figure 10-4 Preview volume configuration changes (non-API robot)



See [“Updating the NetBackup volume configuration with a robot's contents”](#) on page 527.

Previewing volume configuration changes for a robot

Use the procedure in this topic to preview any volume configuration changes for a robot.


See [“About previewing volume configuration changes”](#) on page 523.

See [“Robot inventory options”](#) on page 529.

To preview the volume configuration changes for a robot

- 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 **Preview volume configuration changes**.

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 barcode rules, and so on.

- 6 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, click **Advanced Options**.
 See [“About configuring media name and attribute rules”](#) on page 453.
- 7 To inject any media that is in the media access port before the preview operation, click **Empty media access port prior to update**.
- 8 Click **Start** to begin the inventory preview.

About updating the NetBackup volume configuration

The **Update volume configuration** robot inventory option 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 529.

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


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

[Figure 10-5](#) is an example for an ACS robot. Results for other API robots are similar.

Robot inventory update returns an error if it encounters unsupported characters in the volume serial number or media identifier from API robots.

See [“Volume update prerequisites”](#) on page 526.

Figure 10-5 Update volume configuration for API robot report



See [“Updating the NetBackup volume configuration with a robot's contents”](#) on page 527.

Volume update prerequisites

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

- The robotic library must read barcodes.

- Volumes in the library must have readable barcodes.

You can check the barcode 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 522.

If the robotic library does not support barcodes or the volumes do not have readable barcodes, save the results of the compare operation. The results can help you determine a media ID prefix if you use the **Media Settings** tab of the **Advanced Options** dialog box to assign a prefix.

Updating the NetBackup 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 updating the NetBackup volume configuration”](#) on page 526.

See [“Robot inventory options”](#) on page 529.

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**.

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 barcode rules, and so on.

- 6 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, click **Advanced Options**.
See [“About configuring media name and attribute rules”](#) on page 453.
- 7 To inject any media that is in the media access port before the update operation, click **Empty media access port prior to update**.
- 8 Click **Start** to begin the inventory update.

Robot inventory options

The following table shows the robot inventory options when you use the **NetBackup Administration Console**:

Table 10-6 Robot inventory options

Option	Description
Advanced options	<p>Advanced Options is active if Preview volume configuration changes or Update volume configuration is selected.</p> <p>This button opens the Advanced Robot Inventory Options dialog box, from which you can configure more options.</p> <p>See “About configuring media name and attribute rules” on page 453.</p>
Device host	The Device host option is the host that controls the robot.
Empty media access port prior to update	<p>The Empty media access port prior to update operation is active only for the robots that support that function.</p> <p>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.</p> <p>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 media access port is empty, you are not prompted to place volumes in the media access port.</p> <p>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.</p>
Robot	<p>Use the Robot option to select a robot to inventory.</p> <p>If you selected a robot in the NetBackup Administration Console, that robot appears in this field.</p>
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 517.</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 521.</p>

Table 10-6 Robot inventory options (*continued*)

Option	Description
Preview volume configuration 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 523.
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 NetBackup volume configuration” on page 526.

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 barcode readers
- For the robots that contain media without barcodes

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 non-barcoded 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 the *NetBackup Commands Reference Guide*, available at the following URL:

<http://www.symantec.com/docs/DOC5332>

Table 10-7 `vmphyinv` features, requirements, restrictions, and when to use

Item	Description
vmphyinv 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 barcoded tape libraries because it verifies the contents of the media. ■ Recognizes the NetBackup 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.
vmphyinv 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.
When to use vmphyinv	<p>Use <code>vmphyinv</code> to update the EMM database for NetBackup in the following cases:</p> <ul style="list-style-type: none"> ■ You want to inventory a robot that does not have a barcode reader or that contains non-barcoded media. ■ You insert new media into a robotic library and no NetBackup volume records correspond to the media. Use the slot range or list option of <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 slot range or list 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 532.

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 532.
- Obtains a list of media to mount
See [“About the media that `vmphyinv` mounts”](#) on page 532.
- Mounts the media and reads the tape headers
See [“How `vmphyinv` mounts the media and reads the tape header”](#) on page 533.
- Updates the EMM database
See [“How `vmphyinv` updates the EMM database”](#) on page 535.

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 that `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.

<code>vmphyinv -rn 2 -v moon</code>	Mounts the media in robot 2 and in the volume group moon.
<code>vmphyinv -rn 1 -rcl 2 -number 3</code>	Mounts the media in robot 1 and slot range 2 to 4.
<code>vmphyinv -rn 5 -pn NetBackup -v mars -rcl 2 -number 6</code>	Mounts the media in robot 5, slot range 2 to 7, in volume group mars, and in the NetBackup volume pool.

- NetBackup robot number and a list of media that belong to a specific robot.
For example, if the `-rn robot_number` and `-ml A00001:A00002:A00003` options are specified, only the three specified media are inventoried. If any of these media do not belong to the specified robot, the media are skipped and are not inventoried. To use this option, NetBackup volume records must exist.
- NetBackup robot number and a slot range or list.
Sometimes, media from a different robot or some other source are moved to a robot and the media ID on the tape is unknown. In these cases, specify a slot range option or list option.
With these options, the NetBackup volume record does not need to exist in the EMM database. However, you must specify the density (using the `-d` option).

Note: For a robot that supports multiple media types, specify the density carefully. If you specify the incorrect density, `vmphyinv` cannot complete the mount and permanent drive failure can occur.

The following are some filter examples.

<code>vmphyinv -rn 1 -slot_range 2 10 -d dlt</code>	Mounts the media in slot range 2 to 10 in robot 1.
<code>vmphyinv -rn 0 -slot_list 3:4:5 -d 8mm</code>	Mounts the media in slots 3, 4, and 5 in robot 0.
<code>vmphyinv -rn 2 -slot_range 2 4 -slot_list 5:6:7 -d dlt</code>	Mounts the media in slots 2, 3, 4, 5, 6, and 7 in robot 2.

See [“About the `vmphyinv` physical inventory utility”](#) on page 530.

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).

Note: The default mount timeout is 15 minutes. Specify a different mount time by using the `-mount_timeout` option.

See [“About the media that `vmphyinv` does not recognize”](#) on page 534.

See [“How `vmphyinv` processes cleaning media”](#) on page 534.

About the media that `vmphyinv` does not recognize

If the media is not NetBackup 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 of 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 10-8 `vmphyinv` criteria and actions

Criteria or action	Description
The <code>vmphyinv</code> update 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.■ 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 10-8 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 barcode of an existing volume record needs to be changed.

See [“About the vmphyinv physical inventory utility”](#) on page 530.

Example volume configuration updates

The following examples show different types of volume configuration updates. The examples include only the relevant volume attributes.

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 barcodes.

The following are the attributes for media ID 800001:

Media ID	800001
Media type	8MM cartridge tape
Barcode	TL800001
Media description	TL8 backup volume
Volume pool	NetBackup
Robot type	TL8 - Tape Library 8MM
Volume group	EXB220
Maximum 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
Barcode	TL800001
Media description	TL8 backup volume
Volume pool	NetBackup
Robot type	NONE - Not Robotic
Volume group	NONROB_8MM
Maximum mounts allowed	0 (unlimited)

The new residence information in the EMM database shows a standalone 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 barcode.

Example 2: Adding existing standalone volumes to a robot

The following is an example of how to add a standalone volume that has a barcode to a robotic library that supports barcodes (TL8).

When you move volumes from one robot to another robot, perform separate updates.

The following are the volume attributes for media ID 800021, which has a readable barcode and already exists as a standalone volume:

Media ID	800021
Media type	8MM cartridge tape
Barcode	TL800021
Media description	8MM standalone
Volume pool	None
Robot type	None (standalone)

Volume group NONROB_8MM
Maximum mounts allowed 0 (unlimited)

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 barcode rules YES (selected)
Volume pool NetBackup

Assume that the barcode rules in [Table 10-9](#) exist.

Table 10-9 Example barcode rules

Barcode 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 barcode
<DEFAULT>	DEFAULT	NetBackup	0	Other barcodes

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
Barcode TL800021
Media description 8MM standalone
Volume pool NONE

Robot type	TL8 - Tape Library 8MM
Robot number	0
Robot slot	1
Robot host	shark
Volume group	EXB220
Maximum mounts allowed	0 (unlimited)

The barcode matches the barcode of an existing standalone 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, barcode rules are ignored.

The only setting used on the **Media Settings** tab is the volume group for added or moved volumes. The media type setting was not used because this example was for a single existing volume that already had a media type.

Example 3: Moving existing volumes within a robot

The following is an example of how to move a volume from one slot to another slot within the same robot. The robot supports barcodes and the volume has a readable barcode.

Note: To move volumes within a robotic library, use **Update volume configuration** only if the robotic library supports barcodes and the volumes have readable barcodes. 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
Barcode	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)

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 barcode rules	NO (not selected)
Volume pool	DEFAULT

The resulting volume attributes are the following:

Media ID	800002
Media type	8MM cartridge tape
Barcode	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
Maximum mounts allowed	0 (unlimited)

The updated volume attributes show the new slot number, but all other information is unchanged.

Example 4: Adding new volumes to a robot

The following is an example of how to add new volumes with barcodes to a robot that supports barcodes.

Assume the following:

- The new volume is an 8MM tape with a readable barcode 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 barcode rules	YES (selected)
Volume pool	DEFAULT

The following table contains the example barcode rules.

Table 10-10 Example barcode rules

Barcode 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 barcode

The barcode on the media matches the barcode rule named TL8 and the resulting volume attributes for the new volume are as follows:

Media ID	800002
Media type	8MM cartridge tape
Barcode	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 barcode. 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 barcode rule.

If barcode rules (or barcodes) 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
Maximum mounts allowed	0 (unlimited)

Note: If the robot does not support barcodes or the barcode is unreadable, specify a Media ID prefix on the **Media Settings** tab. Alternatively, specify DEFAULT for the media ID. If you do not, NetBackup does not add new media IDs.

Example 5: Adding cleaning tapes to a robot

A special case exists when you add cleaning tapes. For example, assume that you update a TLD robot.

The tapes you inserted include regular tapes with barcodes that range from DLT00000 to DLT00010 and a cleaning tape with a barcode of CLN001.

The following table contains the example barcode rules.

Table 10-11 Example barcode rules

Barcode tag	Media type	Volume pool	Max mounts/ cleanings	Description
CLN	DLT_CLN	None	30	DLT cleaning

Table 10-11 Example barcode rules (*continued*)

Barcode tag	Media type	Volume pool	Max mounts/ cleanings	Description
DL	DLT	d_pool	200	DLT backup
<NONE>	DEFAULT	None	0	No barcode

Specify the following on the **Media Settings** tab, then run the update.

Media type	DLT
Volume group	STK7430
Use barcode rules	YES (selected)

The barcodes on the regular tapes match the DL barcode rule. The media type of the DL barcode rule matches the Media type on the **Media Settings** tab. The tapes are added as DLT.

The cleaning tape matches the CLN barcode 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 barcode on the volume matches a barcode tag (CLN in this example).
- The media type for the barcode 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 barcode rule (DLT_CLN in this example).

Example 6: Moving existing volumes between robots

When you move volumes from one robot to another and the volumes in both robots are in the same EMM database, perform two separate updates.

These updates move the volumes to stand alone, as an intermediate step, and then to the new robot. Otherwise, NetBackup is unable to update the entries and you receive an "Update request failed" error.

This example assumes that robot 2 is able to read barcodes and the volume has readable barcodes. If not, NetBackup cannot manage the volumes.

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 standalone.
- Perform an Update volume configuration on robot 2. This action updates the configuration to show the volume in robot 2.

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 standalone volume to a TL4 robot. A TL4 robot supports media inventory (detects media presence), but not barcodes.

The following are the attributes for media ID 400021, which already exists as a standalone volume:

Media ID	400021
Media type	4MM cartridge tape
Barcode	-----
Media description	4MM standalone
Volume pool	None
Robot type	NONE - Not Robotic
Volume group	NONROB_4MM
Max mounts allowed	0 (unlimited)

Example 7: Adding existing volumes when barcodes are not used

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
Barcode	-----
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
Maximum mounts allowed	0 (unlimited)

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 barcodes or the robot does not support barcodes. Without a barcode, NetBackup cannot identify the volume and assumes that it is new. The media ID C40000 is generated from the media ID prefix that is 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 530.

Configuring storage units

This chapter includes the following topics:

- [About the Storage utility](#)
- [Creating a storage unit](#)
- [About storage unit settings](#)

About the Storage utility

The data that is generated from a NetBackup job is recorded into a type of storage that NetBackup recognizes.

NetBackup recognizes the following storage configurations, all of which are configured in the **Storage** utility:

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. Storage units can be included as part of a storage unit group or a storage lifecycle policy.

See [“Creating a storage unit”](#) on page 548.

Storage unit groups

Storage unit groups let you identify multiple storage units as belonging to a single group. The NetBackup administrator configures how the storage units are selected within the group when a backup or a snapshot job runs.

See [“About storage unit groups”](#) on page 589.

Storage lifecycle policies

Storage lifecycle policies let the administrator create a storage plan for all of the data in a backup or snapshot.

See [“About storage lifecycle policies”](#) on page 600.

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 1112.

Figure 11-1 Storage Unit node of the Storage utility

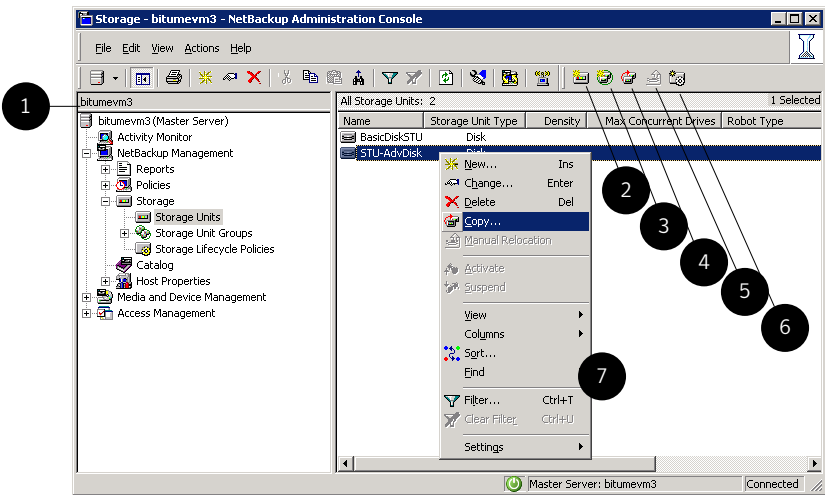


Table 11-1 Storage Unit utility

Number	Description
1	Currently selected master server.
2	Click to create a new storage unit. See “About storage unit settings” on page 561.
3	Click to create a new storage unit group. See “Creating storage unit groups for backups” on page 590.

Table 11-1 Storage Unit utility (continued)

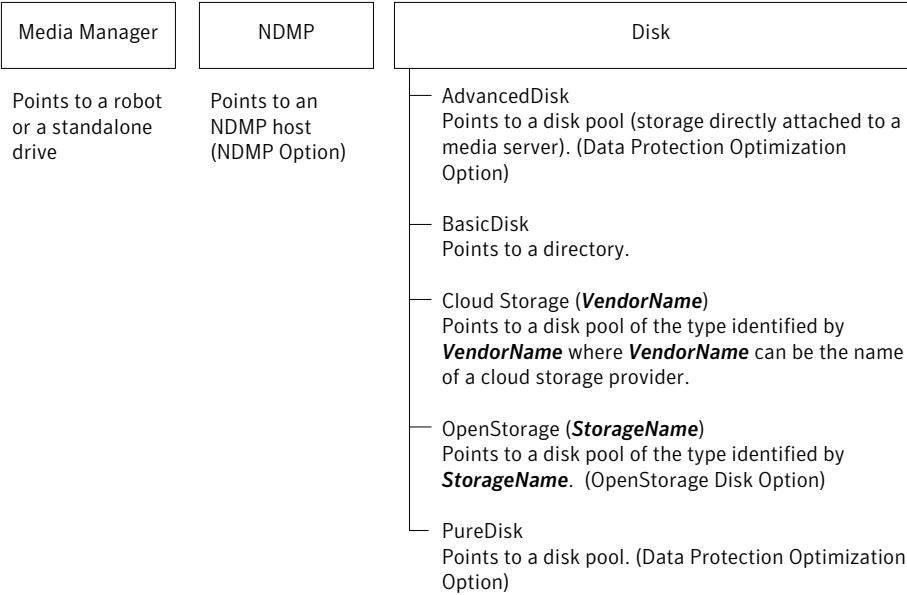
Number	Description
4	Click to copy a storage unit. See “Creating a storage unit by copying a storage unit” on page 551.
5	Click to initiate a data management job when using basic disk staging. See “About basic disk staging” on page 577.
6	Click to create a new storage lifecycle policy. See “Creating a storage lifecycle policy” on page 601.
7	Right-click in the right pane to view the shortcut menu.

Creating a storage unit

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. Storage unit creation is part of several other wizards. However, a storage unit can be created directly from the **Storage** utility in the **NetBackup Administration Console**.

[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



To create a storage unit

- 1 In the **NetBackup Administration Console**, select the **Storage** utility.
- 2 Select **Actions > New > New Storage Unit**.
- 3 Enter a **Storage unit name**.

See “[NetBackup naming conventions](#)” on page 1093.
- 4 Select the **Storage unit type**. The selection specifies the type of storage that the storage unit uses: **Media Manager**, **Disk**, or **NDMP**.
- 5 For disk storage units:
 - Select a disk type from the **Disk type** drop-down menu.
The **Disk type** identifies the type of storage unit destination:

AdvancedDisk The destination is a disk pool.
storage unit

BasicDisk The destination is a path to a volume on a host.
storage unit

Cloud storage unit	<p>The destination is a disk pool of the type that includes a VendorName string. VendorName can be the name of a cloud storage provider.</p> <p>Possible values include Amazon, AT&T, and Rackspace.</p> <p>These values also can contain a _crypt suffix (for example, Amazon_crypt).</p> <p>The _crypt suffix indicates encrypted storage.</p>
NDMP storage	<p>The destination is an NDMP host. The NDMP protocol is used to perform backups and recoveries.</p>
OpenStorage storage unit	<p>The destination is a disk pool of the type that includes a VendorName string. The vendor supplies the StorageName string.</p>
PureDisk storage unit	<p>The destination is a Media Server Deduplication Pool or a PureDisk Deduplication Pool.</p> <p>See the NetBackup Deduplication Guide.</p>

- Select a media server in the **Media server** drop-down menu. The selection indicates that the media server has permission to write to the storage unit.
 - **Absolute pathname to directory** or **Absolute pathname to volume setting**.
See “[Absolute pathname to directory or absolute pathname to volume setting for storage units](#)” on page 561.
 - **Maximum concurrent jobs**
See “[Maximum concurrent jobs storage unit setting](#)” on page 565.
 - **Reduce fragment size**
See “[Reduce fragment size storage unit setting](#)” on page 572.
 - **High water mark**
See “[High water mark storage unit setting](#)” on page 563.
 - **Low water mark**
See “[Low water mark storage unit setting](#)” on page 564.
 - **Enable block sharing**
See “[Enable block sharing storage unit setting](#)” on page 563.
 - **Enable Temporary staging area**
See “[Enable temporary staging area storage unit setting](#)” on page 574.
- 6 For Media Manager storage units, data is written to tape robots and standalone tape drives:

- Select a storage device from the **Storage Device** drop-down menu.
 - Select a media server in the **Media server** drop-down menu. The selection indicates that the media server has permission to write to the storage unit.
 - **Maximum concurrent write drives**
See [“Maximum concurrent write drives storage unit setting”](#) on page 564.
 - **Enable multiplexing**
See [“Enable multiplexing storage unit setting”](#) on page 563.
 - **Reduce fragment size**
See [“Reduce fragment size storage unit setting”](#) on page 572.
- 7 Click **OK** to save the storage unit 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 1093.
- 5 Complete the fields in the **Copy Storage Unit** dialog box.
See [“About storage unit settings”](#) on page 561.

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 > Storage Units**.
- 2 In the right pane, double-click the storage unit you want to change.
To select multiple storage units, hold down the **Control** or **Shift** key.
- 3 Complete the fields on the **Change Storage Unit** dialog box.
See [“About storage unit settings”](#) on page 561.

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 969.
 - 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 970.

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 standalone tape drive, select Media Manager as the **Storage unit type**.

See [“About storage unit settings”](#) on page 561.

Figure 11-3 Media Manager storage unit settings

The screenshot shows the 'Change Storage Unit' dialog box with the following settings:

- Storage unit name:** orbiter-hcart2-robot-tld-0
- Storage unit type:** Media Manager (selected in the dropdown menu)
- On demand only:** ☐ (unchecked)
- Disk type:** (empty dropdown menu)
- Properties section:**
 - Storage device:** tld(0) - hcart2 (selected in the dropdown menu)
 - Robot type:** TLD - Tape Library DLT
 - Density:** hcart2 - 1/2 Inch Cartridge 2
 - Robot number:** 0
- Media server:** orbiter (selected in the dropdown menu)
- Maximum concurrent write drives:** 2 (selected in the dropdown menu)
- Reduce fragment size to:** ☐ (unchecked), 1048576 Megabytes
- Enable Multiplexing:** ☐ (unchecked)
- Maximum streams per drive:** 1 (selected in the dropdown menu)

At the bottom of the dialog are three buttons: **OK**, **Cancel**, and **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 standalone 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 is in use.
 - 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 564.
 - 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.
 - 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.
- Standalone drives with identical densities must be in the same storage unit. If a server contains two 1/4-inch qscsi drives, add a storage unit with **Maximum concurrent write drives** set to 2. The 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 or daemon.
- Standalone drives with different densities must be in different storage units.
- A robot and a standalone 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 11-2](#) describes the different disk types that NetBackup can use as disk media.

Table 11-2 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 Data Protection Optimization Option is licensed.</p> <p>NetBackup assumes the 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 NetBackup AdvancedDisk Storage Solutions Guide.</p>
OpenStorage	<p>An OpenStorage disk type storage unit is used for disk storage, usually provided by a third-party vendor. 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 storage is integrated into NetBackup through an API. The storage vendor partners with Symantec to integrate the storage into NetBackup.</p> <p>The storage host is the storage server. 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 NetBackup OpenStorage Solutions Guide for Disk.</p>
PureDisk	<p>A PureDisk disk type storage unit is used for deduplicated data for a Media Server Deduplication Pool. PureDisk appears as a selection when the NetBackup Data Protection Optimization Option is licensed.</p>

Not all settings are available on each disk storage unit type.

See [“About storage unit settings”](#) on page 561.

Note: Symantec recommends that you do not impose quotas 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 storage lifecycle policies 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 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 computer that hosts the storage
- A NetBackup media server

Disk pool

A collection of disk volumes that are administered as an entity. NetBackup aggregates the disk volumes into pools of storage (a disk pool) you can use for backups.

A disk pool is a storage type in NetBackup. When you create a storage unit, you select the disk type and then you select a specific disk pool.

Configuring credentials for CIFS storage and disk storage units

For Common Internet File System (CIFS) storage with AdvancedDisk and BasicDisk storage units, the following two NetBackup services on Windows computers require matching account credentials:

NetBackup Client Service	The NetBackup Client Service is either <code>bpcd.exe</code> or <code>bpinetd.exe</code> , depending on NetBackup release level. Regardless of the binary file name, the service requires the credentials.
NetBackup Remote Manager and Monitor Service	The NetBackup Remote Manager and Monitor Service binary file name is <code>nbrmms.exe</code> .

The following items describe the requirements for the account and credentials:

- Both of the services must run under the same Windows user account.
- The account must be the same account that the Windows operating system uses for read and write access to the CIFS share.
- Configure the account and the credentials on the media server or media servers that have a file system mount on the CIFS storage. Then, configure Windows so that the two aforementioned services use that account.

If account credentials are not configured properly, NetBackup marks all CIFS AdvancedDisk and BasicDisk storage units that use the UNC naming convention as DOWN.

To configure service credentials for CIFS storage and disk storage units

- ◆ In Windows, configure both the NetBackup Client Service and the NetBackup Remote Manager and Monitor Service so they meet the credential requirements. Those requirements are described previously in this document.

See your Windows operating system documentation for the procedures. How to configure Windows is beyond the scope of the NetBackup documentation.

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 backup policies are created as follows:

- A backup policy named `Policy_gold` has a gold classification. For storage, it is configured to use an SLP named `Lifecycle_Gold`, which has a gold data classification.
- A backup policy named `Policy_silver` has a silver classification. For storage, it is configured to use **Any Available**. That means it can use any available storage unit or any SLP that has a silver classification.

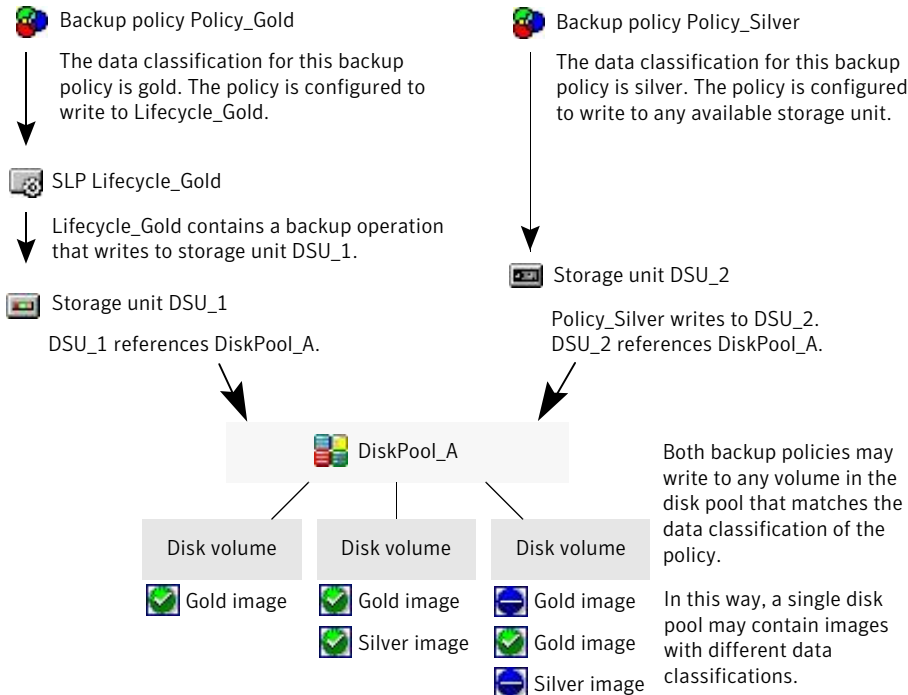
Two storage units are available to both backup policies as follows:

- `DSU_1` is an operation in `Lifecycle_Gold` and references `DiskPool_A`.

- DSU_2 is not in an SLP and 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 600.

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 563.

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:

On Windows: `install_path\netbackup\bin\goodies\nb_updatedssu`
On UNIX: `/usr/opensv/netbackup/bin/goodies/nb_updatedssu`
- Determine the potential free space.

See [“Finding potential free space on a BasicDisk disk staging storage unit”](#) on page 584.
- 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 146.

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 561.

Figure 11-5 NDMP storage unit settings

The screenshot shows the 'New Storage Unit' dialog box with the following settings:

- Storage unit name:** (Empty text field)
- Storage unit type:** NDMP (Selected in dropdown menu)
- On demand only:** (Unchecked checkbox)
- Disk type:** BasicDisk (Selected in dropdown menu)
- Properties:**
 - Storage device:** (Empty dropdown menu)
 - Robot type:** Static
 - Density:** Static
 - Robot number:** Static
- NDMP Host:** (Empty dropdown menu)
- Media Server:** <Any Available> (Selected in dropdown menu)
- Maximum concurrent write drives:** 0 (Value in text field)
- Reduce fragment size to:** (Unchecked checkbox)
- Fragment size:** 1048576 (Value in text field)
- Unit:** Megabytes

Buttons at the bottom: 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 backups to disk. 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 571.

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 1093.

See [“Low water mark storage unit setting”](#) on page 564.

Directory can exist on the root file system or system disk setting for storage units

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 status code 12 (file open failed).

Note: With this setting checked, the root file system or 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 the root file system or on a system drive.
- If the setting is not checked, and the requested directory is to be created on the root file system or on a system drive.

See [“Absolute pathname to directory or absolute pathname to volume setting for storage units”](#) on page 561.

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

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 that the OpenStorage vendor exposes appear in the list.
For PureDisk	The Media Server Deduplication Pools 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 Data Protection Optimization Option needed)
- BasicDisk
- OpenStorage (*vendor name*) (NetBackup OpenStorage Disk Option needed)
- PureDisk (NetBackup Data Protection Optimization Option needed)

Note: The **null_stu** storage unit type is available only when Symantec Support uses the NullOST plug-in to identify and isolate data transfer bottlenecks. The **null_stu** storage unit type is used for troubleshooting purposes only. Do not select **null_stu** as a storage unit type because the data that is written to a null storage unit cannot be restored.

Enable block sharing storage unit setting

The **Enable block sharing** storage unit setting allows the sharing of data blocks that have not changed from one backup to the next. 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.

Caution: For MS-SQL-Server policies, do not enable multiplexing for a schedule that is also configured to backup with multiple stripes. Restores fail when multiplexing is enabled for a schedule that uses more than one stripe.

High water mark storage unit setting

The **High water mark** storage unit setting applies to **BasicDisk** storage units and to disk pools.

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 [“Capacity managed retention type for SLP operations”](#) on page 633.

See [“Maximum concurrent jobs storage unit setting”](#) on page 565.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#).
- [NetBackup Administrator's Guide, Volume II](#).

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 [“Capacity managed retention type for SLP operations”](#) on page 633.

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:

- [NetBackup Deduplication Guide](#).
- [NetBackup Administrator's Guide, Volume II](#).

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 value for **Maximum concurrent write drives**, use the following guidelines:

- Storage unit that contains only standalone 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 standalone 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 as source.)

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 maximum number of jobs that NetBackup can run concurrently is dependent on several factors and is not exclusively regulated by this setting. These factors include the following: the risk of reaching a disk full situation, the scheduling or capacity polling overhead, the media server I/O bandwidth, and various characteristics of the disk storage. No definitive method exists that can predict when a critical limit will be exceeded in a given system.

Note: To specify a **Maximum concurrent jobs** setting of 0 disables the storage unit.

For example, 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 while 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.

- This setting depends on the available disk space and the server's ability to run multiple backup processes. Where disk pools are used, the setting also depends on the number of media servers in the storage unit.

If multiple storage units reference the same disk pool, the number of concurrent jobs that can access the pool is the sum of the **Maximum concurrent jobs** settings on all of the disk storage units. The setting applies to the storage unit and not to the disk pool. Therefore, the job load is automatically spread across the media servers that the storage unit configuration indicates.

- On Windows systems, even with multiple concurrent jobs, the time that is required for a job to complete depends on other factors:
 - The number of other jobs that are started at the same time.
 - The sequence in which the jobs were started.
 - The time that is required to complete each job.

See [“Impact when two disk storage units reference one disk pool”](#) on page 566.

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 561.

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 [“Media multiplexing \(schedule attribute\)”](#) on page 769.

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 11-3 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 11-3 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>
PureDisk (Media Server 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 NetBackup Deduplication Guide.</p>

See [“Use any available media server storage unit setting”](#) on page 575.

See [“Only use the following media servers storage unit setting”](#) on page 570.

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.

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 that are 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 11-4 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	<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>Each media server that accesses the storage 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.
PureDisk media server (Media Server Deduplication Pool)	<p>The media servers function as deduplication servers.</p> <p>NetBackup deduplication must be configured.</p>

See [“Use any available media server storage unit setting”](#) on page 575.

See [“Only use the following media servers storage unit setting”](#) on page 570.

Properties option in the Change Storage Units dialog box

Click **Properties** to display information about the volume or the disk pool, as follows:

Table 11-5 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>
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 <p>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</p> ■ Disk pool <p>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.</p> <p>NetBackup does not assign new jobs to a storage unit in which the disk pool is full. The default setting is 99%.</p>

Table 11-5 Storage Units Properties (*continued*)

Property	Description
Low water mark	The low water mark for the disk pool. Once a disk volume fills to its high water mark, NetBackup attempts to delete enough relocated images to reduce the used capacity of the disk volume down to the low water mark. The low water mark setting cannot be greater than the high water mark setting. Note: The Low water mark setting has no effect unless backups are written through a storage lifecycle policy, using the capacity-managed retention type.
Name	The name of the disk pool.
Number of volumes	The number of disk volumes in the disk pool.
% full	The percentage of storage that is currently in use on the volume. The <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.)
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.

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.)

Maximum fragment size for Media Manager storage units

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.

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.

Maximum fragment size for disk storage units

The default maximum fragment size for a disk storage unit is 524,288 megabytes. To specify a maximum fragment size other than the default, enter a value from 20 megabytes to 524,288 megabytes.

For a **Media Server Deduplication Pool**, you can enter a value from 20 megabytes to 51200 megabytes.

Backups to disk are usually fragmented to ensure that the backup does not exceed the maximum size that the file system allows.

The **Reduce fragment size** setting is intended primarily for storing large backup images on a disk type storage unit.

Note: OpenStorage vendors may have special requirements for the maximum fragment size. Consult the vendor's documentation for guidance.

Note: Basic disk staging units with different maximum fragment sizes may already be configured on a media server from a previous release. Upon upgrade, the disk storage units are not automatically increased to the new default of 524,288 megabytes. To make the best use of upgraded storage units, increase the fragment size on the upgraded storage units.

Robot number 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 574.

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 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 585.

See [“Enable temporary staging area storage unit setting”](#) on page 574.

See [“About basic disk staging”](#) on page 577.

See [“About staging backups”](#) on page 576.

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 1093.

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 555.
Media Manager	See “Media Manager storage unit considerations” on page 553.
NDMP	See “NDMP storage unit considerations” on page 559.

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 573.

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 577.

See [“Staging schedule option in Change Storage Units dialog”](#) on page 573.

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)	<p>The media servers function as deduplication servers.</p> <p>NetBackup deduplication must be configured.</p>

Staging backups

This chapter includes the following topics:

- [About staging backups](#)
- [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](#)
- [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.

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 577.</p> <p>The following storage unit types are available for basic disk staging: BasicDisk 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 or disk staging storage unit can be used in an SLP.</p> <p>See “About storage lifecycle policies” on page 600.</p>

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 689.

Creating a basic disk staging storage unit

Use the following processes to create a basic disk staging storage unit.

To create a basic disk staging storage unit

- 1 In the **NetBackup Administration Console**, select **NetBackup Management > Storage > Storage Units**.
- 2 Click **Actions > New > Storage Unit or New Storage Unit**.
- 3 In the **New Storage Unit** dialog box, name the storage unit.

See [“Storage unit name setting”](#) on page 574.

- 4 Select Disk as the **Storage unit type**.
See “[Storage unit type setting](#)” on page 574.
- 5 Select the **Disk type** of disk storage unit that is to be a disk staging storage unit: BasicDisk.
- 6 Select a media server.
See “[Media server storage unit setting](#)” on page 567.
- 7 Browse or specify 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 561.
- 8 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 561.
- 9 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 565.
- 10 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.

- 11 Check the **Enable temporary staging area** option. Once the option is enabled, create a schedule so that disk staging can occur.

Creating a schedule for a basic disk staging storage unit

The **Disk Staging Schedule** dialog box is similar to the **Add New Schedule** dialog box used to configure schedules in backup policies. The differences appear on the **Attributes** tab.

To define the disk staging schedule

- 1 Click **Staging Schedule**.
- 2 The schedule name defaults to the storage unit name.

In the **Disk Staging Schedule** or the **Disk Staging** dialog box, select the priority that the relocation jobs that are started from this schedule have compared to other types of jobs.

See [“Disk Staging Schedule dialog box”](#) on page 585.
- 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 760.
- 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
 - 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 > 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 578. |
|---|---|

See [“About staging backups”](#) on page 576.

- 3 Click the **Staging Schedule** option.
- 4 In the **Add Schedule** or **Change Schedule** 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 drop-down list. The alternate server can read a backup image originally written by a different media server.
- 7 Select **Multiple copies** and click **Configure**.
- 8 In the **Copies** field, specify the number of copies to create simultaneously. The number must be between 1 and 4.

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 761.

See [“About configuring multiple copies”](#) on page 760.

- 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 689. |
| 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 items give more information about determining disk size for a basic disk staging storage unit.

Minimum disk size

The minimum disk size is the smallest size that is required for the successful operation of the disk staging logic.

The minimum size must be greater than or equal to the largest combined size of the backups that are placed on the storage unit between runs of the disk staging schedule. (In our example, the disk images remain on the disk for one week.)

In this example, the relocation schedule runs nightly, and the largest nightly backup is 700 MB. Symantec recommends that you double this value to allow for any problems that may occur when the relocation schedule runs. To double the value gives the administrator an extra schedule cycle (one day) to correct any problems.

To determine the minimum size for the storage unit in this example, use the following formula:

Minimum size = Max data per cycle × (1 cycle + 1 cycle for safety)

For example: 1.4 GB = 700 MB × (1+1)

Average disk size

The average disk size represents a good compromise between the minimum and the maximum sizes.

In this example, the average nightly backup is 400 MB and the NetBackup administrator wants to keep the images for one week.

To determine the average size for the storage unit in this example, use the following formula:

Average size = Average data per cycle × (number of cycles to keep data + 1 cycle for safety)

2.8 GB = 400 MB × (6 + 1)

Maximum disk size

The maximum disk size is the recommended size needed to accommodate a certain level of service. In this example, the level of service is that disk images remain on disk for one week.

To determine the maximum size for the storage unit in this example, use the following formula:

Maximum size = Max data per cycle × (# of cycles to keep data + 1 cycle for safety)

For example: 4.9 GB = 700 MB × (6 + 1)

Finding potential free space on a BasicDisk disk staging storage unit

Potential free space is the amount of space on a disk staging storage unit that NetBackup could free if extra space on the volume is needed. The space is the total size of the images that are eligible for expiration plus the images ready to be deleted on the volume.

To find the potential free space on a BasicDisk storage unit, use the `bpstulist` and the `nbdevquery` commands 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 <i>BasicDisk</i> .
<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>
```

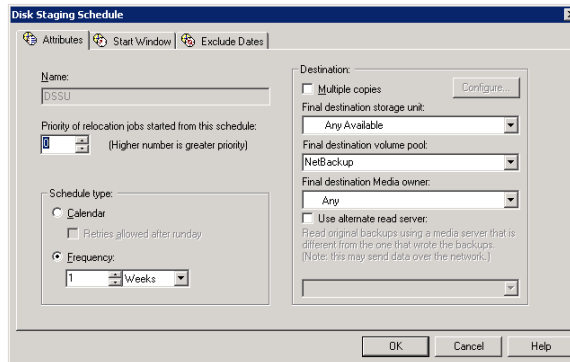
Disk Staging Schedule dialog box

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 12-3 The Attributes tab settings

Attribute	Description
Name	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 to 99999 (highest priority). The default value that is displayed is the value set for Default Job Priorities > Staging on the master server host properties.
Frequency	<p>If the backups that use a disk staging storage unit run more frequently than expected, compare the retention level 1 setting with the Frequency setting. Internally, NetBackup uses the retention level 1 setting for scheduling purposes with disk staging storage units.</p> <p>Make sure that the frequency period is set to make the backups occur more frequently than the retention level 1 setting indicates. (The default is two weeks.)</p> <p>For example, a frequency of one day and a retention level 1 of two weeks should work well.</p> <p>Retention levels are configured in the Retention Periods host properties.</p> <p>See “Retention Periods properties” on page 219.</p> <p>See “Frequency (schedule attribute)” on page 757.</p>

Table 12-3 The Attributes tab settings (*continued*)

Attribute	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.</p> <p>NetBackup continues to free space until the Low water mark is reached.</p> <p>See “Low water mark storage unit setting” on page 564.</p> <p>See “Maximum concurrent write drives storage unit setting” on page 564.</p> <p>See “About staging backups” on page 576.</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 576.</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 12-3 The Attributes tab settings (*continued*)

Attribute	Description
Use alternate read server	<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 standalone 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>

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** or **Manual Relocation to Final Destination** 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 558.

Configuring storage unit groups

This chapter includes the following topics:

- [About storage unit groups](#)
- [Creating storage unit groups for backups](#)
- [Creating storage unit groups for snapshots](#)
- [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 as the storage for a policy in the same way that you specify individual storage units. When you specify a storage unit group, the policy directs backups or snapshots only to those storage units in the designated group.

Storage unit groups can be one of the following types:

- **Backup storage unit groups**
A backup storage unit group contains only the storage units that can contain backups. Furthermore, for **Media Server Deduplication Pool** and third-party disk appliance (OpenStorage) storage, all storage units in the group must be of the same type of storage.
See [“Creating storage unit groups for backups”](#) on page 590.

- **Snapshot storage unit groups**
A snapshot storage unit group contains only the storage units that can contain snapshots. All storage units in the group must have the same **Disk type** selected. See [“Creating storage unit groups for snapshots”](#) on page 592.

NetBackup does not support storage unit groups for the following use cases:

- As a target for optimized duplication.
If you use a storage unit group as a destination for optimized duplication of deduplicated data, NetBackup uses regular duplication.
- As a source of or a target for Auto Image Replication.
- As a target for optimized synthetic backups.
If NetBackup cannot produce the optimized synthetic backup, NetBackup creates the more data-movement intensive synthetic backup.
- As a target for OpenStorage direct-to-tape operations.
See the [NetBackup OpenStorage Solutions Guide for Disk](#).

Creating storage unit groups for backups

The following procedure describes how to create a storage unit group that consists of the storage units that can contain backups.

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. The storage unit group name is case-sensitive.
See [“NetBackup naming conventions”](#) on page 1093.
- 4 For the storage unit group to contain backups, select **Backup** in the drop-down menu.
- 5 Add backup storage units to or remove backup storage units from the group:
 - To add storage units to the group, select the storage units from the **Storage units not in the group** list 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: For **Media Server Deduplication Pool** and third-party disk appliance (OpenStorage) storage, all storage units in the group must be of the same type of storage.

6 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.
 Symantec recommends that you select **Failover** for the following storage types: **AdvancedDisk**, **Media Server Deduplication Pool**, and **OpenStorage (VendorName)**.

- **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 595.

See [“Storage unit selection criteria within a group”](#) on page 594.

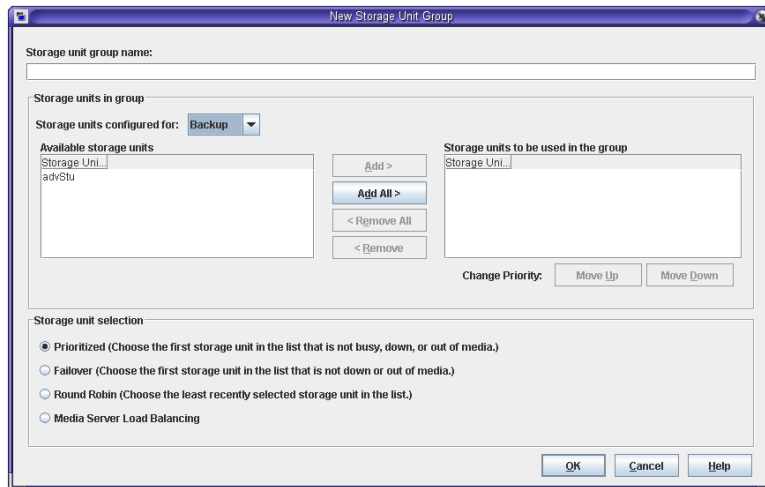
One exception to the selection criteria is in the case of a client that is also a media server with locally connected storage units.

See [“Exception to the storage unit selection criteria”](#) on page 598.

7 Click **OK**.

See [“About storage unit groups”](#) on page 589.

Figure 13-1 Backup storage unit group configuration dialog box



Creating storage unit groups for snapshots

Creating storage unit groups is optional. A snapshot storage unit group must be comprised of storage units that have matching properties.

The following procedure describes how to create a storage unit group that consists of the storage units that can contain snapshots.

To create a snapshot 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. The storage unit group name is case-sensitive.
 See [“NetBackup naming conventions”](#) on page 1093.
- 4 For the storage unit group to contain snapshots, select **Snapshot** in the drop-down menu.
- 5 A storage unit group can contain only those storage units that share similar properties. NetBackup filters the storage units for selection so that dissimilar storage units are not combined in one storage unit group.

Note: The properties of the underlying storage units are read-only. You cannot change the storage unit properties from this dialog box.

Select one or more properties to filter the storage units in the list. Only those storage units that have the selected properties are displayed. For example, select **Replication source** and **Replication target** to display only those storage units that are configured to act as both replication sources and replication targets.

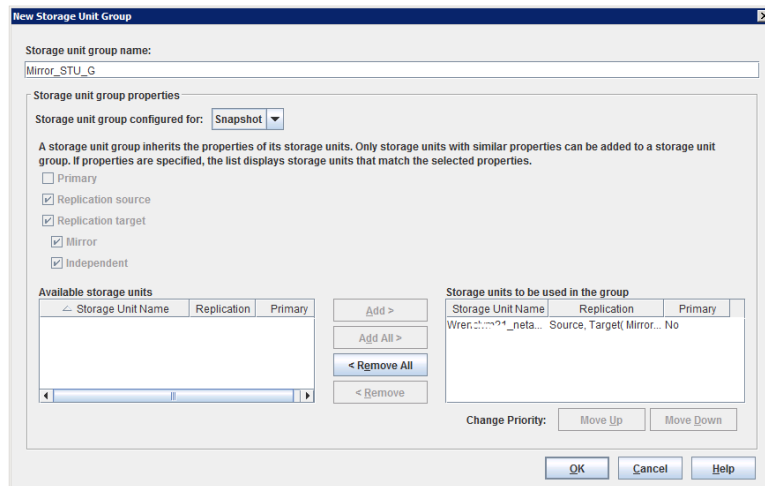
Filter the storage units on the following properties:

- **Primary**
Enable **Primary** to display the storage units that can contain the initial snapshot of primary data.
- **Replication source**
Enable **Replication source** to display the storage units that can serve as a source for a replicated snapshot.
- **Replication target**
Enable **Replication target** to display the storage units that can receive replicated snapshots from a replication source.
 - **Mirror**
Optionally, enable **Mirror** to display the storage units that can serve as a mirrored replication target. (For example, NetApp SnapMirror.)
 - **Independent**
Optionally, enable **Independent** to display the storage units that can act as either a **Mirror** replication target (SnapMirror) or a non-mirror replication target (SnapVault).

For more information about snapshot properties, see the [NetBackup Replication Director Solutions Guide](#).

- 6 Add or remove storage units from the group:
 - To add storage units to the group, select the storage units from the **Available storage units** list and click **Add**.
 - To remove storage units from the group, select the storage units from the **Storage units to be used in the 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.
- 7 Click **OK** to save and close the dialog box.

Figure 13-2 Snapshot storage unit group configuration dialog box



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 595.</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 565.

See [“Exception to the storage unit selection criteria”](#) on page 598.

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.

Note: Symantec recommends that you select **Media server load balancing** for disk staging storage units within a storage unit group.

See [“Other load balancing methods”](#) on page 596.

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 693. 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 594.
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 594.</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 693.</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 the NetBackup capability to limit bandwidth.</p> <p>See “Limit jobs per policy (policy attribute)” on page 693.</p> <p>See “Storage unit selection criteria within a group” on page 594.</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 570.

See [“Storage unit selection criteria within a group”](#) on page 594.

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 594.

Configuring storage lifecycle policies (SLPs)

- [Chapter 14. Configuring storage lifecycle policies](#)
- [Chapter 15. Storage operations](#)
- [Chapter 16. Retention types for SLP operations](#)
- [Chapter 17. Storage lifecycle policy options](#)
- [Chapter 18. Using a storage lifecycle policy to create multiple copies](#)
- [Chapter 19. Storage lifecycle policy versions](#)

Configuring storage lifecycle policies

This chapter includes the following topics:

- [About storage lifecycle policies](#)
- [Creating a storage lifecycle policy](#)
- [Deleting a storage lifecycle policy](#)
- [Lifecycle operation administration using the nbstlutil command](#)

About storage lifecycle policies

A storage lifecycle policy (SLP) is a storage plan for a set of backups. An SLP is configured within the **Storage Lifecycle Policies** utility.

An SLP contains instructions in the form of storage operations, to be applied to the data that is backed up by a backup policy. Operations are added to the SLP that determine how the data is stored, copied, replicated, and retained. NetBackup retries the copies as necessary to ensure that all copies are created.

SLPs 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.

SLPs can be set up to provide staged backup behavior. They simplify data management by applying a prescribed behavior to all the backup images that are included in the SLP. 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.

The **SLP Parameters** properties in the **NetBackup Administration Console** allow administrators to customize how SLPs are maintained and how SLP jobs run.

See “[SLP Parameters properties](#)” on page 239.

Best-practice information about SLPs appears in the following document:


<http://www.symantec.com/docs/TECH208536>

Creating a storage lifecycle policy

A storage lifecycle policy (SLP) is a storage plan for a set of backups. The operations in an SLP are the backup instructions for the data. Use the following procedure to create an SLP that contains multiple storage operations.

To add a storage operation to 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 (Windows)** or **Actions > New > Storage Lifecycle Policy (UNIX)**.



- 3 In the **New Storage Lifecycle Policy** dialog box, enter a **Storage lifecycle policy name**.


See “[NetBackup naming conventions](#)” on page 1093.

- 4 Add one or more operations to the SLP. The operations are the instructions for the SLP to follow and apply to the data that is specified in the backup policy.


If this is the first operation added to the SLP, click **Add**.

If this is not the first operation in the SLP, add an operation that is either hierarchical or non-hierarchical:

To create a hierarchical operation, select an operation to become the source of the next operation. Click **Add**. The new operation is a child of the selected operation. The child is indented under the parent operation.



To create a non-hierarchical operation, do not select an operation. A non-hierarchical operation means that the operation does not have a parent and child relationship with another operation. The new operation is not indented.



- 5 In the **Properties** tab of the **New Storage Operation** dialog box, select an **Operation** type. If you're creating a child operation, the SLP displays only those operations that are valid based on the parent operation that you've selected.

The name of the operation reflects its purpose in the SLP:

- **Backup**
See [“Backup operation in an SLP”](#) on page 608.
- **Backup From Snapshot**
See [“Backup From Snapshot operation in an SLP”](#) on page 610.
- **Duplication**
See [“Duplication operation in an SLP”](#) on page 612.
- **Import**
See [“Import operation in an SLP”](#) on page 614.

- **Index From Snapshot**
 See [“Index From Snapshot operation in an SLP”](#) on page 616.
 - **Replication**
 See [“Replication operation in an SLP”](#) on page 619.
 See [“About NetBackup Auto Image Replication”](#) on page 975.
 - **Snapshot**
 See [“Snapshot operation in an SLP”](#) on page 622.
 See [“About NetBackup Replication Director”](#) on page 1014.
- 6 Configure the properties for the operation.
 See [“Properties tab of the Storage Operation dialog box”](#) on page 643.
 - 7 The **Window** tab displays for the following operation types: **Backup From Snapshot**, **Duplication**, **Import**, **Index From Snapshot**, and **Replication**. If you'd like to control when the secondary operation runs, create a window for the operation.
 See [“Window tab of the Storage Operation dialog box”](#) on page 647.
 - 8 Click the **Advanced** button in the **Properties** tab to display options about how the window should behave if the window closes and a secondary operation is not yet complete.
 - 9 Click **OK** to create the operation.
 - 10 Add additional operations to the SLP as needed. (See step 4.)
 - 11 Change the hierarchy of the operations in the SLP if necessary.
 See [“Modifying the hierarchy of operations in a storage lifecycle policy”](#) on page 604.
 - 12 Click **OK** to create the SLP. NetBackup validates the SLP when it is first created and whenever it is changed.
 See [“Storage lifecycle policy validation dialog box”](#) on page 651.
 See [“Storage lifecycle policy Validation Report tab”](#) on page 652.
 - 13 Configure a backup policy and select a storage lifecycle policy as the **Policy storage**.
 See [“Creating a backup policy”](#) on page 673.
 See [“Policy storage \(policy attribute\)”](#) on page 685.

Modifying the hierarchy of operations in a storage lifecycle policy

In some cases, the hierarchy of operations in an SLP can be modified. Use the arrows to move it in the hierarchy.

Note: The order of the operations at the time that the SLP is saved may differ from the next time the SLP is opened. NetBackup reorders the operations while it stores them in the catalog configuration file. How the hierarchy works is not changed, however, and the parent-child relationships are preserved.

Modify the order of the operation in the operation list if needed.

- Click the arrows to move the operation into the new position.
- Up arrow
 - Swaps the position of the selected operation with the sibling above it, if one exists.
 - Using the up arrow does not change the source of the selected operation. The up arrow also moves the children of an operation and preserves their relationship with the selected operation.
 - The up arrow is disabled if no sibling appears above the selected operation.
- Down arrow
 - Swaps the position of the selected operation with the sibling below it, if one exists.
 - Using the down arrow does not change the source of the selected operation. The down arrow also moves the children of an operation and preserves their relationship with the selected operation.
 - The down is disabled if no sibling appears below the selected operation.
- Right arrow
 - Moves the operation right in the hierarchy, making the sibling above the operation the source for the operation.
 - If no sibling exists above the operation in the hierarchy, the right arrow is disabled. It is always disabled for **Backup** and **Snapshot** operations.
 - Moving the operation to the right does not change the position number of the operation in the list.
 - The right arrow also moves the children of the operation and preserves their relationship with the selected operation.
- Left arrow
 - Moves the operation to the left in the hierarchy, turning the parent into a sibling.
 - The left arrow is enabled for some operations. For the left arrow to be enabled, the selected operation must be either the first or last in a list of siblings.

If the operation 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 operation to preserve the relationship with the operation.

The left arrow is disabled for **Backup** and **Snapshot** operations.

Deleting a storage lifecycle policy

To delete a storage lifecycle policy, use the following procedure. Note that to delete an SLP deletes all versions of the SLP.

To delete a storage lifecycle policy

- 1 Remove the SLP from all backup policies to prevent new backup jobs from writing to the SLP.
- 2 Wait for all in-process backup jobs to the SLP to complete, or cancel the jobs using the **Activity Monitor** or the command line.
- 3 To prevent new jobs or cancel any existing duplication jobs from writing to the SLP, run the following command:

```
nbstlutil cancel -lifecycle name
```

- 4 Use the **Activity Monitor** to cancel in-process jobs that use the SLP.
- 5 Once all of the operations are complete, delete the SLP using one of the following methods:

- The **NetBackup Administration Console**
 - Expand **Storage > Storage Lifecycle Policies**.
 - Select the SLP name.
 - Select **Edit > Delete**.
 - In the **Delete Storage Lifecycle Policies** dialog box, select the SLP name and click **OK**.
- The `nbstl` command

