



VERITAS NetBackup™ 6.0

System Administrator's Guide, Volume I

for UNIX and Linux

N15257B

September 2005

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 1993-2005 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo, and NetBackup are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Portions of this software are derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm. Copyright 1991-92, RSA Data Security, Inc. Created 1991. All rights reserved.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
www.veritas.com

Third-Party Copyrights

For a list of third-party copyrights, see the *NetBackup Release Notes* appendix.

Contents

Preface	xxxv
Getting Help	xxxv
Finding NetBackup Documentation	xxxv
▼ <i>To access the NetBackup online glossary</i>	xxxv
Accessing the VERITAS Technical Support Web Site	xxxvi
Contacting VERITAS Licensing	xxxvii
Accessibility Features	xxxvii
Comment on the Documentation	xxxviii
 Chapter 1. Introduction to NetBackup	1
Overview	1
NetBackup Administration Interfaces	3
NetBackup Administration Console Setup	3
▼ <i>To prepare a CDE (Common Desktop Environment) for NetBackup-Java interfaces</i> 4	
▼ <i>To start the NetBackup-Java Administration Console on a NetBackup-Java capable UNIX system</i> 5	
Running the Java-Based Windows Display Console	6
▼ <i>To start the Windows Display Console</i>	6
Administering Remote Servers	7
Administering Backlevel NetBackup Servers	8
Using the NetBackup Administration Console	9
NetBackup Configuration Wizards	10
Backup, Archive, and Restore	10

Activity Monitor	11
NetBackup Management	11
Reports	11
Policies	12
Storage Units	13
Catalog	13
Host Properties	14
Media and Device Management	14
Access Management	14
NetBackup Administration Console Menus	14
File Menu	14
Edit Menu	15
View Menu	16
Actions Menu	17
Help Menu	17
Standard and User Toolbars	19
Customizing the Administration Console	19
Configuring NetBackup Without Wizards	20
▼ <i>To configure NetBackup without wizards</i>	20
Chapter 2. Managing Storage Units	23
Introduction to Storage Units	24
Viewing Storage Units and Storage Unit Groups	26
Creating, Changing, and Deleting Storage Units	27
Creating a New Storage Unit	27
▼ <i>To use the Device Configuration Wizard</i>	27
▼ <i>To create a storage unit from the Actions menu</i>	27
▼ <i>To create a storage unit by copying an existing storage unit</i>	28
Changing Storage Unit Properties	28
▼ <i>To change storage unit properties</i>	28

Deleting Storage Units	28
▼ <i>To delete storage units</i>	28
NetBackup Naming Conventions	29
Media Manager Storage Unit Considerations	30
Before Adding a Media Manager Storage Unit	32
Disk Storage Unit Considerations	36
Interval Between Capacity Updates	37
Maintaining Available Space on Disk Storage Units	37
NDMP Storage Unit Considerations	38
Storage Unit Properties	39
Absolute Pathname to Directory/Volume	39
Density	40
Disk Type	40
Enable Block Sharing	40
Enable Multiplexing	40
High Water Mark	41
Low Water Mark	41
Maximum Concurrent Write Drives	41
Maximum Concurrent Jobs	42
Media Server	42
NDMP Host	43
On Demand Only	43
Reduce Fragment Size	43
Robot Number	44
Robot Type	45
Staging Relocation Schedule	45
Storage Device	45
Storage Unit Name	45
Storage Unit Type	45
Temporary Staging Area	46

Transfer Throttle	46
About Disk Staging	47
▼ <i>To create a disk staging storage unit</i>	48
Disk Staging Storage Unit Size and Capacity Considerations	49
Minimum Disk Staging Storage Unit Size	49
Disk Staging: Stage I	50
Disk Staging: Stage II	51
▼ <i>To manually initiate a disk staging storage unit relocation schedule</i>	52
Staging Schedule Button	52
Name	52
Priority of Relocation Jobs Started from this Schedule	52
Final Destination Storage Unit	53
Final Destination Volume Pool	53
Use Alternate Read Server	53
Disk Staging Limitations	54
Storage Unit Groups	55
▼ <i>To create a storage unit group</i>	55
▼ <i>To change a storage unit group</i>	57
▼ <i>To delete a storage unit group</i>	57
Chapter 3. Managing Backup Policies	59
Using the Policies Utility	60
Tree and Detail Views	60
Policies Menu Bar	60
Actions Menu	60
Standard and User Toolbars	61
Configuring Backup Policies	61
▼ <i>To create a policy using the Backup Policy Configuration Wizard</i>	62
▼ <i>To create a policy without using the wizard</i>	62
Introduction to Backup Policies	63

Attributes Tab Overview	63
Schedules Tab Overview	63
Clients Tab Overview	64
Backup Selections Tab Overview	64
Disaster Recovery Tab Overview	64
Changing Policies	65
▼ To add or change schedules in a policy	65
▼ To add or change clients in a policy	65
▼ To add or change backup selections in a policy	66
▼ To delete schedules, backup selections, or clients from a policy	67
▼ To copy and paste items	67
▼ To set the general policy attributes	67
What Type of Policy: Policy Attributes Tab	69
Policy Type	69
Policy Storage Unit	71
Notes on Specifying a Storage Unit	72
Policy Volume Pool	73
Volume Pool Override Example	74
Checkpoint Restart for Backup Jobs	75
Checkpoint Frequency	75
Checkpoint Restart Support	75
Checkpoint Restart for Restore Jobs	76
Limit Jobs Per Policy	77
Notes on Limit Jobs Per Policy	77
Job Priority	78
Active. Go Into Effect At	78
Backup Network Drives	79
Setup Example Using UNC Pathnames	79
Setup Example using Backup Network Drives Attribute	80
Follow NFS	81

Notes on Follow NFS	81
Advantages of Using Follow NFS Mounts	81
Disadvantages of Using Follow NFS	82
Cross Mount Points	82
Notes on Cross Mount Points	82
Cases That Can Require Separate Policies	83
How the Cross Mount Points Interacts With Follow NFS	83
Cross Mount Point Examples	84
Compression	85
Advantages of Using Compression	85
Disadvantages of Using Compression	85
How Much Compression Can Be Expected?	85
Encryption	87
Collect Disaster Recovery Information for Intelligent Disaster Recovery	87
Collect Disaster Recovery Information for Bare Metal Restore	87
Collect True Image Restore Information	88
Collect True Image Restore With Move Detection	88
What Happens During True Image Restores	89
Notes On True Image Restores and Move Detection	91
Allow Multiple Data Streams	92
When to Use Multiple Data Streams	92
Keyword Phrase	95
Advanced Client Options	96
When Will the Job Run: Schedules Tab	97
▼ <i>To create or change schedules</i>	97
Schedule Attributes Tab	98
Name	98
Type of Backup	98
Full Backup	99
Cumulative Incremental Backup	99

Differential Incremental Backup	99
User Backup	100
User Archive	100
Application Backup	100
Automatic Backup	100
Automatic Incremental Backup	100
Automatic Cumulative Incremental Backup	100
Automatic Differential Incremental Backup	101
Automatic Full Backup	101
Automatic Vault	101
Vault Catalog Backup	101
More on Incremental Backups	101
Determining Files Due for Backup on Windows Clients	104
Determining Files Due for Backup on UNIX Clients	105
Synthetic Backups	107
Calendar Schedule Type	108
Retries Allowed After Runday	108
Frequency Schedule Type	108
Guidelines for Setting Backup Frequency	109
Backup Frequency Determines Schedule Priority	110
Instant Recovery	110
Multiple Copies	110
▼ <i>To configure a schedule to create multiple copies</i>	111
▼ <i>To configure a disk staging relocation schedule to create multiple copies</i>	112
Override Policy Storage Unit	114
Override Policy Volume Pool	114
Retention	115
Guidelines for Setting Retention Periods	115
Precautions for Assigning Retention Periods	116
Changing Retention Periods	116

Mixing Retention Levels on Backup Volumes	117
Media Multiplexing	117
Final Destination Storage Unit	117
Final Destination Volume Pool	118
Start Window Tab	119
▼ <i>To create a window of time for a schedule</i>	119
Exclude Dates Tab	121
▼ <i>To exclude a date from the policy schedule</i>	121
Calendar Schedule Tab	122
Schedule by Specific Dates	122
▼ <i>To schedule a task on specific dates</i>	122
Schedule by Recurring Week Days	123
▼ <i>To schedule a recurring weekly task</i>	123
Schedule by Recurring Days of the Month	124
▼ <i>To schedule a recurring monthly task</i>	124
How Calendar Scheduling Interacts with Daily Windows	125
Examples of Automatic-Backup Schedules	126
Example 1: Various Automatic Backup Schedules	126
Example 2: Daily Schedules	130
Example 3: Using Various Backup Windows	134
Example 4: Long Backup Window	135
Example 5: Weekend Hours Only	136
Example 6: Full Backup Every Sunday	137
Considerations for User Schedules	139
Planning User Backup and Archive Schedules	139
Creating Separate Policies for User Schedules	139
Using a Specific Policy and User Schedule	141
Example Policies	142
Policy Planning Guidelines for Backups	142
Which Clients Will Be Backed Up: Clients Tab	149

▼ <i>To add a client to a policy</i>	149
▼ <i>To change a client list entry</i>	150
Installing Client Software on Trusting UNIX Clients	150
▼ <i>To install UNIX client software</i>	151
Installing Software on Secure UNIX Clients	152
Installing Software on Windows Clients	152
Configuring a Snapshot Method	153
Which Selections Will Be Backed Up: Backup Selections Tab	154
Overview on Creating Lists for Different Policy Types	154
Backup Selections List for Standard Policies	154
▼ <i>To add or change backup selections for a Standard, Exchange, or Lotus Notes policy</i> 155	
Backup Selections List for Database Policies	156
▼ <i>To create or change backup selections containing scripts for a database policy</i> ..	156
▼ <i>To add templates or scripts to the Backup Selections List</i>	156
Verifying the Backup Selections List	160
▼ <i>To verify a backup selections list</i>	160
Rules for Indicating Pathnames in the Backup Selections List	162
Pathname Rules for Microsoft Windows Clients	162
File Backups	162
Windows Disk-Image (Raw) Backups	164
Microsoft Windows Registry Backup	165
Hard Links to Files (NTFS volumes only)	166
Pathname Rules for UNIX Clients	168
Notes on UNIX Pathnames	168
Symbolic Links to Files or Directories	169
Hard Links to Directories	170
Hard Links to Files	170
UNIX Raw Partitions	172
Backup and Restore of Extended Attribute Files and Named Data Streams ..	174

Ramifications of Backing Up Extended Attributes or Named Data Streams .	176
Restoring Extended Attributes or Named Data Streams	176
▼ <i>To disable the restore of extended attribute files and named data streams</i>	177
Pathname Rules for NetWare NonTarget Clients	178
Pathname Rules for NetWare Target Clients	179
Pathname Rules for Clients Running Extension Products	180
Backup Selections List Directives: General Discussion	181
ALL_LOCAL_DRIVES Directive	181
SYSTEM_STATE Directive	181
Shadow Copy Components:\ Directive	182
Directives for Multiple Data Streams	183
Directives for Specific Policy Types	183
Backup Selections List Directives for Multiple Data Streams	184
NEW_STREAM Directive	184
ALL_LOCAL_DRIVES Directive	188
UNSET and UNSET_ALL Directives	189
Excluding Files from Backups	190
Files Excluded from Backups by Default	190
Excluding Files from Automatic Backups	191
Where Will the Catalog Data Be Located: Disaster Recovery Tab	193
Path	193
Logon	193
Password	193
Send in an E-mail Attachment	194
Identifying Critical Policies	194
Creating a Vault Policy	195
▼ <i>To create a Vault policy</i>	195
Performing Manual Backups	196
▼ <i>To perform a manual backup</i>	196
More About Synthetic Backups	198

Policy Considerations and Synthetic Backups	198
Two Types of Synthetic Backups	201
Synthetic Full Backups	201
Synthetic Cumulative Incremental Backups	202
Recommendations for Synthetic Backups	204
Notes on Synthetic Backups	205
Displaying Synthetic Backups in the Activity Monitor	208
Logs Produced During Synthetic Backups	208
Synthetic Backups and Directory and File Attributes	209
Chapter 4. NetBackup Catalogs	211
What is a NetBackup Catalog?	212
Parts of the Catalog	212
Image Database	213
Image Files	214
Image .f Files	214
NetBackup Relational Database	216
Catalog Protection	217
Catalog Backups	217
Online, Hot Catalog Backup Method	218
▼ <i>To configure an online, hot catalog backup using the Catalog Backup Wizard</i> ..	219
▼ <i>To configure an online, hot catalog backup using the Backup Policy Wizard</i> ..	224
▼ <i>To configure an online, hot catalog backup using the Policy utility</i>	225
Running Online, Hot Catalog Backups Concurrently with Other Backups ..	226
Notes on Catalog Policy Schedules	227
Offline, Cold Catalog Backup Method	228
▼ <i>To configure an offline, cold catalog backup using the Catalog Backup Wizard</i> ..	229
▼ <i>To configure an offline, cold catalog backup using the Actions menu</i>	236
Recovering the Catalog	247
Disaster Recovery E-mails and the Disaster Recovery File	247

Archiving the Catalog	248
Catalog Archiving Process	248
Creating a Catalog Archiving Policy	249
Policy Name	249
Deactivate Policy	249
Type of Backup	250
Retention Level Setting	250
Catalog Archiving Commands	251
Recommendations for Using Catalog Archiving	252
Using Vault with the Catalog Archiving Feature	253
Browsing Offline Catalog Archive	253
Extracting Images from the Catalog Archives	253
▼ <i>To extract images from the catalog archives based on a specific client</i>	253
Using the Catalog Utility	254
Searching for Backup Images	254
Notes on Searching for an Image	256
Messages Pane	257
Verifying Backup Images	257
▼ <i>To verify backup images</i>	257
Viewing Job Results	257
▼ <i>To view or delete a log file</i>	258
Promoting a Copy to a Primary Copy	258
▼ <i>To promote a backup copy to a primary copy</i>	258
▼ <i>To promote many copies to a primary copy</i>	259
▼ <i>To promote a backup copy to a primary copy using bpduplicate</i>	259
Duplicating Backup Images	260
Notes on Multiplexed Duplication	261
Procedure for Duplicating Backups	261
▼ <i>To duplicate backup images</i>	261
Jobs Displayed While Making Multiple Copies Concurrently	264

Importing NetBackup or Backup Exec Images	266
▼ <i>To initiate an import – Phase I</i>	266
▼ <i>To import backup images – Phase II</i>	268
Importing Expired Images	269
Importing Images from Backup Exec Media	270
Host Properties for Backup Exec	270
Considerations Concerning Importing Backup Exec Media	270
Differences Between Importing, Browsing and Restoring Backup Exec and Net-Backup Images	271
Expiring Backup Images	274
▼ <i>To expire a backup image</i>	274
Catalog Maintenance and Performance Optimization	275
Determining Catalog Space Requirements	275
▼ <i>To estimate the disk space required for a catalog backup</i>	275
File Size Considerations	276
Backing Up Catalogs Manually	278
▼ <i>To perform a manual online, hot catalog backup</i>	278
▼ <i>To perform a manual offline, cold catalog backup</i>	279
How Do I Know If a Catalog Backup Succeeded?	279
Strategies to Ensure Successful Catalog Backups	280
About the Binary Catalog Format	281
Catalog Conversion Utility	281
Binary Catalog File Limitations	281
Moving the Image Catalog	282
▼ <i>To move the image catalog</i>	282
Indexing the Catalog for Faster Access to Backups	283
Indexing the Header Files	283
Indexing the Image .f File	283
Compressing and Uncompressing the Image Catalog	286
Uncompressing the Image Catalog	287

▼ <i>To uncompress client records</i>	287
Chapter 5. Viewing NetBackup Reports	289
Introduction to the Reports Utility	290
Reports Menu Bar	291
▼ <i>To run a report</i>	291
Reports Window	292
Report Toolbar	292
Report Contents Pane	292
Shortcut Menus	292
Reports Settings	293
Date/Time Range	293
Client	293
Media Server	293
Job ID	293
Media ID	294
Volume Pool	294
Verbose Listing	294
Run Report	294
Stop Report	294
NetBackup Report Types	295
Status of Backups Report	295
Client Backups Report	296
Problems Report	297
All Log Entries Report	298
Media Lists Report	299
Media Contents Report	302
Images on Media Report	303
Media Logs Report	305
Media Summary Report	305

Media Written Report	306
Using the Troubleshooter Within Reports	307
▼ <i>To run Troubleshooter within Reports</i>	307
Chapter 6. Monitoring NetBackup Activity	309
Introduction to the Activity Monitor	310
Activity Monitor Menu Bar	311
Actions Menu	311
▼ <i>To show or hide column heads</i>	312
▼ <i>To monitor the detailed status of selected jobs</i>	313
▼ <i>To delete completed jobs</i>	313
▼ <i>To cancel uncompleted jobs</i>	313
▼ <i>To restart a completed job</i>	313
▼ <i>To suspend a restore or backup job</i>	314
▼ <i>To resume a suspended or incomplete job</i>	314
▼ <i>To print job detail information from a list of jobs</i>	314
▼ <i>To export Activity Monitor data to a text file</i>	314
▼ <i>To run Troubleshooter within the Activity Monitor</i>	314
Activity Monitor Toolbar	315
Status Bar	315
Setting Activity Monitor Options	316
Jobs Tab	318
Parent Jobs	318
▼ <i>To view job details</i>	319
Daemons Tab	325
Other VERITAS Services	327
▼ <i>To monitor NetBackup daemons</i>	327
▼ <i>To start or stop a daemon</i>	328
Processes Tab	329
Monitoring NetBackup Processes	331

▼ <i>To monitor NetBackup processes</i>	331
Process Details	331
Media Mount Errors	332
Queued Media Mount Errors	332
Cancelled Media Mount Errors	332
Managing the Jobs Database	332
Retaining Job Information in the Database	333
Changing the Default on a Permanent Basis	333
Changing the Default Temporarily	333
bpdjobs Debug Log	335
Customizing bpdjobs Output	335
Chapter 7. Configuring Host Properties	337
Introduction to Host Properties	338
Host Properties Menu Bar	339
Viewing Host Properties	339
▼ <i>To view master server, media server, or client properties</i>	340
Changing Host Properties	340
Interpreting the Initial Settings	340
Selecting Multiple Hosts	342
▼ <i>To simultaneously change the properties on multiple hosts</i>	342
Required Permissions	343
Master Server, Media Server, and Client Host Properties	344
Access Control Properties	344
VERITAS Security Services (VxSS)	344
VxSS Tab within Access Control Properties Dialog	345
VxSS Networks List	345
Add Button	346
Remove Button	347
Authentication Domain Tab within Access Control Properties Dialog	347

Add Button	348
Remove Button	348
Authorization Service Tab within Access Control Properties Dialog	349
Host Name	349
Customize the Port Number of the Authorization Service	349
Authorization Properties	350
User	350
Host	350
Domain\Group	350
Group/Domain Type	351
User must be an OS Administrator	351
Backup Exec Tape Reader Properties	352
Add Button	352
GRFS Advertised Name	352
Actual Client Name	353
Actual Path	353
Change Button	353
Remove Button	353
Bandwidth Properties	354
How Bandwidth Limiting Works	354
Bandwidth Throttle Setting for the Range of IP Addresses	354
From IP Address	355
To IP Address	355
Bandwidth	355
Bandwidth Throttle Settings List	355
Add Button	355
Remove Button	355
Notes on Bandwidth Limiting	355
Busy File Properties	357
Working Directory	357

Operator's E-mail Address	357
Process Busy Files	357
File Action File List	358
Add Button	358
Add to All Button	358
Remove Button	358
Busy File Action	358
Retry Count	358
Clean-up Properties	359
Keep Logs	359
Delete Vault Logs	359
Keep True Image Restoration (TIR) Information	359
Move Restore Job From Incomplete State to Done State	360
Move Backup Job from Incomplete State to Done State	360
Client Attributes Properties	362
Allow Client Browse	362
Allow Client Restore	362
Clients List	362
General Tab	363
Maximum Data Streams	363
Browse and Restore Ability	364
Free Browse	364
Connect Options Tab	364
BPCD Connect Back	364
Ports	365
Daemon Connection Port	365
Windows Open File Backup Tab	366
Add and Remove Buttons	366
Enable Windows Open File Backups for this Client	367
Use VERITAS Volume Snapshot Provider (VSP)	367

Use Microsoft Volume Shadow Copy Service (VSS)	367
Individual Drive Snapshot	367
Global Drive Snapshot	368
Abort Backup on Error	368
Disable Snapshot and Continue	369
Client Name Properties	370
Client Name	370
Client Settings (NetWare) Properties	371
Back Up Migrated Files	371
Uncompress Files Before Backing Up	371
Keep Status of User-directed Backups, Archives, and Restores	371
Client Settings (UNIX) Properties	372
Locked File Action	372
Keep Status of User-directed Backups, Archives, and Restores	372
Reset File Access Time to the Value Before Backup	372
Megabytes of Memory to Use for File Compression	373
Use VxFS File Change Log for Incremental Backups	373
Default Cache Device Path for Snapshots	375
Do Not Compress Files Ending With	375
Add Button	375
Add to All Button	376
Remove Button	376
Client Settings (Windows) Properties	377
General Level Logging	377
TCP Level Logging	377
Wait Time Before Clearing Archive Bit	377
Use Change Journal in Incrementals	378
Incrementals Based on Timestamp	379
Incrementals Based on Archive Bit	380
Time Overlap	380

Communications Buffer	380
User Directed Timeout	381
Maximum Error Messages for Server	381
Keep Status of User-directed Backups, Archives, and Restores	381
Perform Default Search for Restore	381
Encryption Properties	382
Encryption Permissions	382
Enable Encryption	382
Enable Standard Encryption	383
Client Cipher	383
Use Legacy DES Encryption	383
Encryption Strength	383
Encryption Libraries	383
Encryption Key File	384
Exchange Properties	385
Mailbox for Message Level Backup and Restore	385
Enable Single Instance Backup for Message Attachments	386
Exclude Lists Properties	387
Use Case Sensitive Exclude List	387
Exclude List	387
Exceptions to the Exclude List	388
Add Buttons	388
Add to All Buttons	388
Remove Buttons	389
Shared Fields in Exclude Lists	389
Policy	389
Schedule	390
Files/Directories	390
Exclude Lists for Specific Policies or Schedules	390
▼ To create an exclude or include list for a specific policy	390

Syntax Rules for Exclude Lists	392
Traversing Excluded Directories	394
Firewall Properties	396
Default Connect Options	396
Host List	398
Attributes for the Selected Hosts	398
▼ <i>To set up vnetd between a server and a client</i>	400
▼ <i>To set up vnetd between servers</i>	400
▼ <i>To enable logging for vnetd</i>	400
Example Setup for Using the vnetd Port	401
Minimum Master Server Outgoing Connections	402
Minimum Media Server Outgoing Connections	403
Minimum Enterprise Media Server Outgoing Connections	405
Minimum Client Outgoing Connections	406
Minimum Java Server or Windows Administration Console Outgoing Connections	409
Minimum Volume Database Host Outgoing Connections	410
General Server Properties	411
Delay on Multiplexed Restores	411
Check the Capacity of Disk Storage Units	411
Must Use Local Drive	411
Use Direct Access Recovery for NDMP Restores	412
Media Host Override	412
Add Button	412
Add to All Button	413
Change Button	413
Remove Button	413
▼ <i>To force restores to go to a specific server</i>	413
Global Attributes Properties	414
Job Retry Delay	414

Schedule Backup Attempts	414
Policy Update Interval	414
Priority of Restore Jobs	415
Maximum Jobs per Client	415
Maximum Backup Copies	416
Compress Catalog Interval	416
Maximum Vault Jobs	416
Administrator's E-mail Address	417
Setting Up E-Mail Notifications	417
Logging Properties	418
Types of Logging	418
Enable Robust Logging	419
Global Logging Level	419
Process Specific Overrides	420
Diagnostic Logging Levels for NetBackup Services	421
Lotus Notes Properties	422
Path	422
INI File	422
Media Properties	423
Allow Media Overwrite	423
Allow Multiple Retentions Per Media	424
Allow Backups to Span Media	424
Enable SCSI Reserve/Release	425
Enable Standalone Drive Extension	425
Enable Job Logging	425
Media ID Prefix (Non-robotic)	425
Media Unmount Delay	425
Media Request Delay	426
NDMP Global Credentials Properties	427
User Name	427

Password and Confirm Password	427
NetWare Client Properties	428
Network Properties	429
NetBackup Client Service Port (BPCD)	429
NetBackup Request Service Port (BPRD)	429
Announce DHCP Interval	429
Open File Backup (NetWare Client) Properties	430
Enable Open File Backup During Backups	430
OTM Properties	430
Port Ranges Properties	431
Use Random Port Assignments	431
Client Port Window	431
Client Reserved Port Window	432
Server Port Window	432
Server Reserved Port Window	432
Restore Failover Properties	433
Alternate Restore Failover Machines List	433
Add Button	434
▼ <i>To add or change a media server to the Alternate Restore Failover Machine list</i>	434
Change Button	434
Remove Button	434
Retention Periods Properties	436
Value	436
Units	436
Retention Periods List	436
Schedules List	436
Impact Report Button	437
▼ <i>To change a retention period</i>	437
Note on Redefining Retention Periods	438
Servers Properties	439

Master Server	439
Additional Servers	439
Media Servers	439
Restricting Administrative Privileges of Media Servers	440
Multiple Masters Sharing One Enterprise Media Manager Host	441
SharePoint 2003 Properties	443
User ID	443
Password	443
Consistency Check Before Backup	443
Continue with Backup if Consistency Check Fails	443
Timeouts Properties	444
Client Connect Timeout	444
Backup Start Notify Timeout	444
File Browse Timeout	444
Use OS Dependent Timeouts	445
Media Mount Timeout	445
Client Read Timeout	445
Backup End Notify Timeout	446
Media Server Connect Timeout	446
Universal Settings Properties	447
Restore Retries	447
Browse Timeframe for Restores	447
Last Full Backup	448
Use Specified Network Interface	448
Use Preferred Group for Enhanced Authorization	449
Allow Server File Writes	450
Accept Connections on Non-reserved Ports	450
Enable Performance Data Collection	451
Client Sends Mail	451
Server Sends Mail	451

Client Administrator's E-mail	451
UNIX Client Properties	452
UNIX Server Properties	452
NFS Access Timeout	452
VERITAS Products Properties	453
VSP (Volume Snapshot Provider) Properties	454
VSP Overview	454
Logging VSP Messages	455
Cache File Volume List	456
VSP Volume Exclude List	456
Customize Cache Size	459
Cache Size	459
Maximum Cache Size	459
Busy File Wait	461
Busy File Timeout	461
Using VSP with Databases	461
Windows Client Properties	463
Chapter 8. Managing NetBackup	465
Powering Down and Rebooting NetBackup Servers	466
▼ <i>To power down a server</i>	466
▼ <i>To shut down all NetBackup daemons</i>	466
▼ <i>To start up all NetBackup daemons</i>	467
▼ <i>To reboot a NetBackup master server</i>	467
▼ <i>To reboot a NetBackup media server</i>	467
Displaying Active Processes with bpps	467
Administering NetBackup Licenses	469
▼ <i>To access license keys for a NetBackup server</i>	469
▼ <i>To add new license keys</i>	470
▼ <i>To print license key lists</i>	470

▼ <i>To delete license keys</i>	471
▼ <i>To view the properties of one license key</i>	471
▼ <i>To export license keys</i>	472
Using the NetBackup License Utility to Administer Licenses	473
▼ <i>To start the NetBackup License Key utility</i>	473
Administering a Remote Master Server	474
Adding a NetBackup Server to a Server List	474
▼ <i>To add a server to a UNIX server list</i>	475
▼ <i>To add a server to a Windows server list</i>	477
Choosing a Remote Server to Administer	478
▼ <i>To use the Change Server command to administer a remote server</i>	478
▼ <i>To indicate a remote system upon log in</i>	479
Administering via a NetBackup Client	480
Using the Remote Administration Console	481
If You Cannot Access a Remote Server	482
Using the NetBackup-Java Windows Display Console	483
Authorizing NetBackup-Java Users on Windows	483
Restricting Access on Windows	483
Configuring the NetBackup-Java Administration Console	484
NetBackup-Java Administration Console Architectural Overview	484
Authorizing NetBackup-Java Users	486
Authorization File Characteristics	487
Configuring Nonroot Usage	490
Authorizing Nonroot Users for Specific Applications	490
Capabilities Authorization for jbpSA	491
Runtime Configuration Options	491
BPJAVA_PORT, VNETD_PORT	492
FORCE_IPADDR_LOOKUP	492
INITIAL_MEMORY, MAX_MEMORY	494
MEM_USE_WARNING	495

NBJAVA_CLIENT_PORT_WINDOW	495
NBJAVA_CONNECT_OPTION	496
USE_NBJAUTH_WITH_ENHAUTH	496
Configuration Options Relevant to jnbSA and jbpSA	496
Logging Command Lines Used by the NetBackup Interfaces	497
Customizing jnbSA and jbpSA with bp.conf Entries	497
NetBackup-Java Performance Improvement Hints	497
What it Means to be Running the Java Console Locally on a UNIX Platform ...	498
What it Means to be Running the Console Locally on a Windows Platform ...	498
How do I Run a Console Locally and Administer a Remote Server?	498
How do I Make the Console Perform Even Better?	499
Is Performance Better When Remotely Displaying Back or Running Locally? ..	500
Administrator's Quick Reference	502
Managing Client Restores	504
Server-Directed Restores	504
Client-Redirected Restores	505
How NetBackup Enforces Restore Restrictions	505
Allowing All Clients to Perform Redirected Restores	506
Allowing a Single Client to Perform Redirected Restores	506
Allowing Redirected Restores of a Specific Client's Files	507
Redirected Restore Examples	507
Restoring Files and Access Control Lists	512
Restoring Files that Possess ACLs	512
Restoring Files without Restoring ACLs	512
▼ <i>To restore files without restoring ACLs</i>	513
Setting Client List and Restore Permissions	513
Adding Clients to the NetBackup Client Database	513
▼ <i>To create an entry in the client catalog</i>	513
▼ <i>To delete and list entries in the client catalog</i>	514
Setting the List and Restore Permissions	514





Examples	516
Improving Search Times by Creating an Image List	517
Set Original atime for Files During Restores	518
Checkpoint Restart for Restore Jobs	518
Suspending and Resuming a Restore Job	518
Limitations to Checkpoint Restart for Restore Jobs	519
Restoring System State	519
Important Notes on System State	520
▼ <i>To restore the System State</i>	520
Goodies Scripts	522
Server Independent Restores	522
Supported Configurations	522
Methods for Performing Server Independent Restores	524
Method 1: Modifying the NetBackup Catalogs	524
▼ <i>To modify catalogs when the server that wrote the media is available</i>	525
▼ <i>To modify catalogs when the server that wrote the media is not available</i>	525
Method 2: Overriding the Original Server	526
▼ <i>To enable overriding of the original server for restores</i>	526
▼ <i>To manually override the original server for restores</i>	526
Method 3: Automatic Failover to Alternate Server	527
▼ <i>To enable automatic failover to an alternate server</i>	527
Notes on Server Independent Restores	528
Configuring NetBackup Ports	530
Port Descriptions	530
Load Balancing	533
Using NetBackup with Storage Migrator	535
Set a Large Enough Media Mount Timeout	535
Do Not Use the RESTORE_ORIGINAL_ETIME File	535
Do Not Use the Following Client bp.conf File Settings	536

Appendix A. NetBackup Relational Database	537
Installation Overview	538
NetBackup Master Server Installation	540
Relocating the NetBackup Database	540
server.conf	540
databases.conf	541
vxdbms_env.csh, vxdbms_env.sh	542
/bin	542
/charsets	543
/data	543
/lib	544
/log	544
/res	544
/scripts	544
/staging	544
/tix	544
NetBackup Configuration Entry	545
Sybase ASA Server Management	546
▼ <i>To start and stop the ASA daemon</i>	546
▼ <i>To start/stop individual databases</i>	546
Clustering	547
Post-installation Tasks	548
Changing the Database Password	548
▼ <i>To change the database password</i>	548
Moving NBDB Database Files After Installation	549
▼ <i>To move the NBDB and BMRDB database files</i>	549
Adding a Mirrored Transaction Log	550
▼ <i>To create a mirrored transaction log</i>	550
Creating the NBDB Database	551
▼ <i>To manually create the NBDB database</i>	551

Additional create_nbdb Options	552
Backup and Recovery Procedures	553
Using the Online, Hot Catalog Backup Method	554
Using the Offline, Cold Catalog Backup Method	555
Transaction Log Management	556
Catalog Recovery	556
Additional Command Lines for Backup and Recovery of the Relational Databases	
557	
nbdb_backup	557
nbdb_restore	557
Database Unloading Tool	558
Terminating Database Connections	558
▼ <i>To terminate connections</i>	558
Moving the NetBackup Database from One Host to Another	560
Appendix B. NearStore Storage Unit Considerations	563
Required Software, Hardware, and Licenses	564
NearStore and SnapVault Topics	564
NearStore Storage Units and SnapVault Storage Units Cannot Share Volumes .	
564	
Cleaning Up Configured qtrees	564
NearStore SnapVault Snapshot Schedules	565
Advantages of the NearStore Storage Unit	566
NearStore Configuration	568
NearStore Authentication	568
▼ <i>To authenticate the NetBackup media server</i>	568
▼ <i>To create a root NearStore user name and password</i>	568
▼ <i>To create a non-root NearStore user name and password</i>	569
▼ <i>To verify that the NearStore credentials have been entered into the NetBackup EMM</i>	
<i>database</i> 569	
NearStore Disk Storage Unit Properties	570

Viewing the Backup Image	572
Disk Consumption	574
Logging Information	575
Troubleshooting	575
Index	577





Preface

This guide describes how to configure and manage the operation of VERITAS NetBackup™ Server and VERITAS NetBackup Enterprise Server for UNIX and Linux platforms. See the *NetBackup Release Notes* for a list of the hardware and operating system levels that NetBackup supports.

To determine the version and release date of installed software, see the `version` file located here in `/usr/openv/netbackup`

Getting Help

You can find answers to questions and get help from the NetBackup documentation and from the VERITAS technical support web site.

Finding NetBackup Documentation

A list of the entire NetBackup documentation set appears as an appendix in the *NetBackup Release Notes*. All NetBackup documents are included in PDF format on the NetBackup Documentation CD.

For definitions of NetBackup terms, consult the online glossary.

▼ To access the NetBackup online glossary

1. In the NetBackup Administration Console, click **Help > Help Topics**.
2. Click the **Contents** tab.
3. Click **Glossary of NetBackup Terms**.

Use the scroll function to navigate through the glossary.

Accessing the VERITAS Technical Support Web Site

The address for the VERITAS Technical Support Web site is <http://support.veritas.com>.

The VERITAS Support Web site lets you do any of the following:

- ◆ Obtain updated information about NetBackup, including system requirements, supported platforms, and supported peripherals
- ◆ Contact the VERITAS Technical Support staff and post questions to them
- ◆ Get the latest patches, upgrades, and utilities
- ◆ View the NetBackup Frequently Asked Questions (FAQ) page
- ◆ Search the knowledge base for answers to technical support questions
- ◆ Receive automatic notice of product updates
- ◆ Find out about NetBackup training
- ◆ Read current white papers related to NetBackup

From <http://support.veritas.com>, you can complete various tasks to obtain specific types of support for NetBackup:

1. Subscribe to the VERITAS Email notification service to be informed of software alerts, newly published documentation, Beta programs, and other services.
 - a. From the main <http://support.veritas.com> page, select a product family and a product.
 - b. Under Support Resources, click **Email Notifications**.

Your customer profile ensures you receive the latest VERITAS technical information pertaining to your specific interests.
2. Locate the telephone support directory at <http://support.veritas.com> by clicking the **Phone Support** icon. A page appears that contains VERITAS support numbers from around the world.

Note Telephone support for NetBackup is only available with a valid support contract. To contact VERITAS for technical support, dial the appropriate phone number listed on the Technical Support Guide included in the product box and have your product license information ready for quick navigation to the proper support group.

3. Contact technical support using e-mail.

- a. From the main <http://support.veritas.com> page, click the **E-mail Support** icon.
A wizard guides you to do the following:
 - ◆ Select a language of your preference
 - ◆ Select a product and a platform
 - ◆ Provide additional contact and product information, and your message
 - ◆ Associate your message with an existing technical support case
- b. After providing the required information, click **Send Message**.

Contacting VERITAS Licensing

For license information, you can contact us as follows:

- ◆ Call 1-800-634-4747 and select option 3
- ◆ Fax questions to 1-650-527-0952
- ◆ In the Americas, send e-mail to amercustomercore@veritas.com.
In the Asia and Pacific areas, send email to apaccustomercore@veritas.com.
In all other areas, send email to internationallicense@veritas.com.

Accessibility Features

NetBackup contains features that make the user interface easier to use by people who are visually impaired and by people who have limited dexterity. Accessibility features include:

- ◆ Support for assistive technologies such as screen readers and voice input (Windows servers only)
- ◆ Support for keyboard (mouseless) navigation using accelerator keys and mnemonic keys

For more information, see the *NetBackup Installation Guide*.

Comment on the Documentation

Comment on the Documentation

Let us know what you like and dislike about the documentation. Were you able to find the information you needed quickly? Was the information clearly presented? You can report errors and omissions or tell us what you would find useful in future versions of our manuals and online help.

Please include the following information with your comment:

- ◆ The title and product version of the manual on which you are commenting
- ◆ The topic (if relevant) on which you are commenting
- ◆ Your comment
- ◆ Your name

Email your comment to NBDocs@veritas.com.

Please only use this address to comment on product documentation. See “Getting Help” in this preface for information on how to contact Technical Support about our software.

We appreciate your feedback.



Introduction to NetBackup

1

This chapter provides an introduction to NetBackup software, the interfaces available, and the distributed architecture of NetBackup. This chapter contains the following sections:

- ◆ “Overview” below
- ◆ “NetBackup Administration Interfaces” on page 3
- ◆ “NetBackup Administration Console Setup” on page 3
- ◆ “Using the NetBackup Administration Console” on page 9
- ◆ “Configuring NetBackup Without Wizards” on page 20

Overview

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

NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. By carefully scheduling backups, an administrator can achieve systematic and complete backups over a period of time, optimizing network traffic during off-peak hours. The backups can be *full* (backing up all client files) or *incremental* (backing up only the files that have changed since the last backup).

If allowed by the NetBackup administrator, users can backup, 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.)

The Media Manager component of the NetBackup software manages tape and optical disk storage devices. Media Manager is designed so that other VERITAS storage products can also share secondary storage devices. NetBackup provides extensive support for automated storage so human intervention is rarely required.


NetBackup includes both the server and client software:

- ◆ Server software resides on the computer that manages the storage devices.

Overview

- ◆ Client software resides on computer(s) containing data to back up. (Servers also contain client software and can be backed up.)

NetBackup Storage Domain



See the *NetBackup Release Notes* for lists of supported platforms for NetBackup servers and clients.

NetBackup accommodates multiple servers working together under the administrative control of one NetBackup master server. The master server manages backups, archives, and restores. Media servers are directed by the master server and provide additional storage by allowing NetBackup to use the storage devices that they control. Media servers can also increase performance by distributing the network load. A master server may also function as a media server.

During a backup or archive, the client sends backup data across the network to a NetBackup server that manages the type of storage 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 Administration Interfaces

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

- ◆ NetBackup Administration Console

A Java-based, graphical-user interface that is started by running the `jnbSA` command. This is the recommended interface and is the one referred to by most procedures and examples in this manual. (See “NetBackup Administration Console Setup” on page 3. For more details on the console, see “NetBackup-Java Administration Console Architectural Overview” on page 484.)

- ◆ Character-based, menu interface

A character-based, menu interface that is started by running the `bpadm` command. The `bpadm` interface can be used from any terminal (or terminal emulation window) that has a `termcap` or `terminfo` definition. (See Chapter 4 in *NetBackup System Administrator's Guide, Volume II*.)

- ◆ Command line

NetBackup commands can be entered at the system prompt or used in scripts. For complete information on all NetBackup commands, see *NetBackup Commands for UNIX and Linux*. To view the commands online, use the UNIX `man` command.

All NetBackup administrator programs and commands require root-user privileges by default. If it is necessary to have nonroot administrators, see “Configuring Nonroot Usage” on page 490.

It is also possible to display the console on a Java-capable UNIX platform and display it back to a Windows system by using third-party X terminal emulation software.

NetBackup Administration Console Setup

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 system by running `/usr/opensv/java/jnbSA &`

The `jnbSA` command is described in *NetBackup Commands for UNIX and Linux*.

- ◆ On a supported Windows system that has the NetBackup-Java Windows Display Console installed. The Windows Display Console is not automatically installed on Windows systems. Installation is available on the main NetBackup for Windows installation screen.

NetBackup Administration Console Setup

The startup procedures are explained below. For configuration information, see “Configuring the NetBackup-Java Administration Console” on page 484.

Always set the window manager so a window becomes active only when clicked. Do not enable auto-focus, which causes a window to be activated by simply moving the pointer over the window. The NetBackup-Java interfaces do not run properly with auto-focus enabled.

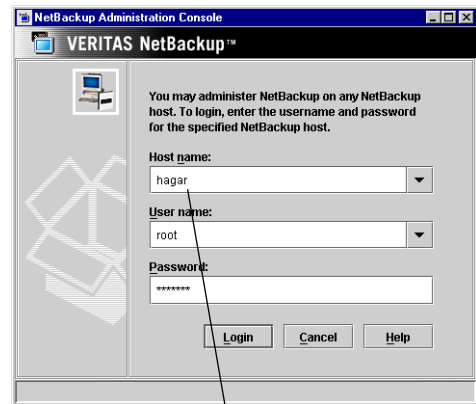
The following are general instructions for correctly setting up the focus on a CDE (Common Desktop Environment) window manager, which is the preferred window manager for NetBackup-Java applications.

▼ To prepare a CDE (Common Desktop Environment) for NetBackup-Java interfaces

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

▼ **To start the NetBackup-Java Administration Console on a NetBackup-Java capable UNIX system**

1. Log in as `root` on the NetBackup client or server where you want to start the NetBackup Administration Console. The client or server must be NetBackup-Java capable.
2. Start the console by entering:
`/usr/opensv/java/jnbSA &`
The login screen appears.
3. Type the name of the UNIX master server host where you initially want to manage NetBackup.



Specified host must be running same NetBackup version as machine where the console is started

Note The NetBackup server or client you specify on the login dialog of the NetBackup-Java console must be running the same version of NetBackup as is installed on the machine where you start the NetBackup-Java console.

4. Specify your user name and password, then click **Login**.
This logs you into the NetBackup-Java application server program on the specified server. The NetBackup Administration Console appears. The console program continues to communicate through the server you specified for the remainder of the current session.
5. Start a utility by clicking on it in the left pane.
6. If you wish to administer another NetBackup server, you can select **File > Change Server** to select a remote NetBackup server on which to make configuration changes.

Note The NetBackup Administration Console supports remote X Windows display only between same-platform systems. For example, assume you are on a Solaris system named *tiger* and the NetBackup-Java software is on a Solaris system named *shark*. Here, you can display the interface on *tiger* by performing an `rlogin` to *shark* and running the command `jnbSA -d tiger`. However, if *shark* were an HP system, you could display `jnbSA` only directly on *shark*.
In addition, the system on which the console is displayed must be running a version

NetBackup Administration Console Setup

of the operating system supported by the console. Refer to the NetBackup release notes for supported versions, including any required patches. The `jnbSA` command is described in *NetBackup Commands for UNIX and Linux*.

Running the Java-Based Windows Display Console

The NetBackup-Java Windows Display Console is provided with NetBackup software. The Windows Display Console offers the user the NetBackup-Java interface in order to administer UNIX NetBackup servers where a NetBackup-Java capable UNIX system is not available. See the *NetBackup Installation Guide* for information on installing the Windows Display Console.

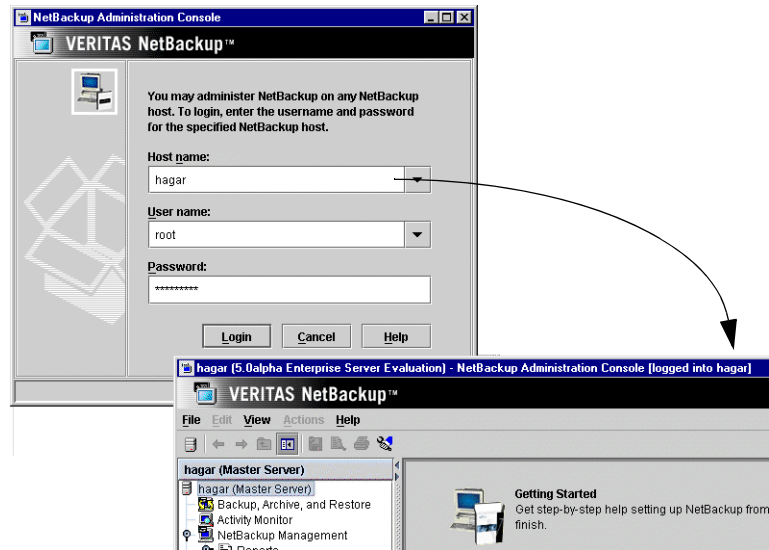
The Windows Display Console can also be used to directly administer a NetBackup UNIX or Windows server. It is also possible to use a point-to-point (PPP) connection between the display console and other servers in order to perform remote administration.

▼ To start the Windows Display Console

1. On a Windows system where the Windows Display Console is installed and configured, select **Start > Programs > VERITAS NetBackup > NetBackup-Java Version 6.0**.
2. The login screen for the NetBackup Administration Console appears, displaying the host name. To perform remote administration, log into another server by typing the name of another host in the **Host name** field, or by selecting a host name from the drop-down list.
3. Type your user name and password. When logging 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. This is not required if the NetBackup host is not a member of a domain.



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 specified in the login screen for the remainder of the current session.

Note The default host is the last host that was successfully logged into. Names of other hosts that have been logged into are available for selection from the drop-down list.

Administering Remote Servers

If a site contains more than one NetBackup master server, the systems can be configured so that multiple servers can be accessed from one NetBackup Administrator Console. Indicating a remote server can be done using one of the following methods:

- ◆ Using the **File > Change Server** menu command.
- ◆ Using the NetBackup-Java Administration Console and indicating a remote system upon NetBackup login.

For more information on remote administration, see “Administering a Remote Master Server” on page 474.

Administering Backlevel NetBackup Servers

There are five general methods to accomplish backlevel administration tasks in NetBackup 6.0. The methods are listed below, though the ordering does not imply any preference.

Earlier Versions of the NetBackup-Java Administration Console on Supported UNIX Platforms

Use the earlier versions of the NetBackup-Java Administration Console installed on the supported UNIX GUI platforms. The earlier versions available in a release are all of those supported in a mixed version environment with the current release, (all those back to and including the last major release version). In the 6.0 release, the 5.0MP4 (or later) and 5.1 versions of the console are available.

Earlier Versions of the NetBackup-Java Administration Console on Windows Platforms

Use the earlier versions of the NetBackup-Java Administration Console installed on the supported Windows platforms.

Remote Display-back from UNIX Servers

Use the UNIX remote display-back capabilities for UNIX servers that can run NetBackup-Java, possibly in conjunction with tools such as Exceed or VNC.

Remote Display-back from Windows Servers

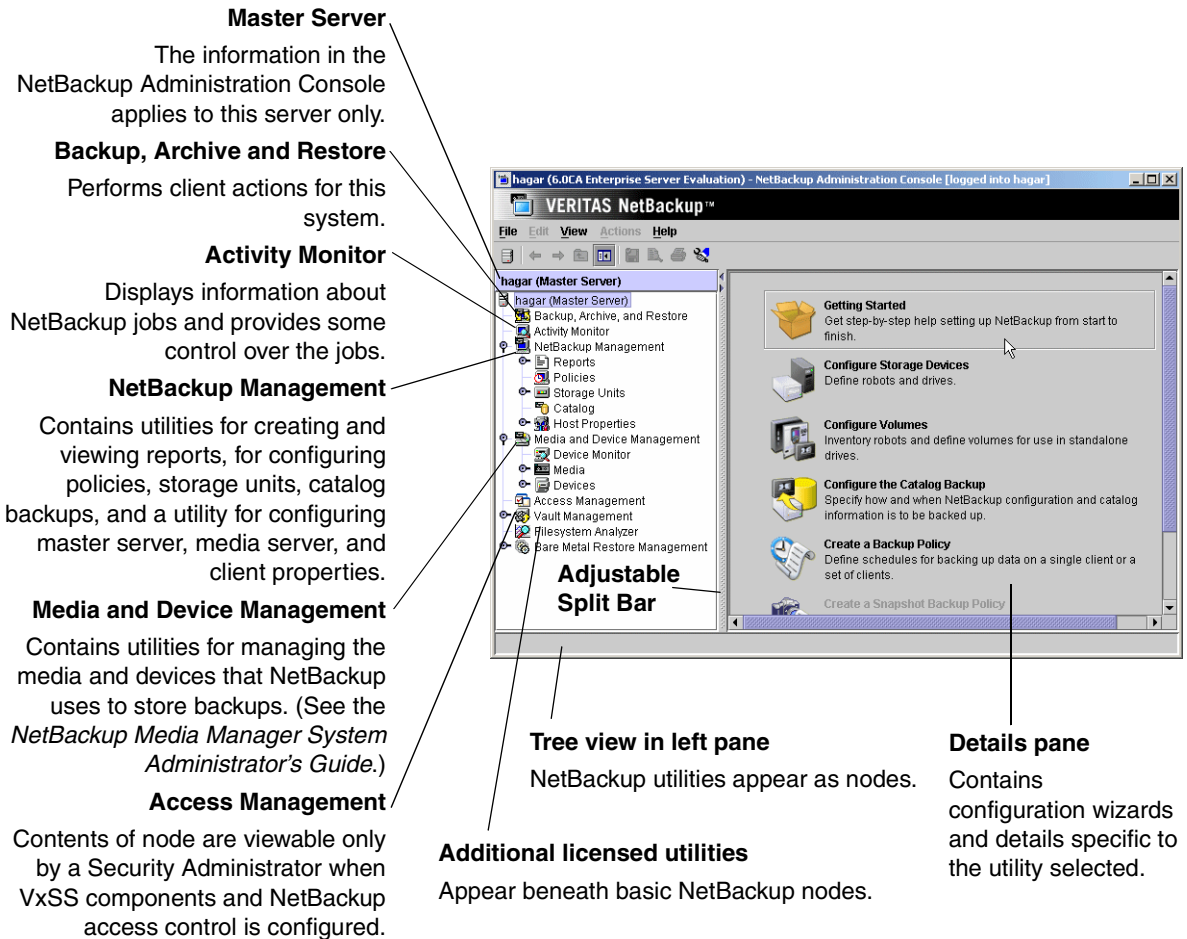
Use the NetBackup Administration Console for Windows with the remote display-back capabilities on Windows NetBackup servers with tools like Windows Terminal Services or Remote Desktop.

At the Console of the Backlevel Server

Use the relevant NetBackup-Java Administration Console from the backlevel server's console.

Using the NetBackup Administration Console

The NetBackup Administration Console provides a graphical user interface through which the administrator can manage NetBackup. The interface can run on any NetBackup-Java capable system.



You may also administer NetBackup through a character-based, menu interface or through a command line. Each method is described in "NetBackup Administration Interfaces" on page 3.

The following sections describe the utilities and menus that appear in the NetBackup Administration Console.

NetBackup Configuration Wizards

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

◆ Getting Started Wizard

Use the **Getting Started Wizard** if you are configuring NetBackup for the first time. It leads you through the necessary steps and other wizards to get you up and running with a working NetBackup configuration. The **Getting Started Wizard** is comprised of the following wizards, which can also be run separately, outside of the **Getting Started Wizard**:

- ◆ Configure Storage Devices
- ◆ Configure Volumes
- ◆ Configure the Catalog Backup
- ◆ Create a Backup Policy

◆ Configure Storage Devices

Use the **Device Configuration Wizard** to guide you through the entire process of configuring a device and storage unit.

◆ Configure Volumes

Use the **Volume Configuration Wizard** to guide you through the entire process of configuring removable media.

◆ Configure the Catalog Backup

Use the **NetBackup Catalog Backup Wizard** to set up your catalog backups, which are essential to recovering your data in case of a server failure or crash.

◆ Create a Backup Policy

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

◆ Recover the Catalog

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

Backup, Archive, and Restore

Use the **Backup, Archive, and Restore** utility to perform backups and archives for this system, and restores for this system and other clients.

Users can back up, archive, and restore files, directories, and formatted raw partitions that reside on their own client computer. A user can restore files at any time but can back up and archive files 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.

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).

Documentation for the NetBackup client is available as online Help from the **Backup, Archive, and Restore** interface.

Activity Monitor

Use the **Activity Monitor** utility to monitor and control NetBackup jobs, daemons, and processes. For more information see Chapter 6, “Monitoring NetBackup Activity” on page 309.

NetBackup Management

This manual describes the applications and utilities listed under **NetBackup Management** in the NetBackup Administration Console tree.

Utilities listed under **Media and Device Management** are described in the *Media Manager System Administrator's Guide*.

The **Access Management** utility is described in the *System Administrator's Guide, Volume II*.

Additional nodes appear if optionally-licensed products such as NetBackup Vault and NetBackup Bare Metal Restore are installed.

The following sections describe items found under **NetBackup Management**.

Reports

Use **Reports** to compile information for verifying, managing, and troubleshooting NetBackup operations. For more information see Chapter 5, “Viewing NetBackup Reports” on page 289.

Policies

Use **Policies** to create and specify the backup policies which define the rules for backing up a specific group of one or more clients. For example, the backup policy specifies when automatic backups will occur for the clients specified in the policy and when users can perform their own backups. The administrator can define any number of backup policies, each of which can apply to one or more clients. A NetBackup client must be covered by at least one backup policy.

The properties of a backup policy include, but are not limited to, the following:

- ◆ General attributes that define the:
 - ◆ *Policy type* of the policy. The policy type determines, for example, whether the clients to be backed up by this policy will be Windows or UNIX clients, or whether this policy will perform an online, hot catalog backup.
 - ◆ *Priority* of backups for this policy relative to backups for other policies.
 - ◆ *Storage unit* to use for backups of clients covered by this policy.
 - ◆ *Volume pool* to use for backups performed according to this policy. A volume pool is a set of volumes that the administrator can assign to specific backup policies or schedules. For example, it is possible to have one volume pool for weekly backups and another for quarterly backups.
- ◆ Schedules that control when backups and archives can occur for the clients.
- ◆ List of client computers covered by the policy.
- ◆ List of files to include in automatic backups of the clients. The backup selection list does not affect user backups because the user selects the files.

As mentioned above, each backup policy has a set of schedules. These schedules control when automatic backups can start and when users can start a backup or archive. Each schedule contains attributes that include:

- ◆ *Type of schedule*. Specify schedules for automatic full or incremental backups or user backups or archives. There are also schedule types that apply only when separately-priced options are installed (for example, a backup schedule for Microsoft Exchange or Oracle databases).
- ◆ *Backup window*. For automatic full or incremental backup schedules, this is the time period when NetBackup can start automatic backups of clients covered by this policy. For user schedules, this is the time period when users can start a backup or archive of their own client.
- ◆ *Frequency*. How often automatic and calendar-based backups should occur and which dates should be excluded from the schedule (dates when backups should not occur).
- ◆ *Retention*. How long NetBackup keeps the data that is backed up by this schedule.

- ◆ *Storage unit.* The storage unit for the data that is backed up by this schedule. This setting, if used, overrides the storage unit specified at the backup policy level.
- ◆ *Volume pool.* The volume pool to use when saving data backed up by this schedule. This setting, if used, overrides the volume pool specified at the backup policy level.

The administrator can also manually start a backup schedule for an automatic full or incremental backup. For example, use for an immediate backup request before scheduled maintenance or when performing an upgrade or applying a new driver. For more information see Chapter 3, “Managing Backup Policies” on page 59.

Storage Units

Use **Storage Units** to display storage unit information and manage NetBackup storage units.

A storage unit group is one or more storage devices of a specific type and density that attach to a NetBackup server. The media can be removable (such as tape) or the media can be a file directory on a hard disk. Removable media can be in a robot or a standalone drive.

The devices in a removable-media storage unit (such as a tape drive) must attach to a NetBackup master or media server and be under control of Media Manager. The administrator first sets up Media Manager to use the drives, robots, and media, then defines the storage units. During a backup, NetBackup sends data to the storage unit specified by the backup policy. During a backup, Media Manager picks a device to which the NetBackup client sends data.

When the storage unit is a directory on a hard disk, the administrator specifies the directory during the storage unit setup and NetBackup sends the data to that directory during backups. Media Manager is not involved.

Storage units simplify administration because once defined, the NetBackup policy points to a storage unit rather than to the individual devices it contains. For example, if a storage unit contains two drives and one is busy, NetBackup can use the other drive without administrator intervention. For more information see Chapter 2, “Managing Storage Units” on page 23.

Catalog

Use **Catalog** to create and configure a special type of backup NetBackup requires for its own internal databases—an *offline, cold catalog backup*. (An *online, hot catalog backup* can also be used, and is configured using a wizard, or through the **Policies** utility.)

These databases, called *catalogs*, are, by default, located on the NetBackup master and media server. The catalogs contain information on every client backup. Catalog backups are tracked separately from other backups to ensure recovery in case of a server crash.

Using the NetBackup Administration Console

Catalog is also used to search for a backup image in order to verify the contents of media with what is recorded in the NetBackup catalog, to duplicate a backup image, to promote a backup image from a copy to the primary backup copy, to expire backup images, or to import expired backup images or images from another NetBackup server. For more information see Chapter 4, “NetBackup Catalogs” on page 211.

Host Properties

Use **Host Properties** to customize NetBackup configuration options. In most instances, no changes are necessary. However, **Host Properties** allows the administrator to customize NetBackup to meet specific site preferences and requirements for master servers, media servers, and clients. All configuration options are described in Chapter 7, “Configuring Host Properties” on page 337.

Media and Device Management

The software that manages the removable media and storage devices for NetBackup is called Media Manager. This software is part of NetBackup and is installed on every NetBackup server. The administrator can configure and manage media through **Media and Device Management** in the NetBackup Administration Console.

The *NetBackup Media Manager System Administrator's Guide* contains information on Media Manager.

Access Management

Customers can protect their NetBackup configuration by using **Access Management** to define who may access NetBackup and what functions a user in a user group can perform. Access Management is described in Chapter 1 in the *NetBackup System Administrator's Guide, Volume II*.

NetBackup Administration Console Menus

The following sections describe menus in the NetBackup Administration Console.

File Menu

The **File** menu contains the following options:

- ◆ **Change Server**—Use to display the configuration for another NetBackup master or media server. The name of the current server appears in the status bar. In order to make **Change Server** available, select the Master Server in the tree on the left side of the NetBackup Administration Console.
- ◆ **New Console**—Opens a new NetBackup Administration Console.
- ◆ **New Window from Here**—Opens another window in addition to those that are already open. If you are currently in Activity Monitor and select **New Window From Here**, the new window will open to Activity Monitor.
- ◆ **Adjust Application Time Zone**—Allows you to adjust the time zone for the administration of remote NetBackup hosts. The default time zone for the console is that of the host on which the console is started, not the host specified (if different) in the console login dialog. (See Chapter 3 in the *NetBackup System Administrator's Guide, Volume II*.)
- ◆ **Export**—Save configuration information or data to a file concerning the selected client, server, policy, host properties, storage unit, storage unit group, or device monitor information.
- ◆ **Page Setup**—Displays the Page Setup dialog to enter printer specifications.
- ◆ **Print Preview**—Displays the print preview of the dialog or pane currently in focus.
- ◆ **Print**—Displays the standard Print dialog: specify the range of pages to be printed, the number of copies, the destination printer, and other printer setup options.
- ◆ **Close Window**—If more than one NetBackup window is open, **Close Window** closes only the currently selected window and does not exit NetBackup. If only one NetBackup window is open, you will exit NetBackup.
- ◆ **Exit**—Closes the NetBackup application and all NetBackup Administration Consoles or windows opened through NetBackup.

Edit Menu

The menu items on the **Edit** menu differ depending on which utility is selected in the tree view:

- ◆ **Delete**—Deletes the selected item.
- ◆ **Change**—Displays a dialog where you can specify changes to the selected item.
- ◆ **Copy**—Copies selected items to the clipboard.
- ◆ **Find**—Use the **Find** option (**Control+F**) to highlight rows in the Details pane that meet specific criteria. For information on using any of the **Find** tabs, click **Help**.

Note The **Find** option finds data in hidden columns. The **Find** option will not find data in hidden rows. That is, rows hidden due to filtering.

View Menu

The **View** menu contains the following options:

- ◆ **Show Toolbar**—Use the **Show Toolbar** option to display or hide the standard NetBackup toolbar.
- ◆ **Show Tree**—Use the **Show Tree** option to display or hide the nodes in the left pane of the NetBackup Administration Console.
- ◆ **Alternate Table Row Color**—Enable to have alternate list entries in the right pane of the NetBackup Administration Console appear with a lightly shaded background.
- ◆ **Back**—Use the **Back** option to return to previously selected window panes, moving backwards.
- ◆ **Forward**—Use the **Forward** option to return to previously selected window panes, moving forwards.
- ◆ **Up One Level**—Use the **Up One Level** option to select the next higher node in the tree.
- ◆ **Options**—Customize specific utilities in the Options dialog:

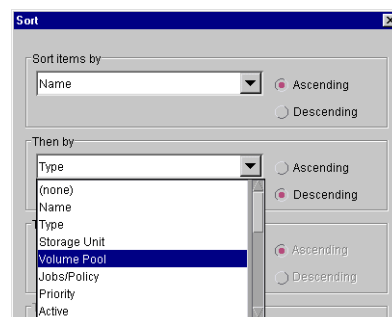
To customize **Activity Monitor**, see “Setting Activity Monitor Options” on page 316.

To customize **Device Monitor** and **Devices**, see the *Media Manager System Administrator's Guide*, Chapter 2.

- ◆ **Refresh**—Use the **Refresh** option to update the Details view with new information retrieved from the master server(s). You can also elect to refresh the display automatically for Activity Monitor by selecting **View > Options > Activity Monitor**, then set the refresh rate > **Automatically refresh display**.
- ◆ **Column Layout**—Use the **Choose Layout** option to choose the columns you wish to display and the order you wish to view them. The **Column Layout** option is available while using the following applications: **Activity Monitor**, **Policies**, **Media** and **Devices**.

- ◆ **Sort**—Use the **Sort** option to sort data using up to three columns of sorting criteria in the Sort dialog. The **Sort** option is available while using the following applications: **Activity Monitor**, **Policies**, **Media** and **Devices**.

First, select a column to sort on by choosing a column header from the **Sort items by** pull-down list. For additional sorting, make selections from the next pull-down list, and so on. Click **OK** to conduct the sort.



To eliminate the sorting selections, open the Sort dialog and click the **Clear All** button. Then click **OK**.

The Detail view shows an arrow in the column header of sorted information. The arrow indicates whether the column is sorted by ascending or descending order.

Name	Type	Storage Unit	Volume Pool	Jobs/Pol
aliu_save	Oracle	chimchim-...	NetBackup	
chimchim_disk_raw	Standard	chimchim-...	NetBackup	
chimchim_sundrops_...	Standard	sundrops-...	NetBackup	
chimchim_sync_test	Standard	chimchim-...	NetBackup	
chimchim_test	Standard	chimchim-...	NetBackup	

Directions of arrows indicate that ascending or descending sorting criteria is in effect.

Name	Type	Storage Unit	Volume Pool	Jobs/Pol
xp_tpc_test	Standard	<any>	NetBackup	
wobegon_xp_stripe_r...	Standard	<any>	NetBackup	
wobegon_xp_stripe_p...	Standard	<any>	NetBackup	
wobegon_xp_stripe_tbu FlashBackup	<any>	FlashBackup	NetBackup	
wobegon_xp_stripe	Standard	<any>	NetBackup	

Clicking on the column header reverses the sort order, but alters the sort by sorting on only one column.

To reverse the sort order but maintain the multi-column sort, open the Sort dialog and use the **Ascending** or **Descending** radio buttons.

- ◆ **Filter**—Use the **Filter** option to display only those rows that meet specific criteria. All other rows are hidden. **Filter** works differently from **Find**: **Find** highlights the row but does *not* hide any rows.

Actions Menu

The menu items on the **Actions** menu differ depending on which utility is selected in the tree view.

Help Menu

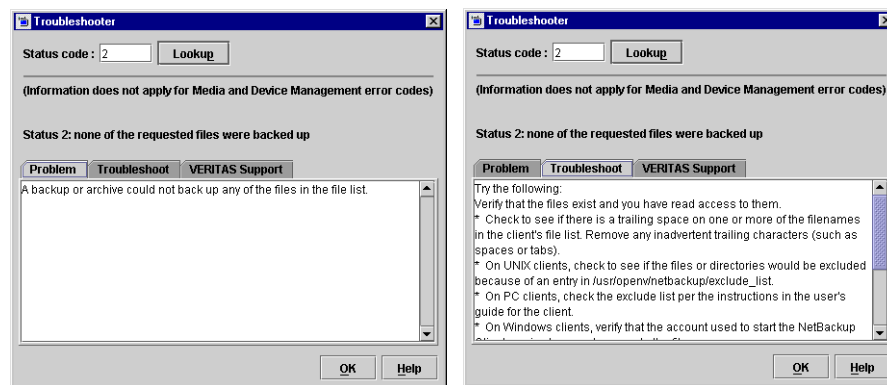
The **Help** menu contains the following options:

- ◆ **Help Topics**—Use the **Help Topics** option to view online help information.
- ◆ **Troubleshooter**—Use the **Troubleshooter** option to open the Troubleshooter dialog.

Using the NetBackup Administration Console

Enter the status code in the **Status Code** field, then click **Lookup**. The dialog contains three tabs:

- ◆ **Problem**—If a valid error code has been entered and looked up, text appears describing why the problem occurred.
- ◆ **Troubleshoot**—Text describes steps to try and correct the problem.
- ◆ **VERITAS Support**—Displays the web site of VERITAS Technical Support. Many codes also display a web address at the VERITAS Technical Support website where you can find additional information pertaining to the status code.



The Troubleshooter is also available from the following locations:

- ◆ From the Activity Monitor, on the **Detailed Status** tab of a job.
- ◆ From the Activity Monitor, on the right-click pop-up menu.*
- ◆ From Reports, by clicking a status code hyperlink.
- ◆ As a button on the Toolbar.


* When using the NetBackup Java applications on a Solaris X86 machine with a two-button mouse, right button pop-up menus can only be popped up using the right button with the **Ctrl** (Control) key as follows:

- a. Press the **Ctrl** key and hold.
- b. Press the second mouse button. A menu appears.
- c. Release the **Ctrl** key.
- d. Select an item from the menu by dragging the cursor to the item and releasing the second mouse button, or releasing the second mouse button, then selecting the menu item with the first mouse button.

- ◆ **License Keys**—Opens a dialog where you can view and modify the license keys for the local computer.

- ◆ **Current NBAC User**

This option is enabled if NetBackup Access Control is configured on your system. The option displays a dialog that lists who you are according to VxSS, the domain you are logged into, the expiration date of your authentication certificate, the type of authentication that you're currently using, and the name of the authentication broker that issued the certificate.



- ◆ **About NetBackup Administration Console**—Displays program information, version number, and copyright information.

Standard and User Toolbars

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

When certain utilities are selected, **Policies** or **Reports**, for example, a second toolbar, called a *user toolbar*, appears. The buttons on the toolbars 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 > Show Toolbar**.

Customizing the Administration Console

The **View** menu contains an **Options** selection that allows you to customize the various NetBackup utilities to suit your preferences (Activity Monitor, Device Monitor, Devices). Click the **Help** button on each tab for more information on the dialog options.

Configuring NetBackup Without Wizards

The easiest way to configure NetBackup is to use the configuration wizards provided.

If configuring NetBackup for the first time, choose the Getting Started Wizard. This wizard steps through the other wizards and completes with a working NetBackup configuration.


If you prefer not to use the available wizards, the following steps explain how to configure NetBackup by using the NetBackup Administration Console. Each step provides references for more information, if needed.

▼ To configure NetBackup without wizards



1. Start the NetBackup Administration Console. (See “NetBackup Administration Interfaces” on page 3.)
2. Complete the addition of storage devices. The Device Configuration Wizard is the preferred method to do this. To perform configuration without using the wizard, see the *Media Manager System Administrator's Guide*.
3. Add the media to be used. In the NetBackup Administration Console, this is performed under **Media and Device Management**. For instructions, see the *Media Manager System Administrator's Guide*.
4. Ensure that all daemons are active on the NetBackup master server (and Enterprise Media Manager server, if the EMM server is not the master server). See “Daemons Tab” on page 325 for information on starting daemons.
5. Configure the storage units. A storage unit is one or more storage devices of a specific type (tape or disk) that attach to a NetBackup server.
 - ◆ For an overview of NetBackup storage units, see “Managing Storage Units” on page 23.
 - ◆ For instructions on adding storage units, see “Creating a New Storage Unit” on page 27.
6. Configure a NetBackup catalog backup. The NetBackup catalogs contain configuration information as well as critical information on client backups. See “NetBackup Catalogs” on page 211.
7. Configure the backup policies for the clients to be backed up. See “Managing Backup Policies” on page 59.

8. Customize host properties to meet site preferences, if desired. In most instances, the NetBackup defaults provide satisfactory results. See “Changing Host Properties” on page 340.
9. Test the configuration as follows:
 - a. If you created user-directed backup or archive schedules, test them by performing a backup or archive from a client. Check the log on the client to see if the operations were successful. If the backup was successful, restore the same files. This verifies the basic backup and restore capabilities.

For instructions on performing user-directed operations, start the NetBackup user interface on the client and follow the instructions that are given in the user’s guide or the online help.
 - b. To test the automatic backup schedules, either wait until the schedules are due or run the schedules manually from the master server. After the automatic schedules have run, check the Backup Status report to ensure that the backups completed successfully for all clients.



Configuring NetBackup Without Wizards



Managing Storage Units

2

Create and manage storage units using the NetBackup **Storage Units** utility.

This chapter explains how to set up storage units for use by NetBackup and contains the following sections:

- ◆ “Introduction to Storage Units” on page 24
- ◆ “Creating, Changing, and Deleting Storage Units” on page 27
- ◆ “Media Manager Storage Unit Considerations” on page 30
- ◆ “Disk Storage Unit Considerations” on page 36
- ◆ “NDMP Storage Unit Considerations” on page 38
- ◆ “Storage Unit Properties” on page 39
- ◆ “About Disk Staging” on page 47
- ◆ “Storage Unit Groups” on page 55

Introduction to Storage Units

A NetBackup storage unit is a storage device attached to a NetBackup server. In order to send backups to a storage device, the administrator needs to define storage units using the Storage Units utility.

There are three types of storage units:

- ◆ Media Manager storage units

This type encompasses the tape robots, standalone tape drives, and optical disk devices—all of which are under the control of Media Manager. Media Manager controls the allocation and mounting of media (called volumes) in the storage devices. (See “Media Manager Storage Unit Considerations” on page 30.)

- ◆ Disk storage units

A disk type storage unit consists of a directory on a disk that stores data. NetBackup permits an unlimited number of disk storage units. (See “Disk Storage Unit Considerations” on page 36.)

There are three types of disk storage units:

- ◆ Basic Disk

- ◆ NearStore, used for Network Attached Storage (NAS). *NearStore* appears as a selection only when the NetBackup Disk Optimization Option is licensed. NearStore storage units cannot be used as part of a storage unit group.