```
nbstl storage_lifecycle_name -delete
```

If the administrator tries to delete an SLP with active images, status code 1519 appears (Images are in process). Wait several minutes and try to delete the SLP again until the error no longer appears.

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.

Lifecycle operation administration using the nbstlutil command

The NetBackup storage lifecycle policy utility command (`nbstlutil`) gives administrators the ability to intervene between pending SLP operations. Specifically, the `nbstlutil` command can be used to cancel, inactivate, or activate the processing of existing SLP-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 14-1 `nbstlutil` details

nbstlutil information	Details
Where to find	<p>The command is found in the following location:</p> <p>On Windows:</p> <pre>install_path\NetBackup\bin\admincmd\nbstlutil</pre> <p>On UNIX:</p> <pre>/usr/opensv/netbackup/bin/admincmd/nbstlutil</pre>
How to use	<p>Use <code>nbstlutil</code> to perform the following administrative actions:</p> <ul style="list-style-type: none"> List the status of SLP-managed images. The EMM table that tracks the status of SLP-processed images can be printed. Support may request this information to troubleshoot an SLP 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 SLP complete. It does not attempt to create any more copies of the backup image. Deactivate (suspend) pending and future SLP 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 SLP operations on selected images or image copies. <p>See the NetBackup Commands Reference Guide for a description of all the options available for <code>nbstlutil</code>.</p>

Table 14-1 nbstlutil details (*continued*)

nbstlutil information	Details
When to use	<p>NetBackup starts a duplication session every five minutes to copy data from a backup operation for a duplication operation. Five minutes is the default frequency of the Image processing interval parameter in the SLP Parameters host properties. 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 is the default frequency of the Extended image retry interval parameter in the SLP Parameters host properties.</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 SLP while the library is repaired. When ready, the SLP can be activated and duplication jobs can begin.</p> <p>Note: Once the job is reactivated, the administrator may want to temporarily change the Extended image retry interval parameter to one hour to begin duplication jobs sooner.</p>

Storage operations

This chapter includes the following topics:

- [Operation types in a storage lifecycle policy](#)
- [Backup operation in an SLP](#)
- [Backup From Snapshot operation in an SLP](#)
- [Duplication operation in an SLP](#)
- [Import operation in an SLP](#)
- [Index From Snapshot operation in an SLP](#)
- [Replication operation in an SLP](#)
- [Snapshot operation in an SLP](#)
- [Creating a hierarchy of storage operations in a storage lifecycle policy](#)


Operation types in a storage lifecycle policy

The **Operation** selections are the instructions in the storage lifecycle policy. The following topics describe the purpose of each operation.

Backup operation in an SLP

Use the **Backup** operation in a storage lifecycle policy to create a backup. All **Backup** operations in a single storage lifecycle policy must be on the same media server.

A **Backup** operation creates a tar-formatted image. To create a snapshot image, select a **Snapshot** operation.


Figure 15-1 Backup operation in the New Storage Operation dialog box**Table 15-1** Backup operation characteristics

Characteristic	Description
Storage unit selection	<p>The selection must be a backup storage unit or a backup storage unit group (see the following note).</p> <p>The selection cannot be a snapshot storage unit or a snapshot storage unit group.</p> <p>Note: If creating an SLP for Auto Image Replication, do not select a storage unit group. Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.</p> <p>See “About NetBackup Auto Image Replication” on page 975.</p>
Child of	<p>A Backup operation cannot serve as the child of any other operation. Therefore, do not click on any other operation in the SLP when adding a Backup operation.</p> <p>See “Creating a storage lifecycle policy” on page 601.</p>
Source for	<p>A Backup operation can be the source of a Duplication operation. (See Figure 15-2.)</p>
Hierarchy notes	<p>If a Backup operation appears in an SLP, it must be the first operation.</p> <p>An SLP can contain a maximum of four Backup operations.</p>

Table 15-1 Backup operation characteristics (continued)

Characteristic	Description
Job type	A Backup operation generates a Backup job in the Activity Monitor .
Window	Backup operations do not offer the option to create an SLP window. See “Window tab of the Storage Operation dialog box” on page 647.

Figure 15-2 SLP that contains a Backup operation



Backup From Snapshot operation in an SLP

Use the **Backup From Snapshot** operation to create a tar-formatted copy of the snapshot. The new copy is a backup copy. The process is sometimes referred to as a *snapdupe* job.

Figure 15-3 Backup From Snapshot operation in the New Storage Operation dialog box


The screenshot shows the 'New Operation' dialog box with the 'Properties' tab selected. The 'Source storage' is 'Wrenchiv..._netapp514004_vault (Replication)'. The 'Operation' is 'Backup From Snapshot'. The 'Destination Storage Attributes' section includes 'Destination storage:', 'Volume pool:' (set to 'NetBackup'), and 'Media owner:' (set to 'Any'). The 'Retention' section includes 'Retention type:' (set to 'Fixed') and 'Retention period:' (set to '2 weeks'). The 'Duplication' section includes 'Alternate read server:' and a checkbox for 'Preserve multiplexing'. A checkbox at the bottom reads 'Postpone creation of this copy until the source copy is about to expire'. An 'Advanced' button is in the bottom right.

Table 15-2 Backup From Snapshot operation characteristics

Characteristic	Description
Storage unit selection	<p>The selection must be a backup storage unit or a backup storage unit group.</p> <p>The selection cannot be a snapshot storage unit or a snapshot storage unit group.</p>
Child of	<p>A Backup From Snapshot operation must use a Snapshot operation as its source.</p> <p>Therefore, click on the Snapshot operation in the SLP when adding a Backup From Snapshot operation.</p> <p>See “Creating a storage lifecycle policy” on page 601.</p>
Source for	<p>A Backup From Snapshot operation can be the source for a Duplication operation. (See Figure 15-4.)</p>
Hierarchy notes	<p>An SLP may contain more than one Backup From Snapshot operation. If the first Backup From Snapshot operation fails with an unrecoverable error, NetBackup does not attempt the second one.</p> <p>Note: The SLP may contain only one Backup From Snapshot operation if the SLP is used by an NDMP policy (or a Standard or MS-Windows policy with NDMP Data Mover enabled).</p>

Table 15-2 Backup From Snapshot operation characteristics (*continued*)

Characteristic	Description
Job type	A Backup From Snapshot operation generates a Backup job in the Activity Monitor .
Window	An SLP window can be created for a Backup From Snapshot operation. See “Window tab of the Storage Operation dialog box” on page 647.

Figure 15-4 SLP that contains a Backup From Snapshot operation

Duplication operation in an SLP

Use the **Duplication** operation to create a copy of a **Backup**, a **Backup from Snapshot**, or another **Duplication** operation. A media server performs the operation and writes the copy.

Note: Use the **Replication** operation to create a copy of a **Snapshot** operation.

See [“Replication operation in an SLP”](#) on page 619.

Figure 15-5 Duplication operation in the New Storage Operation dialog box

The screenshot shows the 'New Storage Operation' dialog box. It has two tabs: 'Properties' and 'Window'. The 'Properties' tab is active. It contains several sections: 'Source' (STU_Adv (Backup)), 'Operation' (Duplication), 'Storage' (Storage unit: STU_Adv, Volume pool: NetBackup, Media owner: Any), 'Retention' (Retention type: Fixed, Retention period: 2 weeks (level 1)), and 'Duplication' (Alternate read server: empty, Preserve multiplexing checkbox). At the bottom left, there is a checkbox labeled 'Postpone creation of this copy until the source copy is about to expire.' with an 'Advanced' button next to it. At the bottom right, there are 'OK', 'Cancel', and 'Help' buttons.


Table 15-3 Duplication operation characteristics

Characteristic	Description
Storage unit selection	<p>The selection must be a backup storage unit or a backup storage unit group.</p> <p>The selection cannot be a snapshot storage unit or a snapshot storage unit group.</p>
Child of	<p>A Duplication operation can be the child of the following operations:</p> <ul style="list-style-type: none">■ Backup operation■ Backup From Snapshot operation■ A Duplication operation <p>Therefore, click on one of these operations in the SLP when adding a Duplication operation.</p> <p>See “Creating a storage lifecycle policy” on page 601.</p>
Source for	<p>A Duplication operation can be the source for a Duplication operation. (See Figure 15-6.)</p>
Hierarchy notes	<p>When a Duplication operation appears in an SLP, it cannot be the first operation.</p>

Table 15-3 Duplication operation characteristics (continued)


Characteristic	Description
Job type	A Duplication operation generates a Duplication job in the Activity Monitor .
Window	An SLP window can be created for a Duplication operation. See “Window tab of the Storage Operation dialog box” on page 647.

Figure 15-6 SLP that contains one Backup operation and two Duplication operations



Import operation in an SLP

Use the **Import** operation as part of Auto Image Replication. An **Import** operation in an SLP indicates that the SLP is to automatically import images into the target master domain. An SLP that contains an **Import** operation is referred to as an Import SLP.

Figure 15-7 Import operation in the New Storage Operation dialog box**Table 15-4** Import operation characteristics

Characteristic	Description
Storage unit selection	An Import operation can import only from a backup storage unit or a backup storage unit group. It cannot import from a snapshot storage unit or a snapshot storage unit group.
Child of	An Import operation cannot serve as the child of any other operation. Therefore, do not click on any other operation in the SLP when adding a Import operation. See “Creating a storage lifecycle policy” on page 601.
Source for	An Import operation can be the source of a Duplication operation. At least one operation in the SLP must use the Target retention retention type.
Hierarchy notes	If an SLP contains an Import operation, it must be the first in the operations list and the only Import operation.
Job type	An Import operation generates an Import job in the Activity Monitor .
Window	An SLP window can be created for an Import operation. See “Window tab of the Storage Operation dialog box” on page 647.

The **Override job priority** option can be selected. It allows administrators to specify a job priority for any import jobs which use this SLP.

Index From Snapshot operation in an SLP

The **Index From Snapshot** operation indexes the contents of existing snapshots. When NetBackup indexes a snapshot, it creates an image .f file in the NetBackup catalog for each snapshot. The presence of an image .f file assists the user when a file needs to be restored from the snapshot, as described in [Table 15-5](#).

The **Backup From Snapshot** operation also creates an image .f file. An **Index from Snapshot** may not be required if a **Backup From Snapshot** occurs frequently enough for the restore needs in your environment. For example, if the **Backup From Snapshot** runs once per week but file restores are required daily, consider using the **Index from Snapshot**.

The actual restore of the snapshot requires that the snapshot be mounted, regardless of whether an **Index from Snapshot** has been performed or not.

Table 15-5 Restore operations

Type of restore	Where performed	Description	Requirements
Live browse restore	<ul style="list-style-type: none">NetBackup Backup, Archive, and Restore interfaceNetBackup OpsCenter Console	The user navigates the directory structure to locate and select the files for restore.	No .f file needs to be present in the NetBackup catalog. During a live browse restore, NetBackup automatically mounts the snapshot so that the user can see what files it contains. Mounting and unmounting the snapshot can be time-consuming.
Search for restore (Also referred to as a Simplified File Restore or an Operational Restore)	NetBackup OpsCenter Console	<ul style="list-style-type: none">The user can search for files based on partial file name. Knowing the name of the client or of the policy that created the backup is not necessary.Offers several filters to narrow search.Search results display all the recovery points available to the user for both backups and snapshots.	The .f file must be present in the NetBackup catalog.

Table 15-6 Index From Snapshot operation characteristics

Characteristic	Description
Storage unit selection	The Index From Snapshot operation does not write data to a storage unit. However, a storage unit selection is needed to select the media server that to be used to access the snapshot. As a best practice, use the storage unit from the Snapshot or Replication operation that is the source for this operation.
Child of	When an Index From Snapshot operation appears in an SLP, it must be the child of a Snapshot or Replication operation. Therefore, click on either a Snapshot or a Replication operation in the SLP when adding an Index From Snapshot operation. See “Creating a storage lifecycle policy” on page 601.
Source for	While an Index From Snapshot operation cannot be the source for any operation, a Replication operation can follow it.
Hierarchy notes	The Index From Snapshot operation can consume system resources and requires that each snapshot be mounted to create the .f file. See “Determining where and when the Index From Snapshot operation occurs” on page 618.
Job type	An Index From Snapshot operation generates an Index From Snapshot job in the Activity Monitor .
Window	An SLP window can be created for an Index From Snapshot operation. See “Window tab of the Storage Operation dialog box” on page 647.

Consider the following items before using the **Index From Snapshot** operation:

- The **Index From Snapshot** operation is supported only in a Replication Director configuration.
- **Standard**, **MS-Windows**, **NDMP**, and **VMware** backup policy types support the use of storage lifecycle policies that contain the **Index From Snapshot** operation.

Note: However, a **Standard** or **MS-Windows** policy with **NDMP Data Mover** enabled is not supported.

- The **Index From Snapshot** operation can run from a full or an incremental schedule. The file entries that are added to the .f file for either schedule are the full set of files since all files can be restored from that snapshot. To do so allows for the most efficient restore, however, more space is consumed in the NetBackup catalog by the .f file.

Determining where and when the Index From Snapshot operation occurs

Including the **Index From Snapshot** operation requires some consideration as the operation can consume system resources and require additional time to perform. For example, to perform the operation can require that a snapshot be mounted or that NetBackup gather content details from the file system to populate the catalog.

To help mitigate the extra resource and time that the operation may take, the system administrator can control when and where the **Index From Snapshot** operation runs:

- Use the storage lifecycle policy **Window** tab to schedule when the **Index From Snapshot** operation can run. Schedule the operation to run when it is least likely to interfere with other jobs.
See [“Window tab of the Storage Operation dialog box”](#) on page 647.
- Use the following points to determine where to position the **Index From Snapshot** operation in the SLP operations list:
 - Each NetBackup environment needs to determine where the operation works best in a specific SLP. To place the **Index From Snapshot** operation too early (toward the top of the operations list), may consume time when the restore capabilities are not needed. To place the operation toward the end of the operations list may cause the administrator to delay a restore until earlier snapshots or replications complete.
 - Use the **Index From Snapshot** operation in an SLP only once. A restore can be performed from any snapshot after one image .*ε* file is created.
 - Any operations list that includes a **Backup From Snapshot** operation does not need an **Index From Snapshot** operation. The **Backup From Snapshot** operation creates an image .*ε* file. The only exception is if the index is needed for restores before the **Backup From Snapshot** operation occurs.
 - An **Index From Snapshot** operation cannot have any dependents. An SLP cannot validate an **Index From Snapshot** operation with children. [Figure 15-8](#) shows an SLP with a valid configuration.
[Figure 15-9](#) is also a valid configuration. A **Replication** operation follows the **Index From Snapshot** operation, but it is not indented. The **Replication** operation is a child of the **Snapshot** operation, not a child of the **Index From Snapshot** operation.
To add a **Replication** operation after an **Index From Snapshot** operation, click on the **Snapshot** operation, and then click **Add**.

Figure 15-8 Example 1 of a valid placement of the Index From Snapshot operation





Figure 15-9 Example 2 of a valid placement of the Index From Snapshot operation



Replication operation in an SLP

Use the **Replication** operation for the following types of replication:

- NetBackup Replication Director to replicate a snapshot, as shown in [Figure 15-10](#).
See [“About NetBackup Replication Director”](#) on page 1014.
- NetBackup Auto Image Replication to replicate a backup, as shown in [Figure 15-11](#).
See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 988.

Table 15-7 Replication operation characteristics

Characteristic	Description
Storage unit selection	<p>Under Destination storage attributes:</p> <ul style="list-style-type: none">■ For Replication Director, select the Storage that is configured to contain replicated snapshots.■ For Auto Image Replication, choose to either:<ul style="list-style-type: none">■ Replicate the backup to storage servers in all target NetBackup domains.■ Replicate the backup to a specific master server in a specific domain. This type of Auto Image Replication is known as Targeted A.I.R.
Child of	<p>Click on the appropriate operation when adding a Replication operation.</p> <p>A Replication operation can be the child of any of the following operations:</p> <ul style="list-style-type: none">■ Snapshot operation for NetBackup Replication Director to replicate a snapshot.■ Another Replication operation.■ Backup operation for NetBackup Auto Image Replication. <p>See “Creating a storage lifecycle policy” on page 601.</p>
Source for	<p>A Replication operation can be the source for the following operations:</p> <ul style="list-style-type: none">■ Replication■ Backup From Snapshot <p>See “Backup From Snapshot operation in an SLP” on page 610.</p>
Job type	<p>A Replication operation generates a Replication job in the Activity Monitor.</p>
Window	<p>An SLP window can be created for a Replication operation.</p> <p>See “Window tab of the Storage Operation dialog box” on page 647.</p>

Figure 15-10 Replication operation following a Snapshot operation

The 'Change Storage Operation' dialog box is shown with the 'Properties' tab selected. The 'Source storage' is 'wrencvm21_primary_stu (Snapshot)' and the 'Operation' is 'Replication'. The 'Destination Storage Attributes' section has 'Storage' set to 'wrencvm21_vault_stu'. The 'Retention' section has 'Retention type' set to 'Fixed' and 'Retention period' set to '2 weeks (level 1)'. There is a checkbox for 'Postpone creation of this copy until the source copy is about to expire.' and an 'Advanced' button. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Figure 15-11 Replication operation following a Backup operation

The 'New Storage Operation' dialog box is shown with the 'Properties' tab selected. The 'Source storage' is 'stu-dp-advdisk (Backup)' and the 'Operation' is 'Replication'. The 'Destination Storage Attributes' section has 'Send the backups to:' with two radio buttons: 'All replication target storage servers (across different NetBackup domains) for the source storage server, cayce' (selected) and 'A specific Master Server:'. Below these are 'Target master server:' and 'Target import SLP:' fields. The 'Retention' section has 'Retention type' set to 'Fixed' and 'Retention period' set to '2 weeks (level 1)'. There is a checkbox for 'Postpone creation of this copy until the source copy is about to expire.' and an 'Advanced' button. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Snapshot operation in an SLP

A **Snapshot** operation creates a point-in-time, read-only, disk-based copy of data. NetBackup provides several types of snapshots, depending on the device where the snapshot occurs.

Use a **Snapshot** operation as the first operation in a storage lifecycle policy for a NetBackup Replication Director configuration.

Figure 15-12 Snapshot operation in the New Storage Operation dialog box




Table 15-8 Snapshot operation characteristics

Characteristic	Description
Storage unit selection	<p>The following topics describe the types of snapshot storage units that can be used as the storage for a snapshot operation:</p> <ul style="list-style-type: none"> ■ See “Primary snapshot storage unit” on page 626. ■ See “Primary + Replication source snapshot storage unit” on page 626. ■ See “Replication source + Replication target snapshot storage unit” on page 627. ■ See “Replication target snapshot storage unit” on page 627. ■ See “Replication source + Replication target + Mirror snapshot storage unit” on page 628. ■ See “Replication source + Replication target + Mirror snapshot storage unit” on page 628. <p>A Storage unit selection is necessary in the following situations:</p> <ul style="list-style-type: none"> ■ If the Snapshot is to be used by a subsequent Replication operation. The storage unit that is specified for the Snapshot operation must be a snapshot-capable storage unit that represents the primary storage. ■ If the SLP contains only one operation and that is a Snapshot operation, specify a storage unit. NetBackup uses that storage unit to determine which media server to use to launch the snapshot job. <p>If neither situation applies to the SLP, the administrator may select No storage unit or may simply make no selection. NetBackup uses the storage unit that is selected for the Backup From Snapshot operation.</p>
Child of	<p>A Snapshot operation cannot serve as the child of any other operation. Therefore, do not click on any other operation in the SLP when adding a Snapshot operation.</p> <p>See “Creating a storage lifecycle policy” on page 601.</p>
Source for	<p>A Snapshot operation can be the source for the following operations:</p> <ul style="list-style-type: none"> ■ Backup From Snapshot ■ Index From Snapshot ■ Replication operation
Hierarchy notes	<p>If a Snapshot operation appears in an SLP, it must be first in the operations list.</p>
Job type	<p>A Snapshot operation generates a Snapshot job in the Activity Monitor.</p>
Window	<p>Snapshot operations do not offer the option to create an SLP window.</p> <p>See “Window tab of the Storage Operation dialog box” on page 647.</p>

See [“About configuring storage lifecycle policies for snapshots and snapshot replication”](#) on page 624.

About configuring storage lifecycle policies for snapshots and snapshot replication

A storage lifecycle policy (SLP) contains instructions in the form of storage operations, to be applied to the data. Operations are added to the SLP that determine how the data is stored and copied or replicated. For example, the NetBackup administrator creates an operation that determines where the data exists as a snapshot, as a replication, or as a duplication. The administrator also determines the retention of the data at each storage unit or storage unit group.

Storage operations that are defined in the SLP use storage units that represent disk pools. An SLP that is configured for snapshots or snapshot replication must contain a specific, hierarchical combination of operations.

The following figure represents an SLP for a replication scenario. In the example, the following operations are used:

- A **Snapshot** operation creates a snapshot.
- A **Replication** operation replicates the snapshot to another volume.
- A **Backup From Snapshot** operation creates a tar-formatted backup from the snapshot.
- A **Duplication** operation copies the backup to tape.

[Table 15-9](#) describes the four types of operations that are required in this example replication scenario.




Table 15-9 Example of a storage lifecycle policy configured for snapshots and snapshot replication

Operation order in SLP	Operation	Description
1	Snapshot	<p>Operation 1 creates a snapshot in the primary storage. The snapshot serves as the source for the other operations in the SLP.</p> <ul style="list-style-type: none">■ The operation must be a Snapshot operation.■ The storage must be a snapshot storage unit that has the following properties set: Primary and Replication source. <p>Note: Only one operation to a Primary storage unit is permitted in an SLP.</p>
2 (Child to operation 1)	Replication	<p>Operation 2 replicates the snapshot that the first operation created.</p> <ul style="list-style-type: none">■ The operation must be a Replication operation and it must be the replication partner to the source storage unit.■ The retention type determines what replication method is used. For example, if retention type Fixed or Expire after copy is selected for a NetApp storage system, the SnapVault replication method is used.■ The storage must be a snapshot storage unit that has the Replication target property set. Since no other replica is created from this operation in this example, it does not need to have the Replication source property set.
3 (Child to operation 2)	Backup From Snapshot	<p>Operation 3 creates a tar-formatted backup copy of the snapshot.</p> <ul style="list-style-type: none">■ The operation must be a Backup From Snapshot operation. This operation creates a backup image from the snapshot.■ The storage must be a backup storage unit.
4 (Child to operation 3)	Duplication	<p>Operation 4 makes a duplicate copy from the tar backup copy. In this example, the copy is duplicated to tape media.</p> <ul style="list-style-type: none">■ The operation must be a Duplication operation. This operation creates a backup copy of the tar-formatted image.■ The storage must be a backup storage unit.

After the SLP is configured for different operations, the NetBackup administrator configures a backup policy that points to the snapshot SLP.

The **SLP Parameters** host properties in the **NetBackup Administration Console** allow administrators to customize how SLPs are maintained and how SLP jobs run.

See “[SLP Parameters properties](#)” on page 239.

Best-practice information about SLPs appears in the following document:


<http://www.symantec.com/docs/TECH208536>

Primary snapshot storage unit

A snapshot operation can use a **Primary** snapshot storage unit. That is, the storage unit represents a disk pool that contains the volumes that have only the **Primary** property set.

Figure 15-13 shows an SLP that contains one primary-only **Snapshot** operation, one **Backup From Snapshot** operation, and one **Duplication** operation. The **Backup From Snapshot** operation is used to create a backup from the snapshot on the primary-only **Snapshot** operation. After the backup is created, it is duplicated to a **Duplication** operation.


Figure 15-13 SLP that contains a Snapshot operation, a Backup From Snapshot operation, and a Duplication operation



Primary + Replication source snapshot storage unit

An SLP operation can use a **Primary + Replication source** snapshot storage unit. That is, the storage unit represents a disk pool that contains volumes that have both the **Primary** property and the **Replication source** property set.

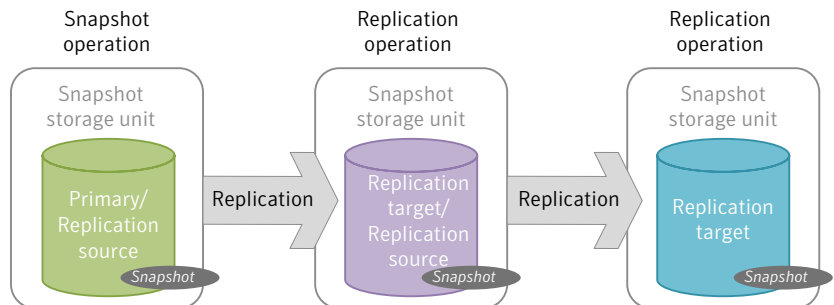
Figure 15-14 shows an SLP that contains a **Primary + Replication source** snapshot storage unit as one operation and one **Replication target** snapshot storage unit as another operation. The **Primary + Replication source** storage unit can replicate to the **Replication target** storage unit.

Figure 15-14 SLP that contains a Snapshot operation and a Replication operation

Replication source + Replication target snapshot storage unit

An SLP operation can use a snapshot storage unit that represents a disk pool that contains volumes that have the following properties: **Replication source** and **Replication target**.

A snapshot storage unit with these properties can serve as both the **Replication source** for another operation in the SLP, and as the **Replication target** for another operation in the SLP.

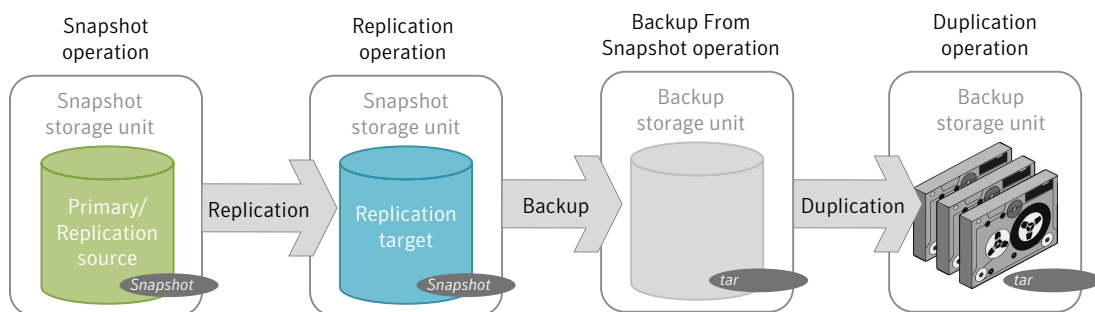
Figure 15-15 SLP that contains a Snapshot operation and two Replication operations

Replication target snapshot storage unit

An SLP operation can use a snapshot storage unit that represents a disk pool that contains volumes that have only the **Replication target** property set.

An operation with this property can serve only as a **Replication target** for another operation in the SLP. It cannot serve as source for a replica, but it can serve as the source for a **Duplication** operation.

Figure 15-16 SLP that contains a Snapshot operation, a Replication operation, a Backup From Snapshot operation, and a Duplication operation



Replication source + Replication target + Mirror snapshot storage unit

An SLP can use a snapshot storage unit that represents a disk pool that contains volumes that have the following properties: **Replication source**, **Replication target**, and **Mirror**.

An operation with these properties can serve as both:

- A **Replication source** in a cascading configuration.
- A mirrored **Replication target** in a cascading configuration. A mirrored **Replication target** must have a forced **Mirror** retention type.

Replication target + Mirror snapshot storage unit

An SLP can use a snapshot storage unit that represented a disk pool that contains volumes that have the following properties: **Replication target** and **Mirror**.

A mirrored **Replication target** must have a forced **Mirror** retention type.

Creating a hierarchy of storage operations in a storage lifecycle policy


The list of operations in the storage lifecycle policy acts as a list of instructions to NetBackup about the data that the SLP protects. In some cases, one operation

depends on another operation. For example, a snapshot may serve as the source for a replication. Or, a backup may serve as the source of a duplication.

The operation hierarchy in the **Storage Lifecycle Policy** dialog box represents a parent and child relationship.


Figure 15-17 is an example of how the operation list uses indentation to indicate the relationship between a parent operation and a child operation.

Figure 15-17 Hierarchical storage operations in an SLP



One copy can be the source for many other copies. Figure 15-18 shows how after the first copy is created, all subsequent copies can be made locally from that source, without tying up network resources.

Figure 15-18 Hierarchical operations



Changing the location of an operation in the hierarchy changes the storage unit that serves as the source for the subsequent copies. Changing the hierarchy cannot change the operation type. (For example, change a backup operation into a duplication operation.)

Retention types for SLP operations

This chapter includes the following topics:

- [Retention types for storage lifecycle policy operations](#)
- [Capacity managed retention type for SLP operations](#)
- [Expire after copy retention type for SLP operations](#)
- [Fixed retention type for SLP operations](#)
- [Maximum snapshot limit retention type for SLP operations](#)
- [Mirror retention type for SLP operations](#)
- [Target retention type for SLP operations](#)

Retention types for storage lifecycle policy operations

The **Retention type** for an operation in a storage lifecycle policy determines how long the data is kept on that storage media.

[Table 16-1](#) describes which retention types are valid selections for the various operations.

Table 16-1 Operation and retention type configurations

Retention type	Backup operation	Snapshot operation	Replication operation	Backup From Snapshot operation	Duplication operation
Fixed	Valid	Valid	Valid	Valid	Valid

Table 16-1 Operation and retention type configurations (*continued*)

Retention type	Backup operation	Snapshot operation	Replication operation	Backup From Snapshot operation	Duplication operation
Expire after copy	Valid	Valid	Valid	Valid	Valid
Maximum Snapshot limit	Invalid	Valid; SLP honors the policy setting.	Invalid	Invalid	Invalid
Mirror	Invalid	Invalid	Valid for snapshot storage only	Invalid	Valid for snapshot storage only
Target retention	Invalid	Invalid	Valid if the first operation in the SLP is an Import and if the storage is of the backup type.	Invalid	Valid if the first operation in the SLP is an Import.
Capacity managed	Valid; AdvancedDisk default; set on the storage server.	Invalid	Invalid	Invalid	Valid; AdvancedDisk default; set on the storage server.

Note: Retention is not associated with the **Index From Snapshot** operation because the operation does not create any copy.

Mixing retention types

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 SLP operations or selecting the storage location for a policy:

- All SLPs 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 an SLP and to the same volume (by the storage unit) directly from a policy.

- Mark all disk storage units that are used with SLPs 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.

Capacity managed retention type for SLP operations

A **Capacity managed** operation means that NetBackup automatically manages the space on the storage, based on the **High water mark** setting for each volume. **Capacity managed** 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 563.

See “[Low water mark storage unit setting](#)” on page 564.

An image copy with a **Capacity Managed** retention is not eligible for expiration until its dependent copies have been created.

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 for all operations in a storage lifecycle policy. If the operating system time is past the date that matches the longest retention period for an image, the image is eligible for deletion.

To see exactly when the storage 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 capacity managed storage 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 longer than the **Desired cache period** indicates.

Rules and recommendations for using the Capacity Managed retention type

Use the following recommendations and rules when configuring storage operations or when selecting the storage location for a policy:

- 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.
- All SLPs 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 storage lifecycle policy and to the same volume (by the storage unit) directly from a policy.
- Mark all disk storage units that are used with SLPs 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.

Capacity managed retention type and disk types that support SIS

Capacity managed is selectable for any disk storage unit that is allowed in an SLP. However, for the disk types that support single-instance storage (SIS), **Capacity managed** functions to various degrees. In order for **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:

- **Media Server Deduplication Pool** storage units
- Some OpenStorage storage units, depending on the vendor characteristics.

Expire after copy retention type for SLP operations

The **Expire after copy** retention indicates that after all direct (child) copies of an image are successfully duplicated to other storage, the data on this storage is

expired. The last operation in the SLP cannot use the **Expire after copy** retention type because no subsequent copy is configured. Therefore, an operation with this retention type must have a child.

Symantec recommends that you not enable **Expire after copy** retention for any storage units that are to be used with SLPs with either of the following: Accelerator or synthetic backups. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.

Note: Although synthetic backups do support the use of storage lifecycle policies, SLPs cannot be used for the multiple copy synthetic backups method.

See [“Using the multiple copy synthetic backups method”](#) on page 860.

If a policy is configured to use an SLP for the backup, the retention that is indicated in the SLP is the value that is used. The **Retention** attribute in the schedule is not used.

An image copy with an **Expire after copy** retention is expired as soon as all of its direct child copies have been successfully created. Any mirrored children must also be eligible for expiration.

Fixed retention type for SLP operations

The **Fixed** retention indicates that the data on the storage is retained for the specified length of time, after which the backups or snapshots are expired.

An image copy with a **Fixed** retention is eligible for expiration when all of the following criteria are met:

- The **Fixed** retention period for the copy has expired.
- All child copies have been created.
- All child copies that are mirror copies are eligible for expiration.

The **Fixed** retention period is always marked from the original backup time of the image. For example, if a tape device is down, causing a 2-day delay in creating a duplicate tape copy, the expiration time of the duplicate copy is not different due to the 2-day delay. The expiration time of the duplicate copy is still x days from the time that the original backup was completed. It does not matter when the copy was created.

Maximum snapshot limit retention type for SLP operations

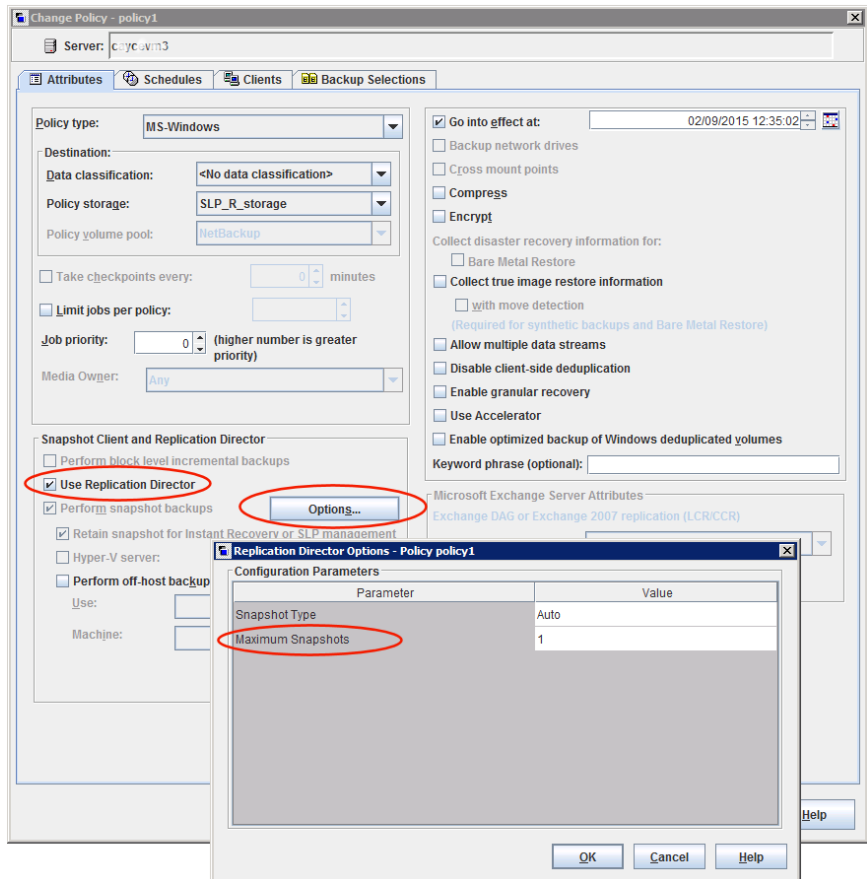
The **Maximum snapshot limit** determines the maximum number of snapshots that can be stored for a particular policy and client pair.

When the maximum is reached, the next snapshot causes the oldest job-complete snapshot to be deleted. A snapshot job is considered to be complete once all of its configured dependent copies are complete. (Dependent copies are created as a result of **Backup From Snapshot**, **Index From Snapshot**, or **Replication** operations.) The practice is referred to as *rotation*. This retention type applies only to snapshots, and not to backups.

For example, Policy P1 contains two clients: C1 and C2. After the policy runs four times, it creates four snapshot images for C1 and four images for C2. If the **Maximum snapshot limit** is set to four, when the policy runs for the fifth time, NetBackup deletes the first snapshot that was created for both C1 and C2 to accommodate the fifth snapshot.

The **Maximum Snapshots** parameter in the **Replication Director Options** dialog box determines the maximum number of snapshots. To access the dialog box, click **Options** in the backup policy.

Figure 16-1 Replication Director Options dialog box



See “Snapshot Client and Replication Director (policy attributes)” on page 740.

Mirror retention type for SLP operations

A mirror replica of a snapshot is eligible for expiration as soon as:

- All immediate child copies are successfully created.
- All immediate child copies that are mirrors are eligible for expiration.

The selection of the **Mirror** retention indicates that NetApp volume SnapMirror is to be used as the replication method. If any non-mirror retention type such as **Fixed** or **Expire after copy** is selected for the **Replication** operation, the NetApp SnapVault replication method is used.

In mirror replication, the replica copy is dependent on the existence of the source. (The source can be the original snapshot or another replica.) Therefore, the retention of the replica depends on the retention of the source. If the source is deleted, the mirror is automatically deleted.

In non-mirror replication, the replica is independent of the source and can have an independent retention. If the source is deleted, the non-mirror replica is not affected and can be used longer than the source. Or, if the replica is deleted first, it is not recreated and the source can be kept longer than the replica.

Target retention type for SLP operations

This setting is used in Auto Image Replication in an Import storage lifecycle policy. Every Import SLP must have at least one operation with a **Target retention**.

The **Target retention** is enforced at the target domain, but the actual retention for the data is specified by the administrator at the source domain.

Target retention indicates that the data at the target master shall use the expiration date that was imported with the image. The date is fixed because the copy must have a fixed retention.

Similar to the **Fixed** retention, an image copy with a **Target retention** retention is eligible for expiration when all of the following criteria are met:

- The **Fixed** retention period for the copy has expired.
- All child copies have been created.
- All child copies that are mirror copies are eligible for expiration.

See [“About NetBackup Auto Image Replication”](#) on page 975.

Storage lifecycle policy options

This chapter includes the following topics:

- [Storage Lifecycle Policy dialog box settings](#)
- [New or Change Storage Operation dialog box settings](#)
- [Storage lifecycle policy validation dialog box](#)
- [Storage lifecycle policy Validation Report tab](#)

Storage Lifecycle Policy dialog box settings

The **New Storage Lifecycle Policy** dialog box and the **Change Storage Lifecycle Policy** dialog box contain the following settings.

Figure 17-1 Storage Lifecycle Policy tab

Change Storage Lifecycle Policy

Storage Lifecycle Policy | Validation Report

Storage lifecycle policy name:
SLP_1_snapshot

Data classification:
<No data classification>

Priority for secondary operations:
0 (higher number is greater priority)

Operation	Window	Storage	Volume Pool	Media Owner	Retention Type	Retention Peri...	Alternate Res...	Target Master
Snapshot		No Storage Unit				Maximum snap...		
Backup From Sn...	Default_24x7...	stu_adv_ca...			Fixed	2 weeks		

↑ ↓ ← →

Add... Change... Remove

State of Secondary Operation Processing

☒ Active

☐ Postponed

☐ Until 12/31/2038 11:59:59 PM

To find impact on Policies associated with this Storage Lifecycle Policy due to change in configuration click here.

Validate Across Backup Policies

OK

Cancel

Help

Table 17-1 Storage Lifecycle Policy tab

Setting	Description
Storage lifecycle policy name	The Storage lifecycle policy name describes the SLP. The name cannot be modified after the SLP is created.

Table 17-1 Storage Lifecycle Policy tab (*continued*)

Setting	Description
Data classification	<p>The Data classification defines the level or classification of data that the SLP is allowed to process. The drop-down menu contains all of the defined classifications as well as the Any classification, which is unique to SLPs.</p> <p>The Any selection indicates to the SLP that it should preserve all images that are submitted, regardless of their data classification. It is available for SLP configuration only and is not available to configure a backup policy.</p> <p>In an Auto Image Replication configuration where the master server domains run different versions of NetBackup, see the following topic for special considerations:</p> <p>See “About the storage lifecycle policies required for Auto Image Replication” on page 988.</p> <p>The Data classification is an optional setting.</p> <p>One data classification can be assigned to each SLP and applies to all operations in the SLP.</p> <p>If a data classification is selected (other than Any), the SLP stores only those images from the policies that are set up for that data classification. If no data classification is indicated, the SLP 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 an SLP with the same data classification.</p> <p>Once data is backed up in an SLP, the data is managed according to the SLP configuration. The SLP defines what happens to the data from the initial backup until the last copy of the image has expired.</p>
Priority for secondary operations	<p>The Priority for secondary operations option is the priority that jobs from secondary operations have in relationship to all other jobs. The priority applies to the jobs that result from all operations except for Backup and Snapshot operations. Range: 0 (default) to 99999 (highest priority).</p> <p>For example, you may want to set the Priority for secondary operations for a policy with a gold data classification 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 695.</p>

Table 17-1 Storage Lifecycle Policy tab (*continued*)

Setting	Description
Operations	<p>Use the Add, Change, and Remove buttons to create a list of operations in the SLP. An SLP must contain one or more operations. Multiple operations imply that multiple copies are created.</p> <p>See “About writing multiple copies using a storage lifecycle policy” on page 654.</p> <p>The list also contains the columns that display information about each operation. Not all columns display by default.</p> <p>For column descriptions, see the following topic:</p> <p>See “New or Change Storage Operation dialog box settings” on page 643.</p>
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 “Creating a hierarchy of storage operations in a storage lifecycle policy” on page 628.</p> <p>See “Modifying the hierarchy of operations in a storage lifecycle policy” on page 604.</p>
Active and Postponed	<p>The Active and Postponed options appear under State of Secondary Operation Processing and refer to the processing of all duplication operations in the SLP.</p> <p>Note: The Active and Postponed options apply to duplication operations that create tar-formatted images. For example, those created with <code>bpduplicate</code>. The Active and Postponed options do not affect the images that are duplicated as a result of OpenStorage optimized duplication, NDMP, or if one or more destination storage units are specified as part of a storage unit group.</p> <p>These options do not apply if either the source media server or the destination media server is not at NetBackup 7.6 or later.</p> <ul style="list-style-type: none"> ■ Enable Active to let secondary operations continue as soon as possible. When changed from Postponed to Active, NetBackup continues to process the images, picking up where it left off when secondary operations were made inactive. ■ Enable Postponed to postpone the secondary operations for the entire SLP. Postponed does not postpone the creation of duplication jobs, it postpones the creation of images instead. The duplication jobs continue to be created, but they are not run until secondary operations are active again. <p>All secondary operations in the SLP are inactive indefinitely unless the administrator selects Active or until the Until option is selected and an activation date is indicated.</p>
Validate Across Backup Policies button	<p>Click this button to see how changes to this SLP can affect the policies that are associated with this SLP. The button generates a report that displays on the Validation Report tab.</p> <p>This button performs the same validation as the <code>-conflict</code> option performs when used with the <code>nbsl</code> command.</p> <p>See “Storage lifecycle policy Validation Report tab” on page 652.</p>

New or Change Storage Operation dialog box settings

The **Storage Operation** dialog box contains two tabs:

- **Properties** tab

The following topic describes the options in the **Properties** tab:

See [“Properties tab of the Storage Operation dialog box”](#) on page 643.

- **Window** tab

Create a window during which a secondary operation can run. The following topic describes the options in the **Window** tab:

See [“Window tab of the Storage Operation dialog box”](#) on page 647.

Properties tab of the Storage Operation dialog box

[Table 17-2](#) describes the options available to configure for the various operations in a storage lifecycle policy. Not all options are available for each operation.

Table 17-2 Properties tab of the Storage Operation dialog box

Setting	Description
Source	Indicates the storage unit that is the source for the operation. The Source displays for the following operation types: Backup From Snapshot , Replication , Duplication , and Index From Snapshot .
Operation	The Operation selection determines which options appear in the dialog box. See “Operation types in a storage lifecycle policy” on page 608. See “About writing multiple copies using a storage lifecycle policy” on page 654.
Retention type	Select a Retention type from the following options: <ul style="list-style-type: none">■ See “Capacity managed retention type for SLP operations” on page 633.■ See “Expire after copy retention type for SLP operations” on page 634.■ See “Fixed retention type for SLP operations” on page 635.■ See “Maximum snapshot limit retention type for SLP operations” on page 636.■ See “Mirror retention type for SLP operations” on page 637.■ See “Target retention type for SLP operations” on page 638. See “Retention types for storage lifecycle policy operations” on page 631.

Table 17-2 Properties tab of the Storage Operation dialog box (*continued*)

Setting	Description
Retention period	<p>Specifies how long NetBackup retains the backup or copy. 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 or copy. After expiration, the files are unavailable for restores. For example, if the retention is set to two weeks, data can be restored from a backup that this schedule performs for two weeks after the backup.</p> <p>See “Retention Periods properties” on page 219.</p>
Target master	<p>Indicates that the copy of the image is to be created in a different master server domain. The target master server manages the storage where the image is to be copied.</p> <p>If Target master is selected for a Replication operation, the operation becomes an operation for Auto Image Replication.</p>
Local storage	<p>Indicate the storage unit to be used.</p> <p>Select from the following storage units:</p> <ul style="list-style-type: none"> Media Manager storage units (tape) Disk storage units (no BasicDisk or disk staging storage units) Storage unit groups (may contain no BasicDisk 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>Storage units or storage unit groups may appear in more than one storage lifecycle policy. Storage units or storage unit groups may be used in a storage lifecycle policy while also being used as standalone units.</p>
Storage unit	<p>Indicate the storage unit to be used.</p> <p>Select from the following storage units:</p> <ul style="list-style-type: none"> Media Manager storage units (tape) Disk storage units (no BasicDisk or disk staging storage units) Storage unit groups (may contain no BasicDisk 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. No storage unit A storage unit selection is necessary if the snapshot is to be used by a subsequent Replication operation or if the SLP contains only one operation. If neither situation applies to the SLP, the administrator may select No storage unit or may simply make no selection. <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 standalone units.</p>
Volume pool	The Volume pool option is enabled for tape storage units or virtual tape libraries (VTL).

Table 17-2 Properties tab of the Storage Operation dialog box (*continued*)

Setting	Description
Media owner	A Media owner is a group of NetBackup servers that are used for a common purpose. By specifying a Media owner , you allow only those media servers to write to the media on which backup images for a policy are written.
Alternate read server	An Alternate read server specifies the name of the server that is allowed to read a backup image originally written by a different media server. This option is available for Duplication operations only.
Preserve multiplexing	<p>The Preserve Multiplexing option is available for the duplication operations that use tape media or virtual tape libraries (VTL). 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 for an Import operation. The job priority that is indicated is the job priority for any import jobs which use this storage lifecycle policy.
Postpone creation of this copy until the source copy is about to expire	<p>Enable this option to defer the job until the source for the duplication is about to expire. When this option is enabled, the job begins 4 hours before the source is to expire. This default can be changed by changing the Deferred duplication offset time in the SLP Parameters host properties.</p> <p>See “SLP Parameters properties” on page 239.</p>

Table 17-2 Properties tab of the Storage Operation dialog box (*continued*)

Setting	Description
Advanced button and Window close preference options	<p>If a window closes and the jobs for an SLP have not completed, NetBackup attempts to suspend the images that are in progress. When the window reopens, NetBackup resumes those jobs at the point they were when suspended.</p> <p>Not all images can be suspended. The image must be the result of a duplication job where both the source and target of the duplication resides on either an AdvancedDisk or Media Manager storage unit.</p> <p>In addition, the duplication job must meet the following conditions:</p> <ul style="list-style-type: none">■ Both the source media server and the destination media server must have NetBackup 7.6 or later installed.■ The storage units must not be part of a storage unit group.■ The duplications were not created using optimized duplication, NDMP duplication, or OpenStorage duplication. <p>See “Duplication operation in an SLP” on page 612.</p> <p>Note: The closing of the window does not stop preliminary operations for some jobs. For example, NetBackup continues to extend the catalog for Exchange Granular Recovery after the window closes, but does not start duplications.</p> <p>Images that result from all other operations (such as Replication operations), are not suspended.</p> <p>Click the Advanced button to display the Window close preference options. The selections apply to the images that NetBackup does not suspend automatically.</p> <p>Select what NetBackup should do if images are not completed by the time the window closes and if the images cannot be suspended:</p> <ul style="list-style-type: none">■ Finish processing the active images. The window closes, but NetBackup continues to process the active images until they are finished. NetBackup does not begin to process any other images until the window reopens.■ Cancel the processing of the active images. The window closes and NetBackup immediately stops processing the active images. When the window reopens, NetBackup begins to process the images where it left off.

Figure 17-2 Windows close preference selections for secondary operations

New Storage Operation

Properties | **Window**

Source storage: No Storage Unit (Snapshot)

Operation: Backup From Snapshot

Destination Storage Attributes

Destination storage: STU-ACS-robot

Volume pool: NetBackup

Media owner: Any

Retention

Retention type: Fixed

Retention period: 2 weeks (level 1)

Duplication

Alternate read server:

☐ Postpone creation of this copy until the source copy is about to expire.

Advanced

Window close preference

After the window closes, NetBackup will not start processing any new images for this operation. Images currently being processed will be suspended, if possible.

Images unable to be suspended will be handled as follows:

☒ Finish processing the active images

☐ Cancel the processing of the active images

When the window reopens, processing will resume for images, which were suspended, cancelled or never started.

OK Cancel Help

Window tab of the Storage Operation dialog box

The **Window** tab appears for secondary operations in a storage lifecycle policy.

Creating a window for a secondary operation is optional. However, creating a window can better define when the job for the operation can run. In this way, the job from a secondary operation does not interfere with jobs of a higher priority, such as backup jobs. Without a window defined, the job for an operation can run at any time, on any day.

Figure 17-3 Window tab for secondary operations in a storage lifecycle policy

New Operation

Properties **Window**

☒ **Select from saved windows** Default_24x7_Window

☐ **Create new**

Window name: Default_24x7_Window View Impact Report

⚠ This is a default duplication window and cannot be modified.

Start Window **Exclude Dates**

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Thu												
Fri												
Sat												

To define a time window during which jobs can start, select a day, and enter a start time. Then enter either the end day and time, or the duration.

Start day: Sunday Start time: 00:00:00

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Start Day' list.

End day: Monday End time: 00:00:00

To change the increments available for selecting start times or end times, change the value in the Resolution field.

Duration (days hours:minutes): 1 00:00:00 Resolution: 5 min

Delete Clear Duplicate Undo

OK Cancel Help

Table 17-3 Window tab of the Storage Operation dialog box

Setting	Description
Select from saved windows	<p>You can either assign an existing window to the operation or create a new window for the operation.</p> <p>To use an existing window, select this option and then select a window from the drop-down menu.</p>
Create new	Select this option to create a new window for this operation to use.
Window name	Enter a name for the new window.
View Impact Report option	This option generates an Impact Report which lists the names of the storage lifecycle policies that currently use the window. The Impact Report also lists the operation that uses the window, and the source and the destination storage for the operation.

Table 17-3 Window tab of the Storage Operation dialog box (*continued*)

Setting	Description
Start Window tab	<p>The Start Window grid is grayed out and cannot be modified if the Default_24x7_Window is selected.</p> <p>The Start Window grid is active if a saved window is selected or when a new window is created.</p> <p>If the Start Window grid is changed for a saved window, click the View Impact Report option to display information about other operations in other SLPs that use the window.</p> <p>See “Creating a new window for a storage lifecycle policy operation” on page 649.</p>
Exclude Days tab	<p>Use the Exclude Days tab to exclude specific dates from a window.</p> <p>See “Excluding days from a window for a storage lifecycle policy operation” on page 650.</p>

Creating a new window for a storage lifecycle policy operation

To create a new window for SLP operations

- 1 In the **Window** tab of the storage operation dialog box, enable **Create new**.
- 2 Select the **Start Window** tab.
- 3 The days of the week appear along the left side of the grid. The time of day appears along the top of the grid in 24-hour time.

To change the increments available for selecting start times or end times, change the value in the **Resolution** field.
- 4 Indicate the opening and closing times of the window for each day. The following lists several methods to do so:
 - Drag the cursor along the Start Window grid on each day you want the window to open and close.
 - Use the drop-down menus to select a **Start day** and an **End day**. Then select a **Start time** and an **End time**.
 - Use the drop-down menu to select a **Start day** and the **Duration** of the window for that day in hours and minutes. Adjust the **Start time** for your environment.

To create multiple time windows:

- First, create one window.

- Click **Duplicate**.
The window is duplicated to any days without existing schedules. Duplication stops when it reaches a day that already contains a window.
 - On days that you do not want the time window to be open, select the window and click **Delete**.
- 5 Use the buttons under the Start Window grid to do the following:
- | | |
|--------------------------------------|---|
| To change the start time or end time | Adjust the Start time or End time . |
| 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 . |
- 6 Click **OK** to save the window and the operation.

Excluding days from a window for a storage lifecycle policy operation

Use the **Exclude Days** tab to exclude specific days from a window. If a day is excluded from a window, jobs do 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 day from the storage lifecycle policy window

- 1 In the **Window** tab, select the name of an existing window from the drop-down menu.
- 2 Select the **Exclude Days** tab.
- 3 Use one or more methods to indicate the days to exclude:
 - Select the day(s) on the 3-month calendar that you want to exclude. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring Week Days**:
 - Click **Select All** to select all of the days in every month for every year.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select a specific day to exclude for every month.
 - Click the column head of a day of the week to exclude that day every month.

- Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to exclude that week every month.
 - To indicate **Recurring Days of the Month**:
 - Click **Select All** to select all of the days in every month.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select that day to exclude each month.
 - Click **Last Day** to exclude the last day of every month.
 - To indicate **Specific Dates**:
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.
The date appears in the **Specific Dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 4 Add additional dates as necessary, then click **OK** to save the window and the operation.

Storage lifecycle policy validation dialog box

The Storage Lifecycle Policy validation dialog box may appear if NetBackup cannot save the SLP as configured because of problems with the operations in the SLP. The dialog box may also appear after the administrator clicks **Validate Across Backup Policies**, before the **Validation Report** tab displays.


The Storage Lifecycle Policy validation dialog box displays the errors that must be corrected before the SLP can be saved. For example, errors regarding the hierarchy of operations in the SLP or errors concerning the storage units that the operations indicate.

The dialog box contains the following information about any validation errors:

Index	The operation in the SLP that contains errors. The index number is the operation's order in the SLP. For example, the second operation has an index number of two.
Operation	The type of operation where the error occurs in the SLP.
Storage	The storage name that is indicated in the operation where the error occurs.
Error code	The NetBackup status code. Use the NetBackup Troubleshooter or the NetBackup Status Codes Reference Guide to determine corrective actions. See "Running the Troubleshooter" on page 50.

Error description The vendor-specific error code and a description of the error.

Figure 17-4 Storage Lifecycle Policy validation dialog box



Storage lifecycle policy Validation Report tab

NetBackup validates the SLP when it is first created and whenever it is changed. The **Validation Report** tab lists the conflicts between the proposed changes to the operations in a storage lifecycle policy and any backup policy that uses the SLP.

Likewise, when a policy is created that indicates an SLP as the **Policy storage**, a similar validation report may display. The report lists any conflicts between the policy and the SLP that it has indicated.

The conflicts that are listed must be resolved in order for a job that references the SLP to run successfully. Use the **Troubleshooter** or the online Help on this tab for a list of common status codes that result from SLP conflicts.

Note: The Request has timed out message may appear in environments with very busy servers.

To increase the timeout to account for the latency in connectivity, increase the `NBJAVA_CORBA_DEFAULT_TIMEOUT` value in the following files:

On Windows: The `setconf.bat` configuration file
(`Install_path\java\setconf.bat`).

On UNIX: The `nbj.conf` configuration file (`/usr/opensv/java/nbj.conf`).


The report checks for the following conflicts between the selected SLP and the backup policies that use it:

- The data classification in the storage lifecycle policy does not match that in the referencing backup policies.
- The SLP contains a **Snapshot** operation, but the referencing backup policies do not have **Perform snapshot backups** enabled.
- The SLP does not contain a **Snapshot** operation, but the referencing backup policies have the **Perform snapshot backups** enabled.
- The SLP contains a **Snapshot** operation, but the referencing backup policies cannot enable the **Retain snapshots for Instant Recovery** option.
- The policy is of the **NBU-Catalog** backup type, but the SLP configuration does not indicate a **CatalogBackup** volume pool.
- The policy is not of the **NBU-Catalog** backup type, but the SLP configuration indicates a **CatalogBackup** volume pool

To generate a validation report using the command line, run the following command:

```
nbstl SLP_name -impact
```

Figure 17-5 Validation Report tab of the Storage Lifecycle Policy dialog box



Using a storage lifecycle policy to create multiple copies

This chapter includes the following topics:

- [About writing multiple copies using a storage lifecycle policy](#)
- [How the order of the operations determines the copy order](#)
- [About ensuring successful copies using lifecycles](#)

About writing multiple copies using a storage lifecycle policy

A storage lifecycle policy can be used to create multiple copies of backups and snapshots.

NetBackup permits only one method to create multiple copies to be in use at one time. Use only 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 the **Policy storage**.
See [“Multiple copies \(schedule attribute\)”](#) on page 760.
- Add multiple **Backup** operations or one or more **Duplication** or **Replication** operations to the operations list of the SLP.
See [“New or Change Storage Operation dialog box settings”](#) on page 643.

The same criteria for creating copies applies to both methods.

The following topics are considerations when storage lifecycle policies are used to create multiple copies.

How the order of the operations determines the copy order


The order in which the operations appear in a storage lifecycle policy determines the copy number.

For example, in [Figure 18-1](#) a lifecycle is configured to create three copies:

- Two copies as part of two different backup operations.
- One copy as part of a duplication operation.

To make sure that copy 1 is written to disk, place the **Backup** operation that writes to a disk storage unit before the **Backup** operation that writes to a tape storage unit.

Figure 18-1 Operation 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 an SLP, all copies must be completed. An SLP 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 an SLP does not allow a copy to be expired before all copy operations in the SLP are complete. NetBackup changes the retention period of a copy to Infinity until all copies are created. After all copies are complete, the retention returns to the level as set in the policy.

To create successful copies, a **Backup** operation may be required to duplicate a backup onto the storage unit for another **Backup** operation.

Consider the following example: The operations list for an SLP contains two **Backup** operations to two storage units (BU_1, BU_2) and three **Duplication** operations.

The backup to BU_1 is successful, but the backup to BU_2 is unsuccessful.

To fulfill the backup on BU_2, NetBackup creates a duplication job from BU_1 to BU_2. The duplication job is in addition to the jobs that are run for the three duplication operations.

Duplication jobs can be controlled by using the `nbstlutil` command.

See [“Lifecycle operation administration using the nbstlutil command”](#) on page 606.

Storage lifecycle policy versions

This chapter includes the following topics:

- [About storage lifecycle policy versions](#)
- [Storage lifecycle changes and versioning](#)
- [When changes to storage lifecycle policies become effective](#)
- [Deleting old storage lifecycle policy versions](#)

About storage lifecycle policy versions

Once a storage lifecycle policy is configured, it runs according to a single configuration or definition. The definition affects the operations once they begin to run as well as the copies once the image is in process.

The ability to create SLP versions lets administrators safely modify a definition without waiting until all of the copies that are associated with the SLP have been processed. Each copy that an SLP manages is tagged with the SLP name and the SLP version number. These two attributes are written into the image header, in the NetBackup image catalog. Whenever an administrator creates or changes an SLP, NetBackup creates a new version (between 0 and n). New jobs use the most recent SLP version.

When a new job is submitted to the Activity Monitor, the job is tagged with the most recent SLP version number. The processing of a copy that is associated with a version remains fixed according to that version of the SLP definition. It is fixed at job time and does not change, unless the administrator uses the `nbstl` command to modify an existing version.

An SLP 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**.
Any change that an administrator makes to an SLP using the **NetBackup Administration Console** creates a new SLP version. The new version is created when the changes to the SLP are committed or saved. The **NetBackup Administration Console** always displays the most recent version.
- Using the `nbstl` command.
If an administrator uses `nbstl` to change an SLP, `nbstl` creates a new version by default.
However, the `nbstl` command contains options to view different versions and to modify the definitions of existing SLP versions without creating a new version. The options are as follows:

<code>-all_versions</code>	Use to display all versions of an SLP 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 SLP 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 operation.
<code>-pool</code>	The volume pool for each operation.
<code>-server_group</code>	The server group for each operation.

When changes to storage lifecycle policies become effective

<code>-rl</code>	The retention level for each operation.
<code>-as</code>	The alternate read server for each operation.
<code>-mpx</code>	The preserve multiplexing option for duplication copies.

Some fields require values for all of the operations in the SLP. Make sure that the number of values that are specified for the fields matches the existing operation count.

For example, in an SLP that contains three operations, to change the value of one, a value must be given for all three operations. Note that the values for all three operations are replaced. To change the value for the second operation, provide the existing values for the first and the third operations.

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 SLP version:

<code>-uf</code>	The type of the operation.
<code>-managed</code>	The retention type for the operation: Fixed, Capacity managed, or Expire after copy.
<code>-source</code>	The source of an operation, used primarily in hierarchical SLP configurations.
<code>-dc</code>	The data classification of an existing version.
	The number of operations. You cannot add an operation or remove an operation from the SLP definitions.

See [“Creating a storage lifecycle policy”](#) on page 601.

You cannot instruct an SLP 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 jobs.

When the `nbstl` command is used to alter an existing storage lifecycle policy version, those changes may not become effective immediately. The images that are managed by the SLP version that was altered may already belong to a job that is Active or Queued, as seen in the Activity Monitor. Once a job is queued, the characteristics (SLP attributes) are fixed for that job and subsequent changes to the definition have no effect. To make changes effective for a backlog of jobs, 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** operation have no effect because the backup job is already underway or completed.
- Changes to a **Duplication** operation do not affect the copies that previous duplication jobs created.
- Changes to a **Duplication** operation do not affect the 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, cancel the applicable duplication jobs. Once the job is canceled, `nbstserv` re-forms and re-submits new duplication jobs for these copies, using the changes to the appropriate version of the SLP.
- Changes to a **Duplication** operation affect the copies that have not yet been created and have not yet been submitted. (That is, they are not yet represented by a duplication job in the Activity Monitor). Your changes become effective for the next duplication session. Whenever `nbstserv` begins a new session, it re-reads the definitions for processing instructions.
- If a duplication job does not complete successfully, unfinished images in the job are submitted as part of a new job. Changes to the version affect the resubmitted job.

Deleting old storage lifecycle policy versions

When a version of a storage lifecycle policy is no longer the active (or most recent) version, the version is subject to deletion. NetBackup automatically deletes the inactive version after all the copies that refer to it have finished processing. When the copies are complete, they are considered SLP-complete.

By default, NetBackup deletes an inactive SLP version after 14 days.

The following parameters in the **SLP Parameters** host properties apply to version deletion:

- **Cleanup interval** (SLP.CLEANUP_SESSION_INTERVAL)

- **Unused SLP definition version cleanup delay**
(SLP.VERSION_CLEANUP_DELAY)

See [“SLP Parameters properties”](#) on page 239.

Configuring backups

- [Chapter 20. Creating backup policies](#)
- [Chapter 21. Synthetic backups](#)
- [Chapter 22. Protecting the NetBackup catalog](#)
- [Chapter 23. About the NetBackup relational database](#)
- [Chapter 24. Managing backup images](#)

Creating backup policies

This chapter includes the following topics:

- [About the Policies utility](#)
- [Planning for policies](#)
- [Creating a backup policy](#)
- [Adding or changing schedules in a policy](#)
- [Changing multiple policies at one time](#)
- [Copying or moving policy items to another policy](#)
- [Copying a policy to create a new policy](#)
- [Copying a schedule into the same policy or different policy](#)
- [Deleting schedules, backup selections, or clients from a policy](#)
- [Policy Attributes tab](#)
- [Schedules tab](#)
- [Schedule Attributes tab](#)
- [Start Window tab](#)
- [Excluding days from a schedule](#)
- [Include Dates 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](#)

- [About the Clients tab](#)
- [Backup Selections tab](#)
- [VMware Policy tab](#)
- [Hyper-V Policies 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 provide the following instructions for a backup:

What type of client to back up.	See “Policy Attributes tab” on page 680.
Where to store the backup.	See “Policy Attributes tab” on page 680.
When and how frequently to perform the backup.	See “Schedules tab” on page 743.
Which clients to back up.	See “About the Clients tab” on page 790.
Which client files and directories to back up.	See “Backup Selections tab” on page 796.

Using the Policies utility

To navigate the Policies utility

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2

To display the policy details of a single policy:	In the center pane, select a policy name. The policy details display in the right pane.
To open a policy:	In the center pane, double-click on the policy name. The Change Policy dialog box opens.

To display specific configuration information:

In the center pane, click on the tree element next to the policy name to expand the policy configuration areas:

- **Attributes**
- **Schedules**
- **Clients**
- **Backup Selections**

In the center pane, under a policy name, select one of the configuration areas to see a detailed view of that area.

To display information about all policies on the current master server: Double-click **Summary of All Policies**.

Figure 20-1 Policies utility in the NetBackup Administration Console

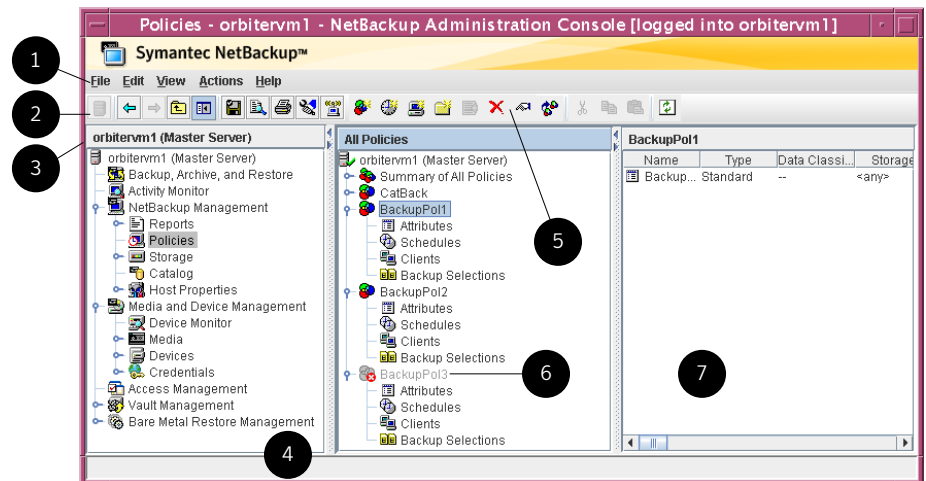


Table 20-1 Policies utility in the NetBackup Administration Console

Number	Description
1	The menu toolbar.
2	The standard NetBackup toolbar. See “Standard and blue toolbars” on page 45.
3	The name of the currently selected master server.
4	Tree view in the left pane.

Table 20-1 Policies utility in the NetBackup Administration Console (*continued*)

Number	Description
5	The user toolbar is specific to the Policies utility. See “Standard and user toolbars” on page 45.
6	Select a subnode from Summary of All Policies to display all possible node attributes in the right pane. For example, Schedules displays a list of all schedules.
7	Details pane.

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 20-2 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>

Table 20-2 Steps for planning policies (*continued*)

Step	Action	Description
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 “Windows example of one client in multiple policies” on page 668.</p>
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.

Table 20-2 Steps for planning policies (*continued*)

Step	Action	Description
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 669.</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 769.</p>
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>/usr</code>, <code>/h001</code>, and <code>/h002/projects</code> on UNIX client1 takes too much time, create a new policy for <code>/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 150.</p> <p>The Multiplexing and Allow multiple data streams policy attributes also allow backup policies to be processed in parallel.</p> <p>See “Allow multiple data streams (policy attribute)” on page 711.</p>

See [“About the Policies utility”](#) on page 664.

See [“Policy Attributes tab”](#) on page 680.

Windows 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 20-3 One client in multiple policies

Policies	Client	Files	Storage
Policy1	client1	C:\ D:\User D:\h001 E:\h002\Projects	8mm
Policy2	client1 client1	E:\h002\DevExp E:\h002\DesDoc	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 20-4 Policy attributes that affect how clients are grouped in policies

Attribute	Description
Policy Type	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. See “Policy type (policy attribute)” on page 681.
Destination	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 . See “Data classifications (policy attribute)” on page 684. See “Policy storage (policy attribute)” on page 685. See “Policy volume pool (policy attribute)” on page 688.
Job Priority	This attribute determines the priority for the backups of all of the clients in the policy. See “Job priority (policy attribute)” on page 695.
Follow NFS	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. See “Follow NFS (policy attribute)” on page 696.
Cross mount points	This attribute lets NetBackup cross file system boundaries for all clients in the policy. See “Cross mount points (policy attribute)” on page 699.

Table 20-4 Policy attributes that affect how clients are grouped in policies
(continued)

Attribute	Description
Backup Network Drives	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.) See “Backup Network Drives (policy attribute)” on page 697.
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 703.

About Microsoft DSFR backups

Microsoft Distributed File System Replication (DFSR) service is a multi-master replication engine that is used to keep folders synchronized on multiple servers. To do this, each of these independent DFSR servers maintains a database which is essential in providing this service but is independent of the DFSR data that is being replicated. Although NetBackup protects this database, as well as the DFSR data, there are some things to consider when restoring both.

A DFSR server maintains information in its database that is used by other DFSR servers to identify and validate replication changes that may be needed. One bit of information is a globally unique version number (GVSN), which increases monotonically with every replicated file or folder alteration. It can be thought of as a version or generation number for each DFSR database on each replicated volume.

A problem occurs if a DFSR server is restored to an earlier database version. The other servers, which monitor and remember each other’s numbers, will be confused when presented with the older number. At this point replication is impossible and will stop until this is corrected. Consider carefully if your restore actually needs to include the DFSR database, in addition to the DFSR data. This is likely to happen if you have a virtualized DFSR server and you restore it to a snapshot that had been taken earlier. Please refer to the Microsoft documentation on managing and using DFSR for additional details.

In an environment that has DFSR, two changes occur in NetBackup, as follows:

- To preserve data integrity, the folder or folders that host the Shared Replica DFSR data are excluded automatically by NetBackup from normal file system backups.

- The top-level DFSR shared folders become part of the Shadow Copy Components. Therefore, the data is snapped consistently by Windows Volume Shadow Copy Service (VSS) before each backup.

Microsoft supports only the VSS writer for DFSR managed data backups. The VSS writer stops and restarts the DFS Replication service automatically. Symantec recommends that you schedule the backups to coincide with a period of low activity. (If you stop the replication service manually, Microsoft change journal problems may occur. Specifically, Update Sequence Number (USN) Journal wrap may occur.)

Symantec recommends that you back up DSFR data based on the amount of data under DSFR control as described in [Table 20-5](#).

Table 20-5 Microsoft DSFR backup recommendations

Amount of data	Symantec recommendation
Less than 50 GBs	<p>Configure one policy as follows:</p> <ul style="list-style-type: none"> ■ Choose the DFSR server host as the client. See “Adding or changing clients in a policy” on page 791. ■ Choose ALL_LOCAL_DRIVES as the Directive in the Backup Selections for the policy. The ALL_LOCAL_DRIVES directive includes the Shadow Copy Components:\ automatically. See “Adding backup selections to a policy” on page 797. <p>One policy can back up the data within a reasonable time window.</p>

Table 20-5 Microsoft DSFR backup recommendations (*continued*)

Amount of data	Symantec recommendation
More than 50 GBs	<p>Configure one backup policy for each DSFR server, and in that policy specify only the replication folders. A policy for each host's replication data ensures that the DSFR data is backed up within a reasonable time window.</p> <p>For each DSFR server host, do the following:</p> <ul style="list-style-type: none"> ■ Create a global exclude list for All Policies and All Schedules. Exclude the following DSFR top-level folder: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolder The global exclude list ensures that the DFSR components are not backed up accidentally by other backup policies for the client. See "Exclude Lists properties" on page 128. ■ Create a backup policy for the DSFR data, as follows: <ul style="list-style-type: none"> ■ For the client, specify the DFSR server host. For the servers that are hosted in a cluster, specify the DFSR cluster name rather than the local host name. See "Adding or changing clients in a policy" on page 791. ■ For the Backup Selections for the policy, specify the absolute path to each of the top-level DFSR folders on that host. The following is an example path: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolders\folder_name <i>Tip:</i> Use the Backup, Archive, and Restore interface to browse the Shadow Copy Components for the DfsrReplicatedFolders folder. The BAR interface shows the path to each DSFR folder that you need to enter as a backup selection. See "Adding backup selections to a policy" on page 797. ■ For the backup policy, create an exception to the exclude list and specify the top-level DFSR directory, as follows: Shadow Copy Components:\User Data\Distributed File System Replication\DfsrReplicatedFolders For the exception Policy, specify the backup policy for the DSFR data. Also specify All Schedules for the Schedule. If DFSR is hosted in a cluster, create the exception for each host in the cluster. The exception ensures that the Shadow Copy Components DFSR paths are included for backup after NetBackup processes the global exclude list. See "Adding an exception to the exclude list" on page 133.

The DSFR servers must be configured for **Windows Open File Backup**. The snapshot provider must be **Volume Shadow Copy Service**. Configure these properties in the **NetBackup Administration Console** master server host properties **Client Attributes** tab.

See ["Windows Open File Backup tab of the Client Attributes properties"](#) on page 93.

During a backup, Windows writes the following event ID messages to the application event log of a DFSR host:

```
Event ID=1102
Severity=Informational
The DFS Replication service has temporarily stopped replication
because another application is performing a backup or restore
operation. Replication will resume after the backup or restore
operation has finished.
```

```
Event ID=1104
Severity=Informational
The DFS Replication service successfully restarted replication
after a backup or restore operation.
```

To restore DFSR data, use the NetBackup **Backup, Archive and Restore** client interface to browse the **Shadow Copy Components** for the files or folders to restore, as follows:

```
Shadow Copy Components:\User Data\Distributed File System
Replication\DfsrReplicatedFolders\folder_name
```

The following Symantec HOWTO provides a more detailed procedure about configuring DSFR backups:

<http://www.symantec.com/docs/HOWTO65638>

Creating a backup policy

The easiest method to set up a backup policy is to use the **Policy Configuration Wizard**. This wizard guides you through the setup process by automatically choosing the best values for most configurations.

Not all policy configuration options are presented through the wizard. For example, calendar-based scheduling and the **Data Classification** setting. After the policy is created, modify the policy in the **Policies** utility to configure the options that are not part of the wizard.

Note: Do not use the Policy Configuration Wizard to configure policies for Replication Director.

Using the Policy Configuration Wizard to create a backup policy

Use the following procedure to create a backup policy with the Policy Configuration Wizard.

To create a backup policy with the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Create a Policy** to begin the **Policy Configuration Wizard**.
- 3 Select the type of policy to create:
 - A policy to back up **File systems, databases, or applications**.
 - A policy to protect **VMware or Hyper-V** virtual machines.
 - A policy to back up **NDMP hosts**.
 - A policy to back up **Oracle** data.
- 4 Click **Next** to start the wizard and follow the prompts.

Click **Help** on any wizard panel for assistance while running the wizard.

Creating a backup policy without using the Policy Configuration Wizard

Use the following procedure to create a backup policy in the **NetBackup Administration Console** without using the Policy Configuration Wizard.

To create a policy without the Policy Configuration Wizard

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **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 1093.
- 4 If necessary, clear the **Use Policy Configuration Wizard** check box.
- 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 680.
 See [“Schedules tab”](#) on page 743.
 See [“About the Clients tab”](#) on page 790.
 See [“Backup Selections tab”](#) on page 796.

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 Expand the policy name in the middle pane, then select **Schedules**.
- 3 Perform one of the following actions:

Add a schedule

On the **Actions** menu, click **New > Schedule**.

Copy a schedule and paste it into another policy

- Expand the policy which contains a schedule that you'd like to copy.
- In the right pane, right-click the schedule and select **Copy**.
- Expand the policy where you'd like to paste the schedule.
- In the right pane, right-click anywhere in the schedule area and select **Paste**.

You can also copy a schedule with the **Edit > Copy to New** option:

See [“Copying a schedule into the same policy or different policy”](#) on page 679.

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 Days** tab, and **Include Dates** tab (when applicable).

See [“Schedule Attributes tab”](#) on page 744.

See [“Start Window tab”](#) on page 776.

See [“Excluding days from a schedule”](#) on page 780.

See [“Include Dates tab”](#) on page 781.

- 5 Click **OK**.
- 6 To add more schedules, click **Add** or **New** and repeat step 4.
If this schedule is the last schedule, click **OK** or **Close**.

Changing multiple policies at one time

Use the following procedure to change more than one NetBackup policy at the same time.

Note: You can change **Attributes**, **Clients**, and **Schedules** of multiple policies at one time. However, you cannot change **Backup Selections** of multiple policies at one time.

To change multiple policies

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management** and select **Policies**.
- 2 In the middle pane, expand **Summary of All Policies** and select the node that you want to change.
- 3 In the right pane, either hold down the **Shift** key and select contiguous rows, or hold down the **Ctrl** key and select discontinuous rows for multiple selection.

Note:

You can change different schedules across different policies at one time. If you are changing multiple schedules of different policy types, the **Attributes** tab is disabled. A tri-state is displayed for an attribute that has different values for multiple policies you selected.

4 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.

5 Specify a value, enable or disable an attribute, or enter text for the attributes you want to change. Any change that you make 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 680.

See [“Schedule Attributes tab”](#) on page 744.

See [“About the Clients tab”](#) on page 790.

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

You can copy or move entire policies, attributes, schedules, clients, and backup selections from one policy 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 middle pane, select either **Attributes**, **Schedules**, **Clients**, or **Backup Selections** of a policy that you want to copy or move.
- 3 In the right pane, based on your selection in the previous step, select the attributes, schedules, clients, or backup selections of a policy that you want to copy or move.
- 4 Do one of the following:

To copy an item

On the **Edit** menu, click **Copy** or, press **Alt+C**.

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 In the middle-pane, select the policy item to which you want to paste the copied items.
- 6 From the **Edit** menu, click **Paste**.

When you paste items with the same name, you are provided an option to either copy and replace the existing item, or copy but keep the existing item, or do not copy.

Note: If the schedules do not match the policy type, the schedules are not copied. The action is indicated in a dialog box.

When you paste Attributes, the existing Attributes of the policy are always replaced. Whereas, when you paste Backup Selections, the Backup Selection is always copied to the policy, and not replaced.

The copying or moving feature is also applicable to instances and instance groups of Oracle and SQL type of backup policies.

Copying a policy to create a new policy

Use the **Copy to New** option to save time creating policies. This option is especially useful for the policies that contain many of the same policy attributes, schedules, clients, or backup selections.

To copy a policy to create a new one

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the middle pane, select the policy to copy.
- 3 On the **Edit** menu, click **Copy to New**
- 4 In the **Copy a Policy** dialog box, enter the name of the policy that you want to copy. You can indicate a policy other than the one that is selected
- 5 Enter the name for the new policy.
- 6 Click **OK**. The only difference between the new policy and the copied policy is the name.

Copying a schedule into the same policy or different policy

Use the **Copy to New** option to save time creating schedules. Use this option to copy a schedule into the same policy or different policy.

To copy a schedule to create a new one

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the middle pane, expand a policy and select the **Schedules** node that contains the schedule that you want to copy.
- 3 In the right pane, select the schedule that you want to copy.
- 4 On the **Edit** menu, click **Copy to New**
- 5 In the **Copy Schedule** dialog box, enter the name of the new schedule.
- 6 Use the menu to select the name of the policy to which you want to paste the schedule. You can paste the schedule into the same policy or a different policy.
- 7 Click **OK**. The **Change Schedule** dialog box opens for the new schedule.

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 Expand the policy name in the middle pane, then select **Attributes**, **Schedules**, **Backup Selections**, or **Clients**.
- 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.

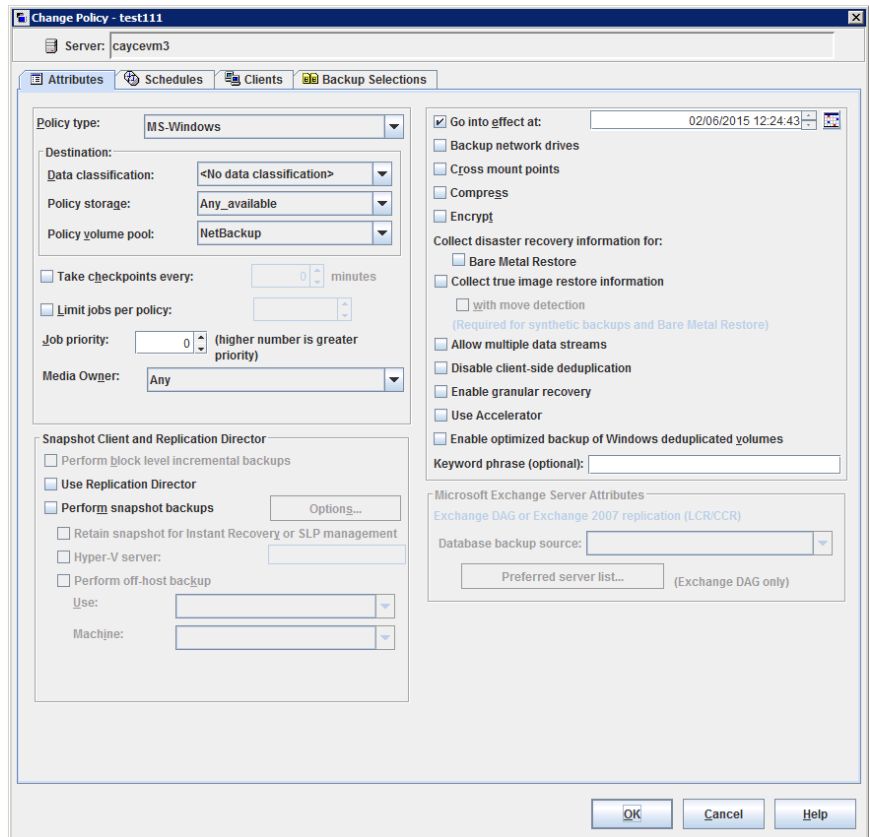
When a client is deleted from the client list, the NetBackup client software is not deleted or uninstalled from the client. Backups for the client can be recovered until the backups expire. Also, when a file is deleted from a backup selection list, the actual file is not deleted from the client.

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 20-2](#) shows the **Attributes** tab of a NetBackup policy.

Figure 20-2 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. For example, **NBU-Catalog** is not used for client backups.

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.

When you change the policy type of an existing policy, other selections or features of the policy may become invalid as well.

[Table 20-6](#) describes all the types of NetBackup policies.

Table 20-6 NetBackup policy types

Policy type	Description
DataTools-SQL-BackTrack (UNIX only)	Use for the policies that contain only clients with the NetBackup for SQL-BackTrack agent. 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 for DB2 agent. For information on setting up this policy type, see the guide for this option.
FlashBackup (UNIX only)	Combines the speed of raw-partition backups with the ability to restore individual files. This policy type applies to UNIX clients only. Requires the Enterprise Client license. For information on setting up this type of policy, see the NetBackup Snapshot Client Guide .
FlashBackup-Windows (Windows only)	Combines the speed of raw-partition backups with the ability to restore individual files. This policy type applies to Windows clients only. Requires the Enterprise Client license. For information on setting up this type of policy, see the NetBackup Snapshot Client Guide .
Hyper-V	<p>To back up the virtual machines that reside on Windows Hyper-V servers, by means of on-host or off-host backups. This policy type requires the Enterprise Client license. Users can upgrade pertinent policies to the Hyper-V policy type using one of the following methods:</p> <ul style="list-style-type: none"> Change the policy type in the NetBackup Administration Console for one policy at a time. Change the policy type for multiple policies at one time using the <code>bpplconvert</code> script that is located in the following location: On Windows: <code>Install_path\Veritas\NetBackup\bin\goodies</code> On UNIX: <code>usr/openv/netbackup/bin/goodies</code> <p>For information on setting up this type of policy, see the NetBackup for Hyper-V Guide.</p>
Informix-On-BAR (UNIX only)	Use for the policies that contain only clients with the NetBackup for Informix agent. For information on setting up this policy type, see the guide for this option.

Table 20-6 NetBackup policy types (*continued*)

Policy type	Description
Lotus-Notes	Use for the policies that contain only clients with the NetBackup for Lotus Notes agent. For information on setting up this policy type, see the guide for this option.
MS-Exchange-Server (Windows only)	Use for the policies that contain only clients with the NetBackup for Exchange Server agent. For information on setting up this policy type, see the guide for this option. This policy type applies to Windows clients only.
MS-SharePoint (Windows only)	Use to configure a policy NetBackup for SharePoint Server. This policy type applies to Windows clients only.
MS-SQL-Server	Use for the policies that contain only clients with the NetBackup for SQL Server agent. For information on setting up this policy type, see the guide for this option.
MS-Windows	Use for the policies that contain only Windows clients of supported Windows operating system levels. Standard and MS-Windows policy types are the only policy types that support the following options: <ul style="list-style-type: none"> ■ Checkpoint restart for backups or restores See "Take checkpoints every __ minutes (policy attribute)" on page 689. ■ Synthetic backups See "Policy type (policy attribute)" on page 681. ■ Collect disaster recovery information for Bare Metal Restore See "Collect disaster recovery information for Bare Metal Restore (policy attribute)" on page 707.
NBU-Catalog	Use for catalog backup jobs.
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.
Oracle	Use for the policies that contain only clients with the NetBackup for Oracle agent. For information on setting up this policy type, see the guide for this option.
SAP	Use for the policies that contain only clients with the NetBackup SAP agent. For information on setting up this policy type, see the guide for this option.

Table 20-6 NetBackup policy types (*continued*)

Policy type	Description
Standard	<p>Use for UNIX clients (including supported Mac clients), except for those clients that are covered by specific products, such as Oracle.</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 ■ Synthetic backups ■ Collect disaster recovery information for Bare Metal Restore
Sybase	<p>Use for the policies that contain only clients with the NetBackup for Sybase agent. 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>
VMware	<p>For backup of any virtual machines that reside on VMware vSphere servers, by means of off-host backups. This policy type requires the Enterprise Client license.</p> <p>Users can upgrade pertinent policies to the VMware policy type using one of the following methods:</p> <ul style="list-style-type: none"> ■ Change the policy type in the NetBackup Administration Console for one policy at a time. ■ Change the policy type for multiple policies at one time using the <code>bpplconvert</code> script that is located in the following location: On Windows: <code>Install_path\Veritas\NetBackup\bin\goodies</code> On UNIX: <code>usr/openv/netbackup/bin/goodies</code> <p>For information on setting up this type of policy, see the NetBackup for VMware Guide.</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 639.

See [“Data Classification properties”](#) on page 110.

See [“Creating a Data Classification”](#) on page 112.

See [“About storage lifecycle policies”](#) on page 600.

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.

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.

See [“About staging backups”](#) on page 576.

Storage unit	<p>Select the name of the storage unit that is to be the storage destination for the policy data. It can be disk or tape-based.</p> <p>If it is configured to do so, the storage unit determines which type of disk staging is used for the policy.</p> <p>See “Creating a storage unit” on page 548.</p>
--------------	---

Storage lifecycle policy	<p>Select the name of the storage lifecycle policy that is to be the storage destination for the policy data.</p> <p>The drop-down 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.</p> <p>See “Global Attributes properties” on page 150.</p> <p>If it is configured to do so, the storage lifecycle policy determines which type of disk staging is used for the policy.</p> <p>If setting up snapshot replication with Replication Director, select a storage lifecycle policy that contains a snapshot-capable storage unit.</p> <p>See “About storage lifecycle policies” on page 600.</p>
Storage unit group	<p>Select the name of the storage unit group that is to be the storage destination for the policy data.</p> <p>See “About storage unit groups” on page 589.</p>
Any Available	<p>If Any Available 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.</p> <p>NetBackup does not select a null_stu storage unit if Any Available is selected. A null_stu storage unit is created only when Symantec Support uses the NullOST plug-in to identify and isolate data transfer bottlenecks.</p>

Figure 20-3 Icons indicate type of storage

-  Storage unit (tape device)
-  Storage unit (disk)
-  Storage unit group
-  Storage lifecycle policy

Note: If different storage is selected for the **Override policy storage** option on the **Schedule Attributes** tab, that selection overrides the **Policy storage** attribute.

See [“Override policy storage \(schedule attribute\)”](#) on page 764.

See [“Considerations for selecting a destination for Policy storage”](#) on page 687.

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 20-7

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 764. ■ 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 551. See “On demand only storage unit setting” on page 570.
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 180.</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 507. See “About scratch volume pools” on page 443. ■ 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 551. See “On demand only storage unit setting” on page 570.

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 764.

See [“Example of overriding the policy volume pool”](#) on page 689.

The following table describes the default volume pools that NetBackup defines.

Table 20-8 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 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 20-9 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 NetBackup volume pools](#)” on page 441.

See “[Adding a volume pool](#)” on page 507.

See “[About scratch volume pools](#)” on page 443.

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, we 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. This is often quicker rather than restarting 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 for the **MS-Windows** or **Standard** policy types. 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 692.

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 150.

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 20-10 Support for checkpoint restart

Item	Description
Basic disk staging	<p>Checkpoint restart is supported for Stage I. Checkpoint restart is not supported for Stage II.</p> <p>See “About basic disk staging” on page 577.</p> <p>See “About staging backups” on page 576.</p>
MS-Windows (policy type)	<p>The following pertain to Windows clients:</p> <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). <p>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 backup resumes are backed up again during the next incremental backup.</p>
Multiple copies (schedule attribute)	<p>Checkpoint restart is supported for the policies that are configured to create multiple backup copies.</p> <p>See “Multiple copies (schedule attribute)” on page 760.</p> <p>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.</p>
Snapshot Client (policy attribute)	<p>Checkpoint restart is supported for use with local or alternate client backups. However, the following policy attributes are not supported:</p> <ul style="list-style-type: none"> ■ Block Level incremental backups ■ Media server Copy ■ Third-Party Copy Device ■ Instant Recovery backup <p>See “Snapshot Client and Replication Director (policy attributes)” on page 740.</p>
Standard (policy type)	<p>Checkpoint restart is supported for UNIX clients.</p>
Synthetic backups (schedule attribute)	<p>Checkpoint restart is not supported.</p> <p>See “Synthetic backup (schedule attribute)” on page 754.</p>

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	<p>This Clean-up host property indicates the number of days that a failed restore job can remain in an Incomplete state.</p> <p>See “Clean-up properties” on page 80.</p>
Restore retries	<p>This Universal Setting host property specifies the number of attempts that a client has to restore after a failure.</p> <p>See “Universal Settings properties” on page 249.</p>

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 689.

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 20-4 Limit jobs per policy attribute



Client A and Client B backups can occur concurrently and back up to different devices

This attribute operates differently for VMware policies, depending on how the policy selects virtual machines.

For more information, see the [NetBackup for VMware Administrator's Guide](#).

[Table 20-11](#) describes the factors that affect the number of concurrent backup jobs that NetBackup can perform.

Table 20-11 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>

Table 20-11 Factors affecting the number of concurrent backup jobs (*continued*)

Item	Description
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>
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 703.</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 are examples of the jobs that 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 1020.</p> <p>This attribute operates differently for VMware policies, depending on how the policy selects virtual machines.</p> <p>For more information, see the NetBackup for VMware Administrator’s Guide.</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 113.

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 511.

See [“Configuring a NetBackup server group”](#) on page 367.

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 **NetBackup 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. Clear the check box to deactivate the policy.

See [“Creating a catalog archiving policy”](#) on page 890.

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 254.

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 20-12 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 701.</p> <p>See “Cross mount points (policy attribute)” on page 699.</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>

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.

<i>master1</i>	NetBackup master server
<i>win_client</i>	Windows NetBackup client
<i>win_PC</i>	Windows computer (not necessarily a NetBackup client)
<i>TestData</i>	A shared folder on <i>win_PC</i>

Table 20-13 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> .
Step 2	Add the folder name to the policy	Add \\win_PC\TestData to the file list of the policy. This step is not necessary if the policy is only used for user-directed backups.
Step 3	Configure the NetBackup Client Service	<p>Perform the following actions:</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. <p>See “Configuring the NetBackup Client Service” on page 1092.</p>
Step 4	Perform a backup	<p>Backups run as scheduled or when a manual backup is performed.</p> <p>See “Performing manual backups” on page 842.</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.

<i>master1</i>	NetBackup master server
<i>win_client</i>	Windows NetBackup client
<i>win_PC</i>	Windows computer (not necessarily a NetBackup client)
<i>share</i>	A shared folder on <i>win_PC</i>

Table 20-14 Using Backup Network Drives 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> , and check Backup network drives in the policy attributes tab.
Step 2	Configure the NetBackup Client Service	<p>Perform the following actions:</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. <p>See “Configuring the NetBackup Client Service” on page 1092.</p>
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>Backups run as scheduled or when a manual backup is performed.</p> <p>See “Performing manual backups” on page 842.</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.

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 20-15 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 701.</p> <p>See “Follow NFS (policy attribute)” on page 696.</p>
Backup selection entries	<p>The following backup selection entries behave in the following manner for Windows and UNIX computers when the Cross mount points attribute is used:</p> <ul style="list-style-type: none"> ■ <code>/</code> Valid for UNIX clients. For Windows clients, the forward slash expands to <code>ALL_LOCAL_DRIVES</code>. ■ <code>:\</code> Valid for Windows clients. For UNIX clients, this entry creates a status 69 (Invalid filelist specification). ■ <code>*:\</code> Valid for Windows clients. For UNIX clients, this entry creates a status 69 (Invalid filelist specification).
UNIX raw partitions	<p>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 Cross mount points is enabled.</p>
<code>ALL_LOCAL_DRIVES</code> directive	<p>Do not use Cross mount points in policies on UNIX systems where you use the <code>ALL_LOCAL_DRIVES</code> directive in the backup selection list.</p> <p>Enabling Cross mount points can cause multiple backups of mounted volumes.</p>

Table 20-15 Considerations for Cross mount points (policy attribute) (*continued*)

Item	Description
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 disk storage unit 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 20-16](#) summarizes the possible results.

Table 20-16 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 20-5](#).

Figure 20-5 Example configuration of client disks




Table 20-17 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.

Note: The **Compression** attribute is independent of the VxFS compression state.

See [“About the backup and restore of compressed files on VxFS file systems”](#) on page 1163.

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 [“Media multiplexing \(schedule attribute\)”](#) on page 769.

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 20-18](#) describes the various factors to consider when you choose to use **Compression**.

Table 20-18 Considerations regarding 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	File size has no effect on the amount of compression. However, it takes longer to compress many small files than a single large one.
Client resources that are required	Compression requires client computer processing unit time and as much memory as the administrator configures.
Effect on client performance	Compression uses as much of the computer processing unit as available and affects other applications that require the computer processing unit. For fast CPUs, however, I/O rather than CPU speed is the limiting factor.
Files that are not compressed	<p>NetBackup does not compress the following files:</p> <ul style="list-style-type: none"> 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: <div> <div>.arc</div> <div>.gz</div> <div>.iff</div> <div>.sit.bin</div> <div>.arj</div> <div>.hqx</div> <div>.pit</div> <div>.tiff</div> <div>.au</div> <div>.hqx.bin</div> <div>.pit.bin</div> <div>.Y</div> <div>.cpt</div> <div>.jpeg</div> <div>.scf</div> <div>.zip</div> <div>.cpt.bin</div> <div>.jpg</div> <div>.sea</div> <div>.zom</div> <div>.F</div> <div>.lha</div> <div>.sea.bin</div> <div>.zoo</div> <div>.F3B</div> <div>.lzh</div> <div>.sit</div> <div>.z</div> <div>.gif</div> <div>.pak</div> </div> 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	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.

Note: If compressed data is written to a storage unit that has deduplication capabilities, the storage unit may not be able to deduplicate the data.

Encryption (policy attribute)

The **Encryption** attribute determines whether the backup should be encrypted by the client. 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 118.

For additional encryption configuration information, see the [NetBackup Security and Encryption Guide](#).

Note: If encrypted data is written to a storage unit that has deduplication capabilities, the storage unit may not be able to deduplicate the encrypted data.

About NetBackup encryption options

NetBackup provides several methods for encrypting backups, as described in the following table.

Table 20-19 NetBackup encryption options

Option	Description
Client encryption	<p>The NetBackup client encryption option is a software-based solution that encrypts the data on the client. The data is encrypted in transit and at rest. Each client manages its own encryption keys.</p> <p>To enable client encryption, select the backup policy Encryption attribute.</p> <p>See “Encryption (policy attribute)” on page 705.</p>

Table 20-19 NetBackup encryption options (*continued*)

Option	Description
Tape drive encryption	<p>With hardware-based tape drive encryption, an encrypting tape drive encrypts the data. The data is encrypted at rest only. The NetBackup Key Management Service manages the encryption keys.</p> <p>See the “Data at rest key management” chapter in the NetBackup Security and Encryption Guide.</p> <p>One method to manage the volumes for hardware-based tape encryption is to use a reserved prefix on the volume pool name. The storage device must have encrypting tape drives. The storage unit must specify the storage device that has the encrypting tape drives. The backup policy must specify the correct storage unit and volume pool.</p> <p>See “About reserved volume pool name prefixes” on page 442.</p>
Software-based media server encryption	<p>With software-based encryption, an agent on the media server encrypts the data. The data is encrypted at rest only. The Media Server Encryption Option uses an external agent that encrypts the data and manages the encryption keys.</p> <p>See the NetBackup Media Server Encryption Option Administrator's Guide.</p>
AdvancedDisk encryption	<p>A plug-in in the NetBackup OpenStorage stack encrypts the data. The data is encrypted at rest only. The NetBackup Key Management Service manages the encryption keys.</p> <p>See the NetBackup AdvancedDisk Storage Solutions Guide.</p>
Cloud storage encryption	<p>A plug-in in the NetBackup OpenStorage stack encrypts the data. The data is encrypted at rest only (by default, NetBackup uses SSL for read and write operations). The NetBackup Key Management Service manages the encryption keys.</p> <p>See the NetBackup Cloud Administrator's Guide.</p>
Media Server Deduplication Pool encryption	<p>The MSDP deduplication plug-in encrypts the data. The data can be encrypted in transit and at rest or at rest only. The NetBackup deduplication plug-in manages the encryption keys.</p> <p>See the NetBackup Deduplication Guide.</p>

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.

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 more information, see the [Bare Metal Restore Administrator's Guide](#).

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.

See [“Synthetic backup \(schedule attribute\)”](#) on page 754.

The following examples show how move detection backs up the files that otherwise would not be backed up:

- On Windows:
 - 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.
- On UNIX:
 - A file that is named `/home/pub/doc` is moved to `/home/spec/doc`. The modification time is unchanged but `/home/spec/doc` is new in the `/home/spec/` directory and is backed up.
 - A directory that is named `/etc/security/dev` is renamed as `/etc/security/devices`. The modification time is unchanged but `/etc/security/devices` is a new directory and is backed up.
 - A file that is named `/home/pub/doc` is installed when extracted from a UNIX `tar` file. The modification time is before the time of the last backup. The file is new in the `/home/pub/` directory and is backed up.
 - A file that is named `docA` is removed, then a file that is named `docB` is renamed as `docA`. The new `docA` has the same name but since its inode number has changed, it 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 `/home/abc/doc/` directory of a UNIX client during a series of backups between 12/01/2011 and 12/04/2011. **Collect true image restore information** was enabled for the policy that performed the backups.

Table 20-20 Sample backups taken before a true image restore

Day	Type of backup	Files that are backed up in <code>/home/abc/doc</code>
12/01/2011	Full	file1 file2 dirA/fileA dirB/fileB file3
12/02/2011	Incremental	file1 file2 dirA/fileA -----
12/03/2011	Incremental	file1 file2 dirA/fileA -----
12/04/2011	User backup	file1 file2 dirA/fileA ----- dirC/fileC file4
12/04/2011	Incremental	file1 file2 ----- file4

Note: Dashes (-----) indicate that the file was deleted before this backup.

A restore of the 12/04/2011 version of the `/home/abc/doc/` directory produces following results:

After a regular restore The restored directory contains all files and directories that ever existed in `/home/abc/doc/` from 12/01/2011 (last full backup) through 12/04/2011:

```
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 12/04/2011 incremental backup.

The restored directory does not include the subdirectories `dirA` and `dirC`, even though they were backed up on 12/04/2011 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 80.
- 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 12/04/2011 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. Each child job displays its own job ID in the **Job ID** column in the **Activity Monitor**. The job ID of the parent job appears in the **Parent Job ID** column, which is not displayed by default. Parent jobs display a dash (-) in the **Schedule** column.

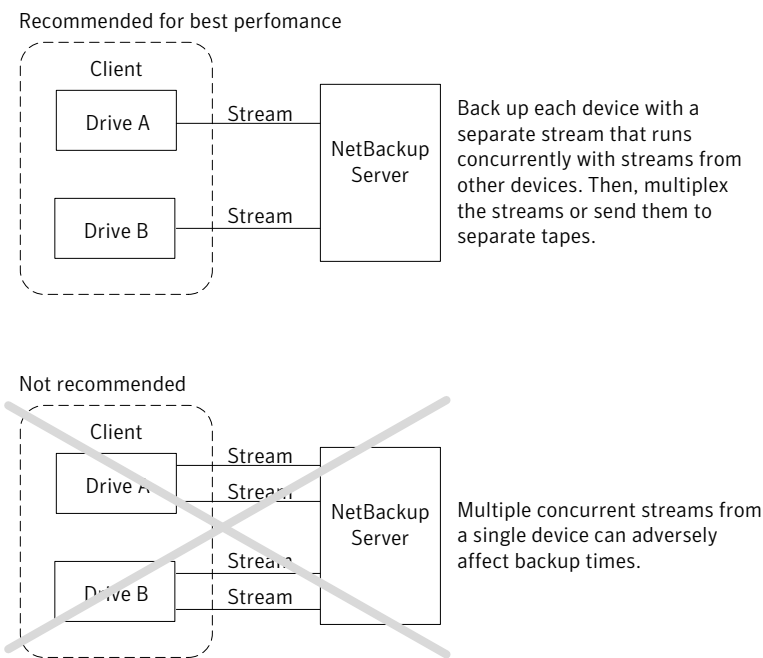
Note: If this attribute is enabled, and a file system is in a client's exclude list, a NetBackup job appears in the **Activity Monitor** for the excluded file system. However, no files in the excluded file system are backed up by the job.

The following table describes the reasons to use multiple data streams.

Table 20-21 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 20-6 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.</p> <p>See “Global Attributes properties” on page 150.</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 20-6 Multiple stream recommendations



The following table describes the aspects of multiple data streams that are adjustable.

Table 20-22 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 818.</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 20-22 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 769.</p>

The maximum jobs settings limit the maximum number of streams as follows:

Table 20-23 Job settings that limit the maximum number of streams

Item	Access method
Maximum jobs per client (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 Global Attributes. <p>See “Global Attributes properties” on page 150.</p> <p>See “Media multiplexing (schedule attribute)” on page 769.</p>
Limit jobs per policy (policy attribute)	<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 693.</p>
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 86.</p>

Job settings also affect the maximum number of streams. The following table describes the interdependency of these settings.

Table 20-24 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 818.

Disable client-side deduplication (policy attribute)

The **Disable client-side deduplication** attribute appears only if the NetBackup Data Protection Optimization Option license key is active.

The clients that are configured for client-side deduplication 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 send their data to a deduplication media server. That server deduplicates the data and then sends it to the storage server.
Disabled	The clients deduplicate their own 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-side 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 90.

See the [NetBackup Deduplication Guide](#).

Enable granular recovery (policy attribute)

The **Enable granular recovery** attribute is available 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

NetBackup does not support the compression or encryption of backups that use Granular Recovery Technology (GRT). When the **Enable granular recovery** option is enabled, the **Compression** option and the **Encryption** option are automatically disabled.

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, see the following:

- [NetBackup for SharePoint Server Administrator's Guide](#)
- [NetBackup for Exchange Server Administrator's Guide](#)
- See "[Active Directory granular backups and recovery](#)" on page 843.

Use Accelerator (policy attribute)

The NetBackup Accelerator increases the speed of full backups. The increase in speed is made possible by change detection techniques on the client. The client uses the change detection techniques and the client's current file system to identify the changes that occurred since the last backup. The client sends the changed data to the media server in a more efficient backup stream. The media server combines the changed data with the rest of the client's data that is stored in previous backups.

If a file or portion of a file is already in storage and has not been changed, the media server uses the copy in storage rather than reading it from the client. The end result is a full NetBackup backup.

Accelerator has the following advantages:

- Reduces the I/O and CPU overhead on the client. The result is a faster backup and less load on the client.
- Creates a compact backup stream that uses less network bandwidth between client and server.
- Creates a full image that contains all data that is needed for restore.

Note: The Accelerator operates differently when used for backup of virtual machines. For full details on the Accelerator for VMware, see the [NetBackup for VMware Administrator's Guide](#).

How the NetBackup Accelerator works

The NetBackup Accelerator creates the backup stream and backup image as follows:

- If the client has no previous backup, NetBackup performs a full backup and creates a track log. The track log contains information about the client's data, for comparison at the next backup.
- At the next backup, NetBackup identifies data that has changed since the previous backup. To do so, it compares information from the track log against information from the file system for each file. For NTFS and ReFS file systems, it also uses the Windows change journal to help identify the data that has changed since the last backup.

Accelerator uses the Windows change journal in two ways: To check for changes in the file system metadata, and to help detect which files have changed since the last backup.

See [“Accelerator and the Windows change journal”](#) on page 719.

- The NetBackup client sends to the media server a tar backup stream that consists of the following: The client's changed blocks, and the previous backup ID and data extents (block offset and size) of the unchanged blocks.
- The media server receives the client's changed blocks and the backup ID and data extents of the unchanged blocks. From the backup ID and file system descriptors, the media server locates the rest of the client's data in existing backups.
- The media server directs the storage server to write the changed blocks and the unchanged blocks in a new full image.

[Figure 20-7](#) shows how an Accelerator backup stream is composed.

Figure 20-7 NetBackup client: Accelerator backup stream

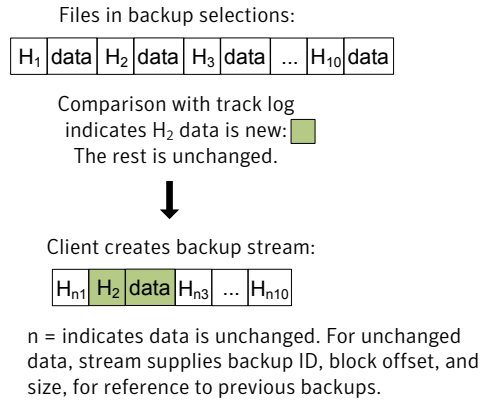


Figure 20-8 shows how the media server constructs a backup image from the Accelerator stream and from previous backups:

Figure 20-8 NetBackup media server constructs backup image

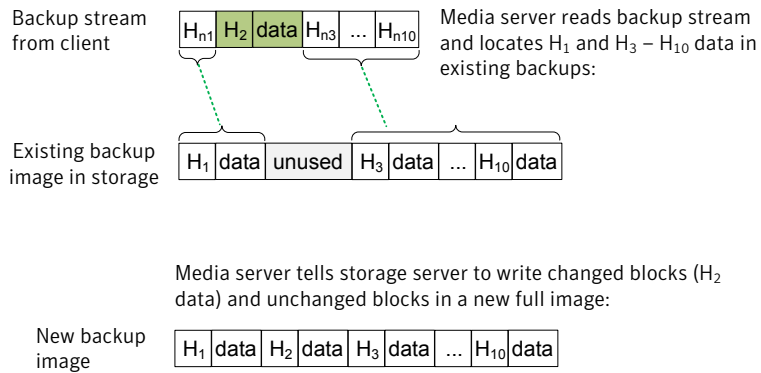
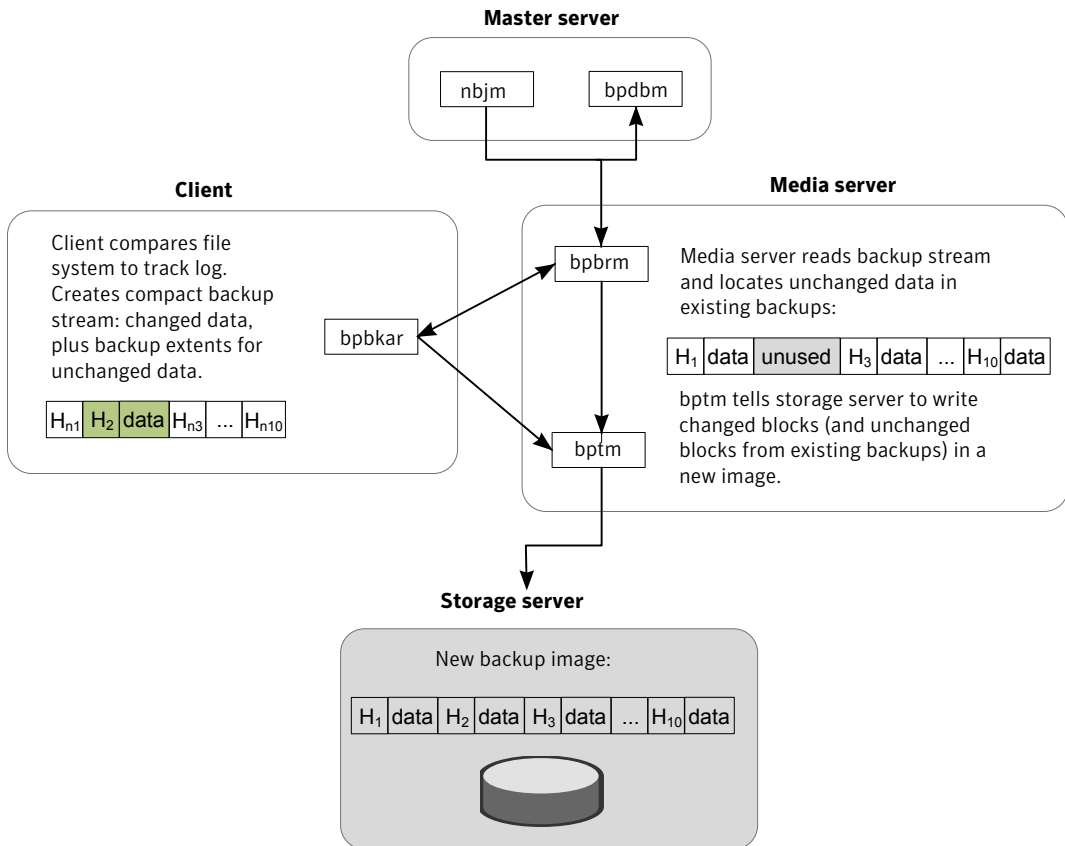


Figure 20-9 recaps Accelerator details in the context of the principal NetBackup processes.

Figure 20-9 Process overview of Accelerator backup



Accelerator and the Windows change journal

For Windows NTFS and ReFS file systems, the Accelerator uses the Windows change journal to help identify the files that changed since the previous backup.

The extent to which the Windows change journal is used depends on the following:

- Whether the **Use Change Journal** option has been enabled for the NetBackup client.
- Whether the policy contains a schedule with the **Accelerator forced rescan** option enabled.

Note: Regardless of the **Use Change Journal** setting or the **Accelerator forced rescan** setting, the Windows change journal always helps determine if a file has changed.

Table 20-25 How Accelerator uses the Windows change journal

NetBackup configuration	Accelerator use of change journal
The Use Change Journal option is not enabled	For full backups and incremental backups, NetBackup examines the metadata for every file to determine if the file has changed. Even without the Use Change Journal option, the Windows change journal helps determine if a file has changed.
The Use Change Journal option is enabled but the Accelerator forced rescan option is not enabled on any of the policy's schedules	<p>The Windows change journal helps determine if a file has changed during both full backups and incremental backups.</p> <ul style="list-style-type: none"> During incremental backups, the Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. During full backups, NetBackup examines the metadata for every file, to determine whether the file has changed. In addition, NetBackup uses the Windows change journal to help determine whether a file has changed.
The Use Change Journal option is enabled and the Accelerator forced rescan option is enabled on at least one of the policy's schedules	<p>The Windows change journal helps determine if a file has changed during both full backups and incremental backups.</p> <ul style="list-style-type: none"> During incremental backups, the Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. During full backups: <ul style="list-style-type: none"> For a full backup schedule that does not enable the Accelerator forced rescan option: The Windows change journal is used exclusively to determine if a file has changed. NetBackup does not examine the file metadata. For a full backup schedule that enables the Accelerator forced rescan option: NetBackup examines the metadata for every file, to determine whether the file has changed. In addition, NetBackup uses the Windows change journal to help determine whether a file has changed.

Accelerator notes and requirements

Note the following about the NetBackup Accelerator:

- NetBackup Accelerator requires the Data Protection Optimization Option license. For the latest information on licensing, contact your Symantec sales or partner representative.

- Supports the disk storage units only. Supported storage includes **Media Server Deduplication Pool**, NetBackup appliance, cloud storage, and qualified third-party OST storage. See the [NetBackup Hardware Compatibility List](#) for supported storage types:
- Storage unit groups are supported only if the storage unit selection in the group is Failover.
- Supports the MS-Windows and Standard policy types. Supports all features of NetBackup that work with the MS-Windows or Standard policy types.

Note: Accelerator for virtual machines operates differently from Accelerator for physical computers. For all details on Accelerator for virtual machines, see the [NetBackup for VMware Administrator's Guide](#).

Note: Accelerator operates differently when used for backup of virtual machines, NDMP for NetApp filers, and Oracle databases.

For full details on Accelerator for VMware, see the [NetBackup for VMware Administrator's Guide](#).

For full details on Accelerator for NDMP for NetApp filers, see the [NetBackup for NDMP Administrator's Guide](#).

For full details on Accelerator for Oracle, see the [NetBackup for Oracle Administrator's Guide](#).

- Accelerator backups ignore the policy attribute that resets the `atime` on files after backup.
By default, NetBackup records the file access time (`atime`) for each UNIX file that it backs up, and then resets the `atime` after the file is backed up. Resetting the `atime` in this way causes the change time (`ctime`) to update as well. This behavior is normally controlled by the following policy attribute: **Reset file access time to the value before backup**. (Default: enabled.)
When NetBackup Accelerator is used to perform backups, NetBackup does not reset the `atime` after the file is backed up, regardless of the policy attribute setting. NetBackup does not reset the `atime` (which avoids updating the `ctime`) because a `ctime` update would require a loss in Accelerator optimization.
See ["Client Settings properties for UNIX clients"](#) on page 99.
- Supports the full backups and incremental backups.
See ["Accelerator backups and the NetBackup catalog"](#) on page 723.
- Supports all platforms, file systems, and logical volumes that NetBackup supports.

- Supports the Windows NTFS and ReFS change journal (**Use Change Journal**) but does not support the VxFS change journal.
- For every policy that enables the **Use Accelerator** option, Symantec recommends at a minimum the following backup schedules:
A full backup schedule with the **Accelerator forced rescan** option enabled.
Another full backup schedule without the **Accelerator forced rescan** option enabled.
See [“Accelerator forced rescan option \(schedule attribute\)”](#) on page 755.
- If **Collect true image restore information** is enabled in the policy, NetBackup performs a regular file system scan to determine the files and folders to include in the backup. It also queries the change journal to help determine which files have changed.
More information is available on the **Use change journal** option and the Accelerator:
See [“Accelerator and the Windows change journal”](#) on page 719.
- If a previous backup of the client does not exist, NetBackup performs a full backup and creates a track log on the client. This initial backup occurs at the speed of a normal (not accelerated) full backup. Subsequent Accelerator backups of the client use the track log for accelerated backup speed.

Note: When you first enable a policy to use Accelerator, the next backup (whether full or incremental) is in effect a full backup: It backs up all files in the **Backup Selections** tab. If that backup was scheduled as an incremental, it may not complete within the backup window. For example, when you upgrade policies to 7.7 and first enable the **Use Accelerator** option, note: Make sure that the next backup window is large enough for a full backup.

- NetBackup retains track logs for future Accelerator backups. Whenever you change the policy's backup selections list, NetBackup does a full non-accelerated backup of the selections that were changed or added in the list. The unchanged backup selections are processed as normal Accelerator backups.
- If the storage unit that is associated with the policy cannot be validated when you create the policy, note: The storage unit is validated later when the backup job begins. If Accelerator does not support the storage unit, the backup fails. In the `bpbrm` log, a message appears that is similar to one of the following:

```
Storage server %s, type %s, doesn't support image include.
```

```
Storage server type %s, doesn't support accelerator backup.
```

- Accelerator requires that the storage have the `OptimizedImage` attribute enabled.
- Symantec recommends that you not enable **Expire after copy** retention for any storage units that are used with storage lifecycle policies (SLP) in combination with Accelerator. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.

Accelerator backups and the NetBackup catalog

Use of Accelerator does not affect the size of the NetBackup catalog. A full backup with Accelerator generates the same catalog size as a full backup of the same data without Accelerator. The same is true of incremental backups: use of Accelerator does not require more catalog space than the same backup without Accelerator.

A potential catalog effect does exist, depending on how often you use Accelerator with full backups. A full backup with Accelerator completes faster than a normal full. It may therefore be tempting to replace your incremental backups with Accelerator full backups. Note: Since a full backup requires more catalog space than an incremental, replacing incrementals with fulls increases the catalog size. When changing your incrementals to fulls, you must weigh the advantage of Accelerator fulls against the greater catalog space that fulls require compared to incrementals.

Configuring Accelerator

The following table outlines the procedure to configure the full backups that use the NetBackup Accelerator.

Table 20-26 To configure Accelerator on full backups

Task	Procedure and notes
Make sure that you have a storage unit that supports Accelerator.	<p>Accelerator supports disk storage units only. Supported storage unit types are PureDisk (Media Server Deduplication Pool and NetBackup appliance), cloud storage, and qualified third-party OST storage.</p> <p>The NetBackup device mapping files contain a complete list of supported storage units (see the next task).</p> <p>Note: Storage unit groups are supported only if the storage unit selection in the group is Failover.</p>

Table 20-26 To configure Accelerator on full backups (*continued*)

Task	Procedure and notes
Update the NetBackup device mapping files if needed.	<p>The NetBackup device mapping files contain all storage device types that NetBackup can use. To add support for the new devices or upgraded devices that support Accelerator, download the current device mapping files from the Symantec support site.</p> <p>See “About the device mapping files” on page 394.</p> <p>See “Downloading the device mapping files” on page 394.</p>
<p>Configure a Standard, MS-Windows or VMware backup policy.</p> <p>Configure a Standard, MS-Windows, VMware, NDMP or Oracle backup policy.</p>	<p>Select the following on the policy Attributes tab:</p> <ul style="list-style-type: none"> ■ A disk storage unit that supports Accelerator. ■ The Use Accelerator option. <p>Accelerator operates differently when used for backup of virtual machines, NDMP for NetApp filers, and Oracle databases.</p> <ul style="list-style-type: none"> ■ For full details on Accelerator for VMware, see the NetBackup for VMware Administrator's Guide. ■ For full details on Accelerator for NDMP for NetApp filers, see the NetBackup for NDMP Administrator's Guide. ■ For full details on Accelerator for Oracle, see the NetBackup for Oracle Administrator's Guide.
To periodically establish a new baseline of change detection on the client, you can select the Accelerator forced rescan option on the Schedule Attribute tab of the policy.	<p>This option provides an additional level of Accelerator change detection in the client's data. This option reduces the speed of Accelerator.</p> <p>See “Accelerator forced rescan option (schedule attribute)” on page 755.</p> <p>See “Accelerator and the Windows change journal” on page 719.</p>

Accelerator messages in the backup job details log

A NetBackup backup that uses Accelerator writes a message similar to the following in the job details log:

```
11/23/2011 10:50:27 AM - Info bpbrm(pid=412) accelerator enabled
```

When NetBackup uses the NTFS or ReFS change journal for the backup, messages similar to the following appear in the job details log:

```
9/24/2011 8:54:14 PM - Info bpbkar32(pid=7868) change journal enabled  
for <C:\>
```


9/24/2011 8:54:14 PM - Info bpbkar32(pid=7868) using change journal data for <C:\>

If the NTFS or ReFS change journal cannot be used, a message of the following form appears in the job details log:

not using change journal data for <backup selection>: <reason ...>

See [“Log messages about the Use Change Journal option and Accelerator”](#) on page 726.

When the **Accelerator forced rescan** option is used, a message similar to the following appears in the job details log:

9/25/2011 5:46:52 PM - Info bpbrm(pid=4136) Accelerator enabled backup with checksum based change detection needs to read each file and calculate the checksum, and will have longer backup time.

See [“Log messages about the Use Change Journal option and Accelerator”](#) on page 726.

NetBackup logs for Accelerator

For log messages about Accelerator, see the following NetBackup log directories.

Table 20-27 NetBackup logs that pertain to Accelerator

Log directory	Resides on
UNIX: /usr/opensv/netbackup/logs/bpbrm Windows: <i>install_path</i> \NetBackup\logs\bpbrm	NetBackup master or media server
UNIX: /usr/opensv/netbackup/logs/bptm Windows: <i>install_path</i> \NetBackup\logs\bptm	NetBackup media server
UNIX: /usr/opensv/netbackup/logs/bpbkar Windows: <i>install_path</i> \NetBackup\logs\bpbkar	NetBackup client

To create the log directories, run the following command on the NetBackup servers and client:

On Windows:

install_path\NetBackup\logs\mklogdir.bat

On UNIX:

/opt/opensv/netbackup/logs/mklogdir

Log messages about the Use Change Journal option and Accelerator

The Accelerator uses the NTFS or ReFS change journal to help identify data changes on the Windows client. The following table describes the change journal messages that may appear in the NetBackup job details log in the Activity Monitor. The left column lists the messages.

For Accelerator backups, these messages indicate various impediments to the use of the Windows change journal. In general, they indicate either of the following:

- That the Accelerator cannot use the change journal for the backup. To use the change journal, you may need to reconfigure the policy.
- That the Accelerator cannot exclusively use the change journal to detect changes in the file system. NetBackup examines the file system metadata for each file, to detect which files have changed. It also uses the change journal to help detect which files have changed.

Note: In the following messages, the variable <%%s> represents the items in your backup selections list.

The change journal messages often appear in pairs. The first message indicates why the change journal cannot be used. For example:

not using change journal data for <%%s>: forced rescan requested

The second message indicates that the Accelerator nevertheless can use the change journal to detect changed files:

not using change journal data for enumeration for <%%s> but will use it for change detection

Table 20-28 Accelerator messages on the Windows Use Change Journal option (job details log)

Message in NetBackup job details	Description
not using change journal data for enumeration for <%%s> but will use it for change detection	<p>NetBackup uses the change journal to help detect which files have changed. NetBackup also examines the file system metadata for each file to help detect which files have changed.</p> <p>Note: Before this message appears, another message explains why NetBackup does not rely entirely on the change journal data to detect changed files.</p>

Table 20-28 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: forced rescan requested	The Accelerator forced rescan option is enabled on the full backup schedule for the policy. As a result, NetBackup cannot rely entirely on the change journal data to detect changed files. NetBackup also examines the file system metadata for each file to help detect which files have changed.
not using change journal data for <%%s>: filter checksum calculation failed	<p>The backup uses a number of filters to determine which files and directories to include in the backup. The filters are the following: NetBackup exclude and include lists, the files and directories that are included in the Shadow Copy Components and in the system state backup, and others.</p> <p>During a backup, a checksum is calculated against the filters. When a new backup runs, the checksum for the current backup is compared against the checksum of the previous backup. If the checksums do not match, the change journal data cannot be used. Instead, Accelerator performs a regular file system scan to determine the files and folders to include in the backup.</p> <p>No action is required. If the filters do not change between backups, the change journal data is used at the next backup.</p>
not using change journal data for <%%s>: unable to validate change journal usage <reason=%%s>	<p>Accelerator cannot use the Use Change Journal option in the following cases:</p> <ul style="list-style-type: none"> ■ No previous backup exists: No baseline update sequence number (USN) was established from the NTFS change journal. ■ The previous backup was not successful. <p>In these cases, the Use Change Journal option cannot be used. Accelerator performs a regular file system scan to determine the files and folders to include in the backup.</p> <p>No action is required. Accelerator uses the Use Change Journal option at the next backup if all conditions are met.</p>
not using change journal data for <%%s>: unable to initialize change journal usage <reason=%%s>	<p>Accelerator cannot use the Use Change Journal option in the following cases:</p> <ul style="list-style-type: none"> ■ Too much activity has occurred on the system (records were purged from the change journal databases before they could be processed). ■ Data corruption occurred. <p>Instead, Accelerator performs a regular file system scan to determine the files and folders to include in the backup.</p> <p>No action is required. When records have been purged, a new baseline is created when the current backup completes. If corruption existed, it is detected and the databases are re-created. Accelerator uses the Use Change Journal option at the next backup.</p>

Table 20-28 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: hard link or reparse point change detected	<p>Accelerator cannot use the Use Change Journal option if a change corresponds to a hard link or to a reparse point. Note that the change may correspond to any number of files and directories and the NTFS change journal does not track all of them.</p> <p>No action is required. If no further changes occur to hard links or reparse points, the Use Change Journal option can be used at the next backup.</p>
not using change journal data for <%%s>: not supported with true image backups	<p>Collect true image restore information or Collect true image restore information with move detection is specified in the policy. To process the files and determine which files have changed, NetBackup performs a regular file system scan to determine the files and folders to include in the backup. It also queries the change journal to help determine which files have changed.</p>
not using change journal data for <%%s>: invalid schedule type	<p>Accelerator does not support the selected schedule type with the Use Change Journal option.</p> <p>The Use Change Journal option is supported for incremental backups (cumulative or differential) or full backups. For full backups, Use Accelerator must be enabled on the policy Attributes tab.</p>
not using change journal data for <%%s>: path must be local and not contain directory junctions and/or symbolic links	<p>In the backup selections, a path contains a reparse point (directory junction or a symbolic link). The Use Change Journal option cannot be used. NetBackup must perform a regular file system scan to back up the directories correctly.</p>
not using change journal data for <%%s>: change journal setup is not complete (may still be in progress)	<p>The Use Change Journal option was recently enabled. After Use Change Journal is first enabled, the NetBackup client may need time to process the NTFS change journal and set up its databases.</p> <p>The Use Change Journal option may be ready at the next backup.</p>
not using change journal data for <%%s>: unable to locate journal data	<p>The Use Change Journal option was recently enabled. After Use Change Journal is first enabled, the NetBackup client may need time to process the NTFS change journal and set up its databases.</p> <p>The Use Change Journal option may be ready at the next backup.</p>

Table 20-28 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: invalid change journal data	<p>When many changes occur on a volume, the NetBackup Client Service may need to increase the size of the internal databases. As a result, the databases may become invalid.</p> <p>After the databases are increased in size and are synchronized with the NTFS change journal, they are marked as valid. The Use Change Journal option is used with the next backup.</p>
not using change journal data for <%%s>: unable to apply snapshot	<p>For a snapshot-based backup, NetBackup uses the change journal databases on the snapshot instead of on the live volume. An error occurred when NetBackup attempted to open the databases on the snapshot.</p> <p>When the next backup runs, NetBackup creates a new snapshot and the databases may be opened without error.</p>
not using change journal data for <%%s>: no previous track log	<p>No previous full backup exists that used Accelerator.</p> <p>NetBackup supports the Use Change Journal option with Accelerator on a full backup only if a previous full backup exists that used Accelerator.</p>
not using change journal data for <%%s>: not supported with regular full backups	<p>The Use Accelerator option is not enabled on the policy.</p> <p>For full backups, the Use Change Journal option requires that Use Accelerator is enabled.</p>
not using change journal data for <%%s>: unable to validate change journal usage <reason=previous backup wasn't a successful backup>	<p>If a backup is partially successful (status code 1), the next Accelerator backup cannot use the Use Change Journal option.</p> <p>NetBackup can use the Use Change Journal option after the next successful backup.</p>
not using change journal data for <%%s>: not supported	<p>A backup selection in the policy is a resource for which the Windows change journal is not supported. Examples of unsupported resources are FAT volumes and FAT32 volumes.</p>

Table 20-28 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: not supported for remote frozen images	The backup has attempted to use a remote frozen image. The Use Change Journal option is not supported with remote frozen images.
not using change journal data for <%%s>: not enabled	The Use Change Journal option is not enabled.
not using change journal data for <%%s>: not configured for use	The Use Change Journal option is not enabled.
not using change journal data for <%%s>: unable to determine volume guid	An internal error occurred. The volumes to back up should be checked to ensure that a volume GUID is associated with each volume.
not using change journal data for <%%s>: snapshot has not been applied (unable to track open files)	To use the change journal data, NetBackup must be able to take a snapshot of the backup volume to correctly identify and handle open files. This error occurs if a backup runs before NetBackup can take a snapshot of the volume. If a snapshot can be taken before the next backup occurs, NetBackup may be able to use change journal data at the next backup.
not using change journal data for <%%s>: invalid policy type	The Use Change Journal option is only supported for Standard or MS-Windows policies.
not using change journal data for <%%s>: list of accelerator based backups does not match list of change journal based backups	<p>On the client, NetBackup keeps track of both Accelerator backups and change journal-based backups. If the lists of these backups do not match, one of the following occurred:</p> <ul style="list-style-type: none"> ■ An Accelerator backup occurred that did not use the change journal. ■ A change journal-based backup occurred that did not use the Accelerator. <p>If either case, the Use Change Journal option cannot be used until the next full backup occurs.</p>

Table 20-28 Accelerator messages on the Windows **Use Change Journal** option (job details log) (*continued*)

Message in NetBackup job details	Description
not using change journal data for <%%s>: previous backup had change journal data that was not complete (missing usn records)	The change journal databases are fixed in size: they can contain only a fixed number of records. If the volume undergoes a lot of activity, records may be purged from the change journal database before a backup can process them. In that case, the Use Change Journal option cannot be used. No action is required. Accelerator uses the Use Change Journal option at the next backup if all conditions are met.
not using change journal data for <%%s>: not supported for non-local volumes / file systems	The backup selection in the policy is not for a local volume. For example, the backup selection specifies a UNC path.
not using change journal data for <%%s>: no forced rescan schedule configured	This message appears only during Accelerator full backups. To use the change journal data during Accelerator full backups, a policy schedule with the Accelerator forced rescan option is required.
not using change journal data for <%%s>: forcing rescan, each file will be read in order to validate checksums	This message appears only during full backups, when the policy is not configured with the Accelerator forced rescan option. Before this message appears, another message explains why the change journal data cannot be used.

About reporting the amount of Accelerator backup data that was transferred over the network

For Accelerator backup reporting, several NetBackup commands can report the amount of data that is transferred over the network for each Accelerator backup. The amount of transferred data is often much less than the size of the Accelerator backup image.

For each Accelerator backup, NetBackup combines the client's (or VM's) changed blocks with the unchanged data from previous backups to synthesize a backup image. However, NetBackup sends only the changed data over the network when the backup occurs. The resulting backup image may be much larger than the amount of backup data that travels the network. For backup reporting, it may be important

to distinguish between the backup image size and the amount of data that was transferred over the network.

For Accelerator backups, the network-transferred data can appear in the output of the following NetBackup commands: `bpdbjobs`, `bpimagelist`, and `bpclimagelist`.

Table 20-29 lists the default location of these commands.

Table 20-29 Default location of `bpdbjobs`, `bpimagelist`, and `bpclimagelist`

Command	Default location
<code>bpdbjobs</code> , <code>bpimagelist</code>	Windows: <code><install_path>\NetBackup\bin\admincmd\ UNIX, Linux <code>/usr/opensv/netbackup/bin/admincmd/</code></code>
<code>bpclimagelist</code>	Windows: <code><install_path>\NetBackup\bin\ UNIX, Linux <code>/usr/opensv/netbackup/bin/</code></code>

The following example uses the `bpimagelist` command to show the results of a backup of `acmevm2`:

```
bpimagelist -backupid acmevm2
```

Example output:

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

In this example, the backup image size in kilobytes is 7799632, and the amount of data that was transferred over the network is 225792.

You can use the following commands to show the amount of data that was transferred over the network for an Accelerator backup.

bpimagelist

```
bpimagelist -backupid backup_id [-l | -L | -json | -json_compact]
```

Brackets [] indicate optional elements, and the vertical bars | indicate that you can choose only one of the options within the brackets.

[Table 20-30](#) describes how the network-transferred data field appears in the `bpimagelist` output:

Table 20-30 **bpimagelist** options that show the amount of network-transferred data for Accelerator backups

bpimagelist option	How the network-transferred data field appears
no option	The field is unlabeled. For example: 225792 See the <code>bpimagelist</code> example output earlier in this topic.
-l	The field is unlabeled (same as no option). For example: 225792
-L	Kilobytes Data Transferred: 225792
-json	"kilobytes_data_transferred": 225792,
-json_compact	"kilobytes_data_transferred":225792,

bpdbjobs

```
bpdbjobs -jobid job_id -report -most_columns
```

or

```
bpdbjobs -jobid job_id -report -all_columns
```

The network-transferred data field appears at the end of the output.

bpclimagelist

```
bpclimagelist -client client_name
```

This command can only show the network-transferred data in the field that normally shows the Accelerator backup image size. To show the network-transferred data with this command, you must configure a NetBackup setting:

See [“Replacing the Accelerator image size with the network-transferred data in NetBackup command output”](#) on page 734.

Additional details on these commands are available in the *NetBackup Commands Reference Guide* or in the man pages.

Replacing the Accelerator image size with the network-transferred data in NetBackup command output

You can configure the output of `bpimagelist`, `bpdbjobs`, and `bpclimagelist` to show the amount of Accelerator backup data that was transferred over the network instead of the backup image size.

The following is the default `bpimagelist` output that shows the Accelerator image size (see the circled value 7799632). The amount of network-transferred data appears farther down in the output (225792):

```
IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 7799632 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
```

You can configure NetBackup command output to show the network-transferred data in the image size field. In the output, the image size value is replaced with the network-transferred data value (see the following example). A script that reads the image size from the command output now reads the amount of network-transferred data.

In the following `bpimagelist` output, the image size field shows the network-transferred data (225792):

```

IMAGE acmevm2 0 0 12 acmevm2 accl_vmware 40 *NULL* root f 0 9 14344
79628 558 2147483647 0 0 225792 28196 1 2 0 accl_vmware_1434479628_FULL.f *NULL
* *NULL* 0 1 0 0 0 *NULL* 0 0 1 0 0 1434479628 1434479628 *NULL* 0 0 0 *NULL* 9
0 0 3398732 0 0 *NULL* *NULL* 0 1434479620 0 0 *NULL* *NULL* 0 0 0 225792
HISTO 0 0 0 0 0 0 0 0 0 0
FRAG 1 -1 3319 76 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 102
4 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 2147483647
0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0
FRAG 1 1 7796313 0 0 0 0 @aaaab acmevm6.acme.com 262144 0 0 -1 10
28 1;PureDisk; acmevm6.acme.com;msdp_dp;PureDiskVolume;0 214748364
7 0 65545 0 0 0 6 0 1434480186 1 1 *NULL* *NULL* 0 0

```

Note: The same change occurs in the labeled output of the commands (such as with the `-L` option of `bpimagelist`). For example, the `Kilobytes` field shows the transferred data value (225792 in the example) rather than the Accelerator backup image size.

To enable the reporting of network-transferred data in the Accelerator image size field of `bpimagelist`, `bpdbjobs`, and `bpclimagelist`

- ◆ Use the `bpsetconfig` command to enable the output change.

To enable this change for the `bpclimagelist` command, enter the following `bpsetconfig` command on the master server. To enable this change for `bpimagelist` or `bpdbjobs`, enter the `bpsetconfig` command on the server where you intend to run `bpimagelist` or `bpdbjobs`.

Refer to [Table 20-31](#) for the `bpsetconfig` command to use based on the type of Accelerator backup that you want to report on.

Table 20-31 To enable the reporting of network-transferred data in the Accelerator image size field of **bpimagelist**, **bpclimagelist**, or **bpdbjobs** output

Type of backup to report on	Enter this command
Incremental VMware Accelerator backups	<p>Windows</p> <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_VMWARE install_path\NetBackup\bin\admincmd\bpsetconfig</pre> <p>UNIX, Linux</p> <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_VMWARE" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>
All VMware Accelerator backups (full and incremental)	<p>Windows</p> <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_VMWARE install_path\NetBackup\bin\admincmd\bpsetconfig</pre> <p>UNIX, Linux</p> <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_VMWARE" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>
All incremental Accelerator backups (physical clients and VMware)	<p>Windows</p> <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> <p>UNIX, Linux</p> <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_INC_ALL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>
All Accelerator backups (full and incremental, physical clients and VMware)	<p>Windows</p> <pre>echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL install_path\NetBackup\bin\admincmd\bpsetconfig</pre> <p>UNIX, Linux</p> <pre>echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED = REPLACE_IMAGE_SIZE_FOR_ACCL_ALL_ALL" /usr/opensv/netbackup/bin/admincmd/bpsetconfig</pre>

To reset the command output to the default setting

- ◆ To disable the reporting of network-transferred data in the Accelerator image size field (return to default), enter the following:

Windows

```
echo REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =  
REPLACE_IMAGE_SIZE_DISABLED |  
install_path\NetBackup\bin\admincmd\bpsetconfig
```

UNIX, Linux

```
echo "REPLACE_IMAGE_SIZE_WITH_DATA_TRANSFERRED =  
REPLACE_IMAGE_SIZE_DISABLED" |  
/usr/opensv/netbackup/bin/admincmd/bpsetconfig
```

Enable optimized backup of Windows deduplicated volumes

Select this attribute to enable optimized backups of optimized files on a Microsoft Data Deduplication file system.

See [“About backups and restores of Microsoft Data Deduplication file systems”](#) on page 738.

If a client has a Microsoft Data Deduplication file system configured, NetBackup performs an optimized backup of optimized files. (Unoptimized files are backed up as full, intact files.) If the client does not have a Microsoft Data Deduplication file system, a normal file backup occurs.

If you do not select this option, NetBackup still backs up the files but does so as *intact* files: the files are fully reconstructed and backed up as complete files. An unoptimized backup of optimized files is not efficient: it takes extra time and extra disk activity to reconstruct each file. However, restores are faster because NetBackup does not reconstruct the files (restores are unoptimized regardless of the type of backup).

Note: For easier management, Symantec recommends creating a separate policy to back up deduplicated volumes.

Only full backups are optimized. Incremental and user backups are backed up as full, intact files.

This option is supported on the Microsoft operating systems that support Data Deduplication file systems. For supported Microsoft operating systems, see the Microsoft documentation.

See [“Configuration for Microsoft Data Deduplication file system backups”](#) on page 740.

See [“Policy Attributes tab”](#) on page 680.

About backups and restores of Microsoft Data Deduplication file systems

Microsoft Data Deduplication file systems store files in one of two different states, as follows:

- Unoptimized files

Files that do not qualify for *optimization* (that is, data deduplication within the file system).
- Optimized files

Files that have been deconstructed, and only their unique data segments are stored on the file system.

NetBackup can do either optimized backups or unoptimized backups of the Microsoft Data Deduplication file systems, as follows:

Table 20-32 Microsoft Data Deduplication file systems backup types

Type	Description
Unoptimized backup	<p>An unoptimized backup is one in which the Enable optimized backup of Windows deduplicated volumes policy attribute is <i>not</i> selected.</p> <p>NetBackup re-assembles the optimized files and backs them up as full, intact files. The storage savings of the Data Duplication file system are not retained in the backup. The optimized files remain optimized on the Data Deduplication file systems. The unoptimized files are backed up as full, intact files.</p> <p>NetBackup supports any type of storage destination for unoptimized backups of Microsoft Data Deduplication file systems.</p>

Table 20-32 Microsoft Data Deduplication file systems backup types (*continued*)

Type	Description
Optimized backups	<p>An optimized backup is one in which the Enable optimized backup of Windows deduplicated volumes policy attribute is selected. Only full backups are optimized. For incremental and user backups, each file to be backed up is reconstructed and backed up in its full form.</p> <p>For optimized files, NetBackup backs up the chunk store and the metadata that maps the files to their segments in the chunk store. The chunk store is a location in the file system that contains the data segments that comprise the files. For unoptimized files, NetBackup backs them up as full, intact files.</p> <p>The following are the attributes for an optimized backup of Microsoft Data Deduplication file systems:</p> <ul style="list-style-type: none"> ■ The policy must be a MS-Windows policy type. ■ The storage destination must be BasicDisk storage. ■ The Enable Optimized Backup for supported Windows File Systems policy attribute must be selected. See “Enable optimized backup of Windows deduplicated volumes” on page 737. ■ The system drive C : cannot be backed up. Microsoft Data Deduplication file systems cannot be used for system drives. <p>Symantec recommends the following actions for backups of Microsoft Data Deduplication file systems:</p> <ul style="list-style-type: none"> ■ Run a Microsoft optimization job and a Microsoft garbage collection job before you run a full backup of deduplicated volumes. One way to do so is to configure a Deduplication Schedule in the Windows Server Manager. ■ If you want to restore entire volumes, configure a disk image backup of the volumes. A disk image backup saves time and ensures that the storage requirement can be met with the existing volume. See “Pathname rules for Windows disk image (raw) backups” on page 806. <p>Optimized backups are supported on the Microsoft operating systems that support Data Deduplication file systems. For supported Microsoft operating systems, see the Microsoft documentation.</p>

To restore optimized backups, use the **Restore from Optimized Backup** option in the **Backup, Archive, and Restore** interface.

Note the following items regarding restores from optimized backups:

- Can only be restored to Microsoft Data Deduplication file systems.
- Restores from optimized backups of Microsoft Data Deduplication files systems using the Microsoft deduplication API are slow. First, the image must be mounted through NFS. Then, the image must be accessed through the Microsoft deduplication API, which requests data in the amounts that may be inefficient. The process to restore an entire drive may take a very long time.

Therefore, NetBackup restores all files as full, intact files (that is, non-optimized). The files are then optimized during the next scheduled file system optimization job, or you can initiate optimization manually. You should ensure that you have adequate storage for the restored files in their unoptimized state.

Configuration for Microsoft Data Deduplication file system backups

For NetBackup to do an optimized backup of Microsoft Data Deduplication file systems, specific Windows software must be installed and configured, as follows:

- Microsoft Server for NFS must be installed on the NetBackup media server. Use Windows Server Manager to add the role. Enable the NFS services. (By default, the NFS services are not enabled.)
- Microsoft Client for NFS must be installed on the Windows host that is to be backed up. Use Windows Server Manager to add the role. Without Client for NFS, restores cannot be performed.
- The Deduplication role must be installed on the Windows host that is to be backed up by using the Windows Server Manager interface.

Also, ensure that the NetBackup Client Service is configured to run as *Administrator*. If not, restores from the optimized backups fail.

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 and Replication Director (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 configuring snapshots, see the following guides:

- [NetBackup Snapshot Client Administrator's Guide](#)
- [NetBackup Replication Director Solutions Guide](#)
- [NetBackup for VMware Administrator's Guide](#)
- [NetBackup for Hyper-V Administrator's Guide](#)

Perform block level incremental backups (policy attributes)

The **Perform block level incremental backups** attribute allows NetBackup to back up only the changed data blocks of VMware virtual machines and Oracle or DB2 database files.

For details, refer to the appropriate NetBackup database agent guide or to the [NetBackup for VMware Administrator's Guide](#).

Use Replication Director (policy attributes)

Enable the **Use Replication Director** attribute when configuring a backup policy for Replication Director. By enabling this policy attribute, NetBackup enables other policy attributes that Replication Director requires:

- **Perform snapshot backups**
Ensures that the policy creates snapshots of the disk array.
- **Retain snapshots for Instant Recovery or SLP management**
Ensures that the policy retains the snapshot after the backup completes.
- **Perform off-host backup**
This option is selected automatically for an **NDMP** policy, along with the following selections:
 - To use **Data Mover**
 - **NDMP** as the **Machine** selection.
- **Replication Director Options**
Click the **Options** button to see the **Replication Director Options** dialog box and the default **Configuration Parameters** as follows:

Snapshot Type	<ul style="list-style-type: none"> ■ Auto (default): The OpenStorage partner uses the best snapshot technology available to that partner to create the snapshot. ■ Differential: The OpenStorage partner creates a snapshot that is completely dependent on the source. This parameter is based on copy-on-write technology. The device creates a cache object to maintain the original blocks of the snapshot when the blocks are modified. ■ Plex: The OpenStorage Partner creates a snapshot that is completely independent of the source snapshot. This option is based on mirror-break-off technology. When a mirror device is attached to the source, the contents of the mirror device is exactly the same as the source device. When the relationship is broken between the two, the mirror device is separated from the source. The mirror device acts as a point-in-time copy. ■ Clone: The OpenStorage Partner creates an independent copy of the volume. The copy process can take some time as the entire copy must be complete. The snapshot that is created is independent of the source.
Maximum Snapshots	<p>Sets the maximum number of Instant Recovery snapshots to be retained at one time. When the maximum is reached, snapshot rotation occurs: The next snapshot causes the oldest to be deleted.</p> <p>Managed by SLP retention is automatically selected if the Fixed or the Expire after Copy retention is currently selected in the SLP.</p>

For additional information about Replication Director configuration, see the [NetBackup Replication Director Solutions Guide](#).

See “[About NetBackup Replication Director](#)” on page 1014.

Validate Policy dialog box

The **Validate Policy** dialog box appears for the backup policies that are configured for Replication Director and are not Exchange, Oracle, or VMware policies. The dialog box appears upon selecting **OK** to save and close the policy.

To ensure that the backup policy can run successfully, NetBackup validates the policy according to the validation level that you select.

Select the validation level for the policy:

- **Complete**
Performs full topology validation on underlying storage with provisioning.
Provisioning dynamically allocates NetApp storage space to each volume or LUN as data is written.
NetBackup checks the storage space on the resource pool members and performs SnapVault and SnapMirror access checks.
- **Basic**

Performs a subset of topology validation on underlying storage without provisioning.

NetBackup checks all policies to ensure that the client can perform a snapshot of the data that is indicated in the **Backup Selections** list.

NetBackup performs license checks, performs SnapVault and SnapMirror access status checks, and checks the CIFS/NFS status.

- **None**

No topology validation or provisioning.

NetBackup does not check the topology and does not provision the underlying storage.

If the policy validation finds no problems, the policy saves and closes. If validation problems are found, NetBackup displays a message that contains an error code and a description.

Perform snapshot backups (policy attributes)

The **Perform snapshot backups** attribute ensures that the policy creates snapshots of the volumes that are indicated in the policy.

Microsoft Exchange Attributes (policy attributes)

The **Microsoft Exchange** attributes let you indicate the database backup source to use for the following:

- Exchange Database Availability Group
- Exchange 2007 replication backup

See the [NetBackup for 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 **Change**.
- To delete a schedule, select the schedule and click **Delete**.

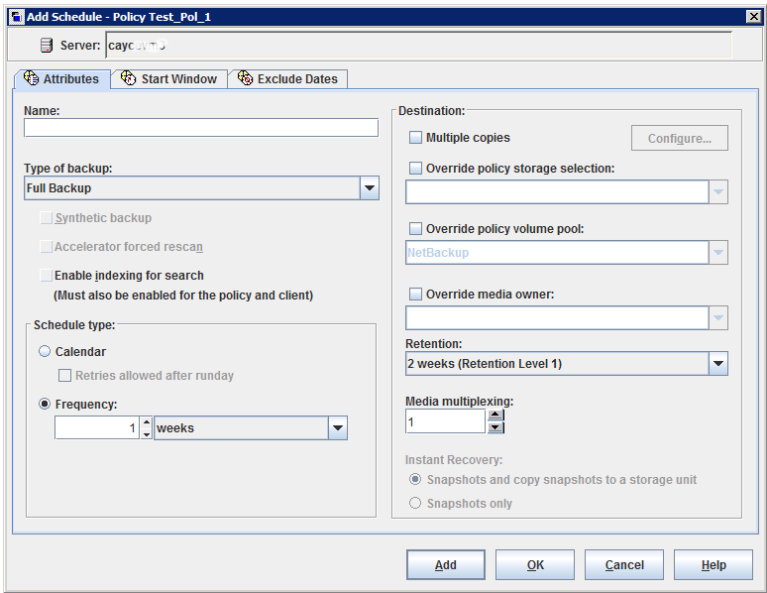
Schedule attributes appear on the following tabs:

Attributes tab	Schedule the time and frequency at which a task runs, along with other scheduled attributes. See “Schedule Attributes tab” on page 744.
Start Window tab	Schedule the time of each day that a task runs. See “Start Window tab” on page 776.
Exclude Days tab	Indicate the days that a job cannot run. See “Excluding days from a schedule” on page 780.
Include Dates 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 “Include Dates tab” on page 781.

Schedule Attributes tab

The schedule **Attributes** tab contains both schedule information and other configuration options, beyond when the job is to run.

Figure 20-10 Schedule Attributes tab



The following topics describe 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 1093.

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 576.

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 20-33](#) describes the types of backups that come standard with NetBackup. [Table 20-34](#) and [Database agent backup types](#) describe other types of backups available in NetBackup.

Table 20-33 Standard backup types

Item	Description
Full Backup	Backs up all of the files that are specified in the backup selections list for the policy. The files are backed up, regardless of when the files were last modified or backed up. Full backups occur automatically according to schedule criteria. If you run incremental backups, you must also schedule a full backup to perform a complete restore. Use this option if you configure a policy for a raw partition backup (formatted partitions only).

Table 20-33 Standard backup types *(continued)*

Item	Description
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 748.</p>
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 748.</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>For a user to be able to initiate a backup, the policy must contain a schedule of the User Backup type.</p> <p>Use this backup type for a catalog archive.</p> <p>See “Considerations for user schedules” on page 753.</p> <p>See “Creating a catalog archiving policy” on page 890.</p>

Table 20-33 Standard backup types *(continued)*

Item	Description
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. The archive then 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>For a user to be able to initiate an archive, the policy must contain a schedule of the User Archive type.</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 20-34](#) describes the types of backups that are available when you install NetBackup Vault.

Table 20-34 NetBackup Vault backup types

Item	Description
Automatic Vault	<p>Applies only to Vault policies. The option does not run a backup, but instead runs the command that is specified in the Vault policy's backup selections list. In this way it starts an automatic, scheduled vault session or vault eject operation. Available only when Vault is licensed.</p> <p>See “Creating a Vault policy” on page 840.</p>
Vault Catalog Backup	<p>Use when the schedule is for a catalog backup policy that Vault uses. Available only when Vault is licensed.</p> <p>If this type is selected, you must configure one of the two schedule attribute combinations or the schedule cannot be saved:</p> <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>

Database agent backup types

Each database agent often has its own set of unique backup types for a schedule. For more information, see the [NetBackup guide](#) that came with the agent.


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.

Note: You must run a full backup before an incremental backup. If no full backup is run, the incremental performs the role of a full backup.

Figure 20-11 shows how data is included in a series of full and differential incremental backups between January 1 and January 4.



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 20-12 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 includes 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 20-12 Full and cumulative incremental example




Table 20-35 describes how to determine the retention of differential and cumulative incremental backups to prevent a gap in backup coverage.

Table 20-35 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 20-36 compares the advantages and disadvantages of using differential or cumulative incremental backups based on possible backup and restore times.

Table 20-36 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 20-37 Example frequencies and retention periods

Backup type	Frequency	Retention period
Full	6 days	2 weeks
Cumulative incremental	2 days	4 days
Differential incremental	1 day	2 days

The schedules that are described in Table 20-37 result in the following series of backups:

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Full	Diff	Cum	Diff	Cum	Diff	Full	Diff

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:

<code>mtime</code>	The file modification time. The file system updates the <code>mtime</code> for a file or directory each time the file is modified. An application can save the <code>mtime</code> of the file before it modifies it. The application then resets it with the <code>utime(2)</code> system call.
<code>atime</code>	The file access time. The file system updates the <code>atime</code> for a file or directory each time the file is accessed (read or write). An application can save the <code>atime</code> of the file before it accesses it. The application then resets it with the <code>utime(2)</code> system call.
<code>ctime</code>	The inode change time. The <code>ctime</code> for a file or directory is updated each time the file or directory's inode changes. (For example, changes due to permissions, ownership, and link-counts changes.) The <code>ctime</code> for a file or directory cannot be saved before a change, and then reset after a change. The <code>ctime</code> of a file or a directory changes when the <code>mtime</code> and <code>atime</code> (changes with the <code>utime(2)</code> system call) is reset.

When NetBackup reads the data for a file that is included in a backup, it does not affect the file modification time. It does affect the access time of the file. For this reason, NetBackup saves the `atime` and `mtime` of the file before it reads the file. Then NetBackup resets the `atime` and `mtime` with the `utime(2)` system call. NetBackup does not cause problems for storage migration products or the administrator scripts that use file access times (`atime`) as criteria for their operations. While this benefit is obvious, a side effect is that it does update the `ctime` of the file.

Note: When NetBackup Accelerator is used to perform backups, NetBackup does not reset the `atime` after the file is backed up. This avoids updating the `ctime`, because a `ctime` update would require a loss in Accelerator optimization.

See [“Accelerator notes and requirements”](#) on page 720.

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.

See [“USE_CTIME_FOR_INCREMENTALS option for NetBackup clients”](#) on page 363.

See [“DO_NOT_RESET_FILE_ACCESS_TIME option for NetBackup clients”](#) on page 349.

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 20-38 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 PropertiesSelect the Backups tab, and specify the backup policy and backup schedule.
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 693.</p> <p>See “Global Attributes properties” on page 150.</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	<p>Use a different storage unit to eliminate conflicts with automatic backups.</p>
Volume pools	<p>Use a different volume pool to manage the media separate from the automatic backup media.</p>
Retention periods	<p>Consider setting the retention period for archives to infinite, since the disk copy of the files is deleted.</p> <p>See “Retention Periods with end dates beyond 2038, excluding Infinity” on page 222.</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)

The **Synthetic backup** schedule attribute allows a backup to be assembled from previous backups. A synthetic backup may be a synthetic full or a synthetic cumulative incremental backup. 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.

See [“About synthetic backups”](#) on page 850.

Accelerator forced rescan option (schedule attribute)

The policy **Schedules** tab contains an option called **Accelerator forced rescan**. This option creates a checksum of the content of each file during backup. It uses the checksums for change detection. It provides a safety net by establishing a new baseline for the next Accelerator backup.

The **Accelerator forced rescan** option detects the following events:

- The file’s data changes but the file’s metadata does not change.
- The file’s metadata becomes corrupted such that it does not indicate that the file has changed.
- A malicious user or application changes the file’s metadata such that it does not indicate that the file has changed.

Note: If **Accelerator forced rescan** is enabled, NetBackup uses the Windows change journal to help determine if a file has changed. NetBackup also examines the file system metadata for each file to help detect which files have changed.

For the most efficient use of Accelerator, Accelerator policies must include at least two full-backup schedules: one full schedule with the **Accelerator forced rescan** option disabled, and another full schedule with **Accelerator forced rescan** enabled. See [Table 20-39](#).

Table 20-39 Required full-backup schedules for each Accelerator policy

Full backup schedules	Notes on schedule frequency
First schedule: Accelerator forced rescan disabled	Configure this schedule to run most of your Accelerator full backups.

Table 20-39 Required full-backup schedules for each Accelerator policy
(continued)

Full backup schedules	Notes on schedule frequency
Second schedule: Accelerator forced rescan enabled	<p>Configure this schedule to run less often than the first full-backup schedule.</p> <p>For example: If the first full-backup schedule runs weekly, run the second schedule (with the Accelerator forced rescan option enabled) every few months. However, the best frequency for this schedule depends upon your environment.</p> <p>Note: If the policy has no schedule that enables the Accelerator forced rescan option, all full backups automatically enable that option and backup performance is reduced.</p>

Note the following about the **Accelerator forced rescan** option:

- The **Accelerator forced rescan** option is grayed out if the **Use Accelerator** option on the **Attributes** tab is not selected.
- Because of the checksum processing on the client, this option reduces backup speed as compared to the **Use Accelerator** option on its own. The speed reduction depends on the client's configuration and its current processing load. If the client is busy with many jobs when Accelerator backup begins, checksum processing can reduce backup speed.
- If the Windows volume is not NTFS or ReFS, and the volume has no full backup schedule that is configured with the **Accelerator forced rescan** option, note: NetBackup uses **Accelerator forced rescan** on every full backup. The **Accelerator forced rescan** option is not enabled in the schedule but is in fact used, as indicated in the NetBackup log messages. This option may reduce the speed of the backup as compared to the previous backups that did not use **Accelerator forced rescan**.
See [Table 20-39](#) for recommended full-backup schedules.
- For an MS-Windows or Standard policy (to back up files and folders), all the data on the client is read. However, NetBackup sends only the changed data over the network to be included in the backup image. Sending only the changed data is similar to a regular Accelerator full backup. Thus, for an Accelerator forced rescan backup, the optimization percentage is similar to a regular Accelerator full backup. Note that the duration of the backup falls somewhere between a non-Accelerator full backup and a regular Accelerator full backup.

Calendar (schedule attribute)

Calendar-based schedules allow administrators to create a job schedule based on a calendar view. Select **Calendar** to display the **Include Dates** tab.

See [“Include Dates tab”](#) on page 781.

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 run day 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 2 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 20-40 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 and Replication Director \(policy attributes\)](#)” on page 740.

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.

The **Instant Recovery** attributes are grayed out if the **Policy storage** option on the Policy **Attributes** tab refers to a storage lifecycle policy. If that is the case, the storage lifecycle policy configuration governs the **Instant Recovery** attributes.

However, the **Override policy storage selection** attribute on the Schedule **Attributes** tab overrides the **Policy storage** option. If a storage unit is selected on the Schedule **Attributes** tab, the **Instant Recovery** attributes become enabled.

See [“Policy storage \(policy attribute\)”](#) on page 685.

See [“Override policy storage \(schedule attribute\)”](#) on page 764.

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.)

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 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 565.

See [“Maximum concurrent write drives storage unit setting”](#) on page 564.

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 654.

Configure Multiple Copies dialog box

The **Configure Multiple Copies** dialog box contains the following options:

Table 20-41 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>To create more than 4 copies, create additional copies at a later time by using duplication.</p>
Priority of duplication job	<p>Indicate the priority that the duplication job (based on this schedule) has over other jobs in the queue (0 to 99999).</p>
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 960.</p>
Storage unit	<p>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 run-time, select Any Available.</p>
Volume pool	<p>Indicate where each copy is to be stored.</p>
Retention schedule	<p>Specify how long NetBackup retains the backups.</p> <p>See “Retention (schedule attribute)” on page 766.</p>
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 689.</p>

Table 20-41 Configure Multiple Copies dialog box (continued)

Field	Description
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. See “Configuring a NetBackup server group” on page 367.

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 > 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 685.

See [“About writing multiple copies using a storage lifecycle policy”](#) on page 654.

- 6 In the **Copies** field, specify the number of copies to be created simultaneously. The number must be between 1 and 4.

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 761.

See [“About configuring multiple copies”](#) on page 760.

- 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 run-time.

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 766.

- 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 689.

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 685.

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 684.

Note: Storage lifecycle policies cannot be selected within the **Configure Multiple Copies** dialog box.

See [“About configuring multiple copies”](#) on page 760.

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 688.

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	<p>Instructs the schedule to override the media owner that is specified as the Media Owner in the policy Attribute tab.</p> <p>Select the new media owner from the list:</p> <ul style="list-style-type: none"> ■ 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 695.

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 2 weeks, data can be restored from a backup that this schedule performs for only 2 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 [“Creating a storage lifecycle policy”](#) on page 601.

See [“Retention Periods properties”](#) on page 219.

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 20-42 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 2-4 weeks. Two to 4 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 2 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 20-43 Suggestions for preventing prematurely expired backups

Item	Description
Retention period	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. For the backups that must be kept for more than one year, set the retention period to infinite.
Full backups and incremental backups	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.
Archive schedules	Set the retention period to infinite.

Table 20-43
 Suggestions for preventing prematurely expired backups *(continued)*

Item	Description
Tape	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.

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 180.

Note: Retention levels can be mixed on disk volumes with no restrictions.

See [“Changing a retention period”](#) on page 220.

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.


Caution: For MS-SQL-Server policies, do not enable multiplexing for a schedule that is also configured to backup with multiple stripes. Restores fail when multiplexing is enabled for a schedule that uses more than one stripe.

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.

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 20-13](#) shows the multiplexed flow of client data to a server.

Figure 20-13 Multiplexed backups



About multiplexing

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 20-44 Circumstances where multiplexing improves performance

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. Also, on UNIX, enable fast-tape positioning (locate block), if it applies to the tape drives in use.

Consider the following configuration settings when using multiplexing.

Table 20-45 Properties and attributes that affect multiplexing

Item	Description	Where to find it
Limit jobs per policy (policy attribute)	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. See "Limit jobs per policy (policy attribute)" on page 693.	<ul style="list-style-type: none">■ In the NetBackup Administration Console, expand NetBackup Management > Policies.■ In the middle pane, double-click the Attributes node of a policy. Or, create a new policy and select the Attributes tab.

Table 20-45 Properties and attributes that affect multiplexing (*continued*)

Item	Description	Where to find it
Maximum jobs per client (host property)	<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 150.</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 567.</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 20-45 Properties and attributes that affect multiplexing (*continued*)

Item	Description	Where to find it
Maximum data streams (host property)	<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 86.</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.
Delay on multiplexed restores (host property)	<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 86.</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.
Media Multiplexing (policy schedule attribute)	<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 769.</p>	<ul style="list-style-type: none"> ■ In the NetBackup Administration Console, expand NetBackup Management > Policies. ■ In the middle pane, double-click the Schedules node of a policy. Or, create a new policy and select the Schedules tab. ■ Click New to create a new schedule and configure the Media Multiplexing option.

Table 20-45 Properties and attributes that affect multiplexing (continued)

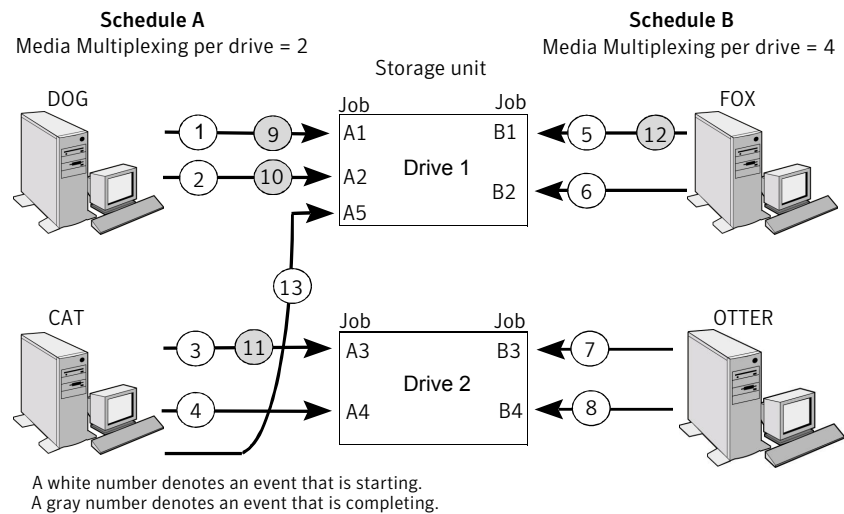
Item	Description	Where to find it
Maximum streams per drive (storage unit setting)	<p>NetBackup can add jobs from more than one schedule to a drive.</p> <p>When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches the Maximum streams per drive limit or the Media Multiplexing limit.</p> <p>See “Maximum streams per drive storage unit setting” on page 567.</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. <p>Or, create a new storage unit.</p> <ul style="list-style-type: none">■ The setting appears on the dialog box that appears.

See “MPX_RESTORE_DELAY option for NetBackup servers” on page 306.

Example of using multiplexing with schedules

Figure 20-14 provides an example of how schedules are affected when multiplexing is active.

Figure 20-14 Multiplexing process scenario



Assume the following about Figure 20-14.

- 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 711.

Table 20-46 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.
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 20-14](#) 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.

See “[MPX_RESTORE_DELAY option for NetBackup servers](#)” on page 306.

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 962.

Start Window tab

The **Start Window** tab provides controls for setting time periods during which NetBackup can start jobs when using a schedule. Time periods are referred to as windows. Configure windows so that they satisfy the requirements necessary to complete a job.

For example, create different 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 policy schedule

Use one of the following procedures to add, change, or delete a time window.

To add or change a time window in the NetBackup Administration Console

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 Expand the policy name in the middle pane, then select **Schedules**.
- 3 Do one of the following:

To add a time window

- Click **Actions > New > Schedule**.
- In the **Add Schedule** dialog box, type the name of a schedule.

To change a time window

In the right pane, double-click the schedule you want to change. The **Change Schedule** dialog box appears.

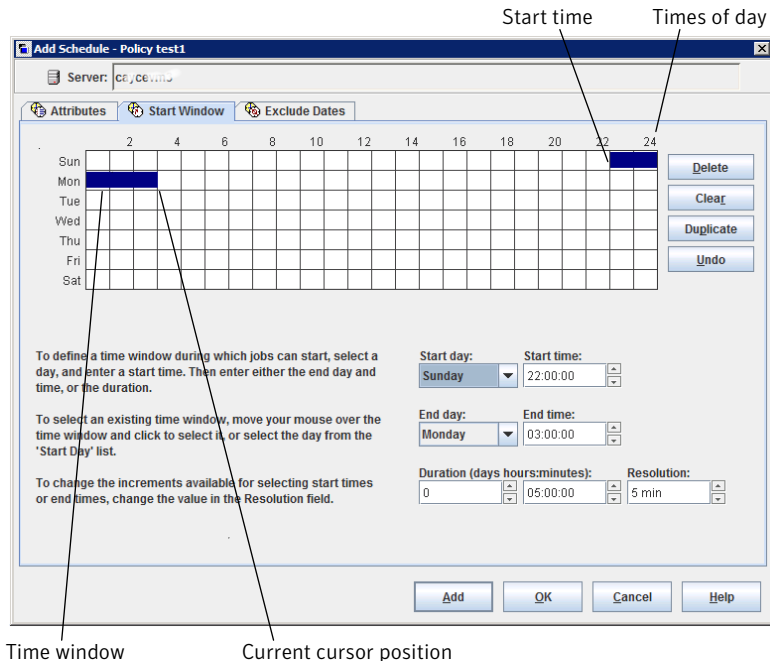
- 4 Click the **Start Window** tab.
- 5 To change the increments available for selecting start times or end times, change the value in the **Resolution** field. You can choose 5, 10, 15, or 30 minutes. For example, a resolution of 10 minutes allows time window adjustments by 10-minute increments.
- 6 To indicate the opening of the time window, do the following:

Drag your cursor in the time table

Click the day and time when you'd like the window to start and drag it to the day and time when you'd like the window to close.

Use the settings in the dialog box

- In the **Start day** field, select the first day that the window opens.
- In the **Start time** field, select the time that the window opens.



7 To indicate the closing of the time window, do one of the following:

Drag your cursor in the time table

Click the day and time when you'd like the window to start and drag it to the day and time when you'd like the window to close.

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

- 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.

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.

8 As necessary, do any of the following:

- Click **Delete**. Deletes the selected time window.
- Click **Clear**. Deletes all time windows from the schedule display.
- Click **Duplicate**. Replicates the time window for the entire week.
- Click **Undo**. Erases the last action.

9 Do one of the following:

- Click **Add**. To save the time window and leave the dialog box open.
- Click **OK**. To save the time window and close the dialog box.

Example of schedule duration

Figure 20-15 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 20-15 Duration example

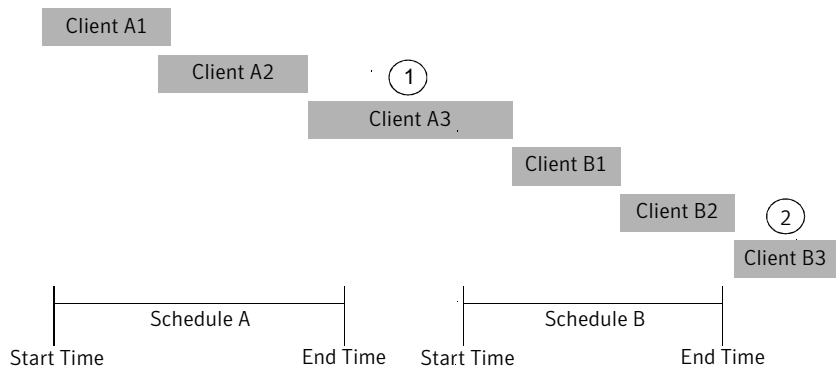


Figure 20-15 illustrates the following points:

- 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 days from a schedule

Use the **Exclude Days** tab to exclude specific days from a schedule for a backup policy. If a day is excluded from a schedule, jobs do 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 day from a schedule

- 1 Use one or more methods to indicate the days to exclude:
 - Select the day(s) on the 3-month calendar that you want to exclude. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring Week Days**:
 - Click **Select All** to select all of the days in every month for every year.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select a specific day to exclude for every month.
 - Click the column head of a day of the week to exclude that day every month.
 - Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to exclude that week every month.
 - To indicate **Recurring Days of the Month**:
 - Click **Select All** to select all of the days in every month.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select that day to exclude each month.
 - Click **Last Day** to exclude the last day of every month.
 - To indicate **Specific Dates**:
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.
The date appears in the **Specific Dates** list.

- To delete a date, select the date in the list. Click **Delete**.
- 2 Add additional dates as necessary, then click **OK** to save the changes.

Include Dates tab

The **Include Dates** 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.

Calendar scheduling with the Include Dates tab

Use the **Calendar** option in the policy **Attributes** tab to create a job schedule based on a calendar view. The **Include Dates** tab lets administrators configure the schedules that run according to specific days, on recurring week days, or on recurring days of the month.

Note: Using the calendar schedule, if a green checkmark does not appear on a day, the day is not included in the schedule.

If **Retries allowed after runday** is enabled, a job could run on a day that is not included in the schedule.

When a new calendar schedule is created with **Retries allowed after runday** enabled, the schedule runs its first job on the next day when the backup window is open. That day may be before the first run day that is included in the schedule.

To use a calendar to schedule run days

- 1 In the **Attributes** tab, enable the **Calendar** attribute.
- 2 Select the **Include Dates** tab.
- 3 Use one or more methods to schedule the days on which jobs can run:
 - Select the day(s) on the three-month calendar that you want jobs to run. Use the drop-down lists at the top of the calendar to change the months or year.
 - To indicate **Recurring Week Days**:
 - Click **Select All** to select all of the days in every month for every year.
 - Click **Deselect All** to remove all existing selections.

- Check a box in the matrix to select a specific day to include for every month.
 - Click the column head of a day of the week to include that day every month.
 - Click the **1st**, **2nd**, **3rd**, **4th**, or **Last** row label to include that week every month.
 - To indicate **Recurring Days of the Month**:
 - Click **Select All** to select all of the days in every month.
 - Click **Deselect All** to remove all existing selections.
 - Check a box in the matrix to select that day to include each month.
 - Click **Last Day** to include the last day of every month.
 - To indicate **Specific Dates**:
 - Click **New**. Enter the month, day, and year in the **Date Selection** dialog box. Click **OK**.
The date appears in the **Specific Dates** list.
 - To delete a date, select the date in the list. Click **Delete**.
- 4 Add additional dates as necessary, then click **OK** to save the included days.

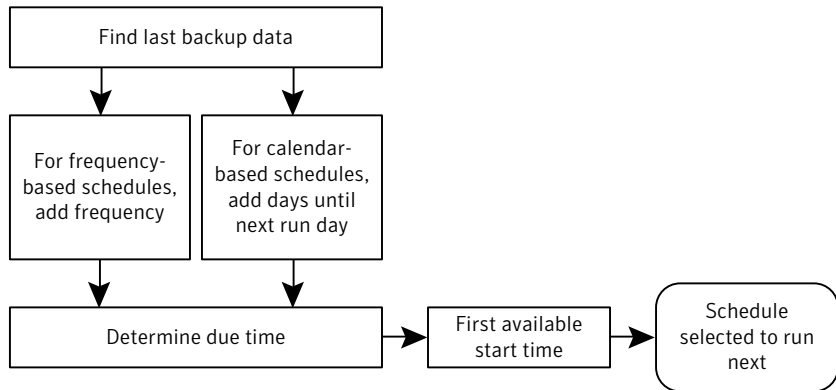
How NetBackup determines which schedule to run next

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 20-16 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 150.

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:

 - 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:

 - Full schedules of the same or longer frequency.
 - Cumulative incremental schedules of the same or longer frequency.
 - Other differential incremental schedules of the same or longer frequency.

Note: To have a longer frequency means that the schedule is configured to run less often.

The comparison rules ensure that no schedule is overlooked for consideration, potentially causing a gap in backup coverage.

Figure 20-17 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.

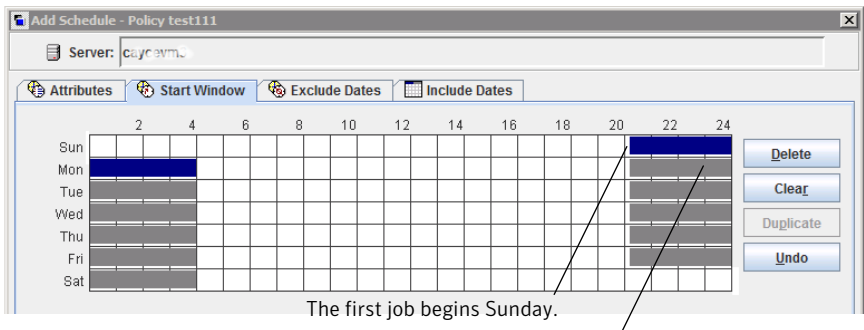
About schedule windows that span midnight

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 the next 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 20-18 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 20-18 Schedule that spans midnight



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.

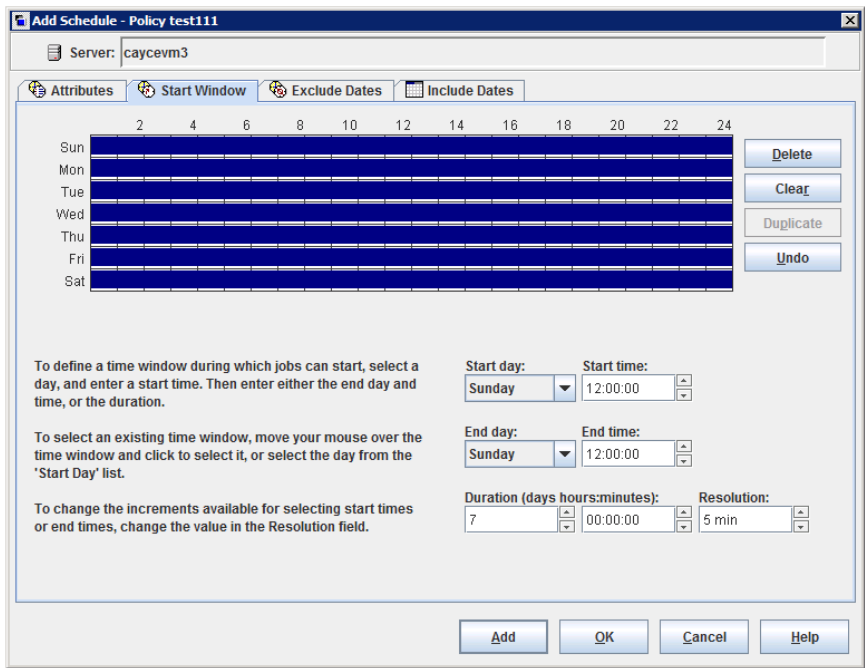
If no other window is scheduled to open later in the day, NetBackup does not wait. If the job has a daily frequency, the job runs again after midnight to meet the daily backup frequency requirement.

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 20-19 shows an open schedule.

Figure 20-19 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 20-20 shows that the backups on a calendar-based schedule would run Monday through Friday.

Figure 20-20 An open schedule that is calendar-based




Figure 20-21 and Figure 20-22 show that the backups based on a frequency-based schedule should run every day of the week, including Saturday and Sunday.

Figure 20-21 An open schedule that is frequency-based

The screenshot shows the 'Add Schedule - Policy test111' dialog box. The 'Server' field is set to 'c:\program'. The 'Attributes' tab is selected. The 'Name' field contains 'freq_daily_sched'. The 'Type of backup' is set to 'Full Backup'. The 'Schedule type' is set to 'Frequency' with a value of '1' and unit 'days'. The 'Destination' section includes options for 'Multiple copies', 'Override policy storage selection', 'Override policy volume pool' (set to 'NetBackup'), and 'Override media owner'. The 'Retention' is set to '2 weeks (Retention Level 1)'. The 'Media multiplexing' is set to '1'. The 'Instant Recovery' section has 'Snapshots and copy snapshots to a storage unit' selected. The 'Add', 'OK', 'Cancel', and 'Help' buttons are at the bottom.

Server: c:\program

Attributes Start Window Exclude Dates

Name: freq_daily_sched

Type of backup: Full Backup

☐ Synthetic backup

☐ Accelerator forced rescan

☐ Enable indexing for search
(Must also be enabled for the policy and client)

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency: 1 days

Destination:

☐ Multiple copies [Configure...](#)

☐ Override policy storage selection:

☐ Override policy volume pool: NetBackup

☐ Override media owner:

Retention: 2 weeks (Retention Level 1)

Media multiplexing: 1

Instant Recovery:

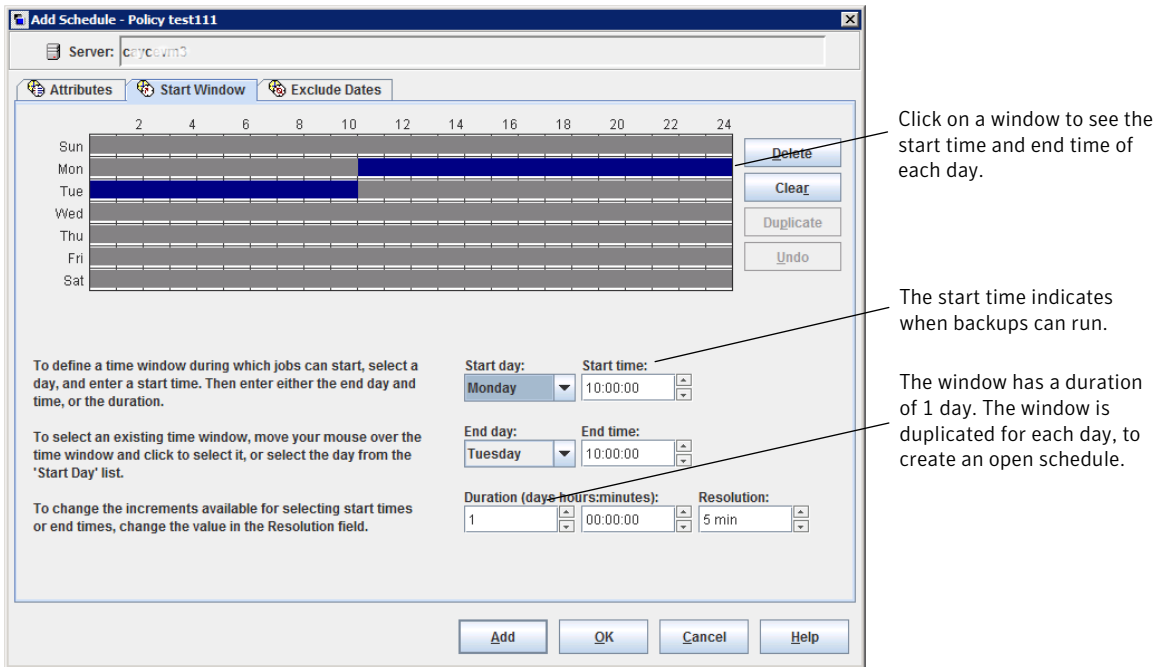
☒ Snapshots and copy snapshots to a storage unit

☐ Snapshots only

Add OK Cancel Help

In [Figure 20-22](#), backups run at 10:00 P.M. nightly based on the start time.

Figure 20-22 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 middle pane, double-click on the policy name where you want to create an open schedule.
- 3 Select the **Schedules** tab.
- 4 Click **New** or **Add** 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 20-47 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 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 150.</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.

On UNIX: You can also use the **Clients** tab to install NetBackup software on trusting UNIX clients.

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 procedures to add, change, or delete clients in an existing NetBackup policy.

To add, change, or delete a client in a policy

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
- 2 Open the policy that you want to change.
- 3 Select the **Clients** tab and perform one of the following actions:

To add a new client Continue to step 4.

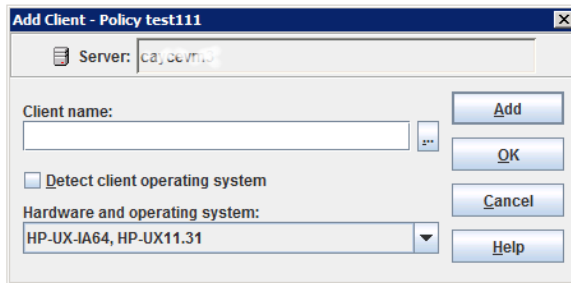
To change an existing client Double-click on the client that you want to change or select the client and click **Change**. The **Change Client** dialog box appears.

Click **OK** to accept the change and close the **Change Client** dialog box. Or, click **Cancel** to close the dialog box without saving the changes.

To delete a client Select a client and click **Delete**. Hold down **Shift** to select multiple clients. A confirmation dialog box appears that lists the clients to be deleted. Click **Yes** to delete the clients or **No** to escape the dialog box.

- 4 To add a new client, click **New**. The **Add Client** dialog box appears.

In the **Add Client** dialog box, enter the host name that you want to add. Or click the **Browse for Computer** button, select the host from the network tree, and click **OK**.



Observe the following rules for assigning client names:

- Use a name by which the server knows the client (one that you can use on the server to `ping` or `telnet` to the client).
- If the client is in multiple policies, use the same name in each policy.
- If the network configuration has multiple domains, use a more qualified name. For example, use `client1.null.com` or `client1.null` rather than only `client1`.
- Add only clients with the hardware and the operating systems that this policy supports. For example, do not add a UNIX 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.
If the hardware and the operating system you want is not in the list, associated client software is not installed on the server. Check the `/usr/opensv/netbackup/client` directory for the directories and software that corresponds to the client you want to install. If the directories or software are not there, rerun the installation script on the server and select the option to install client software.
- Do not use an IP address as a client name in a policy or the backup may fail. Specify a host name instead.

- 5 You can either select the **Detect client operating system** check box or select the appropriate hardware and operating system for the client in the drop-down menu.

Note: This feature is also available when you configure the LiveUpdate policy type and also when you use the Policy Configuration Wizard to configure a policy. To be able to use the auto-detect feature, the master server name must be included in the bp.conf file on the client. This task is accomplished when you enter the master server name during the installation of the client software on the host.

NetBackup is migrating to a simplified client name for Windows clients. The new client selection **Windows** can be used for all supported Windows clients, regardless of the specific Windows version of the client. For Windows 8, Windows 2012 and later clients, choose the **Windows** option. In future releases, specific Windows version choices will be eliminated.

- 6 Perform one of the following:
 - Click **Add** to add the client to the list and leave the dialog box open to add another client.
 - Click **OK** to add the client to the list and close the dialog box.
 - Click **Cancel** to close the dialog box without adding the client.
- 7 When you are finished in the **Clients** tab:
 - Click **OK** to close and save the policy.
 - Click **Cancel** to close the policy without saving any additions or changes.

Installing client software on secure and trusting UNIX clients

A secure UNIX client is a client that does not contain an entry for the NetBackup master server in its `/.rhosts` file. You can install software on a secure UNIX client locally by using the installation DVDs or by using a script.

By contrast, a trusting UNIX client contains an entry for the current master server in its `/.rhosts` file. You can install client software on trusting UNIX clients from the policy **Clients** tab or from the **Actions** menu when the **Policies** utility is selected.

To install NetBackup Windows client software, see the [NetBackup Installation Guide for Windows](#).

Before you can install client software on trusting UNIX clients, the following prerequisites must be met:

- You can install the client software from the **NetBackup Administration Console** only from a UNIX NetBackup server. The server must be the server that was specified in the logon dialog box when the interface was started. This server must also be the master where you currently manage backup policies, and clients must be in a policy on this master.
- Each client that receives the software installation must contain an entry for the current master server in its `/.rhosts` file. If these entries exist, the clients are referred to as trusting clients. The `/.rhosts` entries for the master server are not required for correct operation of NetBackup. The entries can be removed after the client software is installed.

To install client software on trusting UNIX clients

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

Note: If you want to install client software, you cannot click **File > Change Server** to get to another master server. The master server must be the server that you specified in the logon dialog box.

- 2 Select the master server name at the top of the **All Policies** middle pane.
- 3 On the **Actions** menu, click **Install UNIX Client Software**.
- 4 In the **Don't install these clients** box, select the clients in the list that you want to receive the software. Click **Add** to move the client names to the **Install these clients** box.
- 5 Click **Install Client Software** to start the installation.

Client software installation can take a minute or more per client. NetBackup writes the messages in the **Progress** box as the installation proceeds. If the installation fails on a client, NetBackup notifies you but keeps the client in the policy. You cannot stop the installation once it starts.

During installation, NetBackup performs the following actions:

- Copies the client software.
 - Copies from the `/usr/opensv/netbackup/client` directory on the server.
 - Copies to the `/usr/opensv/netbackup` directory on the client.
- Adds the required entries to the client's `/etc/services` and `inetd.conf` files.

- 6 When the install is complete, click **Close**.

To install client software to a different location on the client:

- Create a directory where you want the software to reside.
- Create `/usr/opensv/netbackup` as a link to that directory before installing software.

Browse for Hyper-V virtual machines

On the **Clients** tab, click **New** to select virtual machines.

The following table describes the options that you can use to select Hyper-V virtual machines.

Table 20-48 Options for selecting Hyper-V virtual machines

Option	Description
Enter the VM display name (or VM hostname or VM GUID)	<p>Note: The type of name to enter depends on the Primary VM identifier setting on the Hyper-V tab of the policy.</p> <p>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, the policy validation fails.</p> <p>Make sure the Browse for Virtual Machines option is unchecked.</p>
Browse for Virtual Machines	<p>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).</p> <p>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.</p> <p>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.</p> <p>If NetBackup cannot obtain the IP address of the virtual machine, the IP address is displayed as NONE.</p>
Last Update	<p>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.</p>

Backup Selections tab

The **Backup Selections** tab contains a list of what to back up on each client, host, or instance when NetBackup runs an automatic schedule (for example, a full backup). The list does not apply to user backups or archives, where users select the objects to back up before they start the operation.

The backup selections list can contain the following:

- Paths that identify the location of files or directories
- Directives, which signal NetBackup to perform specific, predefined actions when it processes the selections list
- Scripts or templates
Administrators create scripts to define and control database backups. Templates are used exclusively for Oracle (legacy) and DB2 database backups.
- Database objects

- Wildcards

Certain wildcards can be used in the selections list. See [“Wildcard use in NetBackup”](#) on page 1094.

Windows clients support only the asterisk (*) and the question mark (?) as valid wildcards in the **Backup Selections** tab.

For information on how to use wildcards when you configure backup selections for database agents and other options, refer to the guide for that agent or option.

The list applies to each client (or host, instance, or database) in the policy. Every file on the list does not need to exist on all of the clients. NetBackup backs up the files that it finds that are on the backup selections list. However, each client must contain at least one of the files in the backup selections list. Otherwise, the backup fails with status code 71 (none of the files in the file list exist). Use the Troubleshooter to find the description of a status code. See [“Running the Troubleshooter”](#) on page 50.

The policy type determines what type of backup selections appear in the list. See [Table 20-49](#). See [“Policy type \(policy attribute\)”](#) on page 681.

Table 20-49 Items allowed in the Backup Selections list for specific policy types

Policy type	Items allowed
Standard	Paths and directives
MS-Windows	Paths and directives
Lotus-Notes, MS-Exchange-Server, MS-SharePoint	Paths and directives

Table 20-49 Items allowed in the Backup Selections list for specific policy types
(continued)

Policy type	Items allowed
MS-SQL-Server	
Informix-On-BAR, SAP, Sybase	Scripts
DB2	Scripts and templates
Oracle	For Oracle Intelligent Policies, you select an Oracle database object or other option like a backup of the Fast Recovery Area (FRA). For legacy Oracle policies, you add scripts or templates.
Vault	Vault commands

Adding backup selections to a policy

Use the following procedures 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**, expand **NetBackup Management > Policies**.
- 2 Open the policy where you want to add a backup selection.
- 3 Select the **Backup Selections** tab and perform one of the following actions:

Entering a path to a directory Click **Browse** to browse to a specific client to specify the path to the directory that you want to backup up. Or enter the pathname directly in the **Pathname or directive** field.

The **Browse** button toggles to an **Add** button after a selection is made. Click **Add** to add the path to the selection list.

A path may contain up to 1023 characters.

See [“Pathname rules for Windows client backups”](#) on page 803.

See [“Pathname rules for UNIX client backups”](#) on page 810.

Selecting a directive set or directive

Select or enter a directive or a directive set in the **Pathname or Directive set** drop-down menu.

Click **Add** to add the directive to the selection list.

See [“About the directives on the Backup Selections list”](#) on page 818.

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 or DB2 policies, select a template from the **Script or template** list, or type the name of a template.

Example of an Oracle script on UNIX:

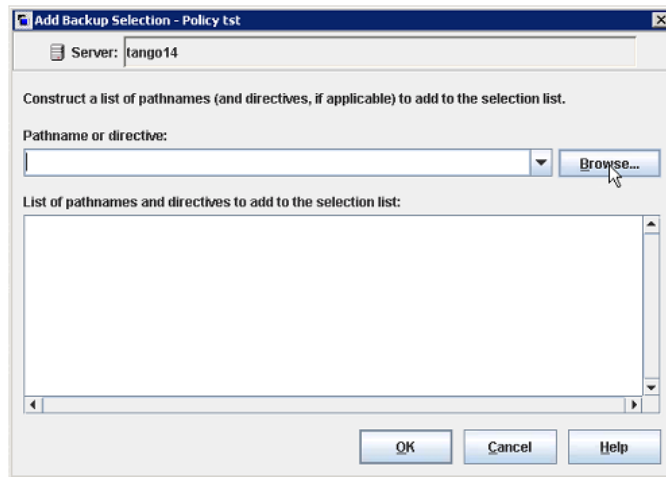
`install_path/retbackup/et/db2/oracle/scripts/man/cold_data_backup.sh`

Example of a DB2 script on UNIX:

`/myscripts/db2_backup.sh`

Click **Add** to add the script or template to the selection list.

You can select multiple directories and files for backup. The policy type determines whether the backup selections list can contain paths, directives, scripts, templates, or a combination.



- 4 Click **OK** to add the selection list to the **Backup Selections** tab in the policy.
- 5 When you are finished in the **Backup Selections** tab:
 - Click **OK** to close and save the policy.
 - Click **Cancel** to close the policy without saving any additions or changes.