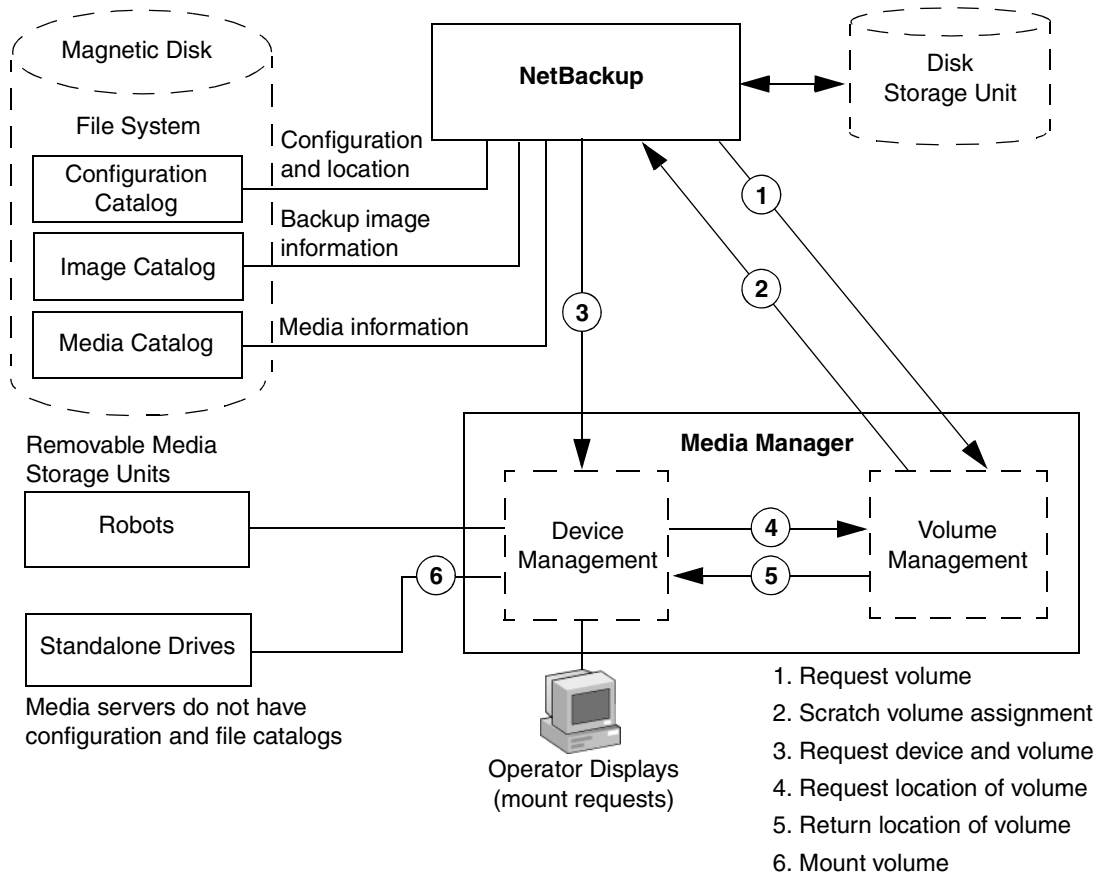
- ◆ SnapVault, used for Network Attached Storage (NAS) and available only with the NetBackup Advanced Client license. SnapVault storage units cannot be used for disk staging or as part of a storage unit group.

Any disk storage unit (except SnapVault) can be used for disk staging. In *disk staging*, the storage unit provides the first storage location in a two-stage process. In this process, client data is backed up to a disk staging storage unit, then, in the second stage, the data is relocated to another storage unit.

- ◆ NDMP storage units

NDMP storage units are controlled by Media Manager but attach to NDMP hosts and require that the NetBackup for NDMP option be installed. (See “NDMP Storage Unit Considerations” on page 38.)

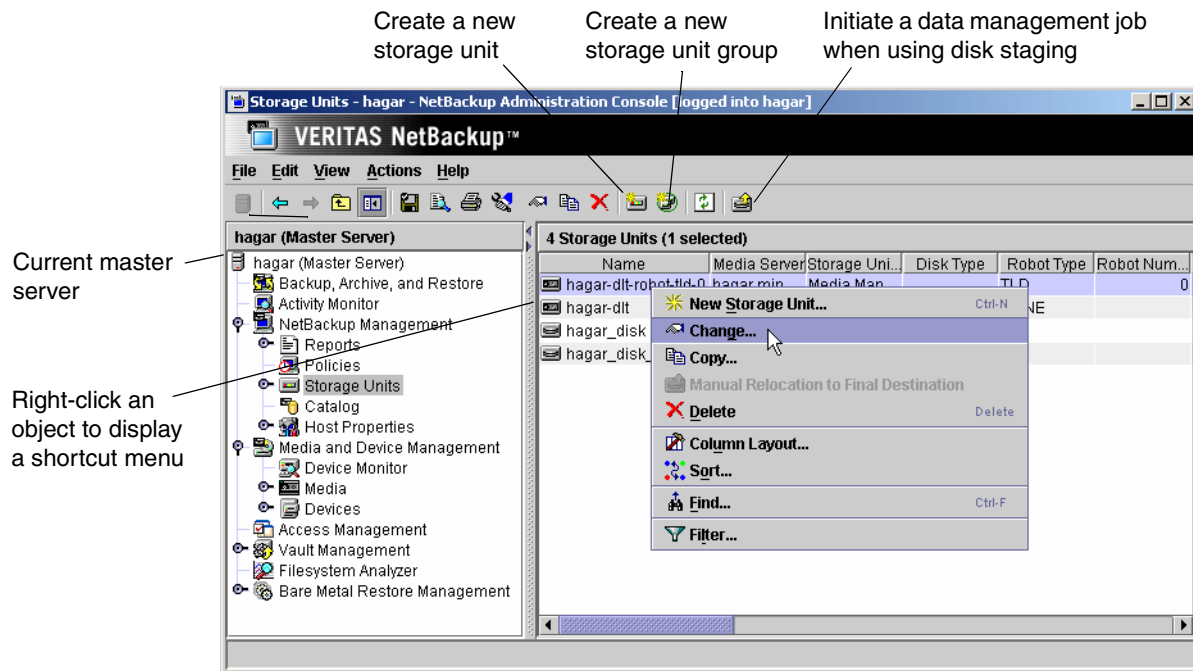
The following figure shows the various components involved and steps in managing the storage of client data.



Introduction to Storage Units

Viewing Storage Units and Storage Unit Groups

In the NetBackup Administration Console, select **NetBackup Management > Storage Units** to display all the storage units for the selected server. All storage units for the selected server display in the Details pane, whether or not the unit is in a storage unit group.



Expand **Storage Units > Storage Unit Groups** to display all the storage unit groups created for the selected server.

Select a storage unit group in the left pane to display all the storage units in the group.

To display storage units and storage unit groups for another NetBackup master server, see “Administering a Remote Master Server” on page 474.

Creating, Changing, and Deleting Storage Units

The following sections contain information on creating and maintaining storage units.

Creating a New Storage Unit

There are several methods to create a new storage unit:

- ◆ Use the Device Configuration Wizard by clicking on **Configure Storage Devices**
- ◆ Create a storage unit using the **Actions** menu
- ◆ Copy an existing storage unit

The easiest way to configure storage units for the first time is to use the Device Configuration Wizard. This wizard guides you through the entire process, simplifying it by automatically choosing settings that work well for most configurations.

▼ To use the Device Configuration Wizard

1. In the NetBackup Administration Console tree, select the **Master Server** or **Media and Device Management**.
2. From the list of wizards in the Details pane, click **Configure Storage Devices** and follow the wizard instructions.

For help while running the wizard, click the **Help** button in the wizard screen.

Note The wizard adds only one disk storage unit if no devices are found.

▼ To create a storage unit from the Actions menu

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Click **Actions > New > Storage Unit**. The New Storage Unit dialog appears.
3. Complete the fields on the New Storage Unit dialog. The options are described in "Storage Unit Properties" on page 39.
4. Click **OK** to add the storage unit to the configuration.

Creating, Changing, and Deleting Storage Units

▼ To create a storage unit by copying an existing storage unit

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Select a storage unit in the Details pane.
3. Click **Actions > Copy Storage Unit**. The Copy Storage Unit dialog appears.
4. Complete the fields in the Copy Storage Unit dialog. The options are described in “Storage Unit Properties” on page 39.

Changing Storage Unit Properties

Best practice guideline suggest that changes to storage units are made only during periods when no backup activity is expected for the policies that will be using the affected storage units. This allows for adjustment time before backups begin and ensures an orderly transition from one configuration to another.

▼ To change storage unit properties

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Double-click the storage unit you wish to change from those listed in the Details pane.
3. Complete the fields on the Change Storage Unit dialog.
The options are described in “Storage Unit Properties” on page 39.

Deleting Storage Units

Deleting a storage unit from the NetBackup configuration does not prevent restoring files that were written to that storage unit.

▼ To delete storage units

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Select the storage unit you wish to delete from those listed in the Details pane. Hold down the Control or Shift key to select multiple storage units.

3. Select **Edit > Delete**. A confirmation dialog appears.
4. Click **OK**.
5. Modify any policy that uses a deleted storage unit to use another storage unit.

NetBackup Naming Conventions

The following set of characters can be used in user-defined names, such as storage units and policies. These characters must be used even when specifying these items in foreign languages.

Do not use a minus as the first character. Spaces are only allowed in a comment for a drive.

- ◆ Alphabetic (A-Z a-z) (names are case-sensitive)
- ◆ Numeric (0-9)
- ◆ Period (.)
- ◆ Plus (+)
- ◆ Minus (-)
- ◆ Underscore (_)

Media Manager Storage Unit Considerations

NetBackup maintains records about the files in the backups and the media where the records are stored. Media Manager manages removable media (for example, tape) and tracks the location of both online and offline volumes.

When sending a backup (or other job) to a Media Manager storage unit, NetBackup requests resources from the Enterprise Media Manager (EMM), then asks Media Manager to mount the volume in a drive.

If a standalone drive does not contain media or if the required volume is not available to a robot, Media Manager displays a mount request. An operator can then find the volume, mount it manually, and assign it to the drive.

Note Media Manager is managed separately and can be used by other applications, such as Storage Migrator. (A tape assigned to Storage Migrator has a different format and would not be usable for NetBackup backup data.)

Notes on Media Allocation

When a volume is allocated to NetBackup, other applications cannot use the volume until the job data on the volume is no longer needed. EMM does not allocate media to an application (including NetBackup), if the media is currently in use.

Media can be assigned to an application by utilizing one of the attributes associated with media, called a *volume pool*. (For more information on volume pools, see Appendix A, "Volume Pools," in the *Media Manager System Administrator's Guide*.) Once the media is associated with a volume pool, and thereby to the application associated with that volume pool, the media won't be selectable for scratch use by another application, as long as all applications are interfacing through NetBackup Media Manager.

NetBackup configurations may involve use of one volume pool, such as *NetBackup*, or configurations may include several volume pools, in order to differentiate media use between differing policies. It is expected that applications are configured to access media

only through NetBackup and Media Manager interfaces, or include other protections, such as verifying ownership of the media by reading the volume header when the media is mounted.

When all NetBackup data on a volume has expired, the volume becomes available to other applications, depending on whether the volume originated from a scratch volume pool. Volumes originally pulled from a scratch volume pool are returned to the scratch pool, and volumes assigned from a specific volume pool are left in that volume pool when unassigned. It is expected that other applications will not explicitly request media from a volume pool that is not associated with the application, based on the configuration set up by the administrator.

The request to Media Manager specifies all the information required to satisfy the mount request, including media ID, drive name, and drive path. This is done for both robotic and standalone media.

The following rules apply when adding Media Manager storage units:

1. *If using NetBackup Enterprise Server:* Add the storage unit to the master server, specifying the media server where the drives attach.

If using NetBackup Server: Add the storage unit to the server where the drives attach. The robotic control must also attach to that server.

2. The number of storage units that you must create for a robot depends on the robot's drive configuration as follows:
 - ◆ Drives with the same density on the same media server must be in the same storage unit. For example, if a robot has 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.
 - ◆ Drives with different densities must be in separate storage units. For example, an STK 9710 library configured in Media Manager as a Tape Library DLT (TLD) can have both half-inch cartridge and DLT drives. Here, you must define a separate storage unit for each density.

Applies only to NetBackup Enterprise Server:

- ◆ Drives on different media servers must be in separate storage units.

Applies only to NetBackup Enterprise Server:

If a robot's drives and robotic control attach to different NetBackup servers, specify the server where the drives attach as the media server. Always specify the same robot number for the drives as is used for the robotic control.

3. Standalone drives with the same density must be in the same storage unit.

Media Manager Storage Unit Considerations

For example, if a server has two 1/4-inch qscsi drives, add a storage unit with **Maximum Concurrent Write Drives** set to 2. Media and device selection logic, part of the Enterprise Media Management (nbemm) daemon, chooses the drive to use when NetBackup sends a backup to this storage unit.

4. Standalone drives with different densities must be in different storage units.
5. A robot and a standalone drive cannot be in the same storage unit.

Before Adding a Media Manager Storage Unit

Before adding a Media Manager storage unit, set up Media Manager to recognize the devices that will be in the storage units. (For device configuration information, see the *Media Manager System Administrator's Guide*.)

As you set up the devices, record the following information from the Media Manager configuration:


Type of Tape Device	Record the Following Information
Robots	<ul style="list-style-type: none"> ♦ The names of the NetBackup servers where the drives attach and the number of drives that attach to each server (<i>Applies only to NetBackup Enterprise Server</i>) ♦ Robot type ♦ Robot number in Media Manager ♦ Media density for the drives in each robot
Standalone tape drives	<ul style="list-style-type: none"> ♦ Media density of each drive ♦ How many drives of each media density are on each NetBackup server

For step-by-step instructions on how to specify this information to NetBackup, see "Creating a New Storage Unit" on page 27.

The following examples show the type of information required by NetBackup for various Media Manager storage unit configurations.

Example 1

The following figure shows a master server containing one drive in a robot and a 1/4 inch SCSI tape drive that is a standalone.



Note TL8 - Tape Library 8mm is the NetBackup name for a device type, not a vendor model number. You must use the NetBackup name when configuring a storage unit. (See "Robot Type" on page 45.)

Each of these devices can be a storage unit. The NetBackup settings required to define these storage units are as follows:

- ◆ 8mm tape drive in the robot

Storage Unit Configuration Setting	Value
Media Server	<i>eel</i>
Robot Type	TL8 - Tape Library 8mm
Robot Number	1
Maximum Concurrent Write Drives	1
Density	8mm - 8mm cartridge

For robots, you must specify the type and number of the robot in which the drives reside.

- ◆ SCSI 1/4 inch tape drive

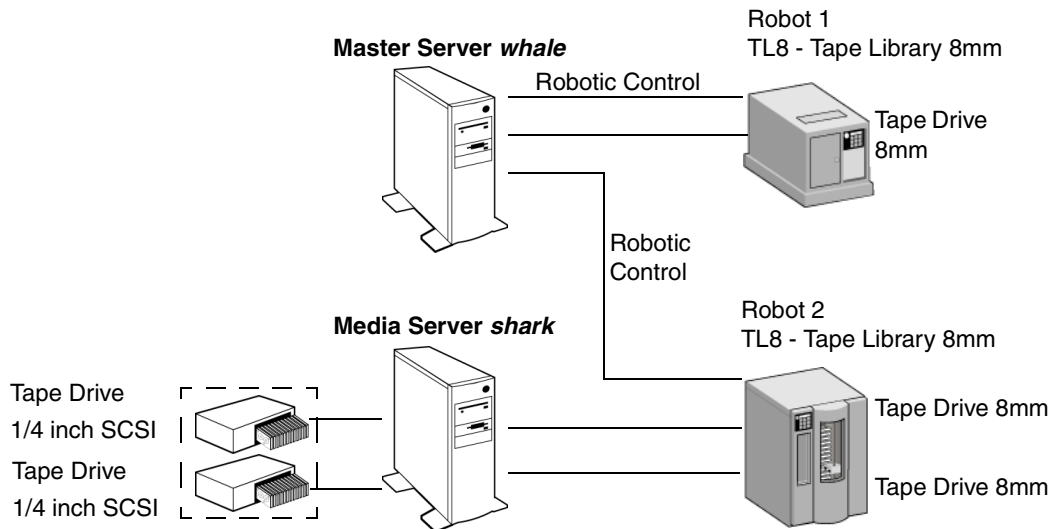
Storage Unit Configuration Setting	Value
Media Server	<i>eel</i>
Robot Type	None
Robot Number	None
Maximum Concurrent Write Drives	1
Density	qscsi - 1/4 inch cartridge

Media Manager Storage Unit Considerations

Example 2

The following example applies only to NetBackup Enterprise Server:

The following figure shows master server *whale*, with a drive in a robot, and media server *shark*, with two drives in a robot and two standalone 1/4 inch SCSI tape drives.



Information Required for Storage Unit on whale

Both the drive and the robotic control for the TL8 - Tape Library 8mm robot attach directly to master server *whale*. The following NetBackup settings are required for this drive to be recognized as a storage unit:

Storage Unit Configuration Setting	Value
Media Server	whale
Robot Type	TL8 - Tape Library 8mm
Robot Number	1
Maximum Concurrent Write Drives	1
Density	8mm - 8mm Cartridge

Master server *whale* also controls the robotics for the TL8 - Tape Library 8mm robot. However, the drives in this robot attach to media server *shark* and therefore the storage unit that contains them must specify *shark* as the media server.

Information Required for Storage Units for Media Server *shark*

For media server *shark*, the two drives in the TL8 - Tape Library 8mm robot can form one storage unit and the two standalone drives can form another storage unit. The following are the NetBackup settings required for these robotic and standalone drives to be recognized as storage units:

◆ 8mm tape drives in robot 2

Storage Unit Configuration Setting	Value
Media Server	<i>shark</i>
Robot Type	TL8 - Tape Library 8mm
Robot Number	2
Maximum Concurrent Write Drives	2
Density	8mm - 8mm Cartridge

The robotic control for the TL8 - Tape Library 8mm is on master server *whale*. However, *shark* must still be the media server for the storage unit because that is where the drives attach. Having the robotic control on one server and drives on another is a valid configuration for this type of robot.

◆ SCSI 1/4 inch tape drives

Storage Unit Configuration Setting	Value
Media Server	<i>shark</i>
Robot Type	None
Robot Number	None
Maximum Concurrent Write Drives	2
Density	qscsi - 1/4 Inch Cartridge

The two standalone 1/4 inch tape drives are of the same density and therefore must be in the same storage unit. If they were of different densities, they would have to each be a separate storage unit.

Disk Storage Unit Considerations

A disk type storage unit is a directory or volume on disk media. Disk media can be one of the following disk types:

- ◆ Basic disk
- ◆ NearStore, used for Network Attached Storage (NAS). *NearStore* appears as a selection only when the NetBackup Disk Optimization Option is licensed.
- ◆ SnapVault, used for Network Attached Storage (NAS) and available only with the NetBackup Advanced Client license

Note While NetBackup allows an unlimited number of disk storage units, do not include the same volume or file system in multiple storage units. This includes disk staging storage units as well.

It is important to select the correct disk type

Before configuring a disk storage unit, configure the disk as explained in the operating system documentation.

Note If a media server of a previous release has disk staging storage units configured, upon upgrade, the storage units are automatically converted to Basic Disk storage units with **This storage unit is a temporary staging area** checked.

Interval Between Capacity Updates

The NetBackup General Server host property **Check the Capacity of Disk Storage Units** property determines how often NetBackup checks disk storage units for available capacity. (See “Check the Capacity of Disk Storage Units” on page 411.)

Maintaining Available Space on Disk Storage Units

It is important that disk storage units be managed to prevent becoming entirely full, causing backups to fail. The following items describe ways in which space could be created for more images on a disk storage unit:

- ◆ Delete files on the volume that are no longer needed.
- ◆ Increase the relocation windows or add resources so all images are relocated in a timely fashion. Once relocation has completed, any image with a .ds extension is automatically deleted, freeing up space. No manual intervention is required.
- ◆ Add new disk space.
- ◆ Set the **High Water Mark** to a value that best works with the size of backup images in your environment.

NDMP Storage Unit Considerations

NDMP Storage Unit Considerations

NDMP storage units are controlled by Media Manager but attach to NDMP hosts. The NDMP hosts must have NetBackup for NDMP option installed.

Create NDMP storage units for drives directly attached to NAS filers. Any drive attached to a NetBackup media server is considered a Media Manager storage unit, even if used for NDMP backups.

Note If a media server of a previous release has remote NDMP storage units configured, upon upgrade of the media server, the storage units are automatically converted to Media Manager storage units.

See the *NetBackup for NDMP System Administrator's Guide* for more information.

New Storage Unit

Storage unit name:

Storage unit type: **NDMP** ☐ On demand only

Disk type: **Basic Disk**

Properties and Server Selection

Storage Device: **TLD(1) - DLT**

Robot type: **TLD - Tape Library DLT**
 Density: **dlt - DLT Cartridge**
 Robot number: **1**

NDMP host: **Bluearc3**

Media server: **Any Available**

Maximum concurrent write drives: **1**

☐ Enable multiplexing ☒ Reduce fragment size to: **1048576** Megabytes

Maximum streams per drive: **1**

OK Cancel Help

Storage Unit Properties

The following sections list and describe storage unit properties. Some properties do not appear for certain storage units types, so the properties are listed alphabetically.

The property values are reflected in the columns of the **Storage Units** details pane.

Absolute Pathname to Directory/Volume

Absolute Pathname to Directory or **Absolute Pathname to Directory** specifies the absolute pathname to a file system or a volume available for disk backups. Enter the pathname directly in the field, then click **Add**. Use any location on the disk, providing there is sufficient space available.

Caution Configuring more than one staging disk storage unit on the same volume or file system can cause problems. Not only will the storage units compete for space, but differing **Low Water Marks** can cause undesired behaviors.

In addition to the platform-specific file path separators (/ and \) and colon (:) within a drive specification on Windows, see “NetBackup Naming Conventions” on page 29.

Directory Can Exist on the Root File System or System Disk

Check this box if it is acceptable for the directory being created in the **Absolute Pathname to Directory** field to exist on the root file system (UNIX) or on a system drive (Windows).

Enabling this setting allows backups to be written to the root file system, possibly causing the the root file system to fill up. Enabling this setting allows the directory to be created automatically.

The job will fail under the following conditions:


- ◆ If the checkbox is clear, and if the directory already exists on the root file system.
- ◆ If the checkbox is clear, and the requested directory is to be created on the root file system.

Properties Button

Click **Properties** to display information about the capacity and remaining space on the volume:

◆ Capacity

The **Capacity** value reflects the total amount of space that the disk storage unit contains, both used and unused.



Storage Unit Properties

◆ Available Storage

The **Available Storage** value reflects available space remaining for storage on the disk storage unit.

Density

The **Density** is determined by the **Storage Device** selection and indicates the media density of the storage unit.

Disk Type

A **Disk Type** is selected for a disk storage unit. A disk storage unit may be a Basic Disk, NearStore (if the NetBackup Disk Optimization Option is licensed), or SnapVault. Any disk type (excluding SnapVault) may serve as a temporary staging storage unit.

For Media Manager storage units, the **Disk Type** does not apply.

For information on SnapVault storage units, see the *Advanced Client System Administration Guide*.

Enable Block Sharing

The **Enable Block Sharing** setting is available on NearStore disk storage units. In order to select NearStore as a storage unit, the NetBackup Disk Optimization Option must be licensed.

The **Enable Block Sharing** setting allows data blocks that have not changed from one backup to the next to be shared. Sharing data blocks can significantly save disk space in the storage unit.

Enable Multiplexing

The **Enable Multiplexing** setting allows multiple backups to multiplex onto a single drive in a storage unit.

The **Maximum Streams per Drive** setting determines the maximum number of concurrent, multiple client backups that NetBackup can multiplex onto a single drive, from 2 to 32. See Chapter 3 in the *NetBackup System Administrator's Guide, Volume II* for more information.

High Water Mark

The **High Water Mark** setting is a threshold that, when reached, signals to NetBackup that the disk storage unit should be considered full. Default: 98%.

NetBackup does not assign new jobs to a storage unit that is considered full. Once the capacity of the storage unit is below the **High Water Mark**, jobs can once again be assigned to the storage unit.

If NetBackup cannot find a storage unit to assign to the job, the job fails.

As the capacity of the disk storage unit grows closer to the **High Water Mark**, NetBackup begins to reduce the number of jobs allowed to write to the storage unit. (The number of jobs is determined by the **Maximum Concurrent Jobs** setting. See “Maximum Concurrent Jobs” on page 42.) Reducing the number of jobs allowed to write to the storage unit helps prevent the situation in which multiple jobs attempt to write to the storage unit at one time, filling it to capacity. Once the storage unit is full, none of the jobs can complete and all the jobs fail due to a disk full condition.

Note The **High Water Mark** is not available on storage units used for disk staging.

Low Water Mark

The **Low Water Mark** pertains only to disk storage units acting as temporary staging storage units. (See “Temporary Staging Area” on page 46.) Default: 80%.

Once the **High Water Mark** is reached, space is created on the disk storage unit until the **Low Water Mark** is met. To do this, NetBackup may copy images to other storage units or expire images (oldest first) to free space. The **Low Water Mark** setting cannot be greater than the **High Water Mark** setting.

Note If a media server of a previous release has disk staging units configured, upon upgrade, the disk storage units are set with the **Low Water Mark** at 100%. To make the best use of the upgraded storage unit, consider adjusting the level.

Maximum Concurrent Write Drives

Maximum Concurrent Write Drives specifies the number of tape drives that NetBackup can use at one time for backups and other job types, in this storage unit. The number of tape drives available is limited to the maximum number of tape drives in the storage device. If a job contains multiple copies, each copy applies toward the **Maximum Concurrent Write Drives** count.

Select the desired number:

Storage Unit Properties

- ◆ For a 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 this storage unit.
- ◆ For a 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.

For example, assume you have two standalone drives of the same density and you specify **1**. In this instance, both tape drives are available to NetBackup but only one drive can be used for backups. This leaves the other tape drive available for restores and other non-backup operations (importing, verifying, and duplicating backups).

Note Specifying 0 effectively disables the storage unit.

Maximum Concurrent Jobs

For disk storage units, **Maximum Concurrent Jobs** specifies the maximum number of backups and other job types, that NetBackup can send to the disk at one time. For example, if three backup jobs are ready to be sent to this storage unit and **Maximum Concurrent Jobs** is set to two, the first two jobs start and the third one waits. If a job contains multiple copies, each copy applies toward the **Maximum Concurrent Jobs** count.

(Default: 1 job. The job count can range from 0 to 200.)

Caution Specifying a **Maximum Concurrent Job** setting of 0 disables the storage unit.

Maximum Concurrent Jobs can be used to balance the load between disk storage units. A higher number of concurrent jobs means that the disk could be busier than if the number is lower.

Maximum Concurrent Jobs corresponds to the **Maximum Concurrent Write Drives** setting for a Media Manager storage unit.

The number to enter depends on the available disk space and the server's ability to comfortably run multiple backup processes. (See "Limit Jobs Per Policy" on page 77.)

Note Increase the **Maximum Concurrent Jobs** setting if the storage unit is being used for online, hot catalog backups as well as non-catalog backups. This ensures that the catalog backup can proceed while regular backup activity is occurring.

Media Server

The following setting applies only to NetBackup Enterprise Server:

The **Media Server** setting specifies the name of the NetBackup media server where the drives in the storage unit attach, or the name of the server that is controlling the disk storage unit.

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.

Concerning Disk Storage

When configuring a disk storage unit, select a single media server.

Concerning NDMP Storage

For NDMP storage, the **Media Server** setting specifies the name of the media server that is to back up the NDMP host. Select the media server from the drop-down menu. Only those media servers that can talk to the specified NDMP storage device are displayed.

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.

NDMP Host

NDMP Host specifies the NDMP tape server that will be used to write data to tape. Select the hostname from the drop-down menu.

On Demand Only

On Demand Only specifies whether the storage unit is available exclusively on demand, that is, only when a policy or schedule is explicitly configured to use this storage unit. Clear the **On Demand Only** check box to make the storage unit available to any policy or schedule.

For SnapVault and NearStore storage units, **On Demand Only** is selected by default and cannot be changed.

Caution If you select **On Demand Only** for all storage units, be sure to designate a specific storage unit for each policy or schedule. Otherwise, NetBackup will be unable to find a storage unit to use.

Reduce Fragment Size

The **Reduce Fragment Size** setting specifies (in megabytes) the largest fragment size that NetBackup can create when storing backups.

Storage Unit Properties

For Media Manager storage units:

The default maximum fragment size for a Media Manager storage unit is 1000 terabyte. To specify a maximum fragment size other than the default, place a check in the **Reduce Fragment Size** check box, then enter a value of 50 to 1,048,575 megabytes.

Fragmenting multiplexed tape backups can speed up restores by allowing NetBackup to skip to the specific fragment before searching for a file. Otherwise, NetBackup starts at the beginning of the multiplexed backup and reads tar headers until finding the desired file.

For disk storage units:

The default maximum fragment size for a disk storage unit is 524,287 megabytes. To specify a maximum fragment size other than the default, enter a value that ranges from 20 to 524,287 megabytes.

Fragmenting disk backups is normally used to ensure that the backup does not exceed the maximum size allowed by the file system. It is intended primarily for storing large backup images on a disk type storage unit. Another benefit of fragmenting backups on disk is increased performance when restoring from images that were migrated by Storage Migrator. For example, if a 500 megabyte backup is stored in 100 megabyte fragments, you can restore a file quicker because Storage Migrator has to retrieve only the specific fragment with the file rather than the entire 500 megabytes.

If an error occurs in a backup, the entire backup is discarded and the backup restarts from the beginning, not from the fragment where the error occurred. (Except for backups where checkpoint restart is enabled. In that case, fragments prior to and including the last checkpoint are retained; the fragments following the last checkpoint are discarded.)

Note If a media server of a previous release has disk storage units configured with a different maximum fragment size, upon upgrade, the storage units are not automatically increased to the new default of 524,287 megabytes. To make the best use of the storage unit, consider increasing the fragment size on upgraded storage units.

Robot Number

The **Robot Number** is determined by the **Storage Device** selection. It is the same robot number used in the Media Manager configuration. For more information on robot numbers, see the *Media Manager System Administrator's Guide*.

Robot Type

The **Robot Type** is determined by the **Storage Device** selection and indicates the type of robot (if any) that the storage unit contains.

For the specific vendor types and models that correspond to each robot type, see the Supported Peripherals section of the NetBackup *Release Notes*.

Staging Relocation Schedule

Click the **Staging Schedule** button to set up the second stage of disk staging. During this stage the backup image is duplicated from the storage unit acting as the disk staging storage unit, to the final destination storage unit. (For more information, see “About Disk Staging” on page 47.)

Storage Device

The **Storage Device** list is a listing of all possible storage devices available. Storage units can be created for the listed devices only.

Storage Unit Name

For the **Storage Unit Name** setting, type a unique name for the new storage unit that describes the type of storage you are defining. This is the name to use when specifying a storage unit for policies and schedules. (See “NetBackup Naming Conventions” on page 29.)

The storage unit name cannot be changed after creation. If this is a Change Storage Unit operation, the **Storage Unit Name** will be inaccessible.

Storage Unit Type

The **Storage Unit Type** setting specifies the type of storage that this storage unit will use:

- ◆ **Disk:** A directory on a hard drive
- ◆ **Media Manager:** A robot or standalone tape drive
- ◆ **NDMP:** For use with NetBackup for NDMP—an optional application that enables NetBackup to use the Network Data Management Protocol (NDMP) to initialize and control backups and restores of Network Attached Storage (NAS) systems that support NDMP

Storage Unit Properties

Temporary Staging Area

If this storage unit is to be used as a temporary staging area for disk staging, click the temporary staging checkbox. If this is a storage unit to be used for staging, set the staging schedule. (See “Staging Relocation Schedule” on page 45.)

In the **Storage Units** details pane, this is indicated in the *Staging* column as yes or no.

Transfer Throttle

The **Transfer Throttle** setting appears for SnapVault storage units only.

The **Transfer Throttle** setting makes it possible to limit the amount of network bandwidth used for the SnapVault transfer, in case bandwidth should be reserved for other applications. Zero (default) means *no network bandwidth limit* for the SnapVault transfer: SnapVault will use all available bandwidth. (Range: 0 to 9999999.)

A value greater than 0 indicates a transfer speed for SnapVault in kilobytes per second. A value of 1, for instance, would set a transfer speed limit for SnapVault of 1 kilobyte per second, which is a very slow transfer rate.

About Disk Staging

Disk staging provides a method for administrators to create images on disk initially, then later copy the images to another media type (as determined in the disk staging schedule). The media type for the final destination is typically tape, but could be disk.

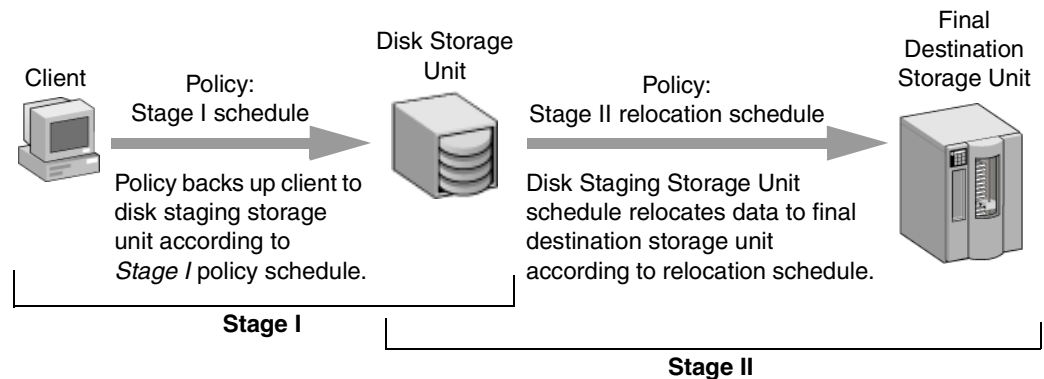
This two-stage process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term, while preserving the advantages of tape-based backups for long term.

Disk staging meets the following objectives:

- ◆ Allows backups when tape drives are scarce.
- ◆ Allows for faster restores from disk.
- ◆ Facilitates streaming to tape without image multiplexing.

Disk staging is conducted in two separate stages:

1. A backup creates an image on the storage unit acting as the disk staging storage unit.
2. A relocation schedule determines when the image from the disk staging storage unit should be relocated to the destination storage unit.



The image continues to exist on both the disk staging storage unit and the destination storage unit. File restores are done from the disk staging storage unit copy, while the destination storage unit copy can be considered the long term copy.

The disk copy continues to exist on the disk staging storage unit until either the copy expires based on the copy's retention period, or until another Stage I process needs space on the disk storage unit.

When a Stage I process detects a full disk staging storage unit, it pauses the backup, finds the oldest image that has been successfully copied to the destination storage unit, and expires this image copy.

About Disk Staging

▼ To create a disk staging storage unit

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Click **Actions > New > Storage Unit**. The New Storage Unit dialog appears.
3. Complete the New Storage Unit dialog:
 - ◆ Name the storage unit. (See “Storage Unit Name” on page 45.)
 - ◆ Select *Disk* as the **Storage Unit Type**.
 - ◆ Select a media server. (See “Media Server” on page 42.)
 - ◆ Enter an absolute pathname to the directory that will be used for storage. (See “Absolute Pathname to Directory/Volume” on page 39.)
 - ◆ Check whether this directory can reside on the root file system or system disk. (See “Directory Can Exist on the Root File System or System Disk” on page 39.)
 - ◆ Enter the maximum concurrent jobs allowed to write to this storage unit at one time. (See “Maximum Concurrent Jobs” on page 42.)
 - ◆ Enter a **Low Water Mark** value. The **High Water Mark** does not apply to disk staging storage units. (See “Low Water Mark” on page 41.)
 - ◆ Check the disk staging checkbox. The checkbox is labeled: **This storage unit is a temporary staging area**. Checking this enables the Staging Schedule button.
 - ◆ Click the Staging Schedule button. The Disk Staging dialog appears.
4. Complete the Disk Staging dialog. The dialog is similar to the scheduling dialog seen when configuring policies. The differences appear on the Attributes tab:
 - ◆ The schedule name defaults to the storage unit name.
 - ◆ Select the priority that relocation jobs started from this schedule will have over other types of jobs. (See “Priority of Relocation Jobs Started from this Schedule” on page 52.)
 - ◆ Select the storage unit where the images will be relocated. (See “Final Destination Storage Unit” on page 53.) (See “Final Destination Storage Unit” on page 53.)
 - ◆ Select the volume pool where the images will be relocated. (See “Final Destination Volume Pool” on page 53.)
 - ◆ Select whether to use an alternate server. (See “Use Alternate Read Server” on page 53.)
5. Click **OK** to accept the disk staging schedule.

6. Click **OK** to add the storage unit.

Disk Staging Storage Unit Size and Capacity Considerations

Leveraging the advantages of disk staging requires that the NetBackup administrator understand the life expectancy of the disk-based image. After the disk-based image is copied to the final destination storage unit, management of the disk-based copy's retention is handed over to the disk staging disk full logic.

Therefore, the size and usage of the file system containing the disk staging storage unit directly impacts the life expectancy of the disk-based image. This is why it is strongly recommended to have a dedicated file system for each disk staging storage unit.

Example: The 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 utilize the disk staging feature. Each night's total incremental backups average from 300 to 500MB. Occasionally a backup contains 700MB. Each following day the disk staging schedule runs and copies the previous night's incrementals to the destination storage unit (tape).

Minimum Disk Staging Storage Unit Size

The minimum disk staging storage unit size represents the minimum size needed for the successful operation of the disk staging logic. The minimum size will not accommodate the desired level of service (as disk images remain on the disk for one week in our example).

The minimum size for the disk staging storage unit must be greater than or equal to the maximum size of backups placed on the storage unit between runs of the disk staging schedule.

In this example, the disk staging schedule runs nightly, and the largest nightly backup is 700MB. VERITAS recommends doubling this value to allow for unanticipated problems running a disk staging schedule. Doubling the value gives the administrator an extra schedule cycle (one day) to correct any problems.

The following formula was used to arrive at the minimum disk staging storage unit size in our example:

Minimum disk staging storage unit size = Max data per cycle * (1 cycle + 1 cycle for safety)

For example: 1.4GB = 700MB * (1+1)

About Disk Staging

Average Disk Staging Storage Unit Size

The average disk staging storage unit size represents a good compromise between the minimum and maximum sizes.

For example, if the average nightly backup is 400MB and the desire is for the images to be kept for one week, the recommended average size is calculated based on the following formula:

$$\begin{aligned} \text{Average Size of disk staging storage unit} &= \\ \text{Average data per cycle} * (\text{number of cycles to keep data} + 1 \text{ cycle for safety}) \\ 2.8\text{GB} &= 400\text{MB} * (6 + 1) \end{aligned}$$

Maximum Disk Staging Storage Unit Size

The maximum disk staging storage unit size is the recommended size needed to accommodate the level of service desired. In this example, the level of service is that disk images remain on disk for one week.

To determine the size, use the following formula:


$$\begin{aligned} \text{Maximum Size} &= \text{Max data per cycle} * (\# \text{ of cycles to keep data} + 1 \text{ cycle for safety}) \\ \text{For example: } 4.9 \text{ GB} &= 700\text{MB} * (6 + 1) \end{aligned}$$

Note When creating a disk staging storage unit, VERITAS strongly recommends dedicating a disk partition/file system to the disk staging storage unit. This allows the disk staging space management logic to operate successfully.

Disk Staging: Stage I

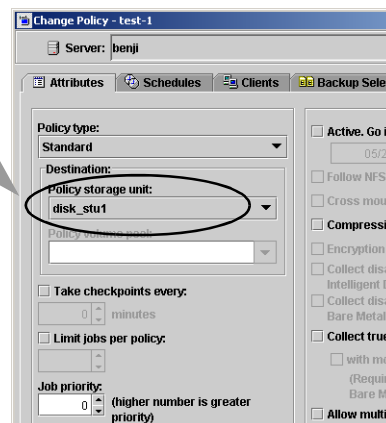
In the first stage of the backup, clients are backed up by a policy that indicates a disk staging storage unit as the **Destination Policy Storage Unit**.

Schedule a policy for Stage I as you would for any other backup.



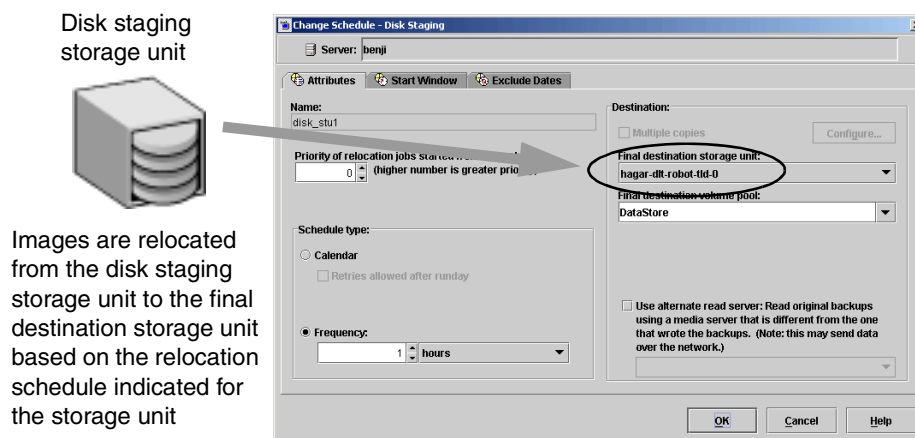
Client

The policy indicates that the client is to back up to a disk storage unit that is acting as a disk staging storage unit.



Disk Staging: Stage II

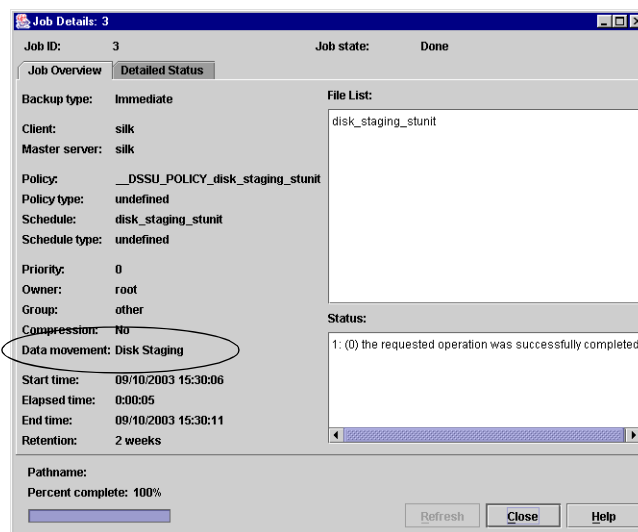
In the second stage of disk staging, images are relocated from the disk staging storage unit to the destination storage unit.



The images are relocated based on the relocation schedule configured during the disk staging storage unit setup, by clicking the **Staging Schedule** button. The button is available only when **Staging Storage Unit** is selected as the storage unit type.

Every time the relocation schedule runs, NetBackup creates a job that acts as a data management job, looking for data that needs to be relocated. The Job Details in the Activity Monitor identify the job as one associated with a disk staging storage unit by listing *Disk Staging* in the job's Data Movement field.

The data management job can also be initiated manually.



About Disk Staging

▼ To manually initiate a disk staging storage unit relocation schedule

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
2. Select a disk staging storage unit in the Details pane.
3. Select **Actions > Manual Relocation to Final Destination** to initiate the schedule.

If the relocation schedule finds data that can be relocated, NetBackup creates a duplication job to relocate the data to the Destination Storage Unit.

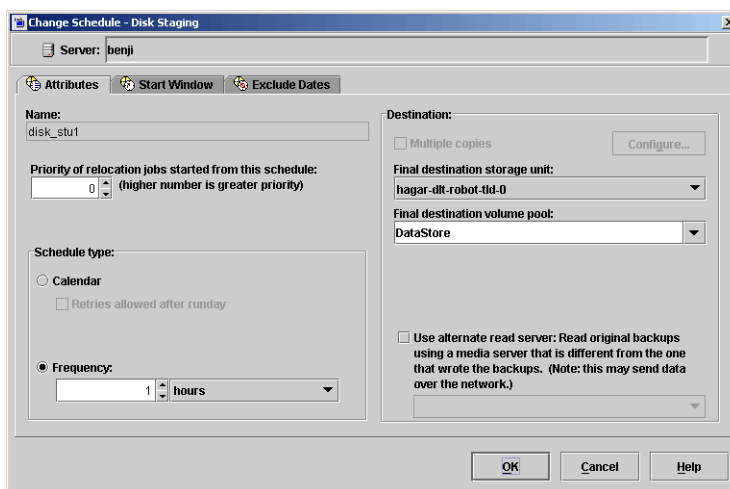
The image then exists both on the disk staging storage unit and the destination storage unit. When the disk staging storage unit becomes full, it is *cleaned* and the oldest images are deleted. (See “Maintaining Available Space on Disk Storage Units” on page 37.)

Staging Schedule Button

Click the **Staging Schedule** button to display the Disk Staging dialog. The Disk Staging dialog is similar to the scheduling dialog seen when configuring policies. The differences appear on the Attributes tab, as described in the following sections:

Name

The schedule name for a disk staging storage unit automatically defaults to the name of the storage unit.



Priority of Relocation Jobs Started from this Schedule

This attribute specifies the priority that NetBackup assigns to relocation jobs for this policy. Range: 0 (default) to 99999 (highest priority).

Final Destination Storage Unit

The **Final Destination Storage Unit** is the name of the storage unit where the images are relocated from the disk staging storage unit.

To relocate images to a Media Manager storage unit, 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, thus determining how many duplication jobs can be launched to handle the relocation. (See “Maximum Concurrent Write Drives” on page 41.)

NetBackup continues to free space until the **Low Water Mark** is reached. (See “High Water Mark” on page 41 and “Low Water Mark” on page 41.)

Final Destination Volume Pool

The **Final Destination Volume Pool** is the name of the volume pool on the final destination storage unit where the images are to be relocated.

If the Final Destination Storage Unit is a Media Manager storage unit (tape), or if *Any Available* is indicated for the Final Destination Storage Unit, the **Final Destination Volume Pool** is selectable. This does not apply if the final destination is a disk storage unit.

Note The schedule created for the disk staging storage unit is not listed under **Schedules** in the NetBackup Administration Console when **Policies** is selected.

Use Alternate Read Server

The Alternate Read Server attribute applies to NetBackup Enterprise Server only.

An alternate read server is a server allowed to read a backup image originally written by a different media server.

The path to the disk or directory must be identical for each media server that is to access the disk.

If the backup image is on tape, the media servers must share the same tape library or the operator must relocate the media.

If the backup image is on a non-SSO robot 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 `bpmedia -oldserver -newserver` or assign a failover media server.

To avoid sending data over the network during duplication, specify an alternate read server that is:

- ◆ Connected to the robot that has the original backups (the source volumes).

About Disk Staging

- ◆ Connected to the robot that contains the destination storage units.

If the destination storage unit is not connected to the alternate read server, data is sent over the network.

Disk Staging Limitations

Disk staging 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 “Checkpoint Restart for Backup Jobs” on page 75.)

Storage Unit Groups

Storage unit groups allow you to identify specific storage devices as a group. A storage unit group name can be specified in a policy, just as individual storage units can be specified. When a storage unit group is used in a policy, only the storage units specified in the group will be candidates for the backup.

The order that storage units are used within a group depends on how the storage unit selection criteria is set. You may elect to use storage units in the order in which they are displayed in the storage unit group dialog; to use each storage unit in turn; or to use the first storage unit in the list that is not full or down and use the remaining storage units for failover purposes only.

The only exception 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 unit groups.

You may have set up a storage unit to be **On Demand Only**. If the storage unit is part of a storage unit group that is needed by a policy, the **On Demand Only** option is satisfied and the device will be used. (See “Policy Storage Unit” on page 71.)

▼ To create a storage unit group

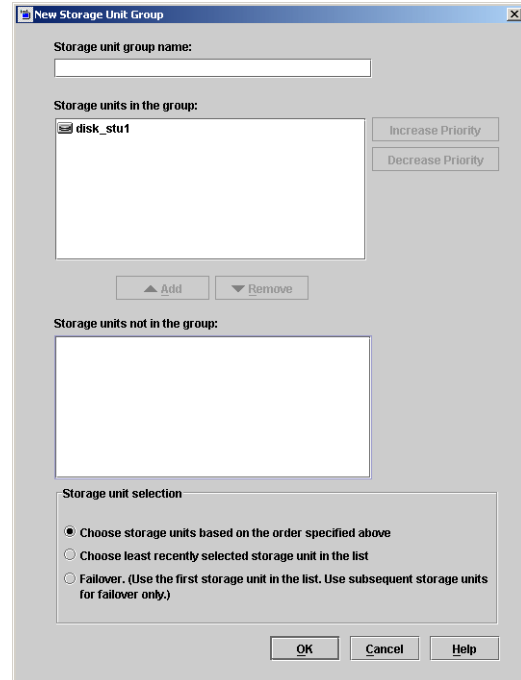
1. In the NetBackup Administration Console, expand **NetBackup Management > Storage Units**.
2. Right-click **Storage Unit Groups** and select **New**. The **New Storage Unit Group** dialog appears.

Storage Unit Groups

3. Enter a storage unit group name for the new storage unit group. (See “NetBackup Naming Conventions” on page 29.)

Note The storage unit group name is case-sensitive.

4. Add to or remove storage units from the group:
 - a. To add storage units to the group, select the storage units from the **Storage units not in the group** list. Click **Add**.
 - b. To remove storage units from the group, select the storage units from the **Storage units in group** list. Click **Remove**.
 - c. Storage units are listed in order of priority: The units at the top of the list having the highest priority in the group. To change the priority of a storage unit, select the storage unit and click **Increase Priority** or **Decrease Priority**.



Caution SnapVault and NearStore storage units cannot be included in storage unit groups.

5. Choose how storage units are to be selected within the group:
 - ◆ Use the storage units in the order displayed in the storage unit group dialog (**Choose storage units based on the order specified above**)(default).
 - ◆ Use the storage unit least recently used (**Choose least recently selected storage within the list**). In effect, the storage units take turns being used.
 - ◆ Use the first storage unit in the list that is not full or down (**Failover**). If the storage unit is only busy, the policy waits to write to it. The other storage units will be used as failovers.

Note VERITAS recommends using the second setting (*least recently selected*) for disk staging storage units within a storage unit group.


6. Click **OK**.

▼ **To change a storage unit group**



1. In the NetBackup Administration Console, expand **NetBackup Management > Storage Units > Storage Unit Groups**.
2. Double-click the storage unit group you wish to change.
3. To add storage units to the group, select the storage units from the **Storage units not in group** list. Click **Add**.
4. To remove storage units from the group, select the storage units from the **Storage units in group** list. Click **Remove**.
5. To change the order of the storage unit list, select the storage unit and click **Increase Priority** or **Decrease Priority**.
6. If desired, change the storage unit selection criteria, which determines the order and how often storage units are selected for use.
7. Click **OK**.

▼ **To delete a storage unit group**

1. In the NetBackup Administration Console, select **NetBackup Management > Storage Units > Storage Unit Groups**.
2. Select the storage unit group you wish to delete from those listed in the Details pane. Hold down the Control or Shift key to select multiple storage units.
3. Select **Edit > Delete**. A confirmation dialog appears.
4. Click **OK**.



Storage Unit Groups



Managing Backup Policies

3

Backup policies define the rules that NetBackup follows when backing up clients. A backup policy can apply to one or more clients. Every client must be covered by at least one backup policy. The best approach to configuring backup policies is to divide clients into groups according to any backup and archiving requirements, then create a policy for each group.

This chapter contains the following sections:

- ◆ “Using the Policies Utility” on page 60
- ◆ “Standard and User Toolbars” and “Configuring Backup Policies” on page 61
- ◆ “Introduction to Backup Policies” on page 63
- ◆ “Changing Policies” on page 65
- ◆ “What Type of Policy: Policy Attributes Tab” on page 69
- ◆ “When Will the Job Run: Schedules Tab” on page 97
- ◆ “Schedule Attributes Tab” on page 98
- ◆ “Start Window Tab” on page 119
- ◆ “Exclude Dates Tab” on page 121
- ◆ “Calendar Schedule Tab” on page 122
- ◆ “Policy Planning Guidelines for Backups” and “Example Policies” on page 142
- ◆ “Which Clients Will Be Backed Up: Clients Tab” on page 149
- ◆ “Which Selections Will Be Backed Up: Backup Selections Tab” on page 154
- ◆ “Rules for Indicating Pathnames in the Backup Selections List” on page 162
- ◆ “Where Will the Catalog Data Be Located: Disaster Recovery Tab” on page 193
- ◆ “Creating a Vault Policy” on page 195
- ◆ “Performing Manual Backups” on page 196
- ◆ “More About Synthetic Backups” on page 198

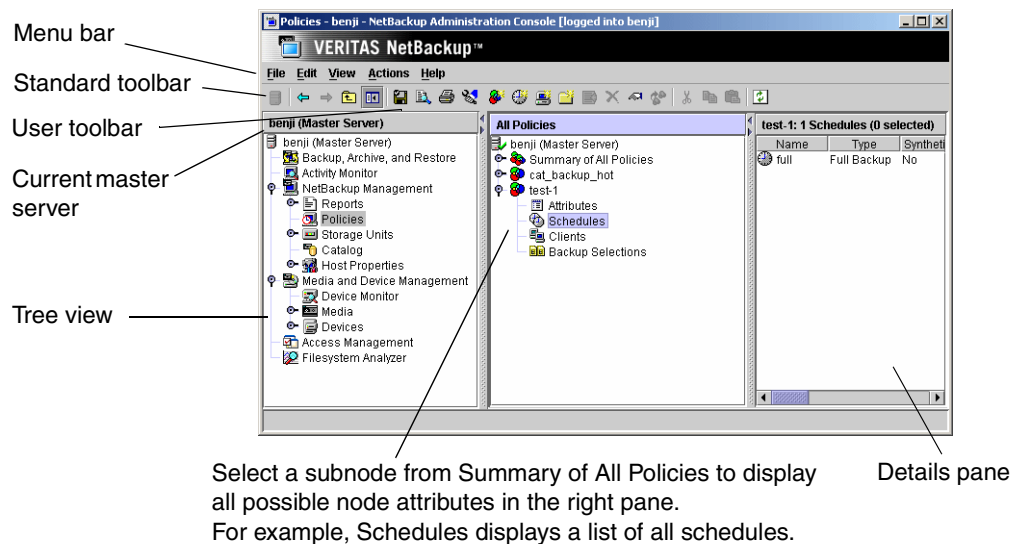
Using the Policies Utility

The **Policies** utility contains tools for configuring and managing policies.

Tree and Detail Views

The center pane labeled **All Policies**, contains a hierarchical view of the policies on the master server that you are currently managing. The Details pane displays a list of all policies with general attribute information for each policy.

Double-click **Summary of All Policies** to expand or collapse the subnodes **Attributes**, **Schedules**, **Clients**, and **Backup Selections**. Select a subnode to display a list of all possible attributes for that node.



Policies Menu Bar

The **Policies** menu bar consists of **File**, **Edit**, **View**, **Actions**, and **Help**. See Chapter 1 for a description of the items found on these menus. The following sections describe additional menu options.

Actions Menu

The **Actions** menu contains the following options when **Policies** is selected.

- ◆ **Activate:** Activates the policy currently selected in the Console tree. A policy must be active for NetBackup to run automatic backups or allow user backups or archives. This setting has no effect on restores.


If the schedule is to be used for a catalog archive, the policy must *not* be active. The **Active** check box must be clear.

For more information on configuring a policy for catalog archiving, see “Creating a Catalog Archiving Policy” on page 249.

- ◆ **Deactivate:** Deactivates the selected policy (see **Activate** above).
- ◆ **Manual Backup:** Displays the Manual Backup dialog when a policy is selected in the Console tree. Select a schedule and client in the Manual Backup dialog then click **OK** to start a manual backup.
- ◆ **Install UNIX Client Software:** Allows installation of client software from the NetBackup Administration Console. (See “Installing Client Software on Trusting UNIX Clients” on page 150.)

Standard and User Toolbars

The user toolbar in the Policies utility contains shortcuts for the following actions:



For information on the standard toolbar, see “Using the NetBackup Administration Console” on page 9.

Configuring Backup Policies

The easiest way to set up a backup policy is to use the Backup Policy Configuration Wizard. This wizard guides you through the setup process, simplifying the process by automatically choosing default values that are good for most configurations.

Note The wizard cannot be used, however, to configure a calendar-based schedule. You can change the schedule to a calendar-based schedule after running the wizard. (See “Calendar Schedule Tab” on page 122.)

Configuring Backup Policies

▼ To create a policy using the Backup Policy Configuration Wizard

1. In the NetBackup Administration Console, select **Master Server** or **NetBackup Management**.
2. From the list of wizards in the Details pane, click **Create a Backup Policy**.
Click **Help** on any wizard screen for assistance while running the wizard.

▼ To create a policy without using the wizard

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. (See “NetBackup Naming Conventions” on page 29.)
Click **OK**.

Note If the schedule is to be used for a catalog archive, the policy must be named *catarc*. For more information on configuring a policy for catalog archiving, see “Creating a Catalog Archiving Policy” on page 249.

Introduction to Backup Policies

Backup policies are configured on four tabs, as described in the following sections.

Attributes Tab Overview

The settings on the **Attributes** tab determine the basic characteristics of all the backups that NetBackup performs according to the selected policy.

The settings include defining:

- ◆ whether the selected policy is currently active or what date and time the policy will go into effect.
- ◆ the type of backup policy, which primarily defines the type of clients the policy can back up.
- ◆ the priority that NetBackup gives to the backups for the selected policy relative to other policies.
- ◆ the storage unit that NetBackup uses by default when backing up clients covered by the selected policy.
- ◆ Advanced Client attributes, if Advanced Client is installed.

For more information on the **Attributes** tab, see “What Type of Policy: Policy Attributes Tab” on page 69.

Schedules Tab Overview

The schedules defined on the **Schedules** tab determine when backups occur. Each schedule also includes various criteria, such as how long to retain the backups.

There are two basic categories of schedules:

- ◆ *Automatic schedules* back up the items listed in the backup selection list on all clients in the policy according to the timetables set up in the schedules.
- ◆ *User schedules* specify the times when users can start user backups and archives from the clients. A user archive is a special type of backup that deletes the files from the user disk if the backup is successful. An archive is useful for freeing disk space while still keeping a copy for future use.

For more information on the **Schedules** tab, see “When Will the Job Run: Schedules Tab” on page 97.

Clients Tab Overview

The client list names the computers that will be backed up according to the selected policy. A client must be on the list of at least one backup policy in order to be backed up. Having a client in more than one backup policy is useful, for example, to back up different sets of files on the client according to different policy rules.

For more information on the **Clients** tab, see “Which Clients Will Be Backed Up: Clients Tab” on page 149.

Backup Selections Tab Overview

The backup selections list names the files, directories, directives, scripts, and templates that NetBackup includes in automatic backups of clients covered by the selected policy.

NetBackup uses the same selection list for all clients backed up according to a policy. All the files and directories do not need to exist on all the clients. NetBackup backs up the files in directories that NetBackup finds.

For more information on the **Selections** tab, see “Which Selections Will Be Backed Up: Backup Selections Tab” on page 154.

Disaster Recovery Tab Overview

The **Disaster Recovery** tab appears for those policies based on the Catalog Backup policy type. The **Disaster Recovery** tab contains options for configuring disaster recovery protection methods for the catalog data.

Note When NetBackup is running with Vault, Vault protects the disaster recovery data by sending it to the Vault site as an email attachment of the Vault report email.

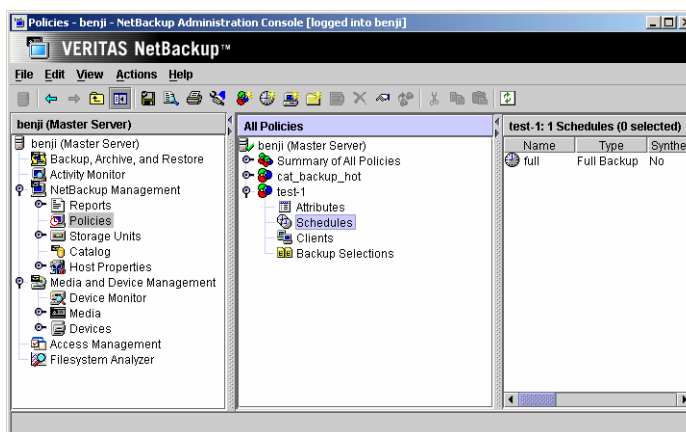
For more information on the **Disaster Recovery** tab, see “Where Will the Catalog Data Be Located: Disaster Recovery Tab” on page 193.

Changing Policies

Try to make changes to policies only during periods when there is no expected backup activity for the affected policies and clients. Preventing this potential conflict lets you make adjustments before backups begin and ensures an orderly transition from one configuration to another.

▼ To add or change schedules in a policy

1. If your site has more than one master server, choose the master server that contains the policy you want to modify.
2. Expand **NetBackup Management > Policies**.
3. Expand the policy name in the middle pane, then select **Schedules**.
4. Perform one of the following actions:
 - ◆ To add a schedule, select **Actions > New > Schedule**. The Add Schedule dialog appears.
 - ◆ To change an existing schedule, double-click the schedule name in the right pane. The Change Schedule dialog appears.
5. Complete the entries in the Attributes tab, Start Window tab, Exclude Dates tab, and Calendar Schedule tab (if Calendar Schedule Type is selected on the Attributes tab). (See “When Will the Job Run: Schedules Tab” on page 97.)




▼ To add or change clients in a policy

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Expand the policy name in the middle pane, then select **Clients**.

Changing Policies

3. Perform one of the following actions:

- ◆ To add a new client, select **Edit > New**. The Add Client dialog appears.
- ◆ To change an existing client, double-click the client name in the right pane. The Change Client dialog appears.




4. Complete the entries in the Add Client or Change Client dialog. (See “To add a client to a policy” on page 149.)

▼ To add or change backup selections in a policy

Note If you are setting up a Vault policy, see “To create a Vault policy” on page 195.

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Expand the policy name in the middle pane, then select **Backup Selections**.
3. Perform one of the following actions:

- ◆ To add a new backup selection, select **Edit > New**.
- ◆ To change an existing backup selection, double-click the backup selection in the right pane.



4. Complete the entries in the New Backup Selections or Change Backup Selections dialog.


If you are unfamiliar with how to specify file paths for your clients, read “Rules for Indicating Pathnames in the Backup Selections List” on page 162 before proceeding.

5. After adding the new backup selection or making changes to an existing selection:
 - ◆ In the New Backup Selection dialog, click **Add**. The new entry appears in the list. After defining all new selections, click **OK**.
 - ◆ In the Change Backup Selection dialog, click **OK**.

▼ To delete schedules, backup selections, or clients from a policy

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.

Note Do not confuse **Cut** and **Delete**. **Cut** copies the selected information to the clipboard, from where you can later paste it. **Delete** does not copy to the clipboard.

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'd like to delete and click the delete button on the toolbar . A confirmation dialog appears.
4. Click **Yes**.

Note Deleting a client from the NetBackup configuration does not delete NetBackup client software from the client. Previous backups for that client can also be recovered up until their expiration date.

Also, deleting a file only deletes the file from the list of files designated for automatic backup. It does not delete the actual file from the disk.

▼ To copy and paste items

You can copy or cut and paste the following items:

- ◆ Copy and paste entire policies
- ◆ Copy and paste schedules

▼ To set the general policy attributes

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Double-click the policy name in the middle pane. The Change Policy dialog appears, containing four policy attribute tabs: **Attributes**, **Schedules**, **Backup Selections**, **Clients**.
3. Select a tab and make any changes.
See the following sections for changes to the **Attributes** tab
 - ◆ “When Will the Job Run: Schedules Tab” on page 97.
 - ◆ “Which Clients Will Be Backed Up: Clients Tab” on page 149.

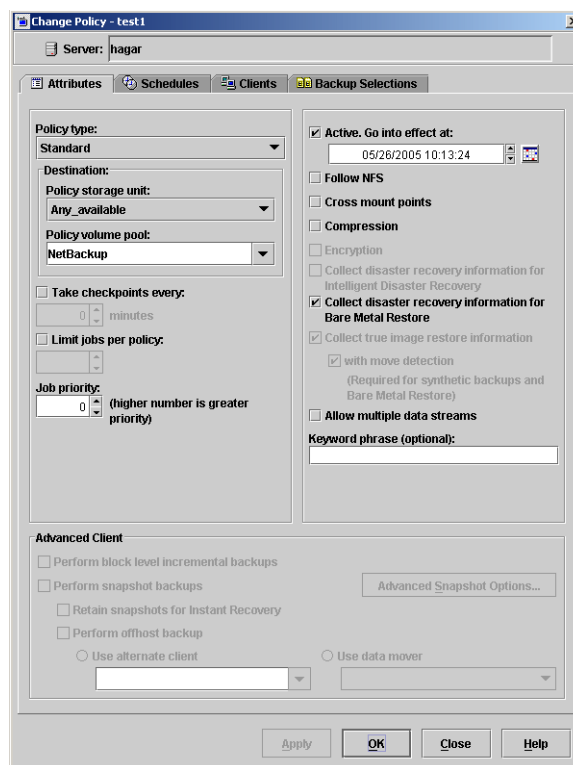
Changing Policies

- ◆ “Which Selections Will Be Backed Up: Backup Selections Tab” on page 154.
- 4. Click **Apply** to save the changes and to keep the dialog open in order to make additional changes. Click **OK** to save the changes and close the dialog.

What Type of Policy: Policy Attributes Tab

The settings on the Attributes tab determine the basic characteristics of all the backups that NetBackup performs according to this backup policy.

The following sections describe the settings on the Attributes tab. Policy attributes are configurable depending on the type of policy and the options installed. For example, **Encryption** is available only when the NetBackup Encryption option is installed.



Policy Type

The **Policy Type** attribute determines the type of clients that can be part of the policy and, in some cases, the types of backups that can be performed on the clients. Select the type of policy from the drop-down list.

If you change the policy type for an existing policy that contains schedules that are invalid for the new policy type, NetBackup prompts you, then either deletes the invalid schedules or, if possible, changes the schedules to an equivalent type.

Policy Type	Description
DataStore	A policy type reserved for use by VERITAS or its partners to provide agents for new applications or databases.

What Type of Policy: Policy Attributes Tab

Policy Type	Description
DB2	Use when the policy will have only clients with the NetBackup for DB2 option. For information on setting up this policy type, see the guide for this option.
Lotus-Notes	Use when the policy will contain only clients with the NetBackup for Lotus Notes option. For information on setting up this policy type, see the guide for this option.
MS-Exchange-Server	Use when the policy will contain only clients with the NetBackup for MS-Exchange option. For information on setting up this policy type, see the guide for this option.
MS-SQL-Server	Use when the policy will contain only clients with the NetBackup for MS-SQL Server option. For information on setting up this policy type, see the guide for this option.
MS-Windows-NT*	Use when the policy will contain only Windows 2000, XP, Windows Server 2003 and/or 5.x NT clients.
NBU-Catalog	Use for hot catalog backup jobs. Allows for a catalog backup while other jobs are running.
NCR-Teradata	Use when the policy will contain only clients with the NetBackup for Teradata option. For information on setting up this policy type, see the guide for this option.
NDMP	Use when the policy will contain only clients with the NetBackup for NDMP option. This policy is available only when the NetBackup NDMP is installed and licensed. For information on setting up this policy type, see the guide for this option.
NetWare	Use when the policy will contain only NonTarget NetBackup Novell NetWare clients (this version uses a Microsoft Windows interface).
Oracle	Use when the policy will contain only clients with the NetBackup for Oracle option. For information on setting up this policy type, see the guide for this option.
Standard*	Use when the policy will contain any combination of the following: <ul style="list-style-type: none"> • UNIX clients (including Mac OS X clients), except those covered by specific such as Oracle. • NetBackup Novell NetWare clients that have the target version of NetBackup software.
Vault	Available only when Vault is licensed. Use as a policy type to schedule and run a Vault job.

Note: The following policy types apply only to UNIX clients.

AFS	Use when the policy will be backing up only AFS file systems on clients. See "Using NetBackup with AFS," in the <i>NetBackup System Administrator's Guide, Volume II</i> for information on setting up these policies.
-----	--

Policy Type	Description
DataTools-SQL-BackTrack	Use when the policy will contain only clients with the NetBackup for DataTools-SQL-BackTrack option. For information on setting up this policy type, see the guide for this option.
FlashBackup	<i>Applies only to NetBackup Enterprise Server:</i> Use when the policy will contain only NetBackup FlashBackup clients on UNIX. This policy is available only when the NetBackup Advanced Client is installed. For information on setting up this policy type, see the <i>Advanced Client System Administration Guide</i> .
FlashBackup-Windows	<i>Applies only to NetBackup Enterprise Server:</i> Use when the policy will contain only NetBackup FlashBackup-Windows clients on Windows. This policy is available only when the NetBackup Advanced Client is installed. For information on setting up this policy type, see the <i>Advanced Client System Administration Guide</i> .
Informix-On-BAR	Use when the policy will contain only clients that are running the NetBackup for Informix option. For information on setting up this policy type, see the guide for this option.
MS-SharePoint	Use to configure a policy for NetBackup for SharePoint Portal Server.
SAP	Use when the policy will contain only clients with the NetBackup for SAP option. For information on setting up this policy type, see the guide for this option.
Split-Mirror	<i>Applies only to NetBackup Enterprise Server:</i> Use when the policy will contain only clients with the NetBackup for EMC option. For information on setting up this policy type, see the guide for this option.
Sybase	Use when the policy will contain only clients with the NetBackup for Sybase option. For information on setting up this policy type, see the guide for this option.

* To utilize CheckPoint Restart for backups, Checkpoint Restart for restores, synthetic backups, or to **Collect Disaster Recovery Information for Bare Metal Restore**, either the **Standard** or **MS-Windows-NT** policy type must be used.

For more details on offhost backup, refer to the *NetBackup Advanced Client System Administrator's Guide*.

Policy Storage Unit

The **Policy Storage Unit** attribute specifies the default storage unit where backups based on the policy will be stored. Select the policy storage unit from the drop-down list.

What Type of Policy: Policy Attributes Tab

If **Any Available** is selected, NetBackup attempts to store data on locally-attached storage units first. NetBackup can be forced to only use a locally attached drive if the **Must use local drive** option is selected.

If a local device is not found or if **Must use local drive** is not selected, then NetBackup tries to find an available storage unit alphabetically.

If **Any Available** is selected, NetBackup uses the first storage unit that meets the following requirements:

- ◆ The storage unit must not be designated as *On Demand Only*
- ◆ The storage unit must have available drives
- ◆ The storage unit must have media available in the required volume pool

An exception to this criteria is the case in which a client is also a media server with locally connected storage units. The locally available storage units take precedence over the sequence based on alphabetical order.

A schedule-level storage unit (when specified) overrides the **Policy Storage Unit** setting. NetBackup uses the default storage unit for all schedules that do not specify a storage unit. (See “Override Policy Storage Unit” on page 114.)

Concerning the Policy Storage Unit and online, hot catalog backups:

- ◆ If the policy is for online, hot catalog backups and you’re indicating a disk storage unit, increase the **Maximum Concurrent Jobs** setting on the storage unit. This ensures that the catalog backup can proceed while regular backup activity is occurring.
- ◆ Online catalog backups must use media servers at version 6.0 or later to store catalog backup data. If your installation contains 5.x and 6.0 media servers hosting disk storage units, do not select *Any Available* for the destination **Policy Storage Unit**, since it is possible that the 5.x media server could be selected.

Notes on Specifying a Storage Unit

- ◆ If your site has only one storage unit or there is no preference for storage:
 - ◆ Specify *Any Available* in the policy for **Policy Storage Unit**
 - ◆ Do not specify a storage unit at the schedule level (See “Override Policy Storage Unit” on page 114.)
 - ◆ In this situation, do not configure all storage units to be *On Demand Only*, or NetBackup will be unable to find an available storage unit for the backups. (See “On Demand Only” on page 43.)
- ◆ If you designate a specific storage unit and the storage unit is unavailable, backups will not run for policies and schedules that require the storage unit.

- ◆ If you have several storage units and want a policy to use more than one but not all of the storage units, select a storage unit group containing the storage units you wish to use.
- ◆ Another method to restrict the storage units used by a policy is the following:
 - a. When configuring volumes in Media Manager, define a volume pool containing volumes available only to the desired storage units.
 - b. In the policy, set **Policy Volume Pool** to the volume pool defined in the previous step.
 - c. For all policies, set **Policy Storage Unit** to *Any Available*.
- ◆ You may have set up a storage unit to be *On Demand Only*. If the storage unit is part of a storage unit group that is needed by a policy, the On Demand Only option is satisfied and the device will be used. (See “On Demand Only” on page 43.)

Policy Volume Pool

The **Policy Volume Pool** attribute specifies the default volume pool for backups of this policy. Select the desired volume pool name from the list of volume pools. Whenever a new volume is required, it is allocated from the volume pool indicated.

Note A schedule-level volume pool, when specified, overrides the policy default set here. (See “Override Policy Volume Pool” on page 114.)

A *volume pool* is a group of media grouped together for use by a single application, protected from access by other applications and users. Volume pools are created in the NetBackup Administration Console in **Media and Device Management > Media > Volume Pools**. Media is assigned to the volume pools for Media Manager storage devices. Disk-type storage devices are not allocated to a volume pool.

NetBackup creates four default volume pools:

- ◆ *None*: The default pool for applications, other than NetBackup and Storage Migrator.
- ◆ *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 (used for online, hot catalog backups) which selects the *CatalogBackup* volume pool by default. Offline, cold catalog backups use media from the *NetBackup* volume pool.
- ◆ *CatalogBackup*: This pool is selected by default for the *NBU-Catalog* policy type. It is used exclusively for online, hot catalog backups. Directing online, hot catalogs to a single, dedicated pool facilitates quicker catalog restores.

What Type of Policy: Policy Attributes Tab

You may want to create additional volume pools. For example:

- ◆ A *Scratch* pool from which NetBackup can automatically transfer volumes when another volume pool does not have media available.
- ◆ An *Auto* volume pool, for use by automatic backups.
- ◆ A *User* volume pool, for use by user backups.

For more information on volume pools, see the *NetBackup Media Manager System Administrator's Guide*.

Volume Pool Override Example

Assume that you want all schedules but one to use the *Backups* pool. The exception is a user-archive schedule that requires the *Archive* pool.

In the policy, set **Policy Volume Pool** to *Backups*. When you set up the schedules for the policy, set **Override Policy Volume Pool** as follows:

- ◆ For schedules that use the *Backups* volume pool, clear **Override Policy Volume Pool**.
- ◆ For the schedule that requires the *Archive* volume pool, select **Override Policy Volume Pool** and specify *Archive* for the pool name.