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 20-50 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 803.</p> <p>See “Pathname rules for Windows disk image (raw) backups” on page 806.</p> <p>See “Pathname rules for Windows registry backups” on page 807.</p> <p>See “Pathname rules for UNIX client backups” on page 810.</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 “About the Reports utility” on page 1056.</p>
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> (on Windows) or in <code>/usr/opensv/netbackup/bin/goodies</code> (on UNIX). 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>On Windows: 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 (check_coverage)” on page 801.</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. For information on `check_coverage`, see the comments in the script.

On Windows: The following example shows the log message that appears when files expected to be on a client are not found.

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:

```
9/1/14 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.

On UNIX: The following table shows examples of the log messages that appear when files expected to be on a client are not found.

Table 20-51 Example UNIX log messages from the File System Backup Coverage Report

Example	Description
Regular expressions or wildcards	<p>Assume that the backup selections list contains a regular expression:</p> <pre>/home1[0123456789]</pre> <p>NetBackup backs up <code>/home10</code> through <code>/home19</code> if both exist.</p> <p>If they are not present, the Problems report or the All Log Entries report displays a message similar to the following:</p> <pre>02/02/14 20:02:33 windows freddie from client freddie: TRV - Found no matching file system for /home1[0123456789]</pre>

Table 20-51 Example UNIX log messages from the File System Backup Coverage Report (*continued*)

Example	Description
Path not present on all clients or wrong path specified	<p>Assume that the backup selections list contains a path named <code>/worklist</code> that is not present on all clients. NetBackup backs up <code>/worklist</code> on the clients where it exists.</p> <p>For other clients, the Problems report or the All Log Entries report displays a message similar to the following:</p> <pre>02/02/14 21:46:56 carrot freddie from client freddie: TRV - cannot process path /worklist: No such file or directory. Skipping</pre> <p>This message occurs if <code>/worklist</code> is not the correct path name. For example, the directory name is <code>/worklists</code>, but <code>/worklist</code> was typed.</p> <p>Note: If the paths seem correct and the message continues to appear, ensure that no trailing spaces appear in the paths.</p>
Symbolic link	<p>Assume the backup selections list names a symbolic link. NetBackup does not follow symbolic links and provides a message in the Problems report or the All Log Entries report:</p> <pre>02/02/14 21:46:47 carrot freddie from client freddie: WRN - /src is only being backed up as a symbolic link</pre> <p>Resolve the symbolic link if you do not intend to back up the symbolic link itself.</p>

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 150.
See ["Limit jobs per policy \(policy attribute\)"](#) on page 693.

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 711.

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 20-52 Pathname rules for Windows client backups

Item	Description
Paths per line	Enter one path per line.

Table 20-52 Pathname rules for Windows client backups (*continued*)

Item	Description
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 and path are case-insensitive.</p> <p>The following example entries would successfully indicate the same directory:</p> <pre>c:\Worklists\Admin\ C:\worklists\admin\ c:\WORKLISTS\Admin\ C:\Worklists\ADMIN\</pre> <p>Note: If a path is listed in the Backup Selections tab more than once, the data is backed up more than once.</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>For Windows clients, wildcards function correctly only when they are placed at the end of the path, in the file or directory name. For example:</p> <pre>C:\abc\xyz\r*.doc</pre> <p>Wildcard characters do not work elsewhere in the path. For example, an asterisk functions as a literal character (not as a wildcard) in the following examples:</p> <pre>C:*\xyz\myfile C:\abc*\myfile</pre> <p>See “Wildcard use in NetBackup” on page 1094.</p>

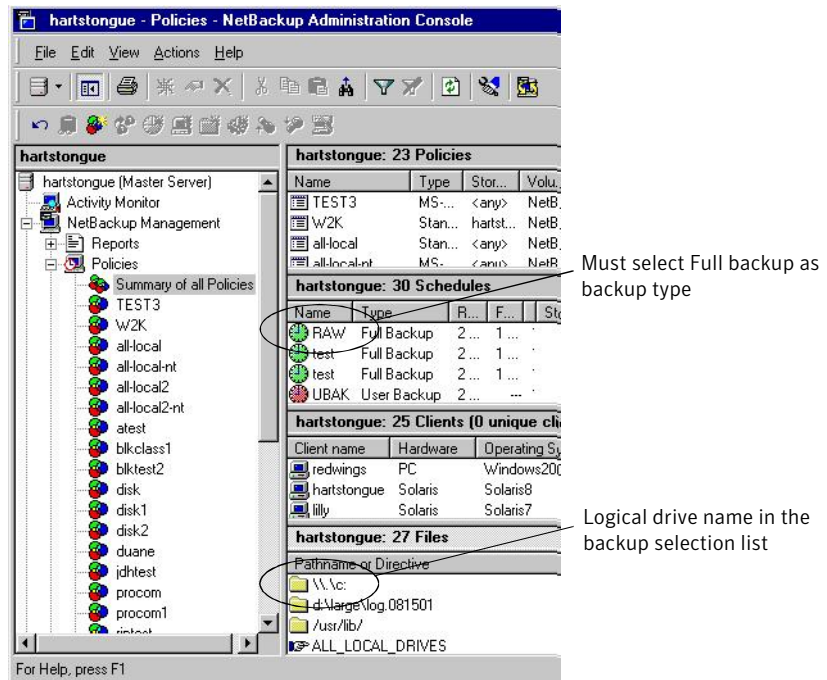
Table 20-52 Pathname rules for Windows client backups (*continued*)

Item	Description
All local drives	<p>To back up all local drives except for those that use removable media, specify the following:</p> <pre>: \</pre> <p>Or</p> <pre>* : \</pre> <p>Or</p> <pre>ALL_LOCAL_DRIVES</pre> <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>
Use of mapped drives	<p>Do not specify a local drive path that is mapped to a CIFS share using the Windows Map Network Drive option.</p> <p>This holds true for a policy that contains multiple clients as well. Do not specify paths that point to different CIFS shares.</p> <p>For example:</p> <pre>P : Q : R :</pre>
Use of UNC path(s)	<p>If a backup policy contains multiple clients that specify a UNC path as a backup selection, the redundant backup copies are created of the same data from different clients.</p> <p>Consider including the host in a policy as a client to be backed up.</p> <p>For example:</p> <pre>\\host_name\dir1</pre>
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 831.</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 832.</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 20-23 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 NetBackup Troubleshooting Guide.</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 20-24 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 20-25 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 20-26 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 20-53 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 20-53 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>For UNIX clients, wildcards can appear anywhere in the path.</p> <p>See “Wildcard use in NetBackup” on page 1094.</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.</pre> <pre>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 696.</p> <p>See “Cross mount points (policy attribute)” on page 699.</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 831.</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 832.</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 77.</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 20-53 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 815.</p>
VxFS extent attributes	On Hewlett-Packard and Solaris SPARC platforms, NetBackup backs up VxFS extent attributes.
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 808.</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 1056.

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 699.</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 696.</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 NetBackup Snapshot Client Administrator's Guide for the differences between Standard and FlashBackup policies.

Windows Server
 Failover Clustering
 (WSFC) environment

The use of FlashBackup in a Windows Server Failover Clustering (WSFC) 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 WSFC resource group.

Note: Early versions of WSFC 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/rzd*`) 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 745.

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.

On UNIX:

Figure 20-27 Example of base file and extended attribute directory and files



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.


On UNIX: In the following example, File 1 is to be restored. Base File 1 currently possesses four extended attribute files.

Figure 20-28 Extended attribute files of Base File 1




On UNIX: The user restores File 1 from a backup that was created when File 1 possessed only three extended attribute files.

Figure 20-29 Backup of File 1



On UNIX: Since **Overwrite existing files** is selected as a restore option, when the user restores File 1, extended attribute files 1, 2, and 3 are overwritten. Extended attribute file 4 remains and is not overwritten.

Figure 20-30 Restore of File 1



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`.

On UNIX: Only the modified GNU `tar` that is supplied with NetBackup can restore the extended attributes or named data streams to a client.

On UNIX: For more information, see the [NetBackup Administrator's Guide, Volume II](#).

Note: Extended attributes and named data streams cannot be compressed.

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 20-54 Summary of directives on the Backup Selections list

Directive	Description	Applicable operating system
ALL_LOCAL_DRIVES	Instructs NetBackup to back up all local drives except for those drives that use removable media. This directive also includes critical system-related components. See “ ALL_LOCAL_DRIVES directive ” on page 822.	All supported systems
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 823.	All supported Windows systems
Shadow Copy Components:\	Instructs NetBackup to back up all writers for the Volume Shadow Copy component. This also implies and/or includes System State:\ if that was not also selected. See “ Shadow Copy Components:\ directive ” on page 824.	All supported Windows systems
Active Directory Application Mode:\	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.	All supported Windows systems
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 825.	Policy type specific
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 830.	All

Table 20-54 Summary of directives on the Backup Selections list (*continued*)

Directive	Description	Applicable operating system
NEW_STREAM	<p>When NEW_STREAM is on the first line of the Backup Selections list, this directive determines how a backup is performed in the following modes:</p> <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 826.</p>	All

ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives

The ALL_FILESYSTEMS directive provides a method to include all file systems and volumes on an NDMP filer in an NDMP backup policy.

You can exclude specific volumes from an ALL_FILESYSTEMS backup selection if you do not want to back up every volume on an NDMP filer. Use the VOLUME_EXCLUDE_LIST directive for this purpose. You may use valid wildcard characters in the VOLUME_EXCLUDE_LIST statement.

Note: The following examples use selections that are specific to NetApp Data ONTAP 7-mode. For specific examples of backup selections for other configurations, refer to the appropriate documentation.

The VOLUME_EXCLUDE_LIST statements must precede ALL_FILESYSTEMS statement. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01
ALL_FILESYSTEMS
```

or

```
VOLUME_EXCLUDE_LIST=/vol/testvol*
ALL_FILESYSTEMS
```

To specify multiple values in a VOLUME_EXCLUDE_LIST statement, separate the values with a comma. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01,/vol/testvol*
ALL_FILESYSTEMS
```


You can also specify more than one `VOLUME_EXCLUDE_LIST` statement with an `ALL_FILESYSTEMS` directive. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Hr_allfiles_vol01
VOLUME_EXCLUDE_LIST=/vol/testvol*
ALL_FILESYSTEMS
```

A `VOLUME_EXCLUDE_LIST` statement may include a maximum of 256 characters. Create multiple `VOLUME_EXCLUDE_LIST` statements if necessary to avoid exceeding the limit of 256 characters. If you specify more than 256 characters, the volume list is truncated. A truncated statement may result in a backup job failure, and the error message `Invalid command parameter(20)` is displayed.

If the backup selection includes read-only volumes or full volumes, an NDMP backup job fails with the status code 20 (`Invalid command parameter(20)`). If you encounter a similar NDMP backup job error, review the `ostfi` logs to identify the volumes for which the failure occurred. You can use `VOLUME_EXCLUDE_LIST` statements with the `ALL_FILESYSTEMS` statement to exclude the read-only volumes and the volumes with insufficient space.

In a NetBackup Replication Director environment where snapshots are replicated to a secondary filer, Symantec recommends that you use storage lifecycle policies to control backups on the secondary filer.

On NetApp 7-mode storage systems, it is generally not recommended for users to store files in `/vol/vol0` because the volume contains filer system files. For this reason, `vol0` should be excluded from the backup if the `ALL_FILESYSTEMS` directive is used in the backup policy. The following is a backup selection list that excludes `/vol/vol0`:

```
VOLUME_EXCLUDE_LIST=/vol/vol0
ALL_FILESYSTEMS
```

- Do not use `ALL_FILESYSTEMS` to backup all volumes on a secondary filer. Inconsistencies may occur when automatically created NetApp FlexClone volumes are backed up or restored. Such volumes are temporary and used as virtual copies or pointers to actual volumes and as such do not need to be backed up.
- If you must back up all volumes on a secondary filer, Symantec recommends that you exclude the FlexClone volumes as well as replicated volumes. For example:

```
VOLUME_EXCLUDE_LIST=/vol/Clone_*
VOLUME_EXCLUDE_LIST=/vol/*_[0-9]
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9]
```

```
VOLUME_EXCLUDE_LIST=/vol/*_[0-9][0-9][0-9]
ALL_FILESYSTEMS
```

This example assumes all FlexClone volumes and only FlexClone volumes begin with `/vol/Clone_`. Adjust the volume specifications appropriately for your environment.

- **VOLUME_EXCLUDE_LIST** applies only to **ALL_FILESYSTEMS**. It does not apply to explicit backup selections or wildcard-based backup selections. If you use the **ALL_FILESYSTEMS** directive in an NDMP policy for Clustered Data ONTAP, you must exclude each selected SVM's root volume using the **VOLUME_EXCLUDE_LIST** directive. Otherwise the backups fail. For more information, refer to the following topic:

Backups from snapshots for NDMP policies fail when the import of a snapshot fails for volumes where logical unit numbers (LUNs) reside with status code 4213 (Snapshot import failed). To avoid this error, use the **VOLUME_EXCLUDE_LIST** directive to exclude any volumes that are used to create LUNs accessed through a storage area network (SAN).

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
- MS-Windows

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 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 711.

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 823.

See [“ALL_LOCAL_DRIVES example: Without multiple data streams”](#) on page 823.

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 823.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 711.

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 711.

System State:\ directive

The `System State:\` can be used on all supported Windows systems.

The `System State:\` directive is needed for the operating system versions which do not support Shadow Copy Components.

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.
- Windows IA64 systems with EFI System partitions.

Note: In the policies that back up Windows 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	<p>Items that the computer does not require to operate. For example, Active Directory Application Mode and Microsoft Distributed File System Replication (DSFR) folders.</p> <p>See “About Microsoft DSFR backups” on page 670.</p>
Other Data	<p>A category that is intended for future NetBackup releases.</p>

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:

- FlashBackup
- NDMP
- Lotus-Notes
- MS-Exchange-Server

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:

On Windows:

```
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`.

On UNIX:

```
NEW_STREAM
/usr
/lib
NEW_STREAM
/home
/bin
```

This backup selection list contains two data streams:

- The `NEW_STREAM` directive at the top of the list starts administrator-defined streaming and the first stream. This stream backs up `/usr` and `/lib`.
- The second `NEW_STREAM` starts a second data stream that backs up `/home` and `/bin`.

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:

On Windows:

D:\Utilities

On UNIX:

/var

after

On Windows:

D:\Datafiles

On UNIX:

/bin

If an incremental backup is due that night, only changed files in D:\Utilities (on Windows) or in /var (on UNIX) are backed up. Add a `NEW_STREAM` directive before D:\Utilities (on Windows) or before /var (on UNIX), to perform a full backup of all files in D:\Utilities (on Windows) or in /var (on UNIX), 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 822.
- 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 1094.

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.

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 4 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 4 hours and a schedule has a frequency of less than 4 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:

On Windows:

```
install_path\NetBackup\bin\admincmd\bpconfig [-prep hours]
```

On UNIX:

```
/usr/opensv/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 825.

See [“Allow multiple data streams \(policy attribute\)”](#) on page 711.

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 832.

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.
- Files or directories beneath a `bind` mount (Linux).

NetBackup automatically excludes the following file system types on most platforms:

- `cdrom` (all UNIX platforms)
- `cacheefs` (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 696.

See [“Cross mount points \(policy attribute\)”](#) on page 699.

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 when, for example, an entire directory is excluded except for one file on the include list.

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 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.

The **Exclude List** or the **Include List** can also be specified through the **NetBackup Administration Console** on the master server.

See [“Exclude Lists properties”](#) on page 128.

UNIX clients Create the exclude and include lists in the following files on the client:

- `/usr/opensv/netbackup/include_list`
- `/usr/opensv/netbackup/exclude_list`

Specific policy

Create an exclude list for a specific policy or for a policy and a schedule combination. Create an `exclude_list` file with a `.policyname` or `.policyname.schedulename` suffix. The following two file examples use a policy that is named `wkstations`. The policy contains a schedule that is named `fulls`:

```
/usr/openv/netbackup/exclude_list.wkstations
/usr/openv/netbackup/exclude_list.wkstations.fulls
```

The first file affects all scheduled backups in the policy that is named `wkstations`. The second file affects backups only when the schedule is named `fulls`.

For a given backup, NetBackup uses a single exclude list—the list that contains the most specific name. For example, if there are files named:

```
exclude_list.wkstations and
exclude_list.wkstations.fulls
```

NetBackup uses only:

```
exclude_list.wkstations.fulls
```

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.

VMware Policy tab

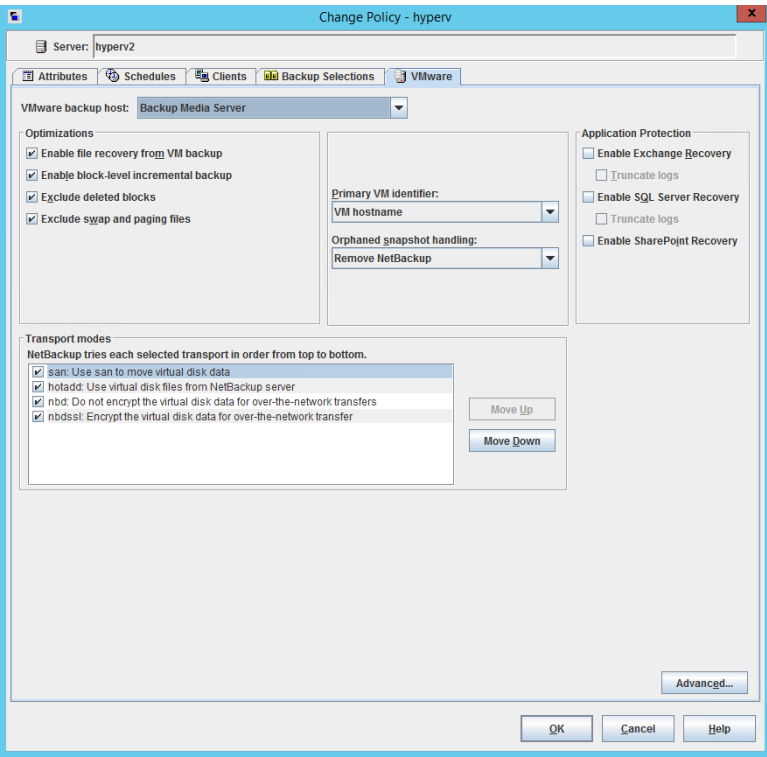
The **VMware** tab appears for policies of the **VMware** policy type.

The following options appear on the **VMware** tab.

Table 20-55 Options on the **VMware** tab of the policy

Option	Description
VMware backup host	Specify the VMware backup host.
Enable file recovery from VM backup	Enable recovery of individual files. You can also recover the entire virtual machine.
Enable block-level incremental backup	Enable block-level backups of the virtual machine.
Exclude deleted blocks	Reduce the size of the backup image by excluding any unused or deleted blocks within the file system on the virtual machine.
Exclude swap and paging files	Reduce the size of the backup image by excluding the guest OS system paging file (Windows) or the swap file (UNIX).
Primary VM identifier	Specify the type of name by which NetBackup recognizes virtual machines when it selects them for backup.
Orphaned snapshot handling	Specify the action that NetBackup takes when a snapshot is discovered before NetBackup creates a new snapshot for the virtual machine backup.
Application Protection	Options to protect applications.
Enable Exchange Recovery	Enable recovery of individual files from Exchange data in the virtual machine backup.
Truncate logs	Truncates the Exchange transaction logs when the backup occurs.
Enable SQL Server Recovery	Enable recovery of individual files from SQL Server data in the virtual machine backup.
Truncate logs	Truncates the SQL Server transaction logs when the backup occurs.
Enable SharePoint Recovery	Enable recovery of individual files from SharePoint data in the virtual machine backup.
Transport modes	Determine how the snapshot data travels from the VMware datastore to the VMware backup host.
Advanced...	Set additional parameters for the VMware backup.

Figure 20-31 VMware tab



For more information about VMware, see the [NetBackup for VMware Guide](#).

Hyper-V Policies tab

The **Hyper-V** tab appears for policies of the **Hyper-V** policy type.

The following options appear on the **Hyper-V** tab.

Table 20-56 Options on the Hyper-V tab of the policy


Option	Description
Optimizations Enable file recovery from VM backup	Allows restore of individual files from the backup. With or without this option, you can restore the entire virtual machine.

Table 20-56 Options on the Hyper-V tab of the policy (*continued*)

Option	Description
Primary VM identifier	Specifies the type of name by which NetBackup recognizes virtual machines when it selects them for backup.
Enable offline backup of non-VSS VMs	Determines whether NetBackup is allowed to perform an offline backup of a virtual machine. This option is intended for guest operating systems that do not support VSS (such as Linux).
Cluster shared volumes timeout	Applies to backups of the virtual machines that are configured in a Microsoft Cluster that uses cluster shared volumes (CSV).
Advanced...	Sets additional parameters for the Hyper-V backup.

For more information about Hyper-V, see the [NetBackup for Hyper-V Guide](#).

Figure 20-32 Hyper-V Policies tab



Disaster Recovery tab

The **Disaster Recovery** tab appears for the **NBU-Catalog** policy type. 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.

Figure 20-33 Disaster Recovery tab




Table 20-57 describes the options on the **Disaster Recovery** tab.

Table 20-57 Options on the Disaster Recovery tab

Option	Description
Path	<p>Browse and 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>Note: The path cannot contain non-ASCII characters.</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 computer. To do so, first mount the UNC location on the master server, and then provide the UNC path to the UNIX computer. <p>On UNIX: The path for the disaster recovery information cannot be to a directory that is on the same partition as <code>/usr/openv/netbackup</code>. If the path is to a location on the same partition as <code>/usr/openv/netbackup</code>, NetBackup displays a status 20 error message. The message states that the disk path is invalid. Change the path on the Disaster Recovery tab to a directory on a different partition.</p>
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 20-57 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> <p><i>email1@domain.com,email2@domain.com</i></p> <p>See “Setting up email notifications about backups” on page 155.</p> <p>On Windows: The <code>nbmail.cmd</code> or <code>mail_dr_info.cmd</code> script must be configured (<i>Install_path\NetBackup\bin\goodies\</i>). In addition specify the email addresses in the Disaster Recovery tab.</p> <p>On Windows: 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>On Windows: 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 <i>Install_path\NetBackup\bin\mail_dr_info.cmd</i> 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 <i>Install_path\NetBackup\bin\goodies\nbmail.cmd</i> is 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 IMAP services. <p>Note: On Windows: By default, neither <code>nbmail.cmd</code> nor <code>mail_dr_info.cmd</code> is configured to send email.</p> <p>See “Configuring the nbmail.cmd script on Windows” on page 156.</p> <p>On Windows: For more information on <code>mail_dr_info.cmd</code>, see the NetBackup Administrator's Guide, Volume II.</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 881.
- 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 > 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:

On Windows:

```
install_path\netbackup\bin
```

On UNIX:

```
/usr/opensv/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 On Windows: Select the policy name in the left pane.
- 3 On UNIX: Select the policy name in the middle pane.
- 4 On the **Actions** menu, click **Manual Backup**. (To perform a manual backup, you must enable the **Go into effect at** attribute.)

See [“Go into effect at \(policy attribute\)”](#) on page 695.

If the **Go into effect at** attribute is set for a future date and time, the backup does not run.

- 5 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).

- 6 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

See the Operating System Compatibility List for a list of operating system versions that support Active Directory granular restore:

<http://www.symantec.com/docs/TECH76648>

To perform Active Directory granular backups and restores, ensure that you meet the following requirements:

- 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 1170.
See “[About configuring Services for Network File System \(NFS\) on Windows 2012 and Windows 2012 R2](#)” on page 1171.
See “[About configuring Services for Network File System \(NFS\) on Windows 2008 and Windows 2008 R2](#)” on page 1179.
- 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 (`bpineta`) must run under the domain administrator account on the Active Directory domain controller or ADAM server. By default, `bpineta` runs under the Local System account.
See “[Configuring the NetBackup Client Service](#)” on page 1092.

For information on the media server platforms that support Granular Recovery Technology see the NetBackup compatibility lists on the [NetBackup landing page](#) on the Symantec Support website.

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 73.

To create a policy to allow Active Directory restores

- 1 Check that the NetBackup Legacy Client Service (`bpineta`) is running under the domain administrator account on the Active Directory domain controller. In this case, the Active Directory domain controller is the NetBackup client.
See [“Configuring the NetBackup Client Service”](#) on page 1092.
- 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 845.
- 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 823.
 - See [“Shadow Copy Components:\ directive”](#) on page 824.
 - See [“ALL_LOCAL_DRIVES directive”](#) on page 822.

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

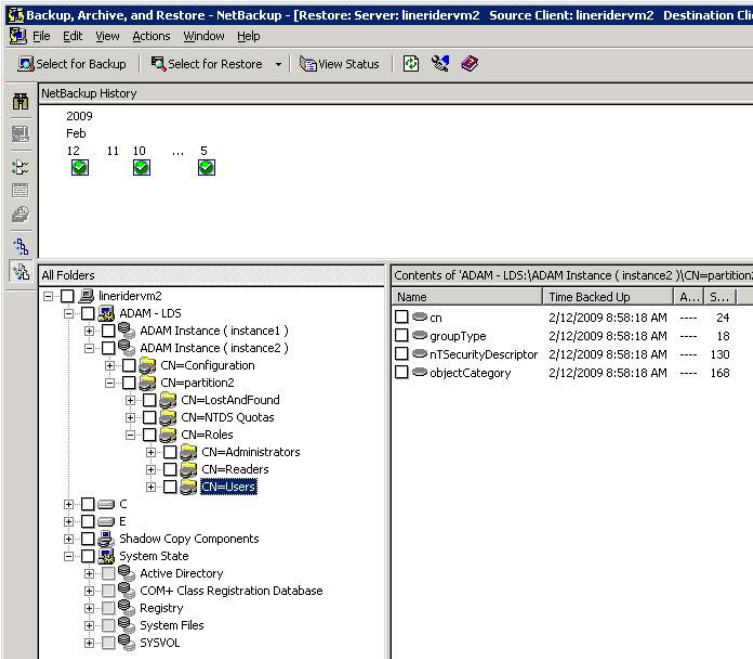
incremental, differential or cumulative. Any other items in the **Backup Selections** list may use a differential or cumulative incremental backup type as indicated. Even though a full backup is forced for an Active Directory backup, normal incremental rules are applied to the non-Active Directory items 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 847.

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 20-58](#) describes potential problems for granular restores.

Table 20-58 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> ■ Computer object credentials change every 30 days and the credentials from the backup may not match the credentials that are stored on the actual computer. When a computer object is restored it is disabled if the userAccountControl property was not preserved in the deleted object. Use the Microsoft Active Directory Users and Computers application to reset the account of a computer object: <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 20-58 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.</p> <p>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](#)
- [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 systems 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 707.
- The schedule that is created for a synthetic backup must have **Synthetic Backup** selected.
See ["Synthetic backup \(schedule attribute\)"](#) on page 754.
- 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 851.
 - 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.
- Symantec recommends that you not enable **Expire after copy** retention for any storage units that are to be used with SLPs with either of the following: Accelerator or synthetic backups. The **Expire after copy** retention can cause images to expire while the backup runs. To synthesize a new full backup, the SLP backup needs the previous backup image. If the previous image expires during the backup, the backup fails.
- 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.
- You can create multiple copies with synthetic backups by using the multiple copies synthetic backup method.
Although synthetic backups do support the use of storage lifecycle policies, SLPs cannot be used for the multiple copy synthetic backups method.
See [“Using the multiple copy synthetic backups method”](#) on page 860.
- Synthetic backups are not supported if any of the component images are encrypted.
- A user-generated backup cannot be used to generate a synthetic image. A backup that is generated from a User Backup schedule or a User Archive schedule cannot be used as one of the components of a synthetic backup.
- Synthetic backups do not support Auto Image Replication.

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 21-1 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>

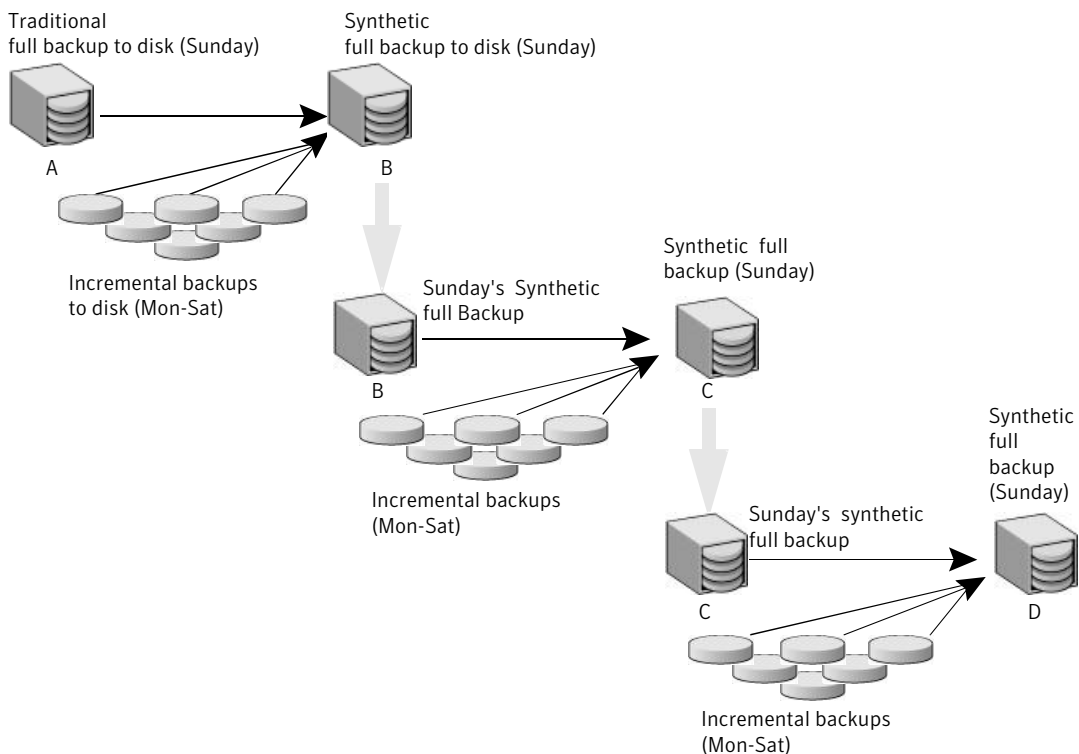
Synthetic full backups

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 21-1](#) 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 21-1 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 855.

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

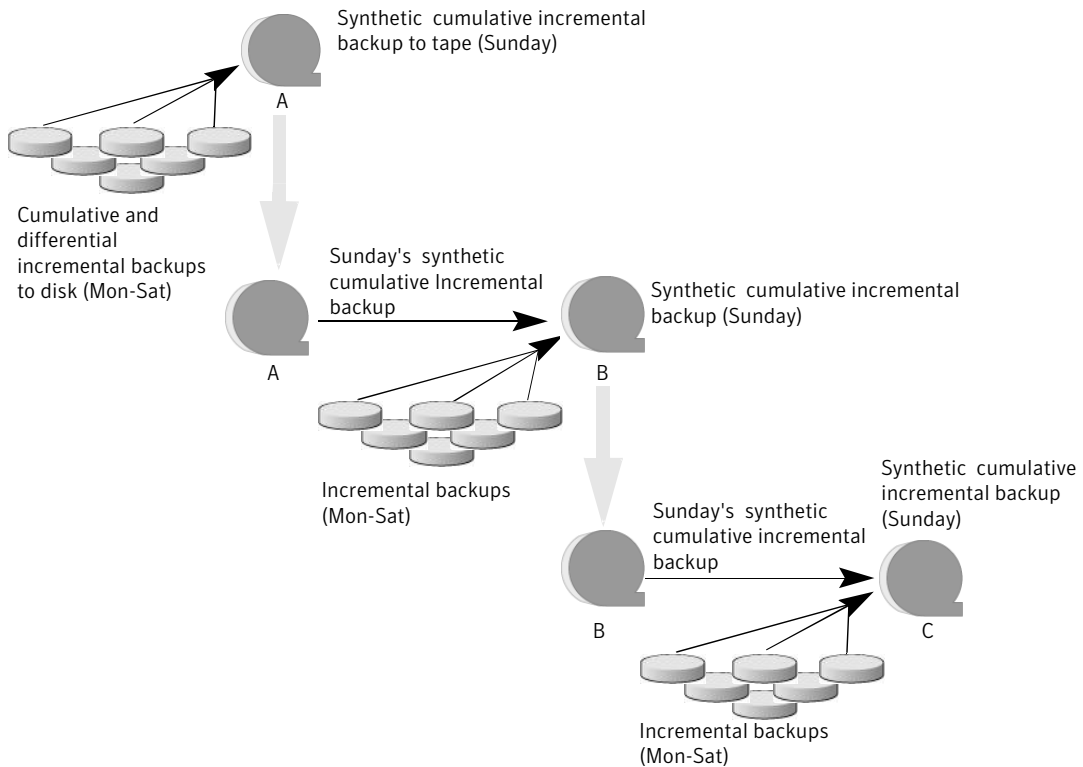
The scenario to create a synthetic, cumulative incremental backup is similar to the scenario to create a synthetic full backup. Remember, a cumulative incremental backup includes all changes since the last full backup.

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 21-2 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 21-2 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 754.

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 707.

Change journal and synthesized backups

If the **Use Change Journal** host property on a Windows client 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 104.

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 707.

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 80.

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 `vnetd`.
- On the media server(s): `bpcd`, `bptm` (if any images are written to or from a tape device), and `bpdm` (if any images are written to or from disk).
Note that several media servers can be involved if the component images are on different nodes.

The log files for synthetic backups are described in the [NetBackup Logging Reference Guide](#).

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.

See [“USE_CTIME_FOR_INCREMENTALS option for NetBackup clients”](#) on page 363.

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 104.

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 21-2](#) 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 21-2 Comparing synthetic copy process with and without method enabled

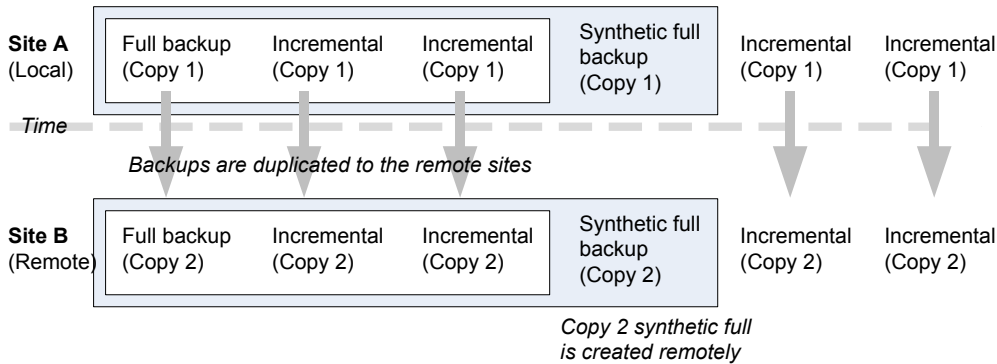
Step	Without using the multiple copy synthetic backups method:	Using the multiple copy synthetic backups method:
1	A full backup is performed at the local site (Site A).	Step 1 remains the same.
2	The full backup is duplicated to the remote site (Site B).	Step 2 remains the same.
3	An incremental backup 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.

Table 21-2 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:
6	A full synthetic backup is produced at Site A.	Step 6 remains the same.
7	The full backup is duplicated to Site B.	A full synthetic backup is produced at Site B from images at Site B. 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 21-3 shows how no extra bandwidth is used to copy the synthetic full backup from Site A to Site B.

Figure 21-3 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:

On Windows:

```
install_path\VERITAS\NetBackup\db\config\multi_synth.policy. schedule
```

On UNIX:

```
/usr/opensv/netbackup/db/config/multi_synth.policy.schedule
```

Configuration variables for multiple copy synthetic backups

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 21-3](#) describes the configuration entries that can be included in the configuration file.

Table 21-3 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. The TARGET_COPY value 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__

Table 21-3 Configuration entries (*continued*)

Entry	Purpose
FAIL_MODE	<p>The second synthetic backup is produced immediately following the first copy synthetic backup if no errors occur during production of the first copy. If an error occurs during the second copy, the <code>FAIL_MODE</code> value specifies the fate of the first copy job and image.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"> ■ <code>FAIL_MODE=ALL</code> ALL means that if the second copy fails, the first copy and its job also fail. (Default.) ■ <code>FAIL_MODE=ONE</code> ONE means that if the second copy fails, the failure does not affect the first copy job.
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"> ■ <code>ENABLED=YES</code> Production of the second copy is enabled. (Default.) ■ <code>ENABLED=NO</code> Production of the second copy is disabled.
SRC_COPY_FALLBACK	<p>Specifies that if a copy by the number given in <code>SRC_COPY</code> or <code>COPY</code> does not exist, the synthetic backup should use the primary backup.</p> <p>The only valid value is the following:</p> <p><code>SRC_COPY_FALLBACK=PRIMARY</code></p>
VOLUME_POOL	<p>Specifies the volume pool for tape media, if one is used. If no volume pool is specified, NetBackup uses the volume pool that is specified in the policy. If a volume pool is entered for disk, the entry is ignored.</p>

Multiple copy synthetic backups 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

NetBackup environments that use the Data Protection Optimization Option license key can benefit from the optimized synthetic backup method. Optimized synthetic backups take advantage of the capabilities of the OpenStorage API.

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.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#)
- [NetBackup OpenStorage Solutions Guide for Disk](#)

Optimized synthetic backups for deduplication

NetBackup environments that use the NetBackup Data Protection Optimization Option license key 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.

For more information, see the following guides:

- [NetBackup Deduplication Guide](#)
- [NetBackup OpenStorage Solutions Guide for Disk](#)

Protecting the NetBackup catalog

This chapter includes the following topics:

- [About the NetBackup catalog](#)
- [Parts of the NetBackup catalog](#)
- [Protecting the NetBackup catalog](#)
- [Recovering the catalog](#)
- [Disaster recovery emails and the disaster recovery file](#)
- [Archiving the catalog and restoring from the catalog archive](#)
- [Estimating catalog space requirements](#)

About the NetBackup catalog

A NetBackup catalog is the internal database that contains 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 Windows 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. It manages and controls access to the following types of data:

- Image metadata (information about backup images and copies (image header files))
- Backup content data (information about the folders, files, and the objects in a backup (.f files))
- NetBackup backup policies.
- NetBackup licensing data.
- The NetBackup error log.
- The client database.

The catalog consists of the following parts:

- Relational database files.
NetBackup uses the SQL Anywhere Relational Database Management System to store information in the NetBackup database (NBDB). The metadata includes information about the data that has been backed up, and about where the data is stored.
In previous releases, much of this information was stored in the image database:
On Windows: `NetBackup/db/images`
On UNIX: `/usr/opensv/netbackup/db/images`
See [“About the NetBackup relational database”](#) on page 870.
- The image database.
The image database contains information about the data that has been backed up.
See [“About the NetBackup image database”](#) on page 868.
- 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 907.
See [“About the databases.conf file”](#) on page 909.

NetBackup is sensitive to the location of the master server components. Running any part of NetBackup (the binaries, the logs, the database, the images) on a network share (NFS, for example) can affect performance of even normal operations.

NetBackup can be CIFS-mounted on SAN or NAS storage as long as the average I/O service times remain less than 20 milliseconds.

The storage must also meet certain conditions to ensure the Sybase requirement for data integrity:

- The order of file writes must be guaranteed.
- When a write request is issued, the write must complete to the physical storage. The write request must not merely be buffered when the SAN or the NAS returns from the write call.

See the following article for more information:

<http://www.symantec.com/docs/TECH158868>

Figure 22-1 Catalog configuration on Windows computers

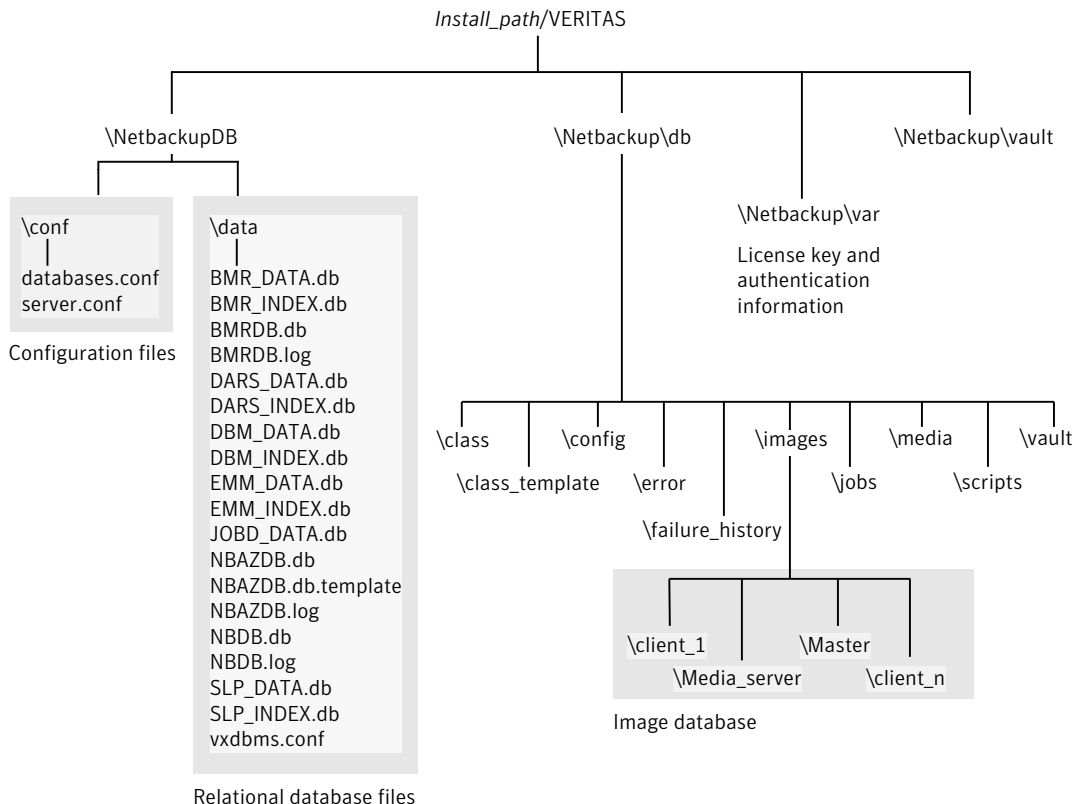
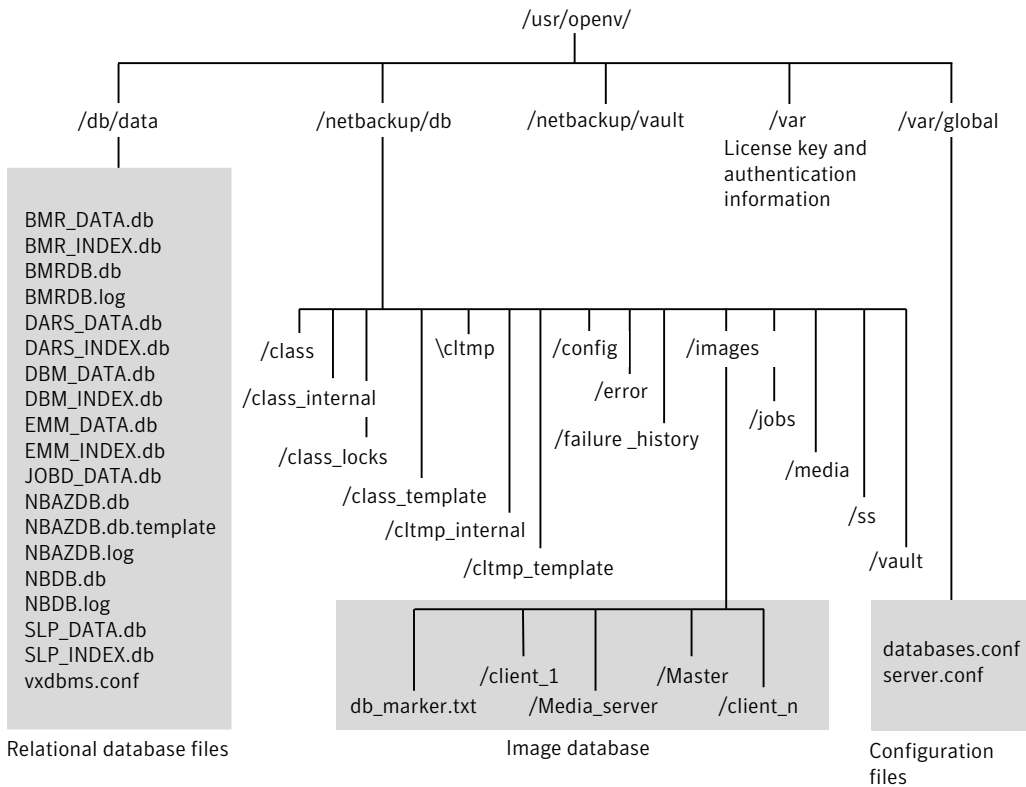


Figure 22-2 Catalog configuration on UNIX computers



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 in the following location:

- **Windows:** Program Files\VERITAS\Netbackup\db\images
- **UNIX:** /usr/openv/netbackup/db/images

The image database contains the following files:

Image files	Files that store only backup set summary information.
.lock files	Used to prevent simultaneous updates on images.
Image .f files	Used to store the detailed information about each file backup.

`db_marker.txt`

Used to ensure that access to the `db` directory is valid when the NetBackup Database Manager starts up. Do not delete this file.

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` (Windows) or `/images` (UNIX) 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 894.

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 896.

The image database component of the NetBackup catalog uses the `.f` files in binary format for Windows, Solaris, 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 [NetBackup Commands Reference Guide](#).

See [“Estimating catalog space requirements”](#) on page 894.

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 contains 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 100 megabytes.

When the backup file of one catalog backup is less than 100 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 100 megabytes.

The following is a UNIX example of an .f file in a single file layout:

```
-rw----- 1 root other  979483 Aug 29 12:23 test_1030638194_FULL.f
```

- Image .f file multiple file layout

When the file information for one catalog backup is greater than 100 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.

See [“About the NetBackup image database”](#) on page 868.

See [“About the NetBackup relational database \(NBDB\) installation”](#) on page 903.

See [“Post-installation tasks”](#) on page 940.

About the NetBackup relational database on Windows

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\BMR_DATA.db* (if BMR is installed)
 - *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_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\JOBDDATA.db*
 - *Install_path\VERITAS\NetBackupDB\data\NBAZDB.db*
 - *Install_path\VERITAS\NetBackupDB\data\NBAZDB.db.template*
 - *Install_path\VERITAS\NetBackupDB\data\NBAZDB.log*
 - *Install_path\VERITAS\NetBackupDB\data\NBDB.db*
 - *Install_path\VERITAS\NetBackupDB\data\NBDB.log*
 - *Install_path\VERITAS\NetBackupDB\data\SEARCH_DATA.db*
 - *Install_path\VERITAS\NetBackupDB\data\SEARCH_INDEX.db*
 - *Install_path\VERITAS\NetBackupDB\data\SLP_DATA.db*
 - *Install_path\VERITAS\NetBackupDB\data\SLP_INDEX.db*

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) or the configuration files to a remote NFS share. CIFS is supported on some SAN or NAS storage.

See [“Parts of the NetBackup catalog”](#) on page 866.

- Configuration files:

- `Install_path\VERITAS\NetBackupDB\data\vxdbms.conf`
- `Install_path\VERITAS\NetBackupDB\conf\server.conf`
- `Install_path\VERITAS\NetBackupDB\conf\databases.conf`

Note: The catalog backup process copies this data to

`Install_path\VERITAS\NetBackupDB\staging` and backs up the copy.

About the NetBackup relational database on UNIX

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

- `/usr/opensv/db/data/BMR_DATA.db` (if BMR is installed)
- `/usr/opensv/db/data/BMRDB.db` (if BMR is installed)
- `/usr/opensv/db/data/BMRDB.log` (if BMR is installed)
- `/usr/opensv/db/data/BMR_INDEX.db` (if BMR is installed)
- `/usr/opensv/db/data/DARS_DATA.db`
- `/usr/opensv/db/data/DARS_INDEX.db`
- `/usr/opensv/db/data/DBM_DATA.db`
- `/usr/opensv/db/data/DBM_INDEX.db`
- `/usr/opensv/db/data/EMM_DATA.db`
- `/usr/opensv/db/data/EMM_INDEX.db`
- `/usr/opensv/db/data/JOBD_DATA.db`
- `/usr/opensv/db/data/NBAZDB.db`
- `/usr/opensv/db/data/NBAZDB.db.template`

- /usr/opensv/db/data/NBAZDB.log
- /usr/opensv/db/data/NBDB.db
- /usr/opensv/db/data/NBDB.log
- /usr/opensv/db/data/SEARCH_DATA.db
- /usr/opensv/db/data/SEARCH_INDEX.db
- /usr/opensv/db/data/SLP_DATA.db
- /usr/opensv/db/data/SLP_INDEX.db

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) or the configuration files to a remote NFS share. CIFS is supported on some SAN or NAS storage.

See [“Parts of the NetBackup catalog”](#) on page 866.

- Configuration files
 - /usr/opensv/db/data/vxdbms.conf
 - /usr/opensv/var/global/server.conf
 - /usr/opensv/var/global/databases.conf

Note: The catalog backup process copies this data to /usr/opensv/db/staging and backs up the copy.

About the Enterprise Media Manager (EMM)

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 that resides on the master server. The NetBackup Resource Broker queries EMM to allocate storage units, drives (including drive paths), and media.

EMM contains the following information:

- Device attributes
- Robotic library and standalone 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
- Log on credentials for storage servers, disk arrays, and NDMP hosts
- Fibre Transport attributes

EMM ensures consistency between drives, robotic libraries, storage units, media, and volume pools across multiple servers. EMM contains information for all media servers that share devices in a multiple server configuration.

The NetBackup scheduling components use EMM information to select the server, drive path, and media for jobs. When the device manager `ltid` starts up, it reads device information from EMM 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 870.

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.

Note: To perform a catalog backup, the master server and the media server must both be at the same NetBackup version.

See the [NetBackup Installation Guide](#) for information about mixed version support.

As additional protection for the catalog, consider archiving the catalog.

See “[Archiving the catalog and restoring from the catalog archive](#)” on page 887.

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 catalog backups

The catalog backup is designed for active environments in which continual backup activity occurs. The catalog backup can be performed while regular backup activity occurs.

The 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 catalog lets you recover either the entire catalog or pieces of the catalog. (For example, the databases separately from the image catalog.)

The catalog backup performs the following tasks:

- Backs up the catalog while continual client backups are in progress.
- Performs a full or an incremental catalog backup.
- Runs the scheduled catalog backups.
- Copies the relational database files to the staging directory and then backs up that directory.
- Catalog backups to tape also include the following items:
 - Spans multiple tapes for a catalog backup.

- Allows for a flexible pool of catalog tapes.
Catalog backups to tape use media from the **CatalogBackup** volume pool only.
- Appends to existing data on tape.

You can configure a catalog backup by using one of the following methods:

- By using wizards:
 - The Catalog Backup Wizard.
See [“Using the Catalog Backup Wizard to configure a catalog backup”](#) on page 876.
 - The Policy Configuration Wizard.
See [“Using the Policy Configuration Wizard to configure a catalog backup”](#) on page 879.

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 881.

Using the Catalog Backup 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.

Note: To perform a catalog backup, the master server and the media server must both be at the same NetBackup version.

See the [NetBackup Installation Guide](#) for information about mixed version support.

To use the Catalog Backup Wizard to configure a catalog backup

- 1 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Configure the Catalog Backup** to launch the **NetBackup Catalog Backup Wizard**.
Click **Help** within any wizard panel for more information on the wizard settings.
- 3 Click **Next** on the **Welcome** panel.

- 4 On the **NetBackup Catalog Backup Policy** panel, select a policy from the list of existing catalog backup policies.
- 5 Or, to create a new catalog backup policy, select **Create a new catalog backup policy**. Click **Next**.
- 6 In the **Policy Name and Type** wizard panel, enter the policy name. Notice that **NBU-Catalog** is automatically selected as the policy type.

Type a unique name for the new policy in the **Add a New Policy** dialog box.

See [“NetBackup naming conventions”](#) on page 1093.

Click **Next**.

- 7 On the **Backup Type** wizard panel, select the backup type. The **User Backup** does not apply for NBU-Catalog policies. Click **Next**.
- 8 On the **Rotation** wizard panel, 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 catalog backups with other backups”](#) on page 884.

Click **Next**.

- 9 In the **Start Window** wizard panel, 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 panel, browse or 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. You must also record this location so that Symantec Catalog can be recovered, if required.

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 panel, 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 886.

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 panel 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 panel displays the total number of catalog backup policies for this master server. Click **Finish** to complete the wizard.
- 16 You may want to add critical policies to the **Critical Policies** list. Specify some policies as critical policies after the **Catalog Backup Wizard** is complete. A policy that is listed on the **Critical Policies** list is considered crucial to the recovery of a site in the event of a disaster.

See [“Adding policies to the Critical Policies list of a catalog backup policy”](#) on page 840.

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 886.

See [“Determining whether or not a catalog backup succeeded”](#) on page 885.

Using the Policy Configuration 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.

Note: To perform a catalog backup, the master server and the media server must both be at the same NetBackup version.

See the [NetBackup Installation Guide](#) for information about mixed version support.

To use the Policy Configuration Wizard to configure a catalog backup

- 1 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**.
- 2 In the right pane, click **Create a Policy** in the right pane to launch the **Policy Configuration Wizard**.

Click **Help** within any wizard panel for more information on the wizard settings.

- 3 Select the **Files systems, databases, applications** option from the list. Click **Next**.

- 4 In the **Policy Name and Type** wizard panel, enter the policy name. Select **NBU-Catalog** as the policy type.

Click **Next**.

See [“NetBackup naming conventions”](#) on page 1093.

- 5 On the **Backup Types** wizard panel, select the backup type. The **User Backup** does not apply for NBU-Catalog policies. Click **Next**.

- 6 On the **Frequency and Retention** wizard panel, 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 catalog backups with other backups”](#) on page 884.

Click **Next**.

- 7 In the **Start Window** wizard panel, 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.

- 8 On the **Catalog Disaster Recovery File** wizard panel, browse or 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. You must also record this location so that NetBackup Catalog can be recovered, if required.

Click **Next**.

- 9 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 886.

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.

- 10 Click **Finish** to complete the wizard.
- 11 You may want to add critical policies to the **Critical Policies** list. Specify some policies as critical policies after the **Policy Configuration Wizard** is complete. A policy that is listed on the **Critical Policies** list is considered crucial to the recovery of a site in the event of a disaster.

See [“Adding policies to the Critical Policies list of a catalog backup policy”](#) on page 840.

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 a catalog backup manually

- 1 In the **NetBackup Administration Console**, expand **NetBackup Management > Policies**.
- 2 Select **Actions > New > Policy**.

- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.

See [“NetBackup naming conventions”](#) on page 1093.

Click **OK**.

- 4 On the **Attributes** tab, complete the following entries:

- Policy Type
Select **NBU-Catalog** as the policy type.
- Policy storage
For disk storage units, increase the **Maximum Concurrent Jobs** storage unit setting to ensure that the catalog backup can proceed during regular backup activity.
See [“Maximum concurrent jobs storage unit setting”](#) on page 565.

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 versions, a media server at a level other than the master server could be selected.

- Policy volume pool
NetBackup automatically creates a **CatalogBackup** volume pool that is selected by default only for **NBU-Catalog** policy types.
 - For other policy attribute descriptions, see the following topic:
See [“Policy Attributes tab”](#) on page 680.
- 5 Select the **Schedules** tab to set up a schedule for a catalog backup.
See [“Concurrently running catalog backups with other backups”](#) on page 884.
See [“About catalog policy schedules”](#) on page 884.
See [“Schedule Attributes tab”](#) on page 744.

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.

- Browse or 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 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 standalone drive and the standalone 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 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 842.

You can also run the `bpbbackup` command from the command line to perform a catalog backup.

More information is available in the [NetBackup Commands Reference Guide](#).

Concurrently running catalog backups with other backups

You can schedule catalog backups to run concurrently with other backup types for the master server.

Make the following adjustments to ensure that the catalog backup can proceed while regular backup activity occurs:

- Set the **Maximum jobs per client** value to greater than one. The property is found in the Global Attributes host properties for the master server.
See [“Global Attributes properties”](#) on page 150.
- Increase the **Maximum concurrent jobs** setting on the storage unit where the backups are sent.
See [“Maximum concurrent jobs storage unit setting”](#) on page 565.

See [“Determining whether or not a catalog backup succeeded”](#) on page 885.

See [“Strategies that ensure successful NetBackup catalog backups”](#) on page 886.

About catalog policy schedules

When you work with catalog policy schedules, consider the following:

- The schedules that are supported in the 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.

How catalog incrementals and standard backups interact on UNIX

A catalog backup policy can include both full catalog backups and incremental catalog backups. However, incremental catalog backups differ from incremental standard backups. Catalog backups use both `mtime` and `ctime` to identify changed data. Standard incremental backups use only `mtime` to identify changed data.

Because of this difference, running a standard policy type backup that includes the `/usr/openv/netbackup/db/images/` directory can adversely affect incremental catalog backups. When standard backups run, they reset the file access time (`atime`). In turn, the reset changes the `ctime` for files and directories. If an incremental catalog backup runs, it sees that the `ctime` has changed and backs up the files. The backup may be unnecessary since the files may not have changed since the last catalog backup.

To avoid additional processing during catalog backups, Symantec recommends the following:

If incremental catalog backups are configured, exclude the NetBackup `/usr/openv/netbackup/db/images/` directory from standard backups.

To exclude that directory, create a `/usr/openv/netbackup/exclude_list` file on the master server.

See [“About NetBackup master server installed directories and files”](#) on page 906.

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.

An email message is sent to the address that is indicated in the **Disaster Recovery** settings for a catalog backup.

Configure this email with the `mail_dr_info.cmd` (on Windows) or the `mail_dr_info` script (on UNIX).

See the [Administrator's Guide, Volume II](#) for more information on setting up this script.

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. These are the only methods 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.
- If you back up your catalogs to disk, 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.

Note: If a catalog backup is on tape, the tape must be removed when the backup is finished or regular backups cannot proceed. NetBackup does not mix catalog and regular backups on the same tape.

Recovering the catalog

Catalog recovery is discussed in the [NetBackup Troubleshooting Guide](#).

Disaster recovery emails and the disaster recovery file

The **Catalog Backup Wizard** and the **Policy Configuration 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 panel. 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.

NetBackup emails the disaster recovery file when the following events occur:

- The catalog is backed up.
- A catalog backup is duplicated or replicated.
- The primary catalog backup or any copy expires automatically or is expired manually.

On Windows: 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 and restoring from the catalog archive

Catalog archiving 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 (files 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 should not be used as a method to reclaim disk space when a catalog file system fills up. In that situation, investigate catalog compression or add disk space to grow the file system.

For additional catalog archiving considerations, see the following topic:

See [“Catalog archiving considerations”](#) on page 893.

To archive the catalog and restore the catalog archive

- 1 Use `bpcatlist` to determine what images are available to be archived.

Running `bpcatlist` alone does not modify any catalog images. Only when the `bpcatlist` output is piped to `bpcatarc` are the `.f` files backed up, and only when the output is piped to `bpcatrm` will the `.f` files be deleted from disk.

To determine what images have `.f` files on disk that can be archived, run the following command. The `catarcid` column indicates whether the `.f` file is not currently backed up (0) or the `catarcid` of the backup of that image.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -online
```

To determine what images have been previously archived and removed from disk, run the following command.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -offline
```

The catalog commands are described in detail in the following topic:

See [“Catalog archiving commands”](#) on page 891.

Note: If catalog archiving has not been previously run, this command should return: `No entity was found.`

For example, to display all images for a specific client before January 1, 2010, run the following command:

```
bpcatlist -client name -before Jan 1 2010
```

To display the help for the `bpcatlist` command run this command.

```
bpcatlist -help
```

Once the `bpcatlist` output correctly lists all the images that are to be archived or deleted, then other commands can be added.

2 Running the catalog archive.

Before running the catalog archive, create a backup policy named **catarc**. The policy is required for the `bpcatarc` command to successfully process images. The name of the policy reflects that the purpose of the schedule is for catalog archiving.

See the following topic for details about configuring the **catarc** policy:

See [“Creating a catalog archiving policy”](#) on page 890.

To run the catalog archive, first run the `bpcatlist` command with the same options used in [step 1](#) to display images. Then pipe the output through `bpcatarc` and `bpcatrm`.

```
bpcatlist -client all -before Jan 1 2010 | bpcatarc | bpcatrm
```

A new job appears in the **Activity Monitor**. The command waits until the backup completes before it returns the prompt. The command reports an error only if the catalog archive fails, otherwise the commands return to the prompt.

The **File List**: section of the Job Details in the **Activity Monitor** displays a list of image files that have been processed. When the job completes with a status 0, the `bpcatrm` command removes the corresponding `.f` files. If the job fails, no catalog `.f` files are removed.

If `bpcatlist` is piped to `bpcatarc` but the results are not piped to `bpcatrm`, then the backup occurs but the `.f` files are not removed from disk. The same `bpcatlist` command can then be rerun and piped to `bpcatrm` to remove the `.f` files.

3 Restoring the catalog archive.

To restore the catalog archive, first use the `bpcatlist` command to list the files that need to be restored. Once `bpcatlist` displays the proper files to restore, run the `bpcatres` command to restore the actual files.

To restore all the archived files from [step 2](#), run the following command:

```
bpcatlist -client all -before Jan 1 2010 | bpcatres
```

This command restores all of the catalog archive files before January 1, 2010.

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.

See [“Go into effect at \(policy attribute\)”](#) on page 695.
- 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:

On Windows: `Install_path\NetBackup\db\images`

On UNIX: `/usr/openv/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 22-1 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, 2012, the following would be entered:</p> <pre>bpcatlist -client all -before Jan 1 2012 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 (0) if the image was never archived. ■ Archived status (S). Indicates whether the catalog was archived (2) or was not archived (1). ■ Compressed status (C). Indicates whether the catalog was compressed (<i>positive_value</i>) or was not compressed (0). ■ 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_097238 Oct 24 10:47:12 2012 ... 973187218 1 0 alpha_097238_UBAK.f alpha_097233 Oct 23 22:32:56 2012 ... 973187218 1 0 alpha_097233_FULL.f alpha_097232 Oct 23 19:53:17 2012 ... 973187218 1 0 alpha_097232_UBAK.f</pre> <p>More information is available in the NetBackup Commands Reference Guide.</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 890.</p>

Table 22-1 Catalog archiving commands (*continued*)

Command	Description
<code>bpcatrm</code>	<p>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 .f files from the online catalog.</p> <p><code>bpcatrm</code> does not remove one .f file unless the file has been previously backed up using the catarc policy.</p>
<code>bpcatres</code>	<p>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 .f files to the catalog. For example:</p> <pre>bpcatlist -client all -before Jan 1 2012 bpcatres</pre>

Catalog archiving considerations

Consider the following items before catalog archiving:

- Perform catalog archiving operations when NetBackup is in an inactive state (no jobs are running).
- Catalog archiving modifies existing catalog images. As a result, it should never be run when the catalog file system is 100% full.
- 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 219.
- Additional time is required to mount the tape and perform the restore of archived .f files.
- There is no simple method to determine to which tape the catalog has been archived. The `bpcatlist -offline` command is the only administrative command to determine what images have been archived. This command does not list what tape was used for the archive. As a result, exercise caution to ensure that the tapes used for catalog archiving are available for restoring the archived catalog images. Either create a separate volume pool to use exclusively for catalog archives or find a method to label the tape as a catalog archive tape.

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:

```
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
- Whether you have enabled the database compression option.

To estimate the disk space that is required for a catalog backup

- 1 Estimate the maximum number of files that each schedule for each policy backs up during a single backup of all its clients.
- 2 Determine the frequency and the retention period of the full and the incremental backups for each policy.

- 3 Use the information from steps 1 and 2 to calculate the maximum number of files that exist at any given time.

For example:

Assume that you schedule full backups to occur every seven days. The full backups have a retention period of four weeks. Differential incremental backups are scheduled to run daily and have a retention period of one week.

The number of file paths you must allow space for is four times the number of files in a full backup. Add to that number one week's worth of incremental backups.

The following formula expresses the maximum number of files that can exist for each type of backup (daily or weekly, for example):

Files per Backup × Backups per Retention Period = Max Files

For example:

A daily differential incremental schedule backs up 1200 files and the retention period for the backup is seven days. Given this information, the maximum number of files that can exist at one time are the following:

$$1200 \times 7 \text{ days} = 8400$$

A weekly full backup schedule backs up 3000 files. 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 on UNIX systems

File system limitations on UNIX 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 gigabytes.
- Set the file size limit for the root user account to unlimited to support large file support.

About the binary catalog format

The catalog in a binary file format has several advantages over the catalog in a text format:

- The catalog is more compact. The binary representations of numbers, dates, and other information, takes up less disk space than the text representations.
- The catalog is much faster to browse and search, especially for large file sizes.
- The catalog supports alternate backup methods without the need to post-process images, which 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 869.

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.

Notes about moving the image catalog

- NetBackup does not support saving the catalog to a remote NFS share. CIFS is supported on some SAN or NAS storage.
See [“Parts of the NetBackup catalog”](#) on page 866.
- 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, on Windows, do not use the `ALTPATH` mechanism to move `install_path\NetBackup\db\error`.
For example, on UNIX, do not move `/usr/opensv/netbackup/db/error`. The catalog backup only follows the symbolic link when backing up the `/images` directory. So, if symbolic links are used for other parts of the NetBackup catalog, the files in those parts are not included in the catalog backup.
- The directory that is specified in the `ALTPATH` file is not automatically removed if NetBackup is uninstalled. If NetBackup is uninstalled, you must manually remove the contents of this directory.

Moving the image catalog between Windows hosts

To move the image catalog on Windows

- 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 883.
- 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 **Daemons** tab in the **Activity Monitor** to stop the Request Manager and the Database Manager daemons. 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 in the **Daemons** tab.

Backups and restores can now resume for the client.

Moving the image catalog between UNIX hosts

To move the image catalog on UNIX

- 1 Check that no backups are in progress by running:

```
/usr/opensv/netbackup/bin/bpps
```

- 2 Stop `bprd` by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

- 3 Stop `bpdbm` by running:

```
/usr/opensv/netbackup/bin/bpdbm -terminate
```

- 4 Create the directory in the new file system. For example:

```
mkdir /disk3/netbackup/db/images
```

- 5 Move the image catalog to the new location in the other file system.

- 6 Create a symbolic link from `/usr/opensv/netbackup/db/images` to the new location in the other file system.

See [“NetBackup file size considerations on UNIX systems”](#) on page 896.

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 150.

Note: Symantec discourages manually compressing or decompressing the catalog backups with the `bpimage -[de]compress` command 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.

Before NetBackup 7.6, NetBackup used the `compress` command on the server to perform compression after each backup session. NetBackup now uses a built-in compression method to compress the image catalog. Image catalog files which were compressed using the older compression method are now migrated transparently to the new method as they are accessed as part of NetBackup operations.

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.

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.

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. Or, on UNIX clients, change the timeout value by changing the `LIST_FILES_TIMEOUT` option in the `bp.conf` file.

Figure 22-3 List Files Timeout General property on the client



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 on Windows

- 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 894.
- 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 150.

- 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`.

To uncompress the NetBackup catalog on UNIX

- 1 Perform the following steps as root on the master server to uncompress the NetBackup catalog.

Verify that the partition where the image catalog resides has enough space to uncompress the client's image records.

- 2 Stop the request daemon, `bprd`, by running:

```
/usr/opensv/netbackup/bin/admincmd/bprreq -terminate
```

- 3 Make sure that `bpdbm` is running:

```
/usr/opensv/netbackup/bin/bpps
```

- 4 Expand **Host Properties > Master Servers**. Open the properties of a host. On the **Global Attributes** properties, clear the **Compress Catalog Interval** check box.

See [“Global Attributes properties”](#) on page 150.

- 5 Set the **Compress Catalog Interval** Global Attributes property to 0.
- 6 Change your working directory to `/usr/opensv/netbackup/bin` and run the command:

```
admincmd/bpimage -decompress -client name
```

- 7 Restart the request daemon, `bprd`, by running:

```
/usr/opensv/netbackup/bin/initbprd
```

- 8 Perform the file restorations from the client.
- 9 Set the **Compress Catalog After** Global Attributes property to its previous value.

The records that were uncompressed for this client are compressed after the next backup schedule.

See [“About image catalog compression”](#) on page 899.

About the NetBackup relational database

This chapter includes the following topics:

- [About the NetBackup relational database \(NBDB\) installation](#)
- [Using the NetBackup Database Administration utility on Windows](#)
- [Using the NetBackup Database Administration utility on UNIX](#)
- [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). 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 \(EMM\)”](#) on page 873.

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) or the configuration files to a remote NFS share. CIFS is supported on some SAN or NAS storage.

See [“Parts of the NetBackup catalog”](#) on page 866.

Installing the NetBackup database

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:

On Windows: `Install_path\VERITAS\NetBackupDB\conf\server.conf`

On UNIX: `/usr/openv/var/global/server.conf`

See [“About the NetBackup server.conf file”](#) on page 907.

- 2 On Windows: 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 915.

- 3 On UNIX: The following entry is added to the `bp.conf` file to set the database location:

`VxDBMS_NB_DATA = /usr/openv/db/data`

- 4 On Windows: 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

- 5 On UNIX: The VxDBMS configuration file for NetBackup is created. This file requires the read and write permissions of `root`:

/usr/opensv/db/data/vxdbms.conf

- 6 The NetBackup database is created:

On Windows: *Install_path\VERITAS\NetBackupDB\data\NBDB.db*

On UNIX: */usr/opensv/db/data/NBDB.db*

- 7 DBA password is set for the NetBackup database in *vxdbms.conf*:

VXDBMS_NB_PASSWORD = encrypted_password

- 8 Additional database files are created with contiguous space pre-allocated:

- The NetBackup system database file that is mentioned in the following step:

On Windows: *Install_path\VERITAS\NetBackupDB\data\NBDB.db*

On UNIX: */usr/opensv/db/data/NBDB.db*

- The EMM database files:

On Windows:

Install_path\VERITAS\NetBackupDB\data\EMM_DATA.db

Install_path\VERITAS\NetBackupDB\data\EMM_INDEX.db

On UNIX:

/usr/opensv/db/data/EMM_DATA.db

/usr/opensv/db/data/EMM_INDEX.db

- The authorization database files for NetBackup Access Control:

On Windows: *Install_path\VERITAS\NetBackupDB\data\NBAZDB.db*

On UNIX: */usr/opensv/db/data/NBAZDB.db*

- The NetBackup transaction log, necessary for recovering the database:

On Windows: *Install_path\VERITAS\NetBackupDB\data\NBDB.log*

On UNIX: */usr/opensv/db/data/NBDB.log*

- 9 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`.)
- 10 The following command initializes the EMM data:
On Windows: `Install_path\VERITAS\Volmgr\bin\tpext.exe`
On UNIX: `/usr/opensv/volmgr/bin/tpext`
See [“About catalog backups”](#) on page 875.

About NetBackup master server installed directories and files

SQL Anywhere is installed in the following directories:

On Windows

`Install_path\VERITAS\NetBackupDB`

`Install_path\VERITAS\NetBackup\bin`

The files in `Install_path\VERITAS\NetBackupDB\conf` can be shared within a cluster.

On UNIX

`/usr/opensv/db`

`/usr/opensv/var/global`

The files in `/global` can be shared within a cluster.

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.

On Windows: The NBDB database files can be moved from their default location in `Install_path\VERITAS\NetBackupDB\data`.

On UNIX: The NBDB database files can be moved from their default location in `/usr/opensv/db/data`.

See [“Moving NBDB database files after installation”](#) on page 941.

Note: NetBackup does not support saving the NetBackup relational database (NBDB, including NBAZDB and EMM) or the configuration files to a remote NFS share. CIFS is supported on some SAN or NAS storage.

See [“Parts of the NetBackup catalog”](#) on page 866.

Note: If Bare Metal Restore is installed, BMRDB must be located on the master server.

See [“Moving NBDB database files after installation”](#) on page 941.

See [“Moving the NetBackup database from one host to another”](#) on page 954.

See [“Moving the NetBackup database files”](#) on page 927.

See [“About the Move Database menu options”](#) on page 937.

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.

On Windows

`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 200M -ch 1365M -cl 200M -wc- -gna 1 -k -os 1M -m -o
"C:\Program Files\Veritas\NetBackupDB\log\server.log" -m
```

On UNIX

The `/usr/openv/var/global/server.conf` file is read when the SQL Anywhere daemon is started.

The SQL Anywhere daemon gets all configuration information from the `server.conf` file, as follows:

```
-n NB_server_name

-x tcpip(LocalOnly=YES;ServerPort=13785) -gp 4096 -gd DBA -gk DBA
-gl DBA -ti 0 -c 200M -ch 1278M -cl 200M -wc- -ua -gna 1 -k -os 1M
-m -o /usr/openv/db//log/server.log -ud
```

NB_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 23-1 Commands used in the server.conf file

Command	Description
-c 200M	Indicates the initial memory that is reserved for caching database pages and other server information. The value may be changed for performance reasons.
-ch 1316M	Indicates the maximum cache size, as a limit to automatic cache growth. The value is calculated automatically based on memory size of the computer.
-cl 200M	Indicates the minimum cache size, as a limit to automatic cache resizing. The value may be changed for performance reasons.
-gd DBA -gk DBA -gl DBA	Indicates that the DBA user is the account used to start, stop, load, and unload data.
-gna 1	Enable automatic tuning of the network database server multiprogramming level.
-gp 4096	Indicates the maximum page size (in bytes) for the database. This parameter is given during database creation.
-k	If -k is specified when database server is started, the database server does not collect Performance Monitor statistics.
-m	Truncates the transaction log when a checkpoint is done. This option provides a way to automatically limit the growth of the transaction log. Disable it using <code>NbDbAdmin.exe</code> (Windows) or <code>dbadm</code> (UNIX) by changing the transaction mode option set to FULL from PARTIAL.
-o <i>Install_path</i> \VERITAS\NetBackupDB\log\server.log	On Windows: 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.
-o /usr/opensv/db/log/server.log	On UNIX: 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.
-os 1M	Specifies a maximum size for the database server message log file.

Table 23-1 Commands used in the server.conf file (*continued*)

Command	Description
-ti 0	Indicates the client idle time that is allowed before shutdown. By default, no idle time is allowed, which prevents the database from shutting down.
-ua	Turns off the use of asynchronous I/O.
-ud	Indicates that the server should run as a daemon.
-wC-	Control whether checksums are enabled on write operations for the database. By default all database pages have checksums, so the database server defaults to -wC-.
-x tcpip(LocalOnly =YES;ServerPort=13785)	Indicates what kind of connections are allowed in addition to shared memory. For example, local TCP/IP connections that use port 13785.

See [“About NetBackup master server installed directories and files”](#) on page 906.

About the databases.conf file

The Windows `Install_path\VERITAS\NetBackupDB\conf\databases.conf` configuration file and the UNIX `/usr/opensv/var/global/databases.conf` configuration file contain 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:

On Windows

```
"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
```

On UNIX

```
"/usr/opensv/db/data/NBAZDB.db" -n NBAZDB
"/usr/opensv/db/data/NBDB.db" -n NBDB
"/usr/opensv/db/data/BMRDB.db" -n BMRDB
```

About the registration.dat file on Windows

This file is created for use with Symantec NetBackup OpsCenter.
It is created in the following location:

`Install_path\VERITAS\NetBackupDB\conf\registration.dat`

About the `vxdbms_env.csh`, `vxdbms_env.sh` scripts (on UNIX)

The scripts `vxdbms_env.csh` and `vxdbms_env.sh` set-up the SQL Anywhere environment as follows:

- `/usr/opensv/db/vxdbms_env.csh`
- `/usr/opensv/db/vxdbms_env.sh`

Other scripts and commands use `vxdbms_env.csh` and `vxdbms_env.sh`.

About the bin directory

On Windows, `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 perform the tasks based on the `nldb` commands. See [“Using the NetBackup Database Administration utility on Windows”](#) on page 917.
To access the utility, the administrator must know the database password. See [“Changing the NetBackup database password”](#) on page 940.
- `create_nldb.exe`
Used during installation and upgrades to create and upgrade the NetBackup database, NBDB.
- `nldb_admin.exe`
Among other things, use `nldb_admin.exe` to change the DBA and NetBackup account passwords, or to start and stop individual databases.
- `nldb_backup.exe`
Use to make an online backup of the SQL Anywhere database files to a file system directory.

Note: Using this command (or the 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.

- `nldb_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.

On UNIX, `/usr/openv/db/bin` contains SQL Anywhere commands and NetBackup-specific commands for administrating NBDB and BMRDB:

- `dbadm`
 This file launches the NetBackup Database Administration utility, which provides administrators with a way to perform the tasks based on the `nbdb` commands. See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931. To access the utility, the administrator must know the database password. See [“Changing the NetBackup database password”](#) on page 940.
- `create_nbdb`
 Used during installation and upgrades to create and upgrade the NetBackup database, NBDB.
- `nbdb_admin`
 Among other things, use `nbdb_admin` to change the DBA and NetBackup account passwords, or to start and stop individual databases.
- `nbdb_backup`
 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`
 Use to change the location of the SQL Anywhere database files from the default location.
- `nbdb_ping`
 Displays the status of the SQL Anywhere database.
- `nbdb_restore`
 Use to recover from an online backup in a file system directory that was created using `nbdb_backup`.
- `nbdb_unload`
 Use to create a dump of all or part of the NBDB database or the BMRDB database schema and data.
- `nbdbms_start_server`
 Use to start and stop the SQL Anywhere daemon.
- `nbdb_upgrade`
 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.

About the contents of the NetBackup directories

The following table describes the contents of the NetBackup directories.

Table 23-2 NetBackup directory contents

Directory	Description
charsets	<p>The following directory contains SQL Anywhere-specific information.</p> <p>On Windows: <code>Install_path\VERITAS\NetBackupDB\charsets</code></p> <p>On UNIX: <code>/usr/opensv/db/charsets</code></p>

Table 23-2 NetBackup directory contents (*continued*)

Directory	Description
lib	On UNIX: The directory <code>/usr/opensv/db/lib</code> contains all the SQL Anywhere shared libraries. The directory also includes ODBC libraries, used to connect to NBDB and BMRDB.
log	The following directory contains the SQL Anywhere server log file <code>server.log</code> that contains only Sybase logs. On Windows: <code>Install_path\VERITAS\NetBackupDB\log</code> On UNIX: <code>/usr/opensv/db/log</code>
res	On UNIX: The directory <code>/usr/opensv/db/res</code> contains SQL Anywhere-specific information.
scripts	The following directory 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. On Windows: <code>Install_path\VERITAS\NetBackupDB\scripts</code> On UNIX: <code>/usr/opensv/db/scripts</code> Note: Do not edit the scripts that are located in this directory.
staging	The directory that is used as a temporary staging area during catalog backup and recovery. On Windows: <code>Install_path\VERITAS\NetBackupDB\staging</code> On UNIX: <code>/usr/opensv/db/staging</code>
tix	On UNIX: The directory <code>/usr/opensv/db/tix</code> contains SQL Anywhere-specific information.
WIN32	On Windows: The directory <code>Install_path\VERITAS\NetBackupDB\WIN32</code> contains SQL Anywhere commands and .dll files.
java	On Windows: The directory <code>Install_path\VERITAS\NetBackupDB\java</code> contains the components necessary to run the NetBackup Administration Console on a Windows computer.
shared	On Windows: The directory <code>Install_path\VERITAS\NetBackupDB\shared</code> is a directory used by Symantec NetBackup OpsCenter.

About the data directory

The following directory is the default location of the NetBackup database, NBDB:

On Windows: `Install_path\VERITAS\NetBackupDB\data`

On UNIX: `/usr/opensv/db/data`

NBDB 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 catalog backup of the SQL Anywhere database.
- `JOB_DATA.db`
 An additional dbspace that contains job records data.
- `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`
 The file that enhances the EMM database performance.
- `SEARCH_DATA.db`
 Contains the metadata that is created while placing holds on backup images.
- `SEARCH_INDEX.db`
 The file that contains the indexes that are created in the `SEARCH_DATA` tables.
- `SLP_DATA.db`
 Contains the run-time information for storage lifecycle policies.
- `SLP_INDEX.db`
 The file that contains the indexes that are created in the `SLP_DATA` tables.
- `vxdbms.conf`
 The file that contains the configuration information specific to the Sybase SQL Anywhere installation:
 On Windows:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_AZ_DATABASE = NBAZDB
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
```

```
AZ_DB_PASSWORD = encrypted_password
VXDBMS_ODBC_DRIVER = NB SQL Anywhere
```

On UNIX:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
VXDBMS_AZ_DATABASE = NBAZDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_NB_DATA = /usr/opensv/db/data
VXDBMS_NB_INDEX = /usr/opensv/db/data
VXDBMS_NB_TLOG = /usr/opensv/db/data
VXDBMS_NB_STAGING = /usr/opensv/db/staging
VXDBMS_NB_PASSWORD = encrypted_password
AZ_NB_PASSWORD = encrypted_password
```

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 randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 940.

- If BMR is installed, the directory also contains: `BMRDB.db`, `BMRDB.log` (transaction log for BMR), `BMR_DATA.db`, `BMR_INDEX.db`

About the NetBackup configuration entry

The `VXDBMS_NB_DATA` registry entry (Windows) or the `bp.conf` entry (UNIX) 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.

On Windows:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\
Config\VXDBMS_NB_DATA
```

On UNIX: `/usr/opensv/netbackup/bp.conf`:

```
VXDBMS_NB_DATA = /usr/opensv/db/data
```

Sybase SQL Anywhere server management

Upon startup, the Sybase SQL Anywhere server uses the SQL Anywhere service to set the server parameters in the `server.conf` file. Then, the service starts the databases that are indicated in the `databases.conf` file.

To start and stop the Sybase SQL Anywhere service on Windows, use one of the following methods:

- In the **NetBackup Administration Console**, select **NetBackup Relational Database Manager** (SQLANYs_VERITAS_NB) in the **Daemons** tab of the **Activity Monitor**.
- From the Windows Service Manager, select **NetBackup Relational Database Manager** (SQLANYs_VERITAS_NB).
- `Install_path\VERITAS\NetBackup\bin\bpdown -e SQLANYs_VERITAS_NB`
- `Install_path\VERITAS\NetBackup\bin\bpup -e SQLANYs_VERITAS_NB`

To start and stop the Sybase SQL Anywhere daemon on UNIX, use one of the following methods:

- Select `NB_dbsrv` in the Activity Monitor in the **NetBackup Administration Console**.
- From the command line:
 - `/usr/opensv/netbackup/bin/bp.kill_all | bp.start_all`
 The Sybase SQL Anywhere daemon is included in the stop command or the start command, which starts and stops all NetBackup daemons.
 - `/usr/opensv/db/bin/nbdbms_start_server ndbms_start_server` starts the SQL Anywhere server if no option is specified.
 - `/usr/opensv/db/bin/nbdbms_start_server -stop -f`
 Stops the server; `-f` forces a shutdown with active connections.
 - `/usr/opensv/db/bin/nbdbms_start_server -stat`
 The `-stat` option tells whether the server is up or down:
 SQL Anywhere Server Ping Utility Version 11.0.1.2044 Ping server successful.
 - `/usr/opensv/db/bin/nbdbms_start_server -h`
 Use `-h` to display usage information about the `nbdbms_start_server`.

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`.

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 the following location:

On Windows: `Shared_drive\VERITAS\NetBackupDB\conf`.

On UNIX: `/usr/openv/var/global`.

Using the NetBackup Database Administration utility on Windows

The NetBackup Database Administration utility is a standalone 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, enter the DBA password. The password is set to a randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 940.

The NetBackup Database Administration utility displays the following information:

Table 23-3 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>vxdbms.conf</code> file. The <code>vxdbms.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 919.</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 925.</p>
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 23-1 shows the **General** tab of the Database Administration utility after a user logs on.

Figure 23-1 General tab of the NetBackup Database Administration utility




Table 23-4 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 23-5 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 23-6 describes the columns in the Index Details dialog box. <p>Click Close after defragmenting the files.</p>
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>

Table 23-4 General tab options (*continued*)

Option	Description
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 23-7 on page 924.</p>

Table 23-4 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 23-4 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>nbdbb_unload -rebuild</code></p>

Table 23-5 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 over 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 23-6 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 23-6 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 23-7 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 23-4 describes how to use the Validate option.</p>

See “[About fragmentation](#)” on page 925.

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.

On Windows, 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.

On Windows, 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 894.

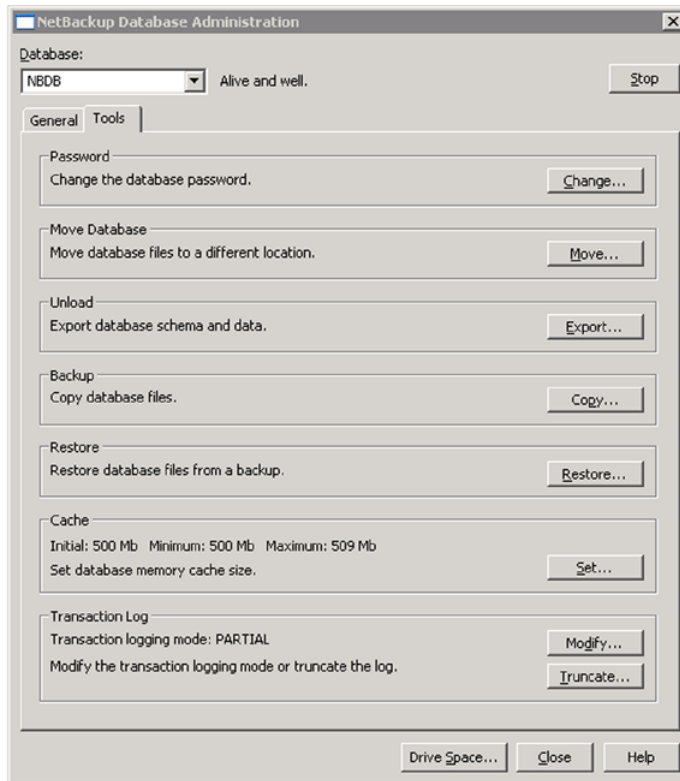
About the Tools tab of the NetBackup Database Administration utility

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 using the NetBackup Database Administration utility” on page 926.
Move Database section	See “Moving the NetBackup database files” on page 927.
Unload section	See “Exporting database schema and data” on page 929.
Backup section	See “Copying or backing up the database files” on page 929.
Restore section	See “Restoring database files from a backup” on page 930.

[Figure 23-2](#) displays the **Tools** tab of the utility.

Figure 23-2 Tools tab of the Database Administration utility



Changing the DBA password using the NetBackup Database Administration utility

To log into the Database Administration utility, you must know the current DBA password.

To change the password for the first time after installation, use the `nbdb_admin` command. The command updates the `vxdbms.conf` file with the new, encrypted string:

See [“Changing the NetBackup database password”](#) on page 940.

To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility.

To change the DBA password from a known password to a new 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 Windows”](#) on page 917.

- 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 Windows”](#) on page 917.

- 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 the critical files that are used to recover the relational databases. Browse to the new location.
Create mirrored transaction log at	<p>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.</p> <p>A mirrored transaction log offers extra protection.</p>
Move mirrored transaction log to	<p>This option is displayed if a mirrored transaction log exists.</p> <p>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.</p>
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 Windows”](#) on page 917.

- 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 927.
Before you rebuild the database.	See “About the General tab of the NetBackup Database Administration utility” on page 919.
Before you add data space.	

Before you add index space.

See [“About the General tab of the NetBackup Database Administration utility”](#) on page 919.

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 Windows”](#) on page 917.

- 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 the backup creates. A copy of all of the database files is made in this directory. This directory is also the location of the database files that the **Restore** option uses.

Note: This is not a catalog backup, performed as part of regular NetBackup operations.

See [“Restoring database files from a backup”](#) on page 930.

- 6 Click **OK**.

Restoring database files from a backup

Use the NetBackup Database Administration utility to restore a database from a backup copy.

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 Windows”](#) on page 917.

- 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 929.

- 6 Click **OK**.

Using the NetBackup Database Administration utility on UNIX

The NetBackup Database Administration utility is a standalone application (`dbadm`) and is installed in the following location:

/usr/opensv/db/bin

To use the NetBackup Database Administration utility, you must be an administrator with root user privileges.

When you start the NetBackup Database Administration utility, enter the DBA password. The password is set to a randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 940.

Enter `y` to change the password or `n` to proceed without changing the password.

After you log on, the NetBackup Database Administration utility displays the following information about the current database:

Table 23-8 NetBackup Database Administration utility properties

Property	Description
Selected database	The selected database: NBDB or BMRDB
Status	The status of the selected database: UP or DOWN
Consistency	<p>The validation state of the selected database: OK, NOT_OK, or DOWN</p> <p>This information can also be obtained running the following command:</p> <pre>nbdb_admin -validate</pre> <p>Note: To perform a full database validation, shut down NetBackup and start only the database service.</p>
Space Utilization	<p>Space Utilization is a measure of the amount of used space as a percentage of the file system space allocated for the database. When the NBDB or the BMRDB databases are initially created, extra space is allocated so that contiguous space is available as needed. As a result, a low Space Utilization is a positive indication unless the database is very large and disk space is in short supply.</p>

The initial screen also displays the following Database Administration main menu:

Table 23-9 Database Administration main menu options

Option	Description
Select/Restart Database and Change Password	<p>This option displays the menu where you can select a database to start or stop, and to change database passwords.</p> <p>See “About the Select/Restart Database and Change Password menu options” on page 933.</p>
Database Space and Memory Management	<p>This option displays the menu where you can perform the following actions:</p> <ul style="list-style-type: none"> ■ Generate a database space utilization report ■ Reorganize fragmented database objects ■ Add free space to the database files ■ Modify database server memory cache settings <p>See “About the Database Space and Memory Management menu options” on page 934.</p>
Transaction Log Management	<p>This option displays the menu where you can truncate the transaction log of the database. You can also change the database's transaction log mode.</p>
Database Validation Check and Rebuild	<p>This option displays the menu where you can validate and rebuild the selected database.</p> <p>See “About the Database Validation Check and Rebuild menu options” on page 936.</p>

Table 23-9 Database Administration main menu options (*continued*)

Option	Description
Move Database	This option displays the menu where you can change the location of the database files. You can also split the files into multiple directories. See “About the Move Database menu options” on page 937.
Unload Database	This option displays the menu where you can unload either the schema or the schema and data from the database. See “About the Unload Database menu options” on page 938.
Backup and Restore Database	This option displays the menu where you can choose the backup and restore options for the database. See “About the Backup and Restore Database menu options” on page 939.
Refresh Database Status	This option refreshes the Status, Consistency, and Space Utilization in the main menu.

Online Help is available from all screens by selecting the **h** option from the main menu.

About the Select/Restart Database and Change Password menu options

Select this option to display the Select/Restart Database and Change Password menu.

The Select/Restart Database and Change Password menu contains the following options.

Table 23-10 Select/Restart Database and Change Password options

Option	Description
NBDB	Select NBDB and then view or modify the database using the other <code>dbadm</code> menu options.
BMRDB	Select BMRDB and then view or modify the database using the other <code>dbadm</code> menu options.
Start Selected Database	Starts the selected database.
Stop Selected Database	Stops the selected database.

Table 23-10 Select/Restart Database and Change Password options (*continued*)

Option	Description
Change Password	<p>Changes the password for the databases. The password is changed for both NBDB and BMRDB, if applicable. Restart the database for the password change to take effect.</p> <p>To log into the Database Administration utility, you must know the current DBA password.</p>

About the Database Space and Memory Management menu options

You can use the Database Space and Memory Management option to perform the following functions:

- To report on database space utilization
- To reorganize fragmented database objects
- To add free space to the database files

The Database Space and Memory Management menu contains the following options.

Table 23-11 Database Space and Memory Management options

Option	Description
Report on Database Space	<p>Select this option to generate a report on space utilization for the currently selected database.</p> <p>The report contains the dbspaces and the physical pathnames of the relational database.</p> <p>For each dbspace, the report displays the name, the amount of free space in KBytes, the file size in KBytes, and the space utilization. The report also displays the transaction log file size and the amount of free space that remains on each of the file systems being used for the database.</p>

Table 23-11 Database Space and Memory Management options (*continued*)

Option	Description
Database Reorganize	<p>Select this option to reorganize fragmented database tables and indexes.</p> <p>These actions are performed from the Database Reorganize menu as follows:</p> <ul style="list-style-type: none"> <p>1) Defragment All</p> <p>This option automatically determines the database tables and indexes that are fragmented. It then uses the SQL Anywhere REORGANIZE command to defragment the tables and compress the indexes.</p> <p>The equivalent command is <code>nbdb_admin -reorganize</code>.</p> <p>2) Table Level Defragmentation</p> <p>This option generates a fragmentation report for each database table. For each table, the report includes the TABLE_NAME, number of ROWS, number of ROW_SEGMENTS, and SEGS_PER_ROW.</p> <p>In addition, a * displays in the ! column for an individual table if it would be automatically selected for reorganization by the Defragment All option.</p> <p>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. The ROW_SEGMENTS value indicates total number of row segments for the table. The SEGS_PER_ROW value shows the average number of segments per row, and indicates whether or not a table is fragmented.</p> <p>A SEGS_PER_ROW value of 1 is ideal, and any value more than 1 indicates a high degree of fragmentation. For example, a value of 1.5 means that half of the rows are partitioned.</p> <p>See “About fragmentation” on page 925.</p> <p>3) Index Compression</p> <p>This option generates a fragmentation report for each database index and lets the administrator select individual indexes to be compressed. For each index the report includes the TABLE_NAME, INDEX_NAME, TYPE of index, LEVEL, and DENSITY. The index TYPE is one of the following values: PKEY (primary key), FKEY (foreign key), UI (unique index), UC (unique constraint), NUI (non-unique index).</p> <p>In addition, a * displays in the ! column for an individual index if it would be automatically selected for reorganization by the Defragment All option.</p> <p>The LEVEL and DENSITY indicate whether or not an index should be reorganized. The number of levels in the index tree determines the number of I/O operations that are needed to access a row using the index. 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>A LEVEL value of 1 is ideal. An index with a LEVEL value of 4 or more or with a level value of 2 or 3 with a density greater than 0.5 is a good candidate for reorganization.</p>

Table 23-11 Database Space and Memory Management options (*continued*)

Option	Description
Add Free Space	<p>Select this option 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>A Rebuild eliminates all free space and then adds back what was initially added when the database was created. If you accidentally add too much free space, a Rebuild can be used to eliminate the space.</p>

About the Database Validation Check and Rebuild menu options

The Database Validation Check and Rebuild option lets you validate and rebuild the currently selected database.

The Database Validation Check and Rebuild menu contains the following options.

Table 23-12 Database Validation Check and Rebuild menu options

Option	Description
Standard Validation	<p>This option lets you validate the indexes and keys on all of the tables in the database. Each table is scanned. 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.</p> <p>This option is equivalent to the command <code>nbdb_admin -validate</code>.</p>
Full Validation	<p>This option performs the same checks as the Standard Validation option. In addition, 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. This option is equivalent to the command <code>nbdb_admin -validate -full</code>.</p> <p>Note: To perform a full database validation, shut down NetBackup and start only the database service.</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>NB_dbsrv</code>, the NetBackup Relational Database Manager). ■ Repeat the validation check using this tool or the <code>nbdb_admin</code> command line utility. <p>If validation errors persist, contact Symantec customer support. The administrator may be asked to rebuild the database using the Database Rebuild option or the <code>nbdb_unload</code> command line utility.</p>

Table 23-12 Database Validation Check and Rebuild menu options (*continued*)

Option	Description
Database Rebuild	<p>This option lets you rebuild the database. This menu option is equivalent to the command <code>nbdb_unload -rebuild</code>. A Database Rebuild results in a complete unload and reload of the database. A new database with all of the same options is built in place. A Database Rebuild may be required if Database Validation errors are reported using the Standard or Full Validation options.</p> <p>During a Database Rebuild, all NetBackup operations are suspended.</p> <p>When you select this option, a message appears which recommends that you exit and create a backup using the Backup Database option before you rebuild the database. You then have the choice of whether to continue or not.</p> <p>See “About the Backup and Restore Database menu options” on page 939.</p>

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.

About the Move Database menu options

The Move Database menu options let you change the location of the database files or split the database files into multiple directories. Changing the location may improve performance in the case of large databases. The database files are moved for both NBDB and BMRDB, if present.

Note: Due to performance issues, NetBackup supports database files only on locally attached drives.

In addition, the Move Database menu lets you create a mirrored transaction log. The transaction logs (NBDB.log and BMRDB.log) are critical files used to recover the relational databases.

For extra protection, a mirrored transaction log can be used. Create the mirrored log in a different directory from the original log.

The Move Database menu contains the following options.

Table 23-13 Move Database menu options

Option	Description
Move Data	Select this option to change the location of the data dbspaces. Changing the location temporarily suspends NetBackup operations.
Move Index	Select this option to change the location of the index dbspaces. Changing the location temporarily suspends NetBackup operations.

Table 23-13 Move Database menu options (*continued*)

Option	Description
Move Transaction Log	Select this option to change the location of the transaction log. Changing the location temporarily suspends NetBackup operations.
Move/Create Mirrored Log	Select this option to change the location of the mirrored transaction log. If the mirrored transaction log does not already exist, it is created in the new location. Changing the location temporarily suspends NetBackup operations.
Stop Using Mirrored Transaction Log	This option is displayed only if a mirrored transaction log is used. Select this option to configure the database server so that it stops using a mirrored transaction log. The existing mirrored transaction log file is deleted. This action temporarily suspends NetBackup operations.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.

About the Unload Database menu options

The Unload Database menu options let you unload either the schema or the schema and data from the relational database.

A SQL command file is created that can be used to rebuild the database. If the data is also included in the unload, a set of data files in comma-delimited format is created.

The Unload Database menu contains the following options.

Table 23-14 Unload Database menu options

Option	Description
Schema Only	This option lets you unload only the database schema. The schema is unloaded as a file that is named <code>reload.sql</code> in the named directory.
Data and Schema	This option lets you unload both the database schema and the data. The data is unloaded as a set of files. One file is created for each database table.
Change Directory option for unloading the database	This option lets you change the directory location for the files that unload options (1) or (2) create.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.

About the Backup and Restore Database menu options

The Backup and Restore Database menu options let you back up the relational database to the specified directory. You can restore from a previously created backup.

Symantec recommends creating a backup copy of the database files in the following situations:

- Before you move the database.
- Before you rebuild the database.
- Before you add data space.
- Before you add index space.
- Before you modify the transaction logging mode from FULL to PARTIAL.
- Before you truncate the transaction log.

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 database only as a precautionary measure.

The Backup and Restore Database menu contains the following options.

Table 23-15 Backup and Restore Database menu options

Option	Description
Online Backup	This option lets you make a copy of the database files while the database is active. Other NetBackup activity is not suspended during this time.
Restore Backup	This option lets you restore from a copy of the database files previously made with either options 1 or 2. The currently running database files are overwritten, and the database is shut down and restarted after the restore is completed.
Change Directory	This option lets you change the directory location for the files that backup options (1) or (2) create. This directory is the source of the files for the restore option (3).

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.

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 NetBackup database password”](#) on page 940.
- Move NBDB and BMRDB database files (possibly to tune performance).
See [“Moving NBDB database files after installation”](#) on page 941.
- Add a mirrored transaction log.
See [“Adding a mirrored transaction log”](#) on page 943.
- Recreate NBDB.
See [“Creating the NBDB database manually”](#) on page 945.

Changing the NetBackup database password

The database password is set to a randomly generated password upon installation. This password is used for NBDB and BMRDB and for all DBA and application accounts. (For example, `EMM_MAIN`.) You can use this procedure to change it to a known password.

The password is encrypted and stored in the `vxdbs.conf` file. The permissions for the `vxdbs.conf` file allow only a Windows administrator or a `root` user to read or write to it.

To change the database password

- 1 Log on to the server as a Windows Administrator or as `root`.
- 2 To change the password for the first time after installation, run the following command. The command updates the `vxdbms.conf` file with the new, encrypted string:

On Windows: `Install_path\NetBackup\bin\nbdb_admin -dba new_password`

On UNIX: `/usr/opensv/db/bin/nbdb_admin -dba new_password`

The password needs to be an ASCII string. Non-ASCII characters are not allowed in the password string.

- 3 To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility. You must know the current password to log into the NetBackup Database Administration utility.

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 917.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.

Requirements to create a new password where NBAC is enabled

If the NetBackup environment has NetBackup Access Control (NBAC) enabled, keep in mind the following items before you run `nbdb_admin -dba new_password`.

- Only the minimum services should be running: `NB_dbsrv`, `nbatd`, `nbazd`, `bprd`, `bpdbm`, `bpjobjd`.
- Users without `root` or administrator permissions need to have read and write permissions to `vxdbms.conf`. Windows users can use the security **Permissions** tab in NBAC to allow a non-administrator user read and write permissions to `vxdbms.conf`.
- To generate `nbdb_admin` logs, make sure that users have read and write permissions to the following directory:
On Windows: `install_path\Veritas\NetBackup\logs\nbdb`
On UNIX: `/usr/opensv/netbackup/logs/nbdb`
- To troubleshoot a situation where only the database server can be started, NBAC must be turned off.

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.

Moving the database files on Windows

To move the NBDB and the BMRDB database files on Windows

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup services by typing the following command:

```
Install_path\VERITAS\NetBackup\bin\bpdwn
```

- 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 927.

- 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.

Moving the database files on UNIX

To move the NBDB and the BMRDB database files on UNIX

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- 3 Start the SQL Anywhere daemon by typing the following command:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

- 4 Use one of the following methods to move the existing data, index, and transaction log files:

- Use the NetBackup Database Administration utility.
See [“About the Move Database menu options”](#) on page 937.

- Type the following command:

```
/usr/opensv/db/bin/nbdb_move  
-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:

```
/usr/opensv/db/bin/nbdb_move -data data_directory  
-index index_directory -tlog log_directory  
-mlog log_mirror_directory
```

- 5 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 6 Perform a catalog backup.

See [“About NetBackup master server installed directories and files”](#) on page 906.

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.

Creating a mirrored transaction log on a Windows computer

To create a mirrored transaction log on a Windows computer

- 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 Create the mirrored transaction log:

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.

Creating a mirrored transaction log on a UNIX computer

To create a mirrored transaction log on a UNIX computer

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- 3 Start the SQL Anywhere daemon by typing the following command:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

- 4 Create the mirrored transaction log:

Type the following command:

```
/usr/opensv/db/bin/nbdb_move -mlog log_mirror_directory
```

To move the existing data, index, transaction log files, and create the mirrored transaction log, type the following command:

```
/usr/opensv/db/bin/nbdb_move -data data_directory-index  
index_directory -tlog log_directory-mlog log_mirror_directory
```

- 5 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 6 Perform a catalog backup.

See [“About catalog backups”](#) on page 875.

See [“Moving NBDB database files after installation”](#) on page 941.

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.

Creating the database on Windows

To create the NBDB database manually on Windows

- 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 2 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.

Creating the database on UNIX

To create the `NBDB` database manually on UNIX

- 1 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- 2 Start the SQL Anywhere daemon by typing the following command:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

- 3 Run the following command:

```
/usr/opensv/db/bin/create_nbdb
```

- 4 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 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.

```
/usr/opensv/volmgr/bin/tpext
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads `EMM` data into the database.

Additional create_nbdb options

In addition to using the `create_nbdb` command to create the NBDB database, you also can use it to perform the following actions. In each command, *NB_server_name* matches the name in `server.conf`.

See [“About the NetBackup server.conf file”](#) on page 907.

- Drop the existing NBDB database and recreate it in the default location by typing the following command:

- On Windows: `create_nbdb -drop [current_data_directory]`

Provide the location of the current NBDB data directory, `current_data_directory`, if the default location is not used.

- On UNIX: `create_nbdb -drop`

The location of the current NBDB data directory is retrieved automatically from the `bp.conf` file.

The `-drop` option instructs NetBackup to drop the existing NBDB database.

- Drop the existing NBDB database and do not recreate by typing the following command:

On Windows: `create_nbdb -db_server NB_server_name
-drop_only [current_data_directory]`

On UNIX: `create_nbdb -drop_only`

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:

- On Windows:

```
create_nbdb -drop [current_data_directory] -data  
data_directory-index index_directory -tlog log_directory  
[-mloglog_mirror_directory]
```

- On UNIX:

```
create_nbdb -drop -data data_directory -index index_directory  
-tlog log_directory [-mloglog_mirror_directory]
```

If the NBDB database files were moved from the default location by using `nbdb_move`, use this command to recreate them in the same location. Specify `current_data_directory`.

If the location of `NBDB.db` changed from the default, `BMRDB.db` must also be recreated. The `BMRDB.db` files must reside in the same location as the NetBackup database files.

About backup and recovery procedures

The 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 Configuration 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 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:

On Windows: `Install_path\NetBackupDB\data\NBDB.log`

On UNIX: `/usr/opensv/db/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:

On Windows: `Install_path\NetBackupDB\data\BMRDB.log`

On UNIX: `/usr/opensv/db/data/BMRDB.log`

It has an optional mirrored log in the following location:

On Windows: `mirrored_log_directory\BMRDB.m.log`

On UNIX: `mirrored_log_directory/BMRDB.m.log`

The BMRDB transaction logs are backed up and truncated during the catalog backup along with the NBDB transaction logs.

Note: If a catalog backup is not run, the logs are not truncated. Truncation must be managed in this manner as it is critical to recovery of the database.

About catalog recovery

Recovery scenarios include the following:

- A full recovery from a complete disaster
Using the **Disaster Recovery** wizard, the databases are restored along with the image catalog to a consistent state.
- A recovery of the database files only
Using `bprecover`, the relational database files and configuration files can be restored and recovered.

Details about catalog recovery scenarios and procedures are available in the [NetBackup Troubleshooting Guide](#).

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 as moving or reorganizing the database files.

Table 23-16 Commands used to back up and recover relational databases

Command	Description
nbdb_backup or nbdb_backup.exe	<p>Use <code>nbdb_backup</code> or <code>nbdb_backup.exe</code> to make a copy of the NBDB database files and the BMRDB database files in a directory. The <code>nbdb_backup</code> command does not truncate the transaction log. Transaction logs are managed only by using the catalog backup.</p> <p>On Windows: <code>Install_path\NetBackup\bin\nbdb_backup.exe [-dbn database_name] -online destination_directory [tlog_only] [-truncate_tlog]</code></p> <p>On UNIX: <code>/usr/opensv/db/bin/nbdb_backup [-dbn database_name] -online destination_directory [tlog_only] [-truncate_tlog]</code></p> <p><code>-dbn database_name</code> only backs up the specified database (NBDB or BMRDB).</p> <p><code>-online destination_directory</code> is the location for the copied files.</p> <p><code>-tlog_only</code> generates only the database transaction log along with the configuration files for the backup.</p> <p><code>-truncate_tlog</code> truncates the database transaction log after the backup completes the backup.</p> <p>Note: Catalog backup activity is usually logged in the <code>nbdb</code> log. If the master server uses NetBackup Access Control (NBAC) in the REQUIRED mode and the EMM database is remote, the logging information appears in the <code>bpdbm</code> log.</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>
nbdb_restore or nbdb_restore.exe	<p>Use <code>nbdb_restore</code> to recover from a database backup that was made using <code>nbdb_backup</code>.</p> <p>On Windows: <code>Install_path\NetBackup\bin\nbdb_restore.exe -recover source_directory -dbn database_name [-ignore_tlog]</code></p> <p>On UNIX: <code>/usr/opensv/db/bin/nbdb_restore -recover source_directory -dbn database_name [-ignore_tlog]</code></p> <p>Logs are recorded in the <code>/admin</code> or <code>\admin</code> directory.</p> <p><code>nbdb_restore</code> restores from the staging directory to recover only the NetBackup database (NBDB). The default is to apply the current transaction log to make the database as consistent as possible with the current <code>db\images</code> or <code>/db/images</code> directory.</p> <p><code>-ignore_tlog</code> does not apply the current transaction log to the NBDB recovery operation.</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>

About the catalog backup process

Normally, a 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 catalog backup process consists of the following process:

- Creates a staging directory where the command can store the temporary copies:
 - On Windows: `Install_path\NetBackupDB\staging`
 - On UNIX: `/usr/opensv/db/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
 - `BMR_DATA.db`
 - `BMRDB.db`
 - `BMRDB.log`
 - `BMR_INDEX.db`
 - `DARS_DATA.db`
 - `DARS_INDEX.db`
 - `DBM_DATA.db`
 - `DBM_INDEX.db`
 - `EMM_DATA.db`
 - `EMM_INDEX.db`
 - `JOBD_DATA.db`
 - `NBAZDB.db`
 - `NBAZDB.db.template`
 - `NBAZDB.log`
 - `NBDB.db`
 - `NBDB.log`
 - `SEARCH_DATA.db`
 - `SEARCH_INDEX.db`
 - `vxdbms.conf`
 - If BMR was installed
- A second child job begins the image catalog backup.
- 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 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.

On Windows:

```
Install_path\NetBackup\bin\nbdb_unload.exe [-dbn database_name] [-t table_list] [-s] destination_directory
```

On UNIX:

```
/usr/opensv/db/bin/nbdb_unload [-dbn database_name] [-t table_list] [-s] destination_directory
```

In the script where:

- `-dbn database_name`
`database_name` is NBDB (default) or BMRDB.
- `-t table_list`
Must list the owner of the table, then the table name. For EMM, the account `EMM_MAIN` owns all tables.
`nbdb_unload -t EMM_MAIN.EMM_Device, EMM_MAIN.EMM_Density`
- `-s`
No data is dumped, only schema.

- *destination_directory*
Specify the location where the dump is created.

Terminating database connections

Before you run `nbdbb_unload`, shut down NetBackup to terminate all active connections to the database. Shutting down NetBackup eliminates any possible concurrency problems.

Terminating database connections on Windows

To terminate database connections on Windows

- 1 Shutdown all NetBackup services by typing the following command:

`Install_path\VERITAS\NetBackup\bin\bpdown`
- 2 In the Windows **Services Manager**, restart the service titled **Adaptive Server Anywhere – VERITAS_NB**. On NetBackup 7.x systems, the service is titled **NetBackup Relational Database Manager**.
- 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 Windows”](#) on page 917.
 - Run `nbdbb_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`

Terminating database connections on UNIX

To terminate database connections on UNIX

- 1 Shutdown all NetBackup daemons by typing the following command:

`/usr/opensv/netbackup/bin/bp.kill_all`
- 2 Start the SQL Anywhere daemon by typing the following command:

`/usr/opensv/netbackup/bin/nbdbms_start_stop start`
- 3 Start only the database server by using

`/usr/opensv/netbackup/bin/nbdbms_start_stop start`

- 4 Use one of the following methods to terminate database connections:
 - Use the NetBackup Database Administration utility.
See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 931.
 - Run `nbdbb_unload` and indicate the outputs (database name, table lists, or schema only) and the destination directory.
- 5 Shut down the database server by using `/usr/opensv/netbackup/bin/nbdbbms_start_stop stop`.
- 6 Stop the SQL Anywhere daemon by typing the following command:
`/usr/opensv/netbackup/bin/nbdbbms_start_stop stop`
- 7 Start all NetBackup daemons by typing the following command:
`/usr/opensv/netbackup/bin/bp.start_all`

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 `nbdbb_backup` to make a physical copy or use `nbdbb_move` to relocate the database files.

Moving the NetBackup database from one host to another

The NetBackup database (NBDB) includes the Enterprise Media Manager (EMM). Both reside together on the master server. In previous releases, NetBackup allowed NBDB and EMM to be moved from the master server to another server.

Support for remote-EMM and shared-EMM server configurations was withdrawn in 7.6. In a remote- or shared-EMM server configuration, the NetBackup relational database (NBDB), the Enterprise Media Manager (EMM), and the Resource Broker (RB) have been moved to a server that is not the master server. Customers who currently use this option should contact Symantec Support. Support engages Symantec Engineering to review the options to disengage this configuration.

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 24-1 Catalog utility options

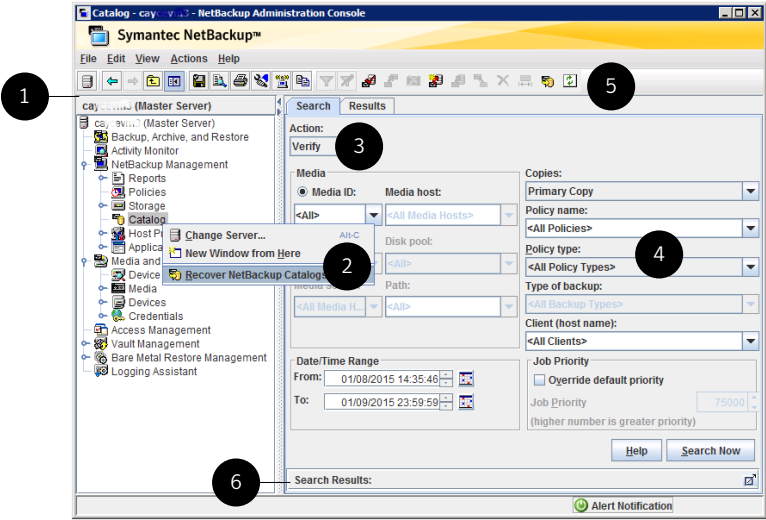


Table 24-1 Host Properties utility

Number	Description
1	The name of the currently selected master server.
2	Right-click Catalog to view the shortcut menu.
3	Select an action to perform in the Catalog utility.
4	Set search criteria, including a specific media and date range.
5	The user toolbar is specific to the Catalog utility. See “Standard and user toolbars” on page 45.
6	Display of search results.

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
The backup is not imported if it begins on a media ID that the initiating backup procedure did not process.
The backup is incomplete if it ends on a media ID that the initiating backup procedure did not process.
See [“About importing backup images”](#) on page 970.

[Table 24-2](#) lists the search criteria for backup images.

Table 24-2 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 959. See “Duplicating backup images” on page 962. See “Expiring backup images” on page 969.
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 Host	On UNIX: Specifies the host name of the media server that produced the originals. Type a host name in the box or select one from the scroll-down list. To search through all hosts, select All Media Hosts .
Media Server	On Windows: 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.

Table 24-2 Catalog utility search properties (*continued*)

Property	Description
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.
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 .
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 113.</p>
Search tab columns	The Search tab columns list information about matching backup images based on the defined search criteria.
Results tab columns	The Results tab displays the progress of the request to verify, duplicate, expire, or import an image.
Messages pane	On UNIX: Displays the messages about the tasks that run as background processes. The pane appears only if there is an informative message or error message for a task. If the task completes normally, the pane does not appear. The Messages pane can be maximized, minimized, or closed.

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, select the image you want to verify and on the **Action** drop-down menu, select **Action > Verify**.

To display information on each file that NetBackup verifies, select **Log all files found in verified image(s)**.

- 4 In the right pane, click the **Results** tab, then select the verification job to view the job results.

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.

Use one of the following methods to promote a copy to a primary copy:

Promote a backup copy to a primary copy using search criteria

[Promoting a backup copy to a primary copy using the NetBackup Administration Console](#)


Promote a copy to a primary copy for many backups using the `bpchangeprimary` command

[Promoting a copy to a primary copy for many backups](#)

Promote a backup copy to a primary copy using the `bpduplicate` command

[Promoting a backup copy to a primary copy using `bpduplicate`](#)

Figure 24-2 Primary copy status



Primary Copy status indicates that the image is now the primary copy

Promoting a backup copy to a primary copy using the NetBackup Administration Console

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 956.

- 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**.

Promoting a copy to a primary copy for many backups

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, 2012:

```
bpchangeprimary -pool SUN -sd 08/01/2012
```

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, 2012:

```
bpchangeprimary -copy 2 -cl client_a -sd 01/01/2012
```

More information is available in the [NetBackup Commands Reference Guide](#).

Promoting a backup copy to a primary copy using bpduplicate

To use bpduplicate to promote a backup copy to a primary copy

- 1 Enter the following command:

On Windows:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bpduplicate  
-npc pcopy -backupid bid
```

On UNIX:

```
/usr/openv/netbackup/bin/admincmd/bpduplicate -npc pcopy  
-backupid bid
```

Where:

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 24-3](#) lists the scenarios in which duplication is possible and scenarios in which duplication is not possible:

Table 24-3 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 654.

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.

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 multiple copies”](#) on page 760.

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.

If the Policy type is 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.
If the Policy type is set to NBU-Catalog .	Specifies that only catalog volume pools are included in the drop-down list.
If the Policy type is 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 `bpxpdate` 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 duplicate 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, select **Preserve Multiplexing** to ensure 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 959.

See [“About multiplexed duplication considerations”](#) on page 967.

About multiplexed duplication considerations

Consider the following items about multiplexed duplication.

Table 24-4 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	<p>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:</p> <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.

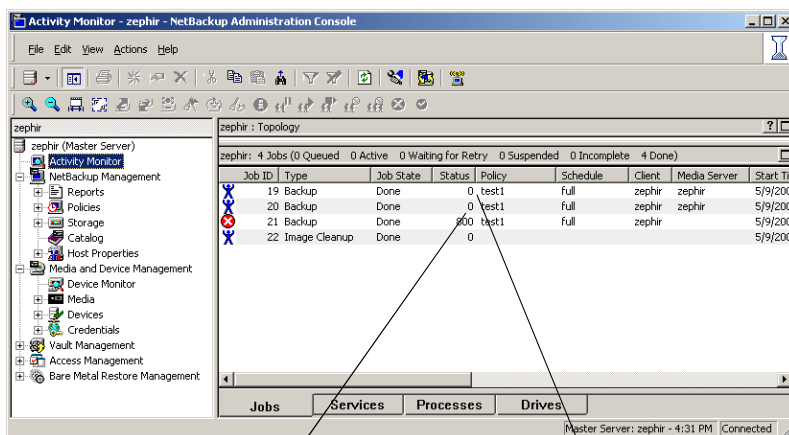
Jobs that appear while making multiple copies

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 lets you 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 an 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 24-3 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, or information about the backup is deleted. 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 956.
- 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.

If the user attempts to manually expire an image or some image copies that are not SLP-complete, the request fails with status code 1573 (Backup image cannot be expired because its SLP processing is not yet complete). Use the NetBackup **Troubleshooter** or the [NetBackup Status Codes Reference Guide](#) to determine corrective actions.

If you use the **NetBackup Administration Console** to expire NetBackup images, the image extension files remain in the catalog. However, NetBackup expires the image extension files, usually within 12 hours.

About importing backup images

NetBackup can import the backups that have expired or the backups from another NetBackup server.

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.

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, 2014, and its retention period is one week, the new expiration date is November 21, 2014.

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 a 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:

On Windows:

```
Install_path \VERITAS\NetBackup\bin\admincmd\bpexpdate  
-d 0 -m mediaID
```

On UNIX:

```
/usr/opensv/NetBackup/bin/admincmd/bpexpdate -d 0 -m  
media_ID
```

- 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 969.

Importing backup images

Use the following procedure to initiate an import of backup images.

To initiate an import of backup images

- 1 To import the images from tape, make the media accessible to the media server so the images can be imported.
- 2 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Catalog**.
- 3 On the **Actions** menu, select **Actions > Initiate Import**.
- 4 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.
- 5 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	In the Media ID field, type the Media ID of the volume that contains the backups to import.
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>For a BasicDisk type, enter or browse to the path to the images in the field provided.</p>

Click **OK**.

- 6 Click **OK** to begin reading the catalog information from the source volume.
- 7 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.

Configuring replication

- [Chapter 25. About NetBackup replication](#)

About NetBackup replication

This chapter includes the following topics:

- [About NetBackup replication](#)
- [About NetBackup Auto Image Replication](#)
- [About NetBackup Replication Director](#)

About NetBackup replication

NetBackup offers two forms of replication:

Backups	<p>Auto Image Replication</p> <p>Use this type of replication to replicate backups from one NetBackup domain to the NetBackup media server in another domain.</p> <p>See “About NetBackup Auto Image Replication” on page 975.</p>
Snapshots	<p>NetBackup Replication Director</p> <p>This type of replication makes use of NetBackup OpenStorage to replicate snapshots on primary storage to the disk arrays of OpenStorage partners.</p> <p>See “About NetBackup Replication Director” on page 1014.</p> <p>For more information, see the NetBackup Replication Director Solutions Guide.</p>

About NetBackup Auto Image Replication

The backups that are generated in one NetBackup domain can be replicated to storage in one or more target NetBackup domains. This process is referred to as Auto Image Replication.

The ability to replicate backups to storage in other NetBackup domains, often across various geographical sites, helps facilitate the following disaster recovery needs:

- **One-to-one model**
A single production datacenter can back up to a disaster recovery site.
- **One-to-many model**
A single production datacenter can back up to multiple disaster recovery sites. See “[One-to-many Auto Image Replication model](#)” on page 977.
- **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.

NetBackup supports the following storage types for Auto Image Replication:

Table 25-1 NetBackup Auto Image Replication storage types

Storage type	Link to more information
Media Server Deduplication Pool	See the <i>NetBackup Deduplication Guide</i> : http://www.symantec.com/docs/DOC5332
An OpenStorage disk appliance	If your storage vendor's product supports replication, you can automatically replicate backup images to a similar device in a different master server domain. See the <i>NetBackup OpenStorage Solutions Guide for Disk</i> : http://www.symantec.com/docs/DOC5332

Notes about Auto Image Replication

- Auto Image Replication is supported between the master server domains that run any of the following versions of NetBackup:
 - NetBackup 7.1.0.4
 - NetBackup 7.5.0.3 and later

For catalog backup images, NetBackup supports Auto Image Replication only between the same release levels of NetBackup.

Both master servers must be at NetBackup version 7.6 or later if Targeted A.I.R. is used. Targeted A.I.R. means that Auto Image Replication is configured to replicate a backup to a specific master server in a specific domain.

- Although Auto Image Replication is a disaster recovery solution, you cannot restore to clients in the primary (or originating) domain from the target master domain.
- Auto Image Replication does not support synthetic backups.
- Auto Image Replication does not support spanning volumes in a disk pool. NetBackup fails backup jobs to the disk pools that span volumes if the backup job is in a storage lifecycle policy that also contains a replication operation.
- Auto Image Replication does not support replicating from a storage unit group. That is, the source copy cannot be in a storage unit group.
- Synchronize the clocks of the master servers in the source and the target domains so that the master server in the target domain can import the images as soon as they are ready. The master server in the target domain cannot import an image until the image creation time is reached. Time zone differences are not a factor because the images use Coordinated Universal Time (UTC).

Process Overview

[Table 25-2](#) is an overview of the process, generally describing the events in the originating and target domains.

NetBackup uses storage lifecycle policies in the source domain and the target domain to manage the Auto Image Replication operations.

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 988.

Table 25-2 Auto Image Replication process overview

Event	Domain in which event occurs	Event description
1	Originating master server (Domain 1)	Clients are backed up according to a backup policy that indicates a storage lifecycle policy as the Policy storage selection. The SLP must include at least one Replication operation to similar storage in the target domain.
2	Target master server (Domain 2)	The storage server in the target domain recognizes that a replication event has occurred. It notifies the NetBackup master server in the target domain.

Table 25-2 Auto Image Replication process overview (*continued*)

Event	Domain in which event occurs	Event description
3	Target master server (Domain 2)	NetBackup imports the image immediately, based on an SLP that contains an import operation. NetBackup can import the image quickly because the metadata is replicated as part of the image. (This import process is not the same as the import process available in the Catalog utility.)
4	Target master server (Domain 2)	After the image is imported into the target domain, NetBackup continues to manage the copies in that domain. Depending on the configuration, the media server in Domain 2 can replicate the images to a media server in Domain 3.

One-to-many Auto Image Replication model

In this configuration, all copies are made in parallel. The copies are made within the context of one NetBackup job and simultaneously within the originating storage server context. If one target storage server fails, the entire job fails and is retried later.

All copies have the same **Target Retention**. To achieve different **Target Retention** settings in each target master server domain, either create multiple source copies or cascade duplication to target master servers.

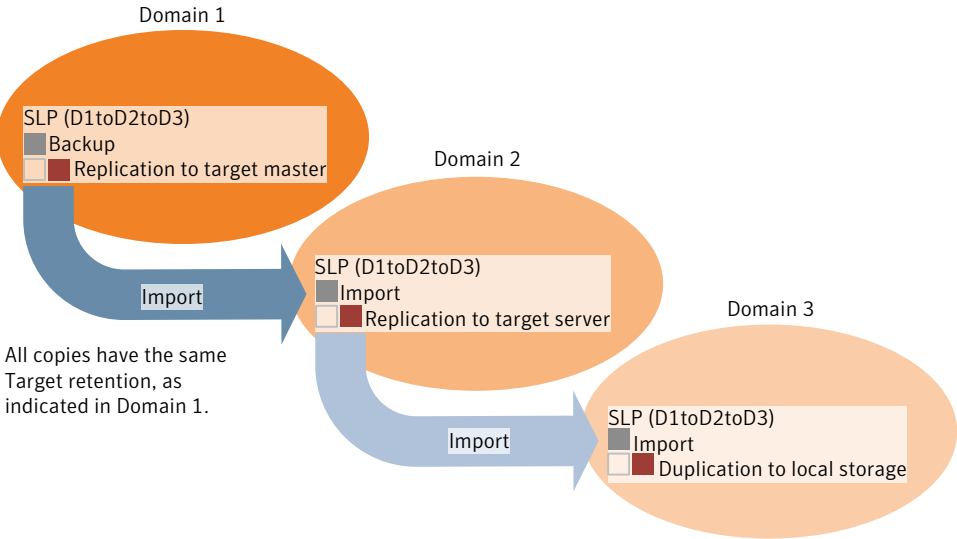
Cascading Auto Image Replication model

Replications can be cascaded from the originating domain to multiple domains. Storage lifecycle policies are set up in each domain to anticipate the originating image, import it and then replicate it to the next target master.

[Figure 25-1](#) represents the following cascading configuration across three domains.

- The image is created in Domain 1, and then replicated to the target Domain 2.
- The image is imported in Domain 2, and then replicated to a target Domain 3.
- The image is then imported into Domain 3.

Figure 25-1 Cascading Auto Image Replication



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 indicates that the master server in Domain 2 is the originating master server. However, after the image is imported successfully into Domain 3, NetBackup correctly indicates that the originating master server is in Domain 1.

The cascading model presents a special case for the Import SLP that replicates the imported copy to a target master. (This master server that is neither the first nor the last in the string of target master servers.)

The Import SLP must include at least one operation that uses a **Fixed** retention type and at least one operation that uses a **Target Retention** type. So that the Import SLP can satisfy these requirements, the import operation must use a **Target Retention**.

Table 25-3 shows the difference in the import operation setup.

Table 25-3 Import operation difference in an SLP configured to replicate the imported copy


Import operation criteria	Import operation in a cascading model
The first operation must be an import operation.	Same; no difference.
A replication to target master must use a Fixed retention type	Same; no difference.
At least one operation must use the Target retention .	Here is the difference: To meet the criteria, the import operation must use Target retention .

The target retention is embedded in the source image.

In the cascading model that is represented in [Figure 25-1](#), all copies have the same **Target Retention**—the **Target Retention** indicated in Domain 1.

For the copy in Domain 3 to have a different target retention, add an intermediary replication operation to the Domain 2 storage lifecycle policy. The intermediary replication operation acts as the source for the replication to target master. Since the target retention is embedded in the source image, the copy in Domain 3 honors the retention level that is set for the intermediary replication operation.

Figure 25-2 Cascading replications to target master servers, with various target retentions



About the domain relationship for replication

The following table describes important configuration differences depending on the devices that are used for NetBackup Auto Image Replication.

Table 25-4 Replication relationships

Storage	Domain Relationship
NetBackup managed storage	<p>For NetBackup managed storage, establish the relationship between the originating domain and the target domain or domains by setting the properties in the source storage server. Specifically, in the Replication tab of the Change Storage Server dialog box of the source storage server.</p> <p>NetBackup managed storage that qualifies for replication is Media Server Deduplication Pool storage.</p> <p>See the <i>NetBackup Deduplication Guide</i>: http://www.symantec.com/docs/DOC5332</p> <p>Before you configure the replication relationship, you can add the target master server as a trusted host.</p> <p>See “About trusted master servers for Auto Image Replication” on page 987.</p>
Third-party vendor disk appliance	<p>For the third-party vendor appliance storage that is exposed through the OpenStorage API, the disk appliance manages the storage. The relationship between the originating domain and the target domain or domains is established by using the storage vendor's tools.</p> <p>The originating NetBackup domain has no knowledge of the storage server in the target domain or domains. When the appliances are configured properly, NetBackup images on the originating 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.</p> <p>NetBackup manages the lifecycle of the backup images but does not manage the storage.</p> <p>Configuring the disk appliance replication relationship is beyond the scope of the NetBackup documentation.</p>

Caution: Choose the target storage server carefully. A target storage server must not also be a storage server for the originating domain.

About the replication topology for Auto Image Replication

For Auto Image Replication, the disk volumes have the properties that define the replication relationships between the volumes. The knowledge of the volume properties is considered the replication topology. The following are the replication properties that a volume can have:

Source	A source volume contains the backups of your clients. The volume is the source for the images that are replicated to a remote NetBackup domain. Each source volume in an originating domain has one or more replication partner target volumes in a target domain.
Target	A target volume in the remote domain is the replication partner of a source volume in the originating domain.
None	The volume does not have a replication attribute.

The following describes the replication topology for the supported storage types:

Table 25-5 Replication topology by storage type

Storage type	Replication topology
Media Server Deduplication Pool	<p>NetBackup exposes the storage for a Media Server Deduplication Pool as a single volume. Therefore, there is always a one-to-one volume relationship for MSDP.</p> <p>You configure the replication relationships in the source domain. To do so, you add target storage servers in the Replication tab of the Change Storage Server dialog box of the source storage server.</p> <p>See the <i>NetBackup Deduplication Guide</i>: http://www.symantec.com/docs/DOC5332</p>
Disk storage devices that support replication and also conform to the Symantec OpenStorage API	<p>Your storage administrator configures the replication topology of the volumes in the storage devices. Based on the volume properties, you create homogeneous disk pools. That is, all of the volumes in a disk pool must have the same properties, and you create the disk pools to match that topology. The disk pools inherit the replication properties from the volumes that you add to them.</p> <p>You should work with your storage administrator to understand the topology so you can create the proper disk pools. You also should work with your storage administrator to understand any changes that are made to the replication topology.</p> <p>NetBackup discovers the topology of the volumes when you configure a disk pool.</p>

NetBackup discovers the replication topology when you configure the replication relationships. NetBackup discovers topology changes when you use the **Refresh** option of the **Change Disk Pool** dialog box.

NetBackup includes a command that can help you understand your replication topology. Use the command in the following situations:

- After you configure the replication targets.
- After you configure the storage server and before you configure disk pools.
- After changes to the volumes that comprise the storage.

See [“Viewing the replication topology for Auto Image Replication”](#) on page 982.

Viewing the replication topology for Auto Image Replication

A volume that is a source of replication must have at least one replication partner that is the target of the replication. NetBackup lets you view the replication topology of the storage.

See [“About the replication topology for Auto Image Replication”](#) on page 980.

To view the replication topology for Auto Image Replication

- ◆ Run the `bpstsinfo` command, specifying the storage server name and the server type. The following is the command syntax:
 - **Windows:** `install_path\NetBackup\bin\admincmd\bpstsinfo -lsuinfo -storage_server host_name -stype server_type`
 - **UNIX:** `/usr/opensv/netbackup/bin/admincmd/bpstsinfo -lsuinfo -storage_server host_name -stype server_type`

The following are the options and arguments for the command:

- | | |
|---|--|
| <code>-storage_server <i>host_name</i></code> | The name of the target storage server. |
| <code>-stype <i>server_type</i></code> | <ul style="list-style-type: none"> ■ For a Media Server Deduplication Pool target, use <code>PureDisk</code>. ■ For an OpenStorage disk appliance, the vendor provides the string for <code>server_type</code>. |

Save the output to a file so that you can compare the current topology with the previous topology to determine what has changed.

See [“Sample volume properties output for MSDP replication”](#) on page 983.

See [“Sample volume properties output for OpenStorage backup replication”](#) on page 984.

Sample volume properties output for MSDP replication

The following two examples show output from the `bpstsinfo -lsuinfo` command for two NetBackup deduplication storage servers. The first example is the output from the source disk pool in the originating domain. The second example is from the target disk pool in the remote master server domain.

The two examples show the following:

- All of the storage in a deduplication disk pool is exposed as one volume: `PureDiskVolume`.
- The `PureDiskVolume` of the deduplication storage server `bit1.datacenter.symantecs.org` is the source for the replication operation.
- The `PureDiskVolume` of the deduplication storage server `target_host.dr-site.symantecs.org` is the target of the replication operation.

```
> bpstsinfo -lsuinfo -storage_server bit1.datacenter.symantecs.org -stype PureDisk
LSU Info:
    Server Name: PureDisk:bit1.datacenter.symantecs.org
    LSU Name: PureDiskVolume
    Allocation : STS_LSU_AT_STATIC
    Storage: STS_LSU_ST_NONE
    Description: PureDisk storage unit (/bit1.datacenter.symantecs.org#1/2)
    Configuration:
    Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
           STS_LSUF_REP_ENABLED | STS_LSUF_REP_SOURCE)
    Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
    Replication Sources: 0 ( )
    Replication Targets: 1 ( PureDisk:target_host.dr-site.symantecs.org:PureDiskVolume )
    Maximum Transfer: 2147483647
    Block Size: 512
    Allocation Size: 0
    Size: 74645270666
    Physical Size: 77304328192
    Bytes Used: 138
    Physical Bytes Used: 2659057664
    Resident Images: 0

> bpstsinfo -lsuinfo -storage_server target_host.dr-site.symantecs.org -stype PureDisk
LSU Info:
    Server Name: PureDisk:target_host.dr-site.symantecs.org
    LSU Name: PureDiskVolume
    Allocation : STS_LSU_AT_STATIC
    Storage: STS_LSU_ST_NONE
```

```
Description: PureDisk storage unit (/target_host.dr-site.symantecs.org#1/2)
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_ACTIVE | STS_LSUF_STORAGE_NOT_FREED |
      STS_LSUF_REP_ENABLED | STS_LSUF_REP_TARGET)
Save As : (STS_SA_CLEARF | STS_SA_IMAGE | STS_SA_OPAQUEF)
Replication Sources: 1 ( PureDisk:bit1:PureDiskVolume )
Replication Targets: 0 ( )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 79808086154
Physical Size: 98944983040
Bytes Used: 138
Physical Bytes Used: 19136897024
Resident Images: 0
```

Sample volume properties output for OpenStorage backup replication

The following examples show sample output from the `bpstsinfo` command for two OpenStorage devices. The first example is the output from the source disk pool that contains the client backups. The second example is from the target disk pool in the remote master server domain.

The two examples show the following:

- Volume `dv01` on storage server `pan1` is the replication source for volume `dv01` on `pan2`.
- Volume `dv02` on storage server `pan1` is the replication source for volume `dv02` on `pan2`.
- Volume `dv03` on both devices has no replication properties.

```
>bpstsinfo -lsuinfo -storage_server pan1 -stype Pan
```

LSU Info:

```
Server Name: pan1
LSU Name: dv01
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
      STS_LSUF_REP_SOURCE)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ( )
```

Replication Targets: 1 (Pan:pan2:dv01)
 Maximum Transfer: 2147483647
 Block Size: 512
 Allocation Size: 0
 Size: 80525455360
 Physical Size: 0
 Bytes Used: 2285355008
 Physical Bytes Used: 0
 Resident Images: 0

LSU Info:

Server Name: pan1
 LSU Name: dv02
 Allocation : STS_LSU_AT_STATIC
 Storage: STS_LSU_ST_NONE
 Description: E:\
 Configuration:
 Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
 STS_LSUF_REP_SOURCE)
 Save As : (STS_SA_IMAGE)
 Replication Sources: 0 ()
 Replication Targets: 1 (Pan:pan2:dv02)
 Maximum Transfer: 2147483647
 Block Size: 512
 Allocation Size: 0
 Size: 80525455360
 Physical Size: 0
 Bytes Used: 2285355008
 Physical Bytes Used: 0
 Resident Images: 0

LSU Info:

Server Name: pan1
 LSU Name: dv03
 Allocation : STS_LSU_AT_STATIC
 Storage: STS_LSU_ST_NONE
 Description: E:\
 Configuration:
 Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED)
 Save As : (STS_SA_IMAGE)
 Replication Sources: 0 ()
 Replication Targets: 0 ()
 Maximum Transfer: 2147483647
 Block Size: 512
 Allocation Size: 0

Size: 80525455360
 Physical Size: 0
 Bytes Used: 2285355008
 Physical Bytes Used: 0
 Resident Images: 0

>bpstsinfo -lsuinfo -storage_server pan2 -stype Pan

LSU Info:

Server Name: pan2
 LSU Name: dv01
 Allocation : STS_LSU_AT_STATIC
 Storage: STS_LSU_ST_NONE
 Description: E:\
 Configuration:
 Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
 STS_LSUF_REP_TARGET)
 Save As : (STS_SA_IMAGE)
 Replication Sources: 1 (Pan:pan1:dv01)
 Replication Targets: 0 ()
 Maximum Transfer: 2147483647
 Block Size: 512
 Allocation Size: 0
 Size: 80525455360
 Physical Size: 0
 Bytes Used: 2285355008
 Physical Bytes Used: 0
 Resident Images: 0

LSU Info:

Server Name: pan2
 LSU Name: dv02
 Allocation : STS_LSU_AT_STATIC
 Storage: STS_LSU_ST_NONE
 Description: E:\
 Configuration:
 Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED | STS_LSUF_REP_ENABLED |
 STS_LSUF_REP_TARGET)
 Save As : (STS_SA_IMAGE)
 Replication Sources: 1 (Pan:pan1:dv02)
 Replication Targets: 0 ()
 Maximum Transfer: 2147483647
 Block Size: 512
 Allocation Size: 0
 Size: 80525455360

```
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0

LSU Info:
Server Name: pan2
LSU Name: dv03
Allocation : STS_LSU_AT_STATIC
Storage: STS_LSU_ST_NONE
Description: E:\
Configuration:
Media: (STS_LSUF_DISK | STS_LSUF_STORAGE_FREED)
Save As : (STS_SA_IMAGE)
Replication Sources: 0 ( )
Replication Targets: 0 ( )
Maximum Transfer: 2147483647
Block Size: 512
Allocation Size: 0
Size: 80525455360
Physical Size: 0
Bytes Used: 2285355008
Physical Bytes Used: 0
Resident Images: 0
```

About trusted master servers for Auto Image Replication

NetBackup provides the ability to establish a trust relationship between replication domains. A trust relationship is optional. The following items describe how a trust relationship affects Auto Image Replication:

Trust relationship	You can select a subset of your trusted domains as a target for replication. NetBackup then replicates to the specified domains only rather than to all configured replication targets. This type of Auto Image Replication is known as Targeted A.I.R.
No trust relationship	NetBackup replicates to all defined target storage servers. You cannot select a specific host or hosts as a target.

You add a trusted master server in the source domain; you specify a remote master server as a trusted host. A trust relationship is reciprocal: If you add host B as a trusted master server for host A, host B then trusts hosts A.

See [“Adding a trusted master server”](#) on page 233.

You select the targets for replication when you configure a storage lifecycle policy. However, before you choose a specific storage server as a replication target, you must create an import SLP in the target domain. Then, you choose the specific target master server and SLP when you create a storage lifecycle policy in the source domain.

See [“About storage lifecycle policies”](#) on page 600.

See [“Creating a storage lifecycle policy”](#) on page 601.

Note: Before you can configure trust relationships for a clustered NetBackup master server, you must enable inter-node authentication on all of the nodes in the cluster. This requirement applies regardless of whether the clustered master server is the source of the replication operation or the target.

See [“Enabling NetBackup clustered master server inter-node authentication”](#) on page 231.

See [“About NetBackup Auto Image Replication”](#) on page 975.

About the storage lifecycle policies required for Auto Image Replication

To replicate images from one NetBackup domain to another NetBackup domain requires two storage lifecycle policies. The following table describes the policies and their requirements:


Table 25-6 SLP requirements for Auto Image Replication

Domain	Storage lifecycle policy requirements
Domain 1 (Source domain)	<p>The Auto Image Replication SLP in the source domain must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The first operation must be a Backup operation to storage that NetBackup supports for replication. Indicate the exact storage unit from the drop-down list. Do not select Any Available. Note: The target domain must contain the same type of storage to import the image. ■ At least one operation must be a Replication operation to storage in another domain that NetBackup supports for replication from the source storage. You can configure multiple Replication operations in an Auto Image Replication SLP. The Replication operation settings determine whether the backup is replicated to all replication targets in all master server domains or only to specific replication targets. See “About trusted master servers for Auto Image Replication” on page 987. See “New or Change Storage Operation dialog box settings” on page 643. ■ The SLP must be of the same data classification as the Import SLP in Domain 2. Note: If the source master server is at a NetBackup version earlier than 7.6 and the master server in the target domain is at NetBackup 7.6 or later: Do not use the data classification of Any. Use a different data classification in the source domain or the Import job fails.
Domain 2 (Target domain)	<p>If replicating to all targets in all domains, in each domain NetBackup automatically creates an Import SLP that meets all the necessary criteria.</p> <p>Note: If replicating to specific targets, you must create the Import SLP before creating the Auto Image Replication SLP in the originating domain.</p> <p>The Import SLP must meet the following criteria:</p> <ul style="list-style-type: none"> ■ The first operation in the SLP must be an Import operation. NetBackup must support the Destination storage as a target for replication from the source storage. Indicate the exact storage unit from the drop-down list. Do not select Any Available. ■ The SLP must contain at least one operation that has the Target retention specified. ■ The SLP must be of the same data classification as the SLP in Domain 1. Matching the data classification keeps a consistent meaning to the classification and facilitates global reporting by data classification. <p>See the following topic for more information about Replication operation configuration: See “Replicating to all inter-domain targets or to a specific target” on page 990.</p>

[Figure 25-3](#) shows how the SLP in the target domain is set up to replicate the images from the originating master server domain.

Figure 25-3 Storage lifecycle policy pair required for Auto Image Replication

SLP on master server in the source domain



SLP that imports the copies to the target domain

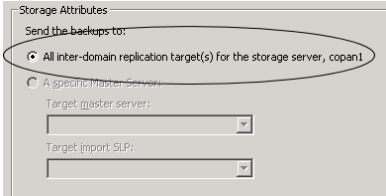
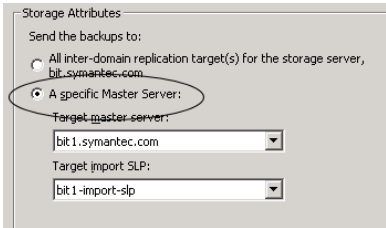
Note: Restart `nbstserv` after you make changes to the underlying storage for any operation in an SLP.

Replicating to all inter-domain targets or to a specific target

With Auto Image Replication, you can replicate backups to all configured replication targets or to a subset of all configured replication targets. To replicate to specific master server domains, you must first configure trusted master servers.

See [“About trusted master servers for Auto Image Replication”](#) on page 987.

Table 25-7 Configuration differences between replicating to all inter-domain targets or to a specific target

Replication objective	Auto Image Replication SLP configuration	Import SLP configuration
Replicate the backup to all configured master server domains.	<p>Create an SLP in the originating domain.</p> <ul style="list-style-type: none"> The first operation must be a Backup operation. The SLP must include a Replication operation. To copy to all domains, select All inter-domain replication target(s). 	<p>The Import SLPs are created automatically in all domains.</p> <p>See Figure 25-4 for a representation of this scenario.</p>
Replicate the backup to targets in specific NetBackup master server domains.	<p>In this case, first create the Import SLPs before the SLP in the originating domain.</p> <ul style="list-style-type: none"> The first operation must be a Backup operation. The SLP must include a Replication operation. Select A specific Master Server and indicate the domain of the target master server. 	<p>The Import SLPs are not created automatically.</p> <p>Note: Create the Import SLP before creating the Auto Image Replication SLP in the originating domain.</p> <p>Create an Import SLP in each target domain.</p> <p>The Import SLP must have an Import operation as the first operation, but can contain other operations as needed.</p> <p>Figure 25-5 represents this scenario.</p>

Additional requirements for Auto Image Replication SLPs are described in the following topic:

See [“About the storage lifecycle policies required for Auto Image Replication”](#) on page 988.

Figure 25-4 Replicating from one domain to all inter-domain master servers


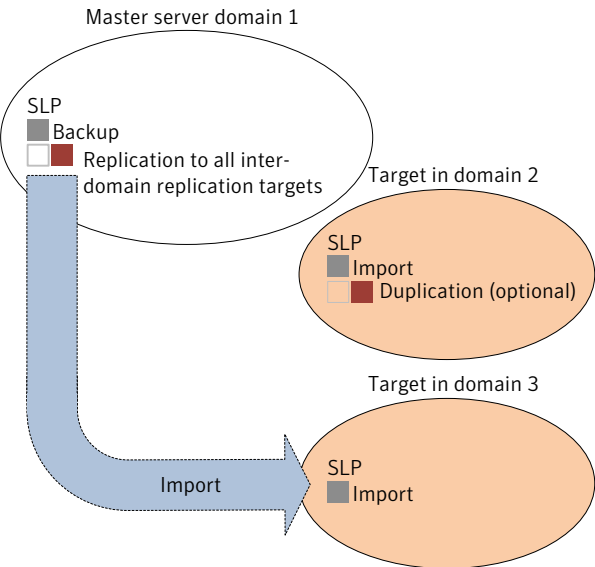


Figure 25-5 represents replication to a specific master domain target.

Figure 25-5 Replicating from one domain to a specific inter-domain master server



To replicate copies to a specific domain, make sure that the target domain master server is a trusted master server of the originating domain.

See [“Adding a trusted master server”](#) on page 233.

Customizing how nbstserv runs duplication and import jobs

The NetBackup Storage Lifecycle Manager (`nbstserv`) runs replication, duplication, and import jobs. Both the Storage Lifecycle Manager service and the Import Manager service run within `nbstserv`.

The NetBackup administrator can customize how `nbstserv` runs jobs by changing the default of certain SLP-related configuration options.

See [“SLP Parameters properties”](#) on page 239.

Using NetBackup Access Control (NBAC) with Auto Image Replication

If Auto Image Replication is configured for two domains and NetBackup Access Control (NBAC) is used, it must be used in both the source domain and the target domain. The configuration for the master servers must be either `USE_VXSS = REQUIRED` or `USE_VXSS = AUTOMATIC`. (However, the setting may be `REQUIRED` in one domain and `AUTOMATIC` in the other.

Auto Image Replication is not supported between master server domains where one master server is configured to use NBAC and NBAC is disabled on the other master server. That is, the configuration for one master server is `USE_VXSS = AUTOMATIC` or `USE_VXSS = REQUIRED` and on the other master server it is `USE_VXSS = PROHIBITED` (disabled).

The following configuration is necessary if NBAC is used in the master server domains:

- In the source master server domain:
The administrator should make sure that the target master server has the permissions set correctly before configuration for the operation begins.
- In the target master server domain:
The security administrator in the target domain must give the administrator in the source domain the correct set of permissions. The source domain administrator needs Browse, Read, and Configure permissions on the following objects: **HostProperties**, **DiskPool**, and **DevHost**.
The source domain administrator can be added as a member to any existing group which has all three permissions.

Consider the following example:

Two NBAC domains each contain a master server:

- Replication source NBAC domain: *DomainA* contains *Master-A*
- Replication target NBAC domain: *DomainB* contains *Master-B*

NBAC is enabled in both the domains. (If NBAC is used in one domain, it must be used in the other domain.)

For *UserA* to create an Auto Image Replication SLP with *Master-B* as the target, *UserA* needs permission on *Master-B* to do so.

A security administrator (*UserB*) in *DomainB* must create a user group (*NB_InterDomainUsers*, for example) and give Browse, Read, and Configure permissions in the following areas:

- **HostProperties**
- **DiskPool**
- **DevHost**

The security administrator in *DomainB* (*UserB*) then assigns *NB_InterDomainUsers* to *DomainA\UserA* using the `bpnbaz -AddUser` command.

Auto Image Replication setup overview

The following table is an overview of the setup process for Auto Image Replication, describing the actions that are required.

Table 25-8 Auto Image Replication setup overview

Step	Action	Description
1	Configure the storage servers	<p>Configure the storage servers for your storage type.</p> <p>See the <i>NetBackup Deduplication Guide</i> or the <i>NetBackup OpenStorage Solutions Guide for Disk</i>:</p> <p>http://www.symantec.com/docs/DOC5332</p>
2	Configure the disk pools.	<p>Configure the disk pools for your storage type.</p> <p>To replicate images from one domain to another requires that suitable storage be configured in each domain. NetBackup must support the storage for replication.</p> <p>See the <i>NetBackup Deduplication Guide</i> or the <i>NetBackup OpenStorage Solutions Guide for Disk</i>:</p> <p>http://www.symantec.com/docs/DOC5332</p>

Table 25-8 Auto Image Replication setup overview (*continued*)

Step	Action	Description
3	Configure the storage units.	Configure the storage units in both the originating domain and the target domain. See “Creating a storage unit” on page 548.
4	Define the relationship between the domains.	Define the relationship between the domains so that the originating domain knows where to send the data. See “About the domain relationship for replication” on page 980. See “About trusted master servers for Auto Image Replication” on page 987.
5	Configure the storage lifecycle policies.	Configure the necessary storage lifecycle policies. See “About the storage lifecycle policies required for Auto Image Replication” on page 988. See “Creating a storage lifecycle policy” on page 601.
6	Configure and run the backup policy in the originating domain.	The backup policy must indicate the configured SLP as the Policy storage selection. See “Creating a backup policy” on page 673.

How to resolve volume changes for Auto Image Replication

When you open the **Change Disk Pool** dialog box, NetBackup loads the disk pool properties from the catalog. NetBackup queries the storage server for changes when you either click the **Refresh** button in the **Change Disk Pool** dialog box or when you configure a new disk pool for the storage server.

Symantec recommends that you take the following actions when the volume topology change:

- Discuss the changes with the storage administrator. You need to understand the changes so you can change your disk pools (if required) so that NetBackup can continue to use them.
- If the changes were not planned for NetBackup, ask your storage administrator to revert the changes so that NetBackup functions correctly again.

NetBackup can process changes to the following volume properties:

- Replication Source
- Replication Target
- None

If these volume properties change, NetBackup can update the disk pool to match the changes. NetBackup can continue to use the disk pool, although the disk pool may no longer match the storage unit or storage lifecycle purpose.

The following table describes the possible outcomes and how to resolve them.

Table 25-9 Refresh outcomes


Outcome	Description
No changes are discovered.	No changes are required.
NetBackup discovers the new volumes that you can add to the disk pool.	The new volumes appear in the Change Disk Pool dialog box. Text in the dialog box changes to indicate that you can add the new volumes to the disk pool.
The replication properties of all of the volumes changed, but they are still consistent.	<p>A Disk Pool Configuration Alert pop-up box notifies you that the properties of all of the volumes in the disk pool changed, but they are all the same (homogeneous).</p> <div></div> <p>You must click OK in the alert box, after which the disk pool properties in the Change Disk Pool dialog box are updated to match the new volume properties</p> <p>If new volumes are available that match the new properties, NetBackup displays those volumes in the Change Disk Pool dialog box. You can add those new volumes to the disk pool.</p> <p>In the Change Disk Pool dialog box, select one of the following two choices:</p> <ul style="list-style-type: none">■ OK. To accept the disk pool changes, click OK in the Change Disk Pool dialog box. NetBackup saves the new properties of the disk pool. NetBackup can use the disk pool, but it may no longer match the intended purpose of the storage unit or storage lifecycle policy. Change the storage lifecycle policy definitions to ensure that the replication operations use the correct source and target disk pools, storage units, and storage unit groups. Alternatively, work with your storage administrator to change the volume properties back to their original values.■ Cancel. To discard the changes, click Cancel in the Change Disk Pool dialog box. NetBackup does not save the new disk pool properties. NetBackup can use the disk pool, but it may no longer match the intended use of the storage unit or storage lifecycle policy.

Table 25-9 Refresh outcomes (continued)

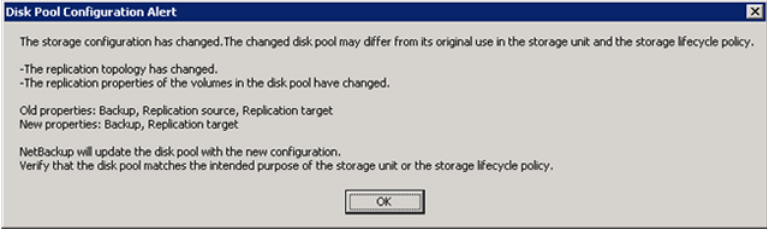

Outcome	Description
The replication properties of the volumes changed, and they are now inconsistent.	<p>A Disk Pool Configuration Error pop-up box notifies you that the replication properties of some of the volumes in the disk pool changed. The properties of the volumes in the disk pool are not homogeneous.</p> <div></div> <p>You must click OK in the alert box.</p> <p>In the Change Disk Pool dialog box, the properties of the disk pool are unchanged, and you cannot select them (that is, they are dimmed). However, the properties of the individual volumes are updated.</p> <p>Because the volume properties are not homogeneous, NetBackup cannot use the disk pool until the storage configuration is fixed.</p> <p>NetBackup does not display new volumes (if available) because the volumes already in the disk pool are not homogeneous.</p> <p>To determine what has changed, compare the disk pool properties to the volume properties.</p> <p>See “Viewing the replication topology for Auto Image Replication” on page 982.</p> <p>Work with your storage administrator to understand the changes and why they were made. The replication relationships may or may not have to be re-established. If the relationship was removed in error, re-establishing the relationships seem justified. If you are retiring or replacing the target replication device, you probably do not want to re-establish the relationships.</p> <p>The disk pool remains unusable until the properties of the volumes in the disk pool are homogenous.</p> <p>In the Change Disk Pool dialog box, click OK or Cancel to exit the Change Disk Pool dialog box.</p>

Table 25-9 Refresh outcomes (continued)

Outcome	Description
NetBackup cannot find a volume or volumes that were in the disk pool.	<p>A Disk Pool Configuration Alert pop-up box notifies you that an existing volume or volumes was deleted from the storage device:</p> <div><div><div><div>Disk Pool Configuration Alert</div><div></div><div><p>An existing volume in this disk pool cannot be found on the storage device and is no longer available to NetBackup. The volume might be offline or deleted. If deleted, any data on that volume is lost.</p><p>Volume(s) deleted: dv02</p><p>Refer to documentation for information on how to resolve this issue.</p><div>OK</div></div></div></div></div> <p>NetBackup can use the disk pool, but data may be lost.</p> <p>To protect against accidental data loss, NetBackup does not allow volumes to be deleted from a disk pool.</p> <p>To continue to use the disk pool, do the following:</p> <ul style="list-style-type: none">■ Use the <code>bpimmedia</code> command or the Images on Disk report to display the images on the specific volume.■ Expire the images on the volume.■ Use the <code>nbdevconfig</code> command to set the volume state to DOWN so NetBackup does not try to use it.

Removing or replacing replication relationships in an Auto Image Replication configuration

Auto Image Replication replicates backups from a storage server in the source domain to storage servers in one or more target NetBackup domains. If a storage server needs to be removed or replaced from such a relationship, the involved domains need to make proper preparations to stop or to redirect replication. That is, to remove the replication relationship.

For example, consider a scenario where cascading Auto Image Replication is used across 3 domains. To remove the storage server in domain 2 (which serves as the destination of the import from Domain 1) preparations need to be made in all three domains. The preparations include modifying storage lifecycle policies and removing the storage server from the replication topology.

Figure 25-6 Removing target storage server example

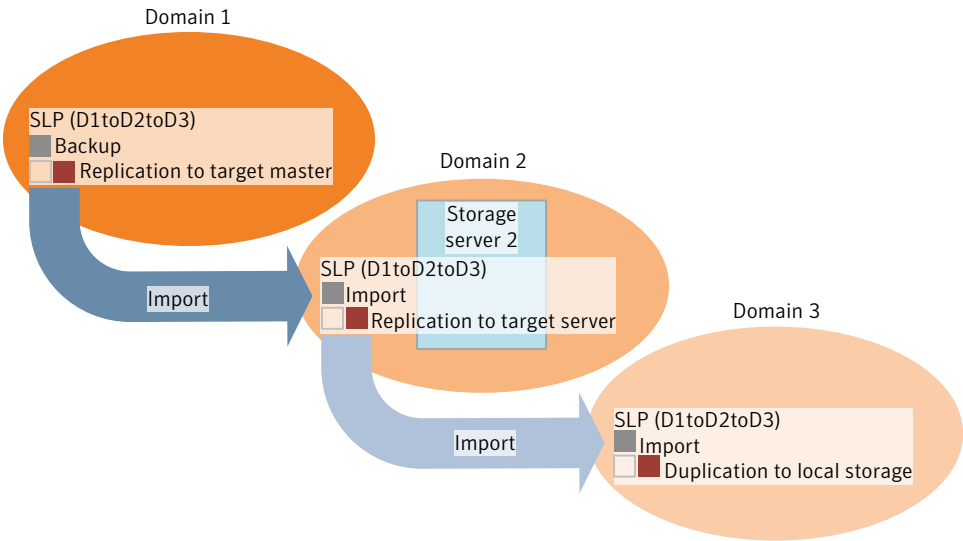


Table 25-10 contains topics that address processes that are involved in removing or replacing relationships in an Auto Image Replication configuration:

Table 25-10

Topic	Reference
Adding a replication relationship between two storage servers.	See “Adding or removing a replication relationship between two storage servers” on page 1000.
Removing a replication relationship between two storage servers.	
Removing a replication relationship between a domain and a storage server.	See “Removing all replication relationships between a domain and a storage server” on page 1000.
Replacing a replication relationship between a domain and a storage server.	See “Replacing all replication relationships between a domain and a storage server” on page 1001.
Removing all replication relationships involving a storage server.	See “Removing or replacing all replication relationships involving a storage server” on page 1005.
Replacing all replication relationships involving a storage server.	

Adding or removing a replication relationship between two storage servers

Before adding a replication relationship, you must understand the consequences of this action:

- Any classic (or non-targeted) Auto Image Replication storage lifecycle policies that replicate from the source domain storage server will replicate to the target domain storage server.

Before removing a replication relationship, you must understand the consequences of this action:

- Any classic (or non-targeted) A.I.R. SLPs that replicate from the source domain storage server will no longer replicate to the target domain storage server.
- Any targeted A.I.R. SLPs that replicate from the source domain storage server to the target domain storage server will fail. These SLPs should be removed or modified so that they no longer replicate to the target domain.

Complete the following steps to add or remove a replication relationship between two storage servers:

To change (add or remove) a replication relationship

- 1 This step depends on whether OpenStorage or MSDP storage is used:
 - For OpenStorage storage, contact your storage administrator to change the replication relationship, as this procedure differs between vendors.
 - For MSDP storage, the relationships can be changed in the source domain by the NetBackup administrator.
- 2 After the relationship is changed, update the disk pools in both domains to reflect the topology changes:

In the NetBackup Administration Console, expand **Media and Device Management > Devices > Disk Pools**. Select and refresh the disk pools. The `nbdevconfig -updatedp` command can also be used. (See the *OpenStorage Solutions Guide* for information about the replication topology for Auto Image Replication.)

Removing all replication relationships between a domain and a storage server

Complete the following steps to remove all replication relationships between two domains. The procedure refers to domain D1 and a storage server in domain D2:

To remove all replication relationships between a domain and a storage server

- 1 In domain D1:

Run the following command on the master server:

```
nbdecommission -list_ref -oldserver  
storage_server_name-machinetype replication_host
```

The output lists:

- All of the storage servers in the domain that reference the deprecated storage server as a replication target or a replication source.
- All of the targeted A.I.R. replication SLPs that reference an SLP that imports to the deprecated storage server.

Note: If an SLP has in-process images, either wait until those images are complete or cancel them before decommissioning the storage server. Note that this includes SLPs of all versions. Use the SLP utility command (`nbstlutil`) to cancel the processing of existing SLP-managed images.

See [“Lifecycle operation administration using the nbstlutil command”](#) on page 606.

- 2 Remove the replication operations from the targeted A.I.R. storage lifecycle polices that were found in step 1.

If these SLPs are no longer necessary, the SLPs can be deleted now.

- 3 Remove the replication relationships that were found in step 1.

See [“Adding or removing a replication relationship between two storage servers”](#) on page 1000.

Any remaining A.I.R. storage lifecycle polices in both domains that are no longer necessary can be deleted now.

- 4 Run the following command to decommission the storage server:

```
nbdecommission -oldserver storage_server_name-machinetype  
replication_host
```

Replacing all replication relationships between a domain and a storage server

Note: If the storage server to be replaced is involved in a classic (non-targeted) Auto Image Replication configuration, the replacement storage server *must not* have any additional replication target relationships.

The procedure refers to two domains: D1 and D2. Auto Image Replication is configured to occur from storage servers in D1 to D2. Also, from D2 to D1.

A new storage server (S2) is added to D2 so that D2 now contains two storage servers (S1 and S2). Relationships to S1 must be replaced with relationships to S2.

[Table 25-11](#) lists the configuration before the switch to D2. [Table 25-12](#) lists the configuration after the changes have been made.

Table 25-11 Example configuration before changes

Domains	Storage servers	Storage lifecycle policies
D1	Contains several storage servers but they are not necessary in the example.	<ul style="list-style-type: none">■ BACKUP_D1_REPLICATE_D2 This SLP replicates to S1 in D2 using the target import SLP IMPORT_S1.■ IMPORT_D1 This SLP imports to a storage server in D1.
D2	S1 (To be deprecated) S2 (New in D2)	<ul style="list-style-type: none">■ BACKUP_D2_REPLICATE_D1 This SLP replicates to D1 using the target import SLP IMPORT_D1.■ IMPORT_S1 This SLP imports to D2 and stores on S1.

Figure 25-7 Topology before changes





Table 25-12 Example configuration after changes

Domains	Storage servers	Storage lifecycle policies
D1	Contains several storage servers but they are not necessary in the example.	<ul style="list-style-type: none">■ BACKUP_D1_REPLICATE_D2 This SLP replicates to S2 in D2 using the target import SLP IMPORT_S2■ IMPORT_S1 This SLP imports to a storage server in D1.
D2	S1 (To be deprecated) S2 (New in D2)	<ul style="list-style-type: none">■ BACKUP_D2_REPLICATE_D1 This SLP replicates to D1 using the target import SLP IMPORT_D1.■ IMPORT_S2 This SLP imports to D2 and stores on S2.

Figure 25-8 Topology after changes



In the following procedure, note that steps must be performed in specific domains.

Table 25-13 Replacing all replication relationships between a domain and a storage server

Domain	Step	Description
In domain D2:	1	<p>Run the following command to determine the replication relationships of storage server S1:</p> <pre>bpstsinfo -lsuinfo -storage_server storage_server_name -stype storage_server_type</pre> <p>For example:</p> <pre>bpstsinfo -lsuinfo -storage_server S1 -stype replication_host</pre>
	2	<p>For each replication target relationship that is found in step 1, add a corresponding relationship between storage server S2 and the target in D1.</p> <p>See “Adding or removing a replication relationship between two storage servers” on page 1000.</p>
	3	<p>Copy all import SLPs that import to storage server S1. Modify the new SLPs to import to storage server S2.</p> <p>For example, copy IMPORT_S1. Modify the SLP to import to storage server S2. Change the SLP name to reflect that it will import to S2: IMPORT_S2.</p> <p>Note: Do not delete the existing import SLPs yet. All of the SLPs that are no longer performing any function can be deleted later in this procedure.</p>
In domain D1:	4	<p>Run the following command on the master server:</p> <pre>nbdecommission -list_ref -oldserver storage_server_name-machinetype replication_host</pre> <p>For example:</p> <pre>nbdecommission -list_ref -oldserver S1 -machinetype replication_host</pre> <p>The output lists:</p> <ul style="list-style-type: none"> ■ All of the storage servers in the source domain (S1) that reference the deprecated storage server as a replication target or replication source. ■ All of the targeted A.I.R. replication SLPs that reference an SLP that imports to the deprecated storage server. <p>Note: If an SLP has in-process images, either wait until those images are complete or cancel them before decommissioning the storage server. Note that this includes SLPs of all versions. Use the SLP utility command (<code>nbstlutil</code>) to cancel the processing of existing SLP-managed images.</p> <p>See “Lifecycle operation administration using the nbstlutil command” on page 606.</p>

Table 25-13 Replacing all replication relationships between a domain and a storage server (*continued*)

Domain	Step	Description
In domain D2:		Complete steps 5-6 only if step 4 listed any replication source relationships.
	5	<p>Modify replication SLPs that replicate from S1 to domain D1 to replicate from S2. This includes both non-targeted and targeted A.I.R. SLPs.</p> <p>For example, change the Replication operation in BACKUP_D2_REPLICATE_D1. Change the backup storage destination from S1 to S2.</p>
	6	<p>Remove the replication source relationships that were found in step 4.</p> <p>See “Adding or removing a replication relationship between two storage servers” on page 1000.</p>
In domain D1:		Complete steps 7-10 only if step 4 listed any replication target relationships.
	7	<p>Add replication relationships from each storage server that was listed in step 2 that has a replication target relationship to storage server S2.</p> <p>See “Adding or removing a replication relationship between two storage servers” on page 1000.</p>
	8	<p>Modify the replication operations that were found in step 4. Change the target import SLP to the corresponding import SLP created in step 3.</p> <p>For example, change the Replication operation in BACKUP_D1_REPLICATE_D2. Change the Target import SLP setting from IMPORT_S1 to IMPORT_S2.</p>
	9	<p>Remove the replication target relationships that were found in step 4.</p> <p>See “Adding or removing a replication relationship between two storage servers” on page 1000.</p>
	10	<p>Run the following command to decommission the storage server:</p> <pre>nbdecommission -oldserver storage_server_name -machinetype replication_host</pre> <p>For example:</p> <pre>nbdecommission -oldserver D1 -machinetype replication_host</pre>
In domain D2:	11	<p>Any import SLPs to S1 that are no longer necessary may now be deleted.</p> <p>For example, IMPORT_S1 can be deleted now.</p>

Removing or replacing all replication relationships involving a storage server

To remove or replace all of the replication relationships involving a storage server, completely remove or replace all replication relationships between a domain and

the storage server. This must be done for each domain that is involved in a replication relationship with the storage server.

The following command lists replication target and replication source relationships. The command is useful for determining which domains have replication relationships with the storage server:

```
bpstsinfo -lsuinfo -storage_server storage_server_name -stype
storage_server_type
```

Example: Replacing a storage server in a non-targeted Auto Image Replication configuration

This example walks through the steps necessary to replace a storage server in a simple, classic (non-targeted) Auto Image Replication configuration.

Specifically, to replace an MSDP storage server (D2_MSDP_1) in domain D2 with another MSDP storage server (D2_MSDP_2).

Table 25-14 Example configuration

Domains	Storage servers	Storage lifecycle policies
D1	D1_MSDP	BACKUP_D1
D2	D2_MSDP_1 D2_MSDP_2	No replication-related SLPs

The changes to the replication and the storage lifecycle topologies are tracked throughout the example.

Figure 25-9 Example replication topology before the process


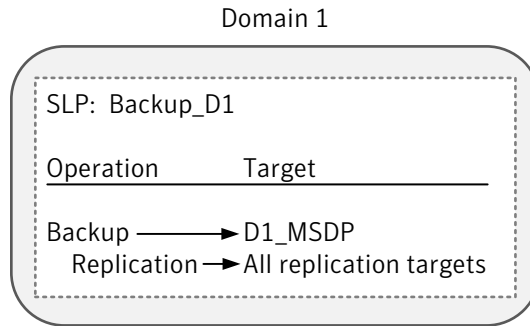



Figure 25-10 Example storage lifecycle policy topology




To replace the storage server D2_MSDP_1 with D2_MSDP_2

- 1 Add a replication target relationship from D1_MSDP to D2_MSDP_2.
- 2 Refresh disk pools in both domains.



- 3 Remove the replication target relationship between D1_MSDP and D2_MSDP_1.

4 Refresh disk pools in both domains.



5 In D1, run the following command to decommission MSDP_D2_1:

```
nbdecommission -oldserver MSDP_D2_1 -machinetype replication_host
```

Example: Replacing a storage server in a cascading, targeted Auto Image Replication configuration

This example walks through the steps necessary to replace an MSDP storage server in a cascading, targeted Auto Image Replication configuration. The changes to the replication and the storage lifecycle topologies are tracked throughout the example.

The environment contains three domains. Each domain contains one or more MDSP storage servers.

Table 25-15 Example configuration

Domains	Storage servers	Storage lifecycle policies
D1	D1_MSDP	BACKUP_D1
D2	D2_MSDP D2_MSDP_2	IMPORT_D2
D3	D3_MSDP	IMPORT_D3

Figure 25-11 Example replication topology


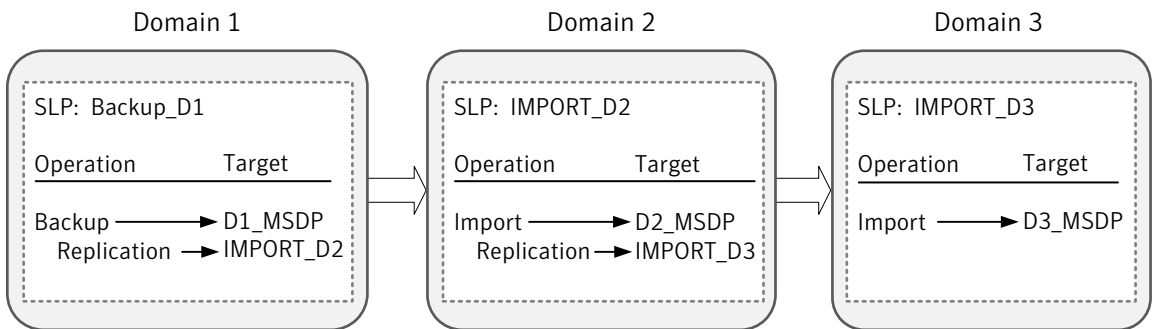



Figure 25-12 Example storage lifecycle policy topology

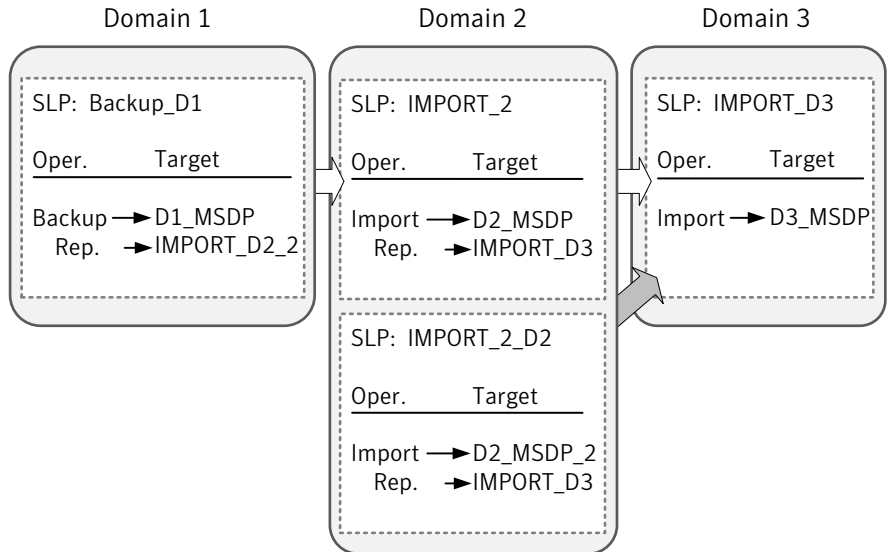


To replace the storage server D2_MSDP with D2_MSDP_2


- 1 Add replication target relationship from D2_MSDP_2 to D3_MSDP.
- 2 Refresh disk pools in both domain D2 and domain D3.



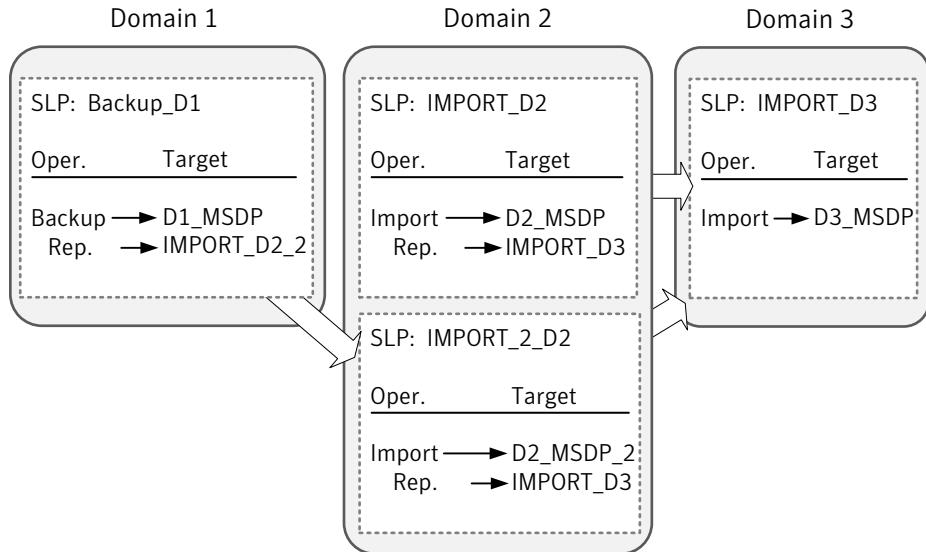
- 3 Copy IMPORT_D2 and modify name to IMPORT_2_D2 and destination storage to a storage unit on D2_MSDP_2.




- 4 Add a replication target relationship from D1_MSDP to D2_MSDP_2.
- 5 Refresh disk pools in both domains.



- 6 Modify BACKUP_D1 Replication Operation Target import SLP to IMPORT_2_D2.



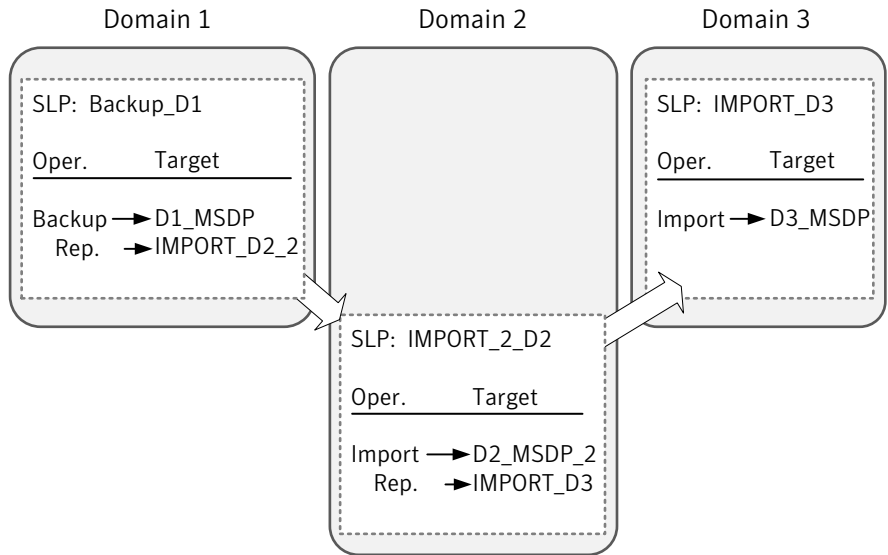
- 7 Remove the replication target relationship from D1_MSDP to D2_MSDP_1.
- 8 Refresh the disk pools in both domains.




- 9 Run the following command in D1:

```
nbdecommission -oldserver MSDP_D2 -machinetype replication_host
```

10 Delete IMPORT_D2.



- 11 Remove the replication target relationship from D2_MSDP_1 to D3_MSDP.
- 12 Refresh disk pools in both domains.



About restoring from a backup at a target master domain

While it is possible to restore a client directly by using the images in the target master domain, do so only in a disaster recovery situation. In this discussion, a disaster recovery situation is one in which the originating domain no longer exists and clients must be recovered from the target domain

Table 25-16 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 the client. 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 the recovery 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 843.
- See [“Enable granular recovery \(policy attribute\)”](#) on page 715.
- See [“Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology \(GRT\)”](#) on page 1187.
- The *NetBackup for Microsoft SharePoint Server Administrator's Guide*:
<http://www.symantec.com/docs/DOC5332>
- The *NetBackup for Microsoft Exchange Server Administrator's Guide*:
<http://www.symantec.com/docs/DOC5332>

Reporting on Auto Image Replication jobs

The Activity Monitor displays both the **Replication** job and the **Import** job in a configuration that replicates to a target master server domain.

Table 25-17 Auto Image Replication jobs in the Activity Monitor

Job type	Description
Replication	<p>The job that replicates a backup image to a target master displays in the Activity Monitor as a Replication job. The Target Master label displays in the Storage Unit column for this type of job.</p> <p>Similar to other Replication jobs, the job that replicates images to a target 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 replicated.</p>

Table 25-17 Auto Image Replication jobs in the Activity Monitor (*continued*)

Job type	Description
Import	<p>The job that imports a backup copy into the target master domain displays in the Activity Monitor as an Import job. An Import job can import multiple copies in one instance. The detailed status for an Import job contains a list of processed backup IDs and a list of failed backup IDs.</p> <p>Note: If the master servers in the source and target domains are not at the same NetBackup version, the following error can occur under certain circumstances: Failed to auto create data classification.</p> <p>This error occurs if the master server in the source domain is at a NetBackup version earlier than 7.6 and the data classification of Any is used. If the master server in the target domain is at NetBackup 7.6, use a different data classification in the source domain or the Import job fails.</p> <p>Note that a successful replication does not confirm that the image was imported at the target master.</p> <p>If the data classifications are not the same in both domains, the Import job fails and NetBackup does not attempt to import the image again.</p> <p>Failed Import jobs 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 an Image Cleanup job. Note that the originating domain (Domain 1) does not track failed imports.</p> <p>See “Running a report” on page 1060.</p>

About NetBackup Replication Director

Replication Director is the implementation of NetBackup OpenStorage-managed snapshots and snapshot replication, where the snapshots are stored on the storage systems of partnering companies. OpenStorage is a Symantec API that lets NetBackup communicate with the storage implementations that conform to the API.

Replication Director uses the functions of the OpenStorage partners to perform the following tasks:

- To share disks so that multiple heterogeneous media servers can access the same disk volume concurrently.
- To balance loads and tune performance. NetBackup balances backup jobs and storage usage among the media servers and disk pools.
- To make full use of disk array capabilities, including fast storage provisioning and almost unlimited storage.
- To use as an alternative to off-site vaulting. Storage replication technology provides an efficient means to send copies of user data (files, applications, databases) to off-site storage as part of a disaster recovery plan.

NetBackup stores snapshots of client data on the volumes that are available to the storage server.

Snapshots represent a point-in-time of primary storage data as captured by the storage hardware. NetBackup can then instruct the storage server to replicate the snapshot from primary volumes to other volumes available to the storage server. The snapshot can be replicated to multiple volumes within the storage server, or to storage outside of the storage server, such as a tape device or other disk storage. Replication Director can accommodate an assortment of scenarios to meet the specific data protection needs of an organization.

Note: Replication is not supported for EMC VNX and Celerra at this time.

Replication Director offers a single NetBackup interface for end-to-end data protection management for the following tasks:

- Unified policy management.
Use the **NetBackup Administration Console** as the one, centralized backup infrastructure to manage the lifecycle of all data.
- Snapshot copy management.
Use NetBackup to manage the entire lifecycle of the snapshot. Replication Director uses OpenStorage with a media server to access the storage server volumes. No image can be moved, expired, or deleted from the disk array unless NetBackup instructs the storage server to do so.
The instruction to perform the initial snapshot comes from an operation in a NetBackup storage lifecycle policy (SLP). You can create one SLP that instructs NetBackup to create the initial snapshot, to replicate the snapshot to several locations, and to indicate a different retention period for each of the replications. Additional instructions (or operations) can be included in the SLP that create a backup from the snapshot, index the snapshot, and more.
- Snapshot copy monitoring.
Use NetBackup OpsCenter to monitor the creation of each copy at each storage location. OpsCenter provides extensive reporting on the entire replication environment.
- Global search and restore.
Recovery is available from any storage device in the environment that is defined to NetBackup. This includes recovery from the primary copy or any replicated copy on disk, or from any duplicated copy on disk or tape.

For more information, see the [NetBackup Replication Director Solutions Guide](#).

Monitoring and reporting

- [Chapter 26. Monitoring NetBackup activity](#)
- [Chapter 27. Reporting in NetBackup](#)
- [Chapter 28. Using the Logging Assistant](#)

Monitoring NetBackup activity

This chapter includes the following topics:

- [About the Activity Monitor](#)
- [Setting Activity Monitor options](#)
- [About the Jobs tab](#)
- [About the Daemons 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](#)

About the Activity Monitor

Use the Activity Monitor in the **NetBackup Administration Console** to monitor and control the following aspects of NetBackup:

Jobs	See “About the Jobs tab” on page 1020. The job details are described in the online Help.
Services or Daemons	See “About the Daemons tab” on page 1027.

Processes See [“About the Processes tab”](#) on page 1039.

Drives See [“About the Drives tab”](#) on page 1044.

- As long as the Activity Monitor is active in the **NetBackup Administration Console**, the `bpj obd` daemon supplies the job activity status to the Activity Monitor.
Updates to the Activity Monitor occur as jobs are initiated, updated, and completed. Without a refresh cycle, updates occur instantaneously.
The status bar appears in the **Jobs** tab, at the top of the Activity Monitor **Details** pane.

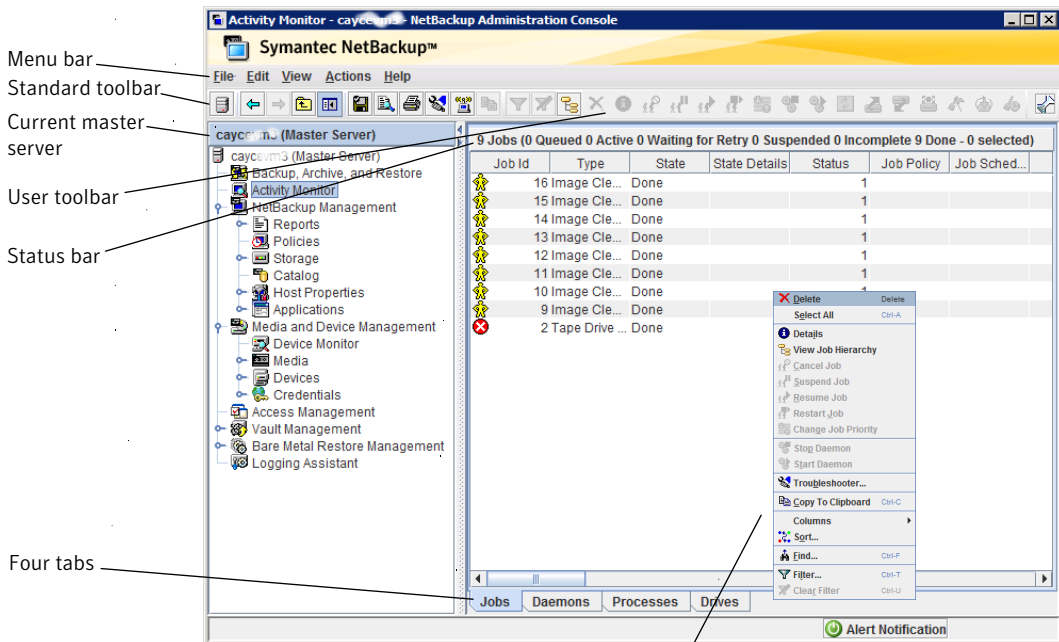
Note: The **Filter** option on the **View** menu is useful for displaying in Activity Monitor only those jobs with specified characteristics. For example, the jobs that were started before a specific date; jobs in the queued state; jobs with status completion codes within a specified range.

The status bar displays the following information:

- 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 or daemons 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 26-1 Activity Monitor



Right-click in pane to view the shortcut menu

Setting Activity Monitor options

The following procedure describes how to set the options for the Activity Monitor in the **NetBackup Administration Console**.

To configure options for the Activity Monitor

- 1 In the **NetBackup Administration Console**, in the menu bar, click **View > Options** and select the **Activity Monitor** tab to access configurable options for the Activity Monitor.
- 2 The following options are available to receive a confirmation warning.

Confirm job deletions	Prompts the user with a confirmation dialog box when a job is deleted.
Confirm job cancellations	Prompts the user with a confirmation dialog box when a job is canceled.
Confirm stop daemons	<p>Enable to display a confirmation dialog box when a daemon is stopped.</p> <p>To discontinue further delete, cancel or stop daemon confirmations, enable the In the future, do not show this warning option.</p>
Maximum details windows	Specifies the maximum number of Activity Monitor job details, daemon details, and the process details windows that can be displayed at one time.
Automatically refresh display	<p>Enable Automatically refresh display to refresh data on the Daemons tab and the Processes tab and the job details elapsed time. Other Jobs tab data refreshes independently of the Auto Refresh setting.</p> <p>Enter the rate (in seconds) at which data refreshes in the Daemons tab and the Processes tab.</p>

- 3 Click **OK** to close the dialog box and apply any changes.

About the Jobs tab

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. The **Jobs** tab also displays details about the jobs. The job details are described in the online Help.

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. The scripts are located in the following directory:

On Windows: `Install_path\VERITAS\NetBackup\bin`

On UNIX: `/usr/openv/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 26-1 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>
Catalog backups	<p>The parent job for catalog backups works with <code>bpdgm</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>A multiple copies job produces one parent job and multiple child jobs. Child jobs that are part of a multiple copies parent job cannot be restarted individually. Only the parent job (and subsequently all the children jobs) can be restarted.</p> <p>See “Multiple copies (schedule attribute)” on page 760.</p>
Multiple data streams	<p>The parent job performs stream discovery and initiates children jobs. A parent job does not display a schedule in the Activity Monitor. Instead, a dash (-) appears for the schedule because the parent schedule is not used and the children schedules may be different. The children jobs display the ID of the parent job in the Activity Monitor.</p>

Table 26-1 Tasks initiated by parent jobs (*continued*)

Task	Description
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 in the Activity Monitor

The following procedure describes how to view job details in the Activity Monitor in the **NetBackup Administration Console**.

To view job details in the Activity Monitor

- ◆ In the **NetBackup Administration Console**, click **Activity Monitor**.

Double-click on a job that is displayed in the **Jobs** tab pane.

The **Job Details** dialog box contains detailed job information on separate tabs:

- **Job Overview** tab
- **Detailed Status** tab
- **Job Hierarchy** tab

To view the details of the previous or next job, click the up or down arrow.

Click **Help** to view the descriptions of the fields.

Not all columns appear by default. In the **NetBackup Administration Console**, click **View > Columns > Layout** or **View > Column Layout** to show or hide columns.

Showing or hiding column heads in the Activity Monitor

The following procedure describes how to show or hide column heads in the Activity Monitor.

To show or hide column heads in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor.
- 2 Select an item in the right-hand pane.
- 3 Click **View > Columns > Layout**. The **Set Column Layout** or **Column Layout** dialog box appears.
- 4 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.
- 5 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.
 - 6 Click **OK** to apply the changes.

Monitoring the detailed status of a selected job in the Activity Monitor

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 in the Activity Monitor

The following procedure describes how to delete a completed job.

To delete completed jobs in the Activity Monitor

- 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**.

When a parent job is deleted, all the children jobs are deleted as well. When a child job is deleted that has no children, that child job alone is deleted.

Canceling a job that has not completed in the Activity Monitor

The following procedure describes how to cancel a job that has not completed.

To cancel a job that has not completed in the Activity Monitor

- 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 a 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 failed (completed) job in the Activity Monitor

The following procedure describes how to restart a job that has completed. Use this procedure to retry a job that has failed.

To restart a completed job in the Activity Monitor

- 1 In the **NetBackup Administration Console**, open the Activity Monitor and select the **Jobs** tab.
- 2 Select the job that 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 in the Activity Monitor

The following procedure describes how to suspend restore or backup jobs.

To suspend a restore or a backup job in the Activity Monitor

- 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 the restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Suspend Job**.

Resuming suspended or incomplete jobs in the Activity Monitor

The following procedure describes how to resume suspended or incomplete jobs.

To resume a suspended or an incomplete job in the Activity Monitor

- 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 the restore jobs that contain checkpoints can be suspended.
- 3 Select **Actions > Resume Job**.

Printing job list information in the Activity Monitor

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 in the Activity Monitor

- 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**.

Exporting Activity Monitor data to a text file

The following procedure describes how to export Activity Monitor data to a text file.

To export Activity Monitor data to a text file

- 1 In the **NetBackup Administration Console**, open the Activity Monitor.
- 2 From any Activity Monitor tab, select **File > Export**.
- 3 Select whether to export all rows or only the rows currently selected.
- 4 Enter the full path to the file where you want the job data to be written, then click **Save**.

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 in the Activity Monitor

- 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.

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 rows that you want to copy.
- 2 Perform one of the following method:
 - Select **Edit > Copy To Clipboard**.
 - Press **Ctrl + C**.
 - Right-click and select **Copy To Clipboard** from the right-click menu.
 - Click the **Copy To Clipboard** icon on the toolbar.
- 3 Paste the selected text into the file (for example, an Excel document).

Note: These methods of copying information from the Activity Monitor also applies to other tables in NetBackup. The table title is not copied.

Changing the Job Priority dynamically from the Activity Monitor

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 26-2 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.

Table 26-2 Change Job Priority options (*continued*)

Option	Description
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.

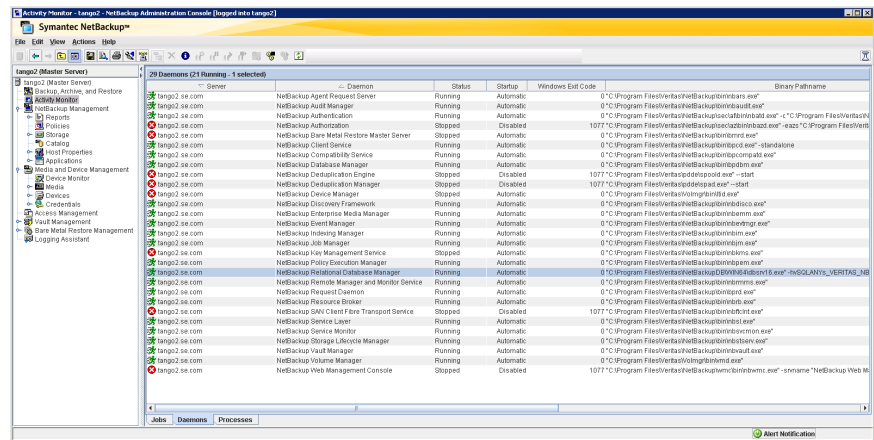
To change the job priority defaults, use the **Default Job Priorities** host properties.

See “[Default Job Priorities properties](#)” on page 113.

About the Daemons tab

In the **NetBackup Administration Console**, the **Daemons** tab in the **Activity Monitor** displays the status of NetBackup daemons on master and media servers.

Figure 26-2 The Daemons tab in the Activity Monitor



Not all columns display by default. Click **View > Column Layout** to show or hide columns.

Table 26-3 describes all of the NetBackup daemons.

Table 26-3 NetBackup daemons

Daemon	Description
NetBackup Agent Request Server (nbars or nbars.exe)	Populates the NetBackup catalog database with database agent metadata and service requests for agents. This service is also responsible for initiating certain actions, such as starting jobs for Oracle cloning.
NetBackup Audit Manager (nbaudit or nbaudit.exe)	The Audit Manager provides the mechanism to query and report on auditing information.
NetBackup Authentication (nbatd or nbatd.exe)	NetBackup Product Authentication validates identities and forms the basis for authorization and access control in NetBackup. The authentication service also generates security certificates that are consumed by various NetBackup components. See “About security certificates for NetBackup media servers” on page 45.
NetBackup Authorization (nbazd or nbazd.exe)	NetBackup Product Authorization provides access control in NetBackup applications.
NetBackup Bare Metal Restore Boot Server Service (bmrbd or bmrbd.exe)	Is present if Bare Metal Restore Boot Server is installed. BMR boot servers manage and provide the resources that are used to rebuild systems.
NetBackup Bare Metal Restore Master Server (bmrmd or bmrmd.exe)	Is present if Bare Metal Restore is installed. Manages restoration data, objects, and servers.
NetBackup BMR MTFTP Services (PXEMTFTP or PXEMTFTP.exe)	Is present if Bare Metal Restore is installed. Provides the PXE protocol services to Bare Metal Restore clients.
NetBackup BMR PXE Service (bmrpxeserver or bmrpxeserver.exe)	Is present if Bare Metal Restore is installed. Provides the TFTP protocol services to Bare Metal Restore clients.
NetBackup Client Service (bpcd or bpcd.exe)	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, bpcd can only be run in standalone mode. On Windows, bpcd always runs under the supervision of bpinetd.exe. NetBackup has a specific configuration parameter for bpcd. If the port number is changed within the NetBackup configuration, the software also updates the port number in the services file.

Table 26-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup CloudStore Service Container (<i>nbcssc</i> or <i>nbcssc.exe</i>)	<p>The CloudStore Service Container is a web-based service container that runs on the media server that is configured for cloud storage. This container hosts different services such as the configuration service, the throttling service, and the metering data collector service. The container requires that an authentication certificate is installed on the media server.</p> <p>See “About security certificates for NetBackup media servers” on page 45.</p>
NetBackup Compatibility Service (<i>bpcompatd</i> or <i>bpcompatd.exe</i>)	Communicates with legacy NetBackup services.
NetBackup Database Manager (<i>bpdbm</i> or <i>bpdbm.exe</i>)	Manages the NetBackup internal databases and catalogs. <i>BPDBM</i> must be running on the NetBackup master server during all normal NetBackup operations.
NetBackup Deduplication Engine (<i>spoold</i> or <i>spoold.exe</i>)	<p>Runs on the NetBackup deduplication storage server host to store and manage deduplicated client data. <i>spoold</i> stands for storage pool daemon; do not confuse it with a print spooler daemon.</p> <p>Active only if the NetBackup Data Protection Optimization Option is licensed and the media server is configured as a deduplication storage server.</p>
NetBackup Deduplication Manager (<i>spad</i> or <i>spad.exe</i>)	<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 Data Protection Optimization Option is licensed and the media server is configured as a deduplication storage server.</p>
NetBackup Device Manager (<i>ltid</i> or <i>ltid.exe</i>)	Starts the Volume Manager (<i>vmd</i>), the automatic volume recognition process (<i>avrd</i>), and any robotic processes. Processes the requests to mount and dismount tapes in robotically controlled devices through the robotic control processes. Mounts the volumes on the tape devices in response to user requests.
NetBackup Enterprise Media Manager (<i>nbemm</i> or <i>nbemm.exe</i>)	<p>Accesses and manages the database where media and device configuration information is stored (<i>EMM_DATA.db</i>). <i>nbemm.exe</i> must be running in order for jobs to run.</p> <p>The service cannot be stopped from the Activity Monitor because it receives data that appears in the NetBackup Administration Console. If it is stopped, the console cannot display the data.</p>
NetBackup Event Management Service (<i>nbevtmgr</i> or <i>nbevtmgr.exe</i>)	Provides the communication infrastructure to pass information and events between distributed NetBackup components. Runs on the same system as the NetBackup Enterprise Media Manager.
NetBackup Indexing Manager (<i>nbim</i> or <i>nbim.exe</i>)	Manages the Hold Service.

Table 26-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Job Manager (<code>njbm</code> or <code>njbm.exe</code>)	Accepts the jobs that the Policy Execution Manager (<code>nbpem</code> or <code>nbpem.exe</code>) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.
NetBackup Key Management Service (<code>nbkms</code> or <code>nbkms.exe</code>)	A master server-based symmetric Key Management Service that provides encryption keys to media server BPTM processes.
NetBackup Legacy Client Service (<code>bpnetd</code> or <code>bpnetd.exe</code>)	<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>The service cannot be stopped from the Activity Monitor because it receives data that appears in the NetBackup Administration Console. If it is stopped, the console cannot display the data.</p> <p>Note: On Windows, 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>To configure a BasicDisk storage unit that uses CIFS, <code>nbrmms</code> must share the same logon credentials as <code>bpnetd</code> on the media server.</p> <p>See “Configuring credentials for CIFS storage and disk storage units” on page 556.</p>
NetBackup Legacy Network Service (<code>vnetd</code> or <code>vnetd.exe</code>)	NetBackup Legacy Network Service (<code>vnetd</code> or <code>vnetd.exe</code>)
NetBackup Policy Execution Manager (<code>nbpem</code> or <code>nbpem.exe</code>)	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 Policy/Client task objects are updated.
NetBackup Proxy Service (<code>nbostrpxy</code> or <code>nbostrpxy.exe</code>)	Executes the Open Storage (OST) calls on any host and returns the results to the proxy plugin side. The plugin side returns them to the application. The proxy server (and plugin) provides a network connection between different servers to relay Open Storage calls.
NetBackup Relational Database Manager (<code>NB_dbsrv</code> or <code>dbsrv12.exe</code>)	<p>Manages the NetBackup relational database. The service must be running on the NetBackup master server during all normal NetBackup operations.</p> <p>On Windows, the display name is SQLANYs_VERITAS_NB.</p>

Table 26-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Remote Manager and Monitor Service (<code>nbrmms</code> or <code>nbrmms.exe</code>)	<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, <code>nbrmms</code> must share the same logon credentials as <code>bpineta</code> on the media server.</p> <p>See “Configuring credentials for CIFS storage and disk storage units” on page 556.</p>
NetBackup Remote Network Transport Service (<code>nbrntd</code> or <code>nbrntd.exe</code>)	<p>Manages the socket connections between a NetBackup media server and a client that is configured for resilient communication. This service runs on the NetBackup master server, NetBackup media servers, and clients. NetBackup starts this service when resilient connections are required between hosts. The service stops when resilient connections are no longer required.</p> <p>One instance of the service can process 256 connections. Multiple instances of the service can run simultaneously.</p> <p>See “NBRNTD_IDLE_TIMEOUT option for NetBackup servers” on page 307.</p> <p>See “Resilient Network properties” on page 209.</p>
NetBackup Request Daemon (<code>bprd</code> or <code>bprd.exe</code>)	<p>Processes the requests from NetBackup clients and servers. <code>bprd</code> also prompts NetBackup to perform automatically scheduled backups. <code>bprd</code> must be running on the NetBackup master server to perform any backups or restores.</p>
NetBackup Resource Broker (<code>nbrb</code> or <code>nbrb.exe</code>)	<p>Allocates the storage units, tape drives, and client reservations for jobs. <code>nbrb</code> works with the Enterprise Media Manager (NEMM).</p> <p>The <code>nbrbutil</code> 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 1032.</p>
NetBackup SAN Client Fibre Transport Service (<code>nbftclnt.exe</code>)	<p>Runs on NetBackup SAN clients. Implements the client side of the Fibre Transport (FT) mechanism. The client FT service opens and closes FT connections and manages the FT connections for shared memory data transfers.</p>
NetBackup Service Layer (<code>nbsl</code> or <code>nbsl.exe</code>)	<p>Facilitates the communication between the NetBackup graphical user interface and NetBackup logic. NBSL is required to run Symantec NetBackup OpsCenter which 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 NetBackup Administration Console. If it is stopped, the console cannot display the data.</p>

Table 26-3 NetBackup daemons (*continued*)

Daemon	Description
NetBackup Service Monitor (<code>nbsvcmon</code> or <code>nbsvcmon.exe</code>)	<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 NetBackup Administration Console. If it is stopped, the console cannot display the data.</p>
NetBackup Storage Lifecycle Manager (<code>nbstserv</code> or <code>nbstserv.exe</code>)	<p>Manages the storage lifecycle operations and schedules duplication jobs. Monitors disk capacity on capacity-managed volumes and removes older images when required.</p> <p>The SLP Manager and the Import Manager run within <code>nbstserv</code>:</p> <ul style="list-style-type: none"> ■ The SLP 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. <p>Note: Restart <code>nbstserv</code> after making changes to the underlying storage for any operation in an SLP.</p>
NetBackup Vault Manager (<code>nbvault</code> or <code>nbvault.exe</code>)	<p>Manages NetBackup Vault. <code>NBVAULT</code> must be running on the NetBackup Vault server during all NetBackup Vault operations.</p>
NetBackup Volume Manager (<code>vmd</code> or <code>vmd.exe</code>)	<p>Manages the volumes (tapes) needed for backup or restore and starts local device management daemons and processes.</p>
Symantec Private Branch Exchange (<code>pbx_exchange.exe</code>)	<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>VRTSspbx</code>.</p>

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 Windows: `Install_path\VERITAS\NetBackup\bin\admincmd\nbrbutil`
- On UNIX: `/usr/opensv/netbackup/bin/admincmd/nbrbutil`

For a complete description of `nbrbutil`, see the [NetBackup Commands Reference Guide](#).

[Table 26-4](#) describes the options available to `nbrbutil` command.

Table 26-4 `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 26-5 describes the configuration settings in detail.
<code>-deleteSetting settingname</code>	Deletes a Resource Broker configuration setting that is identified by <code>settingname</code> .
<code>-dump</code>	Dumps all Resource Broker allocation and request lists.
<code>-dumptables [-f filename]</code>	Enables the Resource Broker to log its internal state in the specified file name.
<code>-help</code>	Lists the help for this command.
<code>-listActiveDriveJobs</code> <code>[driveName]</code>	Lists all the active jobs for a drive.
<code>-listActiveJobs</code>	Lists all the active jobs.
<code>-listActiveMediaJobs mediaId</code>	Lists all the active jobs for a media ID (disk or tape).
<code>-listActivePoolJobs poolName</code>	Lists all the active jobs for a volume pool.
<code>-listActiveStuJobs</code> <code>stuName stugroup</code>	Lists all the active jobs for a storage unit or a storage unit group.
<code>-listOrphanedDrives</code>	Lists the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-listOrphanedMedia</code>	Lists the media that is reserved in EMM but has no corresponding allocation in the Resource Broker.
<code>-listOrphanedPipes</code>	Lists the orphaned Fibre Transport pipes.

Table 26-4 nbrbutil options (*continued*)

Option	Description
<code>-listOrphanedStus</code>	Lists the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-listSettings</code>	Lists the configuration settings of the Resource Broker.
<code>-priority <i>priority</i></code>	Changes the request priority.
<code>-release <i>allocationID</i></code>	Release the allocation with the given identifier.
<code>-releaseAllocHolds</code>	Releases the allocation holds caused by allocation errors for drives and media.
<code>-releaseDrive <i>drivename</i></code>	Releases all allocations for the named drive.
<code>-releaseMDS <i>mdsAllocationKey</i></code>	Releases the EMM and the MDS allocations that MDS allocates by the specified identifier.
<code>-releaseMedia <i>mediaid</i></code>	Releases all allocations for the specified volume.
<code>-releaseOrphanedDrive <i>drivekey</i></code>	Releases the drives that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-releaseOrphanedMedia <i>mediakey</i></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 <i>stuName</i></code>	Releases the storage units that are reserved in EMM but have no corresponding allocation in the Resource Broker.
<code>-reportInconsistentAllocations</code>	Reports inconsistent the allocations between the Resource Broker and MDS.
<code>-resetAll</code>	Resets all Resource Broker allocations, requests, and persisted states.
<code>-resetMediaServer <i>mediaserver</i></code>	Resets all Resource Broker EMM and MDS allocations that are related to <i>ltid</i> 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 26-5 lists the parameters for the `nbrbutil -changesettings` option, and describes the use of each.

Use the `nbrbutil` command with the `-changesettings` option to add or change Resource Broker configuration settings.

Table 26-5 `nbrbutil -changesettings` parameters

Parameter	Description
<code>RB_DO_INTERMITTENT_UNLOADS</code>	<p>When the <code>RB_DO_INTERMITTENT_UNLOADS</code> parameter is set to <i>true</i> (default), <code>nbrb</code> 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><code>RB_DO_INTERMITTENT_UNLOADS=true</code></p>
<code>RB_ENABLE_OPTIMIZATION</code>	<p>When the <code>RB_ENABLE_OPTIMIZATION</code> parameter is set to <i>true</i> (default), this entry instructs <code>nbrb</code> to cache states of resource requests.</p> <p><code>RB_ENABLE_OPTIMIZATION=true</code></p>
<code>RB_RESPECT_REQUEST_PRIORITY</code>	<p>When the <code>RB_RESPECT_REQUEST_PRIORITY</code> parameter is set to <i>false</i> (default), <code>nbrb</code> 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 <code>RB_RESPECT_REQUEST_PRIORITY</code> parameter is set to <i>true</i>, <code>nbrb</code> restarts its evaluation queue at the top of the prioritized job queue after resources have been released.</p> <p><code>RB_RESPECT_REQUEST_PRIORITY=false</code></p>
<code>RB_BREAK_EVAL_ON_DEMAND</code>	<p>When a high priority request appears (a tape span request, a subsequent request for a synthetic or a duplication job, or a read request for an optimized duplication), <code>nbrb</code> immediately interrupts the evaluation cycle. <code>nbrb</code> releases and unloads drives, if required, before a new evaluation cycle is started.</p> <p>If the <code>RB_BREAK_EVAL_ON_DEMAND</code> parameter is set to <i>true</i> (default), the evaluation cycle can be interrupted by high priority requests.</p> <p><code>RB_BREAK_EVAL_ON_DEMAND=true</code></p>

Table 26-5 nbrbutil -changesettings parameters (continued)

Parameter	Description
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 <code>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE</code> parameter sets the maximum number of requests that NetBackup allows in that queue. (Default: 100 requests.)</p> <p><code>RB_MAX_HIGH_PRIORITY_QUEUE_SIZE=100</code></p>
RB_RELEASE_PERIOD	<p>The <code>RB_RELEASE_PERIOD</code> parameter indicates the interval that NetBackup waits before it releases a resource. (Default: 180 seconds.)</p> <p><code>RB_RELEASE_PERIOD=180</code></p>
RB_CLEANUP_OBSOLETE_DBINFO	<p>The <code>RB_CLEANUP_OBSOLETE_DBINFO</code> parameter indicates the number of seconds that can elapse between the cleanup of obsolete information in the <code>nbrb</code> database. (Default: 60 seconds.)</p> <p><code>RB_CLEANUP_OBSOLETE_DBINFO=60</code></p>
RB_MPX_GROUP_UNLOAD_DELAY	<p>The <code>RB_MPX_GROUP_UNLOAD_DELAY</code> parameter indicates the number of seconds that <code>nbrb</code> waits for a new job to appear before a tape is unloaded. (Default: 10 seconds.)</p> <p><code>RB_MPX_GROUP_UNLOAD_DELAY=10</code></p> <p>This setting can help avoid unnecessary reloading of tapes and applies to all backup jobs. During user backups, <code>nbrb</code> uses the maximum value of <code>RB_MPX_GROUP_UNLOAD_DELAY</code> and the Media mount timeout host property setting when <code>nbrb</code> unmounts the tape.</p> <p>During restores, Media mount timeout is used, not <code>RB_MPX_GROUP_UNLOAD_DELAY</code>.</p> <p>See "Timeouts properties" on page 247.</p>
RB_RETRY_DELAY_AFTER_EMM_ERR	<p>The <code>RB_RETRY_DELAY_AFTER_EMM_ERR</code> 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><code>RB_RETRY_DELAY_AFTER_EMM_ERR=60</code></p>
RB_REEVAL_PENDING	<p>The <code>RB_REEVAL_PENDING</code> 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><code>RB_REEVAL_PENDING=60</code></p>

Table 26-5 nbrbutil -changesettings parameters (continued)

Parameter	Description
RB_REEVAL_PERIOD	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.) RB_REEVAL_PERIOD=300

Types of NetBackup daemons

The following table describes additional information about NetBackup daemons, found on the UNIX platform.

Standalone daemons	Always run and listen to accept connections. Examples include <code>bpdbm</code> , <code>bprd</code> , <code>bpjobjd</code> , and <code>vmd</code> .
Multiprocess standalone daemons	Splits or forks a child process to handle requests. Examples include <code>bpdbm</code> and <code>bprd</code> .
Single-process standalone daemons	Accept connections and handle requests in the same process.
<code>inetd</code> daemons	<code>inetd(1m)</code> or <code>bpinetd</code> usually launch these NetBackup daemons. Examples include <code>bpcd</code> , <code>bpjava-msvc</code> , and <code>vnetd</code> .

Symantec recommends that you exit all instances of the **NetBackup-Java Administration Console** after restarting daemons in the **Activity Monitor** or by using a command. Then restart the console with the `jnbSA` command.

The `jnbSA` command is described in the [NetBackup Commands Reference Guide](#).

Monitoring NetBackup daemons

The following procedure describes how to monitor NetBackup daemons.

To monitor NetBackup daemons

- 1 In the **NetBackup Administration Console**, select **Activity Monitor**.
- 2 Select the **Daemons** tab.

- 3 Double-click the name of the daemon to view the details.
- 4 In the **Daemon Details** dialog box, click the up or down arrow to see the details of the next or the previous daemon in the list.

For a description of the daemon, click **Help** in the dialog box.

Starting or stopping a daemon

The following procedure describes how to start or stop a NetBackup daemon.

To start or stop a NetBackup daemon

- 1 In the **NetBackup Administration Console**, select **Activity Monitor**.
- 2 Select the **Daemons** tab.
- 3 Select the daemon(s) that you want to start or stop.
- 4 Select **Actions > Stop Selected** or **Actions > Start Selected**.

Or, select **Actions > Start Daemon** or **Actions > Stop Daemon**.

To start or stop daemons requires the necessary user permissions on the system where the daemon runs.

Displaying all media servers in the Activity Monitor

The **Activity Monitor** may not immediately display all media servers in the **Daemons** tab as soon as the media is added. Also, the **Media Servers** tab of the **Storage Server** dialog may not immediately display all available media servers in a cloud environment.

Even though the media servers may not be visible, it does not affect existing media servers or other NetBackup operations.

To display all media servers:

- Close the **NetBackup Administration Console** on the master server.
- Stop and restart the NetBackup Service Layer (NBSL). Restarting NBSL does not affect any ongoing backup or restore jobs.
- Open the **NetBackup Administration Console**. The services of the newly added media servers should be visible in the **NetBackup Administration Console**.

This situation affects only the **Remote Administration Console** on Windows and the OpsCenter interface.

About the Processes tab

In the **NetBackup Administration Console**, the **Processes** tab in the **Activity Monitor** displays the NetBackup processes that run on master and media servers.

Not all columns display by default. Click **View > Column Layout** to show or hide columns.

[Table 26-6](#) lists and describes the NetBackup processes.

Table 26-6 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.
acssel	None	On UNIX: The NetBackup ACS storage server interface (SSI) event logger <code>acssel</code> logs events.
acsssi	None	On UNIX: The NetBackup ACS storage server interface (SSI) <code>acsssi</code> communicates with the ACS library software host. <code>acsssi</code> processes all RPC communications from <code>acsd</code> or from the ACS robotic test utility that is intended for the ACS library software.
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	<p>The NetBackup Client daemon, this process issues requests to and from the master server and the media server to start programs on remote hosts.</p> <p>On UNIX clients, <code>bpcd</code> can only be run in standalone mode.</p> <p>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 also updates the port number in the services file.</p>
bpcompatd	None	The process for the NetBackup Compatibility service.
bpdbm	13721	<p>The process for the NetBackup Database Manager service.</p> <p>The process that responds to queries that are related to the NetBackup catalog.</p> <p>Manages the NetBackup internal databases and catalogs. This service must be running on the NetBackup master server during all normal NetBackup operations.</p>

Table 26-6 NetBackup processes (*continued*)

Process	Port	Description
bpinetd	None	<p>On Windows: The process for the NetBackup Legacy Client Service.</p> <p>The process that provides a listening service for connection requests.</p> <p>Note: To configure a BasicDisk storage unit that uses CIFS, the media server and the following processes must have the same logon credentials: <code>bpinetd</code>, <code>nbrmms</code>, and <code>vnetd</code>.</p> <p>See “Configuring credentials for CIFS storage and disk storage units” on page 556.</p>
bpjava-msvc	13722	The NetBackup-Java application server authentication service program. <code>bpinetd</code> starts the program during startup of the NetBackup-Java applications and authenticates the user that started the NetBackup-Java 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 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.
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 also updates the port number in the services file.</p>
ltid	None	The process for the NetBackup Device Manager service.
nbaudit		The NetBackup Audit Manager runs on the master server. The Enterprise Media Manager (EMM) maintains audit records in the NetBackup database. The act of starting or stopping <code>nbaudit</code> is audited, even if auditing is disabled.
NBConsole	None	The NetBackup Administration Console on the Windows platform.
nbars	None	The NetBackup Agent Request Server service populates the NetBackup catalog database with database agent metadata and services request for agents. This service is also responsible for initiating certain actions, such as starting jobs for Oracle cloning.
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>

Table 26-6 NetBackup processes (*continued*)

Process	Port	Description
nbEvtMgr	None	The process for the NetBackup Event Manager service. 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 (nbemm) only on master servers.
nbfdrv64	None	The process that controls the Fibre Transport target mode drivers on the media server. <code>nbfdrv64</code> runs on the media servers that are configured for NetBackup Fibre Transport.
nbftsrvr	None	The Fibre Transport (FT) server process that runs on the media servers that are 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	The process for the NetBackup Job Manager service. The process that accepts the jobs that the Policy Execution Manager (NBPEM) submits and acquires the necessary resources. The Job Manager then starts the job and informs <code>nbpem</code> that the job is completed.
nbpem	None	The process for the NetBackup Policy Execution Manager service. It creates Policy/Client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, NBPEM is notified and the appropriate Policy/Client tasks are updated.
nbproxy	None	The process that safely allows multi-threaded NetBackup processes to use existing multi-threaded unsafe libraries.
nbrb	None	This process allocates storage units, tape drives, and client reservations for jobs. <code>nbrb</code> works with the Enterprise Media Manager (NBEMM).
nbrmms	None	The process for the NetBackup Remote Manager and Monitor service. It 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, the media server and the following processes must have the same logon credentials: <code>bpinetd</code> , <code>nbrmms</code> , and <code>vnetd</code> . See “Configuring credentials for CIFS storage and disk storage units” on page 556.
nbsl	None	The process for the NetBackup Service Layer service. <code>nbsl</code> facilitates the communication between the graphical user interface and NetBackup logic.

Table 26-6 NetBackup processes (*continued*)

Process	Port	Description
<code>nbstserv</code>	None	<p>The process for the NetBackup Storage Lifecycle Manager. Manages the storage lifecycle policy operations and schedules duplication jobs. Monitors the disk capacity on the volumes that are capacity-managed and removes older images when required.</p> <p>Note: Restart <code>nbstserv</code> after making changes to the underlying storage for any operation in an SLP.</p>
<code>nbsvcmon</code>	None	The process for the NetBackup Service Monitor. Monitors the NetBackup services. When a service unexpectedly terminates, <code>nbsvcmon</code> attempts to restart the terminated service.
<code>nbvault</code>	None	If Vault is installed, the process for the NetBackup Vault Manager service.
<code>ndmp</code>	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.
<code>opr</code>	None	The NetBackup Volume Manager (<code>vmd</code>) starts the <code>opr</code> operator request daemon. This process receives requests to mount and unmount volumes and communicates the requests to the NetBackup Device Manager <code>ltid</code> . The NetBackup Device Manager communicates the requests to the robotics through SCSI interfaces.
<code>spoold</code>	10082	<p>The process for the NetBackup Deduplication Engine service. It runs on the deduplication storage server.</p> <p>Active only if the NetBackup Data Protection Optimization Option is licensed and configured.</p>
<code>tl4d</code>	13713	The <code>tl4d</code> 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.
<code>tl8d</code> <code>tl8cd</code>	13705	<p>The <code>tl8d</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>tl8cd</code>.</p> <p>The <code>tl8cd</code> process communicates with the TL8 robotics through SCSI interfaces.</p> <p>To share the tape library, <code>tl8cd</code> runs on the NetBackup server that provides the robotic control.</p>
<code>tldd</code> <code>tlddcd</code>	13711	<p>The <code>tldd</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>tlddcd</code>.</p> <p>The <code>tlddcd</code> process communicates with the Tape Library DLT robotics through SCSI interfaces.</p> <p>To share the tape library, <code>tlddcd</code> runs on the NetBackup server that provides the robotic control.</p>

Table 26-6 NetBackup processes (*continued*)

Process	Port	Description
tlhd tlhcd	13717	<p>The <code>tlhd</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>tlhcd</code>.</p> <p>The <code>tlhcd</code> process runs on the NetBackup server that provides the robotic control and communicates with the TLH robotics through SCSI interfaces.</p>
tlmd	13716	The <code>tlmd</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.
vmd	13701	The process for the NetBackup Volume Manager service.
vnetd	13724	<p>This process is preserved for backward compatibility. 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 it connects 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, the media server and the following processes must have the same logon credentials: <code>bpinetd</code>, <code>nbrmms</code>, and <code>vnetd</code>.</p> <p>See “Configuring credentials for CIFS storage and disk storage units” on page 556.</p>
vrts-auth-port	4032	The Veritas Authorization Service verifies that an identity has permission to perform a specific task.
vrts-at-port	2821	The Veritas Authentication Service validates, identifies, and forms the basis for authorization and access.
veritas_pbx	1556	The Symantec Private Branch Exchange allows all socket communication to take place while it connects 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.

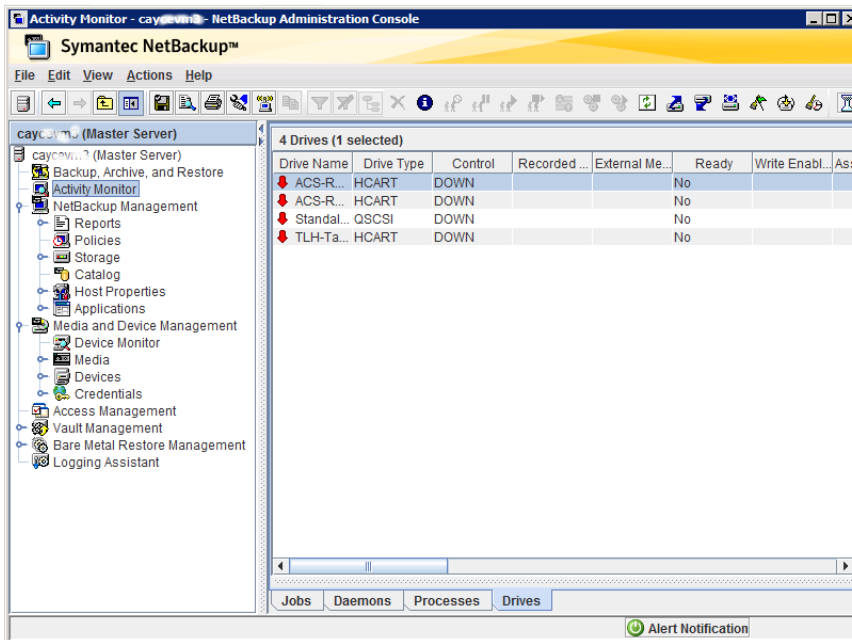
To view the details for a process

- 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 for the selected 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

The **Drives** tab in the Activity Monitor displays the status of NetBackup drives on the monitored server. Use the **Drives** tab to manage drives, device paths, and service requests for operators. Change the status of the drive, clean the tape drive, and perform other operations that you can also perform using the **Device Monitor** under **Media and Device Management**.

Figure 26-3 Activity Monitor Drives tab



To view the details for a drive, double-click the drive in the **Drives** tab pane. For a description of the drive details, click **Help** in the **Drives Details** dialog box.

Monitoring 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 standalone 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 `bpdjobs -clean` command to periodically delete the jobs that are done.

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.

To keep jobs in the jobs database longer than the default of three days, 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.

The `bpdbjobs -clean` is located in the following directory:

On Windows: `install_path\NetBackup\bin\admincmd\bpdbjobs -clean`

On UNIX: `/usr/openv/netbackup/bin/admincmd/bpdbjobs -clean`

Changing the default `bpdbjobs_options` values

Use the one of the following methods to change the default values of the `bpdbjobs_options` on a permanent basis:

- On Windows:

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 | nbsetconfig
```

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 | nbsetconfig
```

Where 192 is the number of hours that successful jobs are kept in the jobs database or Activity Monitor display.

- On UNIX:

Change the entries in the `bp.conf` file.

For example, add the following entry to the `bp.conf` file:

```
KEEP_JOBS_HOURS = 192
```

Where 192 is the number of hours that unsuccessful jobs are kept in the jobs database or Activity Monitor display.

For example, to change the retention of successful jobs, add the following entry:

```
KEEP_JOBS_SUCCESSFUL_HOURS = 192
```

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 on Windows: `BPDBJOBS_OPTIONS` overrides the registry key settings.

If present on UNIX: `BPDBJOBS_OPTIONS` overrides the `bp.conf` settings. To customize the output of `bpdbjobs`, add a `BPDBJOBS_COLDEFS` entry to the `bp.conf` file for each column you want to appear in the output.

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 26-7 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.

In the following example, a batch file (`cleanjobs.bat`) was used on a Windows server. You can copy the script directly from this document and change 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
```

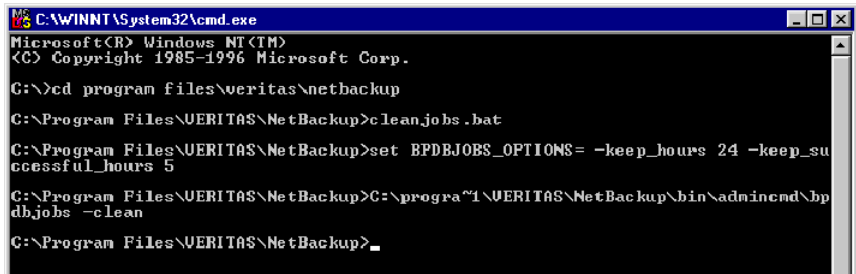
The same script on a UNIX server would look like the following:

```
setenv BPDBJOBS_OPTIONS "-keep_hours 24 -keep_successful_hours 5 -clean"
/usr/opensv/netbackup/bin/admincmd/bpdbjobs ${*}
```