Checkpoint Restart for Backup Jobs

The **Take Checkpoints Every** (checkpoint restart) check box indicates whether NetBackup will take checkpoints during backup jobs based on this policy at the frequency indicated.

Taking checkpoints during a backup is beneficial if a backup based on this policy fails. Without **Take Checkpoints Every** enabled, a failed backup based on this policy is restarted from the beginning of the job. By taking checkpoints periodically during the backup, NetBackup can retry a failed backup from the last checkpoint rather than restarting the entire job.

The number of times that NetBackup will automatically reattempt a failed backup is configured by the **Schedule Backup Attempts** property located in the master server Global Attributes host properties. (See “Schedule Backup Attempts” on page 414.)

Only policy types *MS-Windows-NT* (for Windows clients) and *Standard* (for UNIX clients) support this policy attribute.

Note Although NetWare clients can use the Standard policy type, checkpoint restart for backups is not supported on NetWare clients.

Checkpoint Frequency

How often NetBackup takes a checkpoint during a backup is configurable. (Default: 15 minutes.) The administrator determines on a policy-by-policy basis how to balance more frequent checkpoints with the likelihood of less time lost in resuming a backup (because of more checkpoints). If the frequency of checkpoints impacts performance, consider increasing the interval time (time between checkpoints).

Checkpoint Restart Support

- ◆ **Multiple Copies:** Checkpoint restart is supported for policies configured to create multiple backup copies. If a copy is configured to allow other copies to continue the job if the copy fails and a subsequent checkpoint occurs, and if **Take Checkpoints Every** is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.
- ◆ **VERITAS Volume Snapshot Provider (VSP):** Checkpoint restart is supported for use with VSP. (See “VSP (Volume Snapshot Provider) Properties” on page 454.)
- ◆ **Advanced Client:** Checkpoint restart is supported for use with local or alternate client backups. However, other methods are not supported: Block Level Incremental Backups, Media Server Copy, Third-Party Copy Device, and Instant Recovery backups.

What Type of Policy: Policy Attributes Tab

- ◆ Disk staging storage units: Checkpoint restart is supported for use in Stage I of disk staging, during which data is backed up to disk. (See “Disk Staging: Stage I” on page 50.) Checkpoint restart is unavailable in the Stage II storage unit policy, during which data is relocated to another storage unit.
- ◆ On Windows clients:
 - ◆ System State backups: No checkpoints are taken during the backup of a System State.
 - ◆ Windows Disk-Image (raw) backups: No checkpoints are taken during a Windows disk-image backup.
 - ◆ Single-instance Store (SIS): No checkpoints are taken for the remainder of the backup after NetBackup encounters a Single-instance Store.
 - ◆ When an incremental backup is resumed and then completes successfully, the archive bits are cleared for the files backed up since the job was resumed, but not for the files backed up prior to the resume. This means the files backed up prior to the resume will be backed up again on the next incremental backup.
- ◆ Synthetic backups: Checkpoint restart is not supported for use with synthetic backups in the current NetBackup release.
- ◆ Checkpoints are not taken for a user archive schedule. If resumed, the user archive restarts from the beginning.
- ◆ NetBackup decides when a new job should be started instead of resuming an incomplete job. NetBackup will start a new job in the following situations:
 - ◆ If a new job is due to run.
 - ◆ If the time since the last incomplete backup has been longer than the shortest frequency in any schedule for the policy.
 - ◆ If the time indicated by the Clean-up property, **Move Backup Job from Incomplete State to Done State**, has passed.
 - ◆ For calendar scheduling, if another run day has arrived.

Checkpoint Restart for Restore Jobs


There is no attribute to check for checkpoint restart for restore jobs. By default, NetBackup automatically takes checkpoints during restore jobs. (See “Checkpoint Restart for Restore Jobs” on page 518.)

A backup that has been checkpointed does not need to be restored with a checkpointed restore job. Conversely, a checkpointed restore job does not need to be restored from a checkpointed backup.

Limit Jobs Per Policy

The **Limit Jobs Per Policy** attribute limits the number of jobs that NetBackup will perform concurrently using this policy. By default, the check box is clear, allowing NetBackup to perform an unlimited number of backup jobs concurrently for this policy. Jobs may be limited by other resource settings.

A configuration could contain enough devices to make it possible for the number of concurrent backups to affect performance. To specify a lower limit, select the check box and specify a value from 1 to 999.



Notes on Limit Jobs Per Policy

The number of concurrent backup jobs that NetBackup can perform depends on the following:

- ◆ **Limit Jobs Per Policy** does not prevent concurrent jobs if the jobs are from different policies.
For example, if there are three policies and each has **Limit Jobs Per Policy** set to 2, NetBackup can start two jobs from each policy and have a total of six running at one time.
- ◆ Parent jobs do not count toward the limit; only the children jobs count. Examples of jobs that produce a parent and children jobs include multistreamed jobs, catalog backups, Advanced Client snapshots, or Bare Metal Restore jobs. (For more information, see "Parent Jobs" on page 318.)
- ◆ Number of storage devices available and multiplexing limits. To process more than one backup job at a time, your configuration must include more than one storage unit, or a storage unit with enough drives to perform more than one backup at a time, or storage units configured to multiplex. With removable media devices such as tape drives, this 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.

What Type of Policy: Policy Attributes Tab

- ◆ Server speed. Too many concurrent backups interfere with the performance of the server. The best number depends on the hardware, operating system, and applications that are running.
- ◆ Network loading. The available bandwidth of the network determines how many backups can occur concurrently. If you encounter loading problems, consider backing up over multiple networks or using compression.

A special case exists when backing up a client that is on the same machine as the server. Here, network loading is not a factor because you do not use the network. Client and server loading, however, is still a factor.

- ◆ Multiplexing. If you use multiplexing, set **Limit Jobs Per Policy** high enough to support the specified level of multiplexing.

Lower values can limit multiplexing within a policy if there are jobs from different schedules within that policy. For example, if **Limit Jobs Per Policy** is at 2 and an incremental backup schedule is due to run for four clients, only two are backed up at a time, regardless of multiplexing settings.

Job Priority

The **Job Priority** attribute specifies the priority that NetBackup assigns to backup jobs for this policy. When a drive becomes available, NetBackup assigns it to the first client in the highest priority policy.

To set the priority, enter a number in the **Job Priority** box. Higher values have higher priority (maximum: 99999; default: 0).

Active. Go Into Effect At

To activate the policy, select the **Active** attribute check box. The policy must be active for NetBackup to run use the policy.

The **Go Into Effect** field specifies when this policy may begin to schedule backups. For example, if today is Monday and you enter Wednesday at 12:00 AM, the policy will not run until that time or later. This is useful for configuring a series of policies in advance of when you want the policies to become active.

To deactivate a policy, clear the **Active** check box. To resume backups, recheck the **Active** box, making sure that the **Go Into Effect** date and time is set to the current time or the time when you want to resume backups.

If the schedule is to be used for a catalog archive, the policy must *not* be active. The **Active** check box must be clear. For more information on configuring a policy for catalog archiving, see “Creating a Catalog Archiving Policy” on page 249.

Backup Network Drives

Note The **Backup Network Drives** attribute applies only to certain policy types.

The **Backup Network Drives** attribute was intended to be used on single user systems, Win95, Win98 and ME. These operating systems are not supported with NetBackup 6.0 and the preferred method for backing up data from a machine other than a NetBackup client is to use UNC pathnames. UNC pathnames are more precise and indicate exactly what should be backed up.

When using **Backup Network Drives** or UNC pathnames, the network (shared) 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 where you are backing up data shared from another machine.

Backup Network Drives must be checked when backing up 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 Since it is not possible to back up mapped drive letters, mapped drive letters do not appear when browsing for backups.

Setup Example Using UNC Pathnames

Assume that:

- ◆ *wildrice* is the NetBackup master server
- ◆ *buck* is a Windows NetBackup client
- ◆ *pepper* is a Windows computer (not necessarily a NetBackup client) and has a shared folder named *TestData*
- ◆ A user wants to back up the folder *TestData* on *pepper* through *buck*.

The steps to perform are as follows:

1. On *wildrice*, the NetBackup master server, create a policy for *buck*.
2. Add \\pepper\TestData to the file list of the policy. This is not needed if the policy is only used for user-directed backups.
3. On *buck*, the NetBackup client:

What Type of Policy: Policy Attributes Tab

- a. Change the NetBackup Client Service on *buck* to either Start Up or Log On with the same account as the user that will do 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.
 - b. Stop and start the NetBackup Client Service so the new account takes effect.
4. A user backup can be done by expanding the **Network** node in the Backup, Archive, and Restore client interface to *pepper* and selecting *TestData*.
5. Scheduled backups will run when scheduled or a manual backup can be performed.

Setup Example using Backup Network Drives Attribute

Assume that:

- ◆ *wildrice* is the NetBackup master server
- ◆ *buck* is a Windows NetBackup client
- ◆ *pepper* is a Windows computer (not necessarily a NetBackup client) and has a shared folder named *share*
- ◆ A user wants to back up the folder named *share* on *pepper* through *buck*.

The steps to perform are as follows:

1. On *wildrice*, the NetBackup master server, select **Backup Network Drives** in the policy to be used for the backup.
2. On *buck*, the NetBackup client:
 - a. Change the NetBackup Client Service on *buck* to either Start Up or Log On with the same account as the user that will do 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.
 - b. Stop and start the NetBackup Client Service so the new account takes effect.
 - c. Create a `bpstart_notify.bat` file that will map a drive on *buck* to `\\pepper\share`
To do this, enter the command:
`net use X: \\pepper\share`
where X: is the mapped drive letter.
3. A user backup can be done by expanding the **Network** node in the Backup, Archive, and Restore client interface to *pepper* and selecting *TestData*.

4. Scheduled backups will run when scheduled or a manual backup can be performed.

Follow NFS

Note The **Follow NFS** attribute applies only to UNIX clients in certain policy types. NetBackup allows selecting it only in those instances.

The **Follow NFS** attribute specifies that you want NetBackup to back up or archive any NFS mounted files that are named in the backup selection list, or by the user in the case of a user backup or archive. Clear the box to prevent the back up or archive of NFS mounted files.

Notes on Follow NFS

- ◆ The behavior of the **Follow NFS** attribute depends on the **Cross mount points** setting (explained later in this chapter).
- ◆ **Follow NFS** has no effect on raw partitions. NFS file systems mounted in a raw partition are not backed up, nor can you back up raw partitions from other machines using NFS mounts to access the raw partitions. The devices are not accessible on other machines through NFS.

Note NetBackup does not support raw partition backups on unformatted partitions.

- ◆ **Follow NFS** causes files in Automounted file systems to be backed up. To exclude automounted directories while allowing backup of other NFS mounts, add an entry for the automounter's mount directory to the exclude list on the client.

Advantages of Using Follow NFS Mounts

Following NFS mounts eliminates the need to locate and log on to the systems where the files actually reside. If the files are mounted on the NetBackup client, you can back up, archive, and restore them by working from the NetBackup client, providing you have the necessary permissions on the NFS mount. One use for this capability is to back up systems that are not supported by NetBackup client software.

Disadvantages of Using Follow NFS

Generally, do not back up NetBackup clients over NFS. It is best to back up and archive files on the NFS server where the files physically reside. NFS backups have lower performance and sometimes problems are encountered. Multiple backups may result if files are backed up at the host where they physically reside and also by local NFS clients that mount the files.

If you select **Follow NFS**, consider using the policy for only the files and clients that you back up or archive over NFS.

Note If **Follow NFS** is not selected, the backup process still reads the client's mount table and evaluates each item in the table, resolving any links to the true pathname. This is necessary so NetBackup can accurately avoid backing up files that reside on NFS-mounted file systems.

When evaluating the mount table, if NetBackup cannot access an NFS file system with the five second default, it assumes the file system to be unavailable. To change the five second default, change the UNIX master server host property, **NFS Access Timeout**. (See "NFS Access Timeout" on page 452.)

Cross Mount Points

The following information applies specifically to UNIX clients.

The **Cross Mount Points** attribute controls whether NetBackup crosses file system boundaries during a backup or archive on UNIX clients or whether NetBackup enters volume mount points during a backup or archive on Windows clients.

- ◆ If you select **Cross Mount Points**, NetBackup backs up or archives all files and directories in the selected path, regardless of the file system. For example, if you specify `root (/)` as the file path, NetBackup backs up `root (/)` and all files and directories under it in the tree. Usually, this means all the client's files, other than those available through NFS.
- ◆ If you clear **Cross Mount Points**, NetBackup backs up or archives only files and directories that are in the same file system as the selected file path. This lets you back up a file path such as `root (/)` without backing up all the file systems that are mounted on it (for example, `/usr` and `/home`).

Notes on Cross Mount Points

- ◆ **Cross Mount Points** has no effect on UNIX raw partitions. If the raw partition that is being backed up is the root partition and has mount points for other file systems, the other file systems are not backed up even if you select **Cross Mount Points**.

- ◆ Do not use **Cross Mount Points** in policies where you use the `ALL_LOCAL_DRIVES` directive in the backup selection list.

Cases That Can Require Separate Policies

In some cases, it is best to create separate policies depending on whether you want to cross mount points. For example, to back up the root file system without also backing up file systems mounted on it, create a policy where **Cross Mount Points** is not selected and the backup selection list contains only `root (/)`. Place other file systems in another policy or policies.

To back up all the data on a client, create a policy where **Cross Mount Points** is selected and the backup selection list includes `root (/)`.


How the Cross Mount Points Interacts With Follow NFS

To back up NFS-mounted files, select **Follow NFS**. The table below summarizes the behavior of **Cross Mount Points** and **Follow NFS**:

Cross Mount Points	Follow NFS	Resulting Behavior
No	No	No crossing of mount points. This is the default.
No	Yes	Back up NFS files if the file path is (or is part of) an NFS mount.
Yes	No	Cross local mount points but not NFS mounts.
Yes	Yes	Follow the specified path across mount points to back up files and directories (including NFS), regardless of the file system where they reside.

Cross Mount Point Examples

The next two examples illustrate the concepts mentioned above. In these examples, assume the client disks are partitioned as shown below.



Here, the client has `/`, `/usr`, and `/home` in separate partitions on disk d1. Another file system named `/home/njr` exists on disk d2 and is mounted on `/home`. In addition, disk d3 contains a directory named `/net/freddie/home` that is NFS-mounted on `/net/freddie`.

Example 1

Assume that you clear **Cross Mount Points** and **Follow NFS** and have the following entries in the backup selection list:

```
/
/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 is processing. It does not back up `/home/njr` or `/net/freddie/home`.

Example 2

Assume that you select **Cross Mount Points** and **Follow NFS** and include only `/` in the backup selection list.

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

The **Compression** attribute specifies that software compression be used for backups made by this policy. Select the check box to enable compression. (Default: no compression.)

Note NetBackup allows you to select **Compression** for the policy types where it applies.

Advantages of Using Compression

Compression reduces the size of a backup by reducing the size of files in the backup. In turn, this decreases the amount of media required for storage. Because the compression and subsequent expansion is performed on the client, compression also decreases the amount of data going over the network and the network load.

Disadvantages of Using Compression

Compression increases computing overhead on the client and increases backup time due to the time required to compress the files. The lower transfer rate associated with compression on the client reduces the ability of some tape devices (notably 8mm) to stream data, thus causing more wear on those devices than would otherwise occur.

The savings in media and network resources, however, still 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.

How Much Compression Can Be Expected?

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 repeating non-unique strings. Some data types are more favorable to compression.

Note When compression is not used, it is normal to receive slightly more data at the server than what exists on the client due to client disk fragmentation and file headers added by the client. (Consider running either the `du` command to tell how much space a file occupies or the `df` command to tell how much free disk space is available.)

Compression Specifications

Types of data that compress well:	Programs, ASCII files, and unstripped binaries (typically 40% of the original size).
--	--

What Type of Policy: Policy Attributes Tab

Compression Specifications (continued)

Best-case compression:	Files composed of repeating, nonunique strings can sometimes be compressed to 1% of their original size.																																				
Types of data that do not compress well:	Stripped binaries (usually 60% of original size).																																				
Worst-case compression:	Files that are already compressed become slightly larger if compressed again. On UNIX clients, if this type of file exists and it has a unique file extension, exclude it (and others with the same extension) from compression by adding it under the UNIX Client host properties (see “Do Not Compress Files Ending With” on page 375).																																				
Effect of file size:	File size has no effect on the amount of compression. It takes longer, however, to compress many small files than a single large one.																																				
Client resources required:	Compression requires client computer processing unit time and as much memory as the administrator configures.																																				
Effect on client speed:	Compression uses as much of the computer processing unit as available and affects other applications that require the computer processing unit. For fast CPUs, however, I/O rather than CPU speed is the limiting factor.																																				
Effect on total backup time:	On the same set of data, backups can take three or more times as long with compression.																																				
Files that are not compressed:	<p>NetBackup does not compress:</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 ending with suffixes specified with the COMPRESS_SUFFIX = <i>suffix</i> option in the <code>bp.conf</code> file.- On UNIX clients, files with the suffixes as shown below: <table><tr><td>.arc or .ARC</td><td>.gz or GZ</td><td>.iff or .IFF</td><td>.sit.bin or</td></tr><tr><td>.arj or .ARJ</td><td>.hqx or .HQX</td><td>.pit or .PIT</td><td>.SIT.bin</td></tr><tr><td>.au or .AU</td><td>.hqx.bin or</td><td>.pit.bin or</td><td>.tiff or .TIFF</td></tr><tr><td>.cpt or .CPT</td><td>.HQX.BIN</td><td>.PIT.BIN</td><td>.Y</td></tr><tr><td>.cpt.bin or</td><td>.jpeg or .JPEG</td><td>.scf or .SCF</td><td>.zip or .ZIP</td></tr><tr><td>.CPT.BIN</td><td>.jpg or .JPG</td><td>.sea or .SEA</td><td>.zom or .ZOM</td></tr><tr><td>.F</td><td>.lha or .LHA</td><td>.sea.bin or</td><td>.zoo or .ZOO</td></tr><tr><td>.F3B</td><td>.lzh</td><td>.SEA.BIN</td><td>.z or .Z</td></tr><tr><td>.gif or .GIF</td><td>.pak or .PAK</td><td>.sit or .SIT</td><td></td></tr></table>	.arc or .ARC	.gz or GZ	.iff or .IFF	.sit.bin or	.arj or .ARJ	.hqx or .HQX	.pit or .PIT	.SIT.bin	.au or .AU	.hqx.bin or	.pit.bin or	.tiff or .TIFF	.cpt or .CPT	.HQX.BIN	.PIT.BIN	.Y	.cpt.bin or	.jpeg or .JPEG	.scf or .SCF	.zip or .ZIP	.CPT.BIN	.jpg or .JPG	.sea or .SEA	.zom or .ZOM	.F	.lha or .LHA	.sea.bin or	.zoo or .ZOO	.F3B	.lzh	.SEA.BIN	.z or .Z	.gif or .GIF	.pak or .PAK	.sit or .SIT	
.arc or .ARC	.gz or GZ	.iff or .IFF	.sit.bin or																																		
.arj or .ARJ	.hqx or .HQX	.pit or .PIT	.SIT.bin																																		
.au or .AU	.hqx.bin or	.pit.bin or	.tiff or .TIFF																																		
.cpt or .CPT	.HQX.BIN	.PIT.BIN	.Y																																		
.cpt.bin or	.jpeg or .JPEG	.scf or .SCF	.zip or .ZIP																																		
.CPT.BIN	.jpg or .JPG	.sea or .SEA	.zom or .ZOM																																		
.F	.lha or .LHA	.sea.bin or	.zoo or .ZOO																																		
.F3B	.lzh	.SEA.BIN	.z or .Z																																		
.gif or .GIF	.pak or .PAK	.sit or .SIT																																			

Encryption

The **Encryption** attribute is selectable only if the NetBackup Encryption option is installed and configured. When the **Encryption** attribute is selected, the server encrypts the backup for the clients listed in the policy. See the *NetBackup Encryption System Administrator's Guide* for more information.

Collect Disaster Recovery Information for Intelligent Disaster Recovery

The **Collect Disaster Recovery Information for Intelligent Disaster Recovery** attribute specifies whether or not NetBackup should collect the information required for Intelligent Disaster Recovery (IDR) during backups of Windows clients using this policy. (See "Configuring NetBackup Policies for IDR" in the *NetBackup System Administrator's Guide, Volume II*.)

Collect Disaster Recovery Information for Bare Metal Restore

The **Collect Disaster Recovery Information for Bare Metal Restore** attribute specifies that the BMR client agent runs on each client before every backup to save the configuration information of the client. The Activity Monitor displays this activity as a separate job.

Bare Metal Restore is a separately-priced option. For more information, see the *Bare Metal Restore System Administrator's Guide for UNIX, Windows, and Linux*.

Only policy types *MS-Windows-NT* (for Windows clients) and *Standard* (for UNIX clients) support this policy attribute. On master servers licensed for BMR, **Collect Disaster Recovery Information for Bare Metal Restore** is enabled by default when creating a MS-Windows-NT or Standard policy.

Collect True Image Restore Information

Note The **Collect True Image Restore Information** attribute applies only to certain policy types. NetBackup allows selecting it only in those instances.

The **Collect True Image Restore Information** attribute specifies that NetBackup will start collecting the information required to restore directories to contain what they had at the time of any incremental (or full backup) that the user chooses to restore. Files that were deleted before the time of the selected backup are not restored. Otherwise, for example, a restore based on the date of an incremental includes all files backed up since the last full backup, including those that were deleted sometime during that period.

NetBackup starts collecting the true image restore information beginning 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 actually changed.

NetBackup does not provide true image restores based on the time of a user backup or archive. It does, however, use the backups from user operations for a true image restore, if they are more recent than the latest automatic full or incremental.

For true image incremental backups to include files that were moved, renamed, or newly installed in the directories, also select **With Move Detection**.

Note **Collect True Image Restore Information With Move Detection** must be selected to create synthetic backups. For more information on configuring synthetic backups, see “Synthetic Backups” on page 107.

Collect True Image Restore With Move Detection

The **Collect True Image Restore With Move Detection** attribute specifies that true image incremental backups include files that were moved, renamed, or newly installed from a tar or zip archive. (Depending on how the files were packaged and how they were installed, some newly installed files will not be backed up by non-TIR incrementals.)

Without move detection, NetBackup skips these files and directories because their modification times are unchanged. With move detection, NetBackup compares path names and inode numbers with those from the previous full or incremental backup. If a name or inode number is new or changed, the file or directory is backed up.

Note This attribute must be selected to create synthetic backups.

The following are examples where using move detection backs up files that otherwise would not be backed up:

- ◆ A file named `/home/pub/doc` is moved to `/home/spec/doc`. Here, the modification time is unchanged but `/home/spec/doc` is new in the `/home/spec/` directory and is backed up.
- ◆ A directory named `/etc/security/dev` is renamed as `/etc/security/devices`. Here, the modification time is unchanged but `/etc/security/devices` is a new directory and is backed up.
- ◆ A file named `/home/pub/doc` is installed by extracting it from a UNIX tar file. Here, the modification time is before the time of the last backup but the `doc` is new in the `/home/pub/` directory and is backed up.
- ◆ A file named `docA` is removed and then a file named `docB` is renamed as `docA`. Here, the new `docA` has the same name but its inode number changed so it is backed up.

NetBackup starts collecting information required for move detection beginning with the next full or incremental backup for the policy. This first backup after setting the attribute always backs up all files, even if it is an incremental.

Move detection consumes space on the client and can fail if there is not enough disk space available.

What Happens During True Image Restores

The following table lists the files backed up in the `/home/abc/doc/` directory during a series of backups between 12/01/2004 and 12/04/2004. Assume that **Collect True Image Restore Information** was turned on for the policy that did the backups.

What Type of Policy: Policy Attributes Tab

Day	Type of Backup	Files Backed Up in /home/abc/doc							
12/01/2004	Full	file1	file2	dirA/fileA	dirB/fileB	file3			
12/02/2004	Incremental	file1	file2	dirA/fileA	-----	-----			
12/03/2004	Incremental	file1	file2	dirA/fileA	-----	-----			
12/04/2004	User backup	file1	file2	dirA/fileA	-----	-----	dirC/fileC	file4	
12/04/2004	Incremental	file1	file2	-----	-----	-----	-----	file4	

Note Dashes (-----) indicate that the file was deleted prior to this backup.

Assume that you are going to restore the 12/04/2004 version of the /home/abc/doc/ directory.

- ◆ If you perform a regular restore, the restored directory contains all files and directories that ever existed in /home/abc/doc/ from 12/01/2004 (last full backup) through 12/04/2004:

```
file1
file2
dirA/fileA
dirB/fileB
file3
dirC/fileC
file4
```

- ◆ If you perform a true image restore of the 12/04/2005 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on 12/04/2005:

```
file1
file2
file4
```

NetBackup does not restore *any* of the files deleted prior to the 12/04/2005 incremental backup.

The restored directory does not include the dirA and dirC subdirectories, even though they were backed up on 12/04/2005 with a user backup. NetBackup did not restore these directories because they did not exist at the time of the incremental backup, which was the reference for the true image restore.

Notes On True Image Restores and Move Detection

- ◆ Because the additional information that NetBackup collects for incrementals is the same as for a full backup, incremental backups take much more disk space when you are collecting true image restore information. Adding move detection requires even more space.
- ◆ You can set the period of time that NetBackup keeps the true image restore information by setting **Keep TIR Information** on the Global properties dialog. (See “Keep True Image Restoration (TIR) Information” on page 359.)
- ◆ Incremental backups are slower for a policy where true image restore information is being collected.
- ◆ If the indexing feature is being used, be aware that the INDEX files consume much more space when true image restore information is being collected. (This warning applies only to ASCII image catalogs—the binary image catalog does not need INDEX files.) For more information, see “Indexing the Catalog for Faster Access to Backups” on page 283.
- ◆ You can perform true image restores only on directories that were backed up by a policy for which NetBackup is collecting true image restore information.

If you intend to restore an entire file system or disk by using a true image restore, ensure that all the desired directories are backed up by a policy that is collecting true image restore information.
- ◆ For true image restores, you can list and select only directories. In true image restore mode, the client-user interface does not display individual files. Refer to the online Help in the Backup, Archive, and Restore client interface for more information on performing true image restores.
- ◆ A true image restore preserves files that are currently in the directory but were not present when the backup was completed. In our previous example, assume you created a file named file5 after the incremental backup occurred on 12/04/2004, but before doing the restore. In this case, the contents of the directory after the restore is:

```
file1  
file2  
file4  
file5
```

Allow Multiple Data Streams

The **Allow Multiple Data Streams** attribute specifies that, depending on the directives or scripts/templates (of database policy types) in the backup selection list, NetBackup can divide automatic backups for each client into multiple jobs, with each job backing up only a part of the backup selection list. The jobs are in separate data streams and can occur concurrently.

- ◆ The number of streams (backup jobs) that start for each client and how the backup selection list is divided into separate streams is determined by the directives, scripts, or templates specified in the backup selection list. (See “Backup Selections List Directives for Multiple Data Streams” on page 184.)
- ◆ The total number of streams that can run concurrently is determined by the following settings:
 - ◆ Number of available storage units
 - ◆ Multiplexing settings
 - ◆ Maximum jobs parameters

(See “Adjusting Multiple Data Streams” on page 94.)

Multistreamed jobs consist of a parent job to perform stream discovery, and children jobs for each stream. In the Activity Monitor, the children jobs display the Job ID of the parent job. Parent jobs display a dash (-) in the Schedule column.

Note If **Allow Multiple Data Streams** is in use, and a file system exists in an exclude list for a client, a NetBackup job appears in the Activity Monitor for the file system that was excluded. This is normal behavior and none of the files in the excluded file system will be backed up.

When to Use Multiple Data Streams


Reduce Backup Time

Multiple data streams can reduce the backup time for large backups. This is achieved by splitting the backup into multiple streams and then using multiplexing, multiple drives, or a combination of the two for processing the streams concurrently.


In addition, configuring the backup so each physical device on the client is backed up by a separate data stream that runs concurrently with streams from other devices can significantly reduce backup times.

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 because the heads must move back and forth between tracks containing files for the respective streams.

Recommended for Best Performance



Not Recommended



Reduce Retry Time for Backup Failures

Because the backup streams are completely independent, 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 and NetBackup can restart the failed stream without restarting the others.

For example, assume the backup for a 10 GB partition is split into 5 streams, each containing 2 GB. If the last stream fails after writing 1.9 GB (a total of 9.9 GB backed up), NetBackup retries only the last 2 GB stream. If this 10 GB partition is backed up without multiple data streams and a failure occurs, the entire 10 GB backup must be retried.

The Schedule Backup Attempts host property applies to each stream. For example, if Schedule Backup Attempts is set to 3, NetBackup retries each stream a maximum of three times (see "Schedule Backup Attempts" on page 414.).

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.

What Type of Policy: Policy Attributes Tab

Reduce Administration—More Backups With Fewer Policies

When a configuration contains large file servers with many file systems and volumes, using multiple data streams will provide more backups with fewer policies than are otherwise required.

Adjusting Multiple Data Streams

The two aspects of multiple data streams that you can adjust are the total number of streams and the number of streams that can run concurrently.

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 because the heads must move back and forth between tracks containing files for the respective streams.

Setting the Total Number of Streams

The backup selection list determines the total number of streams that are started. The `NEW_STREAM` directive allows you to explicitly configure a fixed number of streams, or you can have the client dynamically define the streams. (See “Backup Selections List Directives for Multiple Data Streams” on page 184.)

Setting the Number of Streams That Can Run Concurrently

The number of streams that can run concurrently for a policy or client is determined by the following:

- ◆ Storage unit and schedule multiplexing limit (See “Media Multiplexing” on page 117.)
- ◆ Number of drives that are available
- ◆ Maximum concurrent jobs settings for the policy and client

Each storage unit and each schedule has a maximum multiplexing setting. The lower of the two settings is the limit for a specific schedule and storage unit. The maximum number of streams that can be multiplexed is limited to the sum of the multiplexing limits for all drives available in the storage unit and schedule combinations.

For example, assume there are two storage units with 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.

The maximum jobs settings also limit the maximum number of streams:

- ◆ **Maximum Jobs Per Client (Host Properties > Master Servers > Global Attributes)** (see “Maximum Jobs per Client” on page 415)
- ◆ **Limit Jobs Per Policy** (policy attribute)

- ◆ **Maximum Data Streams** (use **Host Properties > Master Servers > Client Attributes** or use the `bpclient` command `-max_jobs` option as shown below) (see “Maximum Data Streams” on page 363)

The maximum job settings are interdependent as follows:

- ◆ If **Maximum Data Streams** is *not* set, the limit is either **Maximum Jobs Per Client** or **Limit Jobs Per Policy**, whichever is lower.
- ◆ If **Maximum Data Streams** is set, NetBackup ignores **Maximum Jobs Per Client** and uses either **Maximum Data Streams** or **Limit Jobs Per Policy**, whichever is lower.

To specify a value for **Maximum Data Streams** with the `bpclient` command:

1. Determine if the client is in the client database on the master server by running the following command on one line:

```
/usr/opensv/netbackup/bin/admincmd/bpclient -client name -L
```

2. If the client is not in the client database, run the following command on the master server on one line:

```
/usr/opensv/netbackup/bin/admincmd/bpclient -client name -add  
-max_jobs number
```

3. If the client is in the client database, run the following command on one line:

```
/usr/opensv/netbackup/bin/admincmd/bpclient -client name -modify  
-max_jobs number
```

Keyword Phrase

The **Keyword Phrase** attribute is a phrase that NetBackup associates with all backups or archives based on this policy. Clients can then list or restore only the backups that have the this phrase associated with them.

Note Only the Windows and UNIX client interfaces support keyword phrases.

You can use the same keyword phrase for more than one policy, making it possible to link backups from related policies. For example, you can use one keyword phrase for full backups and another for incremental backups.

The phrase can be a maximum of 128 characters in length. All printable characters are permitted including spaces and periods. By default, there is no keyword phrase.

Clients can also specify a keyword phrase for a user backup or archive. A user keyword phrase overrides the policy phrase.

What Type of Policy: Policy Attributes Tab

Advanced Client Options

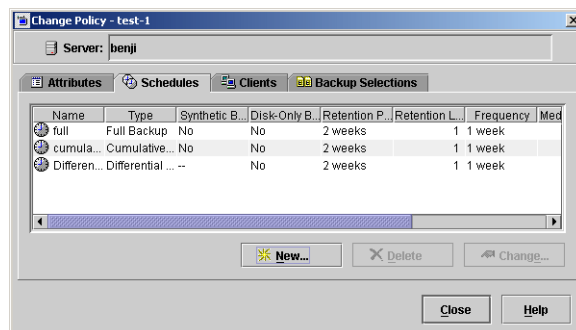
To use the Advanced Client options, you must install and license the Advanced Client option. For more details on offhost backup, refer to the *NetBackup Advanced Client System Administrator's Guide*.

When Will the Job Run: Schedules Tab

The Schedules tab displays the time schedules for the policy selected.

From the policy Schedules tab:

- ◆ Create a new schedule by clicking **New**.
- ◆ Edit a schedule by selecting the schedule and clicking **Change**.
- ◆ Delete a schedule by selecting the schedule and clicking **Delete**.



When creating or editing a schedule, schedule attributes appear on four tabs:

- ◆ **Attributes** tab: Schedule the time and frequency at which a task will run, along with other scheduled attributes.
- ◆ **Start Window** tab: Schedule the time on each day that a task will run.
- ◆ **Exclude Dates** tab: Indicate the dates that you do *not* want a task to run.
- ◆ **Calendar Schedule** tab: Schedule the run days for a task by indicating specific dates, recurring weekdays, recurring days of the month. (This tab appears only when Calendar is selected as the Schedule type.)

Note A policy can contain more than one schedule. VERITAS recommends, however, that you do not mix calendar-based and frequency-based schedule types within the same policy. Under some conditions, combining the schedule types can cause unexpected results.

▼ To create or change schedules

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. In the middle pane, double-click the policy name where you want to change or add a schedule. The **Change Policy** dialog appears.
3. Select the **Schedules** tab. The tab displays the properties of existing schedules. The title bar displays the name of the current policy.
4. Select the schedule you wish to change and click **Change**.
5. The Change Schedule dialog appears containing the **Attributes**, **Start Window**, and optionally, the **Exclude Dates** and **Calendar Schedule** tabs.

Schedule Attributes Tab

6. Make your changes and click **OK**.

Note “To add or change schedules in a policy” on page 65 also provides information on changing existing policies.

Schedule Attributes Tab

The **Attributes** tab contains options that define the backup type, when the backup can occur, and how long the backup image is to be kept. Other attributes such as type of storage and volume pool can also be defined.

Name

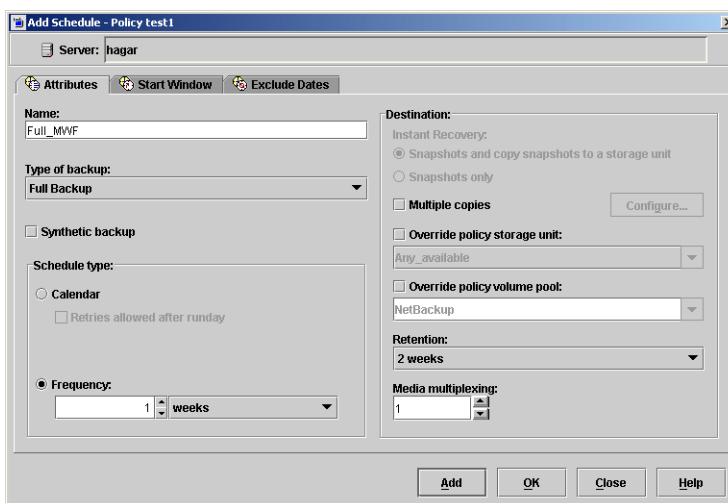
Specify a name for the schedule by typing it in the **Name** field.

(See “NetBackup Naming Conventions” on page 29.)

The schedule name appears on screens and messages about the schedule.

If the schedule is a *relocation schedule*, created as part of configuring a disk staging storage unit, the schedule name cannot be changed, and defaults to the name of the storage unit. For

more information on disk staging storage units, see “About Disk Staging” on page 47.



Type of Backup

The **Type of Backup** specifies the type of backup that the schedule controls. Select a backup type from the drop-down list. The list displays only the backup types that apply to the policy being configured.

If the schedule is a *relocation schedule*, created as part of configuring a disk staging storage unit, there is no backup type selection to be made.

The following sections describe the various backup types.

Full Backup

A full backup backs up all files specified in the backup selections list for the policy, 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 in order to perform a complete restore. If you're configuring a policy for a raw partition backup (formatted partitions only), you must select **Full Backup**.

Cumulative Incremental Backup

A cumulative incremental backup backs up all files specified in the backup selections list for a policy that have changed since the last successful full backup. All files are backed up if no prior backup has been done. Cumulative incremental backups occur automatically according to schedule criteria. A complete restore requires the last full backup and the last cumulative incremental backup.

For more information on incremental backups, see "More on Incremental Backups" on page 101.

Note VERITAS 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 (default). For more information, see "Time Overlap" on page 380.

Differential Incremental Backup

A differential incremental backup backs up all files specified in the backup selections list for the policy that have changed since the last successful incremental (differential or cumulative) or full backup. All files are backed up if no prior backup has been 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 incrementals that have occurred since the last full backup.

For more information on incremental backups, see "More on Incremental Backups" on page 101.

Schedule Attributes Tab

User Backup

A user backup is initiated by the user 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 allowed on the schedule **Start Window** tab.

If the schedule is to be used for a catalog archive, **User Backup** must be selected for the backup type. For more information on configuring a policy for catalog archiving, see “Creating a Catalog Archiving Policy” on page 249.

User Archive

A user archive is initiated by the user through the interface on the client and archives all files that the user specifies. An archive is a special type of backup that first backs up the file and then deletes it from the local disk if the backup is successful. This frees local disk space while still keeping a copy for future use (until the retention period expires). Users can start archives only during the times that you specify in the schedule **Start Window** tab.

Application Backup

An application backup is a backup type that applies to all database agent clients. For more information on configuring schedules for this type of backup, see the NetBackup guide that came with the product.

Automatic Backup

An automatic backup is a backup type for all database agent clients, except NetBackup for Informix and Oracle. For more information on configuring schedules for this type of backup, see the NetBackup guide for the database product.

Automatic Incremental Backup

An automatic incremental backup applies only to NetBackup for Informix clients. For more information on configuring schedules for this type of backup, see the *NetBackup for Informix System Administrator's Guide*.

Automatic Cumulative Incremental Backup

An automatic cumulative incremental backup applies only to NetBackup for Oracle clients. For more information on configuring schedules for this type of backup, see the *NetBackup for Oracle System Administrator's Guide*.

Automatic Differential Incremental Backup

An automatic differential incremental backup applies only to NetBackup for Oracle clients. For more information on configuring schedules for this type of backup, see the *NetBackup for Oracle System Administrator's Guide*.

Automatic Full Backup

An automatic full backup applies only to NetBackup for Informix and for Oracle clients. For more information on configuring schedules for this type of backup, see the *NetBackup for Informix System Administrator's Guide* or *NetBackup for Oracle System Administrator's Guide*.

Automatic Vault

An automatic Vault session applies only to Vault policies. This does not run a backup, but instead runs the vault command 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.

For more information on configuring a Vault policy, see "Creating a Vault Policy" on page 195.

Vault Catalog Backup

Use when the schedule is being created for a Catalog Backup policy that would be used by Vault. Available only when Vault is licensed.

If the schedule is a Vault Catalog Backup type, one of two schedule attribute combinations must be configured or the schedule cannot be saved:

- ◆ Check and configure **Multiple Copies** (see note below), *or*
- ◆ Check **Override Policy Storage Unit** (see note below), **Override Policy Volume Pool** and specify the **Retention**.

Note The selected storage unit selection should not be *Any Available*.


More on Incremental Backups

The following examples show how data is included in a series of full and incremental backups.

Schedule Attributes Tab

Full and Differential Incremental Example

A differential incremental backup backs up data that has changed since the last full or differential incremental backup. The following figure shows how data is included in a series of full and differential incremental backups between January 1 and January 4.




Recovery = Jan 1 (full) + Jan 2 (incr) + Jan 3 (incr) + Jan 4 (incr)

The January 1 full backup includes all files and directories in the policy backup selections list. The subsequent differential incremental backups include only the data that changed since the last full or differential incremental backup. If the disk fails sometime on January 4 (after the backup), the full backup and all three of the incremental backups are required for the recovery.

Full and Cumulative Incremental Example

A cumulative incremental backup backs up data that has changed since the last full backup. The following example shows how data is included in a series of full and cumulative incremental backups between January 1 and January 4. The January 1 full backup includes all files and directories in the policy backup selections list. Each of the cumulative incremental backups include the data 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.



Recovery = Jan 1 (full) + Jan 4 (cumulative incremental)

Retention Requirements for Differential and Cumulative Incremental Backups

Type	Retention Requirement	Comments
Differential	Longer	It is necessary to have the last full backup and all the differential incrementals that have occurred since the last full backup in order to ensure that all files can be restored. Therefore, all the differentials must be kept until the next full backup occurs.
Cumulative	Shorter	Each cumulative incremental backup contains all the changes that have occurred since the last full backup. Therefore, a complete restore requires only the most recent cumulative incremental in addition to the full backup.

Relative Backup and Restore Times for Differential and Cumulative Incremental Backups

Type	Backup Time	Restore Time	Comments
Differential	Shorter	Longer	Less data in each backup, but all differential incremental backups are required since the last full backup for a restore. This results in a longer restore time.
Cumulative	Longer	Shorter	More data in each backup, but only the last cumulative incremental is required for a complete restore (in addition to the full).

It is possible to use a combination of cumulative and differential incremental backups in order to obtain some of 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.)

Backup Type	Frequency	Retention Period
Full	6 days	2 weeks
Cumulative incremental	2 days	4 days
Differential incremental	1 day	2 days

Schedule Attributes Tab

The following set of schedules results in the series of backups shown below:

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8
Full	Differential	Cumulative	Differential	Cumulative	Differential	Full	Differential

Notes about example:

- ◆ Every other day a differential incremental backup occurs, which usually has a minimum backup time.
- ◆ On alternate days, a cumulative incremental backup occurs, which requires more time than the differential backup, but not as much time as a full backup. The differential backup can now be expired.
- ◆ To recover all files requires, at most, two incremental backups in addition to the most recent full backup. This typically 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.

Determining Files Due for Backup on Windows Clients

On Windows clients, NetBackup performs incremental backups of files based on the **Perform Incrementals Based on Archive Bit** setting. This setting is found in the Backup, Archive and Restore client interface, under **File > NetBackup Client Properties**, on the **General** tab.

If the **Perform Incrementals Based on Archive Bit** check box is checked, incremental backups for this client are based on the state of the archive bit of each file. The operating system sets the bit whenever a file is changed and it remains set until cleared by NetBackup. The conditions under which NetBackup clears the bit depend on the type of backup being performed.

- ◆ For a full backup, NetBackup backs up files regardless of the state of their archive bit. After a full backup, the archive bit is always cleared.
- ◆ For a differential incremental backup, NetBackup backs up files that have the archive bit set and have therefore been changed. When the client receives a response from the server indicating that the backup was successful (or partially successful) the archive bits are cleared. This allows the next differential incremental to back up only files that have changed since the previous full or differential incremental backup.
- ◆ For a cumulative incremental backup, NetBackup backs up files that have the archive bit set, but does not clear the archive bits after the backup. This allows the next cumulative incremental to back up not only changed files, but also files that were in this cumulative incremental.

If the **Perform Incrementals Based on Archive Bit** check box is clear, NetBackup includes a file in an incremental backup only if the datetime stamp of the file has been changed since the last backup. The datetime stamp indicates when the file was last backed up.

- ◆ For a full backup, NetBackup backs up files regardless of the datetime stamp.
- ◆ For a differential incremental backup, NetBackup compares the datetime stamp of the file against the last full or incremental backup.
- ◆ For a cumulative incremental backup, NetBackup compares the datetime stamp of the file against the last full backup.

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

Determining Files Due for Backup on UNIX Clients

When performing incremental backups on NetBackup UNIX clients, all relevant files and directories are looked at to determine if they are due for backup based on a reference date (that is, back up all files changed since *date_x*).

UNIX files and directories have three times associated with them:

- ◆ *mtime*: The file modification time. The *mtime* for a file or directory is updated by the file system each time the file is modified. Prior to modifying a file, an application can save the *mtime* of the file, then reset it after the modification using the `utime(2)` system call.
- ◆ *atime*: The file access time. The *atime* for a file or directory is updated by the file system each time the file is accessed (read or write). Prior to accessing a file, an application can save the *atime* of the file, and then reset it after the file access using the `utime(2)` system call.
- ◆ *ctime*: The inode change time. The *ctime* for a file or directory is updated each time the file or directory's inode is changed; examples of this are changing permissions, ownership, link-counts, and so on. The *ctime* for a file or directory cannot be saved before and reset after a change. Another significant fact is that the *ctime* of a file or directory is changed when resetting the *mtime* and *atime* (using the `utime(2)` system call) for the file.

UNIX man pages contain a definition of these attributes.

When NetBackup reads the data for a file that is included in a backup, it does not affect the file modification time, but does affect the access time of the file. For this reason, NetBackup saves the *atime* and *mtime* of the file prior to reading the file, and (by default) resets the *atime* and *mtime* using the `utime(2)` system call. By doing it this

Schedule Attributes Tab

way, NetBackup does not cause problems for storage migration products or administrator scripts that are utilizing 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.

As an option to a NetBackup configuration, customers can choose to have NetBackup not reset the access time of the file after it reads a file. Additionally, customers can choose to have NetBackup use the `ctime` of the file, in addition to the `mtime`, when determining what files to back up in an incremental. Normally, these two options are used together, but there may be sites which 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 is only using the file modification time (`mtime`) to determine files due to be backed up during an incremental backup, it will not detect these moved files. For sites where this is an issue, the `ctime` should also be used (if possible) to determine files due to be included in an incremental backup, using the `bp.conf` attributes `USE_CTIME_FOR_INCREMENTALS` and `DO_NOT_RESET_FILE_ACCESS_TIME`.

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. Using file timestamps, there is no reliable method for determining that files within a moved directory need to be included in an incremental backup.

In either case, these moved files and directories are included in subsequent full backups.

Synthetic Backups

A synthetic full or synthetic cumulative incremental backup is a backup assembled from one previous, traditional (non-synthesized) full backup, *and* subsequent differential backups and/or a cumulative incremental backup. A client can then use the synthesized backup to restore files and directories in the same way that a client restores from a traditional backup.

Synthetic backups can be written to tape or disk storage units, or a combination of both mediums.

For more information on synthetic backups, see “More About Synthetic Backups” on page 198.

Calendar Schedule Type

Calendar-based scheduling allows administrators to select specific days to run a policy. Choosing the **Calendar** schedule attribute causes the **Calendar Schedule** tab to appear in the Change Schedule dialog. For details on calendar-based scheduling, see “Calendar Schedule Tab” on page 122.

If the schedule is a relocation schedule, created as part of configuring a disk staging storage unit, a calendar-based schedule determines which days images are swept from the disk staging storage unit to the final destination storage unit.

Note A policy can contain more than one schedule. VERITAS recommends, however, that you do not mix calendar-based and frequency-based schedule types within the same policy. Under some conditions, combining the schedule types can cause unexpected results.

For details on how calendar-based scheduling works with time windows, see “How Calendar Scheduling Interacts with Daily Windows” on page 125.

Retries Allowed After Runday

Select **Retries Allowed After Runday** to have NetBackup attempt to complete this schedule until the backup is successful. With this attribute selected, the schedule will attempt to do this, even after a specified run day.

Frequency Schedule Type

Using the **Frequency** schedule type, administrators specify how much time must elapse between the successful completion of a scheduled task and the next attempt at the task.

For example, automatic backups for clients using the following schedule. 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.

If the schedule is a *relocation schedule*, created as part of configuring a disk staging storage unit, a frequency-based schedule determines how often images are swept from the disk staging storage unit to the final destination storage unit.

To set the frequency, click in the **Frequency** field and type a number or select a value from the drop-down list. Select a **Frequency** of hours, days, or weeks.

Note **Frequency** does not apply to user schedules because the user can perform a backup or archive whenever the time window is open.

Note A policy can contain more than one schedule. VERITAS recommends, however, that you do not mix calendar-based and frequency-based schedule types within the same policy. Under some conditions, combining the schedule types can cause unexpected results.

Guidelines for Setting Backup Frequency

Choose the backup frequency based on how often you must back up your files to ensure that you can restore critical changes in case of a disk failure. How often the data changes is an important factor in determining backup frequency. For example, determine if files change several times a day, daily, weekly, or monthly. Determine the rate of change by analyzing typical file usage.

Typically, sites perform daily backups to preserve each day's work. This ensures that, at most, only one day's work is lost in case of a disk failure. More frequent backups are necessary when data changes many times during the day and these changes are important and difficult to reconstruct.

Daily backups are usually incremental backups that record the changes since the last incremental or full backup. This conserves resources because incremental backups 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 too many consecutive incremental backups. Too many incremental backups between full backups increases restoration time because of the effort required to merge those incremental backups when restoring files and directories. When setting the frequency for full backups:

- ◆ Choose longer times between full backups for files that seldom change. This uses fewer system resources. It also does not significantly increase recovery time because there should be smaller incremental backups.
- ◆ Choose shorter times between full backups for files that change frequently. This decreases restore time. It can also use less resources because 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 approximately half of the files in a policy selection list change frequently enough to require a full backup every week, but the remaining files rarely change and require only monthly full backups. Here, if all the files are in the same policy, you must perform full backups weekly on all the files. This wastes system resources and media because half the files need full backups only once a month. A better approach is to divide them into two policies, each with the appropriate backup schedule or to consider using synthetic backups.

Backup Frequency Determines Schedule Priority

If more than one automatic schedule is due for a client within a policy, the backup frequency determines the schedule that NetBackup uses:

- ◆ Jobs from the schedule with the lower frequency (longer period between backups) always get higher priority. For example, a schedule with a backup frequency of one year has priority over a schedule with a backup frequency of one month.
- ◆ If NetBackup encounters a backup policy with two schedules (one full, one incremental) that are each due to run, are each within their defined time window, and are each configured with the same frequency value, the schedule that is alphabetically first will be chosen to run.

For example, NetBackup prioritizes the following three schedules in the order shown:

1. monthly_full (frequency is one month)
2. weekly_full (frequency is two weeks)
3. daily_incremental (frequency is one week)

If all three schedules are due for a client, NetBackup adds the job for the monthly full to the worklist and skips the other two.

For an explanation of how NetBackup prioritizes each backup job that it adds to its worklist, see Chapter 5 in the *NetBackup System Administrator's Guide, Volume II*.

Instant Recovery

The **Instant Recovery** options are available under the following conditions:

- ◆ The **Advanced Client** option is licensed and installed. Refer to the *NetBackup Advanced Client System Administrator's Guide*.
- ◆ **Perform Snapshot Backups** is selected.
- ◆ **Retain Snapshots for Instant Recovery** is selected.

Multiple Copies

Using the **Multiple Copies** attribute, NetBackup can create up to four copies of a backup simultaneously, provided that the storage units are on the same media server and there are sufficient resources available for each copy. For example, creating four copies simultaneously in a Media Manager storage unit requires four tape drives.

At one time, this option was referred to as Inline Tape Copy (ITC), but can now accommodate disk storage units.

The **Maximum Backup Copies** property specifies the total number of backup copies that may exist in the NetBackup catalog (2 through 10). NetBackup creates either the number of copies specified under **Multiple Copies**, or the number of copies specified as the **Maximum Backup Copies** property, whichever is smaller. (See “Maximum Backup Copies” on page 416.)

To create more than four copies, additional copies may be created at a later time using duplication.

The storage units used for multiple copies must be configured to allow a sufficient number of concurrent jobs to support the concurrent copies (**Maximum Concurrent Jobs** or **Maximum Concurrent Write Drives** setting).

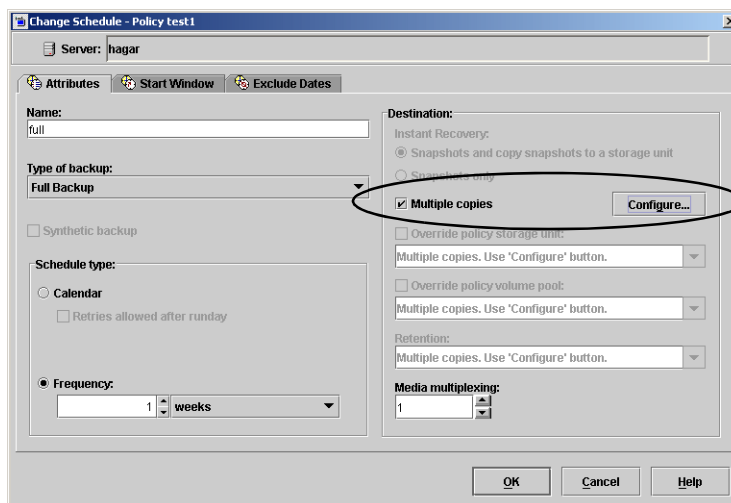
If you create multiple original images simultaneously, the backup time required may be longer than for one copy. Also, if you specify both Media Manager and disk storage units, the duration of disk write operations will match that of slower removable media write operations.

Note The **Multiple Copies** option does not support the following storage types: third-party copies or optical devices.
Also, **Multiple Copies** does not support storage units that use a QIC (quarter-inch cartridge) drive type.

Multiple copies can be created for a regular backup policy or for a relocation schedule, created as part of a disk staging storage unit.

▼ To configure a schedule to create multiple copies

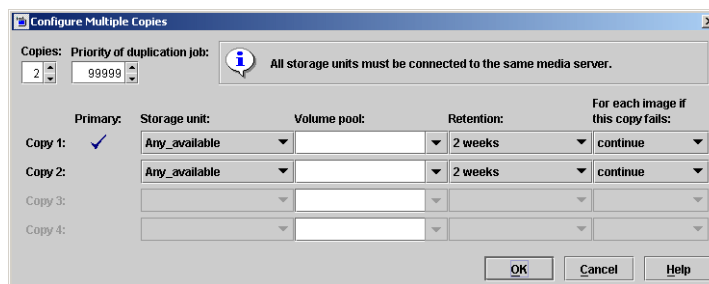
1. Expand **NetBackup Management > Policies**. Double-click an existing policy or select **Edit > New** to create a new policy.
2. Select the **Schedules** tab.
3. Double-click an existing schedule or click **New** to create a new schedule.



Schedule Attributes Tab

4. In the Attributes tab, select **Multiple Copies**, then click **Configure**. The Configure Multiple Copies dialog appears.

5. In the **Copies** field, specify the number of copies to be created simultaneously. The maximum is four, or the number of copies specified by the **Maximum Backup Copies** setting, whichever is smaller.



(See “Maximum Backup Copies” on page 416.)

Copy 1 is the primary copy. If Copy 1 fails for some reason, the first successful copy is the primary copy.

6. Specify the priority that the duplication job will have over other jobs in the queue (0 to 99999).
7. Specify the storage unit where each copy will be stored. Select *Any_Available* for NetBackup to select the storage unit at runtime.
If a Media Manager storage unit has multiple drives, the storage unit can be used for both the original image and copies.
8. Specify the volume pool where each copy will be stored.
9. Select the retention level for each copy. (See “Retention” on page 115.)
10. In the event that the copy does not complete, select whether you’d like the entire job to fail (**fail all copies**), or whether you’d like the remaining copies to continue.

If a copy is configured to allow other copies to continue the job if the copy fails, and if **Checkpoint Restart** is selected for this policy, only the last failed copy that contains a checkpoint can be resumed.

11. Click **OK**, until the policy is saved.

▼ **To configure a disk staging relocation schedule to create multiple copies**

1. Expand **NetBackup Management > Storage Units**.
 - ◆ Double-click an existing disk staging storage unit, or
 - ◆ Select **Actions > New > Storage Unit** to create a new disk staging storage unit.

If creating a new disk staging storage unit, select the **Temporary Staging Area** checkbox and configure the other storage unit selections.

2. Click the **Disk Staging Schedule** button.

3. In the Attributes tab, specify the priority that NetBackup should assign to the duplication jobs. Range: 0 (default) to 99999 (highest priority).

4. Select a schedule type and schedule when you'd like the policy to run.

5. Select whether to use an alternate read server. The server indicated here is allowed to read a backup image originally written by a different media server.
6. Select **Multiple Copies**, then click **Configure**. The Configure Multiple Copies dialog appears.
7. In the **Copies** field, specify the number of copies to be created simultaneously. The maximum is four, or the number of copies specified by the **Maximum Backup Copies** setting, whichever is smaller. (See "Maximum Backup Copies" on page 416.)
Copy 1 is the primary copy. If Copy 1 fails for some reason, the first successful copy is the primary copy.
8. Specify the storage unit where each copy will be stored. If a Media Manager storage unit has multiple drives, it can be used for both the source and the destination.
9. Specify the volume pool where each copy will be stored.
10. Select the retention level for each copy. (See "Retention" on page 115.)
11. In the event that the copy does not complete, select whether you'd like the entire job to fail, or whether you'd like the remaining copies to continue.

Schedule Attributes Tab

If a copy is configured to allow other copies to continue the job if the copy fails, and if **Checkpoint Restart** is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. Click **OK**.

Override Policy Storage Unit

The **Override Policy Storage Unit** attribute specifies whether to use the policy storage unit or another one for this schedule.

- ◆ To override the policy storage unit, select the check box. Choose the storage unit from the drop-down list of previously configured storage units. If the list is empty, no storage units have been configured yet.
- ◆ To use the policy storage unit, do not select the check box. NetBackup uses the policy storage unit you specified with the **Policy Storage Unit** General Attribute. If you did not specify a policy storage unit, NetBackup uses any available storage unit. (See “Policy Storage Unit” on page 71.)

Policy Storage Unit Example

Assume that all the schedules for a policy except one use a Tape Stacker 8MM. The schedule that is the exception requires a Tape Library DLT. Here, you specify Tape Stacker 8MM at the policy level and specify the following on the schedules:

- ◆ For schedules that can use the Tape Stacker 8MM, clear **Override Policy Storage Unit**. When these schedules run, NetBackup uses a Tape Stacker 8MM.
- ◆ For the schedule that requires DLT, select **Override Policy Storage Unit** and select Tape Library DLT. When this schedule runs, NetBackup overrides the policy default and uses the DLT library.

Override Policy Volume Pool

The **Override Policy Volume Pool** attribute specifies whether to use the policy volume pool or another one for this schedule.

- ◆ To override the volume pool specified by the **Policy Volume Pool** General Attribute, select the check box. Choose the volume pool from the list of previously configured volume pools.
- ◆ To use the policy volume pool, do not select the check box. NetBackup uses the volume pool you specified with the **Policy Volume Pool** General Attribute. If you did not specify a policy volume pool, NetBackup uses *NetBackup* as the default, or, for NBU-Catalog policy types, *CatalogBackup*.

Retention

The **Retention** attribute specifies how long NetBackup retains the backups according to a schedule. To set the retention period, select a time period (or level) from the drop-down list. When the retention period expires, NetBackup deletes information about the expired backup, making the files in the backups unavailable for restores. If two weeks is selected, data can be restored from a backup done by this schedule for only two weeks after the backup.

Guidelines for Setting Retention Periods

The retention period for data usually depends on how likely you are to need it after a certain period of time. Some data, such as tax and other financial records, have legal requirements for retention. Other data, such as preliminary documents can probably be expired when the final version is complete.

How long you keep a backup also depends on what you need to recover from it. For example, if day-to-day changes are critical, you must keep all the incremental backups in addition to full backups for as long as you need the data. If incremental backups only track work in progress toward monthly reports, then you can probably expire the incremental backups sooner and rely on the full backups for long term recovery.

When deciding on retention periods, establish guidelines that apply to most of your data. After establishing guidelines, note files or directories that have retention requirements outside of these guidelines and plan to create a separate policy (or policies) for them. For example, placing files and directories with longer retention requirements in a separate policy allows you to schedule longer retention times for them without keeping all the others for the longer time period.

Another consideration for data retention is offsite storage of the backup media. This protects against fires or other disasters that occur at the primary site. Set the retention period to infinite for backups you must retain for more than one year.

- ◆ One method of implementing offsite disaster recovery is to use the duplicate feature to make a second copy for offsite storage.
- ◆ Another approach is to send monthly or weekly automatic full backups to an offsite storage facility. To restore the data, you get the media from offsite storage (a total directory or disk restore with incremental backups requires the last full backup plus all incremental backups).
- ◆ You can also configure an extra set of schedules for the backups to create duplicates for offsite storage.

Regardless of the method you use for offsite storage, ensure that you configure adequate retention periods. You can use the NetBackup import feature to retrieve expired backups but it is easiest just to set an adequate retention period.

Precautions for Assigning Retention Periods

Full backups: Always specify a time period that is longer than the frequency setting for the schedule (where the frequency is how often the backup runs). For example, if the frequency for a full backup is one week, specify a retention period of two to four weeks. This leaves enough margin to ensure that the current full backup does not expire before the next successful full backup occurs.

Cumulative incremental backups: Always specify a time period that is longer than the frequency setting for the schedule. For example, if the frequency setting is one day, then specify a retention period of one week. This leaves enough 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.

Differential incremental backups: Always specify a time period that is longer than the period between full backups. For example, if full backups occur weekly, then save the incrementals for two weeks. A complete restore requires the previous full backup plus all subsequent incrementals.

- ◆ Be certain to assign a retention period that is long enough because NetBackup stops tracking backups when the retention period expires, making it difficult or impossible to recover files.
- ◆ Within a policy, always assign a longer retention period to full backups than to incrementals. Otherwise, it may not be possible to restore all your files.
- ◆ Archive schedules normally use a retention period of infinite.
- ◆ For WORM (write once, read many) optical platters (supported only on UNIX servers), or tape, set the retention to infinite. If infinite is unacceptable because of NetBackup database space limitations, set the retention period to match the length of time that you want to retain the data. For retention periods that are less than infinite, you must delete the WORM media from the Media Manager configuration upon expiration, or Media Manager will reallocate the media for future backups (even though WORM can be written only once).

Changing Retention Periods

Set the default retention periods by selecting **NetBackup Management > Host Properties > Master Server > Double-click on master server > Servers > Retention Periods**. (See “Retention Periods Properties” on page 436.)

The retention periods are indexed to different levels. For example, the default retention period for level 0 is one week. NetBackup also uses the level when determining the volume to use for storing a backup. (See “Mixing Retention Levels on Backup Volumes” on page 117.)

Mixing Retention Levels on Backup Volumes

By default, NetBackup stores each backup on a volume that has existing backups at the same retention level. If a backup has a retention level of 2, NetBackup stores it on a volume with backups at retention level 2. When NetBackup encounters a backup with a different retention level, it switches to an appropriate volume. Because volumes remain assigned to NetBackup until all the backups on them have expired, this approach results in more efficient use of media. One small backup with an infinite retention would prevent a volume from being reused, even if all other backups on the volume have expired.

To mix retention levels on volumes, select **Allow Multiple Retentions per Media** on the Media host properties.

If you keep only one retention level on each volume, do not use any more retention levels than necessary. This consumes resources and also increases the number of volumes required.

Media Multiplexing

The **Media Multiplexing** attribute specifies the number of jobs from this 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.

For more information on configuring multiplexing and the ramifications of using multiplexing, see Chapter 3 in the *NetBackup System Administrator's Guide, Volume II*.

Note Some policy or schedule types do not support media multiplexing and NetBackup does not allow you to select it in those instances.

Final Destination Storage Unit

If the schedule is a relocation schedule, created as part of configuring a disk staging storage unit, a **Final Destination Storage Unit** must be indicated. A **Final Destination Storage Unit** is the name of the storage unit where the images are swept to from the disk storage unit. (For more information on disk staging storage units, see "About Disk Staging" on page 47.)

Schedule Attributes Tab

Final Destination Volume Pool

If the schedule is a relocation schedule, created as part of configuring a disk staging storage unit, a **Final Destination Volume Pool** must be indicated. A **Final Destination Volume Pool** is the name of the volume pool where images are swept from the volume pool on the disk staging storage unit. (For more information on disk staging storage units, see “About Disk Staging” on page 47.)

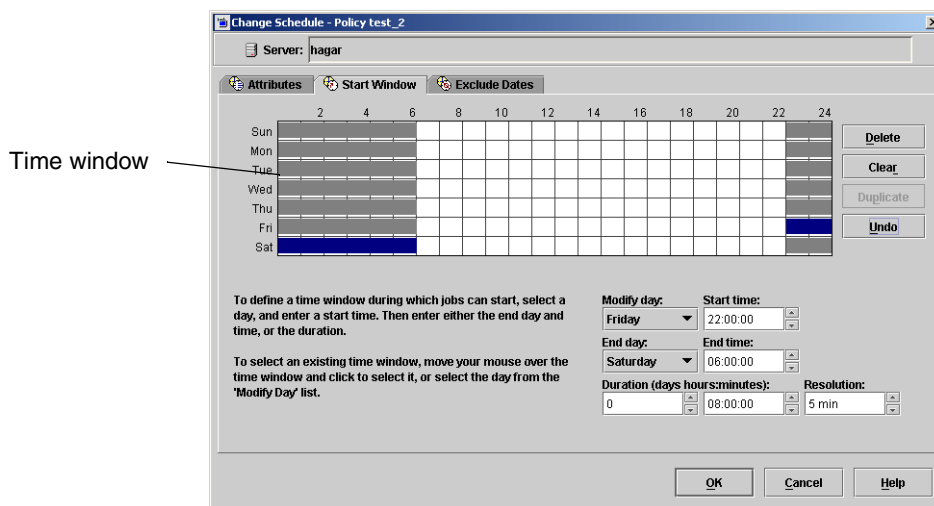
Note The relocation schedule created for the disk staging storage unit is not listed under **Schedules** in the NetBackup Administration Console when **Policies** is selected.

Start Window Tab

The **Start Window** tab provides controls for setting time periods during which NetBackup can start backups, archives, or disk staging relocation when using this schedule. Time periods are referred to as time windows. Configure time windows so that they will satisfy the requirements necessary to complete a task or job. For example, for backups, you can create a different window that opens each day for a specific amount of time, or you can keep the window open all week.

▼ To create a window of time for a schedule

1. Click the **Start Window** tab.
2. To indicate the beginning of the time window during which backups can start:
Click the arrow to the right of **Modify day** and select the first day that the window will be open. Then, click the up and down arrows to the right of **Start time** to select the time the window will open.



3. Indicate how long the time window will remain open by setting a duration time or by choosing an **End day** and **End time**:
 - ◆ To indicate the duration of the time window:
Once you've chosen the opening (or the start) of the window, click the up and down arrows to the right of **Duration (days, hours, minutes)**.
 - ◆ To indicate the close (or the end) of the time window:

Start Window Tab

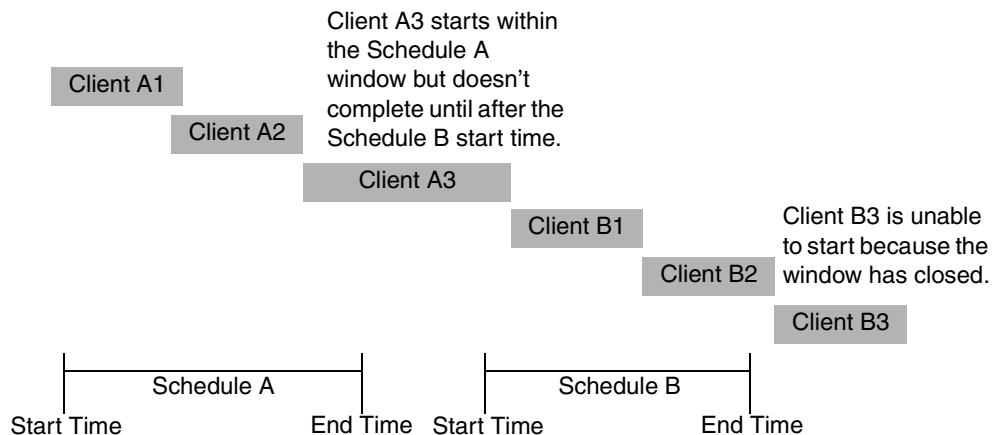
Click the arrow to the right of **End day** and select the last day in the time window. Then, click the up and down arrows to the right of **End time** to select when the time window will end.

Time windows show as bars in the schedule display.

4. If necessary, click a time window to perform actions by the following **Start Window** buttons:
 - ◆ **Delete:** Deletes the selected time window.
 - ◆ **Clear:** Removes all time windows from the schedule display.
 - ◆ **Duplicate:** Replicates the time window for the entire week.
 - ◆ **Undo:** Erases the last action.
5. Click another tab to make additional selections, or click **Add** or **OK** to add the schedule as it is to the Schedule tab.

Duration Example

The figure below represents the effect of schedule duration on two full backup schedules, where the start time for schedule B begins shortly after the end time for previous schedule A. Both schedules have three clients with backups due.



The backup for client A3 in schedule A does not finish until after the schedule B window has opened and does not leave enough time for the schedule B backups. Client B3 must wait until the next time NetBackup runs schedule B.

Client A3 illustrates that, once started, a backup runs to completion even if the window closes while the backup is running.

Exclude Dates Tab

Use the **Exclude Dates** tab to exclude specific dates from a schedule.

The **Exclude Dates** tab displays a 3-month calendar. Use the controls at the top of the calendar to change the month or year. You can exclude specific dates in any month of any year up to and including December 31, 2037.

▼ To exclude a date from the policy schedule

1. Select the **Exclude Dates** tab.
2. Use one of the following methods to indicate a date:
 - ◆ Click the date on the calendar that you wish to exclude. The date appears in the **Exclude Dates** list.
 - ◆ Another method to exclude dates is to click **New**. Then enter the month, day and year in the Date selection dialog. Click **OK**.
3. When you have finished selecting dates, select another tab to make changes or click **OK** to close the dialog.

Calendar Schedule Tab

The Calendar Schedule tab appears when **Calendar** is selected as the Schedule type on the **Attributes** tab of the Schedule dialog. Calendar-based scheduling provides several run day options for use in scheduling when a task will run.

The Calendar Schedule tab displays a 3-month calendar. Use the controls at the top of the calendar to change the month or year.

Schedule by Specific Dates

A task can run on specific dates rather than follow a recurring schedule, and specific dates can be added to a recurring schedule. The **Specific Dates** run day option allows you to schedule specific dates on which your task will run. You can schedule specific dates in any month of any year up to and including December 31, 2037.

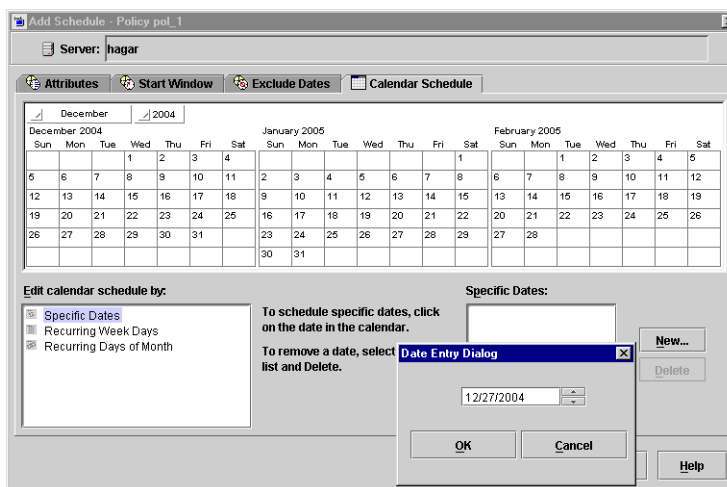
▼ To schedule a task on specific dates

1. In the **Calendar Schedule** tab, select **Specific Dates**.

2. Click on the date in the calendar display or click **New**, enter a date, then click **OK**. The date appears in the calendar schedule list.

3. To remove a date, select it in the calendar schedule list and click **Delete**.

4. When you have finished selecting dates, select another tab to make changes or click **OK** to save and close the dialog.



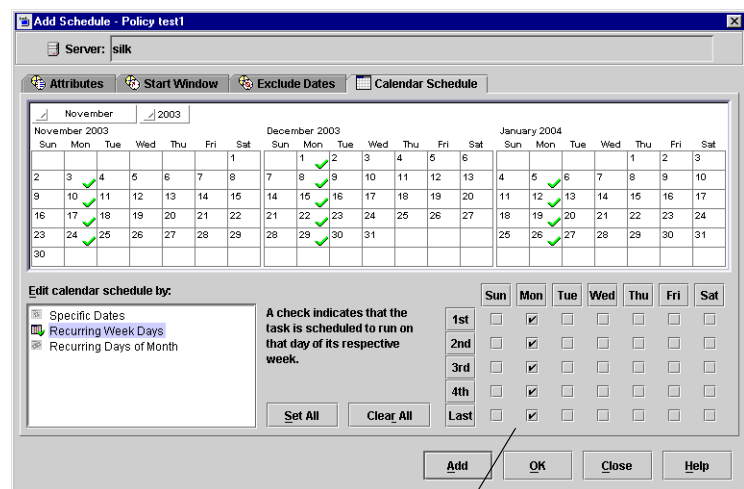
Schedule by Recurring Week Days

The **Recurring Week Days** option provides a matrix that lets you schedule a task for certain days of each week, weeks of each month, or days on particular weeks of the month. For example, use this option to schedule a task on the first and third Thursday of every month. Or, schedule a task that runs the last week in every month.

The week day matrix is not a calendar. It is simply a matrix used to select days and weeks in a month. A check mark entered for a day indicates that the task is scheduled to run on that day of its respective week. By default, no days are selected.

▼ To schedule a recurring weekly task

1. In the **Calendar Schedule** tab, select **Recurring Week Days**.
2. If necessary, select **Clear All** to remove any existing selections from the matrix.
3. Click a check box in the matrix for a particular day to select that day or to clear it.
4. Click the name of the day column header to select or clear the corresponding day for each week of the month.
5. Click a row number to select or clear the entire week.
6. Click the check box for the appropriate day in the **Last** row to schedule a task for the last week of each month, regardless of the number of weeks in the month.
7. When you have finished selecting dates, select another tab to make changes or click **OK** to save and close the dialog.



Schedule by Recurring Days of the Month

The **Recurring Days of the Month** option provides a matrix that you can use to schedule a task for certain days of the month. You can also schedule a task to occur on the last day of the month, regardless of the actual date.

▼ To schedule a recurring monthly task

1. In the **Calendar Schedule** tab, select **Recurring Days of Month**.

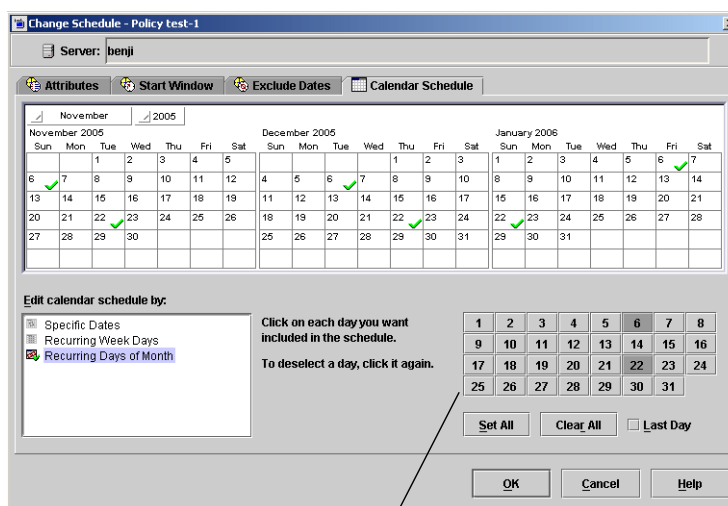
2. To select all calendar dates, click **Set All**.

3. If necessary, select **Clear All** to remove any existing selections from the matrix.

4. Select the button for each day you want included in the run schedule. Clicking the button again will deselect the day.

5. Select the **Last Day** check box if you want to run the schedule on the last day of the month, regardless of the date.

6. When you have finished selecting dates, select another tab to make changes or click **OK** to save and close the dialog.



How Calendar Scheduling Interacts with Daily Windows

Daily windows are taken into account, even when calendar-based scheduling is used. Windows that span midnight, effectively become two separate windows for calendar scheduling. For the first backup after the policy is created, this can sometimes appear as though two backups have run within the same window.

If the calendar schedule indicates that today is a run day, the backup will run once during any window that is open. For example:

1. A new backup policy is created on Monday afternoon. The windows are configured to be open from 6 p.m. until 6 a.m., Sunday through Saturday.
2. In the Calendar Schedule tab, the schedule is set up to run on recurring week days, Monday through Saturday.
3. Since this is a new policy, no backup yet exists based on this policy. And since today (Monday) is a run day, a job will run as soon as the window opens at 6 p.m.
4. At midnight, it is a new day (Tuesday) and there is a window open (until 6 a.m.) so the job is due and will run again. The backups will continue to run soon after midnight from that time forward.

Notice how it is possible for the backup to run just before midnight, then again immediately after midnight. This is valid since both are different run *days* and windows are open at both times (6 a.m. through 6 p.m. every day of the week). Windows that span midnight, effectively become two separate windows for calendar scheduling.

If the desired result is to run jobs at 6 p.m. instead of midnight, use a frequency of one day instead of setting up recurring days in the Calendar Schedule tab.

Examples of Automatic-Backup Schedules

Backups can be scheduled to occur automatically on every day of the week or only on specific days. You can also specify a different backup window for each day.

The days of the week to choose for backups depends on how you want to distribute the backup load. For example, to have all backups occur on Saturday, create a backup window only for Saturday. Leave these values blank for other days.

The best times for automatic backups are usually nights and weekends, when client and network activity is lowest. Otherwise, the backups can adversely affect client and network performance and take longer to complete.

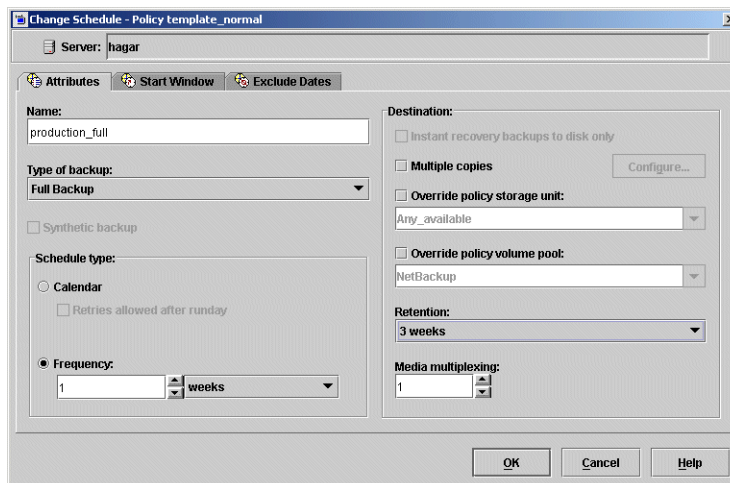
For details on how calendar-based scheduling works with backup windows, see “How Calendar Scheduling Interacts with Daily Windows” on page 125.

Example 1: Various Automatic Backup Schedules

This example shows two approaches for scheduling automatic backups. The first is the recommended method.

Schedule Runs Every Day (recommended method)

The recommended method is to create schedules that run every day of the week.




If the backup for a client does not complete on one day, NetBackup retries it on the next day. This ensures that a retry occurs promptly in case of a failure or lack of time during the first session.

The day of the week when a client is backed up changes if its backup rolls over to the next day.

In this example schedule, full backups can occur on any day of the week but only once every seven days:

If the cycle begins with a full backup on a Monday and completes successfully, the next full backup occurs on the following Monday, seven days later.



If the backup fails on Monday, NetBackup attempts it at the same time each day until it successfully completes. NetBackup can attempt the backup on each subsequent day because the schedule allows backups to occur on any day, but only once during any seven day period. If the backup completes on Tuesday, NetBackup waits seven days from Tuesday for the next backup.

Calendar Schedule Tab

Another Method


This example shows a frequency schedule that allows backups to occur only on specific days. Full backups occur only on Tuesdays and every seven days.

The top screenshot shows the 'Change Schedule - Policy template_normal' dialog box with the following settings:

- Server: hagar
- Name: production_full
- Type of backup: Full Backup
- Schedule type: Frequency (selected)
- Frequency: 1 weeks
- Destination: Any_available
- Retention: 3 weeks
- Media multiplexing: 1


The bottom screenshot shows the 'Calendar' tab with a grid where Tuesday and Wednesday are highlighted. Below the grid, the 'Modify day' is set to Tuesday, 'Start time' is 22:00:00, 'End day' is Wednesday, 'End time' is 08:00:00, 'Duration' is 10:00:00, and 'Resolution' is 5 min.

If the cycle begins with a full backup on a Tuesday and completes successfully, the next full backup occurs on the following Tuesday, seven days later.



Calendar Schedule Tab

If the backup fails on Tuesday, NetBackup must wait until the following Tuesday before trying again.



Calendar Schedule Tab

Example 2: Daily Schedules

The following shows a complete set of frequency schedules that have a backup window available every day. This is the recommended method.

If the backup does not complete on one day, NetBackup tries it again the next day.

Daily Incremental Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Name: production_diff

Type of backup: Differential Incremental Backup

☐ Synthetic backup

Schedule type:
☐ Calendar
☐ Retries allowed after runday

Frequency: 1 days

Destination:
☐ Instant recovery backups to disk only
☐ Multiple copies
☐ Override policy storage unit: Any_available
☐ Override policy volume pool: NetBackup

Retention: 2 weeks

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Thu												
Fri												
Sat												

Delete
Clear
Duplicate
Undo

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Sunday Start time: 22:00:00

End day: Monday End time: 08:00:00

Duration (days:hours:minutes): 0 10:00:00 Resolution: 5 min

Add OK Close Help

Weekly Full Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes Start Window Exclude Dates

Name: production_full

Type of backup: Full Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency:

1 weeks

Destination:

☐ Instant recovery backups to disk only

☐ Multiple copies [Configure...](#)

☐ Override policy storage unit: Any_available

☐ Override policy volume pool: NetBackup

Retention: 6 months

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes Start Window Exclude Dates

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Thu												
Fri												
Sat												

Delete

Clear

Duplicate

Undo

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the "Modify Day" list.

Modify day: Sunday

Start time: 22:00:00

End day: Monday

End time: 08:00:00

Duration (days hours:minutes): 0

Resolution: 5 min

Add OK Close Help

Calendar Schedule Tab

Monthly Full Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes **Start Window** **Exclude Dates**

Name: production_full_monthly

Type of backup: Full Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency: 4 weeks

Destination:

☐ Instant recovery backups to disk only

☐ Multiple copies [Configure...](#)

☐ Override policy storage unit: Any_available

☐ Override policy volume pool: NetBackup

Retention: 3 months

Media multiplexing: 1

[Add](#) [OK](#) [Close](#) [Help](#)

Add Schedule - Policy template_normal

Server: hagar

Attributes **Start Window** **Exclude Dates**

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Fri												
Sat												

[Delete](#) [Clear](#) [Duplicate](#) [Undo](#)

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Sunday **Start time:** 22:00:00

End day: Monday **End time:** 08:00:00

Duration (days hours:minutes): 0 **Resolution:** 5 min

[Add](#) [OK](#) [Close](#) [Help](#)

Quarterly Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes Start Window Exclude Dates

Name: production_full_quarterly

Type of backup: Full Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency:

12 weeks

Destination:

☐ Instant recovery backups to disk only

☐ Multiple copies

☐ Override policy storage unit: Any_available

☐ Override policy volume pool: NetBackup

Retention: 6 months

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes Start Window Exclude Dates

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Thu												
Fri												
Sat												

Delete

Clear

Duplicate

Undo

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the "Modify Day" list.

Modify day: Sunday

Start time: 22:00:00

End day: Monday

End time: 08:00:00

Duration (days:hours:minutes): 0

Resolution: 5 min

Add OK Close Help

Calendar Schedule Tab

Example 3: Using Various Backup Windows

The following is an example of using different backup windows, depending on the day.

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Name: production_diff2

Type of backup: Differential Incremental Backup

☐ Synthetic backup

Schedule type:
☐ Calendar
☐ Retries allowed after runday

☒ Frequency:
 1 days

Destination:
☐ Instant recovery backups to disk only
☐ Multiple copies
☐ Override policy storage unit: Any_available
☐ Override policy volume pool: NetBackup

Retention: 3 weeks

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Calendar grid showing days of the week (Sun-Sat) and hours (2-24). A time window is highlighted for Saturday from 06:00:00 to 10:00:00.

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Friday Start time: 20:00:00

End day: Saturday End time: 06:00:00

Duration (days:hours:minutes): 0 10:00:00 Resolution: 5 min

Delete Clear Duplicate Undo

Add OK Close Help

Example 4: Long Backup Window

The following is an example where the backup window is longer than the period between backups as determined by frequency.

Backups occur according to time elapsed since the last backup and more than one backup can occur for a client during the backup window.

This mode is useful when you want to perform backups twice (or more) daily.

In the following schedule, the backup window spans 7 days and the frequency is 12 hours. A backup is due every 12 hours.

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Name: production_diff3

Type of backup: Differential Incremental Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency: 12 hours

Destination:

☐ Instant recovery backups to disk only

☐ Multiple copies

☐ Override policy storage unit: Any_available

☐ Override policy volume pool: NetBackup

Retention: 3 weeks

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Calendar grid showing a 7-day window (Sun to Sat) highlighted in blue.

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Saturday Start time: 12:00:00

End day: Sunday End time: 12:00:00

Duration (days hours:minutes): 1 00:00:00 Resolution: 5 min

Delete Clear Duplicate Undo

Add OK Close Help

Calendar Schedule Tab

Example 5: Weekend Hours Only

The following example allows full backups to occur only during weekend hours.

The weekend backups are accomplished by having a start time of 8 pm Friday evening and a duration of 60 hours. This allows NetBackup to continue running backups until 8 am Monday morning.

Because the frequency is three days, backups are due again when the schedule starts on the following Friday. If a failure occurs, the administrator can run a manual backup on Monday and the automatic backup is still due on Friday.

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Name: production_full2

Type of backup: Full Backup

☐ Synthetic backup

Schedule type:
☐ Calendar
☐ Retries allowed after runday

Frequency: 3 days

Destination:
☐ Instant recovery backups to disk only
☐ Multiple copies
☐ Override policy storage unit: Any_available
☐ Override policy volume pool: NetBackup

Retention: 3 weeks

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Calendar grid showing days of the week (Sun-Sat) and hours (2-24). Friday and Saturday are highlighted.

Modify day: Friday Start time: 20:00:00

End day: Monday End time: 08:00:00

Duration (days:hours:minutes): 2 12:00:00 Resolution: 5 min

Add OK Close Help

Example 6: Full Backup Every Sunday

The following is an example where a full backup runs every Sunday and cumulative incrementals run on all other days. Each of the cumulative incremental backups contain all files that have changed since the last full backup. This puts more files in each incremental than are present for a differential but it makes restores easier. If a restore is required on Saturday, the Sunday tape and the Saturday tape are needed to do the restore. If this were a differential incremental, then all tapes Sunday through Saturday would be needed.

Full Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Name: production_full3

Type of backup: Full Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency:

7 days

Destination:

☐ Install recovery backups to disk only

☐ Multiple copies

☐ Override policy storage unit: Any_available

☐ Override policy volume pool: NetBackup

Retention: 3 weeks

Media multiplexing: 1

Add OK Close Help

Add Schedule - Policy template_normal

Server: hagar

Attributes | Start Window | Exclude Dates

Calendar grid showing days of the week (Sun-Sat) and hours (2-24). Sunday is highlighted.

Delete Clear Duplicate Undo

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Sunday Start time: 22:00:00

End day: Monday End time: 06:00:00

Duration (days:hours:minutes): 0 08:00:00 Resolution: 5 min

Add OK Close Help

Calendar Schedule Tab

Cumulative Incremental Backups

Add Schedule - Policy template_normal

Server: hagar

Attributes | **Start Window** | **Exclude Dates**

Name: production_cumulative_incremental

Type of backup: Cumulative Incremental Backup

☐ Synthetic backup

Schedule type:

☐ Calendar

☐ Retries allowed after runday

☒ Frequency:

1 days

Destination:

☐ Instant recovery backups to disk only

☐ Multiple copies

☐ Override policy storage unit:

Any_available

☐ Override policy volume pool:

NetBackup

Retention: 3 weeks

Media multiplexing: 1

Add OK Close Help

Change Schedule - Policy test_2

Server: hagar

Attributes | **Start Window** | **Exclude Dates**

	2	4	6	8	10	12	14	16	18	20	22	24
Sun												
Mon												
Tue												
Wed												
Thu												
Fri												
Sat												

Delete

Clear

Duplicate

Undo

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.

To select an existing time window, move your mouse over the time window and click to select it, or select the day from the 'Modify Day' list.

Modify day: Friday

Start time: 22:00:00

End day: Saturday

End time: 06:00:00

Duration (days:hours:minutes): 0

Resolution: 08:00:00

5 min

OK Cancel Help

Considerations for User Schedules

In order for users to perform backups and archives, you must create schedules that allow user backups. There is no requirement, however, to create a policy exclusively for user backups. Restores can be performed at any time and do not use schedules.

Caution 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 manual, references to backups also apply to the backup portion of archive operations unless otherwise noted.

Planning User Backup and Archive Schedules

When planning schedules for user backups and archives, consider the following:

- ◆ The most convenient times for users to perform backups.

If possible, do not permit user backups and archives when automatic backups are occurring. If an automatic backup is running when a user submits a backup or archive, NetBackup queues the user job, unless there is a limiting setting, such as **Limit Jobs per Policy** (a policy attribute) or **Maximum Jobs per Client** (a master server Global Attributes host property). If the automatic backup is long enough, the user job will miss the backup window. Once started, a user job also delays automatic backups and can cause them to miss the backup window.

- ◆ A storage unit. Using a different storage unit can eliminate conflicts with automatic backups.
- ◆ A volume pool. Use a different volume pool if you want to manage the media separate from the automatic backup media.

Caution If the retention period is not long enough and the retention period expires, it can be difficult or impossible to restore the archives or backups.

- ◆ Retention. It is usually best to set the retention period for archives to infinite, since the disk copy of the files is deleted.

Creating Separate Policies for User Schedules

If you create separate policies for user backups or archives (although there is no requirement to do so), the considerations are similar to those for automatic backups. One difference, however, is that no backup selection list is necessary because users select the objects before starting the operation.

Calendar Schedule Tab

The following table shows a set of clients in two user policies.

Policy	Client	Desired Storage	Best Backup Time	Retention
<i>User1</i>	<i>mercury</i>	8mm tape stacker	08:00 to 16:00	Backups - 6 months
	<i>mars</i>			Archives - Infinite
	<i>jupiter</i>			
	<i>neptune</i>			
<i>User2</i>	<i>pluto</i>	8mm tape stacker	12:00 to 20:00	Backups - 6 months Archives - Infinite

- ◆ All clients in policy *User1* have common requirements for user backups and archives.
- ◆ The policy named *User2* was created for *pluto* because the user on this client works from 12 pm to 8 pm (12:00 to 20:00) and therefore requires different backup times.

If NetBackup receives a request for a user backup or archive, it uses the first policy and schedule that it finds that has both of the following:

1. The client for which the user is requesting the operation.
2. A user schedule that:
 - ◆ Specifies the appropriate operation (backup or archive).
 - ◆ Allows the operation to start at the time that the user requests it. If the backup device is busy at the time of the request, NetBackup queues the request and runs the job when the device becomes available (providing the backup window is still open).

For example, assume that at 14:00 (2 pm), a user on the client named *mars* begins a backup of files. NetBackup processes this request as follows:

1. Finds a policy that includes *mars* in its client list and has a user backup schedule that allows a backup to start at 14:00 (2 pm).
2. Performs the backup.

The following policy and schedule meets the criteria for the above request:

Clients	<i>mercury, mars, jupiter, neptune</i>
Files	Applies only to automatic backups

Type of Backup	User backup
Start Time	08:00
Duration	10 hours
Days of Week	All
Retention	6 months
Storage Unit	TS8_1

Using a Specific Policy and User Schedule

To use a specific policy and/or schedule for user backups or archives, perform the following on the client:

- ◆ On Microsoft Windows clients, start the Backup, Archive and Restore client interface. Click **File > NetBackup Client Properties** and select the **Backups** tab. Specify the backup policy and backup schedule.
- ◆ On NetWare target clients, specify the policy and schedule with `backup_policy` and `backup_sched` entries in the `bp.ini` file (see the NetBackup user's guide for the client).
- ◆ On UNIX clients, specify the policy and schedule with `BPARCHIVE_POLICY`, `BPARCHIVE_SCHED`, `BPBACKUP_POLICY`, or `BPBACKUP_SCHED` options in the `bp.conf` file.

Example Policies

Example Policies

The following figures show the clients, backup selection list, and schedules for two example backup policies.

Example 1 specifies that files in `/usr` and `/home` be backed up for the clients *mars*, *jupiter*, and *neptune*. This policy has daily, and weekly automatic schedules and a user backup schedule. All backups go to 8mm tape.

Example Backup Policy 1

Client List	Backup selection list	Schedules		
<i>mars</i> <i>jupiter</i> <i>neptune</i>	<code>/usr</code> <code>/home</code>	Daily Incrementals Run every day between 6 pm and 6 am. Store on 8mm tape. Keep 14 days.	Weekly Fulls Run Mondays every week between 6 pm and 6 am. Store on 8mm tape. Keep one month.	User Backups User can run any day between 8 am and 5 pm. Store on 8mm tape. Keep one year.

Example 2 has different scheduling requirements. For example, this policy has monthly fulls that go to DLT tape.

Example Backup Policy 2

Client List	Backup selection list	Schedules		
<i>pluto</i> <i>mercury</i>	<code>/usr</code> <code>/home</code>	Daily Incrementals Run every day between 6 pm and 6 am. Store on 8mm tape. Keep 14 days.	Weekly Fulls Run Tuesdays every week between 6 pm and 6 am. Store on 8mm tape. Keep one month.	Monthly Fulls Run Sundays every month between 6 pm and 6 am. Store on DLT tape. Keep one year.

Policy Planning Guidelines for Backups

Policies allow you to meet the needs of a wide variety of clients in a single NetBackup configuration. However, taking full advantage of policies for use in backups requires careful planning before starting your configuration. The following procedure provides planning guidelines.

1. Divide clients into groups according to the types of work they perform.

Clients used for similar tasks usually have a high level of commonality in their backup requirements. For example, most clients in an engineering department create the same types of files at similar levels of importance.

In some instances, you can create a single policy for each group of clients. In other cases, you will have to further subdivide the clients and cover them in separate policies, based on their backup requirements as explained later in this procedure.

The table below is the initial grouping for our example. We assume these clients are in the same work group and the initial plan is to cover them all in the same backup policy.

Clients

mercury
mars
jupiter
neptune

2. Gather information about each client. Include information relevant to the backups such as the names, size, and number of files.

In our example client list, *mercury* is a file server and has a large amount of data. To avoid excessively long backup times, we include *mercury* in a separate policy called S1 and the workstations in a policy called WS1. Later, we may find that we need more than one policy for *mercury*, but we will evaluate other factors first. For now, the backup policies are as follows:

Policy	Clients
S1	<i>mercury</i> (file server)
WS1	<i>mars</i> <i>jupiter</i> (workstations) <i>neptune</i>

3. Create backup policies to accommodate special storage requirements.

The storage unit and volume pool settings apply to all files that are backed up by the policy. If files have special storage unit and volume pool requirements, create separate policies for them, even if other factors, such as schedules, are the same.

Policy Planning Guidelines for Backups

In the example below, we create a separate policy (S2) for `/h002/devexp` and `/h002/desdoc` on *mercury* because those files go on DLT tape. Other files on *mercury* go on 8mm tape. If it is necessary to keep backups for some files on separate media, create a policy that specifies a unique volume pool for those backups. Then, add the media for that volume pool.

Policy	Clients	Files	Desired Storage
S1	<i>mercury</i>	/ /usr /h001 /h002/projects	8mm
S2	<i>mercury</i> <i>mercury</i>	/h002/devexp /h002/desdoc	DLT

4. Create additional backup policies if one set of schedules does not accommodate all clients and files. Factors to consider are:

- ◆ Best times for backups to occur. To back up different clients on different schedules, create more policies. For example, create different policies for night-shift and day-shift clients. In our example, we can back them all up during the same hours so additional policies are not necessary.
- ◆ How frequently the files change. For example, if some files change very infrequently in comparison to other files, back them up on a different schedule. To do this, create another policy that has an appropriate schedule and then include the files and clients in that policy.

In our example (see the next table), we place the root (`/`) file system on *mercury* in a different policy (S3). The root (`/`) file system on the workstations is also in a separate policy (WS2).

- ◆ How long backups have to be retained. Each schedule has a retention setting that determines how long NetBackup keeps files that are backed up by the schedule. Because the schedule backs up all the files in the backup selection list, it is best if all files have similar retention requirements. Do not, for example, include files whose full backups must be retained forever, in a policy where full backups are retained for only four weeks.

In our example (see the next table), we place `/h002/desdoc` on *mercury* in a different policy (S4). This is done because `/h002/desdoc` requires full backups every 12 weeks and those backups must be retained for a much longer time than the other files on *mercury*.

Policy Planning Guidelines for Backups

Policy	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
S1	<i>mercury</i>	/usr /h001 /h002/projects	high	8mm	Daily Incr Weekly Full 4 Weeks Full
S2	<i>mercury</i>	/h002/devexp	high	DLT	Daily Incr Weekly Full 4 Weeks Full
S3	<i>mercury</i>	/	low	8mm	Daily Incr 4 Weeks Full
S4	<i>mercury</i>	/h002/desdoc	high	DLT	Daily Incr Weekly Full 4 Weeks Full 12 Weeks Full
WS1	<i>mars</i> <i>jupiter</i> <i>neptune</i>	/usr /people /usr /home /usr /people /var	high	8mm	Daily Incr Weekly Full 4 Weeks Full
WS2	<i>mars</i> <i>jupiter</i> <i>neptune</i>	/ / /	low	8mm	Daily Incr 4 Weeks Full

Policy Planning Guidelines for Backups

5. Create separate policies for clients that require different general attribute settings than other clients. Some attribute settings to consider are:
 - ◆ **Policy Type.** There are several types of backup policies and you must use the correct one for the client. For example, include Windows XP and Windows 2000 clients in a MS-Windows-NT policy.
 - ◆ **Follow NFS.** Select this attribute if a UNIX client has NFS mounted files and you are going to back them up from that client. It is also a good idea to use a separate policy for these clients so problems with NFS do not affect the other clients.
 - ◆ **Cross Mount Points.** Select this attribute if you want NetBackup to cross mount points when backing up the files for UNIX or Windows clients covered by this policy. In some instances, you will not want to cross mount points because it will result in backing up too many files—the UNIX root file system is an example of this.
 - ◆ **Backup Network Drives.** Select this attribute to back up files that the client stores on network drives (applies only to MS-Windows-NT policies).
 - ◆ **Compression.** Set this attribute if you want a client to compress its backups before sending them to the server. Note that the time to compress can increase backup time and make it unsuitable to use for all clients.
 - ◆ **Policy Priority.** Use this attribute to control the order in which NetBackup starts its backups. The client in the higher priority policy is backed up first.

In our example, no extra policies are required because of general attribute settings.

6. Create separate policies as necessary to maximize the benefits of multiplexing.
Using multiplexing for slower clients that produce small backups is a strategy for maximizing drive utilization. However, higher-performance clients that produce long backups are likely to fully utilize drives and not benefit from multiplexing.
7. Evaluate total backup times for each schedule and further subdivide your policies to reduce backup times to an acceptable level.

In our example, backing up `/usr`, `/h001`, and `/h002/projects` on *mercury* takes too much time, so a new policy is created for `/h002/projects`. This new policy (S5) has the same requirements as S1 but it is now possible to back up `/h002/projects` separately thus reducing backup time. The next table shows the final set of backup policies.

In addition to reducing the backup time for each policy, backing up the files with separate policies can reduce the total backup time for the server *mercury*. NetBackup processes files within a backup selection list serially and in the order they appear in the backup selection list. However, separate policies are processed in parallel if

enough drives are available and the maximum jobs attributes are set to allow it. (See “Setting the Number of Streams That Can Run Concurrently” on page 94 for an explanation of maximum jobs settings that also applies to this discussion.)

Multiplexing and Allow Multiple Data Streams also allow processing backup policies in parallel. (See “Using Multiple NetBackup Servers” in Chapter 3 of the *NetBackup System Administrator’s Guide, Volume II* and “Allow Multiple Data Streams” on page 92.)

Policy	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
S1	mercury	/usr /h001	high	8mm	Daily Incremental Cumulative Incremental 4 Weeks Full
S2	mercury	/h002/devexp	high	DLT	Daily Incremental Cumulative Incremental 4 Weeks Full
S3	mercury	/	low	8mm	Daily Incremental 4 Weeks Full
S4	mercury	/h002/desdoc	high	DLT	Daily Incremental Weekly Full 4 Weeks Full Quarterly Full
S5	mercury	/h002/projects	high	8mm	Daily Incremental Weekly Full 4 Weeks Full
WS1	mars	/usr /home	high	8mm	Daily Incremental Weekly Full 4 Weeks Full
	jupiter	/usr /home			
	neptune	/usr /home /var			

Policy Planning Guidelines for Backups

Policy	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
WS2	<i>mars</i>	/	low	8mm	Daily Incremental
	<i>jupiter</i>	/			4 Weeks Full
	<i>neptune</i>	/			



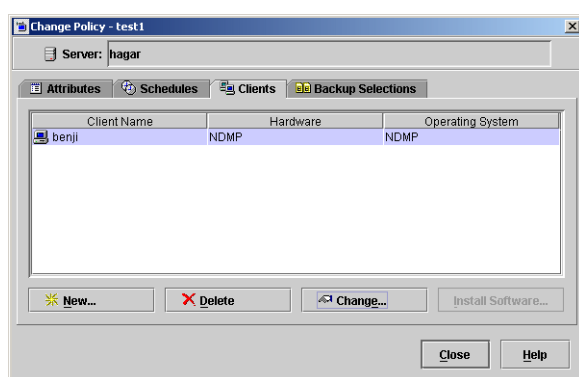
Which Clients Will Be Backed Up: Clients Tab

The Clients tab contains the list of clients that will be backed up or affected by this policy. NetBackup software can be installed on UNIX client machines from the Clients tab.

Note The Clients tab does not appear for Vault policy types.

▼ To add a client to a policy

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Select the Clients tab and click **New**. The Add Client dialog appears.
3. In the **Client Name** field, type the name of the client you are adding.



Observe the following rules for assigning client names:

- ◆ The name must be one by which the server knows the client.
 - ◆ If the client is in multiple policies, use the same name in each policy.
 - ◆ 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 network configuration has multiple domains, use a more qualified name. For example, use *mars.bdev.null.com* or *mars.bdev* rather than just *mars*.
4. Click the **Hardware and operating system** list box, then select the desired entry in the list.

Add only clients with hardware and operating systems that this policy supports. For example, do not add a Novell NetWare client to an MS-Windows-NT policy. If you add the same client to more than one policy, be sure to designate the same hardware and operating system in each of the policies.

Note If the desired hardware and operating system is not in the list, it means that the associated client software is not installed on the server. Check the `/usr/opensv/netbackup/client` directory for the directories and software corresponding to the client you are trying to install. If the directories or software are


Which Clients Will Be Backed Up: Clients Tab

not there, rerun the installation script on the server and choose the option to install client software. (See the NetBackup installation guide that came with your software.)

5. If this is the last client, click **OK**. If you're adding more clients, click **Add**. Click **Close** to cancel changes that you have not yet added and close the Add Client dialog.

▼ To change a client list entry

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies > Summary of all Policies**.
2. In the Details pane, under Clients, double-click the client you wish to change. Or, select multiple clients, then select **Edit > Change**. The Change Client dialog appears.



3. In the **Client Name** field, type the name of the client.

Observe the following rules for assigning client names:

- ◆ If you place the client in multiple policies, use the same name in each policy.
 - ◆ 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 network configuration has multiple domains, use a more qualified name. For example, use `mars.bdev.null.com` or `mars.bdev` rather than just `mars`.
4. Choose the hardware and operating system using the drop-down menu.
Add only clients with hardware and operating systems that the policy will support. For example, do not add a Novell NetWare client to an MS-Windows-NT policy.
 5. Click **OK** to save the change or **Cancel** to discard it.

Installing Client Software on Trusting UNIX Clients

You can install client software on trusting UNIX clients through the NetBackup Administration Console on a UNIX server. Prerequisites are as follows:

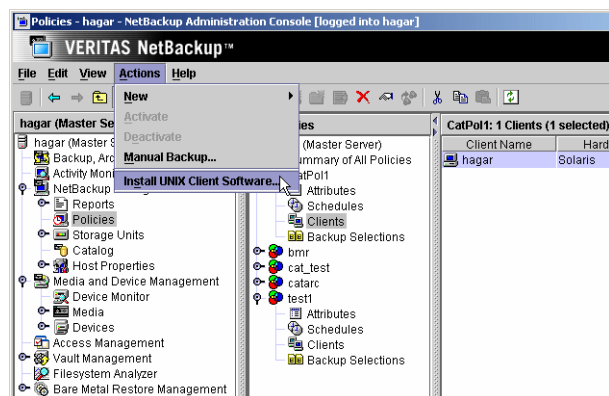
- ◆ You can install the client software only from a UNIX NetBackup server and this server must be the one that you specified in the login dialog when starting the interface. This server must also be the master where you are currently managing backup policies and clients must be in a policy on this master.

For example, assume you want to install clients that are in a policy on a master server named shark. Here, you must have specified shark in the login dialog and therefore be managing NetBackup through the NetBackup-Java Administration Console's application server on this system. shark must also be the master server you are currently managing when you perform the install. In this instance, to install clients for a UNIX master server named tiger you must exit the NetBackup Java interface and restart it, this time specifying tiger in the login dialog.

- ◆ Each client to be installed must have 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 and you can remove them after installing the client software.

▼ To install UNIX client software

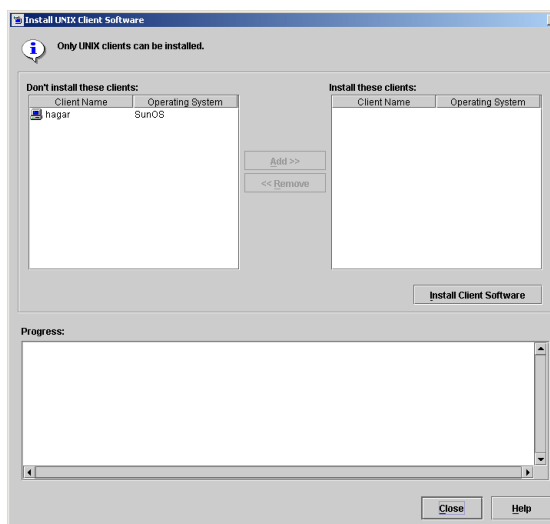
1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**. If you want to install client software, you cannot use the **File > Change Server** command to get to another master server. The master server must be the server that you specified in the login dialog.
2. Select the master server name at the top of the All Policies middle pane.
3. Click **Actions > Install UNIX Client Software**. The Install UNIX Client Software dialog appears.



Which Clients Will Be Backed Up: Clients Tab

4. In the **Don't install these clients** box, select the clients you want to install and click the right arrows. The clients are moved to the **Install these clients** field.
5. Click the **Install Client Software** button to start the installation.

Client software installation can take a minute or more per client. NetBackup writes 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 has started.



During installation, NetBackup does the following:

- ◆ Copies the client software from the `/usr/opensv/netbackup/client` directory on the server to the `/usr/opensv/netbackup` directory on the client.
- ◆ Adds the required entries to the client's `/etc/services` and `inetd.conf` files.

The only way to install client software to a different location on the client is to create the directory where you want the software to reside, then create `/usr/opensv/netbackup` as a link to that directory prior to installing software.

6. When the install is complete, click **Close**.

Installing Software on Secure UNIX Clients

As defined here, a *secure* UNIX client is one that does not contain an entry for the NetBackup master server in its `.rhosts` file. You can install software on clients by using a script or locally on the client from the CD-ROM. For instructions, see the *NetBackup Installation Guide for UNIX*.

Installing Software on Windows Clients

You install NetBackup Windows client software by using the same CD-ROM that contains the server software. For instructions, see the *NetBackup Installation Guide for Windows*.

Configuring a Snapshot Method

The options to configure a snapshot backup method are available only when the Advanced Client option is licensed on a UNIX or Windows server.

For information on configuring snapshots, see the *NetBackup Advanced Client System Administrator's Guide*.

Which Selections Will Be Backed Up: Backup Selections Tab

The backup selections list names the files, directories, directives, scripts, and templates that NetBackup includes in automatic backups of clients covered by the policy. NetBackup uses the same backup selection list for all clients backed up according to the policy.

Selection list entries are processed serially for each client and in the order that they appear in the backup selections, but it is possible to back up multiple clients in parallel if enough drives are available and NetBackup attributes are set to allow it. (See “Setting the Number of Streams That Can Run Concurrently” on page 94.)

All the files listed do not need to exist on all the clients, as NetBackup backs up the files that it finds. However, each client must have at least one of the files in the backup selections list or the client backup will fail with a status 71. The policy backup selections list does not apply to user backups or archives since users select the objects to back up before starting the operation.

Overview on Creating Lists for Different Policy Types

A backup selection list may contain different information, depending on the policy type. For example:


- ◆ Standard, Exchange and Lotus Notes policy types list pathnames and directives. See “Backup Selections List for Standard Policies” on page 154.
- ◆ Depending on the database type, the backup selection list for database policies contains different types of objects. See “Backup Selections List for Database Policies” on page 156.
 - ◆ For Exchange and Lotus Notes, the list contains pathnames and directives.
 - ◆ For MS-SQL-Server, Informix-On-BAR, SAP, and Sybase, the list contains scripts that define and control the database backup, including how the client uses multiple streams.
 - ◆ For Oracle and DB2, the list contains scripts and/or templates.
- ◆ Vault policy types list vault commands. See “Creating a Vault Policy” on page 195.

Backup Selections List for Standard Policies

Standard, Exchange and Lotus Notes policy types list pathnames and directives in the backup selection list.

▼ **To add or change backup selections for a Standard, Exchange, or Lotus Notes policy**

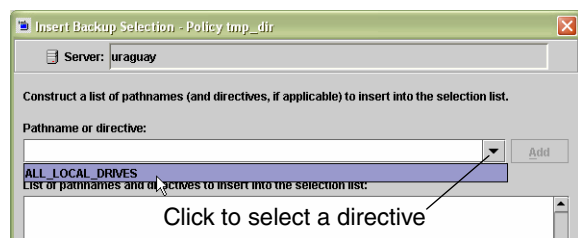
1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Double-click the policy where you wish to change the backup selections list. The Change Policy dialog appears.



3. Click the **Backup Selections** tab.
4. To add an entry, click **New**. The Add Backup Selections dialog appears.
5. Select a pathname or directive:
 - ◆ Type the name of the path in the **Pathname or Directive** field.

For file path information, see “Pathname Rules for UNIX Clients” on page 168 and “Pathname Rules for Microsoft Windows Clients” on page 162

- ◆ Click the drop-down arrow and select a directive in the **Pathname or Directive** field. Click **Add** to include the pathname or directive to the list.



For information on what directives accomplish, see “Backup Selections List Directives: General Discussion” on page 181 and “Backup Selections List Directives for Multiple Data Streams” on page 184 (if the **Allow Multiple Data Streams** general policy attribute is enabled). For separately-priced options, also see the NetBackup guide that came with the option.

Note Pathnames may contain up to 1023 characters.

6. Rearranging the selections in the selection list:
 - ◆ Click **Insert** to add an entry above the one currently selected.
 - ◆ To delete an entry, select the entry and click **Delete**.
 - ◆ To rename an entry, select it and click **Change**. The Change Backup Selection dialog appears. Make your changes and press **OK**.

Which Selections Will Be Backed Up: Backup Selections Tab

7. To verify that the entries on the selections list are accurate, see “Verifying the Backup Selections List” on page 160.

Backup Selections List for Database Policies

For Exchange and Lotus Notes, the backup selections list contains pathnames and directives. For MS-SQL-Server, Informix-On-BAR, SAP, and Sybase, the list contains scripts that define and control the database backup, including how the client uses multiple streams. For Oracle and DB2, the list contains scripts and/or templates.

▼ To create or change backup selections containing scripts for a database policy

1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Double-click the database policy in the Console tree where you wish to change the backup selections lists. The Change Policy dialog appears.
3. Click the **Backup Selections** tab. To add an entry, click **New**. The Add Backup Selections dialog appears.
4. Enter a script into the text box, then click **Add** to add the script to the selection list. Shell scripts require that the full pathname be specified. Be sure that the shell scripts listed are installed on each client specified on the Client tab.
5. Click **OK** to add the items to the **Backup Selections** list.

▼ To add templates or scripts to the Backup Selections List

1. In the NetBackup Administration window, expand **NetBackup Management > Policies**.
2. Double-click the policy in the Console tree where you wish to add or change templates or scripts. The Change Policy dialog appears.
3. Click the **Selections** tab.
4. To add an entry, click **New**. To insert an entry within the current list, select an item and click **Insert**. The Add Backup Selection dialog appears.
5. Specify the backup selections:
 - ◆ **Templates:**

For Oracle policies: From the **Template set** list, choose a template set by operation.

For both Oracle and DB2 policies:

Choose the correct template from the drop-down **Script or template** list, or type the name of a template.

Since templates are stored in a known location on the master server, they do not need to be installed on each client in the Clients list. Enter only the template filename, without a path. For example:

weekly_full_backup.tpl

◆ Shell scripts:

Specify the full pathname when listing scripts, and be sure that the scripts listed are installed on each client in the Client list.

Specifying an Oracle script example:

`install_path/netbackup/ext/db_ext/oracle/samples/rman/cold_d
atabase_backup.sh`

Specifying a DB2 script example:


`/myscripts/db2_backup.sh`

6. To change the order of the backup selections, select one and click **Up** or **Down**.

7. Click **OK** to add the selection to the selection list.

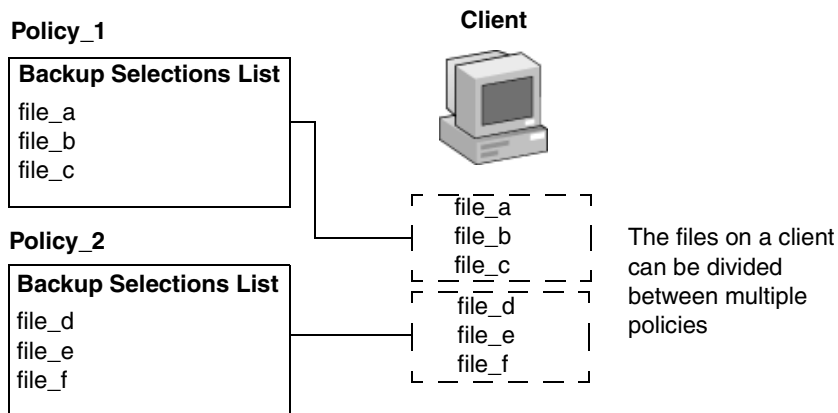
Selection list entries are processed serially for each client and in the order that they appear in the backup selections, but it is possible to back up multiple clients in parallel if enough drives are available and NetBackup attributes are set to allow it. (See “Setting the Number of Streams That Can Run Concurrently” on page 94.)

- ◆ The Global host property, **Maximum Jobs per Client**, and the **Limit Jobs per Policy** policy attribute are set to allow it.
- ◆ Multiple storage devices are available (or you are using multiplexing).



Which Selections Will Be Backed Up: Backup Selections Tab

It is also possible to add a client to multiple policies, then divide the client's files among the backup selections list. This method has the advantage of backing up different files on a client according to different rules. A different schedule can be applied to each policy.



Using multiple policies can also reduce the backup time. When all of a client's files are in the same backup selections list, because NetBackup processes the files serially, the backup can take a long time when there are many files. If the files are divided between different policies, NetBackup can process the policies in parallel, reducing the backup time. The maximum jobs attributes must be set to allow the parallel backups and sufficient system resources must also be available. (See "Setting the Number of Streams That Can Run Concurrently" on page 94 for an explanation of maximum jobs settings that also apply to this discussion.)

Note Understanding disk and controller I/O limitations is important when using multiple policies for a client. For example, if there are two file systems that will 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 way to reduce backup time is to use a single policy in which **Allow Multiple Data Streams** is enabled, 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
```

Which Selections Will Be Backed Up: Backup Selections Tab

The example above produces two concurrent data streams. One has `file_a`, `file_b`, and `file_c`. The other has `file_d`, `file_e`, and `file_f`. (See “Allow Multiple Data Streams” on page 92.)

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, because the heads must move back and forth between tracks containing files for the respective streams.

The **Backup Selections** tab contains a list of files and directories that NetBackup includes in automatic backups of clients covered by a policy. You may also enter a backup selections list *directive* that causes NetBackup to perform specific actions when processing the files in the list. For database backups, scripts and templates are used to define and control the specific type of database backup. (See the NetBackup database guide for more information.)

Note If you are setting up the backup selections list for a Vault job, see “To create a Vault policy” on page 195.

Verifying the Backup Selections List

After creating or modifying a backup selections list, complete the following procedure to make sure that the file paths for the specified clients are correct.

▼ To verify a backup selections list

1. Check all entries to ensure you have followed the file path rules for the clients you are backing up. Also, verify the syntax for any directives that are included in the list. See “Rules for Indicating Pathnames in the Backup Selections List” on page 162.
2. For the first set of backups, check the Problems or All Log Entries reports for warning messages (see examples below) and run the `check_coverage` script (located in `/usr/opensv/netbackup/bin/goodies`).

This step can reveal mistakes that result in not backing up files because the files are not found. The status code for a backup does not always indicate this type of error because NetBackup does not require all paths in the backup selections list to be present on all clients. This allows you to have a generic list that multiple clients can share. Requiring all entries to match for a successful backup would result in more policies, unless all clients had identical filesystems.

If a path is not found, NetBackup logs a trivial (TRV) or warning (WRN) message, but can still end the backup with a status code 0 (success). This is desirable because it eliminates error status codes for files that are not expected to be on a client. However, it means you must check the logs or use the `check_coverage` script to ensure that files are not missed due to bad or missing backup selections list entries.

The examples below show the log messages that appear when files are not found. For information on using `check_coverage`, see the comments in the script.

Example 1: Regular Expressions or Wildcards

Assume the backup selections list contains a regular expression such as:

```
/home1[0123456789]
```

Here, NetBackup backs up `/home10` through `/home19` if they are present. If they are not present, the Problems or All Log Entries report shows a message similar to the following:

```
02/02/04 20:02:33 windows freddie from client freddie: TRV - Found no
matching file system for /home1[0123456789]
```

Example 2: Path Not Present on All Clients or Wrong Path Specified

Assume the backup selections list contains a path named `/worklist` that is not present on all clients. Here, NetBackup backs up `/worklist` on the clients where it exists. For other clients, the Problems or All Log Entries report shows a message similar to the following:

```
02/02/04 21:46:56 carrot freddie from client freddie: TRV - cannot
process path /worklist: No such file or directory. Skipping
```

This message would also occur if `/worklist` were not the correct path name. For example, if the directory name is `/worklists` but you typed `/worklist`.

Note If the paths seem correct and the message still appears, ensure there are no trailing spaces in the paths.

Example 3: Symbolic Link

Assume the backup selections list names a symbolic link. NetBackup does not follow symbolic links and provides a message such as the following in the Problems or All Log Entries report:

```
02/02/04 21:46:47 carrot freddie from client freddie: WRN - /src is
only being backed up as a symbolic link
```

Here, you must resolve the symbolic link if you do not intend to back up the symbolic link itself.

Rules for Indicating Pathnames in the Backup Selections List

The following sections discuss rules for specifying pathnames for each type of NetBackup client:

- ◆ “Pathname Rules for UNIX Clients” on page 168.
- ◆ “Pathname Rules for Microsoft Windows Clients” on page 162.
- ◆ “Pathname Rules for NetWare NonTarget Clients” on page 178.
- ◆ “Pathname Rules for NetWare Target Clients” on page 179.
- ◆ “Pathname Rules for Clients Running Extension Products” on page 180.

Pathname Rules for Microsoft Windows Clients

The following sections describe conventions used to specify backups for Windows clients.

File Backups

Microsoft Windows pathname conventions, UNIX pathname conventions, or a combination of the two can be used in the backup selections list.

Using Microsoft Windows Conventions

- ◆ Enter one pathname per line.
- ◆ Start all pathnames with the drive letter followed by a colon (:) and a backslash (\). The drive letter can be either upper or lower case: `c : \`

When explicitly specifying an entire volume in the file list, it is necessary to append a backslash (\) to the entry to ensure that all data is protected on that volume.

Here is an entry that correctly specifies an entire volume for backup:

```
c : \
```

This entry incorrectly specifies an entire volume for backup:

```
c :
```

- ◆ Precede each component in the path with a backslash.

If the last component in the path is a directory, also follow it with a backslash (\). The trailing backslash is not required but serves as a reminder that the pathname is to a directory instead of a file: `c : \users\net1\`

If the last component is a file, include the file extension and omit the backslash from the end of the name: `c:\special\list.txt`

- ◆ Upper and lower case letters in the pathname must match those in the pathname on the client. The only exception is the drive letter, which can be either upper or lower case: `c:\Worklists\Admin\`
- ◆ You can use the same wildcard characters as in Windows pathnames: `* ?`

To back up all files ending with .doc: `c:\Users*.doc`

To back up all files named log01_03, log02_03, and so on:
`c:\system\log??_03`

- ◆ To back up all local drives except for those that use removable media, specify: `: \` or `*: \` or `ALL_LOCAL_DRIVES`

Drives that are not backed up include floppy disks, CD-ROMs and drives that are located on remote systems but mounted on a system through the network.

The following is an example of a backup selection list that uses the Microsoft Windows conventions:

```
c:\
d:\workfiles\
e:\Special\status
c:\tests\*.exe
```

- ◆ By default, NetBackup does not back up the files described in “Files Excluded from Backups by Default” on page 190.
- ◆ Exclude specific files from backups by creating an exclusion list on the client. See “Excluding Files from Automatic Backups” on page 191.

UNIX Conventions are Permitted on Windows

NetBackup permits you to use UNIX conventions in the backup selection list for Windows clients. UNIX conventions are similar to those for Microsoft Windows, except for the following:

- ◆ Start each line with a forward slash (/).
- ◆ Omit the colon (:) after the drive letter.
- ◆ Specify / to back up all local drives except for those that are removable: /

The following example uses the UNIX conventions:

```
/c/
/d/workfiles/
/e/Special/status
/c/tests/*.exe
```

Rules for Indicating Pathnames in the Backup Selections List

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.

When performing a disk-image backup, be sure to select **Full Backup** for the backup type. Any other backup type will not work for backing up a disk-image.

To specify a disk-image backup, add the logical name for the drive to the policy backup selection list. The format in the following example backs up drive C.

```
\\.\c:
```

Disk-images can be included in the same backup selection list with other backups:

```
\\.\c:
d:\workfiles\
e:\Special\status
HKEY_LOCAL_MACHINE:\
```

To restore the backup, the user first chooses **Select for Restore > Restore from Normal Backup**.

When a user lists the backups from which it can choose, the disk image appears as a file with the same name that was specified in the backup selection list. In this example:

```
\\.\c:
```

After selecting the disk image source, the user enters the destination in the following format:

```
\\.\drive:
```

Where *drive* is the location where the partition will be restored. The leading forward slash is important. For more information, see the online help in the Backup, Archive, and Restore client interface.

Notes on Disk-Image Backups

- ◆ NetBackup first attempts to use Windows Open File Backup methods. If that fails, NetBackup locks the logical drive, ensuring that no changes occur during the backup. If there are open files on the logical drive, a disk-image backup is not performed.
- ◆ Before backing up or restoring a disk-image, all applications that use a handle to the partition must be shut down, otherwise the operation will fail. Examples of such applications are Windows Explorer or Norton Antivirus.

Ensure that there are no active COW (Copy On Write) snapshots in progress. If there is an active COW snapshot, the snapshot process itself will have a handle open to the volume.

- ◆ NetBackup does not support raw partition backups on unformatted partitions.

Microsoft Windows Registry Backup

Backup for Disaster Recovery

To ensure successful recovery in case of a disk failure, always back up the entire registry. That is, back up the directory that contains the entire registry. On most Windows systems, this directory is located at:

```
%systemroot%\system32\config
```

Where %systemroot% is the directory where Windows is installed.

For example, if Windows 2000 is installed in the `c:\winnt` directory, then including any of the following paths will accomplish the backup:

- ◆ `c:\winnt\system32\config` (backs up the entire config directory)
- ◆ `c:\` (backs up the entire C drive)
- ◆ `ALL_LOCAL_DRIVES`
- ◆ `System_State:\` (applies to Windows 2000/XP)

For Windows 2003 systems, enter:

- ◆ `Shadow Copy Components:\`

Caution To ensure a successful recovery of the registry in case of disaster, *do not* include individual registry files or HKEY entries in the same backup selection list that is used to back up the entire registry. If you are using a NetBackup exclude list for a client, do not exclude any registry files from your backups.

See the Disaster Recovery chapter in the *NetBackup Troubleshooting Guide for UNIX and Windows* for instructions on restoring the registry in the case of a disk failure.

Back Up Individual HKEYs (do not use for disaster recovery)

As mentioned above, do not include HKEY entries in the same policy backup selection list used to back up the entire registry. However, if you want the ability to restore individual keys within the registry, create a separate policy and then specify the desired HKEYs in the backup selection list for that policy. The following is an example HKEY entry for a policy backup selection list:

```
HKEY_LOCAL_MACHINE:\
```

Remember, you cannot perform a disaster recovery by restoring HKEYs. In addition, backups and restores will be slower than backing up the registry as a whole.

Rules for Indicating Pathnames in the Backup Selections List

Hard Links to Files (NTFS volumes only)

A hard link is a directory entry for a file. Every file can be considered to have at least one hard link. On NTFS volumes, 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). The actual file is indicated by a Volume Serial Number (VSN) and a File Index which is unique on the volume. Collectively, the VSN and File Index are referred to as the file ID.


During a backup, if the backup selection list includes hard-linked files, the data is backed up only once, using the first file name reference found in the directory structure. If a second or subsequent file name reference is found, it is backed up as a link to the name of the first file. This means you get only one backup copy of the data, regardless of whether you include one or multiple hard links. You can include any of the paths that are hard links to the data in order to back up the data.

During a restore, if all of the hard-link references are restored, the hard-linked files still point to the same file ID as the other files to which they are linked. However, if you do not restore all the hard links, you can encounter anomalies as shown in the following examples.


Example 1

Assume there are three hard links named L1, L2, and L3 that are pointing to the same data as shown in the figure below.


1. During a backup of L2 and L3, L2 is encountered first and backed up, then L3 is backed up as a link to L2. The three files are all hard linked to the same data.



2. Next, the original copies of L2 and L3 are backed up to tape, then deleted, leaving only L1 on the disk.



3. During a subsequent restore, you restore L2 and L3. The restored files, however, do not point to the same file ID as L1. Instead, they



Rules for Indicating Pathnames in the Backup Selections List

are assigned a new file ID 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 L1. The duplication occurs because the backup does not associate L2 and L3 with L1.

Example 2

Assume in example 1, that you attempt to restore only L3. Here, NetBackup cannot link L3 to L2 because L2 does not exist. Since the restore can complete only if it can link to L2, L2 is automatically restored by a secondary restore request to the NetBackup server that has the data. If you restore L2 by itself, there is no problem.

Pathname Rules for UNIX Clients

The following sections describe conventions used to specify backups for UNIX clients.

- ◆ Enter one pathname per line. NetBackup supports a maximum path length of 1023 characters on UNIX clients.
- ◆ Start all pathnames with a slash (/).
- ◆ The following meta or wildcard characters are allowed:

```
*
?
[ ]
{ }
```

For example:

```
/home/. [a-zA-Z0-9] *
/etc/*.conf
```

- ◆ To use meta or wildcard characters literally, precede them with a backslash (\). Assume that the brackets in the following pathname are used as literal characters:

```
/home/abc/fun[ny]name
```

Precede the brackets with a backslash:

```
/home/abc/fun\[ny\]name
```

Note A backslash (\) acts as an escape character only if it precedes a meta or wildcard character. NetBackup normally interprets a backslash literally and it is a legal character to use in pathnames.

- ◆ If there are one or more trailing spaces in a backup selection list entry and a matching entry is not found on the client, NetBackup deletes trailing spaces and checks again. If a match is still not found, NetBackup skips the entry and logs a message similar to one of the following in the NetBackup All Log Entries or Problems report:

```
TRV - cannot process path pathname: No such file or directory.
Skipping
TRV - Found no matching file system for pathname
```

Notes on UNIX Pathnames

- ◆ Pathnames that cross mount points or that the client mounts through NFS can affect the backup configuration. Before creating a backup selection list, familiarize yourself with the **Follow NFS** and **Cross Mount Points** attributes. (See “Follow NFS” on page 81 and “Cross Mount Points” on page 82.)

Rules for Indicating Pathnames in the Backup Selections List

- ◆ NetBackup can back up operating system, kernel, and boot files. NetBackup cannot, however, create bootable tapes. Consult your system documentation to create a bootable tape.
- ◆ By default, NetBackup does not back up the files described in “Files Excluded from Backups by Default” on page 190.
- ◆ Exclude specific files from backups by creating an exclusion list on the client. (See “Excluding Files from Automatic Backups” on page 191.)
- ◆ The **Busy File Settings** host properties for UNIX clients offers alternatives for handling busy and locked files. (See “Busy File Properties” on page 357.)
- ◆ On Hewlett-Packard, AIX, Sequent, and Solaris 2.5 (and later) platforms, NetBackup backs up access control lists (ACLs).
- ◆ NetBackup can back up (and restore) Sun PC NetLink files.
- ◆ On IRIX 6.x and Digital Alpha platforms, NetBackup backs up extended file attributes.
- ◆ On IRIX platforms, NetBackup backs up and restores extended attributes attached to XFS file system objects.
- ◆ On DEC OSF/1 platforms, NetBackup backs up and restores extended attributes attached to files on AdvFS and UFS file systems.
- ◆ By default, NetBackup backs up and restores Solaris 9 extended attribute files. The FlashBackup single file restore program (`sfr`) does not restore extended attribute files. (See “Backup and Restore of Extended Attribute Files and Named Data Streams” on page 174.)
- ◆ By default, NetBackup backs up and restores VxFS 4.0 named data streams. The FlashBackup single file restore program (`sfr`) does not restore extended attribute files. (See “Backup and Restore of Extended Attribute Files and Named Data Streams” on page 174.)
- ◆ On Hewlett-Packard and Solaris 2.5 (and later) platforms, NetBackup backs up VxFS extent attributes.

Symbolic Links to Files or Directories

For symbolic (soft) links, include the pathname to the source file in the backup selections list in order to back up the actual data. If a file is a symbolic link to another file, NetBackup backs up only the link, not the file to which the link points. This prevents multiple backups of the source file.

Because symbolic links are restored only as a symbolic link to the source file, you must restore the source file along with the link in order to get the data.

Rules for Indicating Pathnames in the Backup Selections List

Note If NetBackup restores a symbolic link as root, NetBackup changes the owner and group back to the original owner and group. When NetBackup restores a UNIX symbolic link as a nonroot user, NetBackup sets the owner and group for symbolic links to the owner and group of the person doing the restore. This does not cause problems because when the UNIX system checks permissions, NetBackup uses the owner and group of the file to which the symbolic link points.

Hard Links to Directories

On most UNIX systems, only the root user can create a hard link to a directory. Some systems do not permit hard links and many vendors recommend avoiding using these links.

NetBackup does not back up and restore hard-linked directories in the same manner as files:

- ◆ 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.

Hard Links to Files

A hard link differs from a symbolic link in that it is not a pointer to another file, but is actually two directory entries pointing to the same inode number.


During a backup, if the backup selection list includes hard-linked files, the data is backed up only once, using the first file name reference found in the directory structure. If a second or subsequent file name reference is found, it is backed up as a link to the name of the first file. This means only one backup copy of the data is created, regardless of the number of hard links. Include any of the hard links to the data in order to back up the data.

During a restore, if all of the hard-link references are restored, the hard-linked files remain pointed to the same inode as the other files to which they are linked. However, if you do not restore all the hard links, you can encounter anomalies as shown in the following examples.


Example 1

Assume there are three hard links named L1, L2, and L3 that are pointing to the same data as shown in the figure below.


1. The three files are all hard linked to the same data.



2. L2 and L3 are backed up to tape, then deleted from the disk.



3. When L2 and L3 are restored, the data cannot be associated with the original file and are assigned a new inode number.



1. During a backup of L2 and L3, L2 is encountered first and backed up, then L3 is backed up as a link to L2.
2. Next, the original copies of L2 and L3 are both deleted, leaving only L1 on the disk.
3. During a subsequent restore, you restore L2 and L3. The restored files, however, do not point to the same inode as L1. Instead, they are assigned a new 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 L1. The inode duplication occurs because the backup does not associate L2 and L3 with L1.

Rules for Indicating Pathnames in the Backup Selections List

Example 2

Assume in example 1, that you attempt to restore only L3. NetBackup cannot link L3 to L2 because L2 does not exist. The restore fails, leaving an error message in the progress log. If you restore L2 by itself, there is no problem.

UNIX Raw Partitions

Caution Save a copy of the partition table before performing raw partition backups so you have the copy for reference prior to a restore. To restore the raw partition, a device file must exist and the partition must be the same size as when it was backed up. Otherwise, the results of the restore are unpredictable.

Notes On UNIX Raw Partition Backups

- ◆ Use raw partition backups only if you can ensure that the files have not changed in any way during the backup or, in the case of a database, if you can restore the database to a consistent state by using transaction log files.
- ◆ Do not perform archives of raw partitions on any client. An archive backs up the raw partition, then deletes the device file associated with the raw partition. The file system does not recover the space used by the raw partition.
- ◆ Before backing up file systems as raw partitions, unmount the file system to allow buffered changes to be written to the disk and to prevent the possibility of the file system changing during the backup. You can use the `bpstart_notify` and the `bpend_notify` scripts to unmount and remount the backed-up file systems.
- ◆ The **Cross Mount Points** policy attribute has no effect on raw partitions. If the root partition is being backed up as a raw partition and has mount points for other file systems, the other file systems are not backed up, even if you select **Cross Mount Points**.

The same is true for the **Follow NFS** policy attribute. NFS file systems mounted in a raw partition are not backed up. Nor can you back up raw partitions from other machines by using NFS mounts to access the raw partitions. The devices are not accessible on other machines through NFS.

- ◆ For disks managed by disk volume managers such as VERITAS Volume Manager (VxVm), specify the logical partition names.
- ◆ For clients in a FlashBackup policy, refer to the *NetBackup Advanced Client System Administrator's Guide* (backup selection list and cache section) for the differences between Standard and FlashBackup policies.

When to Use Raw Partition Backups

If there are no file systems to back up and the disks are used in raw mode (such as with some databases), back up the disk partitions as raw partitions. When backing up databases as raw partitions, you can use the `bpstart_notify` and `bpend_notify` scripts to do the preprocessing and postprocessing necessary to back up the databases.

You can also perform a raw partition backup of a disk partition used for file systems. A disadvantage of this method is that you must restore the entire partition to recover a single file (unless you are using FlashBackup). To avoid overwriting the entire partition, use the redirected restore feature to restore the raw partition to another raw partition of the same size, and then copy individual files to the original file system.

Raw partition backups are also useful for backing up entire disks. Since the overhead of the file system is bypassed, a raw partition backup is usually faster. The size of the raw partition backup will be the size of the entire disk, regardless of whether the entire disk is used.

Specifying UNIX Raw Partitions in the Backup Selection List

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:

```
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

Caution Do not specify wildcards (such as `/dev/rsd*`) in pathnames for raw partition backups. Doing so can prevent the successful restore of entire devices, if there is overlap between the memory partitions for different device files.

You can include raw partitions in the same backup selection list as other backups. For example:

```
/home  
/usr  
/etc  
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

Note NetBackup does not distinguish between full and incremental backups when backing 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; although, character special device files are often faster because character devices avoid the use of the buffer cache for accessed disk data. To obtain the optimum backup speed for raw partition backups, test both a block and character special device file to ensure the best choice for your platform.

Rules for Indicating Pathnames in the Backup Selections List

Ensure that you are specifying the actual block- 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.

Selecting a Schedule Backup Type for a UNIX Raw Partition

When performing a raw partition backup, be sure to select **Full Backup** for the Type of Backup from the Schedules tab. Any other backup type will not work for backing up raw partitions. (See “Type of Backup” on page 98.)

Backup and Restore of Extended Attribute Files and Named Data Streams

NetBackup can back up and restore the following file attributes:

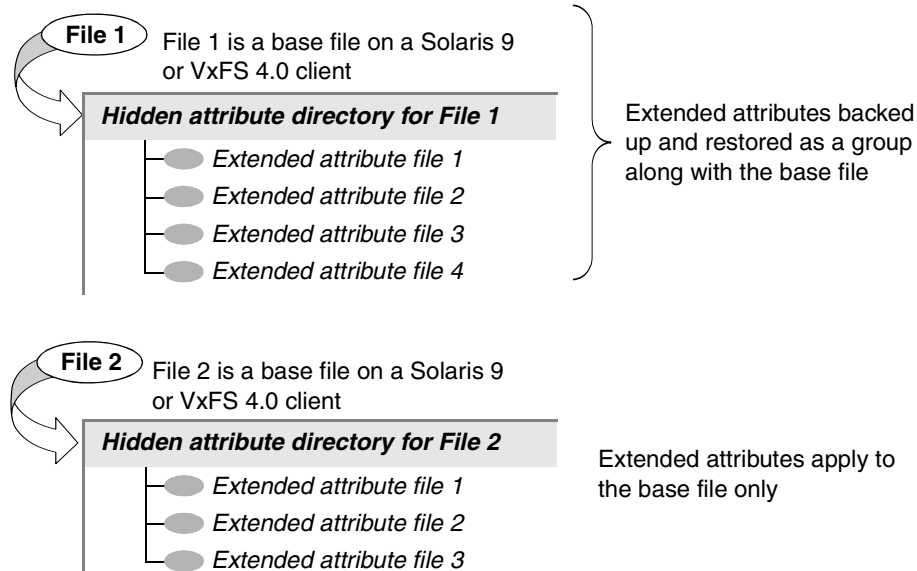
- ◆ Extended attribute files of the Solaris 9 UNIX File System (UFS) and temporary file system (TMPFS)
- ◆ Named data streams of the VxFS 4.0 file system

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 file. The hidden directory is stored within the file system, but can only be accessed via the base file to which it is related. To view which files have extended attributes on Solaris 9 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.

Example of Base File and Extended Attribute Directory and Files



NetBackup Client, Media Server, and Master Server Versions

For backing up and restoring named data streams and Solaris 9 extended attributes:

◆ A NetBackup client:

- ◆ Named data streams can be *restored* to VxFS 4.0 clients only.
- ◆ Extended attributes can be *restored* to Solaris 9 clients only.

A client must be at NetBackup version 5.0 or later in order to *back up* and *restore* VxFS 4.0 named data streams and Solaris 9 extended attributes.

◆ A NetBackup media server:

Restores: Only NetBackup media servers at 5.0 or later can *restore* VxFS 4.0 named data streams and Solaris 9 extended attributes.

Backups: A NetBackup media server of any version can successfully back up named data streams and Solaris 9 extended attributes.

◆ A NetBackup master server:

A NetBackup master server of any version can back up and restore named data streams and Solaris 9 extended attributes.

Rules for Indicating Pathnames in the Backup Selections List

Ramifications of Backing Up Extended Attributes or Named Data Streams

Be aware that the presence of a large number of extended attribute files or named data streams may cause some degradation in backup and restore speed since the base file and all associated files are backed up.


This is especially true in the case of incremental backups, during which NetBackup checks the mtime or ctime of each file individually.

Restoring Extended Attributes or Named Data Streams


If **Overwrite Existing Files** is selected as a restore option in the Backup, Archive, and Restore client interface, and a file possessing extended attributes or named data streams is being restored, any existing attribute files or named data streams for that base file are replaced with the restored files.

In the following example, the user is restoring File 1:


Base File 1 currently possesses four extended attribute files.



The user restores File 1 from a backup that was created when File 1 possessed only three extended attribute files.



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.



If an attempt is made to restore:

- ◆ the extended attribute files to any non-Solaris 9 client, or
- ◆ named data streams to any non-VxFS 4.0 client,

Rules for Indicating Pathnames in the Backup Selections List

an error message appears in the Restore Monitor, informing the user that the extended attributes or named data streams could not be restored. NetBackup then continues with the restore job.

▼ **To disable the restore of extended attribute files and named data streams**

To disable the restore of extended attribute files (on Solaris 9 clients) and named data streams (on VxFS 4.0 clients), add an empty file named `IGNORE_XATTR` to the client in the following directory:

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

File `IGNORE_XATTR` was formerly known as `IGNORE_XATTR_SOLARIS`.

Note Only the modified GNU `tar` that is supplied with NetBackup is able to restore the extended attributes or named data streams to a client. (See “Reading Backup Images with `tar`” in *NetBackup System Administrator’s Guide, Volume II*.)

Note Extended attributes and named data streams cannot be compressed.



Pathname Rules for NetWare NonTarget Clients

For NetWare systems that are running the NonTarget version of NetBackup client software, specify the pathnames in the following form:

/SMDR/TSA/TS/resources/directory/file

Where:

- ◆ *SMDR* (Storage Management Data Requestor) is the name of the NetWare file server that is running the SMDR.NLM used for backups. (NLM means NetWare-loadable module.)
- ◆ *TSA* (Target Service Agent) is a NetWare software module that prepares the data for backup or restore by the SMDR. There are different types of TSAs, depending on the data. For example, there are TSAs for NetWare file systems and DOS workstations.
- ◆ *TS* is the Target Service, which is the NetWare entity that has the data being handled by the selected TSA. For example, with the DOS TSA (tsasms.com) it is a DOS Workstation. In the case of a NetWare file system TSA, it is the system with the NetWare file systems to be backed up.
- ◆ *resources* are the specific resources on the target service. For example, it can be NetWare file systems such as BINDERY, SYS, and USER.
- ◆ *directory/file* is the directory and file that are in the resource (if it is a path to a specific file).

Observe the following rules for paths:

- ◆ Give the server access to each path or the scheduled backup will fail. To provide this access, use the **Allowed Scheduled Access** command on the **Backup** menu in the NetBackup interface on the NetWare client. For more information, see the *NetBackup for Novell NetWare Client System Administrator's Guide*.
- ◆ Enter one pathname per line.
- ◆ Start all pathnames with a slash (/).
- ◆ Precede each component in the path with a slash.

If the last component in the path is a directory, follow it with a slash (/). The trailing slash is not required but is a reminder that the path is to a directory instead of a file.

/TILE/TILE.NetWare File System/TILE/SYS/DOC/

If the last component is a file, include the file extension and omit the slash from the end of the name.

/TILE/TILE.NetWare File System/TILE/SYS/DOC/TEST.TXT

- ◆ All components in a pathname must show upper and lower case letters as they appear in the actual pathname on the client.

Rules for Indicating Pathnames in the Backup Selections List

- ◆ Wildcard usage is the same as when specifying files for Windows clients.
- ◆ To back up all NetBackup for NetWare clients that are in this policy, enter a slash (/) by itself on a line.

/

- ◆ To back up an entire NetBackup for NetWare client, enter a slash (/) followed by the client name and a slash.

/TILE/

The following example backs up SYS, BINDERY, and USER file systems under the file system TSA on the client named tile:

```
/TILE/TILE.NetWare File System/TILE/SYS/
/TILE/TILE.NetWare File System/TILE/BINDERY/
/TILE/TILE.NetWare File System/TILE/USER/
```

Note that the **Allowed Scheduled Access** command on the **Backup** menu in the NetBackup interface on the NetWare client must also specify access to these paths. See the *NetBackup for Novell NetWare Client System Administrator's Guide*.

Pathname Rules for NetWare Target Clients

For NetWare clients that are running the target version of NetBackup client software, use the following format for the pathnames:

/target/

Where *target* is the name of a target defined on the NetBackup for NetWare client. See the *NetBackup Administrator's Guide for Novell NetWare Clients*.

- ◆ Enter one target per line.
- ◆ Start all target names with a slash (/).
- ◆ All target names must be in upper case.
- ◆ Wildcard usage is the same as for Windows clients.

The following example backs up the targets: NETWARE, SYSTEM, and BINDERY:

```
/NETWARE/
/SYSTEM/
/BINDERY/
```

Rules for Indicating Pathnames in the Backup Selections List

Pathname Rules for Clients Running Extension Products

Pathname rules for NetBackup clients that are running separately-priced extension products, such as Advanced Client or NetBackup for MS-Exchange, are covered in the NetBackup guide for the extension product.

Backup Selections List Directives: General Discussion

The backup selections list for a policy can contain directives that signal NetBackup to perform specific actions when processing the files in the selections list.

The available directives depend on the policy type and whether the **Allow Multiple Data Streams** attribute is enabled for the policy. The following example is a backup selections list that contains the `NEW_STREAM` directive. The example is from an MS-Windows-NT policy that has **Allow Multiple Data Streams** enabled:

```
NEW_STREAM
D:\Program Files
NEW_STREAM
C:\Winnt
```

The actions that the `NEW_STREAM` directive causes are explained in “Backup Selections List Directives for Multiple Data Streams” on page 184.

ALL_LOCAL_DRIVES Directive

Use the `ALL_LOCAL_DRIVES` directive to back up all local drives except for those drives that use removable media. The `ALL_LOCAL_DRIVES` directive applies to the following policy types:

- ◆ Standard (except for NetWare target clients)
- ◆ MS-Windows-NT
- ◆ NetWare (NonTarget clients only)

However, using `ALL_LOCAL_DRIVES` for NetWare policy types is not allowable if you are also using **Allow Multiple Data Streams**. (See “`ALL_LOCAL_DRIVES` Directive” on page 188.)

See “Files Excluded from Backups by Default” on page 190 for information on files and directories that NetBackup automatically excludes from backup.