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 26-4 Running cleanjobs.bat on Windows



```
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 completed 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:

On Windows: `install_path\NetBackup\logs\bpdbjobs`

On UNIX: `/usr/openv/netbackup/logs/bpdbjobs`

Note: Before you use a debug log, read the guidelines about legacy logging in the *NetBackup Logging Reference Guide*:

<http://www.symantec.com/docs/DOC5332>

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 your tape drives, disk pools, and service requests for operators, as follows:

Media mounts	See “About media mount errors” on page 1051.
Pending requests and actions	See “About pending requests and actions” on page 1052.
	See “About pending requests for storage units” on page 1052.
	See “Resubmitting a request” on page 1055.
	See “Resolving a pending action” on page 1054.
Tape drives	See “Denying a request” on page 1055.
	See “Changing a drive comment” on page 423.
	See “About downed drives” on page 424.
	See “Changing a drive operating mode” on page 424.
	See “Cleaning a tape drive from the Device Monitor” on page 427.
	See “Resetting a drive” on page 428.
	See “Resetting the mount time of a drive” on page 429.
	See “Setting the drive cleaning frequency” on page 429.
	See “Viewing drive details” on page 430.
	See “Denying a request” on page 1055.

Disk pools

See “[Changing the disk pool state](#)” on page 390.

More information about disk pools is available in the NetBackup guide for your disk storage option:

- The *NetBackup AdvancedDisk Storage Solutions Guide*.
- The *NetBackup Cloud Administrator's Guide*.
- The *NetBackup Deduplication Guide*.
- The *NetBackup OpenStorage Solutions Guide for Disk*.
- The *NetBackup Replication Director Solutions Guide*.

The Guides are available through the following URL:

<http://www.symantec.com/docs/DOC5332>

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 26-8 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 standalone 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 408.

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 the 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

- 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 you licensed a disk pool feature, 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 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

A pending action is similar to a pending request. For a pending action, NetBackup determines the cause of the problem and issues instruction to the operator to resolve the problem.

A media icon identifies a pending action; the icon appears to the left of the request ID. The icon includes a human hand, which depicts that a manual action is required.

Use the following procedure to resolve a pending action.

To resolve a pending action

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**.
- 2 If you licensed a disk pool feature, select the **Drives** tab.
- 3 In the **Pending Requests** pane, select the pending action.
- 4 On the **Actions** menu, select **Display Pending Action** (on UNIX, you can also double-click the 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 1055.

See ["Denying a request"](#) on page 1055.

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 529.

To resubmit a request

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management > Device Monitor**.
- 2 If you licensed a disk pool feature, 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 you licensed a disk pool feature, select the **Drives** tab.
- 3 In the **Pending Requests** pane, select the request.
- 4 On the **Actions** menu, select **Deny Request**.

Reporting in NetBackup

This chapter includes the following topics:

- [About the Reports utility](#)
- [Running a report](#)
- [Copying report text to another document](#)
- [Saving or exporting a report](#)
- [Printing a report](#)

About the Reports utility

Use the **Reports** utility in the **NetBackup Administration Console** 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, manage the report data or select a report to run.

Figure 27-1 Reports utility in the Windows NetBackup Administration Console

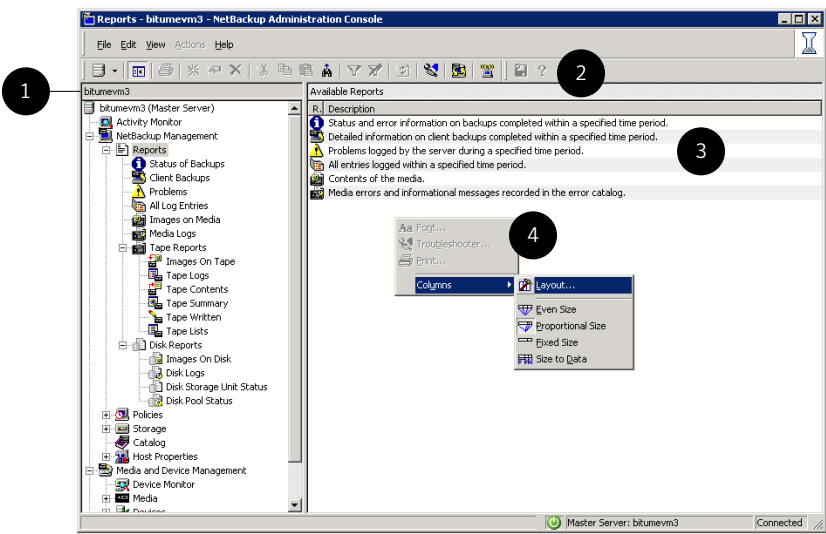


Figure 27-2 Report utility in the Java-based NetBackup Administration Console

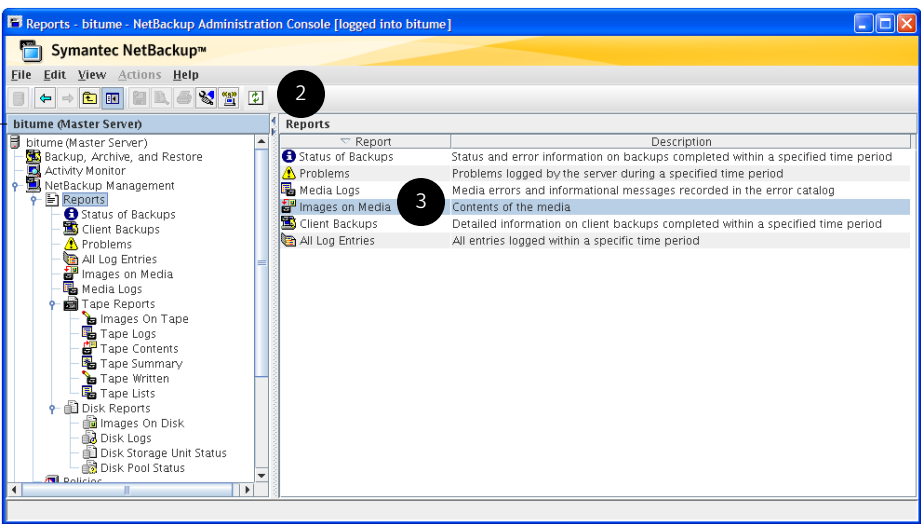


Table 27-1 Reports utility

Number	Description
1	The name of the currently selected master server.

Table 27-1 Reports utility (*continued*)

Number	Description
2	The user toolbar is specific to the Reports utility. See “Standard and user toolbars” on page 45.
3	Report descriptions.
4	Right-click in the right pane to view the shortcut menu.

NetBackup offers many different reports to view information about job activity and media.

Table 27-2 NetBackup Reports

Report name	Description
Status of Backups	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	The Client Backups report shows detailed information about the backups that completed within the specified time period.
Problems	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	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	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	The Media Logs report shows the media errors or the informational messages that are recorded in the NetBackup error catalog.
Images on Tape	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	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.

Table 27-2 NetBackup Reports (*continued*)

Report name	Description
Tape Contents	<p>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.</p> <p>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.</p>
Tape Summary	<p>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.</p> <p>Nonactive media are those with a status of FULL, FROZEN, SUSPENDED, or IMPORTED. Other volumes are considered active.</p> <p>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. 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.</p>
Tape Written	<p>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.</p>
Tape Lists	<p>The Tape Lists report generates information about the volumes that are allocated for backups for the selected media owner or media ID.</p> <p>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.</p>
Images on Disk	<p>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.</p>
Disk Logs	<p>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.</p>
Disk Storage Unit Status	<p>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.)</p> <p>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.</p> <p>Storage units that reference disk groups do not display capacity values.</p>

Table 27-2 NetBackup Reports (*continued*)

Report name	Description
Disk Pool Status	The Disk Pool Status report generates the details of one or more disk pools.

See [“Copying report text to another document”](#) on page 1060.

For information about Vault reports, see the [NetBackup Vault Administrator’s Guide](#).

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 1112.

- 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**.

See [“Copying report text to another document”](#) on page 1060.

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, double-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 to Clipboard**.
- 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**.

See [“Copying report text to another document”](#) on page 1060.

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**.

Using the Logging Assistant

This chapter includes the following topics:

- [About the Logging Assistant](#)
- [Logging Assistant sequence of operation](#)
- [Viewing the Logging Assistant records](#)
- [Viewing the details of a Logging Assistant record](#)
- [Adding a new Logging Assistant record](#)
- [Deleting a Logging Assistant record](#)
- [Setting up debug logging](#)
- [Set minimum debug logging](#)
- [Disabling debug logging](#)
- [Collecting and uploading debug logs](#)
- [Collecting and uploading nbsu information](#)
- [Uploading evidence](#)

About the Logging Assistant

Logging Assistant is a helpful tool that you can use to significantly shorten the time required to set up, collect, and upload debug logs and other information to Symantec Technical Support. Because Logging Assistant automatically performs a number of functions, you can avoid the problems associated with manually logging into NetBackup hosts, creating log directories, and changing logging levels.

Note: Use Logging Assistant under the guidance of Symantec Support.

Logging Assistant uses a series of wizards to help you more quickly troubleshoot a problem. Depending on the category of NetBackup problem that you encounter, the tool suggests the hosts involved in the problem and the logs to enable on those hosts

Note: While collecting NetBackup logs using Logging Assistant, the total NetBackup log size may increase. If you have enabled the new **Keep logs up to GB** option on the **Host Properties > Logging** dialog box and the total NetBackup log size reaches its high water mark, logs are deleted. Logs that you want to retain may also be deleted. To avoid the deletion of logs that you want to retain, you need to disable the **Keep logs up to GB** option while you collect logs using Logging Assistant. Alternatively, you can set the **Keep logs up to GB** option to a value higher than the current value, so that the important logs are not deleted before the log collection is complete.

See [“Logging properties”](#) on page 161.

Logging Assistant lets you upload three categories of data: debug logs, [nbsu](#) diagnostic information, and evidence. Evidence includes any other data that is relevant to solving the problem, such as screen shots, core dumps, and even the debug logs themselves. Each of the categories uses a different wizard.

To use Logging Assistant, make sure that NetBackup 7.6 or later is installed on all hosts. Logging Assistant considers the following as hosts: the master server, one or more media servers (the master server can be the media server as well), and NetBackup clients. No special licensing is required. However, you must have root permissions to use Logging Assistant.

[Table 28-1](#) shows a summary of the main Logging Assistant operations.

Table 28-1 Logging Assistant operations

Operation	Description
Add a new Logging Assistant record.	Add a Logging Assistant record that you use throughout the process of troubleshooting a NetBackup problem. Typically, you associate the record that you create with a failed job that appears in the Activity Monitor. A list of records appears when you select the Logging Assistant node in the right pane of the NetBackup Administration Console . See “Adding a new Logging Assistant record” on page 1070.
View details.	View the details for the selected record.

Table 28-1 Logging Assistant operations (*continued*)

Operation	Description
Delete a Logging Assistant record.	<p>After you upload data to Technical Support and the failed job runs successfully, you can delete the Logging Assistant record.</p> <p>See “Deleting a Logging Assistant record” on page 1072.</p>
Enable debug logging.	<p>Use the Setup Debug Logging Wizard to enable selected NetBackup debug logs and other processes that Technical Support uses to troubleshoot NetBackup problems. Logging Assistant automatically creates the necessary log folders and sets the log levels of the debug logs to the highest verbosity. (The highest log level of many of the debug logs is 5).</p> <p>See “Setting up debug logging” on page 1073.</p>
Set minimum debug logging.	<p>Use the Set Minimum Debug Logging Wizard to set the specified process log levels (verbosity) to a minimum value. This setting still allows info, warning, error, and critical messages (legacy logging), and application logs, diagnostic logs, and some debug logs (unified logging). You may want to set the debug logging to minimum values after you resolve the job failure problem for the selected Logging Assistant record.</p> <p>See “Set minimum debug logging” on page 1077.</p>
Disable debug logging.	<p>Use the Disable Debug Logging Wizard to disable the debug logging after you resolve the job failure problem for the selected Logging Assistant record. The only logs that NetBackup continues to generate are application logs and diagnostic logs.</p> <p>See “Disabling debug logging” on page 1078.</p>
Collect and upload logs.	<p>Use the Collect and Upload Logs Wizard to gather the selected logs onto the master server and upload the data to the Symantec FTP server. The wizard assists you in ensuring that enough disk space is available on the master server to collect the logs.</p> <p>See “Collecting and uploading debug logs” on page 1079.</p>
Collect and upload nbsu.	<p>Use the Collect and Upload nbsu Wizard to gather the <code>nbsu</code> diagnostic information onto the master server and to upload the data to the Symantec FTP server.</p> <p>See “Collecting and uploading nbsu information” on page 1085.</p>
Upload evidence.	<p>Use the Upload Evidence Wizard to upload the evidence that resides in the specified staging location to the Symantec FTP server. Evidence such as screen shots and core dumps can be gathered and uploaded to help in the troubleshooting process.</p> <p>See “Uploading evidence” on page 1088.</p>
Cancel operation.	<p>After you complete a collect and upload wizard, you can cancel the operation if it is currently in progress (for example, if the data is too large and the operation is time-consuming). Check that the progress field of this record displays In progress, then right-click the Logging Assistant record and select Cancel Operation.</p>

Logging Assistant sequence of operation

[Table 28-2](#) shows a typical procedure that uses Logging Assistant to help troubleshoot and solve a NetBackup problem.

Table 28-2 Steps for troubleshooting NetBackup problems

Step	Action	Description
Step 1	Create a Logging Assistant record.	Create a Logging Assistant record to troubleshoot a NetBackup problem. You normally associate the record with a failed job that appears in the Activity Monitor. See “Adding a new Logging Assistant record” on page 1070.
Step 2	Enable debug logging.	Use the Setup Debug Logging Wizard to enable selected NetBackup debug logging that Technical Support uses to troubleshoot NetBackup problems. See “Setting up debug logging” on page 1073.
Step 3	Rerun the failed job.	Go to the Activity Monitor and rerun the failed job. NetBackup generates the debug logs that you have enabled. See “Restarting a failed (completed) job in the Activity Monitor” on page 1024.
Step 4	Collect and upload data.	You can collect and upload debug logs, <code>nbsu</code> diagnostics, and additional evidence. Technical Support may want you to upload the NetBackup Support Utility (<code>nbsu</code>) first to get diagnostic information. Then, if the problem persists you can upload debug logs. You can also collect and upload the evidence that supports the debug logs and <code>nbsu</code> . <ul style="list-style-type: none"> See “Collecting and uploading debug logs” on page 1079. See “Collecting and uploading nbsu information” on page 1085. See “Uploading evidence” on page 1088.
Step 5	Troubleshoot the problem.	Work with Technical Support to read and analyze the debug logs, <code>nbsu</code> diagnostics, and evidence, and correct any problems you detect.
Step 6	Rerun the failed job.	After you take corrective actions, retry the operation. If the job still fails, work with Technical Support to identify and enable additional debug logs using the same Logging Assistant record. Repeat steps 2-6 until the job runs successfully.
Step 7	Disable or set minimum debug logging.	Disable debug logging automatically removes the log directories and disables all debug logging messages. Set minimum debug logging disables all messages except application logs and diagnostic logs. See “Disabling debug logging” on page 1078. See “Set minimum debug logging” on page 1077.

Table 28-2 Steps for troubleshooting NetBackup problems (continued)

Step	Action	Description
Step 8	Delete the record.	Remove the Logging Assistant record from the list of records. See “Deleting a Logging Assistant record” on page 1072.

Viewing the Logging Assistant records

The **Logging Assistant Records** main screen lists all the Logging Assistant records that you create. The records are listed in the right pane of the **NetBackup Administration Console** when you select the **Logging Assistant** node.

The following shows the **Logging Assistant Records** main screen containing one record.

The details for the records (status, description, progress, etc.) are listed in column format. This information is repeated for individual records when you select **View Details** in the **Actions** menu.

Viewing the details of a Logging Assistant record

The **Logging Assistant Records** main screen lists all the Logging Assistant records that you create. The records are listed in the right pane of the **NetBackup Administration Console** when you select the **Logging Assistant** node.

To view the details of a Logging Assistant record

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record that you want to delete.
- 3 Select the record that you want details for.

4 In the **Actions** menu, select **View Details**.

Logging Assistant Record Details

Logging Assistant record ID: 123-456-789

Description: Record_10

Debug logs currently enabled: No

Record status: Logging Assistant record created

Progress:

Staging Directory: C:\Program Files\Veritas\NetBackup\logs\nblastaging

Record creation time: 06/26/2013 12:26:41 PM

Record last modified: 06/26/2013 12:26:41 PM

Debug logging set up time: Logs not enabled

Debug logs upload time: Upload not performed

Associated Hosts and Logs...

OK


The **Logging Assistant Record Details** dialog box displays the following information:

Logging Assistant Record ID	The record name that you enter in the New Logging Assistant Record panel.
Description	The record description that you enter in the New Logging Assistant Record panel.
Debug currently enabled	Yes - you have set up or enabled the debug logs. No - you have not set up the debug logs or you have disabled or reset the logs.

Record status	<p>The following are several of the possible states of the Logging Assistant records:</p> <p>Record created - You have added this record to the Logging Assistant database by using the Add a New Logging Assistant Record panel.</p> <p>Debug logging set up - You have completed the Setup Debug Logging Wizard for this record.</p> <p>Debug Logging partially set up - You have set up debug logging for all the selected hosts and processes by using the Setup Debug Logging Wizard, but not all the hosts and processes have been verified.</p> <p>Debug logging set to minimum - You have completed the Set Minimum Debug Logging Wizard for this record.</p> <p>Debug logging disabled - You have completed the Disable Debug Logging Wizard for this record.</p> <p>Debug logs collected - You have completed the Collect and Upload Debug Logs Wizard, but did not select upload.</p> <p>Logs collected and uploaded - You have completed the Collect and Upload Logs Wizard for this record.</p> <p>nbsu output collected - You have completed the Collect and Upload nbsu Wizard, but did not select upload.</p>
Progress	The progress of the current activity being performed for the Logging Assistant record.
Staging Directory	The staging location where the debug logs and other troubleshooting information is collected and uploaded to the master server.
Record creation time	The date and time when you used the New Logging Assistant Record panel to create this record.
Record last modified	The date and time that you perform the most recent activity on the record by using one of the Logging Assistant wizards.
Debug logging set up time	The date and time when you complete the Setup Debug Logging Wizard.

Debug logs upload time The date and time when you complete the Collect and Upload Debug Logs Wizard.

- 5 Click **Associated Host and Logs** to display the **Associated Hosts and Logs Details** panel. The panel displays the selected hosts for the current record which includes the master server, any media servers, any associated clients, and a list of logs that are enabled on each host.

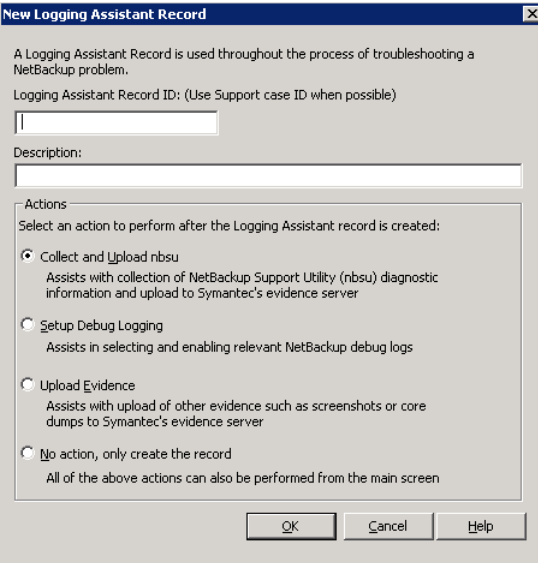


Adding a new Logging Assistant record

The **Add a New Logging Assistant Record** panel contains two fields.

To create a new Logging Assistant record

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 In the **Actions** menu, select **New Logging Assistant Record**. You can also select the operation by using the right-click menu or pressing the Enter key to display the **New Logging Assistant Record** dialog box.



The dialog box is titled "New Logging Assistant Record". It contains the following elements:

- A text box for "Logging Assistant Record ID: (Use Support case ID when possible)".
- A text box for "Description:".
- An "Actions" section with the instruction "Select an action to perform after the Logging Assistant record is created:". It contains four radio button options:
 - ☒ **Collect and Upload nbsu**: Assists with collection of NetBackup Support Utility (nbsu) diagnostic information and upload to Symantec's evidence server.
 - ☐ **Setup Debug Logging**: Assists in selecting and enabling relevant NetBackup debug logs.
 - ☐ **Upload Evidence**: Assists with upload of other evidence such as screenshots or core dumps to Symantec's evidence server.
 - ☐ **No action, only create the record**: All of the above actions can also be performed from the main screen.
- Buttons for "OK", "Cancel", and "Help" at the bottom right.

- 3 In the **Logging Assistant Record ID** text window, type a unique descriptive name for the new policy. Use numeric, alphabetic, plus, minus, underscore, and period characters. Do not use the minus (-) symbol as the first character or leave spaces between characters.

We recommend that you use the Support Case ID as the **Logging Assistant Record ID** to more easily track the record through the troubleshooting process.

- 4 In the optional **Description** field, you can enter a problem summary and the Job ID of the failed job (if applicable).
- 5 The Actions area of this dialog box lists the possible actions that can occur when you click **OK** to leave this dialog:
 - **Collect and Upload nbsu** - Select this option to immediately display the Collect and Upload nbsu Wizard when you leave this dialog. This is often the first action to take after you create the Logging Assistant Record. See ["Collecting and uploading nbsu information"](#) on page 1085.

- Setup Debug Logging - Select this option to immediately display the Setup Debug Logging Wizard after you leave this dialog.
See [“Setting up debug logging”](#) on page 1073.
- Upload Evidence - Select this option to immediately display the Upload Evidence Wizard when you leave this dialog.
See [“Uploading evidence”](#) on page 1088.
- No action - Select this option to return to the main screen when you leave this dialog. The new record appears on the main screen.

Click **OK** to continue. Logging Assistant creates a Logging Assistant record in its database and adds it to the list of records in the Logging Assistant Records pane.

Deleting a Logging Assistant record


After you upload logs to Technical Support and they confirm that the required evidence is available, you can delete the Logging Assistant record.

Warning: Before you delete the record, ensure that logging is disabled or set to minimum.

To delete a Logging Assistant record

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record that you want to delete.

- 3 On the **Actions** menu or on the right-click menu, select **DeleteDelete Record** to display the **Delete Logging Assistant Record** dialog box.



- 4 Click **Yes** in this dialog box if you are sure that you want to delete the selected Logging Assistant record. Make sure that you have uploaded the required logs for this record to the Symantec server.

Setting up debug logging

Use the Setup Debug Logging Wizard to enable selected NetBackup debug logging that Technical Support uses to troubleshoot NetBackup problems. Logging Assistant automatically creates the necessary log folders and sets the verbosity of the pertinent logs to their highest level.

Note: Ensure that each of the selected hosts contains enough available space for the selected debug logs.

To set up debug logging

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 Select the Logging Assistant record in the right pane for which you want to set up debug logging. On the **Actions** menu or on the right-click menu, select **Setup Debug Logging**.
- 3 The **Welcome** panel for the Setup Debug Logging Wizard appears. Click **Next** to continue.
- 4 The **Job Selection** panel appears. Do one of the following:
 - Enable the **Analyze job to identify relevant hosts and debug logging** parameter. Then enter the **Job ID**, and click **Next** to continue. If you do not know the Job ID, go to the Activity Monitor to find the Job ID of the failed job. Logging Assistant identifies the hosts and the pertinent debug logs related to the job.

Based on the Job ID that you enter, Logging Assistant identifies and selects debug logs for hosts and processes that enable the most effective troubleshooting of this problem. The wizard panels that follow this panel show the processes that Logging Assistant has selected.

- For the following scenarios, skip this panel that sets up the logs based on a Job ID, and click **Next**:
 - The problem does not involve a particular NetBackup job.
 - You already know the hosts and debug logs that you want to enable.
 - You want to re-enable the logs that you previously set up for this record. Those log selections appear automatically in the panels that follow.

The following **Host Selection** panel appears:

The screenshot shows the 'Host Selection' panel of the 'Setup Debug Logging Wizard for Logging Assistant Record: 123-456-789'. The panel has a yellow header with the title 'Host Selection' and the subtitle 'Select hosts for Setup Debug Logging'. Below the header, there is a text box with the instruction 'Select host names to set up required debug logging for this Logging Assistant record.'.

There are three main sections for host selection:

- Master Server:** A checkbox labeled 'Setup debug logging on Master Server' is unchecked. Below it is a text box labeled 'Master Server:' containing the text 'orbitervm4'.
- Media Servers:** A checkbox labeled 'Setup debug logging on Media Server' is checked. Below it are two list boxes: 'Available Media Servers' (empty) and 'Selected Media Servers' (containing 'orbitervm4'). Between these list boxes are two buttons: 'Add >' and '< Remove'.
- NetBackup Clients:** A checkbox labeled 'Setup debug logging on NetBackup Client' is unchecked. Below it is a text box labeled 'NetBackup Clients (separate multiple entries with comma):' which is empty.

At the bottom of the panel, there is a text box with the instruction 'To continue, click Next.' and a row of four buttons: '< Back', 'Next >' (highlighted with a dashed border), 'Cancel', and 'Help'.

- 5 Do one or more of the following:
 - Click **Setup debug logging on Master Server** to set up debug logging on the master server. The **Master Server** text window displays the master server.

- Click **Setup debug logging on Media Server(s)** to set up debug logging on the media servers. If the master server is also a media server, you may still want to select it again as a media server. Use the right arrow option (>>) to move a highlighted media server from the **Available Media Servers** window across to the **Selected Media Servers** window.
- Click **Setup debug logging on Client(s). Enter Client names.** to set up debug logging on clients. In the text window under this parameter, type the names of clients on which you want to set up debug logging. Separate the client names by commas. Do not use any spaces between the client names.

Click **Next** to display the panel.

Click **Next** to continue. Logging Assistant validates the specified hosts to establish connectivity and validates the NetBackup versions that are installed on the hosts.

- 6 If you selected **Setup debug logging on Master Server**, the **Setup Debug Logging on Master Server** panel appears.

Setup Debug Logging Wizard for Logging Assistant Record: 12345679

Master Server Debug Logging Selection
Confirm or Select Debug Logging to set up

Confirm (or select) required debug logging (0 problem categories are selected)

Problem Category	Process / Component
<input type="checkbox"/> * Backup	nbjm,nbpem,nbproxy
<input type="checkbox"/> * Duplication	bpdbm,bpduplicate
<input type="checkbox"/> * Job Scheduling	nbjm,nbpem,nbproxy
<input type="checkbox"/> * Network Communication	bpcd,libmessaging,nbjm,px,vnetd
<input type="checkbox"/> * Resource Allocation	bpdbm,nbemn,nbjm,nbproxy,nrb
<input type="checkbox"/> * Restore	bprd,nbjm
<input type="checkbox"/> * Scheduled Application Backup/Restore	bprd,nbjm,nbpem
<input type="checkbox"/> * Snapshot	bprd
<input type="checkbox"/> * User-Directed Application Backup/Rest...	bprd
<input type="checkbox"/> Application - Automatic Schedule Job	nbjm
<input type="checkbox"/> Application - Oracle Guided Recovery (clo...	bprd,dars,nbars,nbsl
<input type="checkbox"/> Application - Oracle RMAN Metadata Coll...	bprd,dars,nbars
<input type="checkbox"/> Backup - BMR	bmr,bmrcommon,bmrconfig,bmrdr
<input type="checkbox"/> Backup - Synthetic	bpdbm,bpsynth
<input type="checkbox"/> Duplication - DSSU	nbpem
<input type="checkbox"/> Duplication - Vault	bpdbm,bpduplicate,vltrun

* prefix denotes a common problem category

☐ If required, specify additional components to set up debug logging (separate multiple entries with comma)

To continue, click Next.

< Back Next > Cancel Help

This panel contains a list of problem categories on the master server and the process names associated with each category of problem. Click the problem category or categories that apply to the failed job.

For example, if you are seeing problems with synthetic backups on the failed job, click **Backup - Synthetic**. The debug logs for `bpcd`, `bpdm`, and `bptm` are enabled.

Click **If required, specify additional components to set up debug logging** to enter components to enable debug logging. For example, if you have Java interface problems, but only want the debug logs for `jnbSA`, instead of all four of the processes list, type `jnbSA` in the text window. Type multiple components separated by commas with no spaces between entries.

Technical Support publishes a list of valid components that you can specify in this text box.

<http://www.symantec.com/docs/TECH204256>

Click **Next** to continue.

- 7 Perform the same actions for the media server (if selected in step 5) and the clients (if selected in step 5).

Click **Next** to continue.

- 8 The **Setup Debug Logging Summary** panel confirms all the items you selected in the previous panels of the Setup Debug Logging Wizard. Review the information. The summary shows on which hosts of your NetBackup system that you chose to enable logs: master server, media servers, and clients or a combination of these items. It also shows the host names and the logs that are enabled on each.

Click **Next** to enable the debug logging. All pertinent logs are set to their highest (most verbose) level.

- 9 The **Setup Debug Logging Results** panel displays the hosts and components that have been successfully enabled and unsuccessfully enabled.

Click **Finish** to return to the **Logging Assistant Records** panel of the **NetBackup Administration Console**.

Set minimum debug logging

The Set Minimum Debug Logging Wizard panel sets to minimum the logging levels of the NetBackup debug logging that the Setup Debug Logging Wizard enabled.

Information about the logging levels that this wizard configures is available in the following document:

<http://www.symantec.com/docs/TECH204256>

To set minimum debug logging

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record whose debug logging you want to set to minimum.
- 3 In the **Actions** menu, select **Set Minimum Debug Log**. The **Welcome** panel for the Set Minimum Debug Logging Wizard appears.

Click **Next** to continue.

- 4 The **Set Minimum Debug Logging Selected Hosts and Processes** panel appears. Use this panel to confirm the host problem categories and processes for all debug logs that were enabled. The hosts can be the master server, one or more media servers, and one or more clients.

Click **Next** to continue.

- 5 The **Set Minimum Debug Logging Results** panel displays the results of the reset operation. The upper window shows the hosts and processes whose debug logs are successfully reset. The lower window shows the unsuccessful debug logging resets.

Click **Finish** to complete the Set Minimum Debug Logging Wizard.

Disabling debug logging

The Disable Debug Logging Wizard disables the NetBackup debug logging that the Setup Debug Logging Wizard enabled. A disabled debug log cannot generate any logging messages.

To disable debug logging

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record whose debug logging you want to disable.
- 3 In the **Actions** menu, select **Disable Debug Logs**. The **Welcome** panel for the Reset Logging Wizard appears.

Click **Next** to continue.

- 4 The **Selected Hosts and Processes** panel appears. Use this panel to confirm the host problem categories and processes for all debug logs that were enabled. The hosts can be the master server, one or more media servers, and one or more clients.

Click **Next** to continue.

- 5 The **Disable Debug Logging Results** panel displays the results of the disable operation. The upper window shows the hosts and processes whose debug logs are successfully disabled.

The lower window shows the results of the unsuccessful debug disable attempts. Logs that other Logging Assistant records use cannot be disabled, and are listed as unsuccessful in the lower window.

Click **Finish** to complete the Disable Debug Logging Wizard.

Collecting and uploading debug logs

The Collect and Upload Logs Wizard helps you gather debug logs to a staging area of the master server, then uploads the log data to the Symantec FTP server. The debug logs have previously been enabled in a Setup Debug Logging Wizard session.

To collect and upload debug logs

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record whose debug logs you want to collect and upload.
- 3 In the **Actions** menu, select **Collect and Upload Debug Logs**. The **Welcome** panel for the Collect and Upload Logs Wizard appears.

Click **Next** to continue.

- 4 The **Selected Hosts and Processes** panel appears. Use this panel to confirm the host problem categories and processes for all debug logs that you want to collect and upload. The hosts can be the master server, one or more media servers, and one or more clients.

Click **Next** to continue.

- 5 The **Collect Debug Logs** panel appears.

Collect And Upload Log Wizard for Logging Assistant Record: 123-456-789

Collect debug logs
Provide information to collect debug logs

Debug logs for this Logging Assistant Record were enabled on 06/26/2013 01:59:09 PM.

Specify Log Duration

☒ Collect available logs

☐ Collect logs spanning this time period: Last 24 hours

☐ Collect logs for specific time period:

Date Time Range

From: 6/25/2013 2:19:59 PM

To: 6/26/2013 11:59:59 PM

☐ Collect nbsu output

Click "Calculate" to check required and available disk space on specified host(s) **Calculate...**

To continue, click Next.

< Back Next > Cancel Help

The following describes the parameters on this panel:

Specify Log Duration

In the **Specify Log Duration** area, click one of the following:

- **Collect available logs** gathers all messages for the selected debug logs. Be careful
- **Collect logs spanning this time period** collects logs from one of three periods of time: the last 24 hours, the last 3 days, or the last week.
- **Collect logs for specific time period** collects log messages that are generated within a precise time window. Specify the start date and time and the end date and time. The statement at the top of this panel indicates when the debug logs for this Logging Assistant Record were enabled. This date-time may affect the time window you select.

Collect nbsu output

Click **Collect nbsu output** to collect the NetBackup Support Utility (*nbsu*) information to be bundled and sent with the debug logs. *nbsu* contains a wide range of diagnostic information that Technical Support can use to troubleshoot.

Calculate

Click **Calculate** to display the estimated available disk space and default log location for each selected host. Use this information to determine if an alternate temporary location is necessary.

If you click **Calculate**, go to step 6.

Otherwise, click **Next** and go to step 8.

6 The Calculate Disk Space on Hosts panel appears.

Calculate Disk Space

Required (rough estimation) and available disk space in temporary location on the host(s) is below. To change temporary location to use on a specific host, select the row and click "Change Temporary Location".

Change Temporary Location...


Host Name	Required Di...	Free Disk S...	Temporary Location
orbitervm4	7.7 MB	26787 MB	C:\Program Files\Veritas\NetBac...

Total estimated log size: 7.7 MB

OK Cancel Help

This panel displays the free disk space information and the default temporary location for each of the selected hosts. Verify that the required disk space on the host for the debug logs does not exceed the available free space on that host. If it does, you can select the host entry, then click **Change Temporary Location** to change the log location on the host.

7 The **Change Temporary Location** panel appears.



This panel lets you move debug logs to a different location on a host to avoid running out of disk space. Enter an alternate temporary directory path location for the logs in the text window, then click **Check Free Disk Space**. The **Free disk space on temporary location** entry changes from Unknown to the free space available on the alternate location. Compare that value with the **Estimated required disk space for logs** entry to determine if the space is sufficient. If necessary, repeat until you find an alternate log location with sufficient disk space.

Click **OK** to return to the **Calculate Disk Space** panel.

Click **OK** again to return to the **Collect debug logs** panel.

Finally, click **Next** to continue.

8 The following **Upload Collected Evidence** panel appears.

The screenshot shows a wizard window titled "Collect And Upload Log Wizard for Logging Assistant Record: 123-456-789". The main heading is "Upload collected logs on Symantec server" with the subtitle "Provide information to upload logs". A gear icon is in the top right corner. The text explains that bundled debug logs will be collected in a new directory in the Master Server staging location and that 0.02 MB of disk space is required. A "Check Free Disk Space..." button is provided. The "Master Server staging location:" is set to "in Files\Veritas\NetBackup\logs\nblastaging" with a "Browse..." button. The "Upload collected logs to Symantec server" checkbox is checked. The "Support Case ID:" is "123-456-789" (example: 01234567) and the "Symantec server:" is "USA". There is a text box for a "Brief comment to include in the README (max 256 chars)". The "Clean-up staging location on Master Server after successful upload" checkbox is also checked. At the bottom, it says "To start collect and upload operation, click Next." and has buttons for "< Back", "Next >", "Cancel", and "Help".

The panel lets you enter information necessary to upload debug logs and other evidence such as `nbsu` output.

The following describes the parameters on this panel:

Check Free Space

Note the staging file location in the text window. If you want, change the location. If you are uncertain whether enough free space is available for the staging location, click **Check Free Space** to display the following:

[illegible]

The **Check Free Space on Master Server** panel lets you view all the mount points on the master server and the space available on each. Click **OK** to return to the **Calculate Disk Space on Hosts** panel.

Staging location on master server

Logging Assistant requires a staging location on the master server for the debug logs to be uploaded. If you determine from the **Check Free Space on Master Server** panel that you need to change the staging location, type the new location in the text window.

Upload collected logs to Symantec server

Select **Upload collected logs to Symantec server** if you want to upload the debug logs to Symantec Support for analysis and troubleshooting. Enter the Support Case ID associated with the Logging Assistant record, and enter the Symantec server location (for example, USA). Contact Technical Support to get this information.

Brief comment to be included in Readme

You can add a brief message of up to 256 characters in the text window. The message is included in the Readme file that is uploaded with the debug logs to Symantec Support.

Remove staging directory on master server after upload

Click **Remove staging directory on master server after upload** to remove the staging location and all the uploaded data still on the master server.

Click **Next** to initiate the upload operation.

- 9 The **Collect and Upload Results** panel appears. It shows the results of the upload operation. The progress of the operation is recorded in a file on the master server.

You can monitor the progress of the collect and upload operations by returning to the Logging Assistant main screen. The Progress column of the main screen changes as the operation progresses.

Click **Finish** to complete the Collect and Upload Logs Wizard.

Collecting and uploading nbsu information

Technical Support often wants you to send initial evidence of a problem in the form of the NetBackup Support Utility (nbsu). You use the Collect and Upload nbsu Logs Wizard to collect and upload the nbsu diagnostic information.

To collect and upload nbsu

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record whose nbsu you want to collect and upload.
- 3 In the **Actions** menu, select **Collect and Upload nbsu Logs**. The **Welcome** panel for the Collect and Upload nbsu Logs Wizard appears.

Click **Next** to continue.

4 The **Host Selection** panel appears:

5 Do one or more of the following:

- Click **Collect nbsu on master server** to collect `nbsu` information on the master server. The **Master server** text window displays the master server name.
- Click **Collect nbsu on media server** to collect `nbsu` information from one or more media servers. If the master server is also a media server, you may still want to select it again as a media server. Use the right arrow option (>>) to move a highlighted media server from the **Available Media Servers** window across to the **Selected Media Servers** window.
- Click **Collect nbsu on NetBackup client(s)** to collect `nbsu` information from one or more clients. In the text window under this parameter, type the names of clients on which you want to collect `nbsu` information. Separate the client names by commas. Do not use any spaces between the client names.

Click **Next** to continue.

6 The Upload nbsu panel appears:

Enter the preparatory information required to upload the `nbsu` information.

The following describes the parameters on this panel:

Staging location on master server

Logging Assistant requires a staging location on the master server for the `nbsu` information to be uploaded. If you determine from the **Check Free Space on Master Server** panel that you need to change the staging location, type the new location in the text window.

Upload nbsu

Select **Upload nbsu** if you want to upload the `nbsu` diagnostic information to Technical Support for analysis and troubleshooting. Enter the Support Case ID associated with the Logging Assistant record, and enter the Symantec server location (for example, USA). Contact Technical Support to get this information.

Brief comment to be included in Readme

You can add a brief message of up to 256 characters in the text window. The message is included in the Readme file that is sent with the debug logs to Technical Support.

Remove staging directory on master server after upload

Click **Remove staging directory on master server after upload** to remove the staging location and to free up the memory that the staging location uses on the master server.

Click **Next** to initiate the collection and upload of the `nbsu`.

- 7 The **Collect and Upload nbsu Results** panel lets you view the results of the `nbsu` collect and upload process. The Operation Details window of the panel shows the file on the master server where the progress of the operation is recorded.

Click **Finish** to complete the Collect and Upload `nbsu` Wizard.

Uploading evidence

The Upload Evidence Wizard lets you upload the auxiliary evidence which includes information such as screen shots, core dumps, `nbclogs`, and emails.

To upload evidence

- 1 In the **NetBackup Administration Console**, in the left pane, select the **Logging Assistant** node.
- 2 The right pane of the **NetBackup Administration Console** displays the list of Logging Assistant records. Select the record whose `nbsu` you want to collect and upload.
- 3 In the **Actions** menu, select **Upload evidence**. The **Welcome** panel for the Upload Evidence Wizard appears.

Click **Next** to continue.

4 The Upload Evidence panel appears:

Upload Evidence Wizard for Logging Assistant Record: 123-456-789

Upload evidence to Symantec server
Provide information to upload evidence

Source location on Master Server: Browse...

Support Case ID: (ex. 01234567)

Symantec server:

Brief comment to include in the README (max 256 chars):

To start the upload, click Next.

< Back Next > Cancel Help

This panel lets you enter the preparatory information that is needed to upload evidence.

The following describes the parameters on this panel:

Evidence location on master server

Logging Assistant requires a staging location on the master server for the evidence to be uploaded. Check the default location in the text window. If you need to change the location, type the new location in the window.

Support Case ID

Enter the Support Case ID associated with the Logging Assistant record. Contact Technical Support to get this information.

Brief comment to be included in Readme

You can add a brief message of up to 256 characters in the text window. The message is included in the Readme file that is sent with the evidence to Technical Support.

Symantec server location

Enter the Symantec server location that is associated with the Logging Assistant record. Contact Technical Support to get this information.

Click **Next** to initiate the upload operation.

- 5 The **Upload Evidence Results** panel lets you view the results of the evidence upload operation. The Operation Details window of the panel displays the file on the master server where the progress of the operation is recorded.

Click **Finish** to complete the Upload Evidence Wizard.

Administering NetBackup

- [Chapter 29. Management topics](#)
- [Chapter 30. Accessing a remote server](#)
- [Chapter 31. Using the NetBackup Remote Administration Console](#)
- [Chapter 32. Alternate server restores](#)
- [Chapter 33. Managing client backups and restores](#)
- [Chapter 34. Powering down and rebooting NetBackup servers](#)
- [Chapter 35. About Granular Recovery Technology](#)

Management topics

This chapter includes the following topics:

- [Configuring the NetBackup Client Service](#)
- [NetBackup naming conventions](#)
- [Wildcard use in NetBackup](#)
- [Wildcard characters in backup selections for an NDMP policy](#)
- [How to administer devices on other servers](#)
- [How to access media and devices on other hosts](#)

Configuring the NetBackup Client Service

By default, the NetBackup Client Service is configured on Windows with the **Local System** account. The **Local System** account lacks sufficient rights to perform certain backup and restore operations.

For example, for NetBackup to access CIFS volumes, the account must be changed from **Local System** to an account with access to the CIFS share.

To change the NetBackup Client Service logon account on a Windows computer:

- Open the Windows Services application.
- To change the logon account, stop the NetBackup Client Service.
- Open the properties for the NetBackup Client Service.
- Provide the name and password of the account that has the necessary permissions. For example, change the logon to that of *Administrator*.
- Restart the service.

If the logon property is not changed for the NetBackup Client Service, the policy validation fails with status code 4206.

Situations in which the NetBackup Client Service logon account requires changing

The following list contains situations in which the NetBackup Client Service logon account needs to be changed:

- To access CIFS storage for a storage unit.
- To use UNC paths, the network drives must be available to the service account that the NetBackup Client Service logs into at startup. You must change this account on each Windows client that is backed up that contains data that is shared with another computer.
- During a snapshot: To have read access to the share for backup purposes and write access during restores.
The account must be for a domain user that is allowed to access and write to the share. To verify the account, log on as that user and try to access the UNC path. (For example: `\\server_name\share_name`.)
- For database agents and options, configure the service with a logon account that has the necessary permission or privileges. See the documentation for your agent or option for more information.
- For the database agents that support VMware backups on a NetApp disk array, configure the logon account to one that has access to the disk array.

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.

Note: No spaces are allowed.

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 29-1 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> <p><code>/etc/*.conf</code></p>
?	<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> <p><code>c:\system\log??_03</code></p>

Table 29-1 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 that are used as wildcards in include and exclude lists: Windows clients: Allowed UNIX clients: Allowed■ Brackets that are used as wildcards in policy backup selections lists: Windows clients: Not allowed; the use of brackets in policy backup selections lists causes backups to fail with a status 71. UNIX clients: Allowed
{ }	<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 that 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`

Table 29-2 Placement of wildcards in the path of backup selections

Client type	Examples
<p>For Windows clients, wildcards function correctly only when they are placed at the end of the path, in the file or the directory name.</p> <p>See “Pathname rules for Windows client backups” on page 803.</p>	<p>The following example is allowed:</p> <pre>C:\abc\xyz\r*.doc</pre> <p>Wildcard characters do not work elsewhere in the path. For example, an asterisk functions as a literal character (not as a wildcard) in the following examples:</p> <pre>C:*\xyz\myfile</pre> <pre>C:\abc*\myfile</pre>
<p>For UNIX clients, wildcards can appear anywhere in the path.</p> <p>See “Pathname rules for UNIX client backups” on page 810.</p>	<p>The following examples are allowed:</p> <pre>/etc/*/abc/myfile</pre> <pre>/etc/misc/*/myfile</pre> <pre>/etc/misc/abc/*.*</pre>

See [“Backup Selections tab”](#) on page 796.

Wildcard characters in backup selections for an NDMP policy

You can use wildcard characters in regular expressions or the directive `ALL_FILESYSTEMS` to specify path names in NDMP policy backup selections.

Note: Wildcard characters are not supported in releases of NetBackup before version 7.6. To use wildcard characters in NDMP policy backup selections, both your master server and the media servers on which NetBackup for NDMP is installed must run NetBackup 7.6 or later.

Wildcard characters in regular expressions or directives are valid for streaming and non-streaming NDMP backups.

Note: Directory-level expansion is not supported for some NDMP servers. Some NDMP filer vendors do not have the APIs that NetBackup uses to support wildcard characters lower than the volume level.

If you specify a backup selection using wildcard characters lower than the volume level for these filers, status code 106 is generated. The following message is displayed: **Invalid file pathname found, cannot process request**.

Currently, only NetApp filers support wildcard characters for backup selections lower than the volume level. This support is not available in NetApp clustered Data ONTAP version 8.2. To see the versions of NetApp Data ONTAP that support wildcard characters for backup selections lower than the volume level, refer to the NetBackup Master Compatibility List at the following URL:

www.netbackup.com/compatibility

You cannot use any wildcard characters that also match file names. For example, a backup selection might include `/vol/vol_archive_01/autoit*`. This specification might match a path name such as `/vol/vol_archive_01/autoit_01/`. However, if this specification also matches a file name like `/vol/vol_archive_01/autoit-v1-setup.exe`, the backup job fails with status code 99 because wildcards can specify only path names. The following message is displayed: **NDMP backup failure (99)**.

Table 29-3 Valid wildcard characters for NDMP policy backup selections

Wildcard character	Description
*	<p>Specifies a string match. For example:</p> <p><code>/vol/vol_archive_*</code></p> <p>This form of the path specification matches all paths that begin with the literal characters <code>/vol/vol_archive_</code> and end with any characters.</p> <p>The string match wildcard can also specify multiple variable characters between literal characters as in the following examples:</p> <p><code>/vol/ora_*archive</code> or <code>/vol/ora_*archive*</code></p> <p><code>/vol/ora_vol/qtrees_*archive</code> or <code>/vol/ora_vol/qtrees_*archive*</code></p>
?	<p>Specifies a single-character match.</p> <p><code>/fs?</code></p> <p>This path specification matches all paths that begin with the literal characters <code>/fs</code> and end with any single character. For example, <code>/fs1</code>, <code>/fs3</code>, <code>/fsa</code>, <code>/fsd</code> and so on match the specified pattern <code>/fs?</code>.</p>

Table 29-3 Valid wildcard characters for NDMP policy backup selections
(continued)

Wildcard character	Description
[...]	<p>Specifies an alphanumeric pattern match. For example:</p> <pre>/fs[1-9]</pre> <p>This path specification matches all paths that begin with the literal characters <code>/fs</code> and end with any single numeric character from 1 through 9. For example, <code>/fs1</code>, <code>/fs2</code>, and so on up to <code>/fs9</code> match the specified pattern <code>/fs[1-9]</code>. However, <code>/fs0</code> and <code>/fsa</code> do not match the specified pattern; 0 is out of the specified numeric range, and <code>a</code> is a non-numeric character.</p> <p>The pattern match wildcard can also specify alphanumeric patterns such as <code>/fs[1-5a]</code>. This specification matches <code>/fs1</code>, <code>/fs2</code>, and so on up to <code>/fs5</code> as well as <code>/fsa</code>.</p> <p>Similarly, the pattern match wildcard can also specify patterns like <code>/fs[a-p4]</code>. This specification matches <code>/fsa</code>, <code>/fsb</code>, and so on up to <code>/fsp</code> as well as <code>/fs4</code>.</p> <p>You must use multiple backup selection specifications if the pattern can match more than 10 volume names in a numeric series. For example, you may want to back up 110 volumes that begin with the literal characters <code>/vol/ndmp</code> and are numbered 1 through 110. To include these volumes in a backup selection with wildcards, specify three backup selections with the following wildcard patterns:</p> <ul style="list-style-type: none"> ■ <code>/vol/ndmp[0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with a single numeric character 0 through 9. ■ <code>/vol/ndmp[0-9][0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with the two-digit numeric characters 00 through 99. ■ <code>/vol/ndmp[0-9][0-9][0-9]</code> This pattern matches any volume name that begins with <code>/vol/ndmp</code> and ends with the three-digit numeric characters 000 through 999. <p>Do not specify <code>/vol/ndmp[1-110]</code> in this example. This pattern produces inconsistent results.</p>

Note: Symantec recommends that you do not use a single forward-slash character (/) in an NDMP policy backup selection. This method of including all the volumes on an NDMP filer in the selection is not supported. Instead, use the `ALL_FILESYSTEMS` directive:

See [“ALL_FILESYSTEMS and VOLUME_EXCLUDE_LIST directives”](#) on page 820.

Note: Nested wildcard expressions can result in recursive path name expansion operations that can impact performance, especially for directories that have a very large number of files or directories. An example of nested wildcard expansion is as follows:

```
/vol/fome06/*/*private
```

Note: Wildcard expressions do not span or include a path separator (/).

Note: All backup selections that contain a wildcard expression must start with a path separator (/). An example of a correct wildcard expression is as follows:

```
/vol/archive_*
```

An example of an incorrect wildcard expression is as follows:

```
vol/archive_*
```

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.
- Specify the remote server in the host name field and start the NetBackup-Java console.
See [“Accessing remote servers”](#) on page 1112.

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.

If the master 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 master server.

The following procedure assumes that all other steps to add a media server are accomplished:

See [“Adding a media server”](#) on page 372.

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.

The `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's 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
```

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

<http://www.symantec.com/docs/DOC5332>

Accessing a remote server

This chapter includes the following topics:

- [Prerequisites for accessing a remote server](#)
- [Accessing remote servers](#)
- [About running the NetBackup Administration Console on a NetBackup client](#)
- [About troubleshooting remote server administration](#)

Prerequisites for accessing a remote server

In a NetBackup environment, you can use multiple NetBackup servers to perform various administrative tasks.

For example, consider a NetBackup environment that contains two NetBackup domains. Server_1 in one domain and Server_2 in the other. You can log into the NetBackup Administration Console of Server_1 and then change to Server_2 to administer Server_2. You must perform the following tasks before you can change to a different server:

- From one server, allow access to another server. See [“Allow access to another server”](#) on page 1103.
- Authorize users of one server to access another server. See [“Authorize users of one server to access another server”](#) on page 1105.

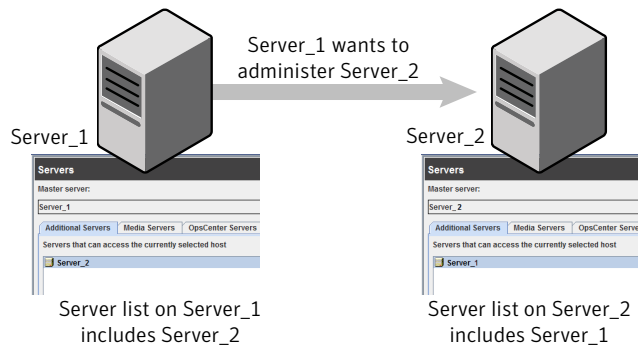
To change to a different server, you can either specify the remote server on the login screen or change to the server from within the NetBackup Administration Console. See [“Accessing remote servers”](#) on page 1112.

Allow access to another server

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. Likewise, the remote host must include the local host in its server list.

Figure 30-1 assumes that server_1 wants to administer server_2.

Figure 30-1 Server lists on local server and remote server




To add a server to a server list

- 1 In the **NetBackup Administration Console**, expand **Host Properties > Master Server**.
- 2 Double-click the name of the server to view the properties.

- 3 Select the **Servers** tab to display the server list.

If the remote server does not appear in the **Additional Servers** list, the current server considers it invalid.

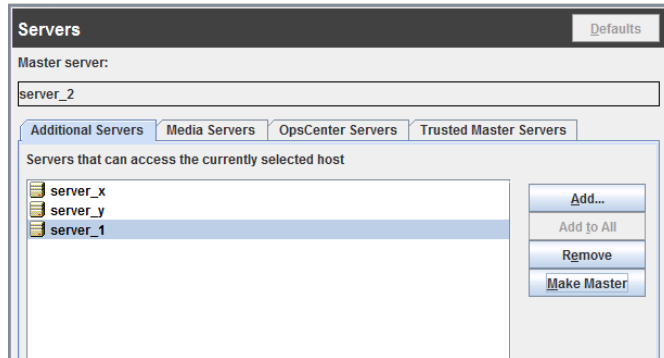


Currently, server_2 allows remote access to two servers: server_x and server_y

- 4 To add a server to the server list, click **Add**.
- 5 In the **New Server** dialog box, type the server name in the field.



- 6 Click **Add** to add the server to the list. Then, click **Close** to close the dialog box. The server appears in the server list.



The `bp.conf` file on every UNIX server contains `SERVER` and possibly `MEDIA_SERVER` entries. The server list in the **Servers** properties dialog box represents these entries. Hosts that are listed as media servers have limited administrative privileges.

- 7 Click **OK** to save the changes.

Authorize users of one server to access another server

You must explicitly authorize users of one server to access another server. Add users of a server to the `auth.conf` file on the server that they are authorized to access. In this example, authorize users of `Server_1` to access `Server_2`. This task must be performed on `Server_2`.

Note: Add only the users that don't have administrative privileges on one server to the `auth.conf` file on another server.

On UNIX, the `auth.conf` file is located at `/usr/opencv/java`.

On Windows, create the `auth.conf` file from the `auth.conf.win.template` file that is located at `C:\Program Files\Veritas\Java`.

See [“Authorization file \(auth.conf\) characteristics”](#) on page 1120.

To authorize users of one server to access another server

- ◆ In the `auth.conf` file on Server_2, add users of Server_1 authorized to access Server_2.

If Server_1 is a Windows host, add the following line in the `auth.conf` file:

For example, `username ADMIN=ALL JBP=ALL`

If Server_1 is a UNIX host, add the following line in the `auth.conf` file:

For example, `root ADMIN=ALL JBP=ALL`

Using the Change Server operation in the NetBackup-Java Administration Console

In previous releases, Windows users administered NetBackup by using any of the following interfaces:

- Windows-based NetBackup Administration Console
- Java-based NetBackup Administration Console

The Windows-based console was removed from NetBackup 7.7. Windows administrators can run either the Java-based console (if the host is a master server) or the Java-based Remote Administration Console (if the host is not a master server).

Because of the differences between how the consoles authenticate users, NetBackup administrators should be aware of how the Change Server operation works when using the Java-based console as compared with the Windows-based console. A number of configuration options are available, depending on the security requirements of the site. (The Change Server operation consists of the administrator selecting **File > Change Server** and selecting a host from the drop-down menu.)

Background information about the Windows Administration Console

The Windows-based NetBackup Administration Console used a host-based authentication model. Any request coming from a mutually trusted console host was processed by the NetBackup server with no further authentication or authorization.

Host-based authentication is based on reverse DNS lookup of the requesting host's IP address. This is then matched against the server list in the master server's configuration file. (On Windows servers this is the registry and on UNIX servers this is the `bp.conf` file.) The host-based trust is created when the server list of a NetBackup host includes the host name of another NetBackup host and, for mutual

trust, vice versa. A host name is added to the server list by using the Server host properties. This process remains the same in NetBackup 7.7 and later.

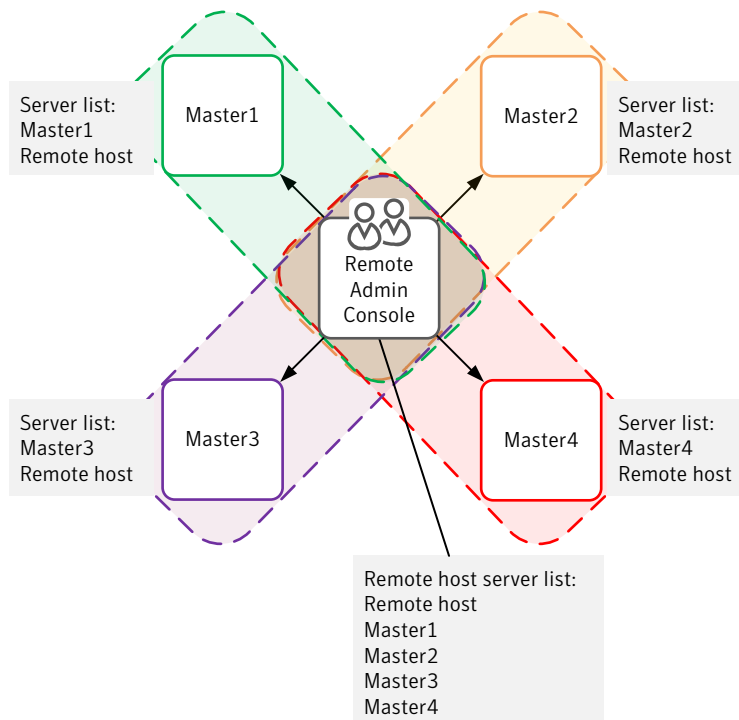
What that meant for the Windows-based console was that the administrator was required to work from a fixed or well-known workstation that was trusted by the NetBackup server. The administrator could connect to the NetBackup server without being asked for any credentials.

How the Change Server operation worked using the Windows-based Administration Console

Since the Windows-based Administration Console used host-based authentication, the Change Server operation was seamless and did not require any user authentication. However, each NetBackup server that was connected to using the Change Server operation was required to extend its host-based trust boundary to include the Administration Console host. There was no need for trust or awareness between the NetBackup servers themselves, but only between the NetBackup server and the Administration Console host. A valid administrator on one NetBackup server might not have had NetBackup administrative access to any other NetBackup server. This is the case in [Figure 30-2](#).

In [Figure 30-2](#), the Windows-based NetBackup Administration Console can access any of the NetBackup master servers because all of the master servers are included in the server list of the Administration Console host. An administrator that logged onto the Master2 host cannot perform NetBackup administration for Master4 because the servers are not mutually trusted. (That is, the Master2 server list does not contain Master4 and the Master4 server list does not contain Master2.)

Figure 30-2 Host-based security model in the Windows Administration Console



The user-based authentication model used by the Java-based Administration Console

The Java-based Administration Console employs a user-based authentication model. In this model, the console authenticates the user by asking for a user name and password. As long as the user credentials are valid and verifiable by the NetBackup master server, the user's operation requests are accepted and processed. The practical impact is that the user must authenticate using the NetBackup master server credentials, but can launch the console on any host.

How the Change Server operation works using the Java-based Administration Console

When the user performs the Change Server operation, the NetBackup server that was initially connected to acts as a proxy for forwarding requests to the other NetBackup server. This requires that the trust boundaries of the NetBackup servers be extended to include one another. As mentioned previously, the trust is created when the server list of one NetBackup server includes the host name of another NetBackup server and vice versa. If there are more than 2 NetBackup servers, then

the trust boundary of each NetBackup server needs to be extended to include other NetBackup servers in the Change Server list.

Combining or extending the trust boundary means that once a user is trusted on a NetBackup server, that user is automatically trusted on the other NetBackup servers within the trust boundary. The more servers in the Change Server list, the larger the trust boundary.

Figure 30-3 shows the most flexible Change Server configuration: A user can log on to any of the NetBackup master servers and perform a Change Server operation to access any other NetBackup master server. The trust boundary is all-inclusive because the server list on each NetBackup server contains all of the other NetBackup servers. In other words, all of the NetBackup servers mutually trust one another.

Figure 30-3 Most flexible model for the Change Server operation

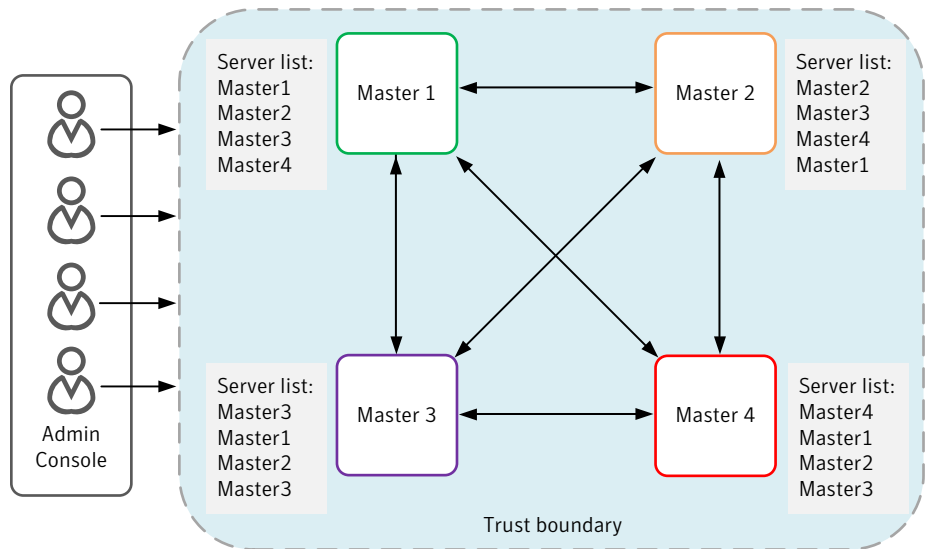


Figure 30-4 shows a more restrictive Change Server configuration: A user must log on to a fixed NetBackup master server, Master1. After successfully connecting to Master1, the user can perform the Change Server operation to access the other NetBackup master servers. In this case, the trust boundary is not all-inclusive. Each NetBackup master server, except Master1, extends its trust boundary to include only Master1. Master1 extends its trust boundary to include all other NetBackup master servers.

Figure 30-4 Restrictive model, offering limited Change Server capability

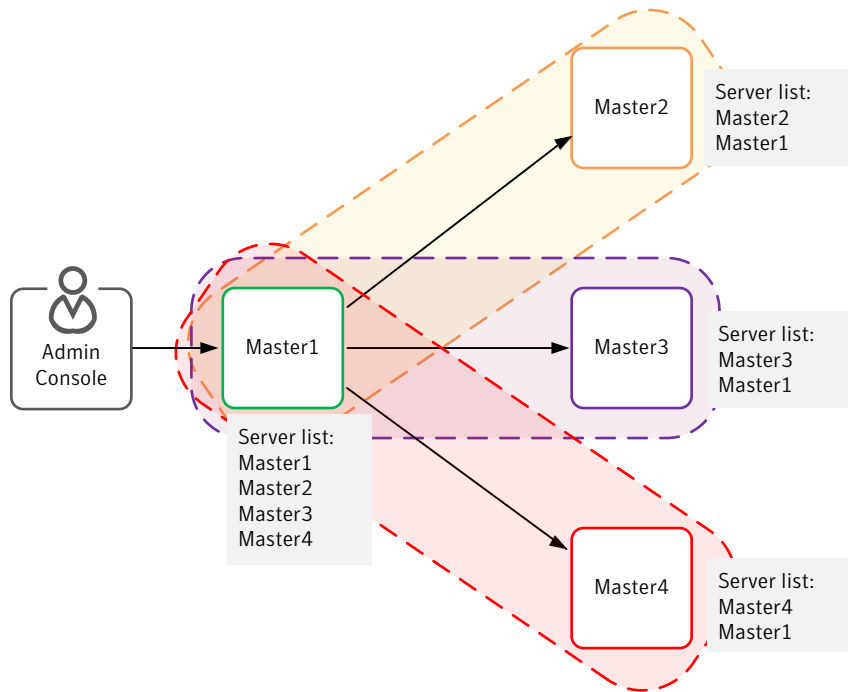



Figure 30-5 shows a model to access multiple NetBackup servers without using the Change Server operation. In this scenario, a user launches separate instances of the Administration Console for each NetBackup server from the same or different hosts. This requires no change to the trust boundary of the NetBackup servers, making this model the most secure and yet flexible.

Figure 30-5 Launch multiple instances of the console: one for each NetBackup server



The following table rates the flexibility and security of the models when configured to use the Change Server operation and when not configured to use it. Note that flexibility and security are on a continuum: the more flexible the model, the less secure it is; the more secure the model, the less flexible it is.

Table 30-1 Comparison of models for flexibility and security concerning the Change Server operation

Attributes	Any-to-any	Many-to-all	One-to-any	Without Change Server
	Any server can use Change Server to access any other server. (Figure 30-3)	Multiple servers, but not all, can use Change Server to access any other server.	Only one server can use Change Server to access any other server. (Figure 30-4)	Change Server is not used. (Figure 30-5)
Flexible	Most flexible	Less flexible: Flexibility increases as the trust boundary grows	Least flexible	More flexible

Table 30-1 Comparison of models for flexibility and security concerning the Change Server operation (*continued*)

Attributes	Any-to-any	Many-to-all	One-to-any	Without Change Server
Secure	Least secure	Less secure: Security decreases as the trust boundary grows	More secure	Most secure

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**.

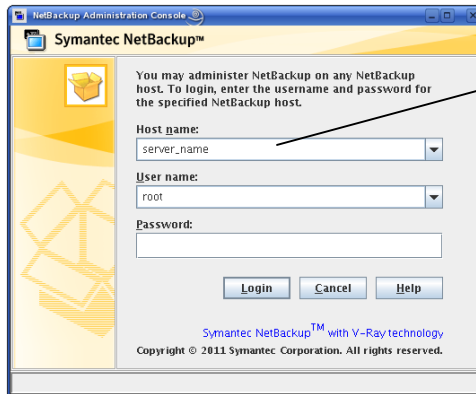
If the server that you want to access is a media server, it must be provisioned with a security certificate.

See [“About security certificates for NetBackup media servers”](#) on page 45.

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 [“Allow access to another server”](#) on page 1103.
- 2 Use one of the following methods to access a remote server:
 - Use the **Change Server** option:
 - Select any main node in the left pane of the **NetBackup Administration Console**. (The option does not appear in the **File** menu unless a main node is selected.)
 - Select the **File > Change Server** menu command.
 - 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.
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*.
If the user has administrative privileges on one server and different privileges on another server, the user is required to reauthenticate.
 - Specify the remote server on the login screen:
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

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 **Remote 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 daemons are running on both hosts.
 - The network connection is valid.
 - The user has administrative privileges on the destination host.
 - To perform a Change Server operation, all servers should be at the same version.
 - The current host is listed in the server list of the destination host.
- See [“Allow access to another server”](#) on page 1103.

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.

If you change the master server list, stop and restart the NetBackup Database Manager (bpdbm) and NetBackup Request Daemon (bprd) to ensure that all appropriate NetBackup processes use the new server entry.

- 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 Remote Administration Console

This chapter includes the following topics:

- [About the NetBackup Remote Administration Console](#)
- [About authorizing NetBackup users](#)
- [Authorization file \(auth.conf\) characteristics](#)
- [About authorizing nonroot users for specific applications](#)
- [About authorizing specific tasks in the Backup, Archive, and Restore user interface](#)
- [About authorizing users on Windows NetBackup server](#)
- [Restricting access to the NetBackup Administration Console on Windows](#)
- [Run-time configuration options for the NetBackup Administration Console](#)
- [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 Remote Administration Console

The **NetBackup Remote Administration Console** is a stand-alone Java-based administration console. This console is installed on a host that does not have NetBackup server software. It is used to monitor NetBackup servers remotely.

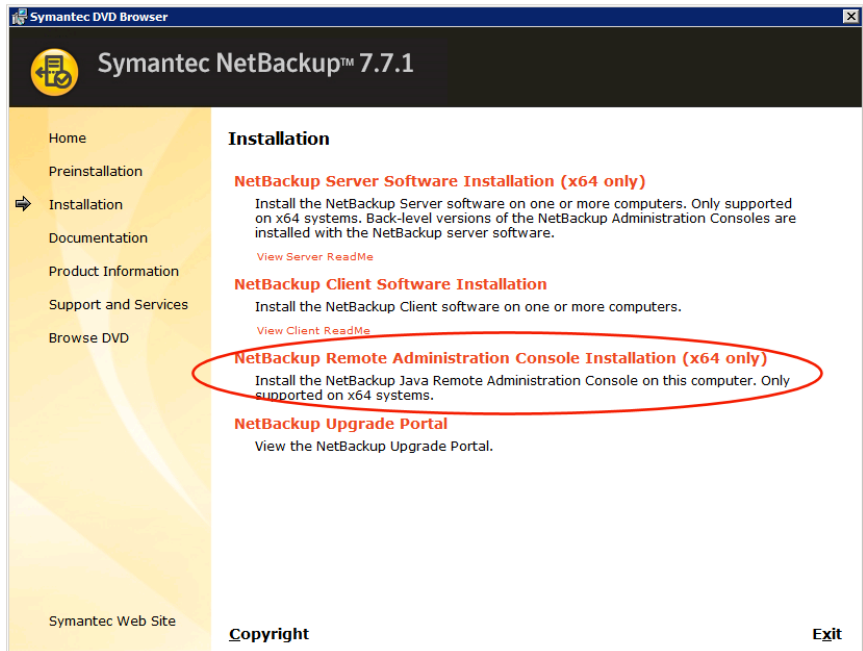
Installing this console installs the **NetBackup Administration Console** and the client software. The presence of the client software enables the computer to be backed up like any other client. No master server software or media server software is installed.

Multiple base versions of the **NetBackup Administration Console** are installed when you install NetBackup server software. For example, when you install or upgrade to NetBackup 7.7.1, base versions from 7.0 through 7.7.1 of the console are all installed.

Note: Symantec recommends that after you install or upgrade NetBackup server software, you should uninstall older versions of the NetBackup Remote Administration Console (Windows and Java) present on the host. If the native **NetBackup Administration Console** for Windows is present, it is automatically uninstalled when you install or upgrade the NetBackup server software.

Note: To upgrade any of the multiple versions of consoles to a patch, you must first install the base version of the Remote Administration Console. Use the installer of the base version to install the Remote Administration Console. You must then upgrade to the corresponding patch of the Remote Administration Console.

Upgrading directly to a patch version of the NetBackup Administration Console from the multiple versions of the consoles is not supported.

Figure 31-1 Installing the NetBackup Remote Administration Console

These processes can be run on two different NetBackup hosts. This distributed application architecture holds true for the **Backup, Archive, and Restore** client interface (`jnbSA`) on UNIX platforms as well.

The administrator first starts the Java-based **NetBackup Administration Console** using one of the following methods:

- Select **Start > Programs > Symantec NetBackup > NetBackup 7.7.1 Administration Console** on the Windows computer on which the NetBackup Remote Administration Console is installed.
- Run the `jnbSA` command on a UNIX computer where NetBackup is installed.

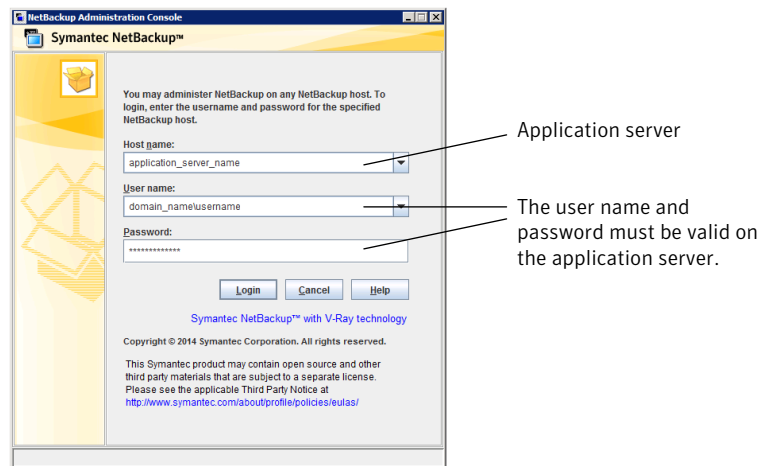
Then the administrator logs on to the application server on the host that is specified in the logon dialog box.

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: 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.

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


Figure 31-2 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).

Note: To perform a Change Server operation, all servers should be at the same version.

Figure 31-3 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 **NetBackup 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 [“Allow access to another server”](#) on page 1103.

See [“Accessing remote servers”](#) on page 1112.

About authorizing NetBackup 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 Administration Console** for specific applications. NetBackup Access Control always takes precedence over the capabilities authorization of the console.

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 logon dialog box.

The `auth.conf` file contains the authorization data for accessing NetBackup applications.

On Windows NetBackup servers, a template file named `auth.conf.win.template` is present at `NB_Install/Java`. Use this template file to create an `auth.conf` file at the same location.

The template file contains an example of giving permissions to a user.

See [“About authorizing nonroot users for specific applications”](#) on page 1123.

On UNIX NetBackup servers, the `auth.conf` file is located at `NB_Install/Java`. This 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.
----------------------	---

To perform remote administration or user operations with `jbpSA`, a user must have valid accounts on the NetBackup UNIX server or client computer.

Note: 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 local administrator group on the NetBackup master server to use all of the NetBackup-Java applications. Other users are allowed to access only **Backup, Archive, and Restore**.

Authorization file (auth.conf) characteristics

On Windows NetBackup servers, a template file named `auth.conf.win.template` is present at `NB_Install/Java`. Use this template file to create an `auth.conf` file at the same location.

The template file contains an example of giving permissions to a user.

On UNIX NetBackup servers, the `auth.conf` file is located at `NB_Install/Java` 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 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 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 applications and the related administrator-related capabilities.

See [“About authorizing nonroot users for specific applications”](#) on page 1123.

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 the Backup, Archive, and Restore user interface”](#) on page 1124.

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 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 computer that is specified in the host field. The NetBackup-Java application server authenticates the user name and password by using the system password file data for the specified computer. The password must be the same password that was used upon logon at that computer.

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.

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 application. This logon method is useful if there is more than one administrator because it saves the state information for each administrator.

Note: NetBackup 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:

ALL	Indicates that the user has administrative privileges for all of the applications that are listed in this table.
AM	Activity Monitor
BMR	Bare Metal Restore
BPM	Backup Policy Management
BAR or JBP	Backup, Archive, and Restore
CAT	Catalog
DM	Device Monitor
HPD	Host Properties
MM	Media Management
REP	Reports
SUM	Storage Unit Management
VLT	Vault Management

For example, to give a user (`user1`) access only to the Device Monitor and Activity Monitor, add the following entry to the `auth.conf` file:

```
user1 ADMIN=DM+AM
```

In order for a nonroot user to modify the files that the **NetBackup-Java Administration Console** uses, run the `nonroot_admin_nbjava` script. The script changes permissions on the following files:

```
/usr/opensv/java/auth.conf  
/usr/opensv/java/Debug.properties  
/usr/opensv/java/nbj.conf
```

Note: `nonroot_admin_nbjava` is located in
`/usr/openssl/java/nonroot_admin_nbjava`.

About authorizing specific tasks in the Backup, Archive, and Restore user interface

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 1149.
 See [“About client-directed restores”](#) on page 1150.
- 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 31-1 Identifiers for the `JBP` keyword in the `auth.conf` file

Identifier	Description
ENDUSER	Allows the users to perform restore tasks from true image 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 users on Windows NetBackup server

To log on to the NetBackup-Java 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 on a computer running 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.

Note: 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 local administrator group on the NetBackup master server to use all of 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 an `auth.conf` file at `NB_Install_path\java` by using the template available at the same location.

See [“About the NetBackup Remote Administration Console”](#) on page 1116.

Restricting access to the NetBackup Administration Console on Windows

Use the following procedure to restrict access to one or more of the NetBackup applications.

To restrict access to one or more of the NetBackup applications

- 1 Create the following file on the Windows system:

```
NB_Install_path\java\auth.conf
```

Note: Use the template file available at the same location.

The template file contains an example of giving permissions to a user.

- 2 Add an entry in `auth.conf` for each user that accesses NetBackup applications. The existence of this file, along with the entries it contains, prohibits unlisted users from accessing NetBackup 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
```

See [“About authorizing NetBackup users”](#) on page 1119.

Run-time configuration options for the NetBackup Administration Console

On Windows computers, the analogous file containing configuration options for the **NetBackup Administration Console** is `Install_path\java\setconf.bat`

On UNIX computers, file `/usr/opensv/java/nbj.conf` contains configuration options for the **NetBackup Administration Console**. Enter one option per line, following the same syntax rules as exist for the `bp.conf` file.

`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 **NetBackup-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 Windows:

```
SET FIREWALL_IN=
HOST1:PORT1=HOST2:PORT2;IP_ADDR1:PORT3=IP_ADDR2:PORT4
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

On UNIX:

```
FIREWALL_IN= HOST1:PORT1=HOST2:PORT2[;...;HOSTn:PORTn=HOSTm:PORTm]
```

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 `NBMaster.symc.com` is in a trusted network, behind a firewall.
- The IP address of `NBMaster.symc.com` is `10.221.12.55`.
- The **NetBackup Java Administration Console** is installed on `localhost`.
- SSH tunnels exist from `localhost` to `NBMaster.symc.com` as follows:

<code>bpjava-msvc</code> port (default 13722)	<code>localhost:port1</code>
<code>vnetd</code> port (default 13724)	<code>localhost:port2</code>
<code>pbx</code> port (default 1556)	<code>localhost:12345</code>

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 [“VNEDD_PORT”](#) on page 1132.

On Windows systems, use `setconf.bat` to add the option:

```
SET FIREWALL_IN=
NBMaster.symc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
SET FIREWALL_IN >> "%NBJDIR%\nbjconf
```

On UNIX systems, add the following line to the `nbj.conf` file:

```
FIREWALL_IN=NBMaster.symc.com:1556=localhost:12345;10.221.12.55:12345=localhost:12345
```

The entry indicates the following:

- The connection to NBMaster.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.

Note: The same options are used if NBMaster.symc.com has a public interface (NBMasterpub.symc.com) that can be reached from the Internet. In this case, the administrator replaces localhost with NBMasterPub.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 all of the interfaces (the **Remote Administration Console**, the **NetBackup 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 **Remote 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 as 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 Activity Monitor in the **NetBackup Administration Console**.

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 **Remote 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 Administration Console** or the **Remote 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.

For more information, see the [NetBackup Installation Guide](#).

The value must be set in the corresponding `setconf.bat` (Windows) or `nbg.conf` (UNIX) configuration option.

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 **Backup, Archive, and Restore** console. This entry also controls how long to keep the log files from the NetBackup-Java console.

The log files are located in `/usr/opensv/netbackup/logs/user_ops/nbjlogs`. The default is three days.

About improving NetBackup-Java performance

The most important factor to consider concerning performance issues for the following interfaces is the platform on which the console runs:

- **Remote Administration Console**
- **NetBackup 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 > Programs > Symantec NetBackup > NetBackup 7.7.1 Administration Console** to start the console.

On supported UNIX platforms, the console is run locally if `jnbSA` or `jbpSA` is entered on the same host on which the console 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 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 computers. For example: the **NetBackup Administration Console** on one computer and the console's application server – `bpjava` processes on another computer.

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 computer 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 computer where the console is running. For example, where the `jnbSA` command is run or the **NetBackup 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 computers. Run NetBackup-Java on a computer 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 computer. (For example, web browsers.) Multiple instances of NetBackup-Java on the same computer have the same effect.
- Run NetBackup-Java on a 1-gigabyte computer 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
```

See [“INITIAL_MEMORY, MAX_MEMORY”](#) on page 1129.

- Consider the amount of physical memory on the relevant computers. Possibly add memory on the host being administered (the console's application server host).
- Consider increasing the swap space to relevant computers:
 - The console host (the host where the console is started)
 - The host being administeredTo increase performance, increase the amount of swap space available to the system where you are running the applications. This is especially true if there is a great deal of other activity on the computer. More swap space can alleviate hangs or other problems that relate to insufficient memory for the applications.
- Consider additional or faster CPUs to relevant computers:
 - 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.

See [“About the jobs database”](#) on page 1045.

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 computer 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 computer 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 computer, 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 31-2 shows the files that contain configuration entries.

Table 31-2 Files containing configuration entries

File	Description
/usr/opensv/java/auth.conf	Authorization options.

Table 31-2 Files containing configuration entries (*continued*)

File	Description
/usr/opensv/netbackup/bp.conf	Configuration options (server and client).
/usr/opensv/java/nbj.conf	Configuration options for the NetBackup-Java Console
/usr/opensv/volmgr/vm.conf	Configuration options for media and device management.
\$HOME/bp.conf	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** check box.

- 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	<p>Select Absolute date and indicate the month and day</p> <p>To begin DST on April 5, set as follows:</p>
Begin DST on the first occurrence of a day in a month	<p>Select First day of week in month. Indicate the day of the week and the month.</p> <p>To begin DST on the first Monday in April, set as follows:</p>
Begin DST on the first occurrence of a day in a month and after a specific date	<p>Select First day of week in month after date. Indicate the day of the week and the month and day.</p> <p>To begin DST on the first Monday after April 5, set as follows:</p>
Begin DST on the last occurrence of a day in a month	<p>Select Last day of week in month. Indicate the day of the week and the month.</p> <p>To begin DST on the last Thursday in April:</p>
Begin DST on the last occurrence of a day in a month and before a specific date	<p>Select Last day of week in month before date. Indicate the day of the week and the month and day.</p> <p>To begin DST before April 30, set as follows:</p>

- 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 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 1147.

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 32-1](#) and [Figure 32-2](#) 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.

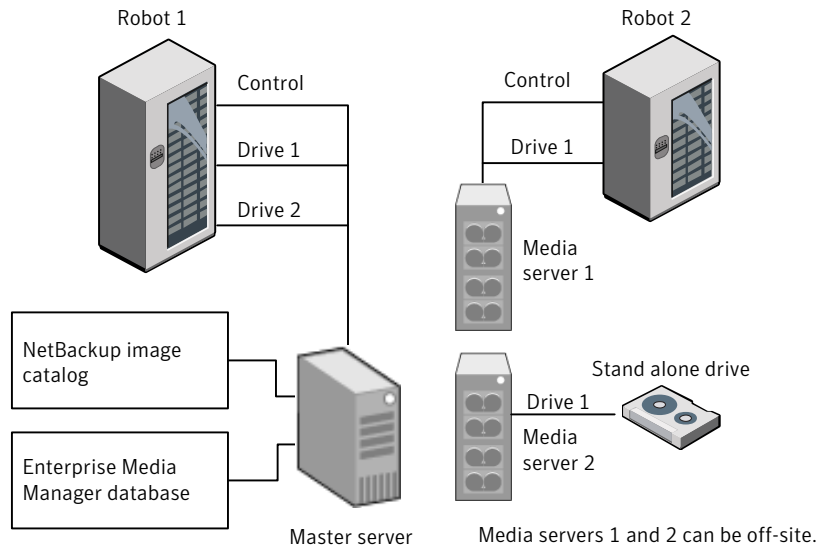
Figure 32-1 NetBackup servers that share robotic peripherals



Assume the following in [Figure 32-1](#):

- 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 32-2 NetBackup servers with separate non-shared peripherals



Assume the following in [Figure 32-2](#):

- 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 1143.
- Override the original server for restores.
See [“Overriding the original server for restores”](#) on page 1144.
- Enable automatic failover to an alternate server

See [“About enabling automatic failover to an alternate server”](#) on page 1146.

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 administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -m media_id
-newserver hostname -oldserver hostname
```

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpmedia -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 administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpimage.exe -id media_id -newserver hostname
-oldserver hostname
```

- As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpimage -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 146.

- 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 that were 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 216.

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 standalone drives of the same type.
When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the master server and `bptm` on the original server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The original server is down.
- The original server is up but `bpcd` on that server does not respond. (For example, if the connection is refused or access is denied.)
- The original server is up and `bpcd` is fine, but `bptm` has problems. (For example, if `bptm` cannot find the required tape.)

Note: The failover uses only the failover hosts that are listed in the NetBackup configuration. By default, the list is empty and NetBackup does not perform the automatic failover.

Failing over to an alternate server

Use the following procedure to enable automatic failover to an alternate server.

To enable automatic failover to an alternate server

- 1 Modify the NetBackup configuration on the master server as follows:
 - By using the **NetBackup Administration Console**:
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/openv/netbackup/bp.conf` file:

```
FAILOVER_RESTORE_MEDIA_SERVERS =  
failed_host host1 host2 ... hostN
```


Where:
`failed_host` is the server that is not operational.
`host1 ... hostN` are the servers that provide failover capabilities.

When automatic failover is necessary for a given server, NetBackup searches through the relevant `FAILOVER_RESTORE_MEDIA_SERVERS` list. NetBackup looks from left to right for the first server that is eligible to perform the restore.

There can be multiple `FAILOVER_RESTORE_MEDIA_SERVERS` entries and each entry can have multiple servers. However, a NetBackup server can be a `failed_host` in only one entry.
- 2 Stop and restart the NetBackup Request daemon on the master server.

Expiring and importing media for alternate server restores

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.

On Windows: 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
```

On UNIX: To identify the media in a specific media spanning group, run the following command as root on the NetBackup master server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpimmedia -spangroups -U -mediaid media_id
```

To display all media in all spanning groups, omit `-mediaid media_id` from the command.

Managing client backups and restores

This chapter includes the following topics:

- [About server-directed restores](#)
- [About client-redirected restores](#)
- [About restoring the files that have Access Control Lists \(ACLs\)](#)
- [About setting the original atime for files during restores on UNIX](#)
- [Restoring the System State](#)
- [About the backup and restore of compressed files on VxFS file systems](#)
- [About backups and restores on ReFS](#)

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:

- On 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** check box.
- On 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.

On UNIX: No progress log is produced if the `bp.conf` file of the requesting server does not contain an entry for the server that performs the restore. (A progress log is an entry in the **Task Progress** tab of the **Backup, Archive, and Restore** client interface.)

On UNIX: Without the entry in the `bp.conf` file (or the server list), the restoring server has no access to write the log files to the requesting server.

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 **Remote Administration Console** to access the Java version of the **Backup, Archive, and Restore** interface from a Windows computer. The **Remote Administration Console** is sometimes referred to as the Windows Display Console.

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:

- On Windows clients:
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 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**.

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.

To do so, first create an `altnames` directory on the NetBackup master server where the backup policy for the clients resides. Place an empty `No.Restrictions` file inside of the directory.

- On Windows:
`Install_path\NetBackup\db\altnames\No.Restrictions`
Do not add a suffix to the files in the `altnames` directory.
- On UNIX:
`/usr/obj/usr/lib/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: The `altnames` directory can present a breach of security, so use it only in limited circumstances. 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.

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.

To do so, create an `altnames` directory on the NetBackup master server where the policy that backed up the other client(s) resides. Place an empty `peername` file inside of the `altnames` directory where `peername` is the client to possess restore privileges.

- On Windows:

```
Install_path\NetBackup\db\altnames\peername
```

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/peername
```

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.

Allowing redirected restores of a specific client's files

The NetBackup administrator can permit a single client to restore the backups that belong to another specific client.

To do so, create an `altnames` directory on the NetBackup master server of the requesting client in the following location:

- On Windows:

```
Install_path\NetBackup\db\altnames\peername
```

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/peername
```

Then, create a `peername` file inside of the directory 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 or the requesting client must be a member of an existing NetBackup policy.

- On Windows: `Install_path\NetBackup\db\images\client_name`
- On UNIX: `/usr/opensv/netbackup/db/images/client_name`

Note: Not all file system types on all computers 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 computer does not support symbolic links nor does it support names greater than 14 characters long. You may want to restore a file to a computer that doesn't support all the features of the computer 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.
- On Windows: *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 the restores that are made by using the command line, not the **Backup, Archive, and Restore** client interface.

Note: On Windows: You must have the necessary permissions to perform the following steps.

On UNIX: You must be a root user for any of the steps that must be performed on the NetBackup server. You may also need to be a root user to make the changes on the client.

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:

- On Windows: Log on to 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
```

- On UNIX: Log on as root on the NetBackup server.

Perform one of the following actions:

- Edit `/usr/openv/netbackup/db/altnames/client1` so it includes the name of *client2*. Or,
- Run the `touch` command on the following file:

```
/usr/openv/netbackup/db/altnames/No.Restrictions
```

Note: The `No.Restrictions` file allows any client to restore files from *client2*.

- Log on to *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.

Figure 33-1 Example restore from a 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:

```
01/29/12 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`.

On Windows: Determine the peer name, then create the following file on the NetBackup master server:

```
Install_path\NetBackup\db\altnames\peername
```

On UNIX: Run the `touch` command on the following file:

```
/usr/opensv/netbackup/db/altnames/peername
```

In this example, the file is:

On Windows:

```
Install_path\NetBackup\db\altnames\client1.dvlp.null.com
```

On UNIX:

```
/usr/opensv/netbackup/db/altnames/client1.dvlp.null.com
```

Edit the *peername* file so that it includes the client names.

For example, if you leave file `client1.dvlp.null.com` empty, *client1*, *client2*, and *client3* can all access the backups that correspond to their NetBackup client name setting.

See [“Allowing a single client to perform redirected restores”](#) on page 1152.

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 [“Allowing redirected restores of a specific client’s files”](#) on page 1152.

Note that this example requires no changes on the clients.

Restore the files.

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 Windows:

- 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\mmddyy.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 actions:
 - 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 **Logging**. Clear the **Keep logs for days** 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.

On UNIX:

- On the NetBackup master server, add the `VERBOSE` entry and a logging level to the `bp.conf` file. For example:

```
VERBOSE = 3
```

- Create the debug log directory for `bprd` by running the following command:


```
mkdir /usr/opensv/netbackup/logs/bprd
```
- On the NetBackup server, stop the NetBackup Request Daemon, `bprd`, and restart it in verbose mode by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdrege -terminate
/usr/opensv/netbackup/bin/bprd -verbose
```

Restart `bprd` to ensure that `bprd` logs information regarding client requests.

- On *client1*, try the file restore.
- On the NetBackup server, identify the peer name connection that *client1* used. Examine the failure as logged on the All Log Entries report or examine the `bard` debug log to identify the failing name combination:


```
/usr/opensv/netbackup/logs/bprd/log.date
```
- Take one of the following on the NetBackup server:

- Enter the following commands:


```
mkdir -p /usr/opensv/netbackup/db/altnames touch
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

These commands allow any client access to *client2* backups by changing its NetBackup client name setting to specify the *client2*.
 - Run the touch command on the following file:


```
/usr/opensv/netbackup/db/altnames/peername
```

The command allows *client1* access to any *client2* backups by changing its NetBackup client name setting to specify *client2*.
 - Add *client2* to the `/usr/opensv/netbackup/db/altnames/peername` file. The addition to the *peername* file allows *client1* access to the backups that were created on *client2* only.
 - On *client1*, change the NetBackup client name setting in the user interface to match what is specified on *client2*.
 - Restore the files to *client1*.
 - Do the following:
 - Delete the `VERBOSE` entry from the `/usr/opensv/netbackup/bp.conf` file on the master server.
 - Delete `/usr/opensv/netbackup/logs/bprd` and the contents.
 - To return the configuration to what it was before the restore, do the following:
 - Delete `/usr/opensv/netbackup/db/altnames/peer.or.hostname` (if it exists)
 - Delete `/usr/opensv/netbackup/db/altnames/No.Restrictions` (if it exists)
 - On *client1*, restore the NetBackup client name setting to its original value.
- See [“Allowing redirected restores of a specific client’s files”](#) on page 1152.
- See [“Allowing a single client to perform redirected restores”](#) on page 1152.

About restoring the files that have Access Control Lists (ACLs)

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.

By default, the NetBackup-modified GNU tar (`/usr/opensv/netbackup/bin/tar`) restores ACLs along with file and directory data.

However, in some situations the ACLs cannot be restored to the file data, as follows:

- Where the restore is cross-platform. (Examples: An AIX ACL restored to a Solaris client or a Windows ACL restored to an HP client.)
- When a `tar` other than the NetBackup modified `tar` is used to restore files.

In these instances, NetBackup stores the ACL information in a series of generated files in the `root` directory using the following naming form:

`.SeCuRiTy.nnnn`

These files can be deleted or can be read and the ACLs regenerated by hand.

Note: If performing an alternate restore where the original directory was ACL-enabled, the alternate restore directory must also be ACL-enabled. If the alternate restore directory is not ACL-enabled, the restore is not successful.

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 standalone 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**.

About setting the original atime for files during restores on UNIX

During a restore, NetBackup sets the `atime` for each file to the current time by default. You can elect to have NetBackup set the `atime` for each restored file to the value the file had when it was backed up. To do so, create the following file on the client:

```
/usr/opensv/netbackup/RESTORE_ORIGINAL_ETIME
```

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.
- 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**, either `bpinetd` on Windows or `inetd` on UNIX, 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 To restore a system state backup using an incremental backup, select the full backup and one or more differential-incremental or cumulative-incremental backups.



- 6 From the **Actions** menu, select **Restores**.
- 7 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.

- 8 Click **Start Restore**.

- 9 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.

- 10 Restart the system before performing subsequent restore operations.

If you booted into **Directory Services Restore Mode** on a domain controller, restart into normal mode when the restore is complete.

About the backup and restore of compressed files on VxFS file systems

NetBackup can back up and restore VxFS-compressed files, maintaining the compression state when the target volume supports file system compression. Future releases will expand this capability to other file systems.

Upon backup of files on a VxFS file system, a message displays in the **Activity Monitor** whenever NetBackup encounters a compressed file:

```
Compress flag found for 'file_name'.
```

Upon restore, NetBackup restores the files to a VxFS file system in their compressed form.

If the restore is to a non-VxFS file system, NetBackup restores the files in an uncompressed form. The following message displays in the **Progress** tab of the **Backup, Archive, and Restore** client interface:

```
File 'file_name' will not be restored in compressed form. Please refer to the Release Notes or User Guide.
```

The message appears only for the first file that cannot be restored in its compressed form.

Note: The compression messages display if the verbose level is 1 or greater.

About backups and restores on ReFS

Microsoft Resilient File System (ReFS) support in NetBackup is automatic and requires no additional configuration.

To restore optimized backups, use the **Restore from Optimized Backup** in the **Backup, Archive, and Restore** interface. You can select individual files for restore.

NetBackup does not support a redirected restore of a Microsoft Resilient File Systems (ReFS) file system.

[Table 33-1](#) lists the ReFS-to-NTFS backup and restore combinations and the success of each.

Table 33-1 ReFS backup and restore

Between file systems	Backups	Restores
ReFS to ReFS	Successful	Successful
ReFS to NTFS	Successful	Successful
NTFS to ReFS	Successful	Limited success For successful restores: <ul style="list-style-type: none">Restore NTFS backups to NTFS file system.Remove all non-supported ReFS items.Copy the files to an ReFS file system.

Powering down and rebooting NetBackup servers

This chapter includes the following topics:

- [Powering down and rebooting NetBackup servers](#)
- [Shutting down and starting up all NetBackup services and daemons](#)
- [Rebooting a NetBackup server](#)
- [Rebooting a NetBackup media server](#)
- [About displaying active processes with `bpps` on UNIX](#)
- [About displaying robotic processes with `vmops` on UNIX](#)

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**. Click the **Jobs** tab and make sure that no jobs are running.
- 2 Click the **Daemons** tab and right-click the NetBackup Request daemon, `bprd`.
Select **Stop daemon** to stop additional job activity and to let current activity end.
- 3 Right-click any daemons that are still running and select **Stop daemon**.

- 4 From the command line, run:

On Windows:

```
Install_path\VERITAS\NetBackup\bin\admincmd\bprdreq -terminate
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

bprdreq does not run on a media server.

- 5 Run the system shutdown command.

The installation process copies the appropriate startup and shutdown script from `/usr/opensv/netbackup/bin/goodies` to `/init.d` and creates links to it from the appropriate `/rc` directory.

Use system startup scripts to begin the Media Manager and NetBackup daemons when the system boots up. Use shutdown scripts to terminate the daemons at system shutdown.

The [NetBackup Installation Guide](#) contains more information about the startup and shutdown scripts.

- 6 On Windows, run:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- 7 Power down the server.

Shutting down and starting up all NetBackup services and daemons

To shut down and start all NetBackup services and daemons, enter the following commands from a command line:

On Windows:

- To shut down all NetBackup services:

```
Install_path\VERITAS\NetBackup\bin\bpdown
```

- To startup all NetBackup services:

```
Install_path\VERITAS\NetBackup\bin\bpup
```

On UNIX:

- To shut down all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- To startup all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.start_all
```

All open Java NetBackup Administration Console sessions need to be restarted and reconnected if NetBackup services are restarted or if a NetBackup server has been rebooted.

Rebooting a NetBackup server

Use the following procedure to reboot a NetBackup server.

To reboot a NetBackup master server

- 1 Restart the system.
- 2 On Windows: 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.
- 3 On UNIX: Ensure that `bprd`, `bpdgm`, and `vmd` are up by running the following script:

```
/usr/opensv/netbackup/bin/bpps -a
```

- 4 On UNIX: Start all NetBackup daemons:

```
/usr/opensv/netbackup/bin/bp.start_all
```

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 On Windows: 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.

- 3 On UNIX: Start `ltid` if it is not already running:

From the **NetBackup Administration Console**, do the following:

- Click **Activity Monitor**, then select the **Processes** tab.
- Right-click `ltid` and select **Start Daemon**.

- 4 On UNIX: From the command line, run:

```
/usr/opensv/volmgr/bin/ltid
```

About displaying active processes with bpps on UNIX

NetBackup provides the `bpps` script to determine which NetBackup processes are active on a UNIX system.

`bpps` is located in the following directory:

```
/usr/opensv/netbackup/bin/bpps
```

The following is example output:

```
root  310 0.0  0.0  176  0 ?  IW Oct 19  15:04 /usr/opensv/netbackup/bin/bpdbm
root  306 0.0  0.0  276  0 ?  IW Oct 19  2:37 /usr/opensv/netbackup/bin/bprd
```

To prevent `bpps` from displaying processes you do not want to check, add the processes to an exclude list. See the comments within the script for more information.

To display both NetBackup and Media Manager options, run:

```
/usr/opensv/netbackup/bin/bpps -a
```

About displaying robotic processes with vmps on UNIX

The `vmps` script shows the Media Manager daemons and robotic processes that are active on a UNIX system.

To run this script, use the following command:

```
/usr/openv/volmgr/bin/vmps
```

In the following sample, the second column contains the process IDs for the processes.

root	303	0.0	0.2	136	264	?	S	Feb 11	4:32	ltid	-v
root	305	0.0	0.0	156	0	?	IW	Feb 11	0:54	vmd	-v
root	306	0.0	0.0	104	0	?	IW	Feb 11	0:15	tl8d	-v
root	307	0.0	0.0	68	56	?	S	Feb 11	12:16	avrd	
root	310	0.0	0.0	116	0	?	IW	Feb 11	0:07	tl8cd	-v

The status for the `nbemm` command is not shown in the output of `vmps`. The `nbemm` status is shown in the output of the `bpps` script.

About Granular Recovery Technology

This chapter includes the following topics:

- [About installing and configuring Network File System \(NFS\) for Active Directory Granular Recovery](#)
- [About configuring Services for Network File System \(NFS\) on Windows 2012 and Windows 2012 R2](#)
- [About configuring Services for Network File System \(NFS\) on Windows 2008 and Windows 2008 R2](#)
- [Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology \(GRT\)](#)
- [Configuring a different network port for NBFSD](#)

About installing and configuring Network File System (NFS) for Active Directory Granular Recovery

NetBackup Granular Recovery leverages Network File System, or NFS, to read individual objects from a database backup image. Specifically, the NetBackup client uses NFS to extract data from the backup image on the NetBackup media server. The NetBackup client uses “Client for NFS” to mount and access a mapped drive that is connected to the NetBackup media server. The NetBackup media server handles the I/O requests from the client through NBFSD.

NBFSD is the NetBackup File System (NBFS) service that runs on the media server. NBFSD makes a NetBackup backup image appear as a file system folder to the NetBackup client over a secure connection.

Network File System, or NFS, is a widely recognized, open standard for client and server file access over a network. It allows clients to access files on dissimilar servers through a shared TCP/IP network. NFS is typically bundled with the host operating system. NetBackup uses Granular Recovery Technology (GRT) and NFS to recover the individual objects that reside within a database backup image, such as:

- A user account from an Active Directory database backup
- Email messages or folders from an Exchange database backup
- A document from a SharePoint database backup

Multiple NetBackup agents that support GRT (for example, Exchange, SharePoint, and Active Directory) can use the same media server.

About configuring Services for Network File System (NFS) on Windows 2012 and Windows 2012 R2

To restore individual items from the Active Directory, you must configure Services for NFS on the NetBackup media server and all Active Directory domain controllers or ADAM/LDS hosts.

Table 35-1 Configuring NFS on Windows 2012 and Windows 2012 R2

Step	Action	Description
Step 1	Configure NFS on the media server.	<p>On the media server do the following:</p> <ul style="list-style-type: none"> ■ Stop and disable the ONC/RPC Portmapper service, if it exists. ■ Enable NFS. See “Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 media server” on page 1172. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1185. ■ Stop the Client for NFS service. See “Disabling the Client for NFS on the media server” on page 1183. <p>Note: If the Active Directory domain controller or ADAM/LDS host resides on the media server, do not disable the Client for NFS.</p> <ul style="list-style-type: none"> ■ Configure the portmap service to start automatically at server restart. Issue the following from the command prompt: <pre>sc config portmap start= auto</pre> This command should return the status [SC] ChangeServiceConfig SUCCESS.

Table 35-1 Configuring NFS on Windows 2012 and Windows 2012 R2
(continued)


Step	Action	Description
Step 2	Configure NFS on all Active Directory domain controllers or ADAM/LDS hosts.	<p>On all Active Directory domain controllers or ADAM/LDS hosts, do the following:</p> <ul style="list-style-type: none"> ■ Enable NFS on the clients. See “Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 client” on page 1176. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1185.

Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 media server

To restore individual items from a backup that uses Granular Recovery Technology (GRT), you must enable Services for Network File System (NFS). When this configuration is completed on the media server, you can disable any unnecessary NFS services.


To enable Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 media server

- 1 Open the Server Manager.
- 2 From the **Manage** menu, click **Add Roles and Features**.




- 3 In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.

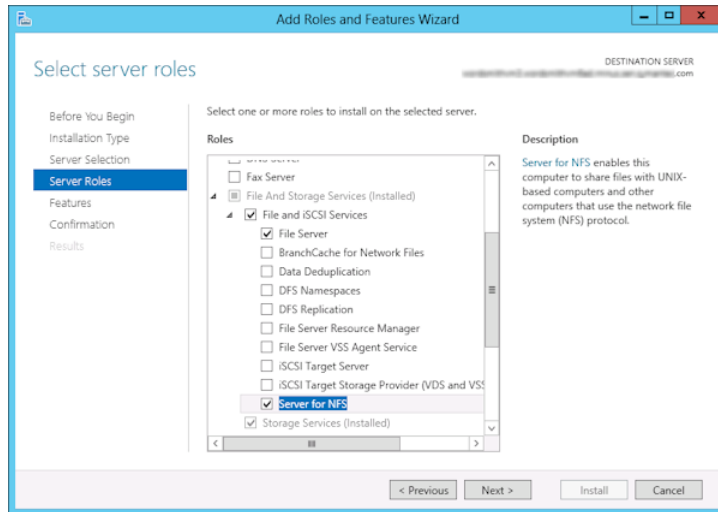
- 4 On the **Select installation type** page, select **Role-based or feature-based installation**.




- 5 Click **Next**.
- 6 On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.



- 7 On the **Server Roles** page, expand **File and Storage Services** and **File and iSCSI Services**.
- 8 Click **File Server** and **Server for NFS**. When you are prompted, click **Add Features**. Click **Next**.



- 9 If the media server is also an Active Directory domain controllers or ADAM/LDS host, on the **Features** page, click **Client for NFS**. Click **Next**.



10 On the **Confirmation** page, click **Install**.

11 Disable any unnecessary services, as follows:


- 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 service.
See [“Disabling the Server for NFS”](#) on page 1185.
- For a host that is only the NetBackup media server, you can disable the Server for NFS and the Client for NFS services.
See [“Disabling the Server for NFS”](#) on page 1185.
See [“Disabling the Client for NFS on the media server”](#) on page 1183.

Enabling Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 client

To restore individual items from a backup that uses Granular Recovery Technology (GRT), you must enable Services for Network File System (NFS). When this configuration is completed all the Active Directory domain controllers or ADAM/LDS hosts, you can disable any unnecessary NFS services.


To enable Services for Network File System (NFS) on a Windows 2012 or Windows 2012 R2 client

- 1 Open the Server Manager.
- 2 From the **Manage** menu, click **Add Roles and Features**.




- 3 In the Add Roles and Features Wizard, on the **Before You Begin** page, click **Next**.

- 4 On the **Select installation type** page, select **Role-based or feature-based installation**.




- 5 Click **Next**.
- 6 On the **Server Selection** page, click **Select a server from the server pool** and select the server. Click **Next**.



- 7 On the **Server Roles** page, click **Next**.

- 8 On the **Features** page, click **Client for NFS**. Click **Next**.



- 9 On the **Confirmation** page, click **Install**.

About configuring Services for Network File System (NFS) on Windows 2008 and Windows 2008 R2

To restore individual items from the Active Directory, you must configure Services for NFS on the NetBackup media server and all Active Directory domain controllers or ADAM/LDS hosts.

Table 35-2 Configuring NFS in a Windows 2008 or Windows 2008 R2 environment


Step	Action	Description
Step 1	Configure NFS on the media server.	<p>On the media server do the following:</p> <ul style="list-style-type: none"> ■ Stop and disable the ONC/RPC Portmapper service, if it exists. ■ Enable NFS. See “Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2” on page 1181. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1185. ■ Stop the Client for NFS service. See “Disabling the Client for NFS on the media server” on page 1183. Note: If the Active Directory domain controller or ADAM/LDS host resides on the media server, do not disable the Client for NFS. ■ Configure the portmap service to start automatically at server restart. Issue the following from the command prompt: <pre>sc config portmap start= auto</pre> This command should return the status [SC] ChangeServiceConfig SUCCESS.
Step 2	Configure NFS on all Active Directory domain controllers or ADAM/LDS hosts.	<p>On all Active Directory domain controllers or ADAM/LDS hosts, do the following:</p> <ul style="list-style-type: none"> ■ Enable NFS. See “Enabling Services for Network File System (NFS) on Windows 2008 or Windows 2008 R2” on page 1181. ■ Stop the Server for NFS service. See “Disabling the Server for NFS” on page 1185.
Step 3	Install the hotfix for Client for NFS on all Active Directory domain controllers or ADAM/LDS hosts.	<p>On all Active Directory domain controllers or ADAM/LDS hosts, install the hotfix for Client for NFS. The hotfix is available at the following location: http://support.microsoft.com/kb/955012</p> <p>Note: Important Windows Vista hotfixes and Windows Server 2008 hotfixes are included in the same packages. However, the hotfix Request page lists only <i>Windows Vista</i>. To request the hotfix package that applies to one or both operating systems, select the hotfix that is listed under <i>Windows Vista</i> on the page. Always refer to the <i>Applies To</i> section in articles to determine the actual operating system for each hotfix.</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 (NFS). 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**.



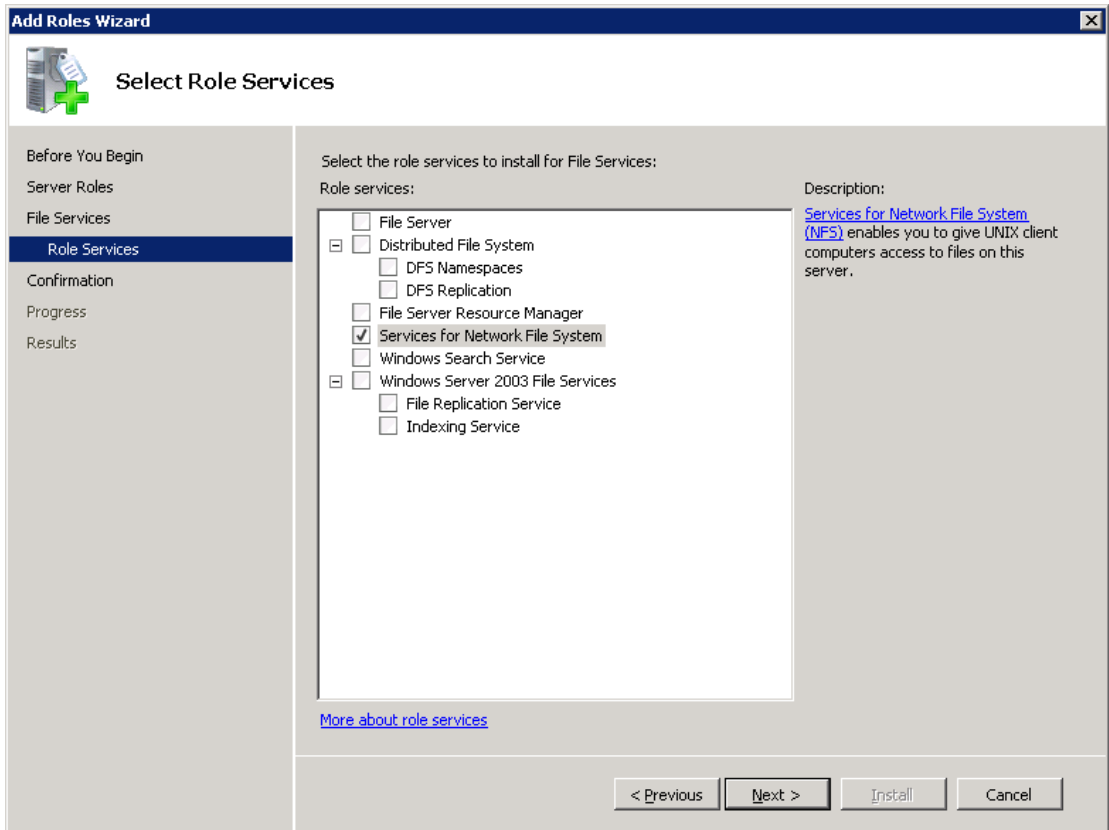
- 3 In the Add Roles Wizard, on the **Before You Begin** page, click **Next**.

- 4 On the **Select Server Roles** page, under **Roles**, check the **File Services** check box. Click **Next**.

Note: If a role service is already installed for the File Services role, you can add other role services from Roles home page. In the File Services pane, click **Add Role Services**.



- 5 On the **Files Services** page, click **Next**.
- 6 On the **Select Role Services** page, do the following:
 - Uncheck **File Server**.
 - Check **Services for Network File System**.
 - Click **Next** and complete the wizard.



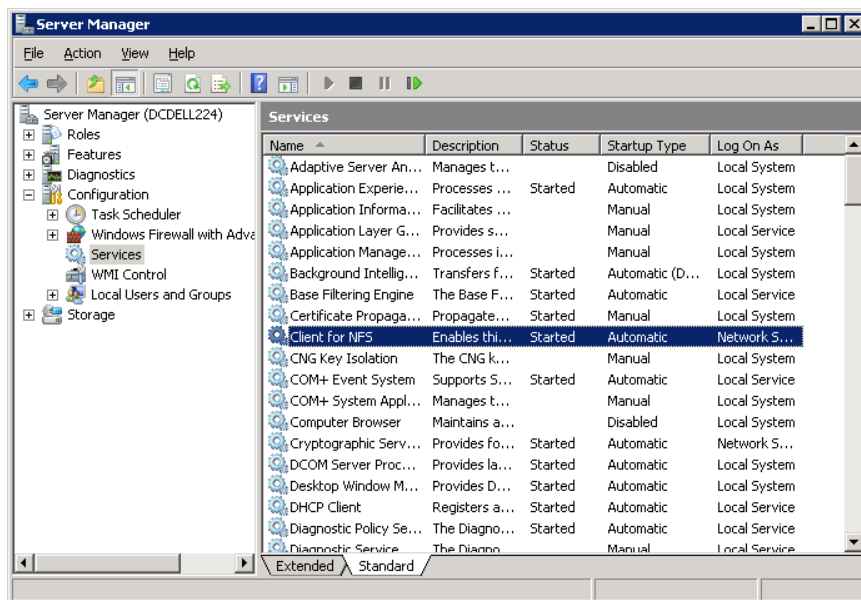
- 7 For each host in your configuration, choose from one of the following:
- If you have a single host that functions as both the media server and the Active Directory domain controllers or ADAM/LDS host, you can disable the Server for NFS.
 - For a host that is only the NetBackup media server, you can 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.

Configuring a UNIX media server and Windows clients for backups and restores that use Granular Recovery Technology (GRT)

To perform backups and restores that use Granular Recovery Technology (GRT), perform the following configuration if you use a UNIX media server and Windows clients:

- Confirm that your media server is installed on a platform that supports granular recovery.
See the [Software Compatibility List](#).
- No other configuration is required for the UNIX 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 a Windows 2012 or Windows 2012 R2 media server”](#) on page 1172.
 See [“Enabling Services for Network File System \(NFS\) on a Windows 2012 or Windows 2012 R2 client”](#) on page 1176.
 See [“Enabling Services for Network File System \(NFS\) on Windows 2008 or Windows 2008 R2”](#) on page 1181.
- You can configure a different network port for NBFSD.
 See [“Configuring a different network port for NBFSD”](#) on page 1188.

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 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`

Index

Symbols

- .ds files 559
- .f files in catalog 616, 618, 868–869
- .lck catalog files 868
- .SeCuRiTty.nnnn files 1159

A

A.I.R. *See* Auto Image Replication

Absolute pathname to directory/volume storage unit setting 561

Accelerator

- and the NetBackup catalog 723
- configuring 723
- description 717
- does not reset atime 721, 752
- Expire after copy retention type 635, 852
- Forced rescan schedule attribute 105, 352
- forced rescan schedule attribute 755
- logs for 725
- messages in the backup job details log 724
- notes and requirements 720
- policy attribute 716
- replacing image size with network-transferred data in command output 734
- reporting network-transferred data in command output 731

Accept connections on non reserved ports property 252, 263

Access Control

- auditing user information without enabling 329
- authorizing users 1120, 1125
- changing DBA password 941
- host properties 67–70, 72, 331
- to a server or client 361
- utility in Administration Console 44
- within the NetBackup Administration Console 1119

Access Control Lists (ACLs) 811, 816, 859, 1158

Access Control Lists (ACLs) for GPFS 351–352

Access Management utility 44

ACS robot type 393

ACS, LSM, Panel, Drive tape drive configuration option 415

ACSLs host robot configuration option 406

Active Directory

- Application Mode (ADAM) 844–845, 848
- granular recovery 843–845
- host properties 73–74
- restoring objects 846

Active secondary operations 642

Activity Monitor

- about the utility 1017
- bpdbjobs command 1049
- BPDBJOBS_OPTIONS environmental variable 1047
- canceling uncompleted jobs 1023
- copy job information 1026
- Daemons tab 1027
- deleting completed jobs 1023
- detailed job status 1023, 1038
- disabling job logging 183
- displaying Backup jobs 610, 612
- displaying Duplication jobs 614
- displaying Import jobs 615
- displaying Index From Snapshot jobs 617
- displaying Replication jobs 620
- displaying Snapshot jobs 623
- filtering column heads 1022
- problem displaying media servers 1038
- restarting a completed job 1024
- resuming suspended jobs 1025
- saving job data to a file 1025
- setup options 1019
- starting and stopping daemons 1038
- stopping jobs 1023
- suspending a job 1024
- using the Troubleshooter 50

ADAM (Active Directory Application Mode). *See* Active Directory Application Mode (ADAM)

adjust time zone 1137

administrator

- account for NetBackup Client Service 1092
- email address property 153

- administrator (*continued*)
 - nonroot 1123
 - AdvancedDisk disk storage units 555–556
 - All log entries report 168, 1058
 - ALL_FILESYSTEMS 820
 - Allow backups to span tape media property 182
 - Allow client browse property 84
 - Allow client restore property 84
 - Allow media overwrite property 181, 262
 - Allow multiple data streams
 - set policy attribute 711
 - when to use 712
 - Allow multiple retentions per media property 182, 768
 - Allow server file writes property 61, 251, 349
 - ALLOW_MEDIA_OVERWRITE 262
 - ALLOW_NON_RESERVED_PORTS 262
 - alternate
 - client restores 341
 - media types 446
 - read server for SLP operations 645
 - server, expiring and importing media for 1147
 - server, failing over to 1146
 - Alternate restore failover machines host
 - properties 217
 - altnames file
 - No.Restrictions file 1151
 - peername file 1152
 - Always property in Fibre Transport host properties 141
 - Announce DHCP interval property 190
 - ANSI format 181
 - AOS/VS format 181
 - API robots 465, 476
 - archive bit 105–106, 746, 750, 857
 - archive jobs 101, 355
 - asterisk as wildcard 1094
 - atime 101, 349, 364, 752, 885, 1160
 - auth.conf file
 - capabilities identifiers 1124
 - description 1120
 - entries for specific applications 1123
 - overview 1119
 - authentication certificates. *See* security certificates
 - Authentication Domain tab 67–68
 - AUTHENTICATION_DOMAIN 263
 - Authorization Service tab 67, 69
 - AUTHORIZATION_SERVICE 265
 - Auto create A.I.R. Import SLP (SLP parameter) 243
 - Auto Image Replication
 - about 975
 - Auto Image Replication (*continued*)
 - and trusted master servers 987
 - Backup operation in source domain 976
 - Maximum size per A.I.R. replication job (SLP parameter) 241
 - nbstserv 993
 - no disk spanning support 976
 - Non-targeted 1000–1001, 1006
 - overview 975
 - role of Backup operation in source domain 609
 - role of Import operation 614
 - role of Replication operation 619
 - supported versions 975
 - synchronizing clocks of master servers 976
 - Target master option 620, 644
 - Target retention type 638
 - Targeted 232, 620, 976, 987, 1000–1001, 1008
 - topology of storage 980
 - TRUSTED_MASTER 327
 - using the Make Master option 237
 - Auto log off timeout option 46, 176
 - Auto snapshot type 742
 - Auto-discovery streaming mode 828
 - AUTO_ADD_ALL_ALIASES_FOR_CLIENT 258, 266
 - Automated Cartridge System
 - drive information 415
 - automatic
 - failover to an alternate server 1146
 - path correction 434
 - Volume Recognition service (avrd) 1039
 - Automatic Volume Recognition process (avrd)
 - and the NetBackup Device Manager 438
 - automounted directories 697
 - avrd (Automatic Volume Recognition process) 1039
- ## B
- Backup
 - end notify timeout property 249, 276
 - job 610, 612
 - network drives policy attribute 697
 - operation 608, 641
 - start notify timeout property 248, 277
 - Backup From Snapshot operation 610, 618
 - Backup option for log files during full backups
 - property 126
 - Backup Selections list
 - ALL_FILESYSTEMS directive 820
 - VOLUME_EXCLUDE_LIST directive 820
 - wildcard characters 1096

Backup, Archive, and Restore 1124

backups

- activating policy 695
- Client backups report 1058
- creating copies 654, 762
- cumulative Incremental Backup 746
- differential incremental backups 746
- duplicating 962
- duration of schedule window 778–779
- expiring 969
- frequency setting 758
- full 745
- how NetBackup determines when files are due 750–751
- importing 970
- manual 843
- NetBackup database files 929, 939
- network drives 697
- NFS mounted files 680, 696
- off-site storage 768
- policy types 745
- raw partitions on Windows 691, 806
- replacing Accelerator image size with network-transferred data in command output 734
- reporting network-transferred data for Accelerator 731
- selections list
 - pathname rules 803, 806–807, 810, 818
 - verifying 799
- send email notification about 155, 158, 160
- Status of Backups report 1058
- status report 168
- user archive 747
- user directed 746
 - schedules 753
- verifying 959
- windows 776

Bandwidth host properties 75–77, 299

barcodes 500

- actions for 466
- adding media type mapping entries 483
- advantages of 464
- best practices 465
- checking 464
- comparing robot contents with volume 521
- configuring rules 467
- rules 466
- settings 470

Bare Metal Restore (BMR) 108, 707, 870, 904, 1021, 1028

Bare Metal Restore service 1039

basic disk staging

- creating a storage unit 578
- Final destination media owner 587
- Final destination storage unit 587
- Final destination volume pool 587
- priority of duplication jobs 581
- relocation schedule 573, 585, 745
- storage units
 - checkpoint restart 691
 - low water mark 564
 - size recommendations 582
- Use alternate read server attribute 581, 588

BasicDisk storage units 644

- configuring 389
- credentials for CIFS 556, 1030–1031, 1040–1041, 1043
- media description 555
- spanning within storage unit groups 598

BE-MTF1 format 181

bind mount on Linux 831

BitLocker Drive Encryption 121

BLAT mail utility 160

block level incremental backups 691, 741

Blowfish cipher encryption 344

bmrbd (NetBackup Bare Metal Restore Boot Server Service) 1028

bmrbd (NetBackup Bare Metal Restore Master Server) 1028

BMRD (NetBackup Bare Metal Restore Master Server) 1039

BMRDB.db

- configuration entry 915
- in catalog 870
- relocating 907, 941

bmrpxeserver (NetBackup BMR PXE Service) 1028

bp.conf file

- about entries 259
- configuring 62
- configuring to use ctime 753
- customizing jnbSA and jbpSA 1132
- entries for Activity Monitor 1046
- indicating database location 904
- NetBackup-Java Administration Console
 - configuration entries 1126
- personal file
 - for UNIX nonroot user 259, 366

- bp.conf file (*continued*)
 - personal file (*continued*)
 - for UNIX root user 259
- BPARCHIVE_POLICY 258, 334, 753
- BPARCHIVE_SCHED 258, 334, 753
- bpbackup command 883
- BPBACKUP_POLICY 258, 335, 753
- BPBACKUP_SCHED 258, 336, 753
- BPBRM logging property 170
- BPBRM_VERBOSE 267
- bpcatarc command 892
- bpcatlist command 892
- bpcatres command 893
- bpcatrm command 893
- bpcd client process 557
- BPCD connect-back property 92, 145
- bpcd daemon 1028, 1039
- BPCD port setting on client 189
- BPCDD (NetBackup Client Service) 1028
- bpchangeprimary command 961
- BPCOMPATD (NetBackup Compatibility Service) 1029, 1039
- bpconfig command 829
- bpdjobs
 - adding entries to bp.conf file 1046
 - adding registry key 1046
 - command 1049
 - debug log 1049
- BPDBJOBS_COLDEFS 258, 268
- BPDBJOBS_OPTIONS environmental variable 1047
- BPDBM (NetBackup Database Manager) 941, 1029, 1039
- BPDBM logging property 170
- BPDBM_VERBOSE 272
- BPDM logging property 170
- bpend_notify script 249
- BPEND_TIMEOUT 276
- bpexpdate command 966
- bpgetconfig 62
- BPINETD (NetBackup Legacy Client Service) 1030
- bpinetd client process 557, 1040
- bpjava-msvc service 1040
- bpjava-susvc service 1040
- bpjobd (NetBackup Jobs Database Management daemon) 941
- BPJOB logging property 170
- bpjobd process 1040
- bpsps script 1168

- BPRD
 - logging property 170
 - NetBackup Request Daemon 1031
 - NetBackup Request Daemon, with NBAC 941
 - port setting on client 189
 - process 1040
- BPRD_VERBOSE 273
- bpsetconfig 62, 259
- bpstart_notify script 248–249, 277, 280
- BPSTART_TIMEOUT 248–249, 276–277, 280
- bpstinfo command 982
- bpsynth log 859
- bptestnetconn utility 205
- BPTM logging level property 170, 275
- BPTM_VERBOSE 274
- Browse and restore ability property 89
- buffer size for data transfer 106
- Busy action property 79
- Busy file host properties 78–79, 337–340
- BUSY_FILE_ACTION 337
- BUSY_FILE_DIRECTORY 338
- BUSY_FILE_NOTIFY_USER 78, 339
- BUSY_FILE_PROCESSING 78, 340

C

- cacheefs file systems, excluding from backup 831
- Calendar schedule type 757, 781
- canceling uncompleted jobs 1023
- Capacity managed retention type 633–634
- case sensitivity
 - in exclude list 130
 - in NetBackup names 1093
 - in storage unit and disk pool names 584
 - in storage unit names 590, 592
 - in volume pool names 508
- catalog
 - and use of Accelerator 723
 - archiving 696, 891
 - files 868
 - recovery 50, 875, 886
 - utility 955
- catalog backups
 - about 875
 - adding critical policies to 840
 - archiving 887, 890–891, 893–894
 - compressing image catalog 899
 - determining success of 885
 - image files 869, 970
 - manual backup 883

- catalog backups *(continued)*
 - master server offline 90
 - Maximum concurrent jobs setting 565
 - media server version 875–876, 879
 - moving client images 896
 - multiple file layout 870
 - overview 865
 - parent and child jobs 1021
 - policy type 837
 - running concurrently with other backups 154
 - schedules for 884
 - single file layout 869
 - space required 894
 - strategies 886
 - uncompressing 900
 - wizard 49, 876
- Catalog cleanup wait time property 82
- CatalogBackup volume pool 442, 653, 688
- cataloging snapshots 616, 618
- CDE (Common Desktop Environment) 36
- cdrom file system, excluding from backup 831
- certificates. *See* security certificates
- Change journal
 - and synthetic backups 858
 - determining if enabling is useful 107
 - guidelines 108
 - problems when stopping the replication service manually 671
- Change Server operation 53–54, 63, 174, 794, 1060, 1099, 1106, 1108–1109, 1111, 1113, 1118
- Check the capacity of disk storage units property 147, 559
- CHECK_RESTORE_CLIENT 258, 277
- Checkpoint restart
 - and disk staging 578
 - and synthetic backups 853
 - backup jobs 690
 - configuration option for maximum incomplete status 294
 - for restore jobs 692
 - Move job from incomplete state to done state property 81–82
- CIFS
 - credentials for BasicDisk storage units 1030–1031, 1040–1041, 1043
 - credentials for disk storage units 556
 - share in the Backup Selections tab 805
 - volumes 1092
- cipher types for NetBackup encryption 120
- Clean-up host properties 80, 294, 692
- cleaning
 - drives 427, 1045
 - frequency 414
- Cleaning Frequency tape drive configuration
 - option 414
- Cleanup interval (SLP parameter) 242
- Client administrator's email property 252
- Client Attributes host properties 84, 87, 91
- Client backups report 1058
- Client cipher property 120
- Client connect timeout property 247, 278
- Client name property 82
- Client port window property 193
- Client read timeout property 100, 249, 280
- Client sends mail setting 252, 365
- Client Service, NetBackup 1092
- CLIENT_CONNECT_TIMEOUT 278
- CLIENT_NAME 260, 340
- CLIENT_PORT_WINDOW 278
- CLIENT_READ_TIMEOUT 248–249, 277, 279–280
- CLIENT_RESERVED_PORT_WINDOW 281
- clients
 - adding a client to the client database 85
 - adding and removing clients 93
 - adding to a policy 791
 - bp.conf options 259
 - choosing a policy type 681
 - deleting from policy 680
 - exclude and include lists 138
 - exclude files list 128–129, 132–133, 832
 - include files list 136
 - installing 793
 - maximum jobs 152
 - moving image catalog 896
 - name 1151
 - peername 1151
 - secure 793
 - setting host names 792
 - taking offline 88–89
 - trusting clients 794
- Clone snapshot type 742
- Cloud Settings tab 222
- cloud storage configuration wizard 49
- cloud storage server
 - credentials for 370
- CloudStore, NetBackup 1029
- clustering 65, 865, 906, 917, 1141
 - master server inter-node authentication 231

- Collect disaster recovery information for Bare Metal Restore policy attribute 707
 - Collect true image restore information (TIR) property 709
 - Collect true image restore information (TIR) with move detection property 81, 707, 710, 857
 - Communications buffer size property 106
 - compatibility list
 - for NDMP 1097
 - Compress catalog interval property 152, 899
 - COMPRESS_SUFFIX 342
 - Compression policy attribute
 - considerations regarding compression 703
 - disabled by Enable granular recovery policy attribute 716
 - compression used by VxFS 703, 1163
 - concurrent jobs
 - on client 152
 - per policy 693
 - CONNECT_OPTIONS 144, 282
 - Consistency check before backup host property 238
 - Continue with backup if consistency check fails property 126
 - Coordinated Universal Time (UTC) 976
 - copies
 - creating using Catalog duplicating option 962
 - creating using storage lifecycle policies 654, 762
 - option in the Configure Multiple Copies dialog box 761
 - primary 761, 960, 966
 - third party 760
 - copy
 - a policy 678
 - a policy schedule 679
 - NetBackup database files 929, 939
 - policy backup selections, clients, schedules 677
 - report text to another document 1060
 - Copy-on-write
 - snapshots 807
 - technology 742
 - correcting paths of devices automatically 434
 - cpio format 181
 - Credential Access host properties 108
 - credentials 370
 - about NDMP 370
 - Critical Policies list 839–840, 879, 881, 883
 - cross mount points
 - effect with UNIX raw partitions 700
 - interaction with Follow NFS policy attribute 701
 - cross mount points *(continued)*
 - policy attribute 813
 - policy setting 699
 - separate policies for 700
 - CRYPT_KEYFILE 347
 - CRYPT_KIND 343
 - CRYPT_LIBPATH 346
 - CRYPT_STRENGTH 345
 - ctime 816
 - cumulative incremental backups 746, 748
 - curly brackets as wildcards 1095
- ## D
- Daemon connection port property 93, 146
 - Daemon port only property 146
 - daemons
 - check with vmps 1169
 - checking processes 1168
 - monitoring 1038
 - starting and stopping 1038
 - tab in the Activity Monitor 1027
 - types 1037
 - DAS drive name tape drive configuration option 415
 - data
 - deduplication 704–705, 737
 - movers 556
 - data classifications
 - creating 112
 - host properties 110
 - in storage lifecycle policies 600, 641
 - policy attribute 684, 686, 764
 - use of Any 989, 1014
 - Data Mover off-host backup method 611, 617, 741
 - Data Protection Optimization Option license key 864
 - Database Administration utility 910–911, 917, 931
 - Database manager process (bpdbm) 1039
 - database schema, exporting 929, 938
 - database, for jobs 1045
 - database-extension clients, adding file paths for 818
 - DataStore
 - policy type 682
 - volume pool 481, 688
 - DataTools-SQL-BackTrack policy type 76, 682
 - datetime stamp 751
 - Daylight savings time 1138
 - DB2 policy type 76, 682, 797
 - db_marker.txt 869
 - DBA password, changing 910, 926, 934, 940
 - dbadm (Database Administration utility) 931

- DBR format 181
- deactivating policy 695
- debug logging levels 166
- decommissioning a media server 378, 385
- deduplication optimized synthetic backup method 864
- Deduplication property 89–90
- deduplication storage server
 - credentials for 370
 - defining target for Auto Image Replication 980
- Default cache device path for Snapshots property 101
- Default Job Priorities host properties 113, 695
- DEFAULT_CONNECT_OPTIONS 144, 282–283, 285
- Deferred duplication offset time (SLP parameter) 243, 645
- Defragment NetBackup database files 921, 935
- Delay on multiplexed restores property 147
- deleting
 - all devices for a media server 375
 - device hosts 378
 - drives 428
 - license keys 55
 - robots 409
 - schedules 779
 - schedules, backup selections, or clients from a policy 679
 - server groups 369
 - storage lifecycle policies 605
 - storage unit groups 594
 - storage units 552
 - volume group 511
 - volume pools 508
- Density storage unit setting 562
- Detailed job status 1023, 1038
- Device Configuration Wizard 49
- device hosts
 - removing 378
 - robot configuration option 404
- Device Monitor
 - add drive comment 423
 - assigning requests 1053
 - display pending requests 1052
 - resubmit request 1055
- devices
 - configuration wizard 401
 - configuring 395
 - discovery 397
 - file 407
 - file permissions 439
 - managing remotely 1099
- devices (*continued*)
 - mapping files 394
- devpts file system, excluding from backup 831
- DHCP setting on client 190
- differential incremental backups 746, 748
- Differential snapshot type 742
- Direct Access Recovery (DAR) 148
- Directory can exist on the root file system or system
 - disk setting 561
- directory junctions on UNIX 812
- Disable client-side deduplication policy attribute 715
- DISABLE_JOB_LOGGING 287
- DISABLE_STANDALONE_DRIVE_EXTENSIONS 288, 452
- DISALLOW_BACKUPS_SPANNING_MEDIA 289
- DISALLOW_CLIENT_LIST_RESTORE 289
- DISALLOW_CLIENT_RESTORE 290
- DISALLOW_SERVER_FILE_WRITES 251, 348
- disaster recovery
 - file 878, 886
 - information 153
 - sending e-mails 839
 - tab 837, 840
 - using the Make Master option 237
- disk
 - array, credentials for 108
 - image backups 806
 - spanning 183, 598, 687
 - staging storage units 596
- Disk logs report 1059
- disk pool
 - changing the state 390
- Disk pool status report 1060
- Disk pool storage unit setting 562
- disk pools
 - AdvancedDisk 389, 1051
 - cloud 389, 1051
 - configuration wizard 49
 - deduplication 389, 1051
 - description 556
 - OpenStorage 389, 1051
 - Replication Director 390, 1051
- Disk resource multiplier (SLP parameter) 243
- disk staging 576
- Disk storage unit status report 1059
- Disk type storage unit setting 562
- disk-image backups
 - checkpoint restart 691

- Distributed Application Restore Mapping
 - host properties 117
- Do not compress files ending with property 101, 342
- DO_NOT_RESET_FILE_ACCESS_TIME 349
- domains
 - replicating backups to another. *See* Auto Image Replication
- down a device 424
- Drive is in a robotic library tape drive configuration option 413
- Drive name tape drive configuration option 412
- Drive path options 418
- Drive Status tape drive configuration option 414
- drives
 - access permission 439
 - adding comment 423
 - cleaning 414, 427, 430, 1045
 - drive status 414
 - monitoring 1045
 - name rules 400
 - name rules, configuring 416
 - replacing 435
 - running diagnostics 432
 - servicing requests 1052
 - types 413
 - updating firmware 437
- DSFR. *See* Microsoft Distributed File System Replication
 - Shadow Copy Components:\ directive 825
 - snapshot provider 96
 - Windows Open File Backups 95
- duplicate backups
 - becoming a primary copy 965
 - creating 962
 - restoring from 960
- Duplication
 - job 614
 - optimized 642, 646
 - SLP operation 612
- Duplication Manager 1032
 - See also* Storage Lifecycle Manager service (nbstserv)
- Duplication to remote master. *See* Auto Image Replication
- DUPLICATION_SESSION_INTERVAL_MINUTES 245
- dynamically-allocated ports 194
- E**
- EFI System partitions 824
- ejecting volumes 498
- Email
 - address for administrator of this client 252
 - disaster recovery 839
 - notifications 153, 155–156, 158, 160
 - send from client 252, 365
 - send from server 252
- EMMSERVER 258, 291
- Enable block sharing storage unit setting 563
- Enable encryption property 120
- Enable granular recovery policy attribute 715, 845
- Enable job logging property 183
- Enable multiplexing storage unit setting 563
- Enable optimized backup of Windows deduplicated volumes 737
- Enable robust logging property 171
- Enable SCSI reserve property 182
- Enable single instance backup for message attachments property 126
- Enable standalone drive extension property 183, 288
- Enable standard encryption property 120
- Encrypting File System (EFS) 121
- Encryption host properties 119, 121, 343, 852
- Encryption policy attribute 705
 - about NetBackup encryption options 705
 - disabled by Enable granular recovery policy attribute 716
- English error log 293
- Enterprise Media Manager (EMM) 231, 318, 554, 870, 1029, 1040, 1141, 1143–1145
 - database 903
 - removing a device host 378
 - service 372, 374, 873
- Enterprise Vault Hosts properties 123
- Enterprise Vault properties 122
- Enterprise-Vault policy type 76
- erasing media 493
- error codes. *See* status codes
- escape character on UNIX 1095
- Exchange credentials property 127
- Exchange granular proxy host property 126
- exclude
 - cdrom file system from backup 831
 - days from a backup policy schedule 780
 - days from an SLP window 649–650
 - files and directories from backup 129, 132–133
 - files and directories from UNIX backups 128
 - files and directories from Windows backups 128
 - files list 138, 832

- exclude (*continued*)
 - list syntax 133
- Exclude list host properties 130–131
- exclude lists
 - configuring the exclude list on a Windows client from a UNIX master server 136
 - for specific policies and schedules 833
 - UNIX example 135
 - Windows example 135
- exclude_list 832
- Expire after copy retention type 635, 659, 852
- expiring backups 969
- export
 - database schema and data 929, 938
 - host properties 66
 - license keys 55
 - reports 1061
- Extended attribute files
 - disabling the restore of 818
 - Solaris 9 812
 - to ignore during backup 350
- Extended attribute files for GPFS
 - to ignore during backup 351
- Extended image retry interval (SLP parameter) 242, 607

F

- Fail all copies when creating multiple copies 582, 763
- Fail property in Fibre Transport host properties 141
- failover
 - media server to alternate media server(s) 216
 - servers, adding or changing 218
 - storage unit selection in group 595
 - to an alternate server 1146
- FAILOVER_RESTORE_MEDIA_SERVERS 292
- Federated SharePoint configurations 239
- Fibre Transport (FT)
 - host properties 139
 - media servers 228, 316
- File browse timeout property 248, 356
- File Change Log (FCL) 101–102
- file lists
 - extension clients 818
 - links on UNIX 809
 - raw partitions 812
 - UNIX clients 810
 - UNIX files not backed up 811, 831
 - Windows clients 803
- file names (not allowed in Backup Selections list) 1097

- File system backup coverage report 800
- files
 - .SeCuRiTty.nnnn 1159
 - /.rhosts 794
 - catalog space requirements 894
 - excluded from backups 831
 - excluding from backup 128–129, 132–133
 - linked, UNIX 812
 - NFS mounted 680, 696
 - No.Restrictions 1151
 - NOTES.INI 179
 - peername 1152
 - redirected restores 1153
 - restrictions on restores 1151
- FilesNotToBackup list 833
- Final destination
 - media owner 587
 - storage unit 587
 - volume pool 587
- Firewall host properties 142, 285
- FIREWALL_IN 1126
- Fixed retention type 635
- FlashBackup 812–814
 - policy type 76, 682
 - Windows policy type 76, 682
- Follow NFS mounts
 - cross mount points 697
 - raw partitions 697, 813
- Follow NFS policy attribute 696, 701
- Force interval for small job (SLP parameter) 242
- FORCE_IPADDR_LOOKUP 1128
- FQDN or IP Address property in Resilient Network
 - host properties 211
- Free browse property 89
- Frequency schedule attribute 757
- frozen media 496–497
- full backups 745, 854

G

- General level logging property 105
- General server host properties 147, 149
- GENERATE_ENGLISH_LOGS 258, 293
- get_license_key license key utility 56
- Getting Started Wizard 49
- Global attributes host properties 151, 153
- Global logging level property 165, 330
- globally unique version number (GVSN) 670
- Go into effect at policy attribute 695
- GPFS (General Parallel File System) 351–352

Granular recovery of Active Directory objects 843
 Granular Recovery Technology (GRT) 148, 646, 715
 Group images across SLPs (SLP parameter) 243
 Group Policy Objects 849

H

hard and soft links
 hard links in UNIX directories 809
 hard links NTFS volumes 809
 restoring 1150
 Hardware Compatibility List (HCL) 393, 406
 High water mark storage unit setting 563
 HKEYS, backing up 808
 host
 device 33
 properties
 changing in a clustered environment 65
 exporting 66
 permission to change 61
 host credentials. *See* credentials
 How long to retry failed import jobs (SLP parameter) 243
 Hyper-V policy type 682, 835
 Hyper-V virtual machines
 browse 795

I

IBM device number tape drive configuration option 415
 If this copy fails option 761
 IGNORE_ACL touch file 351
 IGNORE_XATTR 258, 350–352
 image .f files 616, 618, 868
 image catalog
 compressing 152
 moving 896
 Image cleanup property 82
 Image processing interval (SLP parameter) 242, 607
 IMAGE_EXTENDED_RETRY_PERIOD_IN_HOURS 607
 images
 changing primary copy 960
 duplicating 962
 moving client catalog 896
 on disk report 1059
 on media report 1058
 restoring from duplicate 960
 verifying 959
 Impact Report 220, 648

Import

Images Wizard 50
 job 615
 Manager 993, 1032
 See also Storage Lifecycle Manager service (nbtstserv)
 operation 989
 SLP operation 614
 IMPORT_EXTENDED_RETRY_SESSION_TIMER 245
 IMPORT_SESSION_TIMER 245
 importing backups 970–971
 Inactive media 1059
 Include Dates tab 781
 include files list 136
 include list, on client 138, 832
 include_list 832
 INCOMPLETE_JOB_CLEAN_INTERVAL 294
 Incrementals based on
 archive bit property 106
 timestamp property 105
 Independent property 593
 Index From Snapshot operation 616, 618, 632
 Infinite retention 222
 Informix policy type 682, 797
 INFORMIX_HOME 258, 353
 INI file, for Lotus Notes 179
 Initial browse search limit property 251
 INITIAL_BROWSE_SEARCH_LIMIT 295, 1132
 INITIAL_MEMORY 1129, 1134
 inject volume into robot
 multiple volumes 529
 robot inventory 498
 Inline copy option 760, 963, 968
 Inline Tape Copy 760
 Inode change time 363
 installing and configuring Network File System (NFS) 1170
 installing client software
 on secure clients 793
 on trusting clients 793
 Instant Recovery
 Advanced Backup method 691
 Backups to disk only setting 759
 Inter-domain Master Servers 229, 327, 993
 Internet Assigned Numbers Authority (IANA) 195, 1132
 inventory and compare robot contents 517
 IP Address Family Support host property 192
 IP_ADDRESS_FAMILY 192, 296

IPv4

- addresses 202–203
- IP_ADDRESS_FAMILY option 296
- networks, limiting bandwidth 75, 300
- PREFERRED_NETWORK option 308

IPv6

- addresses 202–203
- and client names 83, 792
- IP_ADDRESS_FAMILY option 296
- networks, limiting bandwidth 245, 323
- PREFERRED_NETWORK option 308

J**Java**

- auth.conf file 1120
- authorizing users 1119
- directory 1122
- performance improvement hints 1134
- Virtual Machine (JVM) 1129

Java Remote Administration Console 1117

- configuration options 1126

jbpSA 1117, 1120–1122, 1124, 1129, 1132–1133**jnbSA 36, 1037, 1132****Job Manager logging property 171****Job retry delay property 151****JOB_PRIORITY 297****jobs**

- canceling 1023
- Concurrent per disk storage unit 565
- database 1045
- filtering 1022
- maximum per client 152
- maximum per policy 693
- priority for policy 695
- restarting 1024
- resuming 1025
- setting default priority 113
- SLP_MultipleLifecycles 245
- suspending 1024
- viewing in the Activity Monitor 1020, 1022

JVM (Java Virtual Machine) 1129**K****Keep logs for days property 168****Keep logs up to GB property 169****Keep status of user-directed backups 101, 107, 355****Keep true image restoration information property 81****Keep vault logs for property 170****KEEP_DATABASE_COMM_FILE 258, 354****KEEP_LOGS_DAYS 355, 1132–1133****KEEP_LOGS_SIZE_GB 302****KeysNotToRestore list 833****Keyword phrase policy attribute 740****L****labeling media 501–502****legacy logging 165, 171, 1065****library**

- name robot configuration option 406

- sharing 399

license keys 52–56**LIFECYCLE_PARAMETERS file 240****Limit jobs per policy setting 693, 714–715, 754, 771****LIMIT_BANDWIDTH 299, 324****limiting bandwidth 75–76****links**

- UNIX hard-linked directories 809

- UNIX symbolic 812

Linux bind mount 831**LIST_FILES_TIMEOUT 355, 900****LiveUpdate 90****LMCP device file robot configuration option 406****load balancing methods 596****Locked file action property 100, 356****LOCKED_FILE_ACTION 356****logging**

- Accelerator 725

- bpdbm 950

- bpsynth 859

- deleting logs after a set time 168

- disabling for NetBackup processes 171

- host properties 161

- legacy 161

- nbdb 950

- NetBackup Administration Console

- add data and timestamp to logs 47

- configuring 46

- enable command logs 47

- log file directories 47

- log file name format 48

- log file size 48

- enable command line logs 47

- off of NetBackup automatically 176

- redirected restore 253

- unified 162

Logging Assistant

- adding a new record 1070

- Logging Assistant (*continued*)
 - debug logging 1073, 1077–1079
 - deleting a record 1072
 - NetBackup Support Utility (nbsu) 1085
 - new in the NetBackup Administration Console 44
 - overview 1063
 - uploading evidence 1088
 - using for troubleshooting 1066
 - viewing records 1067
- Logging for critical processes property 167, 171
- Logging host properties 267, 272–274, 330
- Login Banner Configuration host properties 173
- login banner text, removing 175
- long erase 494
- Lotus Notes
 - host properties 177
 - policy type 76, 683
- Lotus-Notes
 - policy type 796
- Low water mark storage unit setting 561, 564
- ltid. *See* NetBackup Device Manager (ltid)
- ltid (NetBackup Device Manager) 1040

M

- Mac OS X 684
- mail notifications
 - administrator email address 252
 - Disaster Recovery attachment
 - sending 839
 - USEMAIL on UNIX clients 365
- mail_dr_info.cmd 887
- Make Master option 237
- Managed by SLP retention 742
- manual backups
 - NetBackup catalogs 883
 - policy for 843
- Master Compatibility List 1097
- master servers
 - changing 237
 - inter-node authentication for clustering 231
 - rebooting 1168
- Match directive for Preferred Network host
 - properties 198, 203, 205, 207
- MAX_LOGFILE_SIZE 162, 261, 300
- MAX_MEMORY 1129, 1134
- MAX_NUM_LOGFILES 162, 261, 301
- maximum
 - jobs per client 152
 - jobs per policy 693
 - maximum (*continued*)
 - vault jobs property 153
- Maximum concurrent FT connections property in Fibre
 - Transport host properties 140
- Maximum concurrent jobs storage unit setting 565
- Maximum concurrent write drives setting 564
- Maximum concurrent write drives storage unit
 - setting 760
- Maximum data streams property 88, 714–715
- Maximum error messages for server property 107
- Maximum images for snapshot replication job (SLP
 - parameter) 241
- Maximum images per A.I.R. Import job (SLP
 - parameter) 241
- Maximum jobs per client property 714–715
- Maximum number of logs to restore property 178
- Maximum size per A.I.R. replication job (SLP
 - parameter) 241
- Maximum size per duplication job (SLP
 - parameter) 241
- Maximum snapshot limit retention type 636
- Maximum Snapshots parameter 636, 742
- Maximum streams per drive storage unit setting 567, 769
- media
 - active 1059
 - ejection timeout period 499
 - erasing 493
 - formats 446
 - frozen 496–497
 - host override property 149
 - host properties 180
 - ID prefix (non-robotic) property 184
 - ID, generation rules 472, 475
 - ID, prefix for update robot 458
 - inactive 1059
 - labeling 502
 - log entries report 168, 1058
 - mount errors 1051
 - pools. *See* volume pools
 - request delay property 184, 305
 - server connect timeout property 249
 - suspend and unsuspend 506
 - type when not an API robot 460
 - types 445
 - unmount delay property 184
- media access port (MAP) functionality 497
- Media host properties 262, 288–289, 304

- Media Manager device daemon (ltid). *See* NetBackup Device Manager (ltid)
- Media mount timeout property 248
- Media owner policy attribute 695
- Media request delay property 305
- Media server copy advanced backup method 691
- Media Server Deduplication Pool
 - changing the state 390
- Media server load balancing storage unit selection in group 595–596
- Media server storage unit setting 568
- media servers
 - activate or deactivate 371
 - adding 372
 - adding a media server to the Alternate restore failover machine list 218
 - decommissioning 378, 380–381, 385
 - delete all devices from 375
 - moving a robot and its media 409
 - previewing references to 384
 - problems displaying 1038
 - rebooting 1168
 - registering 373
 - registering with the master server 1099
 - Restore failover host properties 216
- media sharing
 - about 511
 - configuring unrestricted 512
 - configuring with a server group 512
- MEDIA_REQUEST_DELAY 305
- MEDIA_SERVER 305
- MEDIA_UNMOUNT_DELAY 304
- Megabytes of memory property 101, 358
- MEM_USE_WARNING 1130
- metadata attributes for GPFS 351–352
- Microsoft Data Deduplication file systems
 - about backups 738
 - about restores 739
 - configuration for backups 740
 - enable backups of 737
- Microsoft Distributed File System Replication
 - About backups 670
 - Shadow Copy Components:\ directive 825
 - snapshot provider 96
 - Windows Open File Backups 95
- Microsoft Exchange Attributes 743
- Microsoft Volume Shadow Copy Service (VSS) 74, 96
- Microsoft Windows Backup 833
- Microsoft Windows User Access Control (UAC) 38
- Minimum images per A.I.R. Import job (SLP parameter) 241
- Minimum size per duplication job (SLP parameter) 241
- Mirror property 593, 628
- Mirror retention type 637
- Mirror-break-off technology 742
- mirrored transaction log, creating 943
- mixing retention levels on tape volumes 768
- mklogdir shell script 162
- mklogdir.bat 162
- mntfs file system, excluding from backup 831
- monitoring
 - NetBackup drives 1045
 - NetBackup processes 1043
- mount
 - points 699
 - requests, pending 1052
- move
 - a robot and its media 409
 - backup job from incomplete state to done state
 - property 82
 - NBDB database files 941
 - NetBackup database files 927, 937
 - restore job from incomplete state to done
 - state 692
 - restore job from incomplete state to done state
 - property 81
- Move backup job from incomplete state to done state
 - property 690
- MPX_RESTORE_DELAY 306
- MS-Exchange-Server policy type 683, 796
- MS-SharePoint policy type 76, 683, 796
- MS-SQL-Server policy type 76, 683, 797
- MS-Windows policy type 683, 796
- MSDP
 - storage servers 327
- MTF format 181
- mtime 816
- multihomed server example 199
- multiple
 - data streams 712, 1021
 - file layout for NetBackup catalogs 870
 - installations 178
- multiple copies
 - checkpoint restart 691
 - creating using a policy schedule 762
 - creating using SLPs 655
 - creating using storage lifecycle policies 654

- multiple copies (*continued*)
 - criteria for creating 760
 - fail all copies 582, 763
 - parent and child jobs 1021
 - setting 760
 - synthetic backups method 635, 852, 860–863
- multiplexing (MPX)
 - and synthetic backups 852
 - demultiplexing 776
 - Maximum jobs per client property 772
 - preserving 645
 - set for schedule 769
 - use with Enable block sharing 563
- multistreaming backups 784, 852
- Must use local drive property 148, 307

N

- named data streams
 - disabling the restore of 818
 - to ignore during backups 350
- naming conventions 1093
- nb_updatedssu script 559
- NBAC (NetBackup Access Control). *See* NetBackup Access Control (NBAC)
- nbars (NetBackup Agent Request Server) 1028, 1040
- nbatd (NetBackup Product Authentication) 941, 1028
- nbaudit (NetBackup Audit Manager) 1028
- nbaudit (NetBackup Audit Service) 1040
- nbazd (NetBackup Product Authorization) 941, 1028
- NBAZDB 905, 915
 - See also* NetBackup Authorization database
- NBDB.db
 - configuration entry 915
 - creating manually 945
 - in catalog 870
 - installation overview 904
 - moving from one host to another 954
 - relocating 906, 941
- NbDbAdmin.exe (Database Administration utility) 917
- nbdecommission command 385
- NBEMM (NetBackup Enterprise Media Manager) 1029, 1040
- nbemmcmd command 62
- nbevtmgr (NetBackup Event Management Service) 1029
- nbEvtMgr process 1041
- nbfsd port 1188
- nbftsvr process 1041
- nbgetconfig 62

- nbim (NetBackup Indexing Manager) 1029
- nbj.conf 652, 1126
- NBJAVA_CLIENT_PORT_WINDOW 1130
- NBJAVA_CORBA_DEFAULT_TIMEOUT 652, 1131
- NBJAVA_CORBA_LONG_TIMEOUT 1131
- NBJM (NetBackup Job Manager) 171, 1041
- nblog.conf file 163
- nbmail.cmd script 153, 155, 887
- nbostrpxy (NetBackup Proxy Service) 1030
- NBPEM (NetBackup Policy Execution Manager) 171, 1030, 1041
- nbproxy process 1041
- NBRB (NetBackup Resource Broker) 171, 1031
- nbrb process 1041
- nbrbutil configuration utility 1032
- NBRMMS (NetBackup Remote Management and Monitor Service) 557, 1031, 1041
- NBRNTD (NetBackup Remote Network Transport Service) 1031
- NBRNTD_IDLE_TIMEOUT 307
- nbsetconfig 62, 259
- NBSL (NetBackup Service Layer) 1031, 1038
- nbssl process 1041
- nbstutil (lifecycle utility) command 606
- nbstserv process 244, 990, 993, 1032, 1042
- nbsvcmon process 1042
- NBU-Catalog policy type 653, 683, 688, 837
- NBVAULT (NetBackup Vault Manager) 1032, 1042
- NCR-Teradata policy type 76
- NDMP
 - ALL_FILESYSTEMS directive in backup policies 820
 - credentials for 44, 108, 370
 - Data Mover off-host backup method 741
 - Direct Access Recovery for restores 148
 - drives 148
 - global credentials 186
 - host name robot configuration option 406
 - host storage unit setting 569
 - hosts 186, 396
 - policy type 76, 611, 683
 - storage units 559, 574, 642, 646, 760
 - wildcard characters in NDMP policy backup selections 1096
- NetApp
 - disk arrays 127, 1093, 1097
 - FlexClone volumes 821
 - policy validation for 742

- NetApp (*continued*)
 - SnapMirror or SnapVault replication method 625, 637
- NetBackup
 - Authorization database 914
 - client service 189
 - Client Service (bpcd) 1092
 - request service port (BPRD) 189
- NetBackup Accelerator. *See* Accelerator
- NetBackup Access Control (NBAC) 35, 40, 44, 67, 329, 905, 914–915, 941, 950, 993, 1119
- NetBackup Administration Console
 - administering different versions 39
 - available consoles 34
 - configuration options 1130
 - improving performance 1133–1134
 - remote console 38
 - setting up on UNIX 36
 - starting 37
 - when to restart 1167
- NetBackup Audit Manager 1040
- NetBackup Authorization database 903, 905, 915
- NetBackup Client Service (BPCD) 1028
- NetBackup Cloud Storage 49
- NetBackup CloudStore 1029
- NetBackup compatibility lists 395
- NetBackup Compatibility Service (BPCOMPATD) 1029
- NetBackup database files
 - adding space 920
 - backing up 929, 939
 - changing DBA password 926, 934
 - defragmenting space 920
 - exporting database schema 929
 - free and used space 920
 - moving 927, 937
 - rebuilding 923, 937
 - restoring 930, 939
 - validating 922, 936
- NetBackup Database Manager (BPDBM) 1029
- NetBackup Deduplication Engine 370
- NetBackup Device Manager (ltid)
 - about 438
 - in the Activity Monitor 1029
- NetBackup Discovery Service (nbdisco) 44, 359
- NetBackup for Informix 353
- NetBackup Java Administration Console
 - configuration options 1126
- NetBackup Java Administration Console configuration options 1167
- NetBackup Java, set up for 36
- NetBackup Job Manager (NBJM) 171, 1030
- NetBackup Key Management Service (NBKMS) 1030
- NetBackup Legacy Client Service (BPINETD) 1030
- NetBackup Legacy Network Service (vnetd) 1043
- NetBackup Monitor Service 1032
- NetBackup Policy Execution Manager (NBPEM) 171, 1030
- NetBackup Product Authentication and Authorization AUTHORIZATION_SERVICE 265
- NetBackup product Authentication and Authorization AUTHENTICATION_DOMAIN 263
- NetBackup Product Improvement Program 323
- NetBackup Remote Management and Monitor Service (NBRMMS) 1031
- NetBackup Remote Network Transport Service (NBRNTD) 1031
- NetBackup Request Daemon (BPRD) 1031
- NetBackup Request Service Port (BPRD) property 189
- NetBackup Resource Broker (NBRB) 115, 171, 1031–1032
- NetBackup Scalable Storage 224–225
- NetBackup Service Layer (NBSL) 1031, 1041
- NetBackup Storage Lifecycle Manager 1032
- NetBackup Support Utility (nbsu) 1085
- NetBackup Vault Manager (NBVAULT) 1032
- NetBackup Volume Manager (VMD)
 - the Activity Monitor 1032
- NetBackup Volume Manager (vmd)
 - and the NetBackup Device Manager 438
- NetBackup volume pool 688
- NetBackup-Java Administration Console
 - configuration options 1126, 1128–1132
 - improving performance 1135
 - restricting access to 1125
 - running locally 1134
 - when to restart 1167
- network
 - addresses, prohibiting 207
 - drives, backing up 697
- Network Attached Storage (NAS) 560
- Network Attributes tab 67, 70, 72
- Network File System (NFS), described 1170
- Network host properties 189
- Network Settings host properties 190, 192
- Never property in Fibre Transport host properties 141

NEW_STREAM

- file list directive 826

NFS (Network File System)

- Follow NFS policy attribute 696, 701
- NFS access timeout property 254, 358
- no disk spanning 183

NFS services 740**NFS_ACCESS_TIMEOUT 358****No storage unit selection 644****non reserved ports 252, 263****Non-targeted A.I.R. 1000–1001, 1006****None volume pool 688****nonroot administration for specific applications 1123****NTFS Deduplication 737****null_stu storage unit type 563, 686****NulOST plug-in 563, 686****O****offline**

- master server and catalog backups 90
- taking clients 88–89

OLD_VNETD_CALLBACK 359**On demand only storage unit setting 570, 598, 634****Only directive for Preferred Network host**

- properties 198, 205, 208

open schedules 786**OpenStorage**

- optimized duplication 315, 642, 646
- storage server. *See* NetBackup OpenStorage Solutions Guide for Disk

OpenStorage Disk Option 555, 562, 600**OpenStorage disk storage units 555, 591****OpenStorage optimized synthetic backup method 864****operating mode of tape drive, changing 424****Operational Restore, OpsCenter 616****operations, scheduling windows for 647****Operator's email address property 78, 340****OpsCenter 323, 909, 913, 1015, 1031, 1038****OpsCenter Operational Restore 616****OpsCenter servers tab 228, 230****Oracle**

- Instances and Instance Groups 44
- Intelligent Policy 44
- policy type 683–684, 742, 797

Override default job priority

- for Catalog jobs 114, 958
- for Media Contents report 114
- for Media contents report 1059
- for queued or active jobs 1026

Override policy

- storage selection setting 764
- volume pool setting 765

Overwrite existing files property 322, 327, 817**P****pagefile.sys 807****parent jobs 711, 1020**

- in Activity Monitor Jobs tab 1020
- Limit jobs per policy setting 694
- parent_end_notify script 1021
- parent_start_notify script 1021

parent_end_notify script 1021**parent_start_notify script 1021****password, changing 910, 926, 934, 940****paths**

- changing for a tape drive 425
- correction, enabling automatic 435
- separators 561

PBX (Symantec Private Branch Exchange) 195, 1043**PBX_PORT 1132****PC NetLink files 811****PDDO storage servers 327****peername files 1151–1152****pending actions 1055****Perform consistency check before backup with**

- Microsoft Volume Shadow Copy Service (VSS) property 126

Perform default search for restore property 107**Perform incrementals based on archive bit 750****Perform off-host backup policy attribute 741****Perform snapshot backups policy attribute 741, 743****permissions**

- for device access 439
- to change NetBackup properties 61

physical inventory utility 530**Plex snapshot type 742****Point-in-time copy 742****policies**

- activating 695
- backup types 745
- changing multiple policies at one time 676
- changing properties 674–675, 677, 680, 790–791
- creating 673, 841
- for Active Directory granular restores 844
- Include Dates tab 781
- media owner attribute 695
- overview 664

- policies (*continued*)
 - planning 666
 - setting priority 113, 695
 - types 681
 - user schedules 753
 - validation 742
 - volume pool policy setting 688–689
 - Policy Configuration Wizard 50, 879
 - Policy Execution Manager
 - Logging property 171
 - Policy storage policy attribute 685, 687, 882
 - Policy update interval property 152, 783, 958
 - policy validation 742
 - Port Ranges host properties 193–194
 - Port, Bus, Target, LUN configuration option 407
 - ports
 - allow operating system to select non reserved port 194
 - dynamically-allocated 194
 - non reserved 252, 263
 - Postponed secondary operations 642
 - power down NetBackup servers 1165
 - Preferred Network host properties 195, 197, 199, 204–205, 313
 - Preferred property in Fibre Transport host
 - properties 141
 - PREFERRED_NETWORK 308
 - prelabel media 502
 - preprocess interval 829
 - Preserve multiplexing option 645
 - previewing a media server's references 384
 - primary copies 960, 965
 - Primary copy setting in the Configure Multiple Copies
 - dialog box 761
 - Primary property 593, 626
 - print
 - job detail information 1026
 - job list information 1025
 - license key 54
 - list of SLP-processed images 606
 - reports 1061
 - Prioritized storage unit selection in group 595
 - priority
 - of a job 113, 695
 - of duplication jobs 581
 - of relocation jobs started from this schedule setting 586
 - Priority for secondary operations 641
 - Priority of duplication job option 761
 - Private Branch Exchange 1032, 1043
 - Problems report 168, 1058
 - proc file system, excluding from backups 831
 - Process busy files property 78, 340
 - processes
 - check with vmps 1169
 - monitoring 1043
 - show active 1168
 - Product Improvement Program 323
 - Prohibited directive for Preferred Network host
 - properties 198, 203, 205, 207
 - properties
 - changing on multiple hosts 63
 - exporting 66
 - viewing 63
 - PureDisk
 - Storage Option 562
 - Storage Pool Authority (SPA) 562
 - storage units 634
 - PXEMTFTP (NetBackup Bare Metal Restore MTFTP Services) 1028
- ## Q
- Query Builder 215
 - question mark as wildcard 1094
 - quick erase 494
 - quotas on file systems 556
- ## R
- random ports, setting on server 193
 - RANDOM_PORTS 311
 - raw partitions
 - backing up 691, 745, 806
 - backups on UNIX 812, 814
 - Follow NFS policy attribute 697
 - restoring 806
 - RE_READ_INTERVAL 312
 - rebooting NetBackup servers 1167–1168
 - recovering the catalog 875, 886
 - redirected restores 814, 1151
 - Reduce fragment size storage unit setting 572
 - ReFS file system 1164
 - register a media server 373
 - registered ports 194
 - registry, backup/restore 808
 - Reload NetBackup database 923, 937
 - reload.sql 952, 954
 - relocation schedule 580, 586, 588, 745, 757

- remote
 - access, allowing 1103
 - device management 1099
 - server, troubleshooting access 1113
- Remote Administration Console 1117
 - configuration options 1126
 - configuring 1130, 1150
 - displaying media servers in the Activity Monitor 1038
 - improving performance 1133
 - installing 38
 - on a Windows client 1113
 - upgrading 38
- removing a device host 378
- REORGANIZE command to defragment NetBackup
 - database 921, 935
- Replace a process level token 127
- replacing a drive 435
- replication
 - between NetBackup domains. *See* Auto Image Replication
 - job 620
 - method 625, 637
 - source property 593, 626–628
 - target property 593, 627–628
 - to an alternate NetBackup domain. *See* Auto Image Replication
- Replication Director
 - Getting Started Wizard, unsupported 49
 - Policy Configuration Wizard, unsupported 50, 673
 - policy validation for 742
 - Replication operation 619
 - SLP selection in policy 686
 - Solutions Guide 370, 593, 1015
 - use of Index From Snapshot operation in SLP 617
 - Use Replication Director policy attribute 741
- Replication operation 619
- REPORT_CLIENT_DISCOVERIES 258, 359
- reports
 - All log entries report 1058
 - Client backups report 1058
 - copying to another document 1060
 - Disk logs report 1059
 - Disk pool status report 1060
 - Disk storage unit status report 1059
 - Images on Disk report 1059
 - Images on media report 1058
 - Media log entries report 1058
 - printing 1061
 - Problems report 1058
 - running 1060
 - saving 1061
 - Status of backups report 1058
 - Tape contents report 114, 1059
 - Tape lists report 1059
 - Tape logs report 1058
 - Tape summary report 1059
 - Tape written report 1059
 - using the Troubleshooter 50
 - utility in Administration Console 1056
- reports *(continued)*
- requests
 - assigning 1053
 - denying 1055
 - resolving pending tape 1053
- REQUIRED_INTERFACE 308, 313
 - See also* Preferred Network host properties
- REQUIRED_NETWORK 308
- reset
 - file access time property 101, 721
 - mount time 429
- Resiliency property in Resilient Network host
 - properties 211
- Resilient connection
 - Resilient Network host properties 209
- Resilient Network host properties 209
 - FQDN or IP Address property in 211
 - Resiliency property in 211
- RESILIENT_NETWORK 314
- Resource Broker (NBRB) 115, 1032
- Resource Broker logging property 171
- Resource Limit host properties 214
- restarting jobs 1024
- Restore Failover host properties 216
- restore jobs
 - move restore job from incomplete state to done state 692
- Restore retries
 - checkpoint restart 692
 - property 250, 361
- RESTORE_RETRIES 360
- restores
 - adjust time zone for 1137
 - alternate server 1140
 - directed from the server 1149
 - from a specific backup copy 581, 763, 1140

- restores *(continued)*
 - keeping progress reports 101, 355
 - NetBackup database files 930, 939
 - overriding the original server 1144–1145
 - raw partition 806
 - redirected 216, 1151
 - registry on Windows clients 808
 - server independent 1140
 - symbolic links on UNIX 812
 - System State 1160
 - using a specific server 149
- RESUME_ORIG_DUP_ON_OPT_DUP_FAIL 258, 315
- resuming suspended jobs 1025
- Retain snapshots for Instant Recovery or SLP
 - management policy attribute 741
- retention levels
 - for archiving catalogs 891
- retention periods
 - changing 220
 - guidelines for setting 766
 - mixing on tape volumes 182, 768
 - precautions for setting 767
 - redefining in the host properties 219
 - setting 644, 766
 - volumes 221
- retention types for SLP operations
 - Capacity managed 633–634
 - Expire after copy 635
 - Fixed 635
 - Maximum snapshot limit 636
 - Mirror 637
 - mixing 632
 - Target 638
- retire a media server. *See* decommissioning a media server
- Retries allowed after runday 781
- Retries allowed after runday policy setting 757
- Retry count property 79
- Retry restores, setting 250
- Reverse Host Name Lookup host property 191–192, 315
- REVERSE_NAME_LOOKUP option 192, 315
- robot configuration
 - changing 407
 - robot number option 404
 - robot type option 404
- Robot control host robot configuration option 407

- Robot control is attached to an NDMP host robot
 - configuration option 405
- Robot control is handled by a remote host robot
 - configuration option 405
- robot control options 405
- Robot device path robot configuration option 407
- Robot device robot configuration option 407
- Robot drive number tape drive configuration
 - option 415
- Robot is controlled locally by this device host robot
 - configuration option 405
- Robot number robot configuration option 404
- Robot number storage unit setting 573
- Robot operating in manual mode
 - configuring 408
- Robot type robot configuration option 404
- Robot type storage unit setting 573
- robot types 393
- Robotic device file robot configuration option 407
- Robotic library tape drive configuration option 414
- robots
 - adding 402
 - compare contents 517
 - configuring 395
 - device file 407
 - device host configuration option 404
 - inventory 514–515
 - moving to new media server 409
 - running diagnostics 430
 - unsupported characters in media ID 526
 - update procedure 527
 - updating the volume configuration 526
- Round robin storage unit selection in group 595
- RS-MTF1 format 181

S

- SAP policy type 76, 683, 797
- save a report 1061
- Scalable Storage host properties 222, 224–225
- Scalable Storage, NetBackup 224–225
- Schedule backup attempts property 152, 690, 712, 790
- schedules
 - adding to a policy 675
 - backups on specific dates 781
 - creating a time window 777
 - default for user backups 336
 - deleting a time window 779
 - determining due time 783

- schedules *(continued)*
 - excluding dates 780
 - frequency setting 757
 - how NetBackup determines which schedule to run 782
 - naming 745
 - open 786, 789
 - overview 743
 - priority 758
 - recalculating 783
 - retention periods
 - guidelines 766
 - setting 766
 - specify multiplexing 769
 - Start Windows tab 776
 - storage unit/storage lifecycle policy 764
 - type of backup 745
 - user backup or archive 753
 - volume pool 765
 - windows that span midnight 785, 790
- scheduling secondary operations in an SLP 647
- scratch
 - pool and WORM media 449
 - pool, adding 443, 508
 - volume pool 688
- scripts 1021
 - bpdbjobs example 1048
 - bpps 1168
 - bpstart_notify 248–249, 277, 280
 - vmpps 1169
- SCSI
 - long erase 493
 - pass-through command 397
 - persistent reserve
 - drive path override 418
 - quick erase 494
 - reserve, configuring 182
 - reserve/release
 - drive path override 418
- search_data 914
- security certificates
 - about 45
 - for cloud storage 45
 - generated by 1028
 - generating 374
 - to access remote servers 1112
- SeCuRiTty.nnnn files 1159
- sendmail 155
- Serial Number tape drive configuration option 414
- SERVER 361
 - vm.conf entry 1100
- server
 - directed restores 1124
 - allowing access 1103
 - alternate server restores 1140
 - directed restore 1149
 - host properties 227, 230
 - media servers 305, 357
 - using 1103
 - independent restores 216, 1140
 - list, adding a server 1103
 - NetBackup
 - controlling access 361
 - power down 1165
 - rebooting 1165
 - SERVER configuration option 316
 - server groups
 - configuring 367
 - deleting 369
 - for media sharing 367
 - Server list
 - on UNIX client 361, 363
 - Server reserved port window host property 194
 - Server sends mail property 252
 - SERVER_CONNECT_TIMEOUT 319
 - SERVER_PORT_WINDOW 194, 283, 320
 - SERVER_RESERVED_PORT_WINDOW 283, 320
- servers
 - removing from server list 231
- Service Manager 916
- setconf.bat 652
- setconf.bat file 1126
- Shadow Copy Components directive 824
- Shadow Copy Service 74, 96
- shared drives
 - configuration wizards 396
- SharePoint Server
 - consistency checks options 239
 - Federated configurations 239
 - properties 237
- shut down NetBackup
 - daemons 1166
 - services 1166
- Simplified File Restore 616
 - See *also* OpsCenter Operational Restore
- single file
 - layout for NetBackup catalogs 869

- single file *(continued)*
 - restore program
 - FlashBackup 812
- Single-Instance Storage (SIS) 126, 634, 691, 704–705
- SKIP_RESTORE_TO_SYMLINK_DIR 321
- slot number
 - for volume 478
- SLP Parameters host properties 239, 244, 625
- SLP Windows 44, 643, 647
- SLP.AUTO_CREATE_IMPORT_SLP 243
- SLP.CLEANUP_SESSION_INTERVAL 242, 660
- SLP.DEFERRED_DUPLICATION_OFFSET_TIME 243
- SLP.DISK_RESOURCE_MULTIPLIER 243
- SLP.DUPLICATION_GROUP_CRITERIA 243
- SLP.IMAGE_EXTENDED_RETRY_PERIOD 242
- SLP.IMAGE_PROCESSING_INTERVAL 242, 245
- SLP.JOB_SUBMISSION_INTERVAL 242, 245
- SLP.MAX_IMAGES_PER_IMPORT_JOB 241
- SLP.MAX_IMAGES_PER_SNAPSHOT_REPLICATION_JOB 241
- SLP.MAX_SIZE_PER_BACKUP_REPLICATION_JOB 241
- SLP.MAX_SIZE_PER_DUPLICATION_JOB 241
- SLP.MAX_TIME_TIL_FORCE_SMALL_DUPLICATION_JOB 242
- SLP.MIN_IMAGES_PER_IMPORT_JOB 241
- SLP.MIN_SIZE_PER_DUPLICATION_JOB 241
- SLP.REPLICA_METADATA_CLEANUP_TIMER 243
- SLP.TAPE_RESOURCE_MULTIPLIER 243
- SLP.VERSION_CLEANUP_DELAY 242, 661
- SLP.WINDOW_CLOSE_BUFFER_TIME 243
- SLP_MultipleLifecycles job 245
- snapdupe 610
 - See also* Backup From Snapshot operation
- Snapshot Client 95, 255, 759, 1021
 - checkpoint restart 691
- Snapshot Client and Replication Director policy
 - attributes 741, 743
- Snapshot operation type 622
- Snapshot Type parameter 742
- Snapshot verification I/O throttle property 125
- snapshots
 - cataloging 616, 618
 - creating copies 654
 - disk array snapshot methods 110
 - job 623
 - Maximum images per snapshot replication
 - job 241
 - Microsoft Volume Shadow Copy Service 74
 - MS-Exchange-Server 125–126
 - operation type 641, 653
 - snapshots *(continued)*
 - Replication Director 127
 - selecting the snapshot provider 85, 93, 95–98
 - storage unit groups for 590, 592
 - Solaris 9 extended attributes 812
 - Source binding 197, 199, 203, 205, 207–208
 - SPC-2 SCSI reserve 182
 - SQL
 - Instances and Instance Groups 44
 - Intelligent Policy 44
 - SQL Anywhere
 - in NetBackup installation 870
 - SQLANYs_VERITAS_NB 916, 1030
 - square brackets as wildcards 1095
 - staging
 - backups 576
 - schedule storage unit setting 573
 - using BasicDisk storage unit 561
 - Standard policy type 796
 - standard toolbar 45
 - start up NetBackup
 - daemons 1167
 - services 1166
 - Start Window tab 776
 - startup text, removing 175
 - status codes
 - 1000 (Client is offline) 90
 - 12 (file open failed) 561
 - 1519 (Images are in process) 605
 - 1573 (Backup image cannot be expired because its SLP processing is not yet complete) 969
 - 4206 (Authentication error occurred) 1093
 - 71 (none of the files in the file list exist) 796
 - Status of backups report 1058
 - stopping jobs 1023
 - Storage device storage unit setting 574
 - Storage Lifecycle Manager service (nbstserv) 244, 990, 993, 1032, 1042
 - storage lifecycle policies 1032
 - See also* Storage Lifecycle Manager service (nbstserv)
 - Active secondary operations 642
 - Alternate read server option 645
 - and synthetic backups 635, 852
 - and the Multiple copies configuration dialog 764
 - batch creation logic 244
 - best practices document 601, 625
 - Capacity managed retention type 633–634
 - cataloging snapshots 616, 618

storage lifecycle policies *(continued)*

- copy number 655
- Data classification setting 641
- deleting 605
- deleting old versions of 660
- deprecated parameters 245
- hierarchy 602, 604, 629
- Impact Report 220, 648
- LIFECYCLE_PARAMETERS file 240
- Local storage option 644
- Media owner option 645
- mixing retention types 632
- operations 601
 - windows 610, 612, 614–615, 617, 620, 623
- operations, windows for 44
- parameters for 239
- Postponed secondary operations 642
- Preserve multiplexing option 645
- Priority for secondary operations 641
- retention types 631, 643, 766
- scheduling secondary operations 647
- SLP_DATA.db 914
- Storage lifecycle policy name 640
- Storage unit option 644
- Target master option 644
- using nbstlutil to administrate lifecycle operations 606
- utility 600
- Validate Across Backup Policies button 642, 652
- versions of 657–659
- Volume pool option 644
- writing multiple copies 654

storage server

- credentials for cloud storage 370
- credentials for deduplication 370
- define target for Auto Image Replication 980
- description 556
- OpenStorage. *See* NetBackup OpenStorage Solutions Guide for Disk

Storage Server Configuration Wizard 49

storage unit groups

- and NetBackup Accelerator 721
- and storage lifecycle policies 642, 646
- create for backups 590
- create for snapshots 592
- deleting 594
- not supported for Auto Image Replication source 609, 976
- storage unit selection criteria 594, 598

storage units

- AdvancedDisk disk type 555
- available storage property of volume 571
- BasicDisk type 555
- capacity property of volume 571
- changing server to manage 1112
- creating 548, 551
- creating a basic disk staging unit 578
- deleting 552
- disk pool comment property 571
- disk storage units 555
- for policy 685, 687
- for schedule 764
- high water mark property of volume 563, 571
- low water mark property of volume 564, 572
- Media Manager type 553
- name setting 574
- NDMP disk type 559
- null_stu 563, 686
- number of volumes property 572
- OpenStorage disk type 555, 634
- percent full property on volume 572
- PureDisk disk type 555, 634
- QIC drive type 760
- raw size property on volume 572
- storage lifecycle policies 557
- type setting 574
- types 591
- usable size property of volume 572
- utility for configuring 547
- vendor-specific 555
- Submission interval (SLP parameter) 242
- subnets 76, 325
- Sun PC NetLink 811
- suspended jobs 82, 1024–1025
- Sybase policy type 684, 797
- Sybase SQL Anywhere
 - dbsrv12.exe 1030
 - management of 916
 - starting/stopping the service 916
 - use in NetBackup 903
- SYBASE_HOME 258, 362
- Symantec Private Branch Exchange 1032, 1043
- symbolic links
 - included in backup selection list 799
 - UNIX 812
- synthetic backups
 - and encryption 852
 - checkpoint restart 691

synthetic backups *(continued)*
 component images 854–855
 deduplication 864
 logs produced during 859
 multiple copy backups method 635, 852, 860
 no Auto Image Replication support 852, 976
 no multiple copy support 761
 no NetBackup change journal support 108
 OpenStorage optimized method 864
 recommendations for using 851
 schedules 755, 784

System State
 backups, checkpoint restart 691
 directive 824
 restoring 1160

T

Take checkpoints every __ minutes (policy attribute) 690

Tape contents report 1059

tape drive configuration
 ACS, LSM, Panel, Drive option 415
 Cleaning Frequency option 414
 DAS drive name option 415
 Drive is in a robotic library option 413
 Drive Status option 414
 IBM device number option 415
 Robot drive number option 415
 Robotic library option 414
 Serial Number option 414

tape drives
 adding 410, 419
 adding a path 417
 assigning requests for tapes 1053
 changing operating mode 424
 changing to a shared tape drive 426
 configuration options 412–413
 configuring 395
 Drive name configuration option 412
 paths, changing 425
 properties, changing 426
 shared, operating modes of 424–425, 427
 TapeAlert 414

Tape lists report 1059

Tape logs report 1058

Tape Media contents report 1059

Tape resource multiplier (SLP parameter) 243

Tape summary report 1059

Tape written report 1059

TapeAlert 414

tar format 181, 610, 612, 624–626, 628

Target master storage option 644

Target retention type 638

Targeted A.I.R. 232, 976, 1000–1001, 1008

TCDebug_TCPP level logging property 105

TELEMETRY_UPLOAD 323

temporary staging area 564, 573–574, 580

third-party copies 760

Third-Party Copy Device Advanced Backup method 691

Throttle Bandwidth host properties 245–246, 323

THROTTLE_BANDWIDTH 245, 323

Time overlap property 106

time zones
 adjustment for restores 1137
 setting Daylight savings time 1138

Timeouts host properties 247, 276, 278–279, 307, 319, 356

tlmd daemon 1043

tmpfs file system, excluding from backup 831

topology of storage 980, 982

tpext utility 946

Transaction log cache path property 178

transaction log, creating 943

traversing directories to back up a file 138

Troubleshooter 50

True Image Restoration (TIR)
 Error code 136 858
 pruning information 858
 with Move Detection 858

True Image Restore (TIR) with Move Detection 108

Truncate log after successful Instant Recovery backup property 126

trusted master servers
 adding 233
 for Auto Image Replication 987
 removing 234
 Trusted Master Servers tab of Servers host properties 229

trusted NetBackup domain. *See* trusted master servers

TRUSTED_MASTER 327

U

ULINK_ON_OVERWRITE 259

UNC path
 checkpoint restart 691
 in the Backup Selections tab 805

- UNC path *(continued)*
 - with CIFS and AdvancedDisk storage units 557
 - with CIFS and BasicDisk storage units 557
- uncompress NetBackup catalogs 900
- unified logging 162, 166, 171, 260–261, 1065
- Universal Settings host properties 195, 249, 263, 349, 361, 365
- UNIX Client host properties 253
- UNIX Client Settings host properties 99, 102, 342, 350, 355–356, 358, 364
- UNIX clients
 - checkpoint restart 691
- UNIX epoch time 222
- UNIX server properties 254, 358
- UNLINK_ON_OVERWRITE 327
- Unload NetBackup database 923, 937
- unloading
 - database schema 929
 - the database schema 938
 - the NetBackup database 952
- UNSET file list directive 830
- UNSET_ALL file list directive 831
- Unused SLP definition version cleanup delay (SLP parameter) 242
- Update Sequence Number (USN) 671
- updating drive firmware 437
- usbdevfs file system, excluding from backup 832
- Use alternate read server attribute 581, 588
- Use case sensitive exclude list host property 130
- Use Change Journal option
 - and synthesized backups 858
 - and use with Accelerator 717, 722, 724, 726
 - guidelines for using 108
- Use defaults from the master server configuration
 - property in Fibre Transport host properties 141
- Use Direct Access Recovery for NDMP restores
 - property 148
- Use legacy DES encryption property 120
- Use non reserved ports property 145
- Use OS dependent timeouts property 248
- Use random port assignments properties 193
- Use reserved ports property 145
- Use specified network interface property 195, 251
- Use VxFS file change log for Incremental backups
 - property 101
- USE_AUTHENTICATION 329
- USE_CTIME_FOR_INCREMENTALS 363
- USE_FILE_CHG_LOG 364
- USE_VXSS 329

- USEMAIL on UNIX clients 365
- user
 - archive backups 747
 - backups 746
 - schedules, planning 753
 - toolbar 45
- User Access Control (UAC) 38
- User directed timeouts property 106
- UTC (Coordinated Universal Time) 976

V

- Validate Across Backup Policies button in SLP 642, 651–652
- Validate NetBackup database 922, 936
- Validate Policy dialog box 742
- validation for backup policies 742
- Validation Report tab 642, 652
- Vault
 - backup type 747
 - catalog archiving 891
 - designating duplicate as the primary 960
 - Maximum vault host property 153
 - parent and child jobs 1022
 - policy type 76, 684, 797, 841
 - process 1042
 - vlteject command 842
 - vltrun command 842
- vendor-specific storage units 555
- VERBOSE 330
- verifying backup
 - images 959
 - selections list 799
- Veritas Security Subsystem (VxSS)
 - USE_VXSS 329, 331, 993
- Veritas Volume Manager (VxVM) 759, 813
- Veritas Volume Snapshot Provider 96
- veritas_pbx (Symantec Private Branch Exchange) 93, 146, 1043
- veritas_pbx port 93, 146
- view properties of a license key 55
- virtual tape libraries (VTL) 644–645
- VLANs 228, 316
- vlteject Vault command 842
- vltrun Vault command 842
- vm.conf file, adding SERVER entries 1100
- VMD (NetBackup Volume Manager)
 - and the NetBackup Device Manager 438
 - the Activity Monitor 1032
- vmd process 1043

- vmphysinv physical inventory utility 530
- vmops script 1169
- VMware
 - Access Host properties 255
 - policy type 684, 742, 833
- vnetd
 - enabling logging for 146
 - NetBackup Legacy Network Service 1030, 1043
 - Only property (for selection of ports) 146
 - Veritas Network Daemon 145
 - with CIFS BasicDisk storage units 1040–1041, 1043
- VNETD_PORT 1132
- volume groups
 - changing name 489–490, 497
 - deleting 511
 - moving 509
 - rules for assigning 444
- volume pool names
 - reserved prefixes 442
- volume pools
 - adding 507
 - and WORM media 448
 - changing attributes 507
 - changing for a volume 490
 - DataStore pool 688
 - deleting 508
 - for schedule 765
 - indicating one for use by a policy 688
 - overview 441
 - properties 508
 - reserved prefixes for names 442
 - scratch 443
- Volume Shadow Copy Service (VSS) 74, 85, 96, 671, 824
- Volume Snapshot Provider (VSP) 85, 95–96
- VOLUME_EXCLUDE_LIST 820
- volumes
 - adding 451, 453, 477
 - assignments 688
 - changing properties 490
 - configuration wizard 49
 - determining retention period 221
 - ejecting 498
 - exchanging 495, 505
 - injecting 498
 - maximum mounts allowed 479
 - moving 489, 503–504
 - preview configuration update 523

- volumes (*continued*)
 - properties 478
 - update volume configuration 504
- VRTSpbx (Symantec Private Branch Exchange) 1032
- VXDBMS_NB_DATA registry entry 915
- VxFS
 - file change log 101
 - named data streams 815
 - restore of compressed files 1163
- vxlogcfg command 163, 166
- VXSS_NETWORK 330–331, 365

W

- Wait time before clearing archive bit property 105–106
- wildcard characters 804, 811, 1094, 1096
 - in Backup Selections list 1096
- Window close buffer time (SLP parameter) 243
- Window tab for secondary operations 44, 643, 647
- Windows 2003 824, 845
- Windows 2008 845
- Windows 2012 793
- Windows 8 793
- Windows Client host properties 257
- Windows Client Settings host properties 104, 107–108, 355
- Windows Disk-Image (raw) backups 691, 806
- Windows Map Network Drive option 805
- Windows Open File Backup host properties 93
- Windows policy type 683
- Windows Server Failover Clustering (WSFC) 814
- Windows Service Manager 916
- Windows User Access Control (UAC) 38
- windows, for storage lifecycle policy operation 44, 610, 612, 614–615, 617, 620, 623
- wizards
 - Catalog Backup 49, 874, 876
 - Catalog Recovery 50
 - Cloud Storage Server Configuration 49
 - Device Configuration 49, 401
 - Disk Pool Configuration 49
 - Getting Started 49
 - Import Images 50
 - Policy Configuration 50, 673, 879
 - Storage Server Configuration 49
 - Volume Configuration 49
- Working directory property 78, 339
- WORM media 449, 451
 - about 447

WORM media *(continued)*
 and the Quantum drive 450
 and unique drive and media types 449
 and volume pools 448
 Disabling WORM volume pool name
 verification 450
 encrypting 448

Y

year 2038 problem 222