SYSTEM_STATE Directive

The `System_State:\` directive is a valid directive only when backing up Windows 2000/XP machines. If the machine is not one of these types, and not Windows 2003 Server, then `System_State:\` will not have any effect.

Windows 2003 Server computers recognize the `System_State:\` directive and behave as if following the `Shadow Copy Components:\` directive. A message informs the user that this directive translation has occurred.

If the machine is Windows 2000/XP, the list of items that get backed up can include the following:

Rules for Indicating Pathnames in the Backup Selections List

- ◆ Active Directory
- ◆ COM+ Class Database
- ◆ Cluster Database
- ◆ IIS Database
- ◆ Registry
- ◆ Boot Files and Protected Files
- ◆ SYSVOL
- ◆ Certificate Server

On Windows 2000, the registry gets backed up in the process of regular file system backups. 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 affects Windows 2003 Server systems that use the Volume Shadow Copy components.

The `Shadow Copy Components:\` directive specifies that all of the Volume Shadow Copy component writers get backed up. Selecting this directive insures that all of the necessary components will be backed up.

The Volume Shadow Copy components include the following:

- ◆ *System State* writers, which can include:
 - ◆ System Files
 - ◆ COM+ Class Registration Database
 - ◆ SYSVOL
 - ◆ Active Directory
 - ◆ Cluster Quorum

Rules for Indicating Pathnames in the Backup Selections List

- ◆ Certificate Services
- ◆ Registry
- ◆ Internet Information Services
- ◆ *System Service* writers, which can include:
 - ◆ Removable Storage Manager
 - ◆ Event Logs
 - ◆ Windows Internet Name Service
 - ◆ Windows Management Instrumentation
 - ◆ Remote Storage
 - ◆ Dynamic Host Configuration Protocol
 - ◆ Terminal Server Licensing
 - ◆ Background Intelligent Transfer Service
- ◆ *User Data* writers, which include items that are not required by the machine to operate. For example, Active Directory Application Mode.
- ◆ *Other Data* writers, a category intended for future NetBackup releases.

Directives for Multiple Data Streams

If the **Allow Multiple Data Streams** general attribute is set for a policy, you can use the following directives in the backup selections list:

- ◆ NEW_STREAM
- ◆ ALL_LOCAL_DRIVES
- ◆ UNSET
- ◆ UNSET_ALL

The rules for using these directives are explained in “Backup Selections List Directives for Multiple Data Streams” on page 184.

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. The policy types that currently have their own backup selections list directives are:

Rules for Indicating Pathnames in the Backup Selections List

- ◆ AFS
- ◆ FlashBackup
- ◆ NDMP
- ◆ Split-Mirror
- ◆ Lotus-Notes
- ◆ MS-Exchange-Server

For example, the following directives can appear only in the backup selections list for an AFS policy:

```
CREATE_BACKUP_VOLUMES
```

```
SKIP_SMALL_VOLUMES
```

Except for AFS, the policy types listed above can be used when their associated separately-priced option is installed.

For information on the other policies and their backup selections list directives, see the NetBackup guide for the option.

Caution Include policy-specific directives only in backup selections lists for the policies that support the directives or errors can occur.

Backup Selections List Directives for Multiple Data Streams

If the **Allow Multiple Data Streams** general attribute is set for the policy, the following directives can be used in the backup selections list to control the way that NetBackup creates backup streams:

- ◆ NEW_STREAM
- ◆ ALL_LOCAL_DRIVES
- ◆ UNSET and UNSET_ALL

Note For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times because the heads must move back and forth between tracks containing files for the respective streams.

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 or absence of `NEW_STREAM` on the first line of the backup selections list determines whether the backup is performed in *administrator-defined* streaming or *auto-discover* streaming mode.

Administrator-defined Streaming Mode

If `NEW_STREAM` is on the first line of the backup selections list, the backup is performed in administrator-defined streaming mode and the following occurs:

- ◆ The backup is split 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 are in the same stream.
- ◆ The end of each stream is defined by the start of a new stream (that is, a `NEW_STREAM` directive).
- ◆ The last stream in the backup selections list is terminated by the end of the backup selections list.

Note 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 because the heads must move back and forth between tracks containing files for the respective streams.

For example, consider the backup selections list below:

```
NEW_STREAM
/usr
/lib
NEW_STREAM
/home
/bin
```

This backup selection list has two data streams.

- ◆ The `NEW_STREAM` at the top of the list invokes administrator-defined streaming and starts 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 you add a backup selections list entry as part of an existing stream, its first backup is according to the next schedule that is due for the policy. If the next backup due is an incremental, then only changed files 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.

Rules for Indicating Pathnames in the Backup Selections List

In the previous example, assume you add `/var` after `/bin`. If an incremental is due that evening, only changed files in `/var` are backed up. However, if you add a `NEW_STREAM` directive before `/var`, then NetBackup performs a full backup of all files in `/var`, regardless of when they were last changed.

Auto-discover Streaming Mode

Auto-discover streaming mode is invoked if `NEW_STREAM` is not the first line of the backup selections list *and* the list contains 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 but splits each drive volume (Windows) or file system (UNIX) into its own backup stream. See “`ALL_LOCAL_DRIVES` Directive” on page 188.
- ◆ If wildcards are used, the expansion of the wildcards results in one stream per wildcard expansion.

If the backup selections list contains neither the `ALL_LOCAL_DRIVES` directive nor wildcards, auto-discover mode is not used and preprocessing is done on the server rather than the client. In this case, each file path in the backup selections list becomes a separate stream.

Auto-discover streaming mode applies to:

- ◆ Standard and MS-Windows-NT policy types, except for NetWare clients.
- ◆ Clients that are running NetBackup 3.2 or later.

With auto-discover, the client determines how many streams are required by preprocessing the backup selections list before the backup begins. The first backup of the policy always includes preprocessing. However, preprocessing does not necessarily occur before every backup and whether it occurs depends on the preprocess interval.

Setting the Preprocess Interval for Auto-discovery

The preprocess interval applies only to auto-discover mode and specifies how often preprocessing occurs. When a schedule is due and auto-discovery is used, NetBackup checks whether the previous preprocessing session occurred within the preprocess interval:

- ◆ If yes, NetBackup does not run preprocessing on the client.
- ◆ If no, NetBackup runs preprocessing on the client and makes required changes to the streams.

If necessary, you can change the interval by using the `bpconfig` command. The default is four hours and is a good value for most sites that run daily backups. If the interval is too long or too short, the following can occur:

- ◆ Too long an interval can result in new streams not being added soon enough and backups can be missed. For example, assume the preprocess interval is set to four hours and a schedule has a frequency of less than four hours. Here, it is possible for a new stream to be omitted from the next backup because the preprocessing interval has not expired when the backup is due.
- ◆ Too short an interval 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 there are a large number of clients that the server must contact for preprocessing.

The form of the `bpconfig` command to use for changing the interval is:

```
/usr/opensv/netbackup/bin/admincmd/bpconfig [-prep hours]
```

For example:

```
/usr/opensv/netbackup/bin/admincmd/bpconfig -prep 12
```

You can set the preprocess interval for immediate preprocessing by specifying `-prep 0`. (Preprocessing occurs prior to every backup.) Specifying `-prep -1` sets the preprocess interval to the default value of 4 hours.

The following example sets the preprocess interval to 12 hours. You can determine the current interval by using the `bpconfig` command with the `-L` option:

```
bpconfig -L
```

(output of the above command)

```
Mail Admin:          *NULL*
Wakeup Interval:     9 minutes
Max Jobs/Client:     8
Backup Tries:        2 in 12 hours
Keep Logs:           3 days
Max drives/master:   0
Maximum Backup Copies: 10
Compress DB Files:   older than 10 days
Media Mnt Timeout:   0 minutes (unlimited)
Shared Timeout:      0 minutes (unlimited)
Display Reports:     24 hours ago
Keep TIR Info:       1 days
Prep Interval:       12 hours
Max Backup Copies:   2
DB Clean Interval:   12 hours
Policy Update Interval: 10 minutes
```

Rules for Indicating Pathnames in the Backup Selections List

Auto-Discover Streaming Mode Example

Assume the selection list has the following entries:

```
/usr  
/lib  
/home/*
```

For this selection list, NetBackup generates:

- ◆ One stream for the `/lib` directory
- ◆ One stream for the `/usr` directory
- ◆ One stream for each subdirectory and file in the `/home` directory because of the wildcard (*)

If the `/home` directory has three subdirectories: `tom`, `dick`, and `harry`, but no files, NetBackup produces a separate stream for each subdirectory: `/home/tom`, `/home/dick`, and `/home/harry`. This is a total of five streams for the backup.

However, if the wildcard is removed from `/home`, as in the following, then auto-discover is not used.

```
/usr  
/lib  
/home
```

In this mode, NetBackup generates only three streams, one for each of the directories in the list. Preprocessing is done on the server instead of the client.

ALL_LOCAL_DRIVES Directive

The `ALL_LOCAL_DRIVES` directive applies only to Standard (except for NetWare target clients), MS-Windows-NT, and NetWare policies where the clients are running NetBackup 3.2 or later software. If used, this directive must be the only entry in the backup selections list for the policy; that is, no other files or directives can be listed.

The action that `ALL_LOCAL_DRIVES` causes depends on whether **Allow Multiple Data Streams** is enabled for the policy.

- ◆ If **Allow Multiple Data Streams** is enabled, the `ALL_LOCAL_DRIVES` directive is valid only if the policy type is Standard (except for NetWare clients) or MS-Windows-NT. In this instance, NetBackup backs up the entire client and splits the data from each drive (Windows) or file system (UNIX) into its own backup stream. NetBackup periodically runs preprocessing on the client to make necessary changes to the streams. See “Setting the Preprocess Interval for Auto-discovery” on page 186.
- ◆ If **Allow Multiple Data Streams** is not enabled, NetBackup backs up the entire client and includes all drives and file systems in the same stream.

Caution Do not select **Cross Mount Points** for policies where you use the `ALL_LOCAL_DRIVES` directive.

Example 1

Assume **Allow Multiple Data Streams** is enabled in auto-discover mode and the client is a Windows system with two drive volumes, C: \ and D: \. The backup selections list contains:

```
ALL_LOCAL_DRIVES
```

For this backup selections list, NetBackup generates:

- ◆ 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.

Example 2

Assume **Allow Multiple Data Streams** is not enabled and the client is a Windows system with two drive volumes, C: \ and D: \. The backup selections list contains:

```
ALL_LOCAL_DRIVES
```

Here, NetBackup backs up the entire client in one data stream that contains the data from both C: \ and D: \.

SYSTEM_STATE is also backed up because SYSTEM_STATE is included in the `ALL_LOCAL_DRIVES` directive.

UNSET and UNSET_ALL Directives

All policy-specific directives that are passed to a client in a stream are passed in all subsequent streams. (See “Directives for Specific Policy Types” on page 183.) The `UNSET` and `UNSET_ALL` directives change this behavior. These directives are recognized only if **Allow Multiple Data Streams** is set for the policy.

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 and included in the current and later streams.

Rules for Indicating Pathnames in the Backup Selections List

UNSET_ALL

UNSET_ALL has the same effect as UNSET but unsets all policy-specific directives that have been defined up to this point in the backup selections list.

Example

Assume you have a backup selections list as shown below. In this 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
```

If you want the `set` command passed to the first two streams but not the last, an `UNSET` or `UNSET_ALL` can be used at the beginning of the third stream to prevent it from being passed to the last stream.

```
NEW_STREAM
set destpath=/etc/home
/tmp
/use
NEW_STREAM
/export
NEW_STREAM
UNSET_ALL [or UNSET set destpath=/etc/home]
/var
```

Excluding Files from Backups

There are a number of files and files states which, by default, are not backed up by NetBackup. You can also exclude specific files from automatic backups by specifying the files or directories in an exclude list on the client.

Files Excluded from Backups by Default

By default, NetBackup does not back up the following files:

- ◆ NFS files or directories, unless you set **Follow NFS**.
- ◆ Files or directories in a different file system if you do not set **Cross Mount Points**.

Rules for Indicating Pathnames in the Backup Selections List

- ◆ Files or directories with path lengths longer than 1023 characters.
- ◆ Files or directories where the operating system does not return inode information (the `lstat` system call failed).
- ◆ Directories that NetBackup cannot access (cannot use the `cd` command to access).
- ◆ On a disk managed by Storage Migrator, migrated files or directories where Storage Migrator does not return inode information (`mig_stat` fails). Note that NetBackup Server does not support Storage Migrator.
- ◆ Socket special files (named pipes are backed up).
- ◆ Locked files when mandatory locking is enabled by an application that currently has the file open.
- ◆ Busy files. If a file is open, NetBackup backs up the last saved version of the file.

NetBackup automatically excludes the following file system types on most platforms:

- ◆ `cdrom` (all UNIX platforms)
- ◆ `cachefs` (AIX, Solaris, SGI, UnixWare)
- ◆ `devpts` (Linux)
- ◆ `mntfs` (Solaris)
- ◆ `proc` (UNIX platforms; does not exclude automatically for AIX, so `/proc` must be added manually to exclude list. If not added manually, partially successful backups may result when using the `ALL_LOCAL_DRIVES` directive on AIX)
- ◆ `tmpfs` (Linux)
- ◆ `usbdevfs` (Linux)

Excluding Files from Automatic Backups

On most NetBackup clients, you can exclude specific files from automatic backups by specifying them in an exclude list on the client.

You can also create an include list to add back in some of the files by using an include list. The include list is useful, for example, if you want to exclude an entire directory except for one file.

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

The method for specifying files in the exclude and include lists depends on the type of client that you are configuring.

Rules for Indicating Pathnames in the Backup Selections List

- ◆ On Microsoft Windows clients, specify exclude and include lists in the Backup, Archive, and Restore client interface: Start Backup, Archive, and Restore and click **File > NetBackup Client Properties**. Go to the **Exclude List** or **Include List** tab. For further instructions, see the NetBackup user's guide for the client.

The **Exclude List** or **Include List** can also be specified through the NetBackup Administration Console on the master server. (See "Exclude Lists Properties" on page 387.)

- ◆ On NetWare target clients, the exclude and include lists are specified when adding the targets. See the NetBackup user's guide for the client.
- ◆ On UNIX clients, you create the exclude and include lists in the following files on the client:

```
/usr/opensv/netbackup/exclude_list
```

```
/usr/opensv/netbackup/include_list
```

Where Will the Catalog Data Be Located: Disaster Recovery Tab

The Disaster Recovery tab displays for *NBU-Catalog* policies only. The tab contains information indicating where the data that is crucial to disaster recovery is to be located.

Note VERITAS recommends saving the image file to a network share or a removeable device. Do not save the disaster recovery information to the local machine.

The screenshot shows a window titled "Change Policy - catalog_2" with a tabbed interface. The "Disaster Recovery" tab is selected. It contains the following fields and controls:

- Server:** A text box containing "hagar".
- Path:** A text box for specifying the directory.
- Logon:** A text box for the username.
- Password:** A text box for the password.
- Send in an E-mail attachment (recommended):** A checked checkbox.
- E-mail address:** A text box for the email address.
- Critical policies:** A list box with a "Name" header and an "Add" button to its right. Below the list box are "Change" and "Delete" buttons.
- Information box:** A message stating: "The disaster recovery file generated for each catalog backup contains information needed to recover the NetBackup catalog. This file also contains media information necessary to recover critical policies. Record the location of this file so that the NetBackup catalog can be recovered if necessary."
- Buttons:** "Apply", "OK", "Close", and "Help" at the bottom.

Path

Enter the path to the directory where the disaster recovery information will be saved. Specify a local directory or NFS share.

Logon

Enter the logon and password information to access the NFS share.

Password

Enter the password needed to log on to the share.

Where Will the Catalog Data Be Located: Disaster Recovery Tab

Note Logon and password fields are enabled:

- When the path begins with \\ indicating a UNC share path.
 - When NetBackup doesn't have write access to the location user specified.
-

Send in an E-mail Attachment

VERITAS recommends sending the disaster recovery report to at least one e-mail address. Enter the e-mail address where the information will be sent. To send the information to more than one address, separate multiple e-mail addresses using a comma:

email1,email2

Note See "Setting Up E-Mail Notifications" on page 417 to ensure that the system is set up to send mail.

Identifying Critical Policies

The policies placed on the **Critical Policies** list are identified as being crucial to getting a site up and running quickly in the event of a disaster. The NetBackup Disaster Recovery report that is generated when online catalog backups are run, lists all of the media used for backups of critical policies, including the most recent full backup. The NetBackup Disaster Recovery wizard warns you if any media for critical policies are not available.

Note Only policies that use incremental or full backup schedules should be designated as critical policies—only media for these backup schedules will be listed in the Disaster Recovery report.

Note When NetBackup is running with Vault, Vault includes the disaster recovery data in the Vault Recovery Report.

Creating a Vault Policy

Creating a Vault policy differs from creating other policies in the following ways:

- ◆ You must specify *Vault* as the policy type.
- ◆ You do not specify clients for Vault policies, therefore the Clients tab does not appear.
- ◆ Rather than specifying files to back up in the backup selection list, specify one of two Vault commands to run: `vltrun` or `vlteject`

When configuring a Vault policy, be sure to specify Vault as the policy type. Instead of entering a directive in the backup selections list, you'll indicate one of two Vault commands. There are no clients specified in Vault policies.

▼ To create a Vault policy

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. 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**. Complete the schedule.

Note The Clients tab does not appear for Vault policy types.

6. On the Backup Selections tab, enter one of two Vault commands:
 - ◆ Use `vltrun` to specify the robot, vault name, and profile for the job. The `vltrun` command accomplishes all the steps necessary to select, copy, and eject media. If the vault profile name is unique, use the following format:


```
vltrun profile_name
```

If the vault profile name is not unique, use the following format:

```
vltrun robot_number/vault_name/profile_name
```
 - ◆ Use the `vlteject` command to eject media and/or generate reports for Vault sessions that have been completed already and for which media has not been ejected. For example:


```
vlteject -eject -report [-vault vault_name [-sessionid id]]  
[-auto y|n] [-eject_delay seconds]
```

Performing Manual Backups

Both commands are located in the following directory:

`/usr/opensv/netbackup/bin/`

For more information on Vault names, profile names, and command usage, see the *Vault System Administrator's Guide*.

7. Click **OK**.

Performing Manual Backups

You can perform immediate manual backups of selected automatic backup schedules and clients within a policy. A manual backup is useful for situations such as:

- ◆ Testing a configuration.
- ◆ Backing up a client that missed the regular backup.
- ◆ Backing up a client before installing new software to preserve the old configuration.
- ◆ Preserving records before a special event such as a company split or merge.
- ◆ Backing up quarterly or yearly financial information.

In some cases, it may be useful to create a policy and schedule that you use only for manual backups. Do this by creating a policy with a single schedule that has no backup window (and therefore never runs automatically).

▼ To perform a manual backup

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies** and select the policy name in the middle pane.
2. Select **Actions > Manual Backup**. (The policy must be set to Active for this command to be available.) The Manual Backup dialog appears.

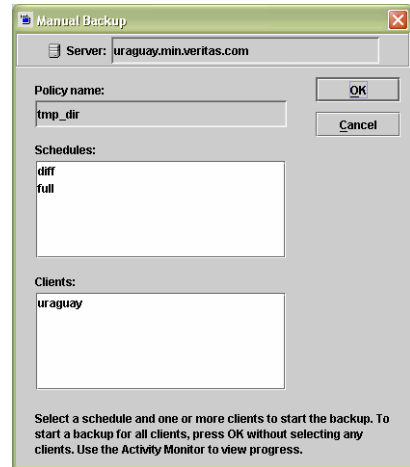
Note Not only does the policy need to be Active, but if **Go into effect** is set on the policy to a future date and time, the backup will not run.

3. In the Manual Backup dialog, 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 and cannot be manually backed up because they do not have a backup selection list (the user selects the files).

4. Click **OK** to start the backup.



More About Synthetic Backups

There are a number of reasons why implementing synthetic backups may be useful in your NetBackup configuration. Synthetic backups can be written to tape or disk storage units, or any combination thereof.

Processing Takes Place on Master and Media Server(s) Instead of Client

One advantage of synthesizing a full backup is that all of the processing takes place on the master and media server and not the client. During a traditional full backup, all files are copied from the client to a master or media server, even though those files may not have changed since the last incremental backup.

When creating a synthetic full backup, NetBackup takes full advantage of the fact that new or changed files have already been copied to the media server during the last incremental backup. NetBackup does not require that the client even be running in order to combine the full and incremental backups on the media server to form a new, full backup. The new, full synthetic backup is an accurate representation of the clients' file system at the time the most recent full backup has been run.

Reduces Network Traffic

Another benefit is that files are transferred over the network only once, reducing network traffic. After the full and incremental backup images have been combined into a synthetic full (or cumulative) backup, the tapes or disk containing the component images can be recycled or reclaimed, thereby reducing the number of tapes or disk space in use.

Supports Disk Environments

Synthetic backups can be run in environments comprised exclusively of disk storage units.

Uses Drives More Effectively

Synthetic backups can be written to tape or disk storage units, or a combination of both mediums. If using tape, backups can be synthesized when drives are not generally in use. For example, if backups occur primarily overnight, the drives can be busy synthesizing full backups during the day.

Policy Considerations and Synthetic Backups

Selecting the Synthetic Backup Option

The **Synthetic Backup** option is available under the following conditions:

- ◆ The policy type must be either *Standard* or *MS-Windows-NT*.
- ◆ The **Collect True Image Restore Information With Move Detection** option must be selected on the Policy Attributes tab. (See “Collect True Image Restore With Move Detection” on page 88.)
- ◆ The schedule created for a synthetic backup must have the **Synthetic Backup** option selected. (See “Schedules to Include in a Policy for Synthetic Backups” on page 199.)
- ◆ The master servers, media servers, and clients must all have NetBackup version 5.0 or later installed in order to synthesize backups.
- ◆ One of the following, or a combination thereof, must be available:
 - ◆ Disk storage unit(s) with adequate space available.
 - ◆ Tape library(s) with multiple drives for reading/writing.(See “Disk Storage Unit Considerations” and “Tape Storage Unit Considerations” on page 205.)

Schedules to Include in a Policy for Synthetic Backups

A policy for synthetic backups must contain at least three types of schedules:

- ◆ At least one traditional, full backup must be run successfully to create a full image.
The synthetic backup jobs will fail 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 will receive a status code of 1 (partially successful) for a policy that contains full or incremental synthetic backup schedules, but no incremental backup schedules.

Remember that since the synthetic backup synthesizes all of the incremental backups to create a new full or cumulative backup image, the synthetic backup is only as current as the last incremental backup.

Note If you are configuring a synthetic cumulative backup and the clients are archive bit-based (default), use only differential incremental backups for the traditional, non-synthesized backups.

- ◆ One full and/or one cumulative backup schedule with the **Synthetic Backup** option selected.

More About Synthetic Backups

Adding Clients to a Policy for Synthetic Backups

Every time a client is added to a policy that will be used for synthetic backups, the client must have a traditional, full backup created for it before a synthetic backup is possible. Upon adding a client to the policy, you must run a manual traditional full backup.

Since **Collect True Image Restoration (TIR) with Move Detection** is required for synthetic backups, all clients included in the policy must support TIR.

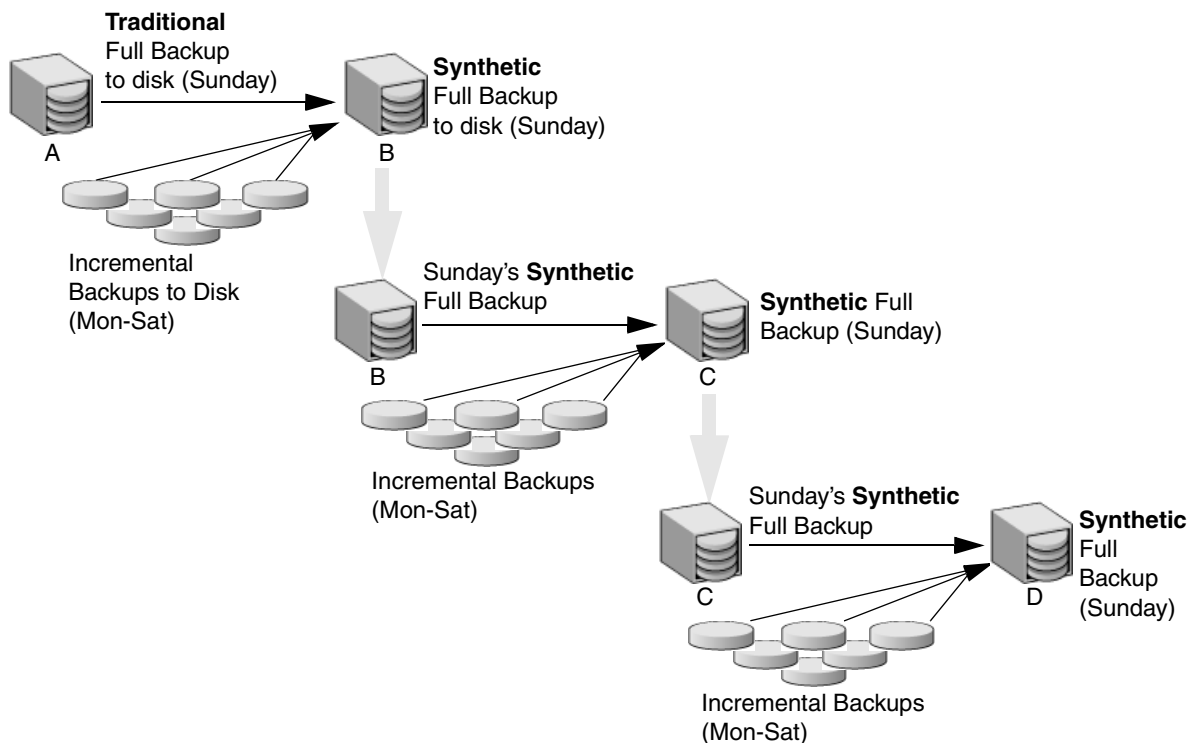
Two Types of Synthetic Backups

Two types of synthetic backup images can be created: synthetic full and cumulative synthetic. The images used to create the synthetic image are known as *component* images. For instance, the component images in a synthetic full are the previous full image and the subsequent incremental images.

Synthetic Full Backups

For a discussion of synthetic cumulative incremental backups, see “Synthetic Cumulative Incremental Backups” on page 202.

The following figure illustrates the creation of synthetic full backups (B, C, D) from an existing full backup (A) and the incremental backups between full backups.



The traditional full backup (A) and the incremental backups are created in the traditional manner—by scanning, then copying data 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.

More About Synthetic Backups

Synthetic Full Backup Usage Example

1. Create a *Standard* or *MS-Windows-NT* 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.
2. Make certain that the traditional full backup runs. If, for some reason, the backup does not complete, run the backup manually.
3. Per schedule, run daily differential incremental backups for the clients throughout the week. The last incremental for the week will run on Saturday.
4. Per schedule, run synthetic full backups for the clients on subsequent Sundays.

Note The synthetic full backups in this scenario will be only as current as the Saturday incremental backup.

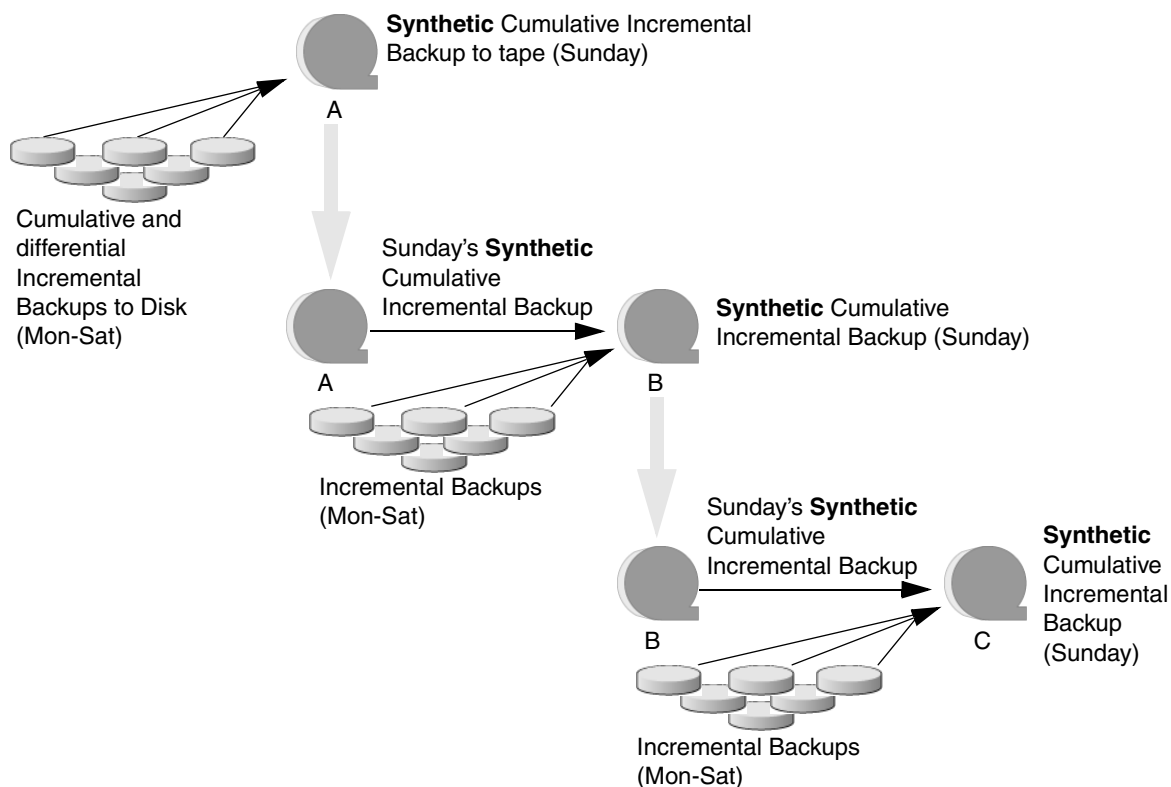
Synthetic Cumulative Incremental Backups

The scenario for creating a synthetic cumulative incremental backup is similar to that of creating 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 taken since the last cumulative backup.
- ◆ The last cumulative incremental backup. If no cumulative incremental backup is available, just the differential incremental backups are used for the synthetic image.

The following figure illustrates the creation of synthetic cumulative incremental backups (A, B, C) from the latest cumulative incremental backup and subsequent differential incremental backups.



Synthetic Cumulative Backup Usage Example

1. Create a *Standard* or *MS-Windows-NT* 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.
2. Make certain that the traditional full backup runs. If, for some reason, the backup does not complete, run the backup manually.

More About Synthetic Backups

3. Per schedule, run daily differential incremental backups for the clients throughout the week. The last incremental for the week will run on Saturday.
4. Per schedule, run synthetic cumulative incremental backups for the clients on subsequent Sundays.

Note The synthetic cumulative backups in this scenario will be only as current as the Saturday incremental backup.

Recommendations for Synthetic Backups

Scenario in Which Synthesized Backups Would Be Most Beneficial

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 incrementals will be too large and a synthetic backup would not be any more helpful than creating a traditional full backup.

Refrain From Multiplexing Backups that Will Be Synthesized

While synthesizing multiplexed backups is possible, it is not recommended because of its inefficiency. Synthesis of multiplexed client images requires multiple passes over the source media—one per client.

Synthesized Backups and Multistreaming

Performance issues similar to those encountered while multiplexing synthesized backups problems will occur if multiple streams are selected for synthesized backups. Multiple stream performance issues can be improved by backing up to disk whenever possible.

Reducing 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, consider running an incremental backup just prior to the synthetic backup.

Note Only frequency-based scheduling allows an incremental backup and a synthetic backup to run on the same day. (See “Frequency Schedule Type” on page 108.) Calendar-based schedules are dependant on one another, so if a daily incremental schedule runs earlier in the day, the synthetic cumulative backup will not run later that same day (00:00:00–23:59:59). (See “Calendar Schedule Type” on page 108.)

Disk Storage Unit Considerations

Disk-based images are more efficient for synthesizing. For example, while synthesizing a backup, NetBackup processes the newest component images first, followed by sequentially older images. When two or more component images have been written to the same tape, the tape movement may be somewhat inefficient compared to disk-based images.

Tape Storage Unit Considerations

- ◆ When performing tape backups, a tape other than the tapes where full and differential images reside, is required for the formation of a synthetic image.
- ◆ The maximum drive usage applies only to the drive needed for writing the synthetic backup. If any of the component images reside on tape, an additional drive is needed for reading.
- ◆ If a single tape drive device is used to generate synthetic images, component images (full, differential, or cumulative images) should be placed in a hard drive location first. In that way, a synthetic image can be generated with the single tape drive device.

Notes on Synthetic Backups

Test Findings Regarding Synthetic Backups

- ◆ The time it takes to run a synthetic full backup does not increase significantly over time.
- ◆ Synthetic full backups are generated more quickly when built from disk-based incremental backups. If the synthetic full backup is also generated on disk, the run time is even faster. The disk copy could then be duplicated to tape.

Test Findings Regarding Restores from Synthetic Backups

- ◆ The time 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 whether restoring 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, depending on the pattern of file changes within the directory.

More About Synthetic Backups

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 has the newest files at the front of the media and the unchanged files at the end. Over time, this introduces the potential for fragmentation of a single directory across the synthetic full image.

Note that this 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.

General Notes

- ◆ Synthetic backups are supported on all NetBackup server platforms.
- ◆ The option to create multiple copies is not allowed for synthetic backups.
- ◆ Synthetic backups are not supported if any of the component images are encrypted.
- ◆ A user-generated backup image cannot be used to generate a synthetic image. In other words, an image generated from a User Backup schedule or a User Archive schedule cannot be used as one of the components of a synthetic backup.

Synthetic Backup Jobs Create Two Sets of Catalog Files

When a synthetic backup job is run, two sets of catalog files are created: an image file and one or more .f files.

1. The first set is named using the timestamp of the most recent incremental + 1. This set represents the actual synthetic backup image which is as recent as the most recent incremental.
2. The second set is named using the current timestamp. This set is used to mark the time the synthetic backup job was run. It does not contain any file data.

Do not manually remove any of these catalog files. The catalog files are automatically expired after the retention period as specified in the schedule for the policy. (See “Retention” on page 115.) The two sets of catalogs have the same expiration times.

For example:

Catalog after running incremental backup jobs:

```

XDisk_1064417510_INCR
XDisk_1064417510_INCR.f

XDisk_1064420508_INCR
XDisk_1064420508_INCR.f

XDisk_1064421708_INCR
XDisk_1064421708_INCR.f

```

After running synthetic backup job:

First set:	XDisk_1064421709_FULL XDisk_1064421709_FULL.f	Synthetic full backup image
Second set:	XDisk_1064424108_FULL	Current time

Timestamp differences

The catalog for a synthetic image usually has a timestamp one second larger later than the most recent incremental component image. The timestamp may be more than one second larger if there were possible image name conflicts.

True Image Restore and Synthesized Backups

Since the **Collect True Image Restore with Move Detection** policy property is required for synthetic backups, all clients included in the policy must support TIR.

The TIR information in the image catalog is normally *pruned* (removed) after the number of days indicated in the master server **Clean-Up** host property, **Keep True Image Restoration (TIR) Information**. (See “Keep True Image Restoration (TIR) Information” on page 359.)

However, if a synthetic full and/or synthetic cumulative schedule has been defined in the policy, the TIR information will not be pruned from the component images until a subsequent traditional or synthetic full or cumulative backup image has been generated successfully.

For example, if the **Keep True Image Restoration (TIR) Information** host property specifies that TIR information is to be pruned from the catalog after two days, on the third day the TIR information will be pruned only if a traditional or synthetic full backup image has been generated.

If the TIR information has been pruned from one or more component images and you accidentally expire the most recent synthetic image, if you try to rerun the synthetic backup job, it will automatically restore the TIR information to the catalog. In case the TIR

More About Synthetic Backups

information cannot be restored due to bad, missing, or vaulted media, the synthetic backup job will fail with error code 136 (*TIR info was pruned from the image file*). If the problem is correctable, you can run the synthetic backup again.

Checkpoint Restart and Synthesized Backups:

If Checkpoint Restart (**Take Checkpoints** setting on the policy Attributes tab) is indicated for the policy, the backups produced with the synthetic backup schedule will not be checkpointed. Selecting **Take Checkpoints** for synthetic backups has no effect.

Change Journal and Synthesized Backups:

If the Windows client host property, **Use Change Journal in Incrementals**, is enabled, the property will have no effect when the client is backed up using the synthetic backup schedule. (See “Use Change Journal in Incrementals” on page 378.)

Displaying Synthetic Backups in the Activity Monitor

A synthetic job is distinguished from a traditional full backup by the notation 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 program `bpsynth` to manage the synthetic backup process. `bpsynth` plans how the synthetic backup will be built from the previous backup images.

`bpsynth` then schedules tape drive resources needed for the synthetic backup (if any). If the required resources cannot be obtained, the job fails with a status code indicating that a resource is needed.

If the resources can be obtained eventually but not immediately, the synthetic job waits until the resources become available. This could happen if a drive needed is currently being used by a backup or a restore, or by another synthetic backup job.

`bpsynth` then calls program `bpcoord` to coordinate reading the necessary files from each of the component images, one image at a time. (`bpcoord` cannot be run directly from the command line.)

`bpsynth` passes information to programs `bptm` and `bpdm` to cause tape and disk images to 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. See the *NetBackup Troubleshooting Guide for UNIX and Windows* for more information.

On the master server: `bpsynth`, `bpcoord`, `bpdbm` and the log files located in `/usr/opensv/logs` as described in the *NetBackup Troubleshooting Guide*.

On the media server(s): `bptm` (if any tape images), `bpdm` (if any disk images), `bpcd`

Note that several media servers may be involved if the component images are on different nodes.

However, one `bpsynth/bpcoord` pair will be used for each stream or client. This can be inefficient with tape images since each `bpsynth/bpcoord` pair needs a tape drive to write the new image. Also, each `bpsynth/bpcoord` pair may use the same component image volumes. One may have to complete before the next one proceeds.


Synthetic Backups and Directory and File Attributes

In order for a synthetic backup to include changes made to directory and file attributes (for example, access control lists (ACLs)), the change must first be picked up by a component incremental backup.



- ◆ 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 will not be reflected in an incremental backup, and therefore not in a synthetic full backup.

To include ACL changes in backups, for each UNIX client, enter `USE_CTIME_FOR_INCREMENTALS` in the `bp.conf` file on the client. (See "USE_CTIME_FOR_INCREMENTALS" in the *NetBackup System Administrator's Guide, Volume II*.)

- ◆ Windows: For each Windows client, select **Incrementals: Based on Archive Bit**. (**NetBackup Management > Host Properties > Clients > Selected client(s) > Windows Client**.)



More About Synthetic Backups



NetBackup Catalogs

4

This chapter describes the role that the catalog plays in a NetBackup environment and explains how to configure a catalog backup. The various functions that a user can perform in the Catalog utility are also explained.

This chapter contains the following sections:

- ◆ “What is a NetBackup Catalog?” on page 212
- ◆ “Catalog Protection” on page 217
- ◆ “Recovering the Catalog” on page 247
- ◆ “Disaster Recovery E-mails and the Disaster Recovery File” on page 247
- ◆ “Archiving the Catalog” on page 248
- ◆ “Using the Catalog Utility” on page 254
- ◆ “Catalog Maintenance and Performance Optimization” on page 275

What is a NetBackup Catalog?

NetBackup catalogs are internal databases that contain information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and storage devices.

NetBackup requires the catalog information in order to restore backups so it is extremely important to configure a catalog backup before using NetBackup for regular client backups, and to schedule the catalog backups to occur on a regular basis thereafter. Without regular catalog backups, you risk losing your regular backups if there is a problem with the disk that contains the catalogs.

Parts of the Catalog

The NetBackup catalog resides on the disk of the NetBackup master server. The catalog consists of the following parts:

- ◆ Image database: The image database contains information about what has been backed up. It is by far the largest part of the catalog. (See “Image Database” on page 213.)
- ◆ NetBackup data stored in relational databases: This includes the media and volume data describing media usage and volume information which is used during the backups. (See “NetBackup Relational Database” on page 216)
- ◆ NetBackup configuration files: Policy, schedule and other flat files used by NetBackup.

Catalog Configuration (Default)

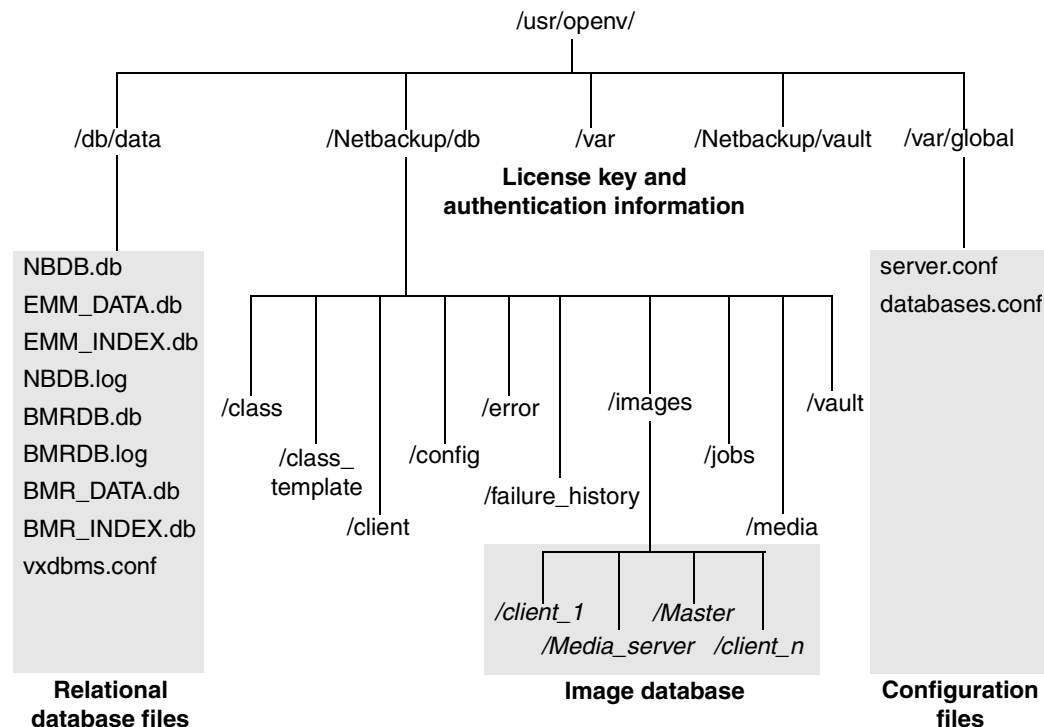


Image Database

The image database (`/usr/opensv/netbackup/db/images`) contains subdirectories for each client backed up by NetBackup, including the master server and any media servers.

The image database contains:

- ◆ Image files (files which store only backup set summary information)
- ◆ Image .`tf` files (files which store the detailed information of each file backup)

The image database is the largest part of the NetBackup catalog and consumes about 99% of the total space required for the NetBackup catalog. While most of the subdirectories in the NetBackup catalogs are relatively small, `/images` can grow to several tens, or even hundreds of gigabytes. Typically, it is the image database on the master server that may grow too large to fit on a single tape, an important consideration for offline, cold catalog backups. Image database growth depends on the the number of clients, policy schedules, and the amount of data backed up.

What is a NetBackup Catalog?

To determine catalog space requirements, see “Determining Catalog Space Requirements” on page 275.

If the image catalog becomes too large for the current location, consider moving the image catalog to a file system or disk partition that contains more space. See “Moving the Image Catalog” on page 282 for more information.

The image database component of the NetBackup catalog uses the .f files in binary format for Windows, Solaris, HP_UX, Compaq Tru64 UNIX, AIX, and Linux platforms.

The image database of existing catalogs can be upgraded to binary format using the catalog conversion utility, `cat_convert` as described in “Catalog Conversion Utility” on page 281.

Image Files

Each image file is an ASCII file, generally less than 1 kilobyte in size, containing only backup set summary information. For example, the backup ID, the backup type, the expiration date, fragment information, and disaster recovery information.

Image .f Files

The binary catalog may contain one or more image .f files. This type of file is also referred to as a files-file. The image .f file may be large because it contains the detailed backup selection list for each file backup. Generally, image files range in size from 1 kilobyte to 10 gigabytes.

Whether the catalog contains one .f file or many .f files is determined by the file layout. 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

When the detailed backup file information of one catalog backup is less than 4 megabytes in size, NetBackup stores the information in a single image .f file. The image .f file is always greater than or equal to 72 bytes, but less than 4 megabytes.

The following is an 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 detailed backup file information of one catalog backup is greater than 4 megabytes, the information is stored in multiple .f files: one main image .f file plus nine additional .f files.

Separating the additional .f files from the image .f file and storing the files in the `catstore` directory improves performance while writing to the catalog.

The main image .f file is always exactly 72 bytes.

```
-rw----- 1 root other      72 Aug 30 00:40 test_1030680524_INCR.f
-rw----- 1 root other      804 Aug 30 00:08 catstore/test_1030680524_INCR.f-list
-rw----- 1 root other 1489728 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgDir0
-rw----- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgExtraObj0
-rw----- 1 root other 1280176 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgFile0
-rw----- 1 root other     192 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgHeader0
-rw----- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgNDMP0
-rw----- 1 root other 9112680 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgRecord0
-rw----- 1 root other 2111864 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgStrings0
-rw----- 1 root other      11 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgUserGroupNames0
```

NetBackup Relational Database

NetBackup installs Sybase Adaptive Server Anywhere 9.0.1 during the master server installation, as a private, non-shared server for the NetBackup database (NBDB). The same installation of Sybase ASA is used for the optionally-licensed product, Bare Metal Restore (BMR) database (BMRDB). The BMR database is created during the BMR installation process.

As part of the catalog backup, the database and configuration files for the NBDB and BMRDB databases are protected:

- ◆ Database files:
 - ◆ `/usr/opensv/db/data/NBDB.db`
 - ◆ `/usr/opensv/db/data/EMM_DATA.db`
 - ◆ `/usr/opensv/db/data/EMM_INDEX.db`
 - ◆ `/usr/opensv/db/data/NBDB.log`
 - ◆ `/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_DATA.db` (if BMR is installed)
 - ◆ `/usr/opensv/db/data/BMR_INDEX.db` (if BMR is installed)
- ◆ 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.

For information on the use of Sybase ASA in the NetBackup catalog, see Appendix A, “NetBackup Relational Database” on page 537

Description of NetBackup Relational Database

NBDB database: The NBDB database contains media and device information managed by the Enterprise Media Manager (nbemm) and the resource utilization information managed by the Resource Broker (nbrb).

BMR database: The BMRDB database contains information managed by Bare Metal Restore.

Catalog Protection

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, the catalog must be protected by a particular type of backup--a *catalog backup*. A catalog backup backs up catalog-specific data as well as produces disaster recovery information.

As additional protection for the catalog, you may consider archiving the catalog as well. (See "Archiving the Catalog" on page 248.)

Note The "Recommended Backup Practices," section in Chapter 7 of the *NetBackup Troubleshooting Guide* provides helpful setup information that can aid in simplifying disaster recovery. Since the catalog plays a critical role in the NetBackup environment, much of the information concentrates on catalog considerations.

Catalog Backups

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.

Note Configure a catalog backup *before running any regular backups*.

Note If any NetBackup utilities or other methods have been used to relocate portions of the NetBackup catalog, these changes should be noted as they become part of subsequent catalog backups. In the event that a catalog recovery is needed, the same alterations will need to be implemented prior to the recovery of the catalog.

Choose the catalog backup method that works best for your environment:

- ◆ Online, hot catalog backup (recommended method)

This type of catalog backup is for highly active NetBackup environments in which continual backup activity is typically occurring. It is considered an *online, hot* method because it can be performed while regular backup activity is taking place. This type of catalog is policy-based and can span more than one tape. It also allows for incremental backups, which can significantly reduce catalog backup times for large catalogs.

Online, hot catalog backups use media from the *CatalogBackup* volume pool only.

For more information, see "Online, Hot Catalog Backup Method" on page 218.

- ◆ Offline, cold catalog backup

Catalog Protection

This type of catalog backup is for NetBackup environments in which there are periods when no backup activity is occurring. It is considered an *offline, cold* backup because it should not be run when regular backup activity is taking place. For Sybase ASA, the databases (NBDB and BMRDB) are shut down during the backup. This type of catalog backup must fit on a single tape.

Offline, cold catalog backups use media from the *NetBackup* volume pool only.

For more information, see “Offline, Cold Catalog Backup Method” on page 228.

For additional information that applies to both types of catalog backups, see “Strategies to Ensure Successful Catalog Backups” on page 280.

Online, Hot Catalog Backup Method

The online, hot catalog is new in NetBackup 6.0. It is policy-based, which means that it has all of the scheduling flexibility of a regular backup policy. This catalog backup type is designed for use in highly active NetBackup environments where there is usually backup activity taking place and the catalog size is large.

The online, hot catalog backup:

- ◆ Can back up the catalog while continual client backups are in progress.
- ◆ Can span multiple tapes for a catalog backup.
- ◆ Allows for a flexible pool of catalog tapes.
- ◆ Can perform a full or incremental catalog backup.
- ◆ Can restore the catalog to a different location.
- ◆ Can run scheduled catalog backups.
- ◆ Offers a wizard to automate the catalog recovery process or a guided command line tool.
- ◆ Appends to existing data on tape.
- ◆ Can be duplicated.

Configure an online catalog backup using one of the following methods:

- ◆ By using the Catalog Backup Wizard. (See “To configure an online, hot catalog backup using the Catalog Backup Wizard” on page 219.)
- ◆ By using the Backup Policy Configuration Wizard. (See “To configure an online, hot catalog backup using the Backup Policy Wizard” on page 224.)
- ◆ By selecting the *NBU-Catalog* type when creating a backup policy. (See “To configure an online, hot catalog backup using the Policy utility” on page 225.)

▼ To configure an online, hot catalog backup using the Catalog Backup Wizard

Note Online, hot catalog policies (policy type *NBU-Catalog*) write only to media in the *CatalogBackup* volume pool. This procedure assumes that there is a configured storage device and media available in the *CatalogBackup* volume pool. See the *Media Manager System Administrator's Guide* for more information about adding media to a volume pool.

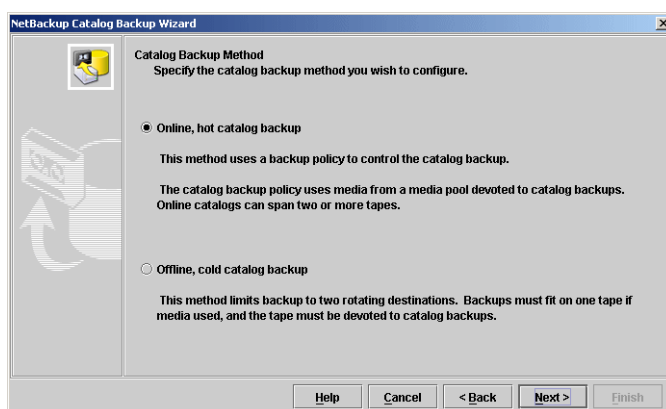
1. Launch the **NetBackup Catalog Backup Wizard** by clicking **Configure the Catalog Backup** in the right pane. The wizard is visible when either **Master Server** or **NetBackup Management** is selected in the left pane.

Click **Help** within any wizard screen for more information on the wizard settings.

2. Click **Next** on the Welcome screen.

3. On the Catalog Backup Method screen, select to configure an online, hot catalog backup.

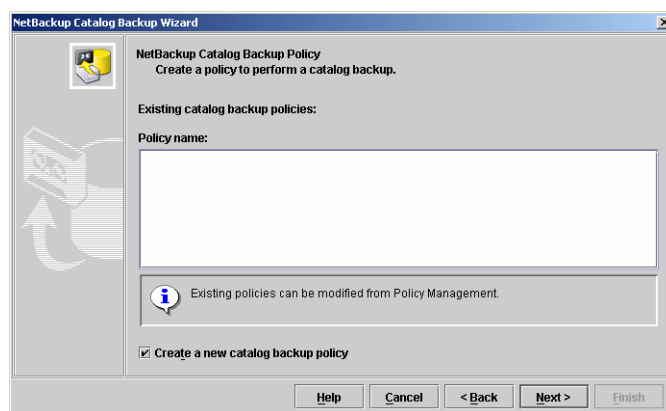
Click **Next**.



4. On the Catalog Backup Policy screen, select a policy from the list of existing catalog backup policies.

Or, to create a new catalog backup policy, select **Create a new catalog backup policy**.

Click **Next**. This launches the Policy Name and Type screen of the Backup Policy wizard.



Catalog Protection

5. 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. (See “NetBackup Naming Conventions” on page 29.)

Click **OK**.

Backup Policy Configuration Wizard

Policy Name and Type
Specify the policy name.

Policy name:
CatPol1

The policy type determines the types of clients that can be backed up by this policy or the type of backups that this policy will perform on those clients.

Policy type:
NBU-Catalog

Help Cancel < Back Next > Finish

6. Select the backup type.

(**User Backup** does not apply for *NBU-Catalog* policies.)

Backup Policy Configuration Wizard

Backup Type
Specify the types of backups.

☒ Full Backup
Backs up all the catalog related data.

☒ Incremental Backup
Backs up all changed files specified in the selection list.

☒ Differential (files changed since last full or incremental backup)

☐ Cumulative (files changed since last full backup)

☐ User Backup
Allows users to initiate backups on their own.

Help Cancel < Back Next > Finish

7. Select the rotation schedule.

The selection **After each backup session** refers to a period when no regular backup policy is running.

By default, a frequency-based schedule is selected. This ensures that in busy environments where there are usually jobs running, the catalog backup will be run.

Backup Policy Configuration Wizard

Rotation
Select a rotation for backup and retention.

Start a full backup every:
1 weeks

Retain full backups for:
2 weeks

Start an incremental backup

☒ By frequency, every:
1 days

Retain incremental backups for:
2 weeks

☐ After every backup session

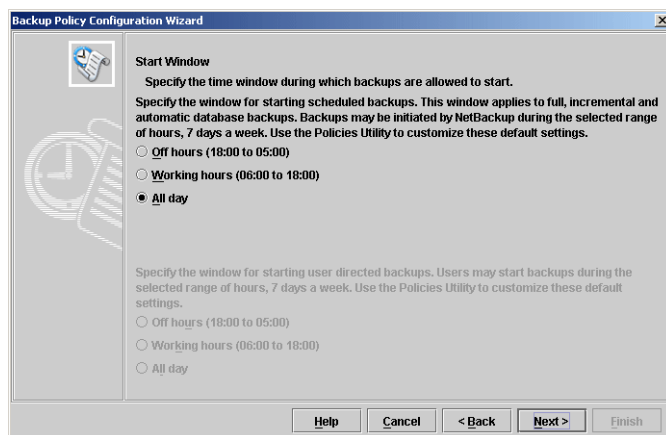
To configure a calendar based backup schedule, modify this policy with the Policies Utility in the NetBackup Administration Console.

Help Cancel < Back Next > Finish

Note See “Running Online, Hot Catalog Backups Concurrently with Other Backups” on page 226 if online, hot catalog backups are scheduled to run concurrently with other backup types for the master server.

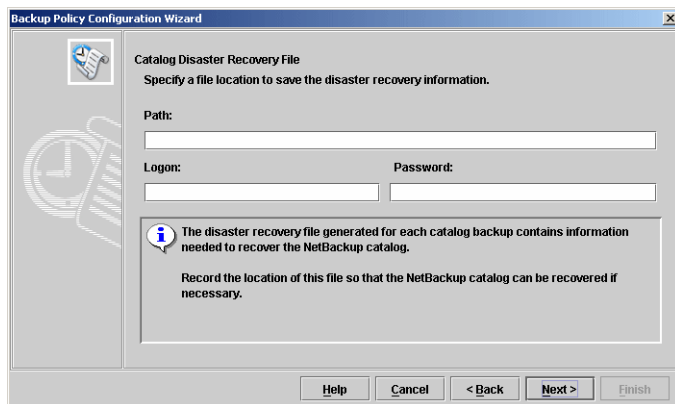
8. The Start Window screen helps you define a window of time during which the catalog backup will start. The scheduled windows (**Off hours**, **Working hours**, **All day**) are preset in the wizard, but can be customized using the **Policies** utility.

User Window selections are disabled, as regular users (those who are not NetBackup administrators) cannot start catalog backups.



9. Select where each disaster recovery image file will be saved on disk. The image file contains the disaster recovery information.

Note VERITAS recommends saving the image file to a network share or a removeable device. Do not save the disaster recovery information to the local machine.



Enter the path to the directory where the disaster recovery information will be saved. Specify a NFS share.

If necessary, enter the logon and password information to access the NFS share.

Catalog Protection

Note If the logon information is not valid, NetBackup returns a message requesting that the user reenter the logon and password information and/or clear the alternate location option to continue.

Logon and password fields are enabled:

- When the path begins with \\ indicating a UNC share path.
- When NetBackup doesn't have write access to the user specified location.

10. VERITAS strongly recommends configuring your NetBackup environment to send the disaster recovery information to a NetBackup administrator. This backup-specific information is sent after every catalog backup.

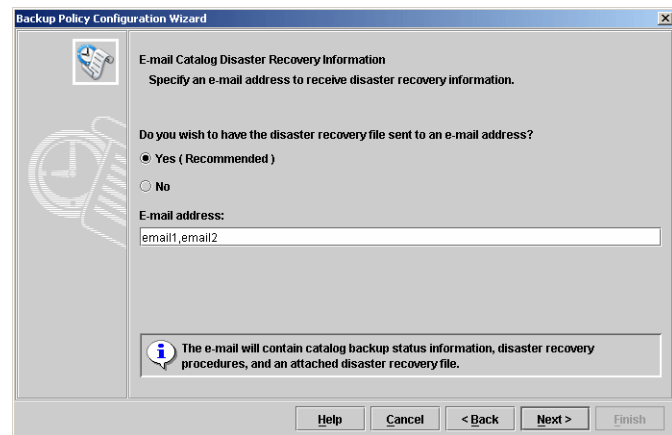
To send the information to more than one administrator, separate multiple e-mail addresses using a comma:

email1,email2

Make sure that e-mail notification is enabled in your environment. (See "Send in an E-mail Attachment" on page 194.)

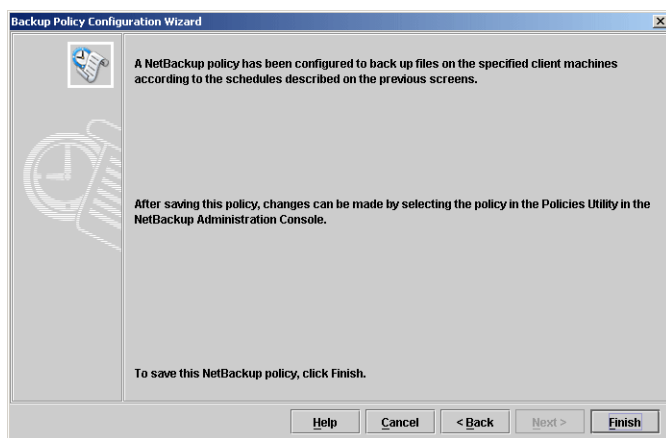
For more information on the e-mail that is sent and the attached disaster recovery file, see "Disaster Recovery E-mails and the Disaster Recovery File" on page 247.

Note The disaster recovery e-mail is not sent to the address specified in the Global Attributes property. The **Administrator's E-mail Address** Global Attributes property specifies the address(es) where NetBackup sends notifications of scheduled backups or administrator-directed manual backups.



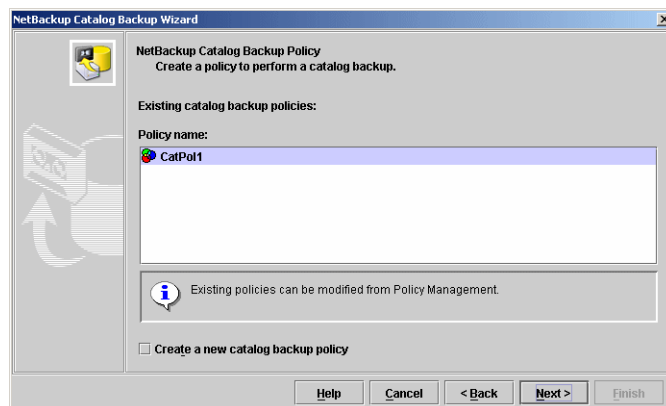
11. The last screen of the policy wizard lets you know that once the policy is created, you can make changes in **NetBackup Management > Policies**.

Click **Finish** to create the policy.



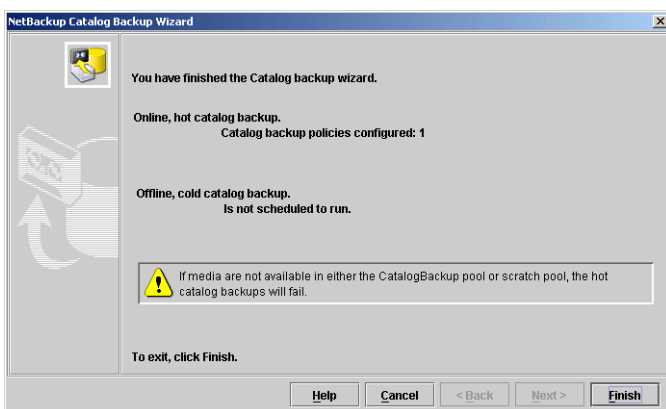
12. The Catalog Backup Wizard resumes, listing the new policy.

13. Click **Next** to finish the Catalog Backup Wizard.



14. The final Catalog Backup Wizard screen displays the total number of catalog backup policies configured for this master server.

15. Click **Finish** to complete the wizard.



Catalog Protection

▼ To configure an online, hot catalog backup using the Backup Policy Wizard

Using the Backup Policy Wizard to create an online, hot catalog backup policy is similar to using the Catalog Backup Wizard.

Note Online, hot catalog policies (policy type *NBU-Catalog*) write only to media in the *CatalogBackup* volume pool. This procedure assumes that there is a configured storage device and media available in the *CatalogBackup* volume pool. See the *Media Manager System Administrator's Guide* for more information about adding media to a volume pool.

1. Launch the **Backup Policy Configuration Wizard** by clicking **Create a Backup Policy** in the right pane. The wizard is visible when either **Master Server** or **NetBackup Management** is selected in the left pane.

Click **Help** within any wizard screen for more information on the wizard settings.

2. Click **Next** on the Welcome screen.
3. Enter the policy name and select *NBU-Catalog* as the policy type. (See "NetBackup Naming Conventions" on page 29.)
Click **OK**.
4. Select the backup type. (**User Backup** does not apply for *NBU-Catalog* policies.)
5. Select the rotation schedule. For *NBU-Catalog* policy types running incremental backups, **Every backup session end** means a period when no regular backup policy is running.
6. The Start Window screen helps you define a window of time during which the catalog backup will start. The windows are preset in the wizard, but can be customized using the **Policies** utility.

User Window selections are disabled because regular users (those who are not NetBackup administrators) cannot start catalog backups.

7. Select where each disaster recovery image file will be saved on disk.

Note VERITAS recommends saving the image file to a network share or a removeable device. Do not save the disaster recovery information to the local machine.

Disk Information

- ◆ Path to the disk:

For UNIX master servers, specify a local directory or NFS share.

For Windows master servers, specify a local directory or UNC path (CIFS Windows share).

- ◆ Logon information used to:
 - ◆ Access the NFS share for UNIX master servers.
 - ◆ Access the Windows share for Windows master servers.

Note If the logon information is not valid, NetBackup returns a message requesting that the user reenter the logon and password information and/or clear the alternate location option to continue.

Logon and password fields are enabled:

- When the path begins with \\ indicating a UNC share path.
- When NetBackup doesn't have write access to the location user specified.

8. VERITAS strongly suggests having the disaster recovery information e-mailed to an NetBackup administrator in your organization. The disaster recovery information is sent after every catalog backup.

9. Click **Finish** to complete the wizard.

▼ **To configure an online, hot catalog backup using the Policy utility**

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. (See "NetBackup Naming Conventions" on page 29.)

Click **OK**.

4. The following fields on the Attributes tab apply to online, hot catalog backups. For information about other fields on this tab, see "What Type of Policy: Policy Attributes Tab" on page 69.

Policy Type: Select *NBU-Catalog* as the policy type.

Policy Storage Unit: If you're indicating a disk storage unit, increase the **Maximum Concurrent Jobs** setting on the storage unit. This ensures that the catalog backup can proceed while regular backup activity is occurring.

Catalog Protection

Note Online catalog backups must use media servers at version 6.0 or later to store catalog backup data. If your installation contains 5.x and 6.0 media servers hosting disk storage units, do not select *Any Available* for the destination **Policy Storage Unit**, since it is possible that the 5.x media 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. Offline, cold catalog backups use media from the *NetBackup* volume pool.

5. Set up a schedule for an online catalog backup as you would for any other policy. For information on the Schedules tab settings, see “When Will the Job Run: Schedules Tab” on page 97.

See the following sections (below) for additional information regarding online catalog backup schedules: “Running Online, Hot Catalog Backups Concurrently with Other Backups” and “Notes on Catalog Policy Schedules.”

Note The Clients tab does not apply to the NBU-Catalog policy and is not displayed.

6. The Disaster Recovery tab appears for NBU-Catalog policies only. The tab contains information regarding the location of data crucial to disaster recovery:
 - ◆ The path to the directory where the disaster recover information will be saved.
 - ◆ The logon and password information to the share.
 - ◆ The e-mail address(es) where disaster recovery reports will be sent.
 - ◆ A critical policy list, containing the names of policies that are backing up critical data. Media containing critical policy backups are listed on the NetBackup Disaster Recovery Report that is generated when the online catalog backup is run.

For information on the Disaster Recovery tab settings, see “Where Will the Catalog Data Be Located: Disaster Recovery Tab” on page 193.

Running Online, Hot Catalog Backups Concurrently with Other Backups

If online, hot catalog backups are scheduled 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 is occurring:

- ◆ In the Global Attributes host properties for the master server, set the **Maximum Jobs per Client** value to greater than one. (See “Global Attributes Properties” on page 414.)
- ◆ On the storage unit where the backups will be sent, increase the **Maximum Concurrent Jobs** setting. (See “Maximum Concurrent Jobs” on page 42.)

Notes on Catalog Policy Schedules

The following schedules are supported in the online, hot catalog backup policy type:

- ◆ Full
- ◆ Differential incremental (depends on a full schedule)
- ◆ Cumulative incremental
- ◆ Session-based differential incremental
- ◆ Session-based cumulative incremental
- ◆ Although it is possible to configure multiple catalog backup policies, VERITAS recommends configuring only one.
- ◆ Online catalog backups must use media servers at version 6.0 or later to store catalog backup data. If your installation contains 5.x and 6.0 media servers hosting disk storage units, do not select *Any Available* for the destination **Policy Storage Unit**, since it is possible that the 5.x media server could be selected.
- ◆ The incremental schedule depends on a full schedule.
- ◆ The least frequent schedule runs if many schedules are due at the same time.
- ◆ There can be more than one session-based incremental schedule in one catalog backup policy:
 - ◆ If one is cumulative and the other is differential, the cumulative runs when the backup session ends.
 - ◆ If both are cumulative or both are differential, the first schedule found runs when the backup session ends.
- ◆ The queued scheduled catalog backup is skipped if there is a catalog backup job still running from the same policy.
- ◆ *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 there is a catalog backup job 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, and should both be considered when duplicating, verifying, or expiring the backup.

Note Additional child catalog jobs are created for 5.x media servers and the BMR database if a remote EMM server is configured.

Offline, Cold Catalog Backup Method

The *offline, cold catalog backup* can be used in NetBackup environments with catalogs small enough to fit onto a single tape. This type of catalog backup is for environments that have time during which no backup activity is taking place.

Configure an offline catalog backup using one of the following methods:

- ◆ By using the Catalog Backup Wizard. (See “To configure an offline, cold catalog backup using the Catalog Backup Wizard” on page 229.)
- ◆ By selecting **Actions > Configure offline NetBackup Catalog Backup**. (See “To configure an offline, cold catalog backup using the Actions menu” on page 236.)

While the offline, cold catalog backup is suitable in many existing NetBackup installations, this type of catalog backup differs from the online, hot catalog backup in a number of ways. The items listed below are the same limitations that existed in previous NetBackup releases. They are provided here for purposes of comparison.

The offline, cold catalog backup:

- ◆ Cannot back up the catalog while any job or catalog operation is running.
- ◆ Cannot span tapes.
- ◆ Is not tracked like a regular backup. Since online, hot catalog backups are run as jobs via a policy, the online, hot catalog backup job is included in the catalog.
- ◆ Cannot be duplicated as regular backups and online, hot catalog backups can be duplicated.
- ◆ Cannot be written to media and given an expiration date.
- ◆ Cannot be an incremental backup.
- ◆ Overwrites any data on a tape. This type of catalog backup starts at the beginning of the tape and does not append to existing data on the tape.

▼ **To configure an offline, cold catalog backup using the Catalog Backup Wizard**

Note Offline, cold catalog policies write only to media in the *NetBackup* volume pool. This procedure assumes that there is a configured storage device and media available in the *NetBackup* volume pool. See the *Media Manager System Administrator's Guide* for more information about adding media to a volume pool.

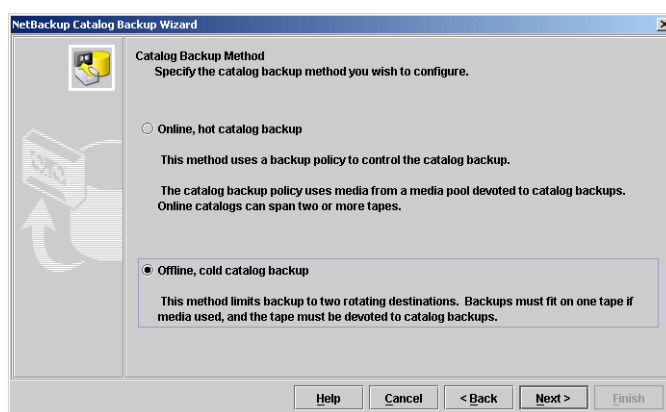
1. Launch the **NetBackup Catalog Backup Wizard** by clicking **Configure the Catalog Backup** in the right pane. The wizard is visible when either **Master Server** or **NetBackup Management** is selected in the left pane.

The wizard is also available by:

- ◆ Selecting the Catalog utility, then selecting **Actions > Configure Offline NetBackup Catalog Backup**.
- ◆ Right-clicking the Catalog utility in the left pane and selecting **Configure Offline NetBackup Catalog Backup**.

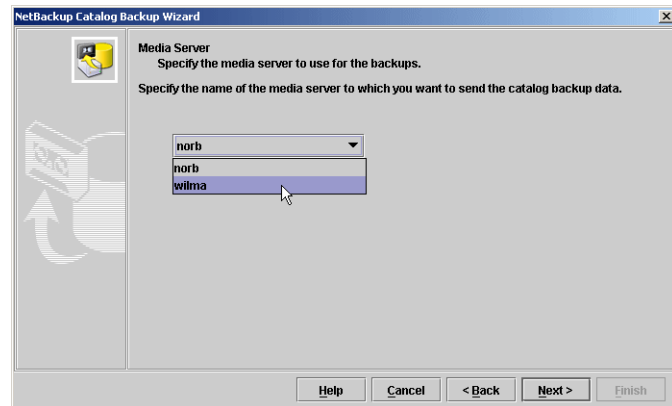
Click **Help** within any wizard screen for more information on the wizard settings.

2. Click **Next** on the Welcome screen.
3. On the Catalog Backup Method screen, select to configure an offline, cold catalog backup. Click **Next**.

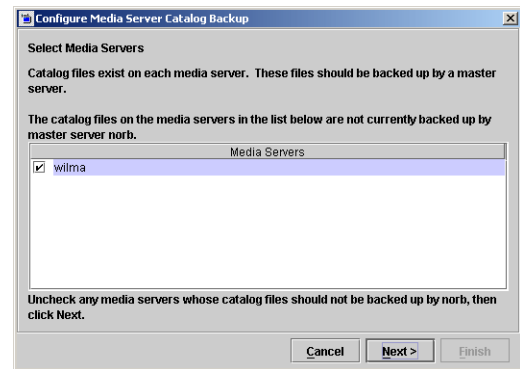


Catalog Protection

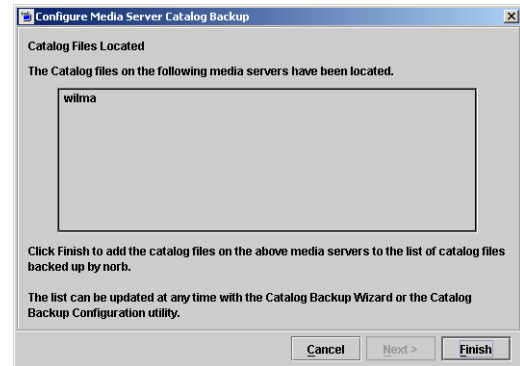
4. Specify the media server where the catalog backups are to be sent.



5. Select the media servers that should be backed up in this catalog backup.



6. NetBackup displays where catalog files have been located.



7. Indicate where the catalog files are located.

The pathnames of the catalogs on the master server are automatically added during installation and generally require no action on your part other than to ensure that the directories are listed.

Verify that the catalogs of the master server and each media server are included. Verify that the absolute path names are correct and are in the correct format.

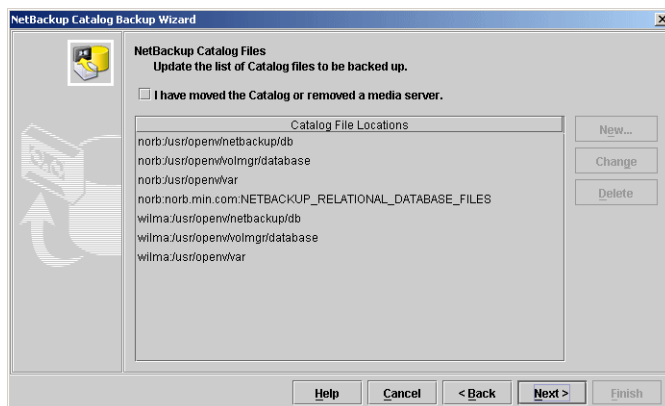
The `NETBACKUP_RELATIONAL_DATABASE_FILES` directive automatically includes the database files in the `/usr/opensv/db/data/` directory, as well as `vxdbms.conf`, `server.conf`, and `databases.conf`. (If the files have been relocated to different directories using `nbdb_move`, the locations will be automatically determined and the files will be included.)

If using NetBackup Access Control (NBAC) in the NetBackup configuration, add the following directives for each host in the NBAC domain:

```
[host:]nbat
```

```
[host:]nbaz
```

Note If the master server using NBAC is a UNIX machine, VERITAS recommends that you do not include the NetBackup master server configuration file (`/usr/opensv/netbackup/bp.conf`) in the offline catalog backup file list. If `bp.conf` is included in the list, it must not be recovered until all other catalog recovery is completed.



If NetBackup cannot find or follow a path, the entire catalog backup fails.

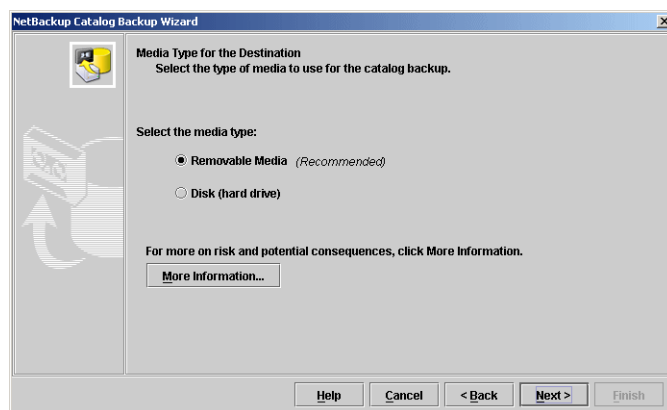
Notes Regarding the Catalog File Locations List

- ◆ If you have moved the location of the catalog on the master server, the new location must be specified.
- ◆ Pathnames to the catalog on media servers are not automatically added during installation and must be added to this list.
- ◆ If `/usr/opensv/netbackup/db/images` is a symbolic link to another filesystem, you *must* specify the true location of the images directory here. Symbolic links do not apply to Windows.

Catalog Protection

- ◆ Do not specify a symbolic link as the final component in a UNIX path or the entire catalog backup will fail.
While NetBackup follows links at other points in the path, NetBackup does not follow a link when it is the final component. If any other part of a listed path is a symbolic link, NetBackup saves the actual path during the backup.
- ◆ The files that are associated with the NetBackup relational databases, NBDB and BMRDB, are automatically included in the catalog backup and don't need to be listed here. Since the locations of these files can change, the pathnames are dynamically generated during the backup.

8. If both removable media and disk is configured, the Media Type for the Destination screen displays. Specify whether the backup will be to removable media or to disk. For this example, we'll select **Removable Media**.

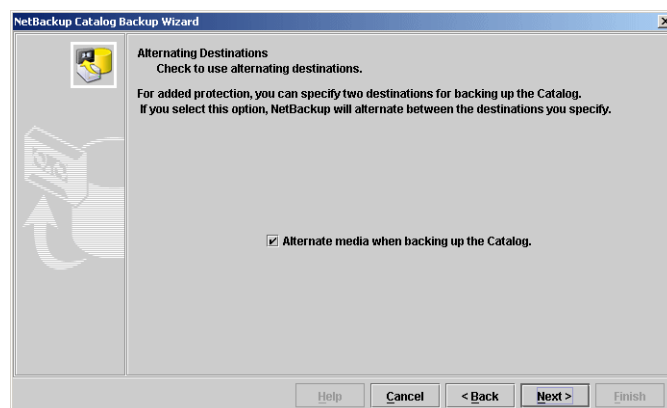


9. Select a destination volume from the drop-down list to store the catalog backup.

Note Note the media ID of the volume selected so you know which volume contains catalog backups.

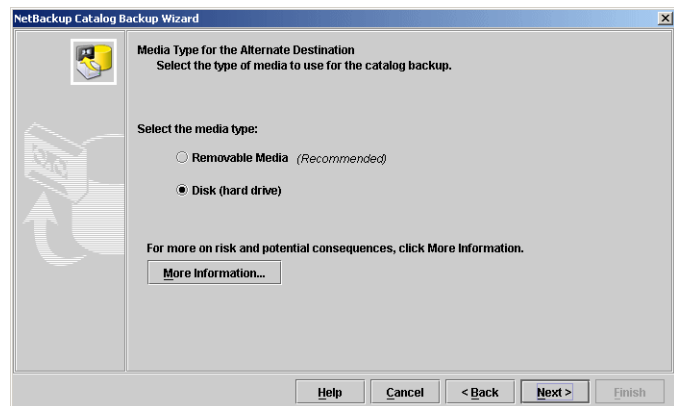
10. You may prefer to back up the catalog to two different locations, alternating for each catalog backup.

Alternating between the two destinations adds protection in the event that a disk or tape containing a catalog backup is lost.



NetBackup will always back up to the media not used for the previous catalog backup.

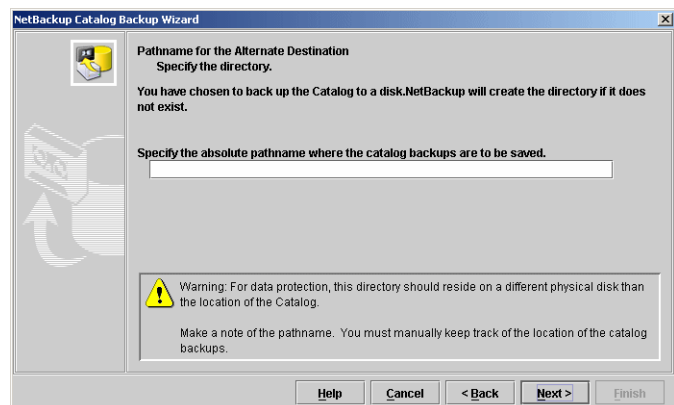
11. Specify whether the alternate destination for the catalog backup will be to removable media or to disk. For this example, we'll select **Disk**.



12. Enter the path to the alternate destination for the catalog backup.

The path can point to:

- ◆ A directory on a disk attached to the master server. NetBackup creates the directory if it does not exist.
- ◆ An NFS-mounted file system or a link to an NFS-mounted file system that grants write access to the root user.
- ◆ A shared directory on another computer. This shared directory must be available to the service account that the NetBackup Client service logs into at startup.



In addition to the platform-specific file path separators (/ and \) and colon (:) within a drive specification on Windows, follow the NetBackup naming conventions. (See "NetBackup Naming Conventions" on page 29.)

Catalog Protection

- 13.** The frequency of catalog backups is configurable. Choose one of the following options:

After each session of scheduled, user, or manual backups

Consider this option if you are sending your catalog backups to a robot or tape stacker, a second standalone tape drive, or to disk.

After each session of scheduled backups

Consider this option if you are sending your catalog backups to a robot or tape stacker, a second standalone tape drive, or to disk, or if there is only one scheduled backup session per day/night.

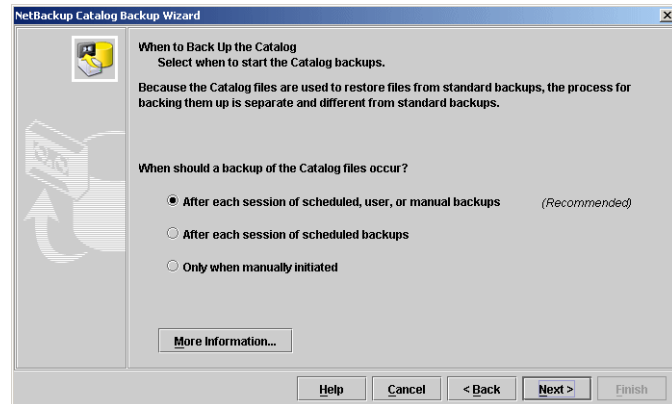
This schedule type requires that no jobs are running in order for the catalog backup to run.

Caution Do not use this schedule in NetBackup environments in which continual backup activity is typically occurring. It could mean that the catalog backup would never have an opportunity to run.

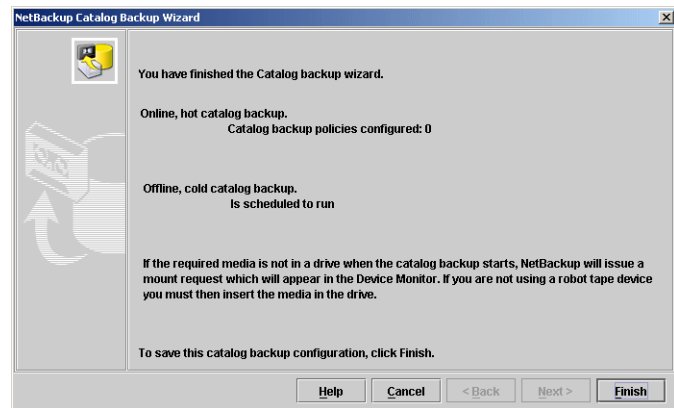
Only when manually initiated

Consider this option if you will be running multiple backup sessions in a single day or night. Be certain to perform a manual catalog backup once a day or after a series of backups.

Caution It is imperative that you back up your catalogs often. If the offline catalog backup files are lost, you lose information about backups and configuration changes that were made between the time of the last offline catalog backup and the time that the disk crash occurred.



14. The final wizard screen displays the total number of catalog backup policies configured for this master server. Click **Finish** to complete the wizard.



Catalog Protection

▼ To configure an offline, cold catalog backup using the Actions menu

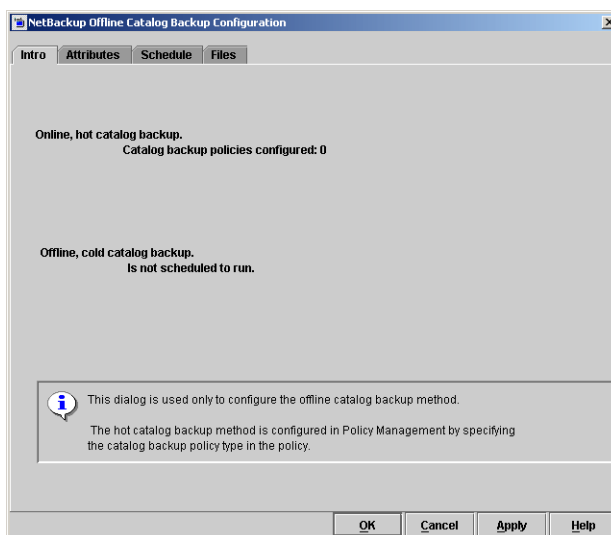
Note Offline, cold catalog policies write only to media in the *NetBackup* volume pool. This procedure assumes that there is a configured storage device and media available in the *NetBackup* volume pool. See the *Media Manager System Administrator's Guide* for more information about adding media to a volume pool.

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. With the **Catalog** utility in focus, select **Actions > Configure Offline NetBackup Catalog Backup**. The Catalog Backup Configuration dialog appears containing four tabs: Intro, Attributes, Schedule, Files.
3. Specify the properties on each tab of the dialog:
 - ◆ “Catalog Intro Tab” on page 236
 - ◆ “Catalog Attributes Tab” on page 237
 - ◆ “Catalog Schedule Tab” on page 241
 - ◆ “Catalog Files Tab” on page 243
4. Click **OK**.

Catalog Intro Tab

There is no information to specify on the Catalog Intro tab. It serves to inform users how many online, hot catalog backup policies are configured and whether or not there is an offline, cold backup scheduled.

The tab also reminds users that there is another type of catalog backup available that might serve their needs better: the online, hot catalog backup. See “Catalog Backups” on page 217.



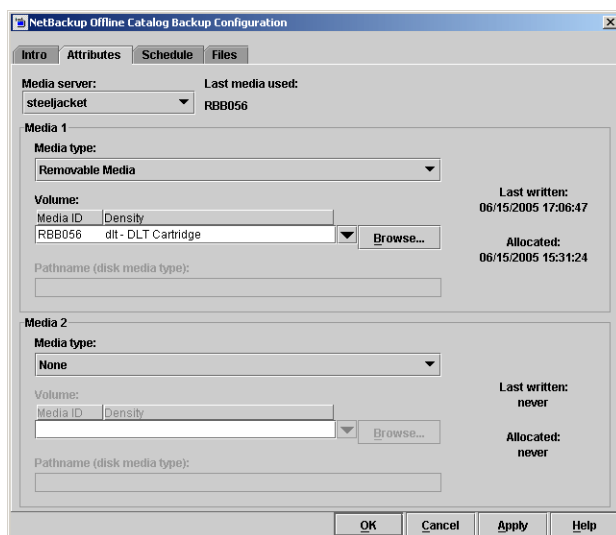
Catalog Attributes Tab

The **Attributes** tab contains general attributes for NetBackup catalog backups.

Media Server

The following setting applies only to NetBackup Enterprise Server:

The **Media Server** setting specifies the name of the media server to which catalog backups will be sent. This defaults to the master server where you are running the NetBackup Administration Console. To choose a server, select one from the drop-down menu. The list shows all servers that have a storage unit defined on the master server where you are changing the configuration.



If you are backing up the catalogs to a media server, modify the NetBackup catalog backup paths on the master server using the **Catalog Files** tab. (See “Catalog Files Tab” on page 243. Also, ensure that the media server was named in the `bp.conf` file on the master server at the time that you started `bprd` and `bpdbm`.

On NetBackup Server, **Media Server** cannot be changed and is the NetBackup server where the catalogs reside.

Last Media Used

The **Last Media Used** setting shows the media ID (for **Removable Media**) or absolute pathname (for disk) that contains the last NetBackup catalog backup. The value in this field is the value that you specified for either Media 1 or Media 2. These are the media that NetBackup alternates between for catalog backups.

Media 1 and Media 2 Areas

The **Media 1 and Media 2 Areas** setting specifies the media to use for the catalog backups. You do not have to assign both Media 1 and Media 2. If you do assign both, NetBackup alternates between the media.

Media Type

The **Media Type** setting specifies the media type. Select one from the drop-down menu:

- ◆ **None:** No media is assigned
- ◆ **Disk:** A directory on a disk drive

Catalog Protection

- ◆ **Removable Media:** A volume that is in a robot or drive under control of Media Manager

Depending on the storage devices that are available, VERITAS recommends the following choices for **Media Type**:

1. If you have a robot or a tape stacker, choose **Removable Media** and use this automated device to store the catalog backups. This is the easiest way to back up your catalogs because NetBackup automatically finds the volume if it is in a robot or tape stacker when the backup is started.
2. If you do not have a robot or tape stacker, but have an extra standalone storage device that you can devote to catalog backups, choose **Removable Media** and use the extra standalone device.
3. If you have only one standalone drive (no robot or tape stacker), the most convenient method is to choose **Disk** for the media type and send the catalog backups to a hard drive. The hard drive that you use for the catalog backup must be different than the hard drive where the catalogs reside. By default, the catalogs are stored in the following locations:

`/usr/opensv/netbackup/db`

`/usr/opensv/var/global`

`/usr/opensv/db/data`

If you choose to back up the catalog to disk, the destination of the catalog backup must be on a different drive.

Caution The safest way to protect your data is to save all backups (including your catalog backup) to removable media, then move a full set of that media to offsite storage on a regular basis. A backup written only to disk will share the same risks as the computer(s) being backed up. A natural disaster (for example, lightning, flood or fire) is more likely to destroy both your primary data and its backups if the backups are only on disk.

If the disks holding the catalogs and the catalog backup are both destroyed, it will be much more difficult to recover your business data. Assuming the backups of your business data are on tape, recovering without the catalog backup means manually importing all of the backup tapes to rebuild the catalogs. This process takes time that you may not want to spend when you need to resume your business activities.

4. If you have only one standalone drive (no robot or tape stacker) and there is not enough space available on a different hard drive, choose **Removable Media**. In this situation, you must back up the catalogs to the same tape drive as the backups of your

business data. This involves swapping tapes in and out of the drive each time the catalogs are backed up. Swapping tapes is not convenient, but is required because NetBackup will not place catalog backups and the backups of your business data on the same tape.

Media ID

If you've chosen **Removable Media**, specify a valid media ID.

The volume you specify must be configured under **Media** in the same manner as other NetBackup volumes. This means the media ID must appear under **Media and Device Management > Media**. The volume must also meet the following requirements:

- ◆ The volume must be in the NetBackup volume pool. To verify, look under **Media** and ensure that the **Volume Pool** column for the media ID displays NetBackup.
- ◆ The volume cannot be currently assigned to NetBackup for backups because NetBackup does not mix catalog backups and regular backups on the same media.

To locate an available volume, expand **Media and Device Management > Media** and find a volume where the **Time Assigned** column is empty and the **Status** column is 0. Once a catalog backup occurs, the **Time Assigned** and the **Status** column for the volume updates.

Note If a column does not appear, size the columns by right-clicking in the pane and selecting **Columns** from the shortcut menu.

The **Last Written** information under Media 1 and Media 2 indicate when the volume specified in the Media ID field was last used. The value is *never* if the volume has never been used for NetBackup catalog backups.

Note If you delete and then add back the media ID for a volume that was used for NetBackup catalog backups, NetBackup changes its Last Written date and time. However, the contents of the volume itself are not altered until the next time the volume is used for a backup.

The **Allocated** information under Media 1 and Media 2 indicate when the media was allocated for NetBackup catalog backups.

Notes on the Media ID

- ◆ To delete the media for Media 1 or Media 2, set the **Media Type** value to None. Do not use backspace to leave the Media ID box blank.
- ◆ If you delete a volume from the catalog backup configuration, Media Manager makes it available for reassignment. This can cause problems if, for example, you temporarily backup the catalog to a different volume.

Catalog Protection

- ◆ You must manually track catalog backup media separately because NetBackup does not keep a record of catalog backup media in its catalogs as it does with other backup media. If NetBackup did track catalog backup media in the catalog, and the disk containing the catalogs crashed, the record would be lost with the catalogs.

A convenient way to track the media is to indicate an e-mail address in the Global Attributes properties. NetBackup sends an e-mail that includes the status of each catalog backup and the media ID that was used. Print the e-mail or save it on a disk other than the disk containing the catalogs. (See “Global Attributes Properties” on page 414.)

If the catalogs are intact, you can also find these media IDs in the Media Manager volume listing. The Status column shows 1 for these volumes. However, these IDs do not appear in the NetBackup media reports.

Pathname (Disk Media Type)

For disk media, this is the path to the directory where you want to store the catalog backup. Type the path in the field. For example:

```
/nb/dbbackup
```

The path can be any of the following:

- ◆ A directory on a disk attached to the master server. NetBackup creates the directory if it does not exist.
- ◆ An NFS-mounted file system or a link to an NFS-mounted file system that grants write access to the root user.

When backing up the catalogs to disk, observe the following precautions:

- ◆ Always back up to a physical disk other than the one containing the catalogs. For example, if your computer has two physical disks and the catalogs are on the first disk, back up the catalogs to the second disk. If you back up the catalogs to the same disk and that disk fails, both the catalogs and the backups are lost and it will be difficult or impossible to restore data for your NetBackup clients. By default, the catalogs are stored in the following locations, so the destination of your catalog backup must be on a different disk:

```
/usr/opensv/netbackup/db
```

```
/usr/opensv/var/global
```

```
/usr/opensv/db/data
```

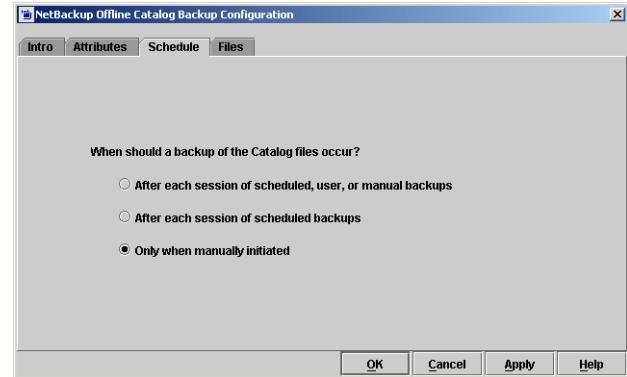
- ◆ Ensure that the disk has adequate space for the catalogs. If the disk fills up, the catalog backups will fail.
- ◆ Ensure that the path is a directory rather than a file. If the path is a file, an error occurs when the backup is done (*not* when you specify the path).
- ◆ The following rule applies to the path you specify:

In addition to the platform-specific file path separators (/ and \) and colon (:) within a drive specification on Windows, follow the NetBackup naming conventions. (See “NetBackup Naming Conventions” on page 29.)

Catalog Schedule Tab

The Catalog **Schedule** tab contains selections concerning when you want to back up the catalogs.

Caution It is essential that you back up your catalogs often. If these files are lost, you lose information about backups and configuration changes that were made between the time of the last catalog backup and the time that the disk crash occurred.



- ◆ After each session of scheduled, user, or manual backups

Backs up the catalogs after any session that results in the creation of at least one successful backup or archive. This includes automatic, manual, and user backups.

- ◆ After each session of scheduled backups

Backs up the catalogs after any automatic backup session that results in at least one successful backup of a client. A backup *does not* occur after a manual backup or a user backup or archive.


This schedule type requires that no jobs are running in order for the catalog backup to run.

Caution Do not use this schedule in NetBackup environments in which continual backup activity is typically occurring. It could mean that the catalog backup would never have an opportunity to run.



- ◆ Only when manually initiated

Does not automatically back up the catalogs. If you elect to back up catalogs manually, select **NetBackup Management > Catalog**. Right-click **Catalog** and select **Back up NetBackup Catalog**.

Caution If you elect to back up catalogs manually, be certain to do so once a day or after every series of backups.




Catalog Protection



Catalog Files Tab

The Catalog **Files** tab contains the absolute pathnames to the catalog files to be backed up.

The pathnames of the catalogs on the master server are automatically added during installation and generally require no action on your part other than to ensure they are listed.



Click to navigate to a directory or file

In the case of NetBackup Enterprise Server, however, where the master server and media servers may reside on different machines, the pathnames to the NetBackup database on the media servers are *not* automatically added during installation and require that you add them to the file list.

The `NETBACKUP_RELATIONAL_DATABASE_FILES` directive automatically includes the database files in the `/usr/opensv/db/data/` directory, as well as `vxdbms.conf`, `server.conf`, and `databases.conf`. (If the files have been relocated to different directories using `nbd_b_move`, the locations will be automatically determined and the files will be included.)

Note The file names in the Catalog Files Location list are case-sensitive. The catalog backup will fail if the entries are typed without regard to case.

On the **Files** tab:

- ◆ To add a pathname, click **New** and type the pathname in the **Catalog File Locations** list.

The pathname format depends on whether the catalog is on a master server or a remote media server. It also depends on whether the backup is sent to the master server or to a remote media server.

- ◆ To change a pathname, select the pathname and click **Change**. Edit the pathname. Click outside the edit box to escape the edit box.
- ◆ To delete a pathname, select the pathname you wish to delete and click **Delete**.

Catalog Protection

Notes Regarding the Catalog File Locations List

- ◆ If using NetBackup Access Control (NBAC) in the NetBackup configuration, add the following directives for each host in the NBAC domain:

```
[host:]nbat
```

```
[host:]nbaz
```

Note If the master server using NBAC is a UNIX machine, VERITAS recommends that you do not include the NetBackup master server configuration file (`/usr/openv/netbackup/bp.conf`) in the offline catalog backup file list. If `bp.conf` is included in the list, it must not be recovered until all other catalog recovery is completed.

- ◆ Make sure there are no invalid paths in the list of catalog files to be backed up, especially if you've moved catalog files, deleted old paths, or added new paths to the catalog backup configuration. If NetBackup cannot find or follow a path, the entire catalog backup fails.
- ◆ On UNIX systems, if `/usr/openv/netbackup/db/images` is a symbolic link to another filesystem, you *must* specify the true location of the images directory here. Symbolic links do not apply to Windows.
- ◆ Do not specify a symbolic link as the final component in a UNIX path or the entire catalog backup will fail.
While NetBackup follows links at other points in the path, NetBackup does not follow a link when it is the final component. If any other part of a listed path is a symbolic link, NetBackup saves the actual path during the backup.
- ◆ The files that are associated with the NetBackup relational databases, NBDB and BMRDB, are automatically included in the catalog backup and don't need to be listed here. Since the locations of these files can change, the pathnames are dynamically generated during the backup.

Absolute Pathnames for Catalogs on the Master Server

The pathnames of the catalogs on the master server are automatically added during installation and, unless you are backing up the catalogs to a media server, require no action on your part other than to ensure they are listed.

The files in the following directory have NetBackup scheduling information, error logs, and all information about files backed up from client workstations:

```
/usr/openv/netbackup/db
```

The files in the following directory contain license key and authentication information:

```
/usr/openv/var
```

If you are backing up the catalogs to a media server, prefix each pathname with the name of the media server.

Absolute Pathnames for Catalogs on Media Servers

If you are backing up catalog files that are on media servers, prefix each pathname with the name of the media server:

```
media_server_name:catalog_backup_path
```

The paths that you must add depend on the version of NetBackup installed on the media server.

- ◆ For UNIX media servers running NetBackup versions earlier than 6.0, add the following two paths:

- ◆ *media_server_name:/usr/opensv/netbackup/db/media*

The files in this directory have information about files that were backed up or archived from client workstations.

- ◆ *media_server_name:/usr/opensv/volmgr/database*

The files in this directory have information about the media and devices being used in the configuration.

- ◆ For UNIX media servers running NetBackup version 5.x, also include the following path:

```
media_server_name:/usr/opensv/var
```

For example, to add the paths for a UNIX NetBackup 5.x media server named *elk*, create the following entries:

```
elk:/usr/opensv/netbackup/db/media  
elk:/usr/opensv/volmgr/database  
elk:/usr/opensv/var
```

Backups of 6.0 Media Servers

For NetBackup 6.0 media servers, the critical media and device data is stored in the NetBackup relational database, NBDB, on the EMM server (generally, the master server). No pathnames should be added for 6.0 media servers to be included in the cold catalog backup.

Paths For Windows NetBackup Media Servers

If you are backing up catalogs that are on Windows NetBackup media servers, prefix each path name with the name of the media server:

```
media_server_name:catalog_backup_path
```

Catalog Protection

For example, to add the paths for a Windows NetBackup 5.x named *mars*, create the following entries (where *Install_path* is the directory where NetBackup is installed):

```
mars:C:Install_path\NetBackup\db  
mars:C:Install_path\Volmgr\database  
mars:C:Install_path\NetBackup\var
```

The files in the `db` directory have NetBackup error logs and all information about files backed up from client workstations.

The files in the `database` directory have information about the media and devices used in the configuration.

Note Remember to use the backslash (\) in the pathnames for a Windows NetBackup server.

Recovering the Catalog

The method used to recover the catalog in a disaster recovery situation depends on the method used to back up the catalog. Catalog recovery from online, hot catalog backups and offline, cold catalog backups are discussed in the *NetBackup Troubleshooting Guide*.

Disaster Recovery E-mails and the Disaster Recovery File

While using the Catalog Backup Wizard to configure an online, hot catalog backup, you're asked whether you'd like the disaster recovery information sent to an e-mail address. If the online catalog backup is configured using the Policy utility, this information appears on the Disaster Recovery tab. (See "Where Will the Catalog Data Be Located: Disaster Recovery Tab" on page 193.)

The disaster recovery e-mail and the accompanying attachment that is sent contain 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 included as an attachment.

If an online, hot catalog backup policy included both full and incremental backups, the attached image file may be a full or an incremental catalog backup. Recovering from an incremental backup will completely recover the entire catalog if you select **Automatically recover the entire NetBackup catalog** on the wizard screen because it references information from the last full backup. It is not necessary to first recover the last full catalog backup, then subsequent incremental backups.

Archiving the Catalog

The catalog archiving feature helps users tackle the problems that large amounts of catalog data can pose: large catalogs require a greater amount of disk space and can be time-consuming to back up. Catalog archiving reduces the size of online catalog data by relocating the large catalog .f files to secondary storage. NetBackup administration will continue to require regularly scheduled catalog backups, but without the large amount of online catalog data, the backups will be faster.

Catalog archiving is available on both UNIX and Windows platforms.

Note When considering whether to archive .f files, note that additional time is required to mount the tape and perform the restore.

Catalog Archiving Process

The following section describes the steps to archive a catalog. Catalog archiving operations must be performed when NetBackup is in an inactive state (no jobs are running).

1. Create a policy named *catarc* to reflect that the purpose of the schedule is for catalog archiving. (See “Creating a Catalog Archiving Policy” on page 249.)
2. Run `bpcatlist` to display images available for archiving.

Note Running `bpcatlist` alone will not modify any catalog images. Only when the `bpcatlist` output is piped to `bpcatarc` and `bpcatrm` will the images be modified and the image .f files removed.

To determine what images are available for catalog archiving, run:

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -online
```

To determine what images have been previously archived run:

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -offline
```

Note This should return a message of *No entity was found* if catalog archiving has not been previously run.

3. Once the `bpcatlist` output correctly lists all the images to be archived, pipe the output through `bpcatarc` and `bpcatrm`. For example:

```
bpcatlist -client all -before Jan 1 2005 | bpcatarc | bpcatrm
```

The command waits until the backup completes successfully before returning the prompt. An error is reported if the catalog archive fails.

The Activity Monitor displays a Job ID for the job. The File List for the job (double-click the job in the Activity Monitor) displays a list of image files that have been processed. When the job completes with a status 0, `bpcatrm` removes the corresponding `.f` files. If the job fails, no catalog `.f` files are removed.

4. To restore the catalog archive:

- a. Use `bpcatlist` to list the files that need to be restored.
- b. Once `bpcatlist` displays the proper files to restore, run `bpcatres` to restore the actual files.

To restore all the archived files from Step 2 above, run:

```
bpcatlist -client all -before Jan 1 2005 | bpcatres
```

This command restores all the catalog archive files prior to Jan 1, 2005.

For more information on the archiving commands, see “Catalog Archiving Commands” on page 251.

Creating a Catalog Archiving Policy


The catalog archiving feature requires the presence of a policy named *catarc* in order to have catalog archiving commands run properly. The policy can be reused for catalog archiving.

Policy Name

Create a new policy named *catarc* that waits until `bpcatarc` activates it. This policy is not run by a user. Instead, `bpcatarc` activates this special policy to perform a catalog backup job, then deactivates the policy after the job is done.

Deactivate Policy

The catalog archive policy must be deactivated. On the Attributes tab, clear the **Active** field.



Archiving the Catalog

Type of Backup

The type of backup indicated for the catalog archive policy must be *User Backup*. This is set in the Change Schedule dialog on the Attributes tab.

Retention Level Setting

Since it may not be necessary to set an infinite retention level, you should be certain to 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.

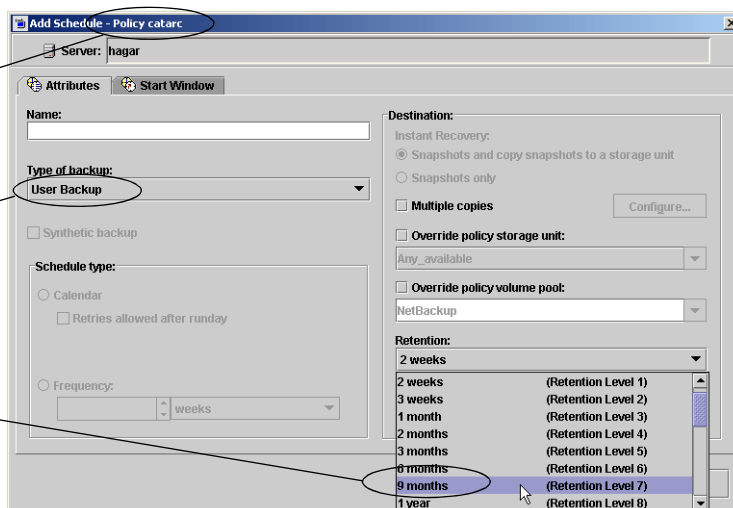
Note Failure to set the retention level of the catalog archive for a time at least as long as the longest retention period of the backups being archived can result in the loss of catalog data.

You may find it useful to set up, then designate a special retention level for catalog archive images.

Policy name must be titled:
catarc

Policy schedule must be
User Backup type

Set to a period of at least
as long as the longest
retention period of backups
being archived



Schedule

A schedule is required for *catarc*. The schedule for *catarc* must include in its window the time *bpcatarc* command is being run. If *bpcatarc* is run outside of the schedule indicated in *catarc*, the operation will fail.

Files

On the Files tab, browse to the directory where catalog backup images are placed:

```
/usr/opensv/netbackup/db/images
```

Clients

On the Clients tab, enter the name of the master server.

Catalog Archiving Commands

The catalog archiving feature relies on three commands to first designate a list of catalog .f files, then archive the files. A fourth command, `bpcatres`, is used to restore the files if necessary.

Create a Catalog List with `bpcatlist`

The `bpcatlist` command queries the catalog data, then lists portions of the catalog based on selected parameters, such as date, client, policy, schedule name, backup ID, the age of the backup image, or the date range of the backup image. `bpcatlist` outputs the formatted image summary information of matched images to standard output.

The other catalog archiving commands, `bpcatarc`, `bpcatrm`, and `bpcatres`, all depend on input from `bpcatlist` via a piped command.

For example, to archive (backup and delete) all of the .f files created prior to January 1, 2005, the following would be entered:

```
# bpcatlist -client all -before Jan 1 2005 | bpcatarc | bpcatrm
```

`bpcatlist` is also used to provide status information. For each catalog, it lists the following information:

- ◆ Backup ID (Backupid)
- ◆ Backup date (Backup Date)
- ◆ Catalog archive ID (Catacid). After an .f file is successfully backed up, a catalog archive ID is entered into the catacid field in the image file.
- ◆ Online status (S), indicating if the catalog is online (1) or deleted from the online media and stored on other media (0)
- ◆ Compressed status (C), indicating if the catalog is compressed (1) or not compressed (0)
- ◆ Catalog file name (Files file)

The following is an example of the `bpcatlist` output, showing all of the backups for client alpha since October 23:

```
# bpcatlist -client alpha -since Oct 23
Backupid      Backup Date      ...Catacid  S C Files file
alpha_0972380832 Oct 24 10:47:12 2000 ... 973187218 1 0 alpha_0972380832_UBAK.f
```

Archiving the Catalog

```
alpha_0972336776 Oct 23 22:32:56 2000 ... 973187218 1 0 alpha_0972336776_FULL.f
alpha_0972327197 Oct 23 19:53:17 2000 ... 973187218 1 0 alpha_0972327197_UBAK.f
```

For detailed information on `bpcatlist`, see `bpcatlist` in *NetBackup Commands for UNIX and Linux*.

Back Up the Catalog with `bpcatarc`

The `bpcatarc` command reads the output from `bpcatlist` and backs up the selected list of `.f` files. After an `.f` file is successfully backed up, a catalog archive ID is entered into the `catarcid` field in the image file. For archiving of the `.f` files to proceed, a policy named *catarc*, based on a User Backup type schedule is required. The schedule for *catarc* must include in its window the time `bpcatarc` command is being run. (See “Creating a Catalog Archiving Policy” on page 249.)

Remove the Catalog with `bpcatrm`

The `bpcatrm` command reads the output from `bpcatlist` or `bpcatarc` and deletes selected image `.f` files from the online catalog if the image file has valid `catarcid` entries.

`bpcatrm` does not remove an `.f` file unless the file has been previously backed up using the *catarc* policy.

Restore the Catalog with `bpcatres`

The `bpcatres` command reads the output from `bpcatlist` and restores selected archived `.f` files to the catalog. For example:

```
# bpcatlist -client all -before Jan 1 2000 | bpcatres
```

Recommendations for Using Catalog Archiving

- ◆ Perform catalog archiving operations when NetBackup is in an inactive state (no jobs are running).
- ◆ To ensure that catalog backup images are not on the same tapes as user backups, create a separate media pool for catalog archives.
- ◆ You may find it useful to set up, then designate, a special retention level for catalog archive images.

To specify retention levels, go to **Host Properties > Master Server > Retention Periods** or see “Retention Periods Properties” on page 436.

Using Vault with the Catalog Archiving Feature

Since the catalog archiving feature uses a regular User Backup schedule in the *catarc* policy, Vault duplicates and vaults the files no differently from other backups.

Browsing Offline Catalog Archive

If a user tries to browse an offline catalog, the user will receive an error message stating that the catalog image .*f* file has been archived. The catalog archiving feature is intended to be used by a NetBackup Administrator only. Use the *bplist* command to determine if a catalog .*f* file is archived.

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 separate archives 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:

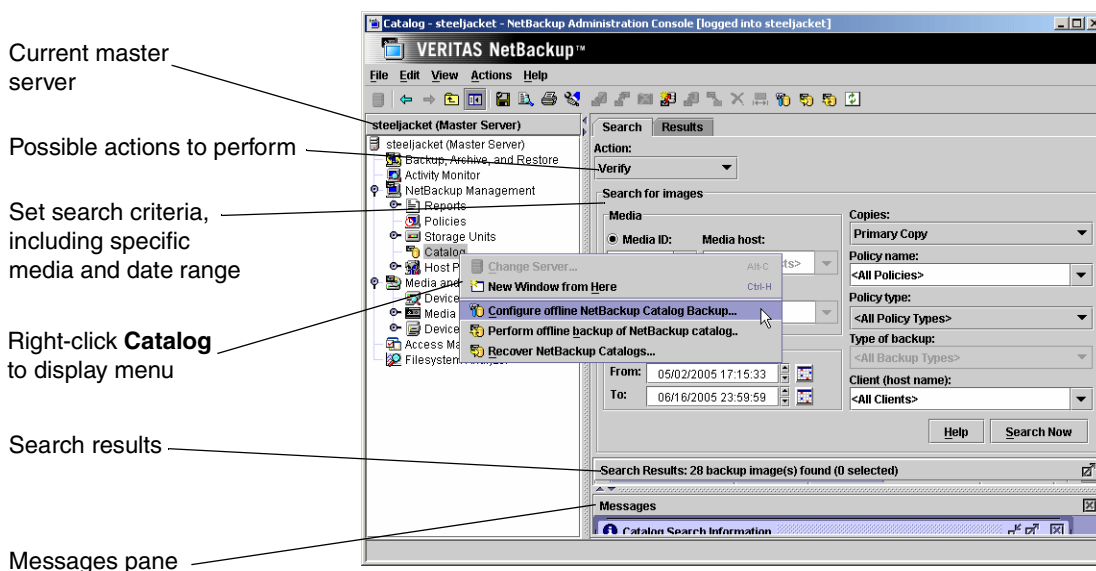
```
bpcatlist -client clientname | bpcatarc | bpcatrm
```
4. If you don't wish to write more images to that client's volume pool, change the volume pool before running another archiving a catalog again.

Using the Catalog Utility

Use the **Catalog** utility to create and configure *catalog backups*, 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 other operations on catalog information. For example:

- ◆ Search for backup images in order 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.



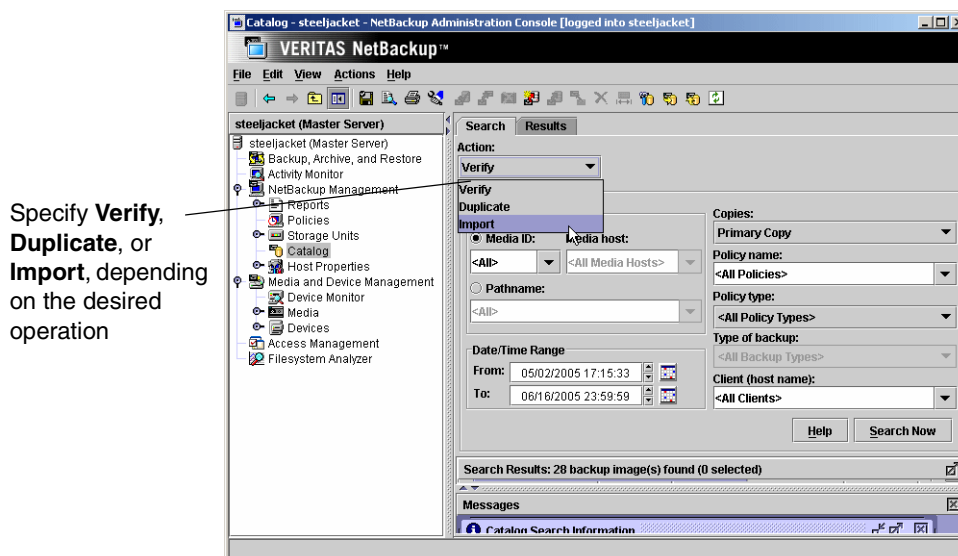
Searching for Backup Images

Use the **Catalog** utility to search for a backup image. You may want to search for a backup image in order to:

- ◆ 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.



Search for backup images using the criteria described in the following table:

Search Criteria for Backup Images

Search Criteria	Description
Action	Select the action that was used to create the image for which you're looking: Verify, Duplicate, Import .
Media ID	The media ID for the volume that contains the desired backups. Type a media ID in the box or select one from the scroll-down list. To search on all media, select <All> .
Media Host	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 .

Using the Catalog Utility

Search Criteria for Backup Images (continued)

Search Criteria	Description
Pathname	To search for an image on a disk storage unit, select Pathname and specify the file path that includes the originals.
Date/time range	The range of dates and times that includes all the backups for which you want to search. The default range is determined by the Global Attributes property, Policy Update Interval . (See "Global Attributes Properties" on page 414.)
Copies	The source you want to search. From the scroll-down list, select either Primary or the copy number.
Policy Name	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)	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	The type of schedule that created the backups for which you are searching. 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 .

Notes on Searching for an Image

When searching for specific kinds of images, note the following:

- ◆ **Duplication image:** If the original is fragmented, NetBackup duplicates only the fragments that exist on the specified volume.
- ◆ **Verification image:** Backups that have fragments on another volume are included, as they exist in part on the specified volume.
- ◆ **Import image:** If a backup begins on a media ID that has not been processed by the first step of "To initiate an import – Phase I" on page 266 it is not imported. If a backup ends on a media ID that has not been processed by the first step of "To initiate an import – Phase I" on page 266, the imported backup is incomplete.

Messages Pane

The **Messages** pane displays messages about a task running as a background process. The pane is displayed only if there is an informative message or error message for the task. If the task completes normally, the pane is not displayed. 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.

Although this operation does not compare the data on the volume with the contents of the client disk, it does read each block in the image to verify that the volume is readable. (However, data corruption within a block could be possible.) NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

▼ To verify backup images

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Set up the search criteria for the image you wish to verify as explained in the “Search Criteria for Backup Images” table. Click **Search Now**.
3. Select the image you wish to verify and select **Actions > Verify**. The Confirm Verify dialog may appear.
To display information on each file that NetBackup verifies, select **Log all files found in verified image(s)**.
4. Click the **Results** tab, then select the verification job just created to view the job results.

Viewing Job Results

The results of verify, duplicate, or import jobs appear in the **Results** tab. The top portion of the dialog displays all existing log files.

To view a log file, select the name of the log from the list. The log file currently displayed appears in the bottom portion of the **Results** dialog. If an operation is in progress, the log file display is refreshed as the operation proceeds.

Using the Catalog Utility

▼ To view or delete a log file

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Click the **Results** tab.
3. Select a log file.

4. Select **View > Full View** to display the entire log file in a screen editor.

Select **Edit > Delete** to delete the log.

You can also right-click the log file and select an action from the scroll-down menu.

Promoting a Copy to a Primary Copy

Each backup is assigned a *primary copy*. NetBackup uses the primary copy to satisfy restore requests. The first backup image created successfully by a NetBackup policy is the primary backup. If the primary copy is unavailable and you have created a duplicate, 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 that remains in the robot should be the primary backup. When a primary backup expires, the next backup (if it exists) is promoted to primary automatically.

▼ To promote a backup copy to a primary copy

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Set up the search criteria for the image you wish to promote to a primary copy. (See “Searching for Backup Images” on page 254.) Be sure that you’ve indicated a copy in the **Copies** field and not **Primary Copy**. Click **Search Now**.
3. Select the image you wish to promote.
4. Click **Actions > Set Primary Copy**.

After promoting to the primary copy, the Primary Status column immediately reads **Yes**.

▼ To promote many copies to a primary copy

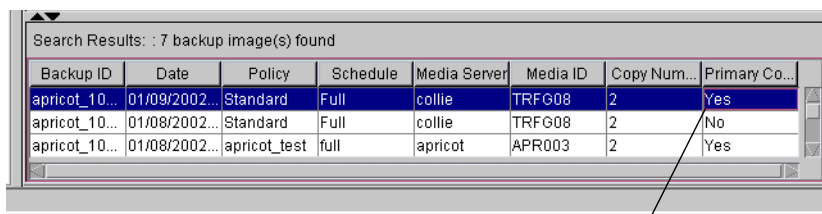
You can also promote many copies to be a primary copy using the `bpchangeprimary` command. For example, the following command will promote all copies on media belonging to the volume pool, SUN, created after 08/01/2002 to be the primary copy:

```
bpchangeprimary -pool SUN -sd 08/01/2002
```

The following command will promote copy 2 of all backups of client oak, created after 01/01/2002 to be the primary copy:

```
bpchangeprimary -copy 2 -cl oak -sd 01/01/2002
```

For more information on `bpchangeprimary`, see *NetBackup Commands for UNIX and Linux*.



Backup ID	Date	Policy	Schedule	Media Server	Media ID	Copy Num...	Primary Co...
apricot_10...	01/09/2002...	Standard	Full	collie	TRFG08	2	Yes
apricot_10...	01/08/2002...	Standard	Full	collie	TRFG08	2	No
apricot_10...	01/08/2002...	apricot_test	full	apricot	APR003	2	Yes

Primary Copy status indicates that the image is now the primary copy

▼ To promote a backup copy to a primary copy using `bpduplicate`

1. Enter the following command:

```
/usr/openv/netbackup/bin/admincmd/bpduplicate -npc pcopy -backupid bid
```

Where:

pcopy is the copy number that will become the new primary copy.

bid is the backup identifier as shown in the Images on Media report.

To find the volume that has the duplicate backup, use the Images on Media report. Specify the backup ID which is known (and also the client name if possible to reduce the search time). The report shows information about both copies. (See “Images on Media Report” on page 303.)

The `bpduplicate` command writes all output to the NetBackup logs so nothing appears in the command window.

After promoting the duplicate to the primary copy, use the Backup, Archive and Restore interface on a client to list and restore files from the backup. For instructions, see the online help in the Backup, Archive, and Restore client interface.

Duplicating Backup Images

NetBackup can create up to 10 copies of unexpired backups. Indicate the number of backup copies in **Host Properties > Master Servers > Global Attributes > Maximum backup copies**. (See “Global Attributes Properties” on page 414.)

NetBackup can create up to four of the copies simultaneously.

An alternative to taking time to duplicate backups is to make multiple copies concurrently. (This feature is sometimes referred to as Inline Tape Copy.) You can create up to four copies simultaneously at backup time. Keep in mind that an additional drive is required for each copy and the destination storage units cannot be optical disk, NDMP, QIC, or third-party copies. The backup time may be longer than for one copy only.

NetBackup does not verify in advance whether the storage units and drives required for the duplicate operation are available for use, only that the destination storage unit exists. The storage units must be connected to the same media server.

The following lists describe scenarios which present candidates for duplication and scenarios where duplication is not possible:

Possible to duplicate backups:

- ♦ 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. This is done with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

Not possible to duplicate backups:

- ♦ while the backup is being 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)
- ♦ of offline NetBackup catalogs.
- ♦ when it is a multiplexed duplicate of the following:
 - FlashBackup
 - NDMP backup
 - Backups from disk type storage units
 - Backups to disk type storage units
 - Nonmultiplexed backups

Note Do not duplicate images while a NetBackup catalog backup is running. This results in the catalog backup not having information about the duplication.

Notes on Multiplexed Duplication

- ◆ When duplicating multiplexed backups, 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 will have a multiplexing factor that is no greater than that used during the original backup.
- ◆ If all backups in a multiplexed group are duplicated to a storage unit that has the same characteristics as the one where the backup was originally performed, the duplicated group will be identical, with the following exceptions:
 - ◆ If EOM (end of media) is encountered on either the source or destination media.
 - ◆ If any of the fragments in the source backups are zero length (occurs if many multiplexed backups start at the same time), then during duplication these zero length fragments are removed.

Procedure for Duplicating Backups

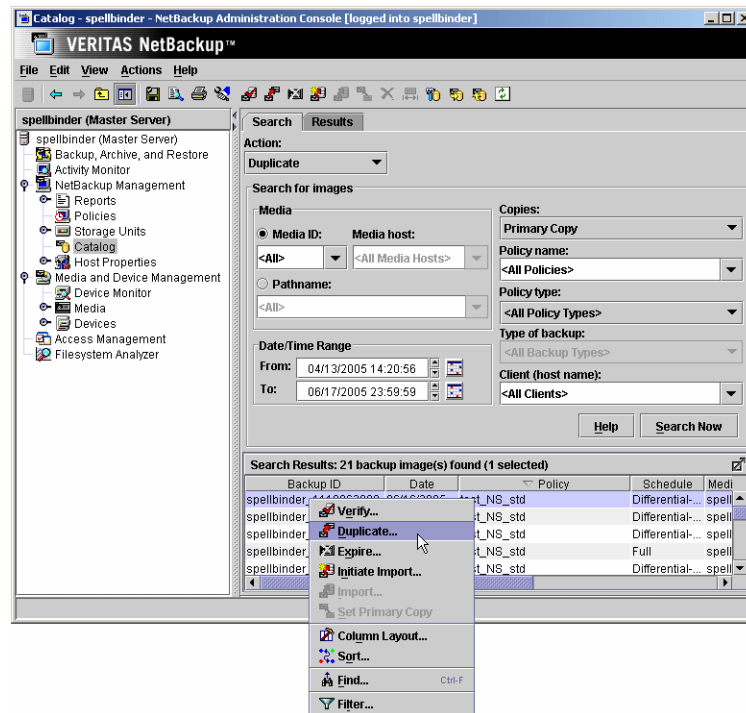
NetBackup can create up to 10 copies of unexpired backups.

▼ To duplicate backup images

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Set up the search criteria for the image you wish to duplicate. Click **Search Now**.
3. Right-click the image(s) you wish to duplicate and select **Duplicate** from the shortcut menu. The **Setup Duplication Variables** dialog appears.

Note If duplicating an online, hot catalog backup, be sure to select all child jobs that were used to create the catalog backup. All jobs must be duplicated in order to duplicate the catalog backup.

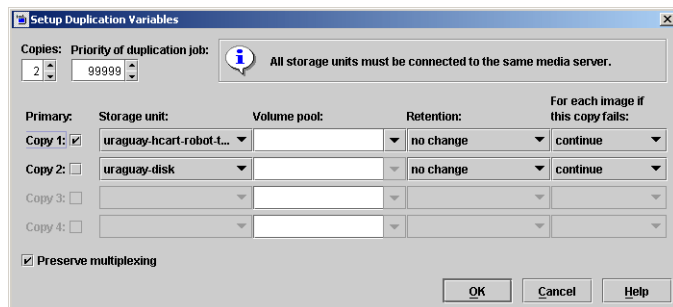
Using the Catalog Utility



4. Specify the number of copies you would like created.

If there are enough drives available, the copies will be created simultaneously.

Otherwise, the system may require operator intervention if, for instance, four copies are to be created and there are only two drives.



5. The primary copy is the copy from which restores will be done. Normally, the original backup will be 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 one chosen is the one with the smallest *copy number*. If the primary is copy 1, when it expires, copy 2 becomes primary. If the primary is copy 5, when it expires, copy 1 becomes primary.)

6. Specify the storage unit where each copy will be stored. If a storage unit has multiple drives, it can be used for both the source and destination.

Note The ability to make multiple copies concurrently (sometimes referred to as Inline Tape Copy) is not supported on the following storage types: NDMP, third-party copies, or optical devices.
Also, this ability is not supported on storage units that use a QIC (quarter-inch cartridge) drive type.

7. Specify the volume pool where each copy will be stored. The volume pool selections are based on the policy type setting that was used for the query:
 - ◆ If the policy type was set to query for *All Policy Types* (default), both catalog and non-catalog volume pools are included in the drop-down list.
 - ◆ If the policy type was set to query for *NBU-Catalog*, only catalog volume pools are included in the drop-down list.
 - ◆ If the policy type was set to query for any other policy type than *NBU-Catalog* and *All Policy Types*, only non-catalog volume pools are included in the drop-down list.

NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as the media ID of the volume that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure 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, such as elapsed time, apply only to the primary. It is the primary copy that NetBackup uses to satisfy restore requests.

- ◆ If *No Change* is selected for the retention period, the expiration date is the same for the duplicate and source copies. You can use the `bpexpdate` 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, 2002 and its retention period is one week, the new copy's expiration date is November 21, 2002.

Using the Catalog Utility

9. Specify whether the remaining copies should continue or fail if the specified copy fails.
10. 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 will have 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.


Preserve Multiplexing does not apply when the destination is a disk storage unit. However, if the source is a multiplexed tape and the destination is a disk storage unit, selecting **Preserve Multiplexing** ensures that the tape is read in only one pass rather than multiple passes.
11. Click **OK** to start duplicating.
12. Click the **Results** tab, then select the duplication job just created to view the job results. (See “Viewing Job Results” on page 257.)

Jobs Displayed While Making Multiple Copies Concurrently

When making multiple copies concurrently (sometimes referred to as Inline Tape Copy), either at backup time or duplication, a parent job plus a job for each copy is displayed.

The parent job displays the overall status, whereas the copy jobs display the status of the copy. This enables you to troubleshoot a problem 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 will be 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 with two copies. The parent job is 579, copy 1 is job 580, and copy 2 is job 581. Copy 1 finished successfully, but copy 2 failed with a 98 status (error requesting media). Since at least one copy finished successfully, the parent job shows a successful (0) status.



Copy 1 was successful,
but Copy 2 failed

Since at least one copy was successful,
the parent job was successful

Importing NetBackup or Backup Exec Images

NetBackup can import backups that have expired, backups from another NetBackup server, or backups that were written by Backup Exec for Windows. (Supports Backup Exec versions 7.0 through 9.1.)

During an import operation, NetBackup recreates NetBackup catalog entries for the backups on the imported volume. Importing is useful for moving volumes from one site to another and for recreating NetBackup catalog entries.

NetBackup supports importing and restoring the following Backup Exec backup types. Please refer to the specific NetBackup manuals for details on the support for each backup type.

- ◆ Windows
- ◆ UNIX
- ◆ NetWare
- ◆ Exchange
- ◆ SQL

NetBackup does not support reading Backup Exec media written by Backup Exec for NetWare.

Importing images is accomplished in two phases. (Importing Backup Exec media requires one additional step):

Note If importing Backup Exec media, run `vmphyinv` to update the Backup Exec media GUID in the NetBackup Media Manager database. This needs to be done only once after creating the media IDs in the NetBackup Media Manager database. See *NetBackup Commands for Windows* for more information about using `vmphyinv`.

- ◆ Phase I:
NetBackup creates a list of expired catalog entries for the backups on the imported volume. No actual import occurs in Phase I.
- ◆ Phase II:
Images are selected for importing from the list of expired images created in Phase I.

▼ To initiate an import – Phase I

The result of initiating Phase I of the import process is to create a list of expired images from which to choose to import in Phase II. No import occurs in Phase I.

1. If importing Backup Exec media, run `vmphyinv` to update the Backup Exec media GUID in the NetBackup Media Manager database. This needs to be done only once after creating the media IDs in the NetBackup Media Manager database. See *NetBackup Commands for UNIX and Linux* for more information about using `vmphyinv`.
2. If importing the images from tape, make the media accessible to the media server from which the images will be imported.
3. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
4. Select **Actions > Initiate Import**. The Initialize Import dialog appears.

- ◆ The **Master Server** field indicates the master server to which you are importing the images.
- ◆ In the **Media Host** field, specify the name of the host that contains the volume you are going to import.
- ◆ Select whether the images to be imported are located on tape or on disk.

If images are on tape:

- ◆ In the **Media ID** field, type the Media ID of the volume that contains the backups you are importing.
- ◆ Check whether or not you're importing password-protected Backup Exec images. Validate the Backup Exec password by typing it again in the field provided.

If images are on disk:

- ◆ Enter the path to the images in the field provided.
- ◆ If the image is on a NearStore server, check the NearStore checkbox and enter the name of the NearStore server in the field provided.


Note When importing from Backup Exec media, if the media is password protected and the user does not provide the password or if an incorrect password is provided, the job fails with an appropriate error, and the logs indicate that either no password, or a wrong password, was provided. If the media is not password protected and the user provides a password, the password provided by the user is ignored.

Using the Catalog Utility

If the password contains non-ASCII characters, the media can be imported only by using the NetBackup Administration Console on Windows or by using the `bpiimport` command. The NetBackup-Java Administration Console cannot be used.

Note If importing an online, hot catalog backup, be sure to import all child jobs that were used to create the catalog backup. All jobs must be imported in order to import the catalog backup.

Click **OK**. The Confirm Initiate Import dialog appears.



5. Click **OK** to start the process of reading the catalog information from the source volume.
6. Click on the **Catalog Results** tab to see NetBackup look at each image on the tape and determine whether or not it has expired and can be imported. The job also displays in Activity Monitor as an Import type. Select the import job log just created to view the job results.

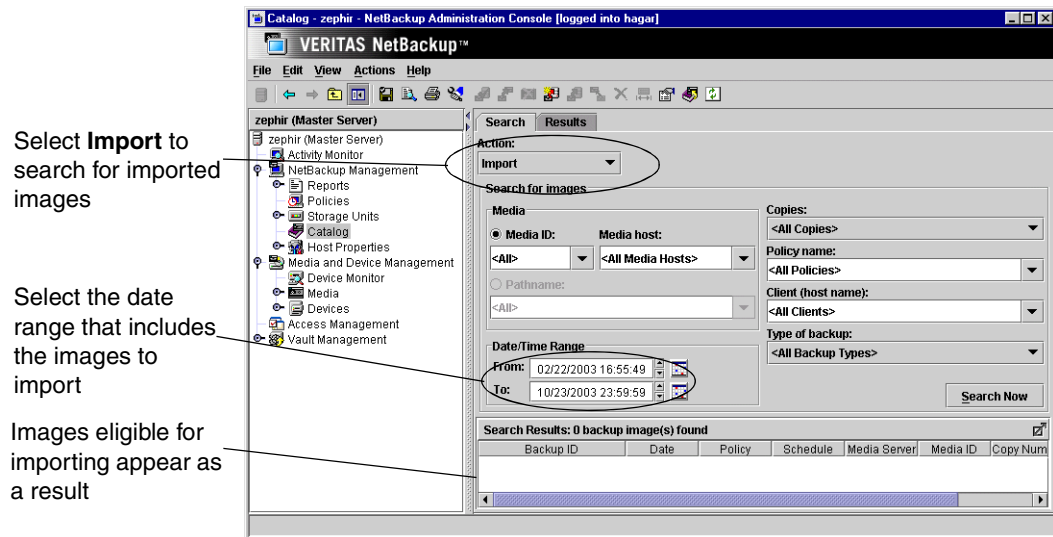
Note Since it is necessary to mount and read the tape at this phase, reading the catalog and building the list can take some time to complete.

▼ To import backup images – Phase II

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.

Note When importing backups that have fragments on multiple tapes, do not start the Import (Phase II) until you have run the Initiate Import (Import Phase I) to read the catalog for all the tapes containing fragments. If this is not done, the import will fail with a message such as: *Unexpected EOF or Import of backup id failed, fragments are not consecutive*.

- Set up the search criteria to find imported images by setting the search action to **Import**. Be sure to select a date range that includes the images you want to import.



- Select the image(s) you wish to import and Select **Actions > Import**. The Confirm Import dialog appears.
- To view the log, click the **Results** tab, then select the import job log just created.

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, 2004 and its retention period is one week, the new expiration date is November 21, 2004.

Notes About Importing Backup Images

- NetBackup can import images stored on disk in this release. However, importing disk images from previous versions of NetBackup, is not allowed.
- NetBackup cannot import disk images that span disks. To avoid spanning storage units, do not use Checkpoint Restart on a backup policy that writes to a storage unit or a storage unit group that contains multiple disk staging storage units.
- You cannot import a backup if an unexpired copy of it already exists on the server where you are trying to import it.
- NetBackup does not direct backups to imported volumes.

Using the Catalog Utility

- ◆ If importing an online, hot catalog backup, be sure to import all child jobs that were used to create the catalog backup. All jobs must be imported in order to import the catalog backup.
- ◆ To import from a volume that has the same media ID as an existing volume (for example A00001) on this server, first duplicate the existing volume to another media ID (for example, B00001). Then, remove information about the existing media ID that is causing the problem (in this example, A00001) from the NetBackup catalog by running the following command:

```
/usr/openv/netbackup/bin/admincmd/bpexpdate -d 0 -m media ID
```

Next, delete the existing media ID that is causing the problem (in this example, A00001) from Media Manager on this server. Finally, add the volume you are importing (the other A00001) to Media Manager on this server. The *Media Manager System Administrator's Guide* contains instructions for deleting and adding volumes.

To avoid this problem in the future, use unique prefix characters for media IDs on all servers.

Importing Images from Backup Exec Media

The clients who will be performing the restore operation must be at NetBackup 5.0 or later. Backup Exec images cannot be restored to clients on platforms not supported by NetBackup versions prior to 5.0.

Host Properties for Backup Exec

The Backup Exec UNIX agent identifies itself to the Backup Exec server using a GRFS advertised name. The advertised name may not have been the same as the real machine name and path.

NetBackup must know what the advertised name is, along with the actual client name and path in order to create accurate .f file paths.

This is done by setting the **GRFS Advertised Name**, **Actual Client**, and **Actual Path** properties in the Backup Exec Tape Reader host properties. If no entries are indicated, NetBackup assumes that the advertised name is the same as the real machine name and the advertised path is the same as the real path. (See "Backup Exec Tape Reader Properties" on page 352.)

Considerations Concerning Importing Backup Exec Media

The following items should be taken into consideration when importing Backup Exec media:

- ◆ It is not possible to restore UNIX data to Windows systems, Windows data to UNIX systems, Windows data to NetWare systems and UNIX data to NetWare systems.
- ◆ Importing from Backup Exec media does not convert or migrate Backup Exec job history, job schedules, or job descriptions to NetBackup.
- ◆ Importing from Backup Exec media does not convert Backup Exec application setup or configuration information to NetBackup.
- ◆ Intelligent Disaster Recovery (IDR) operations using the NetBackup IDR wizard and Backup Exec media is not supported. This includes both local and remote IDR restores.
- ◆ It is not possible to restore Backup Exec backups taken with the Intelligent Image Option.
- ◆ If Backup Exec hardlink backups are redirected and restored to partitions or drives other than the source partition or drive, the hardlinks are not restored, even though the progress log shows that the hardlinks were restored successfully.

Differences Between Importing, Browsing and Restoring Backup Exec and NetBackup Images

There are some differences between Backup Exec and NetBackup when importing, browsing, and restoring images:

Running `vmphyinv`

Importing Backup Exec media requires running `vmphyinv` to update the Backup Exec media GUID in the NetBackup Media Manager database. This needs to be done only once after creating the media IDs in the NetBackup Media Manager database and before running Phase I and Phase II import operations. See *NetBackup Commands for UNIX and Linux* for more information about using `vmphyinv`.

Importing and Restoring QIC Media

To import and restore Backup Exec Quarter Inch Cartridge (QIC) media written with tape block sizes more than 512 bytes, you must use a NetBackup Windows media server. A NetBackup UNIX media server will not work to import and restore the media in this case.

Spanned Media: Importing Differences

When importing a Backup Exec backup which spans multiple media, run a Phase 1 import on the first media of the spanned backup set. Then, run a Phase 1 import on the remaining media of the spanned backup set in any order.

Using the Catalog Utility

This differs from the NetBackup process, where Phase1 import can be run in any order in case the image spans multiple media.

SQL: Browsing and Restoring Differences

Backup Exec SQL images are browsed, then restored using the NetBackup Backup, Archive, and Restore client interface.

NetBackup SQL images are browsed, then restored using the NetBackup SQL interface.

File Level Objects: Browsing and Restoring Differences

When a user selects a Backup Exec file for restoring, the directory where that file is located will also get restored.

When a user selects a NetBackup file for restoring, only that single file is restored.

NetWare: Restoring Differences

NetBackup will not support restoring Backup Exec NetWare non-SMS backups created using the NetWare redirector.

Storage Management Services (SMS) software allows data to be stored and retrieved on NetWare servers independent of the file system the data is maintained in.

NTFS Hard Links, NTFS SIS Files, and Exchange SIS Mail Messages: Restoring

- ◆ When restoring Backup Exec NTFS images, any backed up directory with the name *SIS Common Store* will be restored, whether or not it is the actual NTFS single instance storage common store directory. This occurs even though the file was not specifically selected for restore.
- ◆ When restoring objects from backups which contain NTFS hardlinks, NTFS SIS files or Exchange SIS mail messages, additional objects, which the user did not select for restore, may be sent to the client. These additional objects will be skipped by the client and not restored. Although the objects which the user selected for restore are restored, the job is considered partially successful because some objects (though not selected by the user), were skipped.
- ◆ When redirecting NTFS hard links, NTFS SIS files or Exchange SIS mailboxes for restore:
 - ◆ All or some of the files should be redirected to any location on the source drive, or
 - ◆ all files should be redirected to a single location on a different drive.

For example, if the following hard link or SIS files are backed up:

```
C:\hard_links\one.txt
```

```
C:\hard_links\two.txt  
C:\hard_links\three.txt
```

Upon restore, some or all of the files can be redirected to any location on C:\, or all the files must be redirected to a different drive.

The following combination would be unsuccessful:

```
C:\hard_links\one.txt to a location on C:\  
C:\hard_links\two.txt to a location on D:\
```

If all the files are to be redirected to a different drive, specify that C:\ be replaced with D:\ in the redirection paths.

Unsuccessful: The redirection paths specify that C:\hard_links be replaced with D:\hard_links.

Successful: The redirection paths specify that C:\hard_links be replaced with C:\redir_hard_links.


Expiring Backup Images

To expire a backup image means to force the retention period to expire. When the retention period expires, NetBackup deletes information about the backup, making the files in the backups unavailable for restores without first re-importing.

▼ To expire a backup image

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Set up the search criteria for the image you wish to expire as explained in the table, “Search Criteria for Backup Images” on page 255. Click **Search Now**.
3. Select the image you wish to expire and select **Actions > Expire**.
4. A message appears telling you that once the backups have been expired, they cannot be used for restores.

Select **Yes** to proceed with expiring the image or **No**.



Catalog Maintenance and Performance Optimization

Determining Catalog Space Requirements

NetBackup requires disk space to store its error logs and information about the files it backs up. The maximum amount of disk space that NetBackup requires at any given time varies according to the following factors:

- ◆ Number of files that you are backing up
- ◆ Frequency of full and incremental backups
- ◆ Number of user backups and archives
- ◆ Retention period of backups
- ◆ Average length of full pathname 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 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.

“Example Reference Table for Catalog Requirements” shows that a full backup for policy S1 includes 64,000 files.

2. Determine the frequency and retention period of the full and incremental backups for each policy.
3. Use the information from steps 1 and 2 above to calculate the maximum number of files that exist at any given time.

For example:

Assume you schedule full backups every seven days with a retention period of four weeks and differential incremental backups daily with 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 plus one week’s worth of incrementals.

The following formula expresses the maximum number of files that can exist at any given time for each type of backup (daily, weekly, and so on):

Catalog Maintenance and Performance Optimization

$$\text{Files per Backup} \times \text{Backups per Retention Period} = \text{Max Files}$$

For example:

If a daily differential incremental schedule backs up 1200 files for all its clients and the retention period is seven days, the maximum number of files resulting from these incrementals that can exist at one time are:

$$1200 \times 7 \text{ days} = 8400$$

If a weekly full backup schedule backs up 3000 files for all its clients and the retention period is four weeks, the maximum number of files due to weekly-full backups that can exist at one time are:

$$3000 \times 4 \text{ weeks} = 12,000$$

Obtain the total for a server by adding the maximum files for all the schedules together. The maximum number of files that can exist at one time due to the above two schedules is the sum of the two totals, which is 20,400.

Note For 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 above calculation for the incremental from $1200 \times 7 = 8400$ to $3000 \times 7 = 21,000$. After adding 12,000 for the fulls, the total for the two schedules is 33,000 rather than 20,400.

4. Obtain the number of bytes by multiplying the number of files by the average length of the file's full pathnames and file information.

If you are unsure of the average length of a file's full pathname, use 100. Using the results from the examples in step 3 yields:

$$(8400 \times 150) + (12,000 \times 150) = 3060000 \text{ bytes (or about 2988 kilobytes)}$$

5. Add 10 to 15 megabytes to the total calculated in step 4. This is the average space for the error logs. Increase the value if you anticipate problems.
6. Allocate space so all this data remains in a single partition.

File Size Considerations

- ◆ Some UNIX systems have a large file support flag. Turn the flag ON to enable large file support. For example, AIX disables large file support by default, so the file size limit is 2GB.
- ◆ For UNIX systems, set the file size limit for the root user account to *unlimited* in order to support large file support.

Example: Estimating Required Catalog Backup Space

The following table shows backup schedules, retention times, and number of files for a group of example policies. By substituting the information from this table into the formula from step 3 above, we can calculate the maximum number of files for each policy. The following steps demonstrate this for policy S1:

1. Apply the following formula to policy S1:

Max Files equals:

$$\begin{aligned} & (\text{Files per Incremental} \times \text{Backups per Retention Period}) \\ & + \\ & (\text{Files per Monthly Full Backups} \times \text{Backups per Retention Period}) \end{aligned}$$

2. Substitute values from the following table:

$$1000 \text{ files} \times 30 + 64,000 \text{ files} \times 12 = 798,000 \text{ files}$$

Perform steps 1 and 2 for each policy. Adding the results together shows that the total number of files for all policies is:

$$4,829,600 \text{ files}$$

Multiply the total number of files by the bytes in the average path length and statistics (100 for this example). The total amount of disk space required for file paths is:

$$460.59 \text{ megabytes (1,048,576 bytes in a megabyte)}$$

Adding 15 megabytes for error logs results in a final uncompressed catalog space requirement of:

$$475.59 \text{ megabytes}$$

Example Reference Table for Catalog Requirements

Policy	Schedule	Backup Type	Retention	Number of Files
S1	Daily	Incremental	1 month	1000
	Monthly	Full	1 year	64,000
S2	Daily	Incremental	1 month	1000
	Monthly	Full	1 year	70,000
S3	Daily	Incremental	1 week	10,000
	Weekly	Full	1 month	114,000
	Monthly	Full	1 year	114,000

Catalog Maintenance and Performance Optimization

Example Reference Table for Catalog Requirements (continued)

Policy	Schedule	Backup Type	Retention	Number of Files
S4	Daily	Incremental	1 week	200
	Weekly	Full	1 month	2000
	Monthly	Full	3 months	2000
	Quarterly	Full	Infinite	2000
WS1	Daily	Incremental	1 month	200
	Monthly	Full	1 year	5600
WS2	Daily	Incremental	1 week	7000
	Weekly	Full	1 month	70,000
	Monthly	Full	1 year	70,000

Backing Up Catalogs Manually

Catalog backups typically run automatically because they are configured to do so using one of the catalog backup methods (online, hot or offline, cold). However, a catalog backup can be started manually.

Starting a backup manually is useful in the following situations:

- ◆ To perform an emergency backup. For instance, if problems are anticipated or if the system is being moved and there is no time to wait for the next scheduled catalog backup.
- ◆ If you have only one standalone drive and no robots or tape stacker and are using the standalone drive for catalog backups. In this situation, automatic backups are not convenient because the catalog backup tape must be inserted before each catalog backup and removed when the backup is done. (The tape swapping is necessary because NetBackup does not mix catalog and regular backups on the same tape.)

For directions on starting each catalog method, see “To perform a manual online, hot catalog backup” and “To perform a manual offline, cold catalog backup” on page 279.

▼ To perform a manual online, hot catalog backup

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Select the catalog backup policy you want to run.


3. Select **Actions > Manual Backup**. For more information, see “Performing Manual Backups” on page 196.

Note You can also run the `bpbbackup` command from the command line to perform an online, hot catalog backup. See *NetBackup Commands for UNIX and Linux* for more information on `bpbbackup`.

▼ To perform a manual offline, cold catalog backup

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Select **Actions > Backup NetBackup Catalog** to start the backup. The Backup NetBackup Catalog dialog appears.

The backup is saved to the least recently used of Media 1 and Media 2.
3. Select the master server for which you wish to create a catalog backup and click **OK**.



Note If the volume for the catalog backup is not in a drive, a mount request occurs and all catalog backups must wait for the mount before they can proceed. For a scheduled catalog backup, all other backups must wait until the catalog backup is complete.

Note You can also run the `bpbbackupdb` command from the command line to perform an offline, cold catalog backup. See *NetBackup Commands for UNIX and Linux* for more information on `bpbbackup`.

How Do I Know If a Catalog Backup Succeeded?

The All Log Entries, Problems, and Media Log reports, available from the Reports utility, provide information on NetBackup catalog backups. In addition, you can use:

- ◆ E-mail: An e-mail message is sent to the address indicated in the Disaster Recovery settings for an online catalog backup. This e-mail can be configured with the `mail_dr_info` script. For more information, see “Disaster Recovery E-mails and the Disaster Recovery File” on page 247.
- ◆ The `dbbackup_notify` script can be used to send an e-mail for offline catalog backups. See the *System Administrator's Guide, Volume II* for more information on setting up this script.

Strategies to Ensure Successful Catalog Backups

- ◆ Use only those methods described in this chapter to back up the catalogs. The methods described here are the only operations that can track all relevant NetBackup activities and ensure consistency between the catalog files.
- ◆ Back up the catalogs frequently and on a regular basis. If a catalog backup is lost, you lose information about backups and configuration changes that were made between the time of the last NetBackup catalog backup and the time that the disk crash occurred.
- ◆ Never manually compress the catalogs or NetBackup may be unable to restore the catalogs using `bprecover`.
- ◆ If you back up your catalogs to disk (not recommended), always back up to a different disk than where the catalogs reside. If you back up to the same disk and that disk fails, you will also lose the catalog backups in addition to the catalogs and recovery will be much more difficult. Also, ensure that the disk has enough space for the catalogs or it will fill up and backups will fail.
- ◆ The NetBackup binary image catalog is more sensitive to the location of the catalog. Storing the catalog on a remote file system may have critical performance issues for catalog backups. NetBackup does not support saving catalogs to a remote file system such as NFS or CIFS.

Considerations if running offline, cold catalog backups:

- ◆ If you are using media servers, be sure to manually alter the NetBackup catalog configuration to include the catalogs on the media servers.
- ◆ Keep a hard-copy record of the media IDs where you store the NetBackup catalog backups, or include the administrator's e-mail address in the Global Attributes properties. The e-mail that the administrator receives includes the status of each catalog backup and the media ID that was used. Print the e-mail or save it on a disk other than the disk containing the catalogs. (See "Global Attributes Properties" on page 414.)
- ◆ If sending catalog backups to a robot or tape stacker, a second standalone tape drive, or to disk, choose either of the two automatic backups: **After each session of scheduled, user, or manual backups** or **After each session of scheduled backups**
- ◆ If using a single, standalone tape drive to back up both catalog *and* regular business data, choose either:
 - ◆ **After each session of scheduled backups** if you will be running only one backup session per day or night, or
 - ◆ **Only when manually initiated** if you will be running multiple backup sessions in a single day or night

Because NetBackup will not place catalog and regular backups on the same tape, both methods require you to swap tapes.

The general procedure for catalog backups when you have only one standalone drive is as follows:

- a. Insert the tape configured for catalog backups.
- b. Manually start the backup. (See “Backing Up Catalogs Manually” on page 278.)
- c. When the backup is complete, remove the tape and store it in a safe place.

Caution The catalog backup tape must be removed when the backup is done or regular backups will not occur. NetBackup does not mix catalog and regular backups on the same tape.

About the Binary Catalog Format

Maintaining the catalog in a binary file format has several advantages over maintaining 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 requiring post-processing, improving catalog performance for alternate backup methods.

Catalog Conversion Utility

NetBackup offers a catalog format conversion utility called `cat_convert`. This utility converts ASCII image .`f` files in the image database of NetBackup versions 3.4, 4.0V or 4.5, to the binary format.

Upon installation, NetBackup does *not* convert existing ASCII catalogs to the binary catalog format. However, any new catalogs created will be binary. The `cat_convert` command is described in *NetBackup Commands for UNIX and Linux*.

Binary Catalog File Limitations

There are a few size limitations associated with the binary catalog to keep in mind.

- ◆ The maximum number of files that can be backed up per image:
 $(2^{31}) - 1$ files = 2,147,483,647 files = 7FFFFFFF files

- ◆ The maximum number of different user IDs and group IDs (combined):

$(2^{31}) - 1$ IDs = 2,147,483,647 IDs = 7FFFFFFF IDs

Moving the Image Catalog

Consider moving the image catalog to a file system or disk partition that contains more available space if the image catalog becomes too large for its current location.

Note Because NetBackup does not support saving catalogs to a remote file system, we strongly advise against moving the image catalog to a remote file system such as NFS or CIFS.

▼ To move the image catalog

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.
7. Add the new image-catalog path to the list that is included in NetBackup catalog backups.

Caution Be certain to add the path for the image catalog and not the link name. Otherwise, NetBackup will not back up the new location. In this example, the pathname is `/disk3/netbackup/db/images`.

Indexing the Catalog for Faster Access to Backups

If you have a large number of backups, consider indexing the catalogs in order to reduce the time it takes to restore files.

Indexing the catalog means creating indexes of the backed up files that are recorded in the NetBackup image catalog. NetBackup uses the indexes to go directly to the catalog entry for a file rather than starting the search at the beginning of the catalog entries.

There are two types of catalog indexing:

- ◆ Indexing the image header files. Image indexing benefits both binary and ASCII catalogs.
- ◆ Indexing the image .f file. This type of indexing creates an index of the file history detail up to nine directory levels deep. Since binary catalogs are self-indexing, binary catalogs do not benefit from this type of indexing. (See “Indexing the Image .f File” on page 283.)

Indexing the Header Files

The index image header files, run the following command:

```
bpimage -create_image_list [-client name]
```

Running this command creates the following index files in each client image directory:

```
IMAGE_FILES  
IMAGE_INFO  
IMAGE_LIST
```

To stop image header indexing for a client, remove these files.

Indexing the Image .f File

Use the `index_clients` command to generate indexes for one or all clients, for up to nine levels of directories:

```
/usr/opensv/netbackup/bin/index_clients level client_name
```

Where:

- ◆ *level* is the number of directory levels (1 to 9) to be indexed. The levels refer to the directories from where files were backed up on the client.

For example, if you're searching for `/payroll/smith/taxes/01` and *level* is 2, NetBackup starts the search at `/payroll/smith`. The default is 9.

Catalog Maintenance and Performance Optimization

- ◆ *client_name* is the name of the client of the backups you want to index. The default is all clients.

Run this command to activate indexing for a client. Once activated, indexing is done automatically every time NetBackup cleans up from the previous day's activities.

Notes on Image .f File Indexing

- ◆ Indexing the image .f file applies to ASCII catalogs only.
- ◆ NetBackup does not produce index files for backups that contain less than 200 files.
- ◆ Changing the index level affects future index creation and does not immediately create index files.
- ◆ If you are collecting true image restore information, the INDEX files take much more space for incrementals.

Catalog Index Examples

- ◆ To index client mars to index level 5 (five levels of directories), run:

```
/usr/opensv/netbackup/bin/index_clients 5 mars
```
- ◆ To index selected clients, run a command for each client. (Wildcards are not allowed.) The following commands index clients named *mars*, *jupiter* and *neptune* to index level 5:

```
/usr/opensv/netbackup/bin/index_clients 5 mars
/usr/opensv/netbackup/bin/index_clients 5 jupiter
/usr/opensv/netbackup/bin/index_clients 5 neptune
```
- ◆ To index all NetBackup clients to index level 3, run:

```
/usr/opensv/netbackup/bin/index_clients 3
```
- ◆ To index all NetBackup clients to index level 9, run:

```
/usr/opensv/netbackup/bin/index_clients
```

Space Requirements for Image .f file Indexing

The index files do not require much additional disk space. Regardless of the number of clients, indexing all clients to level 9 requires approximately 1.5 percent more space in the NetBackup catalog than if indexing is not used for any clients.

The index files reside in a directory named:

```
/usr/opensv/netbackup/db/images/clientname/INDEX
```

The indexing level resides in a file named:

```
/usr/opensv/netbackup/db/images/clientname/INDEXLEVEL
```

Disabling Catalog Indexing

To stop NetBackup from generating new INDEX files for a client, delete the INDEXLEVEL file. NetBackup continues to use existing INDEX files.

To temporarily stop using the INDEX files during searches but retain existing index files, change the INDEX directory to INDEX.ignore. When you are done, change INDEX.ignore back to INDEX to resume indexing.

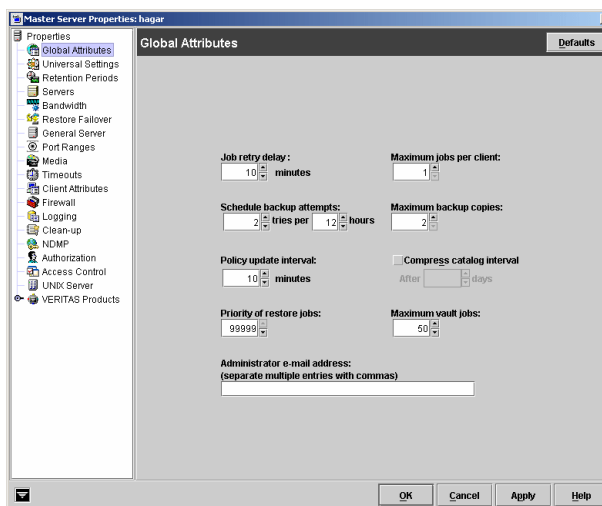
To permanently eliminate INDEX files for a client, delete the INDEX directory and the INDEXLEVEL file.

Compressing and Uncompressing the Image Catalog

The image catalog has information about all client backups and is accessed when a user lists or restores files. NetBackup offers you the option of compressing all or older portions of this catalog. There is no method to selectively compress image-catalog files other than by age.

Control image-catalog compression by setting the Global Attributes property, **Compress Catalog Interval**. This property specifies how old the backup information must be before it is compressed, thereby letting you defer

compression of newer information and not affect users who are listing or restoring files from recent backups. By default, **Compress Catalog Interval** is set to 0 and image compression is not enabled. (See “Global Attributes Properties” on page 414.)



Caution VERITAS discourages manually compressing or decompressing catalog backups using `bpimage - [de]compress` or any other method. If a regular or catalog backup is running while manually compressing or decompressing a catalog backup, this can result in inconsistent image-catalog entries, producing incorrect results when users list and restore files.

If you choose to compress the image catalog, NetBackup uses the `compress` command on the server to perform compression after each backup session, regardless of whether successful backups were performed. The operation occurs while NetBackup is expiring backups and before running the `session_notify` script and the backup of the NetBackup catalogs.

The time to perform compression depends on the speed of your server and the number and size of the files you are compressing. 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 compression. To minimize the impact of the initial sessions, consider compressing the files in stages. For example, you can start by compressing records for backups older than 120 days and then reduce this value over a period of time until you reach a comfortable setting.

Compressing the image catalog can greatly reduce the disk space used as well as the amount of media required to back up the catalog. The amount of space you reclaim varies with the types of backups you perform. Full backups result in a larger percentage of catalog compression than incremental backups because there is normally more duplication of data in a catalog file for a full backup. A reduction of 80% is sometimes 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, you may have to increase the time-out value associated with list requests by changing the `LIST_FILES_TIMEOUT` option in the `bp.conf` file of the client.

Uncompressing the Image Catalog

You may find it necessary to temporarily uncompress all records associated with an individual client (for example, if you anticipate large or numerous restore requests). Perform the following steps as root on the master server:

▼ To uncompress client records

1. 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/bprdreq -terminate
```
3. Verify that `bpdbm` is running by using:

```
/usr/opensv/netbackup/bin/bpps
```
4. Expand **Host Properties > Master Servers**. Open the properties of a host. On the **Global Attributes** properties page, clear the **Compress Catalog Interval** check box. (See "Global Attributes Properties" on page 414.)
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
```

Catalog Maintenance and Performance Optimization

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 will be compressed after the next backup schedule.

Viewing NetBackup Reports

5

Generate reports using the NetBackup **Reports** utility. The reports serve to verify, manage, and troubleshoot NetBackup operations. NetBackup reports display information according to job status, client backups, and media contents. The Troubleshooter is available within the **Reports** utility to help analyze the cause of errors that can appear in a NetBackup report.

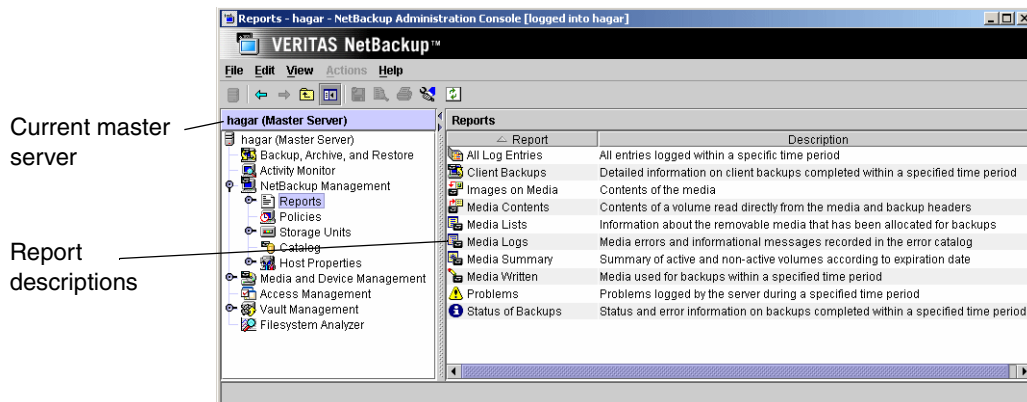
The following topics are discussed in this chapter:

- ◆ “Introduction to the Reports Utility” on page 290
- ◆ “Reports Window” on page 292
- ◆ “NetBackup Report Types” on page 295
- ◆ “Using the Troubleshooter Within Reports” on page 307

Introduction to the Reports Utility

Introduction to the Reports Utility

Once **Reports** is expanded in the NetBackup Administration Console, the Details pane displays a description of all possible reports. Each report type is discussed in “NetBackup Report Types” on page 295.



Reports Menu Bar

The **Reports** menu bar consists of the following menu items:

Option	Description
File	Options Change Server , New Window from Here , Adjust Application Time Zone , Export , Page Setup , Print Preview , Print , Close Window , and Exit are described in the section, "File Menu" on page 14.
Edit	Option Find is described in "Edit Menu" on page 15. Edit Default Time: Opens the Default Time Setting dialog. The setting here determines the Date/Time range for the report, where applicable.
View	Options Show Toolbar , Show Tree , Alternate Table Row Color , Back , Forward , Up One Level , Options , Refresh , Column Layout , Sort , and Filter are described in the section, "View Menu" on page 16.
Actions	No options displayed.
Help	Options Help Topics , Troubleshooter , License Keys , Current NBAC User , and About NetBackup Administration Console are described in the section, "Help Menu" on page 17.

▼ To run a report

1. In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of report types appears.

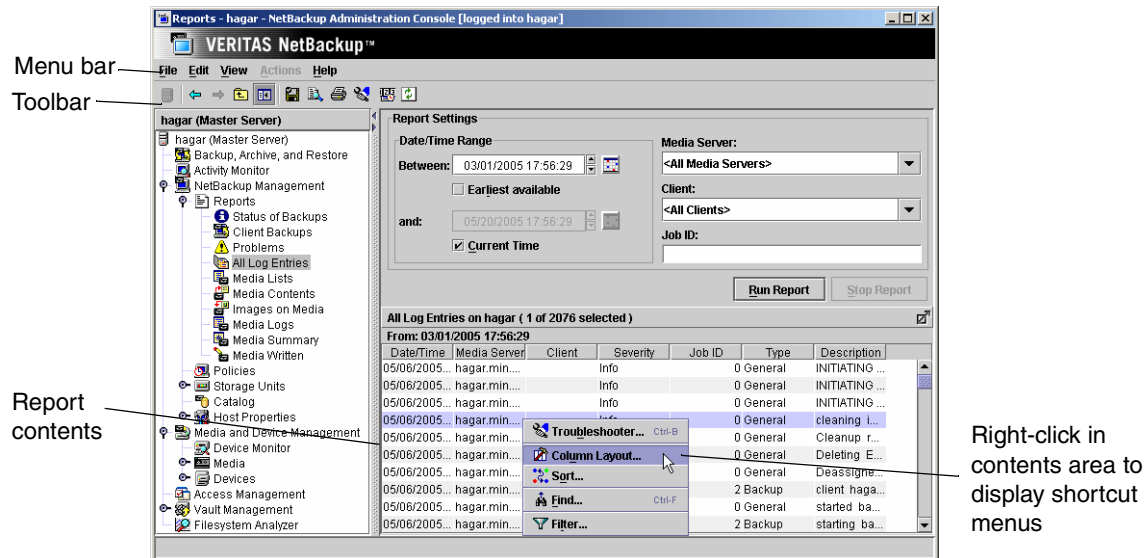
The report information is for the master server that is currently selected. To run a report on a different master server, click **File > Change Server**. (See "Administering a Remote Master Server" on page 474.)

2. Select the name of the report you would like to run. The right pane displays various options for running the report.
3. Select the media server(s) and/or clients on which to run the report and/or select the time period for which the report will run.
4. Click **Run Report**. For a description of the report fields, see "NetBackup Report Types" on page 295.

Reports Window

Reports Window

The Reports window contains a number of methods to make it easier to view report listings and manage report data.



Report Toolbar

The buttons on the toolbar provide shortcuts for menu commands. To display or hide the NetBackup toolbar, click **View > Show Toolbar**.

For information on the standard toolbar buttons, see “Using the NetBackup Administration Console” on page 9.

Report Contents Pane

The lower right pane in the Reports window displays the contents of the report that you’ve run.

Shortcut Menus

To display a list of commands that apply to a list, right-click on a report. Depending on which report you’re viewing, the shortcut list may include:

- ◆ **Column Layout:** Opens the Column Layout dialog where you can show or hide columns. (By default, all columns are not displayed.)
- ◆ **Sort:** Opens the Sort dialog where you can specify sort criteria for the columns.
- ◆ **Find:** Opens the Find dialog, used to find text within the report.
- ◆ **Filter:** Use the Filter option to narrow in on specific data in a table that you wish to view. The controls on the Filter dialog allow you to list rows that match specified criteria.
- ◆ **Troubleshooter:** Launch the Troubleshooter to enter a status code.

Reports Settings

Use the report settings to specify the following criteria for building your report. Not all settings are available for every report type.

Date/Time Range

Specify the time period that you want the report to encompass. By default, the start time is one day before the report is run and the end time is the time the report is run.

Select **Earliest Available** to include the earliest possible data available.

Select **Current Time** to include all data up to the present.

The Clean-up host property, **Keep Logs**, determines the period of time for which the information is available. (See “Keep Logs” on page 359.)

Client

Click the **Client** box and select **All Clients** or the client to which the report will apply.

Media Server

Click the **Media Server** box and select **All Media Servers** or the name of the media server to which the report will apply. The master server that is currently selected and the media servers appear in the report.

Job ID

Specify the Job ID for which you want the report.

Reports Window

Media ID

For media types of reports, specify the media ID or **All Media**. The Media Contents report requires a specific ID.

Volume Pool

For a media summary report, specify the volume pool name or **All Volume Pools**.

Verbose Listing

Select **Verbose Listing** to have NetBackup provide more details in the Media Summary report.

Run Report

Click **Run Report** after you've selected the criteria for a report.

Stop Report

Click **Stop Report** if a report is running, but you don't want to wait for it to finish.

NetBackup Report Types

The following sections describe the contents of NetBackup reports.

Status of Backups Report

The Status of Backups report shows status and error information on jobs completed within the specified time period. If an error has occurred, a short explanation of the error is included. The following table explains the columns in the Status of Backups report:

Status of Backups Report

Column	Meaning
Client	Name of the client for which the backup was performed.
Date/Time	Time the backup began.
Description	Message describing the status.
Job ID	Job ID corresponding to the backup that appears in the Activity Monitor.
Media Server	Media server that controlled the backup.
Policy	Name of the policy that was used to back up the client.
Schedule	Name of the schedule that was used to back up the client.
Status	Completion status of the backup. If the status code is 0, the operation succeeded. If the status code is not 0, right-click on the entry in the table and launch the Troubleshooter to display troubleshooting information.

NetBackup Report Types

Client Backups Report

The Client Backups report shows detailed information on backups completed within the specified time period. The following table explains each field in the Client Backups report.

Client Backups Report

Field	Meaning
Backup Date/Time	Date and time that the backup began.
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Client	Name of the client for which the backup was performed.
Compressed	Yes indicates that the backup was compressed.
Elapsed Time	How much time the backup required.
Encrypted	Yes indicates that the backup is encrypted. Encryption and decryption is possible only with the NetBackup Encryption option.
Extended Security Information	This field is reserved for future use and always displays <i>No</i> .
Expiration Time	Date and time at which NetBackup will expire the record of this backup.
File Restore Raw	Individual file restore from raw. This is set by the corresponding policy attribute if it applies.
File System Only	This field is reserved for future use and always displays <i>No</i> .
Image Dump Level	Applies to NDMP backups. <i>0</i> indicates a full backup and greater than 0 indicates an incremental backup.
Image Type	Shows <i>Regular</i> , if it is a scheduled or user-directed backup, <i>Pre-imported</i> , if phase I of the import process is completed, or <i>Imported</i> , if it is an imported image (phase II of import process complete).
Keyword	Keyword that the user associates with this image at the time of the backup.
Kilobytes	Number of kilobytes in the backup.
Multiplexed	Yes indicates that the backup was multiplexed.

Client Backups Report (continued)

Field	Meaning
Number of Files	Number of files in the backup.
Object Descriptor	This field is reserved for future use and is always empty.
Policy	Name of the policy that was used to back up the client.
Policy Type	Type of policy (for example, <i>Standard</i> , <i>MS-Windows-NT</i> , and so on).
Primary Copy	Primary copy shows which copy (1 or 2) NetBackup uses to satisfy restore requests.
Retention Level	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown is the first level assigned. (See "Retention" on page 115.)
Schedule Name	Name of the schedule that was used for the backup.
Schedule Type	Type of schedule used for the backup (for example, full or incremental).
True Image Restore Available	<i>Yes</i> indicates that NetBackup is collecting true image restore information for this policy.

Problems Report

The Problems report lists the problems that the server has logged during the specified time period. The information in this report is a subset of the information obtained from the All Log Entries report. (See "All Log Entries Report" on page 298.)

NetBackup Report Types

All Log Entries Report

The All Log Entries report lists all log entries for the specified time period. This report includes the information from the Problems report and Media Logs report. This report also shows the transfer rate, which is useful in determining and predicting rates and backup times for future backups (the transfer rate does not appear for multiplexed backups). The following table explains the columns in the All Log Entries report:

All Log Entries Report

Column	Meaning
Client	NetBackup client involved in the event. If the event did not involve a client, the column is blank.
Date/Time	Date when the event began.
Description	Message describing the status.
Job ID	Identifier that NetBackup assigns when performing the job. If the job ID is 0, the event is not related to running a job.
Media Server	Media server that controlled the backup.
Policy	Name of the policy that was used to back up the client.
Process	Process that returned the status.
Schedule	Name of the schedule that was used to back up the client.
Severity	Severity level of the status: Critical, Warning, Error, Info.
Status	Completion status of the backup. If the status code is 0, the operation succeeded. If the status code is not 0, right-click on the entry in the table and launch the Troubleshooter to display troubleshooting information.
Type	Type of status.

Media Lists Report

The Media Lists report shows information for volumes that have been allocated for backups. This report does not show media for disk type storage units or for offline, cold catalog backups of the NetBackup catalogs.

- ◆ For information about backups saved to disk storage units, use the Images on Media report.
- ◆ To track media used for offline, cold catalog backups, keep a hard copy record or configure the **Administrator's E-mail Address** Global Attribute host property. (See "Administrator's E-mail Address" on page 417.) This host property causes NetBackup to send an E-mail that indicates the status of each catalog backup and the media ID that was used. You can then print the E-mail or save it on a disk other than the one that contains the catalogs.

The following table explains the columns in the Media Lists report:

Media Lists Report

Column	Meaning
Allocated Date/Time	Date and time that Media Manager allocated the volume.
Density	Density of the device that produced the backups on this volume.
Expiration Time	Date and time when the backups on the volume expire.
Header Size	Optical header size in bytes.
Images	Total number of backups on the volume.
Kilobytes	Total number of kilobytes on this volume.
Last Offset	Optical offset of the last header.
Last Read Date/Time	Last time a restore was done from this volume.
Last Written Date/Time	Last time the volume was used for backups.
Media ID	Media ID that is assigned when the volume is added to Media Manager.

NetBackup Report Types

Media Lists Report (continued)

Column	Meaning
Media Server	<p>Server where the volumes reside.</p> <p><i>Applies only to NetBackup Enterprise Server:</i> It is possible to have more than one if the master server has media servers and <i>ALL</i> was selected for the server.</p>
Number of Restores	Number of times this volume has been used for restores.
Partner ID	For an optical disk, this is the media ID of the volume on the other side of the platter.
Retention Level	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown is the first level assigned. (See "Retention" on page 115.)
Sector Size	Optical sector size in bytes.
Status	<p>The messages that commonly appear here are the following:</p> <p>Active: The volume is currently in use.</p> <p>Expired: The volume has expired.</p> <p>Frozen: The volume is unavailable for future backups. A frozen volume never expires, even after the retention period ends for all backups on the media. This means that the media ID is never deleted from the NetBackup media catalog and remains assigned to NetBackup. (The <code>bpmmedia</code> command can also be used to manually freeze or unfreeze volumes.)</p> <p>A frozen volume is available for restores. If the backups have expired, the backups first require importing.</p> <p>Imported: The backup was imported to this server. The volume cannot be used for further backups until retention periods for all backups on it have expired. At that time, the imported volume is deleted from the NetBackup media catalog and unassigned from NetBackup.</p> <p>An imported volume is available for restores. If the backups have expired, the backups first require importing.</p> <p>MEDIA_FULL: The media is full and no more backups are written to it. NetBackup sets FULL status if it encounters an end of media (EOM) during a backup. A full volume is unavailable for future backups until the retention period expires for all backups that are on it.</p>

Media Lists Report (continued)

Column	Meaning
	MPX_MEDIA: The media contains multiplexed images.
	Multi-Retlev: (Multiple Retention Level) The volume contains backups of more than one retention level.
	Suspended: The volume cannot be used for further backups until retention periods for all backups on it have expired. (The <code>bpmmedia</code> command can also be used to manually suspend or unsuspend volumes.) A suspended volume is available for restores. If the backups have expired, the backups first require importing.
Valid images	Number of nonexpired backups on the volume. For example, if the volume has 50 backups but only 10 are valid, then the other 40 have expired. If the volume has any multiplexed backups, this field contains <i>MPX</i> .
Volume Pool	A number that corresponds to the volume pool for the media. 0 = None (no volume pool) 1 = NetBackup 2 = DataStore 3 = CatalogBackup

NetBackup Report Types

Media Contents Report

The Media Contents report shows 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 has to be mounted, there will be a longer delay before the report appears.

Note The Media Contents report does not apply to disk type storage units or NetBackup offline, cold catalog backups.

The following table explains the columns in the report.

Media Contents Report

Column	Meaning
Allocated Date/Time	Date and time that Media Manager allocated the volume.
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Block size (in bytes)	Size of the data blocks used to write the backup. When multiplexing is used, the block size can vary between backups on the same volume.
Copy Number	Shows the copy number (1 or 2).
Creation date	Date that NetBackup created the backup.
Expiration Time	Time that the backup expires.
File number	Position of the file, where file 1 is the first. If the volume contains multiplexed backups, it can have multiple files with the same number.
Fragment Number	Greater than 1 only if the backup is split across multiple volumes or if the storage unit maximum fragment size is exceeded.
Media Id	Media ID that is assigned when the volume is added to Media Manager.
Retention Level	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown is the first level assigned. (See "Retention" on page 115.)

Images on Media Report

The Images on Media report lists the contents of the media as recorded in the NetBackup image catalog. You can generate this report for any type of media (including disk) and filter it according to client, media ID, or path.

Note The Images on Media report does not show information for media used for NetBackup offline, cold catalog backups.

The following table explains the columns in the Images on Media report:

Images on Media Report

Column	Meaning
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Blockmap	Indicates whether this fragment is a blockmap (<i>Yes</i> or <i>No</i>).
Block Size	Size of the data blocks used to write the backup. When multiplexing is used, the block size can vary between backups on the same volume.
Client	Name of the client that was backed up.
Copy Number	Greater than 1 only if there are multiple copies.
Compressed	<i>Yes</i> indicates that the backup is compressed.
Density	Density of the device that produced the backup.
Device Written On	Device where the backup was written. This is the drive index configured in Media Manager.
Encrypted	<i>Yes</i> indicates that the backup is encrypted. Encryption and decryption is possible only with the NetBackup Encryption option.
Expiration Date/Time	Expiration date and time for the corresponding copy number; not the expiration of first copy.
File Number	File number on the media.
Fragment Number	Fragment number. <i>IDX</i> (Index file) if the fragment contains true image restore information or is for an individual-file-restore-from-raw backup. <i>TIR</i> indicates that the fragment number is part of a true-image backup.

NetBackup Report Types

Images on Media Report (continued)

Column	Meaning
Kilobytes	Size of the fragment in kilobytes. This value does not include the space for tape headers between backups. A fragment size of 0 is possible in a multiplexed backup.
Policy	NetBackup policy for which the backup was created.
Media Date/Time	Date and time when the copy will expire. Only valid on fragment 1 of a copy.
MediaID	Media ID of the volume that contains the backup image. For disk, it is a pathname.
Media Server	Server with the database that contains this information.
Media Type	Type of storage and can be removable (RMed) or disk (Disk).
Multiplexed	<i>Yes</i> indicates that the copy is multiplexed. Valid for all the fragment numbers.
Number of Files	Number of files in the backup.
Offset	Applies only to optical disk and is the byte offset on the media where the backup image begins. Ignore this value for tapes and magnetic disk.
Policy Type	Type of policy (for example, <i>Standard</i> , <i>MS-Windows-NT</i> , and so on).
Remainder	Bytes written beyond kilobytes filed. Size of fragment is exactly: Kilobytes*1024 + Remainder.
Retention Level	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown is the first level assigned. (See "Retention" on page 115.)
Schedule	Name of the schedule that was used to back up the client.
Schedule Type	Type of backup (full, differential incremental, cumulative incremental, or user-directed).

Media Logs Report

The Media Logs report shows media errors or informational messages that are recorded in the NetBackup error catalog. This information also appears in the All Log Entries report. (See “All Log Entries Report” on page 298.)

Media Summary Report

The Media Summary report summarizes active and nonactive volumes for the specified server according to expiration date. It also shows how many volumes are at each retention level. In verbose mode, the report shows each media ID and the expiration date.

Nonactive media are those with a status of FULL, FROZEN, SUSPENDED, or IMPORTED. Other volumes are considered active. (See “Media Lists Report” on page 299.)

The only expired volumes that appear in this report are those that are FROZEN. NetBackup deletes other expired volumes from the media catalog when it runs backups. An expired volume with other status can show up only if the report is run between the time the volume expires and the time that the next backup is done.

NetBackup Report Types

Media Written Report

The Media Written report identifies volumes that were used for backups within the specified time period. This report does not display volumes used for NetBackup offline, cold catalog backups or volumes used for duplication if the original was created prior to the specified time period.

The following table explains the columns in the Media Written report:

Media Written Report

Column	Meaning
Kilobytes	Number of kilobytes in the backup.
Last Written Date/Time	Date and time when the media was last written.
Media ID	Media ID that is assigned when the volume is added to Media Manager.
Media Server	Server that contains the Enterprise Media Manager database with the records for this volume.
Retention Level	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown is the first level assigned. (See "Retention" on page 115.)
Times Written	Number of times this media was written.


Using the Troubleshooter Within Reports

Use the Troubleshooter within Reports to find explanations and corrective actions based on the NetBackup status code that the job returns.



▼ To run Troubleshooter within Reports

1. Run a report.
2. Right-click a line in the report and select **Troubleshooter** from the shortcut menu.
3. The Troubleshooter dialog appears, stating an explanation of the problem on the Problem tab, and a recommended action on the Troubleshoot tab.

Open the Troubleshooter at any time (**Help > Troubleshooter**), enter a status code, then click **Lookup**.



Using the Troubleshooter Within Reports



Monitoring NetBackup Activity

6

This chapter explains how to use the NetBackup Activity Monitor to perform various functions in order to monitor and troubleshoot NetBackup jobs, daemons, and processes.

This chapter includes the following sections:

- ◆ “Introduction to the Activity Monitor” on page 310
- ◆ “Jobs Tab” on page 318
- ◆ “Daemons Tab” on page 325
- ◆ “Processes Tab” on page 329
- ◆ “Media Mount Errors” on page 332
- ◆ “Managing the Jobs Database” on page 332

Introduction to the Activity Monitor

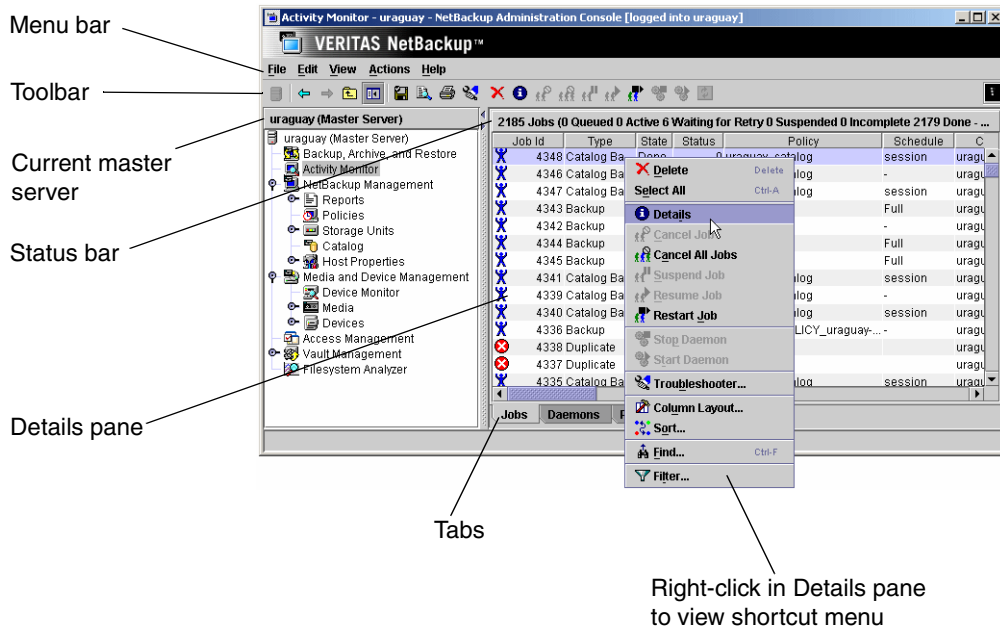
Introduction to the Activity Monitor

Use the Activity Monitor in the NetBackup Administration Console to monitor and control NetBackup jobs, daemons, and processes.

While the Activity Monitor is active in the NetBackup-Java Administration Console, the `bpj obd` daemon continuously supplies the status of NetBackup job activity to the Activity Monitor.

When receiving job activity data from the `bpj obd` daemon, updates to the Activity Monitor occur as jobs are initiated, updated and completed. The updates occur instantaneously because there is no associated refresh cycle.

The Activity Monitor contains the following information:



Activity Monitor Menu Bar

The **Activity Monitor** menu bar consists of **File**, **Edit**, **View**, **Actions**, and **Help**. See Chapter 1 for a description of the items found on these menus.

Note The **Filter** option on the **View** menu is useful for displaying in Activity Monitor only those jobs with specified characteristics. For example, jobs that started before or after a specific date and time; jobs that are in either the active or queued state; jobs that have status completion codes within a specified range.

The following section describes the **Actions** menu options.

Actions Menu

The **Actions** menu contains the following options when **Activity Monitor** is selected.

- ◆ **Details:** Displays detailed information about the job, daemon, or process you select in the list.
- ◆ **Cancel Job:** Cancels uncompleted jobs that you have selected in the Jobs list. A cancelled checkpointed backup or restore job cannot be resumed from the last checkpoint. If you want to be able to resume the job, use **Suspend Job** instead.
- ◆ **Cancel All Jobs:** Cancels all uncompleted backup jobs.
- ◆ **Suspend Job:** Suspends an Active, Queued, or Waiting for Retry checkpointed backup or restore job. An administrator may want to suspend a job to free a resource or to run another job, then resume the suspended job when the resource is available. (See “Move Restore Job From Incomplete State to Done State” on page 360 and “Move Backup Job from Incomplete State to Done State” on page 360.)
- ◆ **Resume Job:** Resumes an Incomplete or Suspended checkpointed backup or restore job from the last checkpoint. When a checkpointed backup is resumed, the backup is resumed on the same media server. If *Any available* is specified as the storage unit group, or if a specific storage unit group is specified, the backup may use a different storage unit on the same media server.

However, a backup job to a tape storage unit cannot be resumed on a disk storage unit, or a disk storage unit to a tape storage unit.

A job may be in an incomplete state indefinitely and may be resumed until the backup or incomplete backup has expired. (See “Move Restore Job From Incomplete State to Done State” on page 360 and “Move Backup Job from Incomplete State to Done State” on page 360.)

Introduction to the Activity Monitor

The same is true for optical storage units mixed with tape devices: If a backup that was originally started on an optical device is resumed on a tape device (or vice versa), the backup will fail with a 174 (Media Manager - system error occurred) status. In this situation, use a specific storage unit, or use storage unit groups to separate the optical and tape devices.

- ◆ **Restart Job:** Restarts a job from the beginning. (The job is not required to be checkpointed.) The job may be restarted on the same media server or a different media server. (See “Checkpoint Restart for Backup Jobs” on page 75.)

Note The job must be marked as *Done* or *Waiting for Retry* in order to restart it. To restart an Incomplete or Suspended job, cancel the job to force it into the Done state.

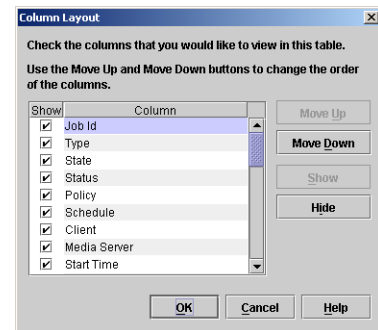
When a job is restarted, a new job ID is assigned to the restarted job. The job details available through the Activity Monitor and various reports record the event in the following format:

timestamp Job manually restarted as *new_jobID*.

- ◆ **Stop Daemon:** Stops daemons that you have selected in the Daemons list.
- ◆ **Start Daemon:** Starts daemons that you have selected in the Daemons list.

▼ To show or hide column heads

1. Open the Activity Monitor.
2. Click **View > Column Layout**. The **Column Layout** dialog appears, showing the current settings.
3. Select the heading you wish to display or hide.
 - ◆ Select the **Show** button to display the heading.
 - ◆ Select the **Hide** button if you do not want to see the column heading.
4. To change the order in which the columns appear, select the column heading, then click the **Move Up** button or the **Move Down** button to reorder the columns.
5. Click **OK** to apply the changes.



▼ **To monitor the detailed status of selected jobs**

1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the job(s) for which you want to view details.
3. Select **Actions > Details**. A Jobs Details dialog appears for each job you selected.

▼ **To delete completed jobs**

1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the job(s) you want to delete.
3. Select **Edit > Delete**. All selected jobs are deleted.

▼ **To cancel uncompleted jobs**

1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the uncompleted jobs you want to cancel. An uncompleted job is one that is in the Queued, Re-Queued, Active, Incomplete, or Suspended state.
3. Select **Actions > Cancel Job**. All selected jobs are cancelled.

Note If the selected job is a parent job, all the children of that parent job are cancelled as well. In most cases, cancelling a child job cancels only that job and allows the other child jobs to continue. One exception to this is when creating multiple copies—cancelling a child job cancels the parent job and all child jobs.

4. To cancel all uncompleted jobs in the jobs list, click **Actions > Cancel All Jobs**.

▼ **To restart a completed job**

1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the completed job you want to restart.
3. Select **Actions > Restart Job**. All selected jobs are restarted. In this case, a new job ID is created for the job. The job details for the original job will reference the job ID of the new job.

Introduction to the Activity Monitor

▼ To suspend a restore or backup job

1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the job you want to suspend.
3. Select **Actions > Suspend Job**. All selected jobs are suspended.

▼ To resume a suspended or incomplete job


1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the suspended or incomplete job you want to resume.
3. Select **Actions > Resume Job**. All selected jobs are resumed.

▼ To print job detail information from a list of jobs

1. Open the Activity Monitor and select the **Jobs** tab.
1. Select a job to print. Hold down the Control or Shift key to select multiple jobs. If no job is selected, all jobs will be printed.
2. Select **File > Print**.

▼ To export Activity Monitor data to a text file

1. 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 pathname of the file where you want the job data to be written, then click **Save**.



If a job fails, use the Troubleshooter on the **Help** menu to find explanations and corrective actions based on the NetBackup status code that the job returns.

▼ To run Troubleshooter within the Activity Monitor

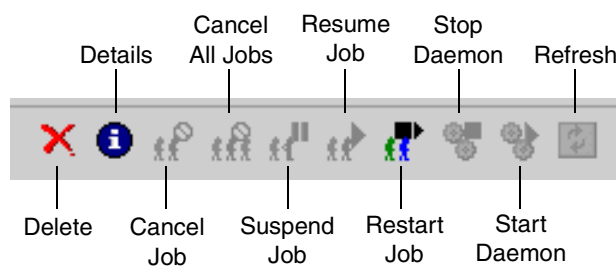
If a job fails, use the Troubleshooter on the **Help** menu to find explanations and corrective actions based on the NetBackup status code that the job returns.

1. Select a job in the Activity Monitor.
2. Select the job you wish to troubleshoot.
3. Open the Troubleshooter using one of the following methods:
 - ◆ Click the **Troubleshoot** icon.
 - ◆ Select **Help > Troubleshooter**.
 - ◆ Open the job details for a job, click the **Detailed Status** tab. Then click **Troubleshooter**.
 - ◆ Right-click on the job and select **Troubleshooter**.
4. The Troubleshooter dialog appears, stating an explanation of the problem on the Problems tab, and a recommended action on the Troubleshoot tab.

If there is no status code entered in the Troubleshooter status code field, enter the status code of the failed job and click **Lookup** to locate the troubleshooting information. You can open the Troubleshooter at any time and enter a status code.

Activity Monitor Toolbar

The buttons on the toolbars provide shortcuts for menu commands. To display or hide the NetBackup toolbar, click **View > Show Toolbar**.



Status Bar

The status bar appears in the Jobs tab, at the top of the Activity Monitor Details pane. The status bar displays the following information:

- ◆ The master server on which the jobs reside.
- ◆ The total number of jobs.

Introduction to the Activity Monitor

- ◆ The number of jobs in each of the job states: Active, Queued, Waiting for Retry, Suspended, Incomplete, and Done.
- ◆ Number of jobs currently selected.
- ◆ Number of daemons running.

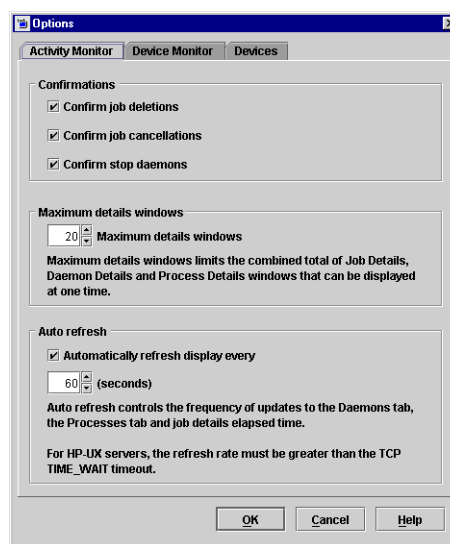
The numbers always reflect the actual number of jobs, even when filtering is used.

Setting Activity Monitor Options

Click **View > Options** and select the Activity Monitor tab to access configurable options for the Activity Monitor.

While working in the Activity Monitor, you may elect to receive confirmation warnings:

- ◆ **Confirm job deletions:** If checked, the user will be prompted with a confirmation dialog when deleting jobs.
- ◆ **Confirm job cancellations:** If checked, the user will be prompted with a confirmation dialog when cancelling jobs.
- ◆ **Confirm stop daemons:** If checked, the user will be prompted with a confirmation dialog when stopping daemons.




To discontinue further confirmations, check **In the future, do not show this warning** when deleting or cancelling a job, or when stopping a daemon.

Set the **Maximum Details Windows** value to specify the maximum number of Activity Monitor job details, daemon details and process details windows that can be displayed at one time.

Check **Auto Refresh** to periodically refresh data on the Daemons tab and the Processes tab and job details elapsed time. Other Jobs tab data is refreshed independently of the **Auto Refresh** setting.

Enter the rate (in seconds) at which data will be refreshed in the Daemons and Processes tabs.

After making any changes, click **OK** to close the dialog and apply the changes.



Introduction to the Activity Monitor



Jobs Tab

The **Jobs** tab displays all jobs that are in process or have been completed for the master server currently selected.

Note Job selection preference is given to jobs from NetBackup 6.0 media servers over media servers of previous versions.

Parent Jobs

For some backup jobs, a parent job is used to perform pre- and post-processing. If the job is a parent job, the Job PID column in the Jobs tab is blank.

A parent job will execute start and end notify scripts if the scripts exist in `/usr/opensv/netbackup/bin/` on the client.

- ◆ PARENT_START_NOTIFY
- ◆ PARENT_END_NOTIFY

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:

- ◆ *Advanced Client:* The parent job creates the snapshot, initiates children jobs and deletes the snapshot when complete.

Children jobs are created if the Advanced Client settings in the policy 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.)

Children jobs are *not* created if the Advanced Client settings are configured to retain snapshots for Instant Recovery, but to create snapshots only. That is, the snapshot will not be backed up to a storage unit, so no children jobs are generated. (**Snapshots only** is selected in the policy Schedule Attributes tab.)

- ◆ *Bare Metal Restore:* The parent job runs `brmsavecfg`, then initiates the backup as a child job. If multistreaming and BMR are used together, the parent job may start multiple children jobs.
- ◆ *Offline, cold catalog backups:* The parent job initiates the `bpbbackupdb` as a child job.
- ◆ *Online, hot catalog backups:* The parent job for hot catalog backups works with `bpdbm` to initiate multiple children backup jobs: a Sybase backup, a file system backup of the master, and if necessary, backups of any 5.x media servers and BMR database.

- ◆ *Multiple copies:* A job producing multiple copies consists of one parent job and multiple children 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. (For setup information, see “Multiple Copies” on page 110)
- ◆ *Multiple data streams:* The parent job performs stream discovery and initiates children jobs. A parent job does not display a schedule in the Activity Monitor. Instead, a dash (-) is displayed for the schedule because the parent schedule is not used and the children schedules may be different. The children jobs display the ID of the parent job in the Activity Monitor.
- ◆ *SharePoint:* The parent job runs a resolver process during which children jobs are started. This process is similar to the stream discovery for multiple data streams. If multiple data streams are enabled, some children jobs may be split into multiple streams.
- ◆ *Vault:* The parent job starts the Vault profile, which, in turn, starts the duplicates as jobs. The duplicates do not appear as children jobs in the Activity Monitor.

▼ To view job details

To view the details for a specific job, double-click on the job in the Jobs tab. The Job Details dialog appears, containing detailed job information on two tabs: a Job Overview tab and a **Detailed Status** tab.

Not all columns are displayed by default. Click **View > Column Layout** to show or hide columns.

Job Details

Name	Purpose
Active Start	The time when the most recent attempt became active.
Active Elapsed	The time since the most recent attempt became active.
Attempt	For Active jobs, the number of the current attempt. For Done jobs, the total number of attempts.
Client	The name of the client associated with the job.
Copy	The copy number when the multiple copies are created.
Current File	For Active jobs, the path of a file that was recently written to the image. If the job is backing up many files, not all of them necessarily appear in this column over the course of the backup.

Jobs Tab

Job Details

Data Movement	Distinguishes between the various types of jobs: Standard (resulting from a regular backup), Synthetic (backup), or Disk Staging. If the Advanced Client is licensed, additional entries can include: Instant Recovery Disk, Instant Recovery Disk and Tape, and Snapshot.
Elapsed Time	The amount of time that has elapsed since the job was initially queued.
End Time	The date and time that the operation completed.
Files	The number of files that have been written.
Job ID	The identifier that NetBackup assigns to each job. The identifier is unique on the server where the job was run.
Job PID	The process ID. If a backup is multiplexed, all jobs associated with the same multiplexed storage unit have the same PID. The Job PID column is blank if it is a generic, or parent, job. For example, multiple data stream backups are initiated by a generic job.
Job State	<p><i>Active:</i> Indicates currently active jobs.</p> <p><i>Done:</i> Indicates completed jobs.</p> <p><i>Incomplete:</i> Indicates backup or restore jobs that have failed with a resumable error. Look in the Activity Monitor to determine if the failed job requires manual intervention. After correcting the problem, the administrator may resume the job.</p> <p>When a job is resumed, it retains the same job ID. A job may remain in the Incomplete state for a limited time before being set to Done, after which the job is no longer resumable.</p> <p><i>Queued:</i> Indicates jobs in the NetBackup queue. A queued restore job is one for which NetBackup is still determining which files are needed. Jobs stay in the queued state until resources are assigned. Backup jobs sent to standalone drives requiring operator intervention remain in the Queued job state until media is provided.</p> <p><i>Waiting for Retry:</i> Indicates jobs that are placed back in the queue to be tried again because the previous attempt was unsuccessful.</p> <p><i>Suspended:</i> Backup or restore jobs that have been suspended by the NetBackup administrator. Suspended jobs do not display a status code.</p>
KB Per Second	The average data transfer rate in kilobytes per second over the length of the current attempt.

Job Details

Kilobytes	The number of kilobytes that have been written.
Master Server	The master server on which the job is run.
Media Server	The NetBackup server controlling the media.
Media to Eject	The number of tapes to be ejected for the selected Vault job. <i>The number may not represent the number of tapes actually ejected.</i> For example, if the Vault profile was configured for manual eject, the tapes may not have yet been ejected. Or, if something went wrong with the device, fewer tapes may actually have been ejected than the number here indicates.
OffHost Type	Indicates the off-loading of backup processing to a separate backup agent executing on another host. Additionally licensed products such as Advanced Client are required for these offhost backups types: <ul style="list-style-type: none"> ♦ Alt. Client (Alternate Client Backup) ♦ 3PC (Third-Party Copy) ♦ MSC (Media Server Copy) ♦ NAS (Network Attached Storage) ♦ Alt. Client 3PC (Alternate Client Backup used with Third-Party Copy) ♦ Alt. Client MSC (Alternate Client Backup used with Media Server Copy) ♦ Alt. Client NAS (Alternate Client Backup used with NAS)
Operation	For Active jobs, this indicates the operation that is currently being performed.
Owner	The owner of the job.
Parent JobID	Indicates the parent job ID of a composite job. Vault, for example, is a composite job which consists of a single parent, followed by multiple child jobs. The Parent JobID number is followed in sequence by the child job ID numbers. Each job is represented in an individual line in the Jobs tab.
% Complete (Estimated)	The percentage of the job that is complete. For backups, it is based on the size of the previous backup for the same policy, client, schedule, and retention period. If there is no previous backup that matches this criteria then NetBackup does not provide an estimate. If the current backup is larger, this indication is 100%. For other types of jobs, the estimate is based on other factors.
Policy	The name of the policy that NetBackup is using to back up the client. If the policy is associated with a disk staging storage unit, the name follows the convention: <code>__DSSU_POLICY_storageunitname</code> .

Jobs Tab

Job Details

Profile	The profile defines the processing to be done by a Vault job. Multiple profiles can be configured for the Vault.
Robot	For a Vault job, the name of the robot with which Vault is associated.
Schedule	The name of the schedule that NetBackup is using to back up the client.
Session ID	The session ID, a unique numeric value, for a Vault session. Session ID assignment starts at 1 the first time a Vault session is run after Vault has been installed. The value increments by one every time a new Vault session runs.
Start Time	The date and time that the first attempt was initially queued.
Status	<p>Status code and text describing the completion status of each job attempt. A status of zero (0) means that the job completed successfully. Any other completion value for status indicates a problem.</p> <p>REQUESTING_RESOURCE displays as a job requiring resources is making an initial request for resources. Examples of jobs requiring resources are: backup jobs, restore jobs, media jobs, and jobs created using <code>tpreq</code>, <code>bplabel</code>, <code>tpclean</code>, and <code>vmphyinv</code>.</p> <p>GRANTED_RESOURCE displays as resources are allocated.</p> <p>AWAITING_RESOURCE displays if the job is queued due to unavailability of resources.</p> <p>Parent job details display three entries for each operation that the parent job executes:</p> <ul style="list-style-type: none"> ♦ the operation step being executed ♦ the status of the operation once it completes ♦ the time the operation ended <p>For example: The Detailed Status tab contains <i>Parent Job</i> as the first operation performed. The status associated with the parent job operation indicates the final job status returned to Job Manager (<code>nbjrm</code>):</p> <pre>BEGIN_OPERATION Parent Job STATUS 0 END_OPERATION 1230065211</pre>
Storage Unit	The name of the storage unit that the job is using.
Type	<p>The type of job running:</p> <ul style="list-style-type: none"> ♦ Alternate backup: A backup job that a NetBackup client performs on behalf of another client.

Job Details

- ♦ **Archive:** An archive job initiated by the user through the interface on the client. Files are backed up, then deleted from the local disk.
- ♦ **Backup:** A backup job.
- ♦ **Catalog backup:** Offline, cold or online, hot catalog backups. Online catalog backups are composite jobs, consisting of a parent job and two or more child jobs.
- ♦ **Cleaning:** Job initiated by `tpclean` to automatically clean tape drives.
- ♦ **DQTS:** Job initiated as part of the Enhanced Device Qualification Tool Suite. DQTS commands verify the proper functioning of tape drives and robotic tape libraries.
- ♦ **Duplicate:** Job initiated through the Catalog utility.
- ♦ **Erase:** Job initiated by `bplabel -L` or through the Media Management **Actions** menu to write a NetBackup label on unassigned media.
- ♦ **Import:** Job initiated through the Catalog utility to retrieve expired backup images.
- ♦ **Inventory:** Job initiated by `vmphyinv` to physically inventory the media contents of a robotic library or standalone drive and update the volume database.
- ♦ **Label:** Job initiated by `bplabel` to write a NetBackup label on specified media.
- ♦ **Media Contents Report:** Initiated by `bpmedialist -mcontents`. For more information on this report, see “Media Contents Report” on page 302.
- ♦ **Restore:** Job restoring data to a client.
- ♦ **Synthetic backup:** Job that constructs a full or cumulative backup image by processing component images (previous backups).
A synthetic full backup is a backup assembled from a previous, traditional (non-synthesized) full backup, and subsequent differential backups and/or a cumulative incremental backup.
- ♦ **Tape Formatting:** Job initiated by `tpformat` to write a volume label (including a media ID) on an optical disk platter.
- ♦ **Tape Request:** Job initiated by `tpreq` to request a tape volume for mounting and associate a file name with the assigned drive.

Jobs Tab

Job Details

- ♦ Vault: A composite job, consisting of a parent job and multiple child jobs.
- ♦ Verify: Job initiated through the Catalog utility or `bpverify` to verify the contents of one or more backups by reading the backup volume and comparing its contents to the NetBackup catalog. A Verify job is a composite job, consisting of a parent job and multiple child jobs.

Note `tpreq`, `tpclean`, `DQTS`, `tpformat`, `vmphyinv` and `bpmedialist` will not initiate a job when executed on a back level media server.

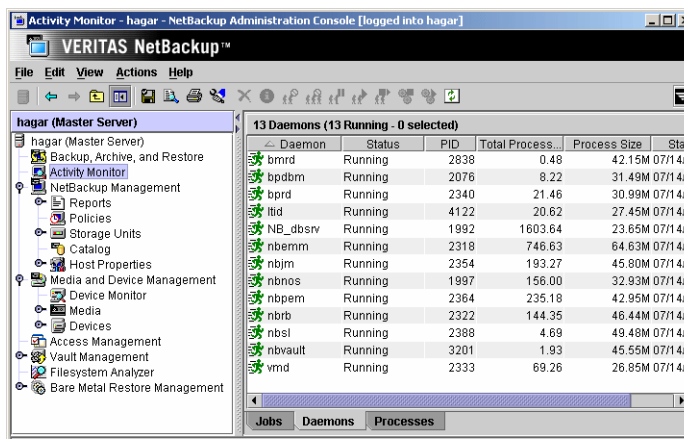
Vault

The name of the logical Vault for a robot configured through the Vault Management node. (Appears for Vault jobs only.)

Daemons Tab

The **Daemons** tab displays the status of NetBackup daemons on the selected master server.

Not all columns are displayed by default. Click **View > Column Layout** to show or hide columns.



NetBackup Daemons

Daemon	Description
Adaptive Server Anywhere - VERITAS_NB (NB_dbdrv)	Sybase ASA database service used by NetBackup for relational DBMS.
NetBackup Bare Metal Restore Master Server (bmrtd)	Appears if Bare Metal Restore is installed.
NetBackup Compatibility Service (bpcompatd)	Service used to communicate with legacy NetBackup services.
NetBackup Database Manager (bpdbm)	Manages the NetBackup internal databases and catalogs. BPDBM must be running on the NetBackup master server during all normal NetBackup operations.
Media Manager Device daemon (ltid)	Stopping and restarting LTID stops and restarts any robotic processes.
NetBackup Enterprise Media Manager (nbemm)	Accesses and manages the database where media and device configuration information is stored (EMM_DATA.db).
NetBackup Client Service (inetd)	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.
NetBackup Job Manager (nbjrm)	Accepts jobs submitted by the Policy Execution Manager (nbpem) and acquires the necessary resources. The Job Manager then starts the job, and informs nbpem that the job is completed.

Daemons Tab

NetBackup Daemons

Daemon	Description
NetBackup Notification Service (nbnos)	Infrastructure service which allows NetBackup components to send and receive events.
NetBackup Policy Execution Manager (nbpem)	Compiles worklist for jobs and determines when jobs are due to run. If a policy is modified or if an image expires, nbpem is notified and the worklist is recompiled.
NetBackup Request Manager (bprd)	Processes requests from NetBackup clients and servers. bprd also prompts NetBackup to perform automatically scheduled backups. bprd must be running on the NetBackup master server in order to perform any backups or restores.
NetBackup Resource Broker (nbrb)	Makes the allocations for storage units, tape drives, client reservations. Works in conjunction with the Enterprise Media Manager (NBEMM).
NetBackup Service Layer (nbsl)	Facilitates communication between the NetBackup graphical user interface and NetBackup logic. NBSL is required in order to run NetBackup Operations Manager (NOM), a managing and monitoring application that works in conjunction with NetBackup.
NetBackup Vault Manager (nbvault)	Manages NetBackup Vault. NBVAULT must be running on the NetBackup Vault server during all NetBackup Vault operations.
NetBackup Volume Manager (vmd)	Keeps track of the location of volumes (tapes) needed for backup or restore.

More About Daemons

Standalone daemons: These NetBackup daemons are always running and listening to accept connections. Examples include bpdbm, bprd, bpjobd, vmd, the robotic services, nbdbd, and visd.

Multi-process standalone daemons: NetBackup daemons that “fork” a child process to handle requests. Examples include bpdbm and bprd.

Single-process standalone daemons: NetBackup daemons that accept connections and handle requests in the same process. Examples include the Media Manager robotic daemons.

inetd daemons: NetBackup daemons that are usually launched by way of inetd(1m) or bpinetd. Examples include bpcd, bpjava-msvc, vopied, and vnetd.

Note After restarting daemons in the Activity Monitor or by using a command, VERITAS recommends exiting all instances of the NetBackup-Java Administration Console, then restarting the console using the `jnbSA` command. (The `jnbSA` command is described in *NetBackup Commands for UNIX and Linux*.)

Other VERITAS Services

There are several daemons that NetBackup uses which do not appear in the Activity Monitor.

VERITAS Services

Service	Description
VERITAS Private Branch Exchange (VxPBX)	VxPBX allows all socket communication to take place while connecting through a single port. VxPBX is installed upon NetBackup installation if it is not already present. VxPBX runs as <code>pbx_exchange.exe</code> on Windows.
NetBackup Monitor Service (nbsvcmon)	Monitors NetBackup services running on the local machine. If a NetBackup service abnormally terminates, the Monitor Service attempts to restart the service. The Monitor Service is the last NetBackup daemon started and the first service to stop.
VERITAS Authentication Service (VRTSat)	VERITAS Authentication Service validates, identities, and forms the basis for authorization and access. One of the VERITAS Security Services (VxSS).
VERITAS Authorization Service (VRTSaz)	VERITAS Authorization Service. One of the VERITAS Security Services (VxSS).

▼ To monitor NetBackup daemons

1. Open the **Activity Monitor** and select the **Daemons** tab.
2. Select the daemon(s) for which you want to view details.
3. Select **Actions > Details**. A Daemons Details dialog appears for each daemon you selected.

Daemons Tab

▼ To start or stop a daemon

1. Open the **Activity Monitor** and select the **Daemons** tab.
2. Select the daemon(s) you want to start or stop.
3. Select **Actions > Start Daemon** or **Actions > Stop Daemon**.

The following table describes fields that appear in the Daemons tab and the Daemon Details dialog:

Daemons Details

Detail	Description
Daemon command	The full path of the command used to start the daemon.
Daemon name	The name of the NetBackup daemon.
Parent process ID	The process ID of the daemon's parent process.
Priority	The priority of the daemon process.
Process ID	The process ID of the daemon.
Process size	The process size of the daemon in kilobytes.
Start time	The date and time when the daemon process was started.
Status	May be <i>Running</i> or <i>Stopped</i> .
Total processor time	The processor time used by the daemon in seconds.
User name	The user name under which the daemon was started.

Processes Tab

The **Processes** tab displays the NetBackup processes running on the selected master server.

Not all columns are displayed by default. Click **View > Column Layout** to show or hide columns.

NetBackup Processes

Process	Description
avrd	The Automatic Volume Recognition process handles automatic volume recognition and label scanning. This allows Media Manager to read labeled tape and optical disk volumes and assign the associated removable media requests to drives.
bmrtd	Process for the NetBackup Bare Metal Restore Master Server daemon.
bpcd	Issues requests to and from the master and media server to start programs on remote hosts.
bpcompatd	Process for the NetBackup Compatibility daemon.
bpdbm	Process for the NetBackup Database Manager daemon. Responds to queries related to the NetBackup catalog.
inetd	Process for the NetBackup Client daemon. Provides a listening service for connection requests.
bpjava-msvc	The NetBackup-Java application server authentication service program. It is started by <code>inetd</code> during startup of the NetBackup-Java GUI applications and authenticates the user that started the NetBackup-Java GUI application.
bpjava-susvc	The NetBackup-Java application server user service program on NetBackup servers. It is started by <code>bpjava-msvc</code> upon successful login via the NetBackup-Java GUI applications login dialog window. <code>bpjava-susvc</code> also services all requests from the NetBackup-Java GUI applications for administration and end-user operations on the host on which the NetBackup-Java application server is running.
bpjobd	Performs queries and updates of the jobs database.
bprd	Process for the NetBackup Request Manager daemon. Starts automatic client backups and responds to client requests for file restores and user backups and archives.

Processes Tab

NetBackup Processes

Process	Description
ltid	Process for the Media Manager Device daemon. See the <i>Media Manager System Administrator's Guide</i> for more information.
NBConsole	The NetBackup Administration Console on the Windows platform.
nbemm	Process for the NetBackup Enterprise Media Manager daemon. Accesses and manages the database where media and device configuration information is stored (EMM_DATA.db).
nbjm	Process for the NetBackup Job Manager daemon. Accepts jobs submitted by the Policy Execution Manager (NBPEM) and acquires the necessary resources. The Job Manager then starts the job, and informs nbpem that the job is completed.
nbnos	Infrastructure service which allows NetBackup components to send and receive events.
nbpem	Process for the NetBackup Policy Execution Manager daemon. Compiles worklist for jobs and determines when jobs are due to run. If a policy is modified or if an image expires, NBPEM is notified and the worklist is recompiled.
nbproxy	Process that allows new multi-threaded NetBackup processes to safely use existing multi-threaded unsafe libraries.
nbrb	Process for the NetBackup Resource Broker daemon. Makes the allocations for storage units, tape drives, client reservations. Works in conjunction with the Enterprise Media Manager (NBEMM).
nbsl	Process for the NetBackup Service Layer daemon. Facilitates communication between the graphical user interface and NetBackup logic.
nbvault	If Vault is installed, the process for the NetBackup Vault Manager daemon.
tl14d, tl18d tl18cd, tldcd tl1dd, tlhd tlhcd, tlmd tshd	Process concerning media control. See the <i>Media Manager System Administrator's Guide</i> for more information.
vmd	Process for the NetBackup Volume Manager daemon.

Monitoring NetBackup Processes

To view the details for a process, double-click on the process in the Processes tab. The Process Details dialog appears.

▼ To monitor NetBackup processes

1. Open the Activity Monitor and select the **Processes** tab.
2. Double-click a process from the process list to view detailed status.

Process Details

The following sections describe fields that appear in the Processes tab and/or the Process Details dialog:

Parent process ID: The unique identifier of this parent process.

Process ID (PID): The unique identifier of this process. Process ID numbers are reused, so they only identify a process for the lifetime of that process.

Process command: The name of the command that initiated the process.

Process name: The name of the currently selected process.

Total Processor Time: Amount of process time (in seconds) that the process has spent.

User ID: The ID of the user who initiated the process.

Process size: The process size of the daemon in kilobytes.

Start time: The date and time when the daemon process was started.

Total user time: Amount of processor time (in seconds) that this process has spent in user mode.

Media Mount Errors

When media is mounted for NetBackup jobs, errors can occur. Depending on the kind of error encountered, a mount request becomes either queued or is cancelled.

Queued Media Mount Errors

When queued, an operator-pending action is created and is displayed in the Device Monitor. This leads to one of the following actions:

- ◆ The mount request is suspended until the condition is resolved.
- ◆ The request is denied by the operator.
- ◆ The media mount timeout is reached.

Cancelled Media Mount Errors

When automatically cancelled, NetBackup tries to select other media to use for backups. (This applies only in the case of backup requests.)

Many conditions lead to the automatic cancelling of the mount request instead of queuing a mount request. This leads to reselection of different media and a stronger likelihood that the backup is not held up.

The following conditions can lead to automatic media reselection:

- ◆ When the requested media is in a DOWN drive.
- ◆ When the requested media is misplaced.
- ◆ When the requested media is write-protected.
- ◆ When the requested media is in a drive not accessible to the media server.
- ◆ When the requested media is in an offline ACS LSM (Automated Cartridge System Library Storage Module). (ACS robot type only.)
- ◆ When the requested media has an unreadable barcode. (ACS robot type only.)
- ◆ When the requested media is in an ACS that is not accessible. (ACS robot type only.)
- ◆ When the requested media has been otherwise determined to be unmountable.

Managing the Jobs Database

NetBackup uses the `/usr/opensv/netbackup/bin/admincmd/bpdbjobs -clean` command to periodically delete done jobs.

By default, the `bpdbjobs` process deletes all done jobs that are more than three days old and retains more recent done jobs until the three-day retention period expires.

If the `bprd` NetBackup request daemon is active, `bprd` starts the `bpdbjobs` process automatically when performing other cleanup tasks. This occurs 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.

Retaining Job Information in the Database

There may be times when it is desirable to keep jobs in the jobs database longer than three days. The default can be changed on a more permanent basis, or temporarily, lasting only until the next reboot or cycling of NetBackup services.

Changing the Default on a Permanent Basis

Since the `bpdbjobs` database resets to default conditions upon reboot or cycling NetBackup Services, you may want a more permanent means of indicating how long to keep jobs in the Activity Monitor.

Add the following entry to the `bp.conf` file:

```
KEEP_JOBS_HOURS = 192
```

Where 192 is the number of hours that all jobs (both successful and unsuccessful) will be kept in the jobs database (or Activity Monitor display).

To retain only successful jobs, add the following entry:

```
KEEP_JOBS_SUCCESSFUL_HOURS = 192
```

Note The retention period values are measured against the time the job ended.

Changing the Default Temporarily

In the absence of a `bp.conf` entry, the `bpdbjobs` process determines how long to retain a job by checking the following locations in the order indicated:

1. The `bpdbjobs` command-line options.
2. The `BPDBJOBS_OPTIONS` environment variable.

Caution Keep in mind that the `bpdbjobs` database resets to default conditions (done jobs deleted after three days) upon reboot or cycling NetBackup Services. If you choose to change the default using a temporary method, you must reinitiate the

Managing the Jobs Database

method after every reboot or each time the NetBackup services are cycled. To change the default on a permanent basis, see “Changing the Default on a Permanent Basis” on page 333.

bpdbjobs Command Line Options

The `bpdbjobs` command interacts with the jobs database to delete or move done job files. The command line options are the first location that the `bpdbjobs` process checks for instructions on retaining jobs.

The `-clean` option causes `bpdbjobs` to delete done jobs older than a specified time period:

```
bpdbjobs -clean [ -M <master servers> ]
[ -keep_hours <hours> ] or [ -keep_days <days> ]
[ -keep_successful_hours <hours> ] or
[ -keep_successful_days <days> ]
```

For example:

```
bpdbjobs -clean -keep_hours 720
```

For a complete description of the `bpdbjobs` command, see *NetBackup Commands for UNIX and Linux*.

BPDBJOBS_OPTIONS Environment Variable

The `BPDBJOBS_OPTIONS` environmental variable provides a convenient way to set job retention options using a script.

The options listed below can be used to determine the length of time NetBackup retains jobs. The options should be entered in lower case in the `BPDBJOBS_OPTIONS` environmental variable:

- ◆ `-keep_hours hours`

Use with the `-clean` option to specify how many hours `bpdbjobs` keeps *unsuccessful* done jobs. Default: 72 hours.

To keep both successful and failed jobs longer than the default of 72 hours, `keep_successful_hours` must be used in conjunction with `keep_hours`

- ◆ `-keep_successful_hours hours`

Use with the `-clean` option to specify how many hours `bpdbjobs` keeps *successful* done jobs. The number of hours can range from 3 to 720 but must be less than or equal to `keep_hours`.

Values outside the range are ignored. Default: 72 hours.

- ◆ `-keep_days days`

Use with the `-clean` option to specify how many days `bpdbjobs` keeps done jobs. Default: 3 days.

- ◆ `keep_successful_days days`

Use with the `-clean` option to specify how many days `bpdbjobs` keeps successful done jobs. Default: 3 days.

This value must be less than the `-keep_days` value.

In the following example, a script (`cleanjobs`) was created which can be copied directly from this document, then changed according to your needs.

- ◆ 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. The correct location of `bpdbjobs` must be indicated. In this example, NetBackup was installed in the default location:

```
setenv BPDBJOBS_OPTIONS "-keep_hours 24 -keep_successful_hours 5
-clean"
/usr/opensv/netbackup/bin/admincmd/bpdbjobs ${*}
```

The `.bat` file can be stored anywhere, as long as it is run from the appropriate directory.

bpdbjobs Debug Log


If you need detailed information on `bpdbjobs` activities, enable the `bpdbjobs` debug log by creating the following directory:

```
/usr/opensv/netbackup/logs/bpdbjobs
```



Note Before using this or other debug logs, read the guidelines in the Debug Logs section of the *NetBackup Troubleshooting Guide for UNIX and Windows*.

Customizing bpdbjobs Output

To customize the output of `bpdbjobs`, add a `BPDBJOBS_COLDEFS` entry to the `bp.conf` file for each column you wish to appear in the output. For more information on the available entries, see the *NetBackup System Administrator's Guide, Volume II*, Chapter 3.



Managing the Jobs Database



Configuring Host Properties

7

This chapter describes the NetBackup property settings and explains how each can be changed for one or more servers or clients. This chapter contains the following sections:

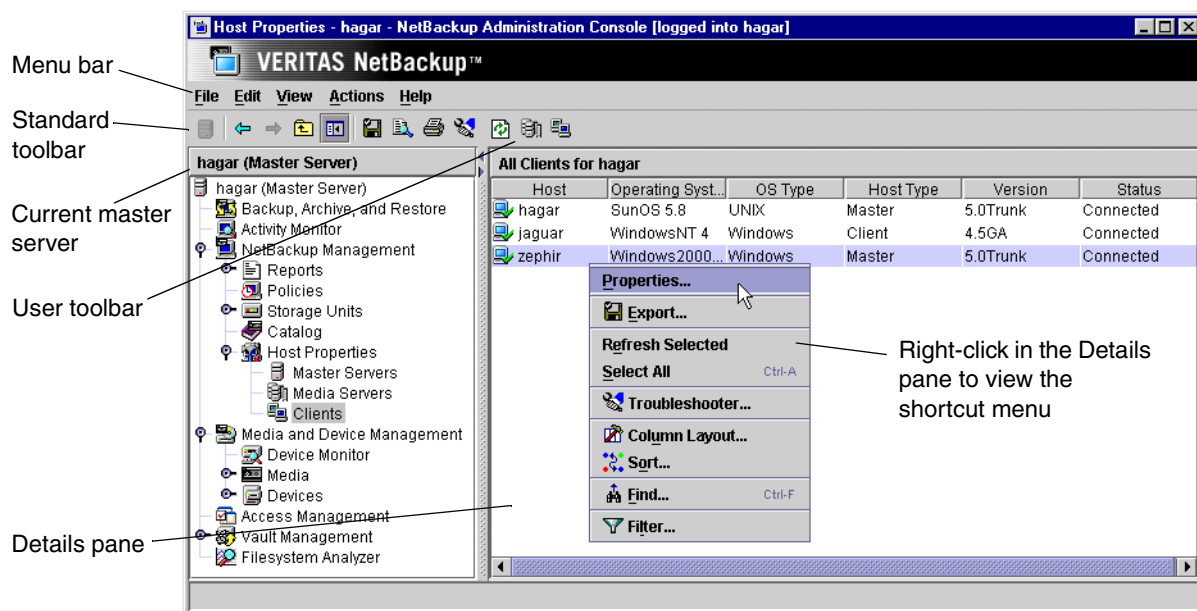
- ◆ “Introduction to Host Properties” on page 338
- ◆ “Changing Host Properties” on page 340
- ◆ “Required Permissions” on page 343
- ◆ “Master Server, Media Server, and Client Host Properties” on page 344

Introduction to Host Properties

Introduction to Host Properties

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

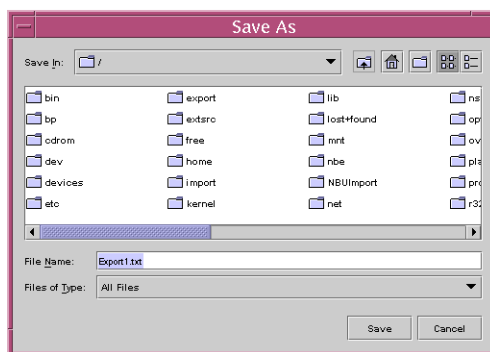
The Host Properties nodes in the Administration Console tree and the Details pane contain the following information:



Host Properties Menu Bar

The **Host Properties** menu bar consists of the following menu items:

Option	Description
File	<p>Options Change Server, New Window from Here, Adjust Application Time Zone, Export, Page Setup, Print Preview, Print, Close Window, and Exit are described in the section, "File Menu" on page 14.</p> <p>Use Export to export the host properties of a host. Expand Master Servers, Media Servers, or Clients and select one or more hosts. Click File > Export. The Save As dialog appears. Enter the full path name and click Save.</p>
Edit	Options Select All and Find are described in the section "Edit Menu" on page 15.
View	Options Show Toolbar , Show Tree , Back , Forward , Up One Level , Options , Refresh Selected , Refresh , Column Layout , Sort , and Filter are described in "View Menu" on page 16.
Actions	<p>The Actions menu contains the following options:</p> <p>Properties: Displays the properties of the host currently selected.</p> <p>Configure Media Server: Select to enter the name of a media server to configure.</p> <p>Configure Client: Select to enter the name of a client to configure. This is a way to configure a client that is not currently included in a policy.</p>
Help	Options Help Topics , Troubleshooter , VERITAS Web Page , License Keys , and About NetBackup Administration Console are described in "Help Menu" on page 17.



Viewing Host Properties

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

Changing Host Properties

▼ To view master server, media server, or client properties

1. In the NetBackup Administration Console, expand **NetBackup Management > Host Properties**.
2. Select **Master Servers**, **Media Servers**, or **Clients**.
3. In the Details pane, click the server or client to view the version and platform. Then, double-click to view the properties.

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

Changing Host Properties

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

Using the NetBackup Administration Console is one way to change the host properties. Another method is by using the `bpgetconfig` command to obtain a list of configuration entries in the `bp.conf` file, then using `bpsetconfig` to change the entries as desired. The commands are described in *NetBackup Commands for UNIX and Linux*. The `bp.conf` entries are described in the *NetBackup System Administrator's Guide, Volume II*.

Interpreting the Initial Settings

The dialogs use specific conventions regarding multiple host selections.

If the focus is on a setting that is set differently between the multiple selected hosts, the following statement appears at the bottom of the dialog: *This value is different on the selected hosts*. This notice is especially helpful regarding differences in text field settings.

Check Box States

The host property check boxes may appear in one of the following three 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 has not been set the same for all selected hosts. When the cursor is moved to such a field, a small notice appears at the bottom of the dialog noting that the value is different on the selected hosts.

States of Multiple Hosts

- ◆ If a dialog contains a **Selected Host** (or similarly named) combo box, all controls on the dialog reflect the values for the host currently selected in the **Selected Host** box.
- ◆ If a dialog does *not* contain a **Selected Host** (or similarly named) combo box, settings of all the selected hosts are combined to arrive at a value that is displayed to the user.

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

Radio Button States

None of the buttons in a radio button group appear selected when multiple hosts are selected. Leaving it in that state keeps the hosts untouched. Selecting any one from the group updates the setting on all selected hosts.

Number Spinner States

A number spinner appears blank when multiple hosts are selected. Leaving it blank keeps the setting untouched on the selected hosts. Changing the value updates the setting on all selected hosts.

Multiple Hosts of Differing Operating Systems

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

For example, if two clients are selected, Linux client apricot and Windows 2000 client grapefruit, neither the Windows Client node nor the UNIX Client node will appear in the Host Properties tree, or any of the sub-nodes. If all the selected hosts are running the same operating systems, the corresponding node and sub-node will appear.


Changing Host Properties

At any time you can choose from the following options:

- ◆ Click **Defaults** to set all the fields on the current dialog to the default values.
- ◆ Click **OK** to apply all changes since **Apply** was last clicked. **OK** also closes the dialog.
- ◆ Click **Cancel** to cancel changes made since the last time changes were applied.
- ◆ Click **Apply** to save changes to all of the properties for the selected host(s).

To make sure that NetBackup uses a changed setting, restart the all daemons and utilities (including the NetBackup Administration Console) to ensure that the new configuration values are used.

- ◆ Click **Help** for information on the properties that appear on the current dialog.



Selecting Multiple Hosts

You may select more than one host in order to change properties on multiple hosts at the same time.

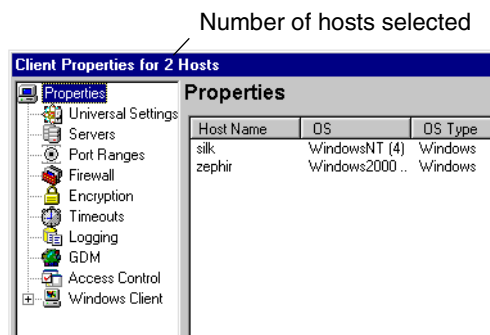
▼ To simultaneously change the properties on multiple hosts

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

The **Properties** dialog appears, displaying the names of the selected hosts that will be affected by subsequent host property changes.

The following information about each selected host is displayed:

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



Required Permissions

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

For example, if you logged on to server *shark* using the NetBackup Administration Console and want to change a setting on a client *tiger*, *tiger* must include *shark* in its Servers List. (See “Adding a NetBackup Server to a Server List” on page 474.)

Note All updates to a destination host (unless it is the same as the host you logged on to using the NetBackup Administration Console) will fail if the target host has placed a check box in **Allow Server File Writes** on the Universal Settings properties. (See “Universal Settings Properties” on page 447.)

Master Server, Media Server, and Client Host Properties

The following sections describe all of the property dialogs that can appear for master servers, media servers, and all supported clients. The description explains if the dialog is available on master servers, media servers, and/or clients. The dialogs are arranged alphabetically.

Access Control Properties

The **Access Control** properties apply to currently selected master servers, media servers, and clients.

VERITAS Security Services (VxSS)

The **VERITAS Security Services** selection determines whether or not the local system uses VxSS.

- ◆ **Required:** Select **Required** if the local system should accept requests only from remote systems using VxSS. Connections from remote systems not using VxSS are rejected. Consider selecting **Required** if all systems are at NetBackup 5.0 or later and maximum security is required.
- ◆ **Prohibit:** Select **Prohibit** if the local system should reject connections from any remote system using VxSS. Consider selecting **Prohibit** if the network is closed and maximum performance is required.
- ◆ **Automatic:** Select **Automatic** if the local system should negotiate with the remote system on whether to use VxSS. Consider selecting **Automatic** if the network contains mixed versions of NetBackup.

VxSS Tab within Access Control Properties Dialog

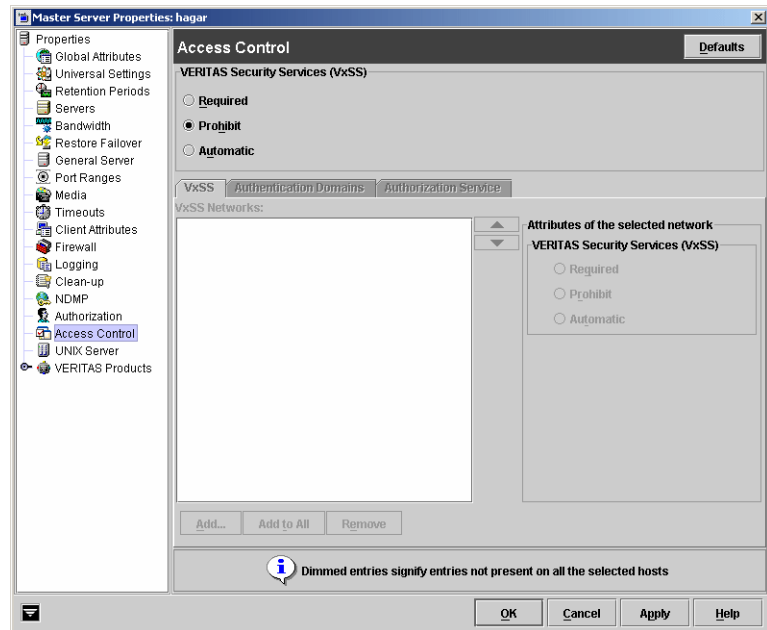
The VxSS tab contains a list of networks that are allowed or (not allowed) to use VxSS with the local system.

VxSS Networks List

The VxSS Networks list indicates whether specific networks can or cannot use VxSS with the local system.

The names on the list are relevant only if the setting above (VERITAS Security Services) is set to **Automatic** or **Required**.

If a media server or client does not define a VxSS network, it will use the VxSS networks of its master server.



Note VERITAS recommends setting the master server VxSS property to **Automatic** until the clients are configured for Access Control. Then, if desired, change the VxSS on the master server to **Required**.

Master Server, Media Server, and Client Host Properties

Add Button

To add a network to the **VxSS Network** list, click **Add**. The **Add VxSS Network** dialog displays, containing the following properties:

Host/Domain

Indicate whether the network to be added is a **Host name** or a **Domain name**.

Host Name/IP

If the network is a host, enter the one of the following:

- ◆ The host name of the remote system. (*host.domain.com*)
- ◆ The IP address of the remote system. (*10.0.0.29*)

Domain Name/IP

If the network is a domain name, enter one of the following:

- ◆ A dot followed by the Internet domain name of the remote systems. (*.domain*)
- ◆ The network of the remote system followed by a dot. (*10.0.0.*)

Bit Count

Select **Bit Count** to indicate that the mask will be 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 **Subnet Mask** to enter a subnet mask in the same the format as the IP address.

Attributes of the Selected Network: VERITAS Security Services

The **VERITAS Security Services** selection determines whether or not the network uses VxSS.

- ◆ **Required:** Select **Required** if the network should accept requests only from remote systems using VxSS. Connections from remote systems not using VxSS are rejected. Consider selecting **Required** if all systems are at NetBackup 5.0 or later and maximum security is required.
- ◆ **Prohibit:** Select **Prohibit** if the network should reject connections from any remote system using VxSS. Consider selecting **Prohibit** if the network is closed and maximum performance is required.
- ◆ **Automatic:** Select **Automatic** if the network should negotiate with the remote system on whether to use VxSS. Consider selecting **Automatic** if the network contains mixed versions of NetBackup.

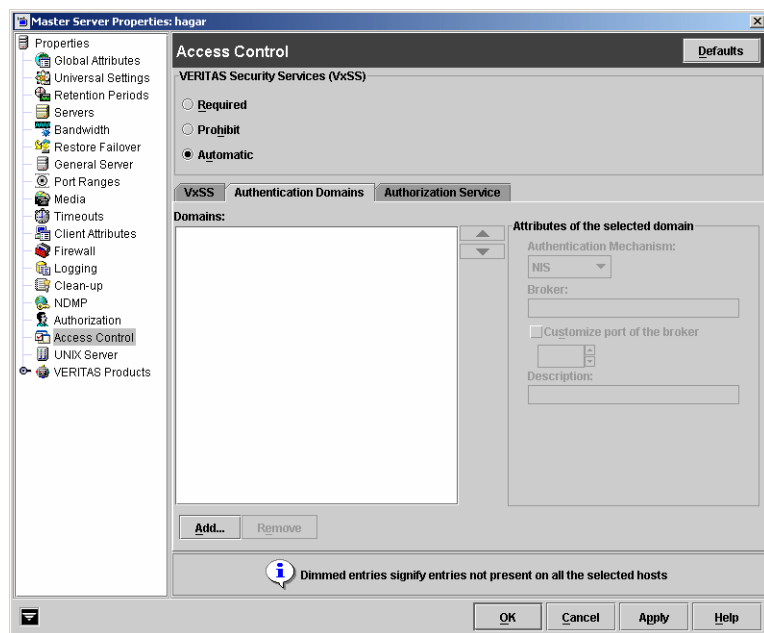
Remove Button

To delete a network, select the network name, then click **Remove**.

Authentication Domain Tab within Access Control Properties Dialog

The Authentication Domain tab contains properties which determine which VxSS authentication broker a machine uses. A master server that uses VxSS must have at least one authentication domain entry.

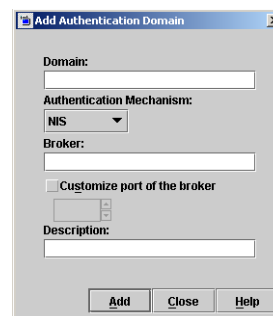
If a media server or client does not define an authentication domain, it will use the authentication domains of its master server.



Master Server, Media Server, and Client Host Properties

Add Button

To add an authentication domain to the domain list, click **Add**. The **Add an Authentication Domain** dialog displays, containing the following properties:



Domain

An Internet or Windows domain name.

Authentication Mechanism

Indicate the authentication mechanism:

NIS: The Network Information Service, version 1.

NIS+: The Network Information Service, version 2.

PASSWD: The local UNIX password file on the specified broker.

VXPD: A VxSS Private Database.

WINDOWS: A Windows Active Directory or Primary Domain Controller.

Note If using a UNIX authentication domain, enter the fully qualified domain name of the host performing the authentication.

Broker

The broker is a machine using an operating system supporting the domain type that has the VxSS Authentication service installed on it.

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

Customize the Port Number of Service

Indicate the port number of the authentication broker, if desired.

Description

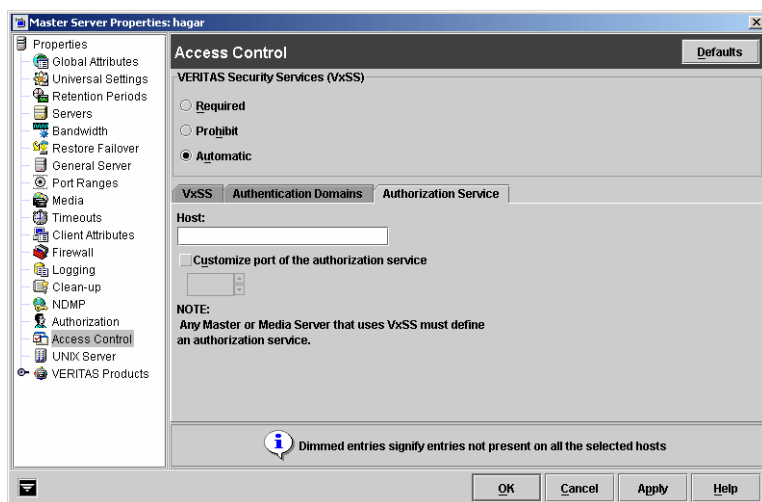
Include a description of the domain, if desired.

Remove Button

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

Authorization Service Tab within Access Control Properties Dialog

The selected **Authorization Service** determines which VxSS authorization service is to be used by the local NetBackup server. The **Authorization Service** tab does not appear as a property for clients.



Note If configuring this tab for a media server using Access Control, you must define the host that will perform authorization.

Host Name

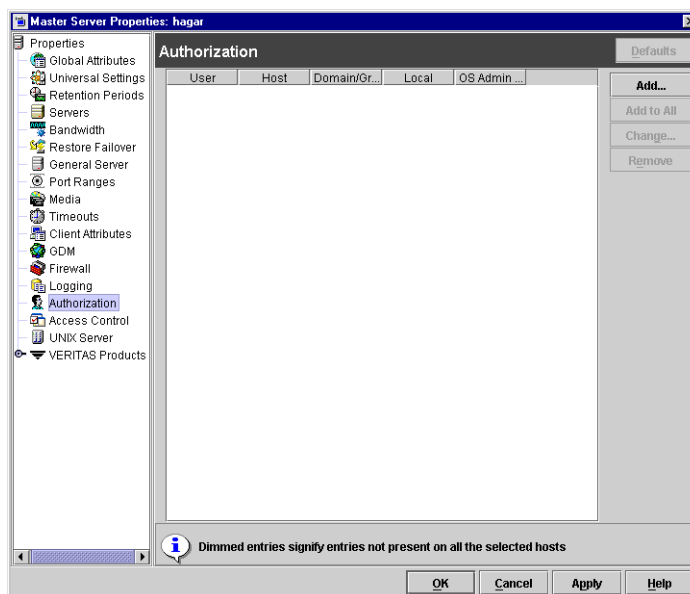
Enter the host name or IP address of the authorization service.

Customize the Port Number of the Authorization Service


To use a non-standard port number, select **Customize the Port Number** and enter the port number of the authorization service.

Authorization Properties

The **Authorization** properties apply to currently selected master servers and media servers.



Click **Add** to add an authorized user, or click **Change** to change the configuration of an existing authorized user. The Add User or Change User dialog appears.



User

In the **User** field, type the name that will identify this user to NetBackup. To indicate any user, enter a single asterisk: *

Host

In the **Host** field, type the name of the remote NetBackup Administration Console host from which this user can use NetBackup. To indicate all hosts, enter a single asterisk: *

Domain\Group

In the **Domain\Group** field, type the Windows domain and group name in the form domain\group or the UNIX local group name or the UNIX netgroup name. Or, enter * to indicate for all groups.

Group/Domain Type

Select whether this user is authorized to use NetBackup in a **Local Group** or a **Network Group**.

User must be an OS Administrator

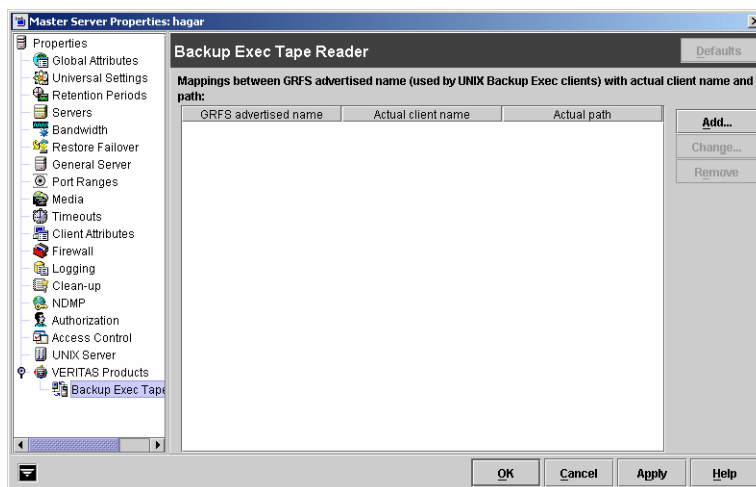
Place a check in the **User must be an OS Administrator** check box to indicate whether the user must be a system administrator of the host from which they are connecting.

For configuration information, see “Enhanced Authentication and Authorization” on page 75 in *NetBackup System Administrator’s Guide, Volume II*.

Backup Exec Tape Reader Properties

The Backup Exec Tape Reader properties apply to currently selected master servers.

The Backup Exec Tape Reader is a feature that enables NetBackup to read media written by Backup Exec. This is done by using a two-phase import process. (See “Importing Images from Backup Exec Media” on page 270.)



Add Button

Click **Add** to enter a GRFS mapping. The Add a GRFS Mapping dialog appears, containing the fields described in the following sections.

GRFS Advertised Name

In order to set the correct client name and paths in Backup Exec UNIX images .f file paths, the master server must be mapped between the **GRFS Advertised Name** (generic file system name) and the actual client name and path.

The **GRFS Advertised Name** uses the following format:

ADVERTISED_HOST_NAME/advertised_path

where *ADVERTISED_HOST_NAME* is the advertised host name and *advertised_path* is the advertised path. The *ADVERTISED_HOST_NAME* should usually be entered in capitals.

The **GRFS Advertised Name** is the name that the Backup Exec UNIX agent (running on the UNIX client machine) used to identify itself to the Backup Exec server. The advertised name may not have been the same as the real machine name and path.

A Backup Exec service had mapped the advertised name to the actual machine name and path, then backed up the *advertised* name and path. When NetBackup imports Backup Exec UNIX backups, the mapping service is not present, so the names and paths must be indicated.

If no entries are indicated in the Backup Exec Tape Reader host properties, NetBackup assumes that the advertised name is the same as the real machine name and the advertised path is the same as the real path.

Actual Client Name

The **Actual Client Name** maps the advertised name to the real machine name.

Actual Path

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

Change Button

Click **Change** to change the selected GRFS entry.

Remove Button

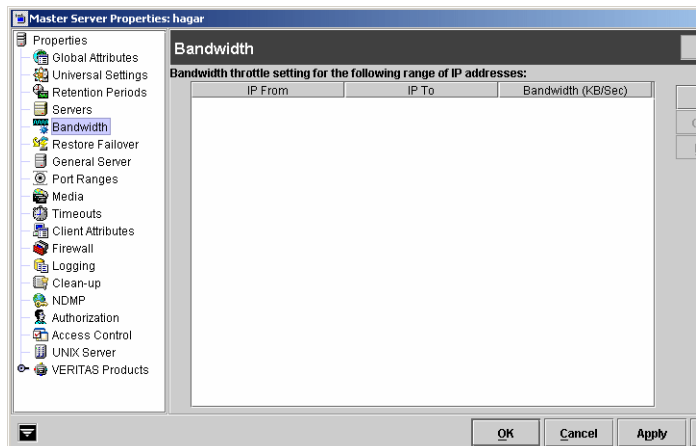
Click **Remove** to remove the selected GRFS entry.

Bandwidth Properties

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

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

Bandwidth limiting only restricts bandwidth during backups.



How Bandwidth Limiting Works

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

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

As the number of active backups increase or decrease on a subnet, NetBackup dynamically adjusts the bandwidth limiting on that subnet. If additional backups are started, the NetBackup server instructs the other NetBackup clients running on that subnet to decrease their bandwidth setting. Similarly, bandwidth per client is increased if the number of clients decreases. Changes to the bandwidth value occur on a periodic basis rather than as backups stop and start. This can reduce the number of bandwidth value changes that are required.

Bandwidth Throttle Setting for the Range of IP Addresses

This area lists the clients in the range of added IP addresses.

From IP Address

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

To IP Address

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

Bandwidth

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

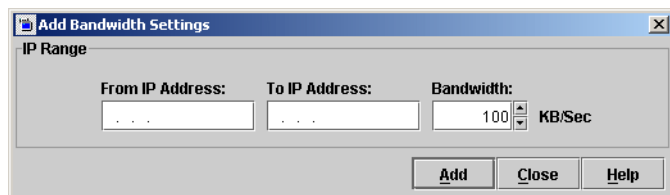
For example, a value of 200 indicates 200 kilobytes per second.

Bandwidth Throttle Settings List

The bandwidth throttle settings list indicates the clients in the range of IP addresses that were added.

Add Button

Click the **Add** button to prepare an entry using the **From**, **To**, and **Bandwidth** fields and add it to the bandwidth table. An entry is added for each of the selected clients.



Remove Button

Click the **Remove** button to remove a selected entry from the bandwidth table.

Notes on Bandwidth Limiting

- ◆ NetBackup does not currently support bandwidth limiting on the following clients:
 - ◆ NetBackup for Oracle clients
 - ◆ NetBackup for DataTools SQL-BackTrack clients

Master Server, Media Server, and Client Host Properties

- ◆ NetBackup for Microsoft SQL-Server clients
- ◆ Bandwidth limiting has no effect on a local backup (where the server is also a client and data does not go over the network).
- ◆ Bandwidth limiting restricts maximum network usage and does not imply required bandwidth. For example, if you set the bandwidth limit for a client to 500 kilobytes per second, the client can use up to that limit. It does not mean, however, that the client requires 500 kilobytes per second.
- ◆ You cannot use bandwidth limiting to load-balance 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.

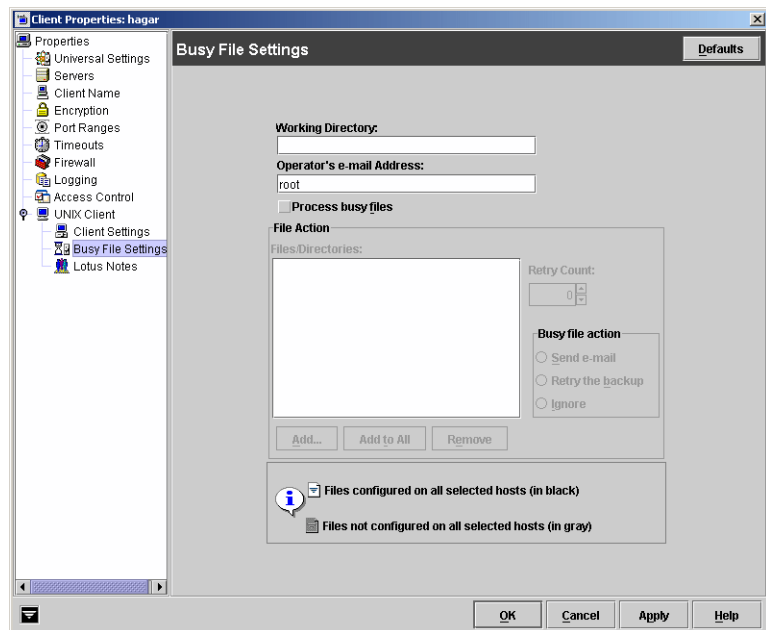
Busy File Properties

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

Working Directory

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

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



Operator's E-mail Address

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

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

Process Busy Files

The **Process Busy Files** property, if checked, causes NetBackup to process busy files according to the settings on this tab, if it determines that a file is changing while it is being backed up. By default, this is not selected and NetBackup does not process the busy files. (See "Busy-File Processing (UNIX Clients Only)" on page 173 in *NetBackup System Administrator's Guide, Volume II*.)

File Action File List

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

Add Button

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

Add to All Button

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

Remove Button

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

Busy File Action

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

- ◆ **Send e-mail:** Directs NetBackup to mail a busy file notification message to the user specified in the **Operator's E-mail Address** field in this dialog.
- ◆ **Retry the Backup:** Directs NetBackup to retry the backup on the specified busy file. The number of times NetBackup will attempt the backup is determined by the **Retry Count** value.
- ◆ **Ignore:** Directs NetBackup to exclude the busy file from busy file processing. The file will be backed up and a log entry indicating that it was busy will appear in the All Log Entries report.

Retry Count

The **Retry Count** property specifies the number of times to attempt the backup. Default retry count: 1.

Clean-up Properties

Clean-up properties apply to:

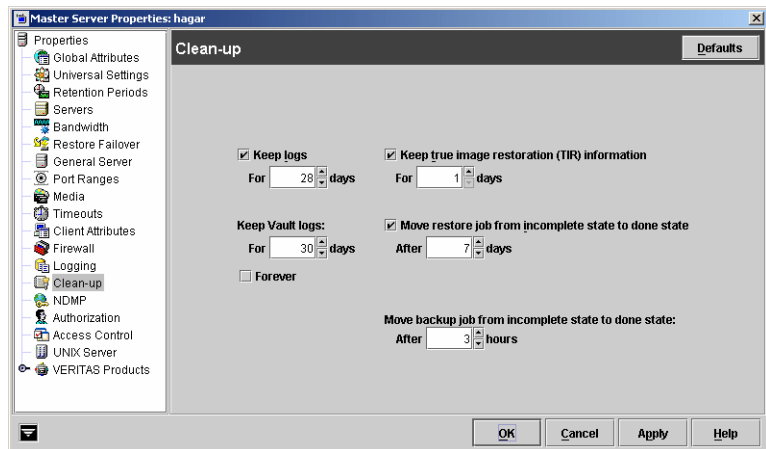
Keep Logs

The **Keep Logs** property specifies the length of time, in days, that the master server keeps its error catalog, job catalog, and debug log information.

NetBackup derives the Backup Status,

Problems, All Log Entries, and Media Log reports from its error catalog, so this attribute limits the time period that these reports can cover. When this time expires, NetBackup also deletes these logs (that exist) on UNIX media servers and UNIX clients.

Specify how many days you'd like 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. Default: 28 days.



Delete Vault Logs

The **Delete Vault Logs** property is enabled if Vault is installed, and specifies the amount of time that the Vault session directories will be kept. Session directories are found in the following location:

```
install_path\netbackup\vault\sessions\vaultname\sidxxxx
```

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

Keep True Image Restoration (TIR) Information

The **Keep True Image Restoration (TIR) Information** property specifies the number of days to keep true image restore information on disk. After the specified number of days, the images are *pruned* (removed). This applies to all policies for which NetBackup is collecting true image restore information. Default: 1 day.

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

Master Server, Media Server, and Client Host Properties

- ◆ Backed up files
- ◆ True image restore information

NetBackup also stores the true image restore information on disk in the `/usr/opensv/netbackup/db/images` directory and keeps it for the number of days specified by this property. Keeping the information on disk speeds up restores. If a user requests a true image restore after the information has been deleted from disk, NetBackup retrieves the required information from the media. The only noticeable difference to the user is a slight increase in total restore time. NetBackup deletes the additional information from disk again after one day.

Move Restore Job From Incomplete State to Done State

The **Move Restore Job From Incomplete State to Done State** property indicates the maximum number of days that a failed restore job can remain in an Incomplete state before 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 utilized, the **Restore Retries** property on the Universal host property dialog allows a failed restore job to be retried automatically. (See “Universal Settings Properties” on page 447 and “Checkpoint Restart for Restore Jobs” on page 518.)

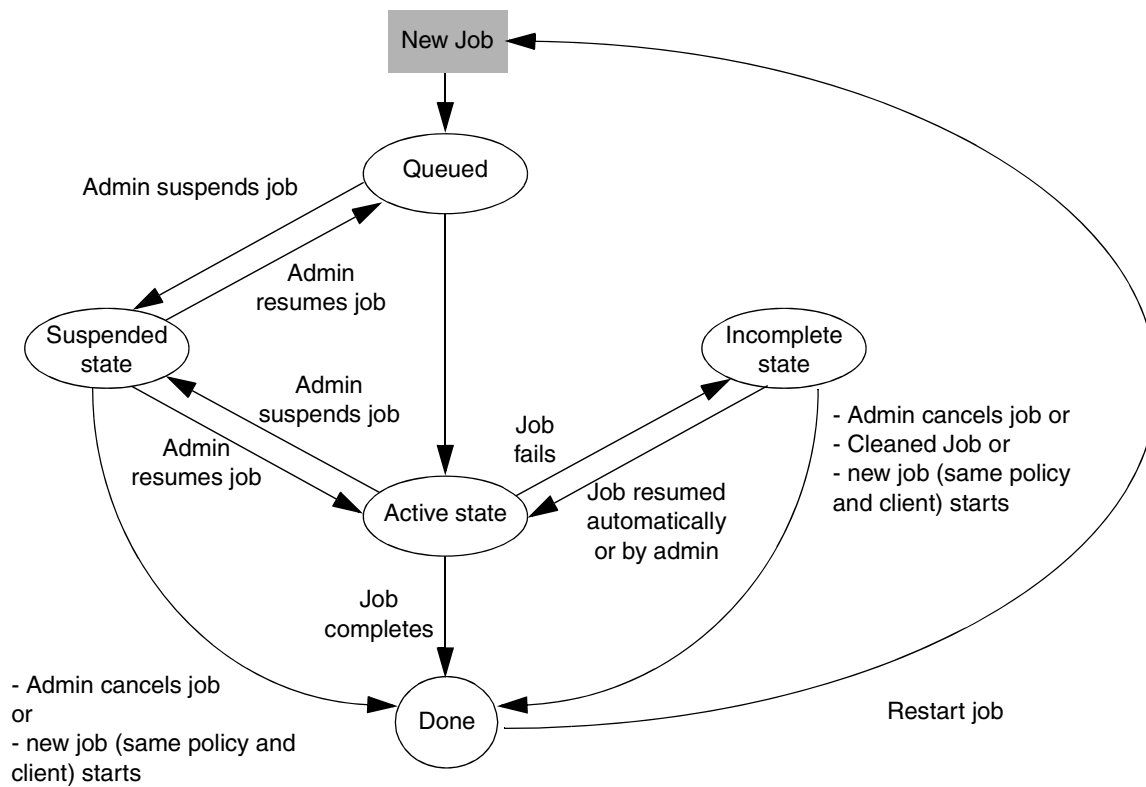
Move Backup Job from Incomplete State to Done State

The **Move Backup Job From Incomplete State to Done State** property indicates the maximum number of hours that a failed backup job can remain in an incomplete state before Activity Monitor shows the job as done. Minimum setting: 1 hour. Maximum setting: 72 hours. Default: 3 hours.

The following figure depicts the different states for a checkpointed backup job:

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

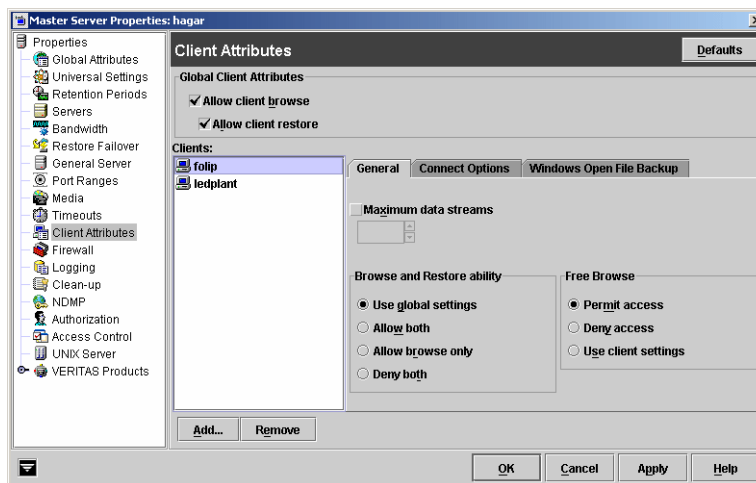
Note A resumed job reuses the same job ID, but a restarted job receives a new job ID. The job details indicate that the job was resumed or restarted.



Client Attributes Properties

Client Attributes properties apply to clients of currently selected master servers. **Client Attributes** contains three subtabs:

- ◆ **General Tab** (described below)
- ◆ **“Connect Options Tab”** on page 364
- ◆ **“Windows Open File Backup Tab”** on page 366



Allow Client Browse

The **Allow Client Browse** property allows all clients to browse files for restoring. This Global client attribute is overridden if, for a particular client(s), the **Browse and Restore Ability** on the General tab in this dialog is set to **Deny both**.

Allow Client Restore

The **Allow Client Restore** property allows all clients to restore files. This Global client attribute is overridden if, for a particular client(s), the Browse and Restore Ability is set to **Allow Browse Only** or **Deny both**.

Clients List

The **Clients** list is a list of clients in the client database on the currently selected master server(s). A client must be in the client database before you are able to change the client properties in the Client Attributes dialog. The client database consists of directories and files in the following directory:

```
/usr/openv/NetBackup/db/client
```

If the desired clients are not listed in the **Clients** list, click **Add** to add clients. To remove a client from the **Clients** list, select the client and click **Remove**.

You can also create, update, list, and delete client entries by using the `bpclient` command located in the following directory:


```
/usr/openv/netbackup/bin/admincmd
```

The name entered here must also match the **Client Name** property for the specific client. If it does not, the client will not be able to browse its own backups. (See “Client Name” on page 370.)

Note If you are using dynamic addressing (DHCP), use the `bpcclient` command to add clients to the client database. (See “Dynamic Host Name and IP Addressing” on page 167 in the *NetBackup System Administrator's Guide, Volume II* for instructions.)

Add Button

Click **Add** to add a client to the client database. Clicking **Add** displays the **Add Client** dialog. Type a client name in the field.



Remove Button

Select a client in the **Clients** list and click **Remove** to delete the selected client from the client database.

General Tab

The following sections describe the properties on the **General** tab within **Client Attributes**. For the properties on the Windows Open File Backup tab, see “Windows Open File Backup Tab” on page 366.

Maximum Data Streams

The **Maximum Data Streams** property specifies the maximum number of jobs allowed at one time for each client selected in the **Client Attributes** host properties tab. (This value applies to the number of jobs on the client, even if multistreaming is not used.)

To change the setting, select **Maximum Data Streams**, then scroll to or enter a value up to 99.

Maximum Data Streams interacts with the **Maximum Jobs Per Client (Host Properties > Master Server > Global Attributes)** and **Limit Jobs Per Policy** (a policy setting) as follows:

- ◆ If **Maximum Data Streams** is *not* set, the limit is either **Maximum Jobs Per Client** or **Limit Jobs Per Policy**, whichever is lower.
- ◆ If **Maximum Data Streams** is set, NetBackup ignores **Maximum Jobs Per Client** and uses either **Maximum Data Streams** or **Limit Jobs Per Policy**, whichever is lower.

Browse and Restore Ability

The **Browse and Restore Ability** property specifies the permissions that clients have for listing and restoring backups and archives. To change the **Browse and Restore Ability** property, select the client(s) in the General tab of the **Client Attributes** dialog and choose the desired action:

- ◆ To use the **Global Client Attribute** settings (“Allow Client Browse” on page 362 and “Allow Client Restore” on page 362), select **Use Global Settings**.
- ◆ To allow users on the selected clients to both browse and restore, select **Allow Both**.
- ◆ To allow users on the selected clients to browse but not restore, select **Allow Browse Only**.
- ◆ To prevent users on the selected clients from browsing or restoring, select **Deny Both**.

Free Browse

This property applies to the privileges allowed to a non-root user logged into the client.

The **Free Browse** property specifies whether the clients selected in the General tab of the **Client Attributes** dialog can list and restore from scheduled backups. (This setting does not affect user backups and archives.)

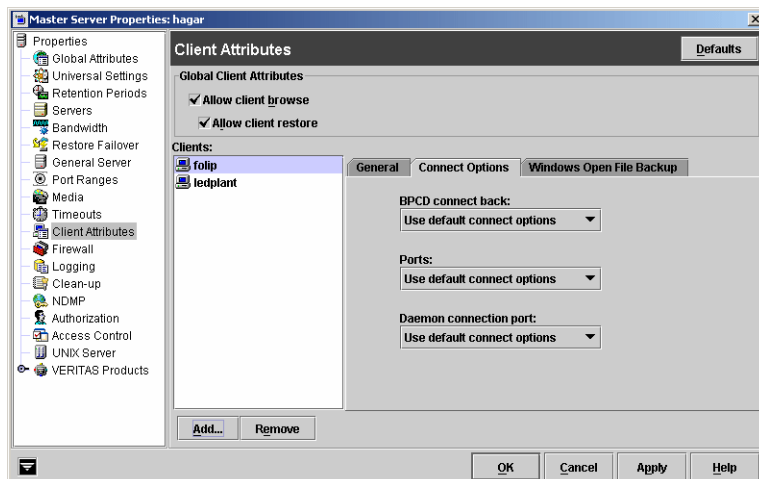
Root users are able to list and restore from scheduled backups as well as user backups regardless of the **Free Browse** setting.

Connect Options Tab

The properties in the **Connect Options** tab describe how a NetBackup server connects to NetBackup clients.

BPCD Connect Back

- ◆ **Use default connect options:**
Use the value



defined in the Firewall host properties of the client's NetBackup server. (See "Default Connect Options" on page 396.)

- ◆ **Random port:** NetBackup randomly chooses a free port in the allowed range to perform the legacy connect-back method.
- ◆ **VNETD port:** NetBackup uses the `vnetd` port number for the connect-back method.

Ports

- ◆ **Use default connect options:** Use the value defined in the Firewall host properties of the client's NetBackup server. (See "Default Connect Options" on page 396.)
- ◆ **Reserved Port:** Use a reserved port number.
- ◆ **Non-reserved port:** Use a non-reserved port number.

Daemon Connection Port

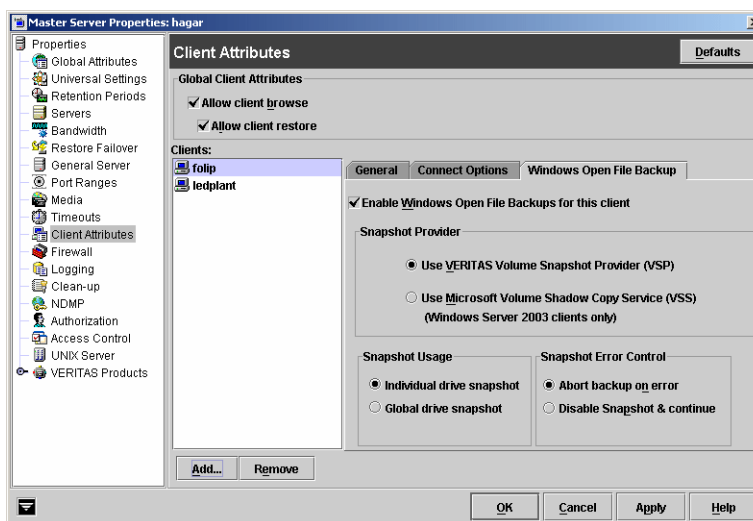
- ◆ **Use default connect options:** Use the value defined in the Firewall host properties of the client's NetBackup server. (See "Default Connect Options" on page 396.)
- ◆ **Automatic:** The daemons on the server will be connected to using `vnetd` if possible. If using `vnetd` is not possible, the connection will be made using the daemon's legacy port number.
- ◆ **VNETD only:** The daemons on the server will be connected to using `vnetd` only. If your firewall rules prevent connecting to the server using the legacy port number, check this option.
- ◆ **Daemon port only:** The daemons on the server will be connected to using only the legacy port number.

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.

Windows Open File Backup Tab

Windows Open File Backup properties apply to selected Windows master servers. The properties appear as a tab on the dialog.

Windows Open File Backup properties specify whether Windows Open File Backup is to be used by a specified client, and whether Volume Snapshot Provider or Volume Shadow Copy Service is to be 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.

Add and Remove Buttons

Click **Add** to add NetBackup clients (5.0 or later) only if you want to change the Windows Open File Backup defaults. By default, no clients are listed and the server uses the following Windows Open File Backup defaults for all Windows NetBackup clients (5.0 or later):

- ◆ Windows Open File Backup is enabled on the client.
- ◆ The snapshot provider for the client is VSP.
- ◆ Snapshots are taken of individual drives as opposed to all drives at once.
- ◆ Upon error, the snapshot is aborted.

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

To make changes to any of the default settings above, add the client name using **Add** and highlight the client name before making changes to the highlighted client's Windows Open File Backup configuration settings in the Windows Open File Backup tab.

Enable Windows Open File Backups for this Client

The **Enable Windows Open File Backups for this Client** property specifies that Windows Open File Backups be used for the clients selected in **Client Attributes**. Add clients to the list only if you want to change the default property settings. (Default: Windows Open File Backup is enabled for all Windows NetBackup clients, 5.0 or later.)

Use VERITAS Volume Snapshot Provider (VSP)

The **Use VERITAS Volume Snapshot Provider (VSP)** property specifies that Volume Snapshot Provider (VSP) be used as the snapshot provider for the clients selected in **Client Attributes**.

VSP is configured for each client using the VSP tab for the client (**Host Properties > Clients > Selected client(s) > Windows Client > VSP**). (See "VSP (Volume Snapshot Provider) Properties" on page 454.)

VSP can be used for Windows NT, Windows 2000, Windows XP and Windows Server 2003 clients. By default, all NetBackup clients (5.0 or later) use VSP as the Windows Open File Backup snapshot provider.

Use Microsoft Volume Shadow Copy Service (VSS)

The **Use Microsoft Volume Shadow Copy Service (VSS)** property specifies that VSS be used to create volume snapshots of volumes and logical drives for the clients selected in **Client Attributes**. VSS can be used for Windows Server 2003 clients only. Configure VSS through the Microsoft's VSS configuration dialogs.

Individual Drive Snapshot

The **Individual Drive Snapshot** property specifies that the snapshot should be of an individual drive. When this property is enabled, snapshot creation and file backup is done sequentially on a per volume basis. For example, assume that drives C and D are being backed up. If **Individual Drive Snapshot** is selected, NetBackup performs the following actions for the backup job:

1. NetBackup takes a snapshot of drive C, backs it up, and discards the snapshot.
2. NetBackup 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 being backed up. This mode is useful when it is not necessary to maintain relationships between files on the different drives. Additionally, this configuration can be used if snapshot creation consistently fails when all volumes for the backup are snapshot at once when the **Global Drive Snapshot** property is enabled. (For example, if one volume in the

volume set has problems meeting the VSP quiet time requirements.) **Individual Drive Snapshot** is enabled by default for all non multi-streamed backups using the Windows Open File Backup option.

Global Drive Snapshot

The **Global Drive Snapshot** property specifies that the snapshot be of a global drive, where all the volumes that require snapshots for the backup job (or stream group for multi-streamed backups) are taken at one time.

For example, assume that drives C and D are being backed up. In this situation, NetBackup performs the following actions:

1. NetBackup takes a snapshot of C and D.
2. NetBackup backs up C, then backs up D.
3. NetBackup discards the C and D snapshots.

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

Note The **Individual Drive Snapshot** and **Global Drive Snapshot** properties only apply to non multi-streamed backups using Windows Open File Backup. All multi-streamed backup jobs share the same volumes snapshots for the volumes in the multi-streamed policy and the volume snapshots are taken in a global fashion (all at once).

Abort Backup on Error

The **Abort Backup on Error** property specifies that a backup aborts if it fails for a snapshot related issue *after* the snapshot is created and while the backup is using the snapshot to back up open or active files on the file system.

The most common reason for a snapshot issue after it has been created and is in use by a backup, is the cache storage filling to capacity. If the backup detects a snapshot issue after it was successfully created and is in use, the backup job aborts with a snapshot error status if **Abort on Error** is checked (default).

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 is only applicable to snapshot errors that occur after the snapshot has been successfully created and is in use by a backup job.

Disable Snapshot and Continue

The **Disable Snapshot and Continue** property 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 the course of the backup—the file may not have been backed up by the backup job and may not be able to be restored.

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

Client Name Properties

The **Client Name** properties apply to a single, currently selected client.

Client Name

The host specified in the **Client Name** field is the NetBackup client name for the selected client. This is the name by which the client is known to NetBackup.

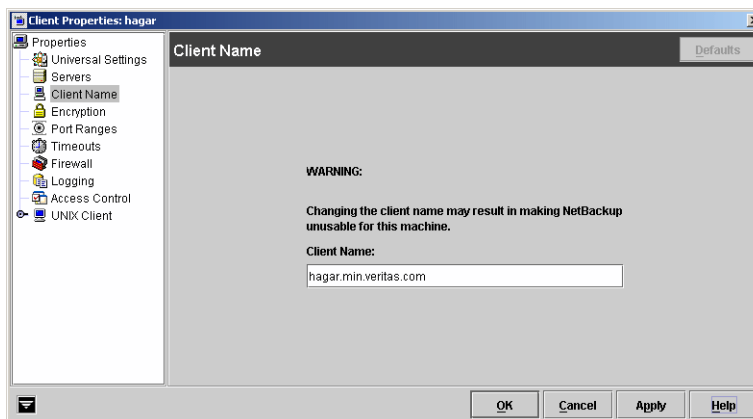
The name must match the name used by the policy that is backing up the client. The only exception is for a redirected restore, where the name must match that of the client whose files are being restored. The client name is initially set during installation.

The name entered here must also match the client name in the Client Attributes dialog for the master server. If it does not, the client will not be able to browse its own backups. (See “Client Attributes Properties” on page 362.)

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

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

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



Client Settings (NetWare) Properties

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

Back Up Migrated Files

The **Back Up Migrated Files** property specifies that files that have been moved to secondary storage will be moved back to primary storage and backed up by NetBackup. If the option is not selected (default), only the metadata for the file is backed up and the file is not moved back to primary storage. The metadata, in this case, is the information that is still in primary storage that marks where the file would be and any information needed to retrieve the file from secondary storage.

Uncompress Files Before Backing Up

The **Uncompress Files Before Backing Up** property specifies that compressed files will be uncompressed before backing up. This is useful if the file will be restored to a version of NetWare that does not support compression. If the option is not selected (default), the file will be backed up in its compressed state.

Keep Status of User-directed Backups, Archives, and Restores

The **Keep Status of User-directed Backups, Archives, and Restores** property specifies the number of days for the system to keep progress reports before automatically deleting the reports. Default: 3 days.

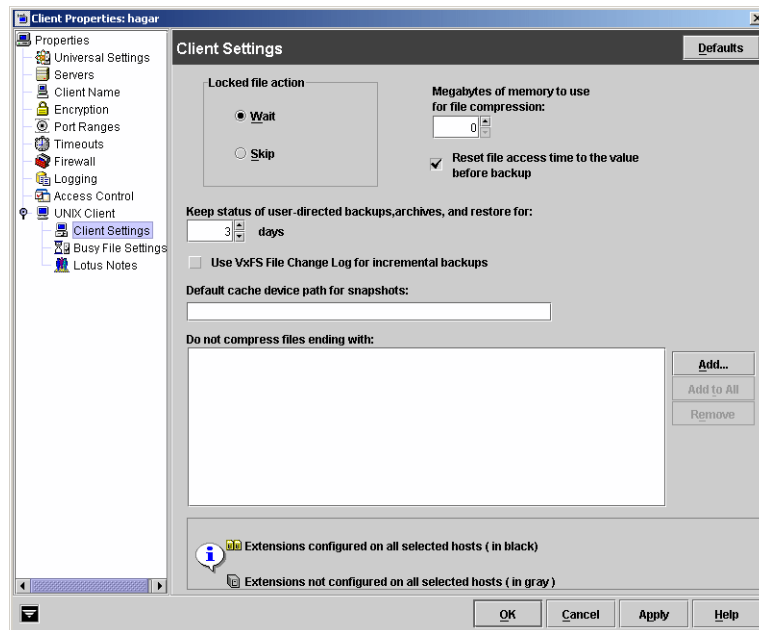
Client Settings (UNIX) Properties

The **UNIX Client** properties apply to currently selected UNIX clients.

Locked File Action

The **Locked File Action** property specifies the behavior of NetBackup when it tries to backup a file that has mandatory file locking enabled in its file mode.

- ◆ **Wait:** By default, NetBackup waits for files to become unlocked. If the wait exceeds the **Client Read Timeout** host property, configured on the master server, the backup fails with a status 41. See “Client Read Timeout” on page 445.
- ◆ **Skip:** NetBackup skips files that currently have mandatory locking set by another process. A message is logged if it was necessary to skip a file.



Keep Status of User-directed Backups, Archives, and Restores

The **Keep Status of User-directed Backups, Archives, and Restores** property specifies the number of days to keep progress reports before the system automatically deletes the reports. Default: 3 days. Minimum: 0. Maximum: 9,999 days.

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

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

Reset File Access Time to the Value Before Backup

The **Reset File Access Time** property specifies that if a file is backed up, its access time (atime) will display the time of the backup. By default, NetBackup preserves the access time by resetting it to the value it had before the backup.

Note This setting affects software and administration scripts that examine a file's access time. DO NOT use this option or `USE_CTIME_FOR_INCREMENTALS` if you are running Storage Migrator on the system. Setting these options causes the atime for files to be updated every time they are backed up. This makes it appear as if the files are frequently used and stops Storage Migrator from selecting them for migration.

Megabytes of Memory to Use for File Compression

Note This option has a reasonable default and should be changed only if problems are encountered.

The **Megabytes of Memory to Use for File Compression** property specifies the amount of memory available on the client to use when compressing files during backup. If you select compression, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to compress code, the greater the compression and the greater the percentage of machine resources used. If other processes also need memory, it is generally best to use a maximum value of 1/2 the actual physical memory on a machine to avoid excessive swapping. Default: 0.

Use VxFS File Change Log for Incremental Backups

The **Use VxFS File Change Log for Incremental Backups** property indicates whether or not NetBackup utilizes the File Change Log (FCL) on VxFS (4.1 or later) clients. This feature is supported on only the Solaris platform in this release. Default: off.

The FCL tracks changes to files and directories in a file system. Changes may include creates, links, unlinks, renaming, data appended, data overwritten, data truncated, extended attribute modifications, holes punched, and file property updates.

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

Recommended Use

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, consider enabling this property for a client containing a very large file system that experiences relatively few changes. The incremental backups of the client could be faster since the policy needs to read only one location, the FCL, to determine what needs to be backed up on the client.

Master Server, Media Server, and Client Host Properties

If, however, there are many changes to one file, or multiple changes to many files, the time saving benefit may not be as great.

Note The VxFS mount point must be specified in the policy backup selections list for this property to take effect. (See “Which Selections Will Be Backed Up: Backup Selections Tab” on page 154.)

Conditions for Use

In order for the **Use VxFS File Change Log** feature to work:

- ◆ The **Use VxFS File Change Log** property must be enabled for every client that wants NetBackup to take advantage of the FCL.
- ◆ The FCL must be enabled on the VxFS client. See the *VERITAS File System Administrator's Guide* for instructions on enabling the FCL on the VxFS client.
- ◆ The **Use VxFS File Change Log** property must be enabled on the client(s) in time for the first full backup. Subsequent incremental backups need this full backup in order to stay synchronized.
- ◆ The VxFS mount point must be specified in the policy backup selections list in some manner:
 - ◆ By specifying ALL_LOCAL_DRIVES.
 - ◆ By specifying the actual FCL mount point.
 - ◆ By specifying a directory at a higher level than the VxFS mount point, provided that **Cross Mount Points** is enabled. (See “Cross Mount Points” on page 82.)

Note If the policy has **Collect True Image Restore Information** or **Collect True Image Restore Information with Move Detection** enabled, the **Use VxFS File Change Log** property on the client is ignored.

Activity Monitor Messages

The Activity Monitor displays messages noting when the file change log is being used during a backup:

Using VxFS File Change Log for backup of *pathname*

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

Keeping the Data Files Synchronized with the FCL

The data files must be in sync with the FCL for this property to work. To keep the data files synchronized with the FCL, avoid turning the FCL on the VxFS client off and on.

Note If any errors are encountered while processing the FCL, NetBackup switches to the normal files system scan. This is displayed in the Activity Monitor.

VxFS Administration

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

Default Cache Device Path for Snapshots

The **Default Cache Device Path for Snapshots** property identifies a raw partition to be used by the copy-on-write process used when either nbu_snap or VxFS_Snapshot are selected as the snapshot method. The partition must exist on all clients included in the policy. For additional information, see the *Advanced Client System Administration Guide*.

Do Not Compress Files Ending With

The **Do Not Compress Files Ending With** list specifies a list of file extensions. During a backup, NetBackup does not compress files with these extensions because the file can already be in a compressed format.

You cannot use wildcards when specifying these extensions. For example, you can specify .A1 but not .A* or .A[1-9]

Files that are already compressed become slightly larger if compressed again. On UNIX clients, if this type of file exists and it has a unique file extension, exclude it (and others with the same extension) from compression by adding it to this list.

Add Button

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

Master Server, Media Server, and Client Host Properties

Add to All Button

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

Remove Button

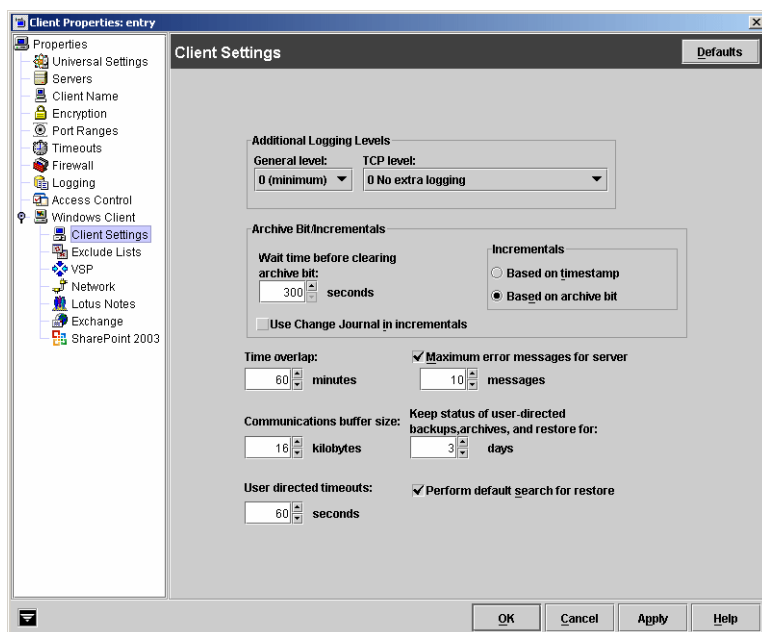
Click the **Remove** button to remove a file ending from the list. To remove a name, either type it in the box or click the browse button (...) and select a file ending. Use commas or spaces to separate names. Then, click the – button.

Client Settings (Windows) Properties

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

General Level Logging

The **General Level Logging** property enables bpinetd, bpbkar, tar, and nbwin logging. Scroll to the desired level of logging. The higher the level, the more information is written. Default: 0.



TCP Level Logging

The **TCP Level Logging** property enables TCP logging. Scroll to the desired level of logging:

- 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/write buffer.

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

Wait Time Before Clearing Archive Bit

The **Wait Time Before Clearing Archive Bit** property specifies the number of seconds the client will wait before clearing the archive bits for a differential incremental backup. The minimum allowable value is 300 (default). The client waits this long for acknowledgment from the server that the backup was successful. If the server does not reply within this time period, the archive bits are not cleared.

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

Use Change Journal in Incrementals

NetBackup offers support for the Microsoft change journal in order to enhance performance of incremental backups on Windows 2000, Windows XP, and Windows Server 2003 systems. By enabling the **Use Change Journal in Incrementals** check box, NetBackup can provide faster incremental backups for NTFS 5 (and later) volumes storing large numbers—possibly millions—of files. **Use Change Journal in Incrementals** is available only when a valid tracker database exists on the applicable volumes. Default: Not enabled.

Enabling **Use Change Journal** automatically enables **Incrementals are based on timestamp**.

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

Determining if enabling change journal support is useful in your NetBackup environment:

Utilizing 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:

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

Unsuitable candidates for enabling NetBackup change journal support:

- ◆ Support for the change journal is intended to reduce scan times for incremental backups by using information gathered from the change journal on a volume. Therefore, enabling NetBackup change journal support is not recommended if the file system on the volume contains relatively few files and folders (hundreds of thousands). 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

- ◆ A NetBackup client utilizing change journal support must belong to only one policy. This avoids the confusion caused by multiple backups setting conflicting update sequence number (USN) information in the permanent record.
- ◆ After selecting **Use Change Journal in Incrementals**, the NetBackup client daemon service must be restarted on the target system. A full backup of the target system must be completed under change journal monitoring to enable change journal-based incrementals.
- ◆ Change journal support is not offered for user-directed backups. The USN stamps for full and incremental backups in the permanent record will not be changed.
- ◆ NetBackup support for change journal works with Checkpoint Restart for restores.
- ◆ Support for change journal is not offered with several NetBackup options or VERITAS products. Enabling the **Use Change Journal in Incrementals** check box in the Windows Client host properties will have no effect while using the following options or products:
 - ◆ True Image Restore (TIR) (See “Collect True Image Restore Information” on page 88.)
 - ◆ True Image Restore with Move Detection (See “Collect True Image Restore With Move Detection” on page 88.)
 - ◆ Synthetic backups (See “Synthetic Backups” on page 107.)
 - ◆ Intelligent Disaster Recovery (IDR) (See the *NetBackup System Administrator's Guide, Volume II*.)
 - ◆ Bare Metal Restore (BMR)

Incrementals Based on Timestamp

The **Incrementals Based on Timestamp** property specifies that files will be selected for backup based on the date that the file was last modified. Selecting **Use Change Journal in Incrementals** automatically selects **Incrementals Based on Timestamp**.

Incrementals Based on Archive Bit

The **Incrementals Based on Archive Bit** property specifies that NetBackup will include files in an incremental backup only if the archive bit of the file is set. The system sets this bit whenever a file is changed and it normally remains set until NetBackup clears it.

A full backup always clears the archive bit. A differential-incremental backup clears the archive bit if the file is successfully backed up within the number of seconds indicated by **Wait Time Before Clearing Archive Bits**. A cumulative-incremental or user backup has no effect on the archive bit.

Clear the **Incrementals Based on Archive Bit** check box to have NetBackup include a file in an incremental backup only if the datetime stamp for the file has been changed since the last backup. For a differential-incremental backup, NetBackup compares the datetime stamp to the last full or incremental backup. For a cumulative-incremental backup, NetBackup compares the timestamp to the last full backup.

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

Note VERITAS 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 (default).

Time Overlap

The **Time Overlap** property specifies the number of minutes to add to the date range for incremental backups when using date-based backups. This value compensates for differences in the speed of the clock between the NetBackup client and server. Default: 60 minutes.

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

Communications Buffer

The **Communications Buffer** property specifies the size (in kilobytes) of the TCP/IP buffers used to transfer data between the NetBackup server and client. For example, specify 10 for a buffer size of 10 kilobytes. The minimum allowable value is 2. There is no maximum allowable value. Default: 16 kilobytes.

User Directed Timeout

The **User Directed Timeout** property specifies the number of seconds that are allowed between the time that a user makes a backup or restore request and when the operation begins. The operation fails if it does not begin within this time period.

There is no minimum or maximum value. Default: 60 seconds.

Maximum Error Messages for Server

The **Maximum Error Messages for Server** property defines the maximum number of times that a NetBackup client will send the same error message to a NetBackup server. For example, if the archive bits cannot be reset on some files, this property limits the number of times the message appears in the logs on the server. Scroll to the desired number. Default: 10.

Keep Status of User-directed Backups, Archives, and Restores

The **Keep Status of User-directed Backups, Archives, and Restores** property specifies the number of days for the system to keep progress reports before automatically deleting them. Default: 3 days.

Perform Default Search for Restore

The **Perform Default Search for Restore** property causes NetBackup to automatically search the default range of backup images and display the backed up folders and files whenever a restore window is opened.

Clear the **Perform Default Search for Restore** check box to disable the initial search. With the property disabled, the NetBackup Restore window does not display any files or folders upon opening. Clicking a backup image, or selecting a range of backup images, starts a search. Default: option is enabled.

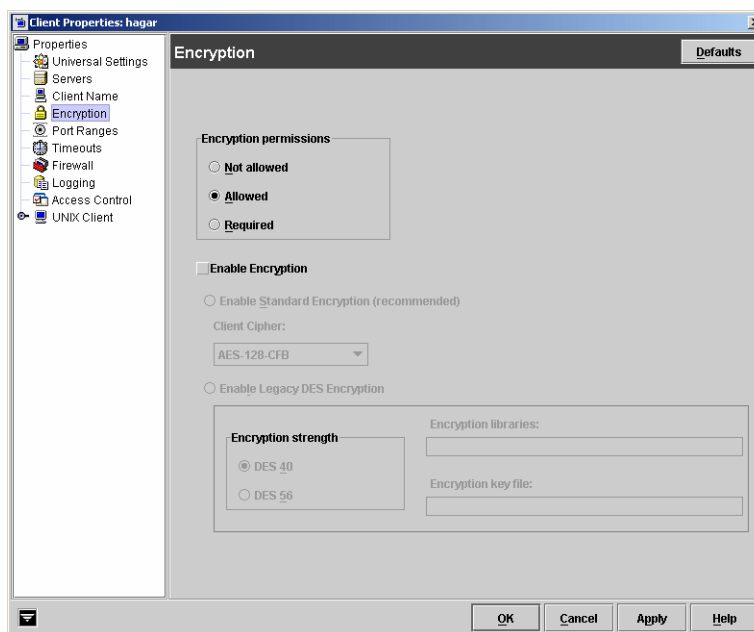
Encryption Properties

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 is hidden.

The separately-priced NetBackup Encryption option must be installed on the client for these

settings (other than **Allowed**) to take effect. For more specific information on the Encryption option, see the *NetBackup Encryption System Administrator's Guide*.



Encryption Permissions

The **Encryption Permissions** property indicates the encryption setting on the selected NetBackup client as determined by the master server. If it is necessary to change this property, click the desired radio button:

- ◆ **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. This is the default setting for a client that has not been configured for encryption.
- ◆ **Required:** Specifies that the client requires encrypted backups. If the server requests an unencrypted backup, the backup job ends due to error.

Enable Encryption

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

Enable Standard Encryption

The **Enable Standard Encryption** property pertains to the 128-bit and 256-bit options of NetBackup Encryption.

If the selected client is running NetBackup 5.1 and is not using Legacy encryption, **Enable Standard Encryption** is automatically selected.

Client Cipher

The following cipher types are available: BF-CFB, DES-EDE-CFB, AES-256-CFB, and AES-128-CFB. AES-128-CFB is the default.

More information on the ciphers file is found in the *NetBackup Encryption System Administrator's Guide*.

Use Legacy DES Encryption

The **Use Legacy DES Encryption** property pertains to 40-bit and 56-bit Data Encryption Standard (DES) NetBackup encryption packages.

If the selected client is running a version of NetBackup earlier than 5.1, **Use Legacy DES Encryption** is automatically selected.

Encryption Strength

The **Encryption Strength** property defines the encryption strength on the NetBackup client when Legacy encryption is being used:

- ◆ **DES_40:** Specifies 40-bit DES encryption. This is the default value for a client that has not been configured for encryption.
- ◆ **DES_56:** Specifies 56-bit DES encryption.

Encryption Libraries

The **Encryption Libraries** property specifies the folder that contains the encryption libraries on NetBackup clients. The default setting is generally sufficient.

The following is the default location:

- ◆ On Windows systems: *install_path*\bin\

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

- ◆ On UNIX systems: /usr/opensv/lib

Master Server, Media Server, and Client Host Properties

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

Encryption Key File

The **Encryption Key File** property specifies the file that contains the encryption keys on NetBackup clients.

The following is the default location:

- ◆ On Windows systems: `install_path\NetBackup\bin\keyfile.dat`
Where *install_path* is the folder where NetBackup is installed and by default is `C:\Program Files\VERITAS`.
- ◆ On UNIX systems: `/usr/opensv/netbackup/keyfile`

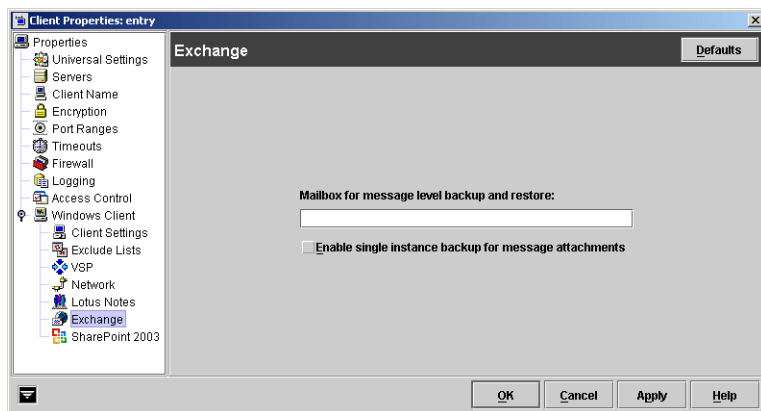
If it is necessary to change the setting, specify the new name.

Exchange Properties

The **Exchange** properties apply to currently selected Windows clients.

The **Exchange** properties contain the setting which defines the mailbox to associate with the NetBackup Client Service account. You must define this mailbox only if the

NetBackup client and NetBackup Microsoft Exchange Server agent software are installed on the Microsoft Exchange Server.



The NetBackup Client Service account must be associated with a valid Exchange mailbox for NetBackup to access the mailboxes and folders during backups and restores. We recommend that you create a uniquely named mailbox for the NetBackup Client service account. If a mailbox is not created for the NetBackup Client service, you can use any existing mailbox on the Microsoft Exchange Server to which the NetBackup Client service account is granted logon rights.

The following section explains the mailbox setting. For more information on this mailbox setting, see the *NetBackup for Microsoft Exchange Server System Administrator's Guide*.

Mailbox for Message Level Backup and Restore

Specifies the mailbox for the NetBackup Client service account. The mailbox can be one of the following:

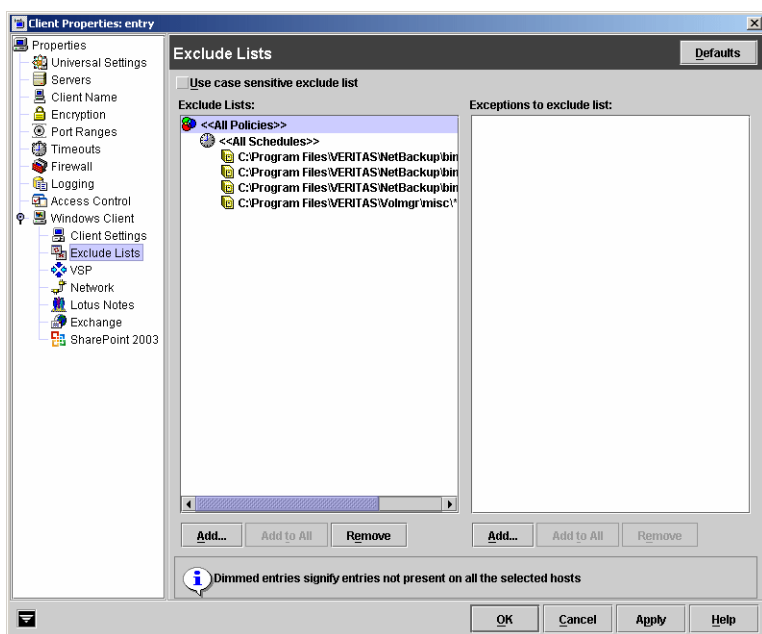
- ◆ An Exchange mailbox name
- ◆ A fully qualified name of the form
`/O=org_name/OU=site_name/CN=server_name/CN=mailbox_name`
- ◆ A mailbox alias

Enable Single Instance Backup for Message Attachments

Microsoft Exchange Server uses single-instance storage (SIS) to store mail messages. This capability in the Exchange Server allows the database to keep one copy of a message attachment sent to multiple users on the same server. To perform SIS backups, check **Enable Single Instance Backup for Message Attachments** on the client where Exchange server is installed.

Exclude Lists Properties

The **Exclude Lists** properties allow you to create and modify exclude lists for Windows clients. An exclude list names policies, schedules, files and directories that you wish to exclude from automatic backups.



Note **Exclude Lists** properties apply only to Windows clients. On NetWare target clients, specify the exclude list (and exceptions) when adding the targets (see the NetBackup user's guide for the client). NetWare NonTarget clients do not support exclude lists. For UNIX clients, see "Excluding Files from Automatic Backups" on page 191.

Use Case Sensitive Exclude List

The **Use Case Sensitive Exclude List** property indicates that the files and directories listed for exclusion/exception are case sensitive.

Exclude List

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

Exceptions to the Exclude List

The **Exceptions to the Exclude List** displays 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 *will* be backed up. This is useful if you want to exclude all files in a directory but one.


Add Buttons

The **Add** button performs different functions, depending on whether it is used from the **Exclude List** or from the **Exceptions to the Exclude List**.

From the Exclude List

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


This means that when the policies on the **Exclude List** run, the files and directories specified on the list *will not* be backed up.



From the Exceptions List

Click **Add** to create an exception to the **Exclude List**. The exception is configured in the **Add Exceptions to Exclude List** dialog, then added to the **Exceptions to the Exclude List**.

This means that when the policies on the **Exceptions to the Exclude List** run, the items on the list *will* be backed up. Effectively, you are adding files back into the backup list of a policy.



Add to All Buttons

The **Add to All** button is enabled only under the following conditions:

- ◆ More than one client is selected for configuration and,
- ◆ a list item is selected that has not been configured on some the selected hosts. (Rather, a grayed-out list item is selected.)

Add to All performs different functions, depending on whether it is used from the **Exclude List** or from the **Exceptions to the Exclude List**.

From the Exclude List

Click **Add to All** to add the selected list item to all currently selected clients. This means that the item will be excluded from the backup list on all selected clients.

From the Exceptions List

Click the **Add to All** button to add the selected list item to the **Exceptions to the Exclude List** of all currently selected clients. This means that when the policies on the **Exceptions to the Exclude List** run, the items on the list *will* be backed up on all selected clients.

Remove Buttons

Remove performs different functions, depending on whether it is used from the **Exclude List** or from the **Exceptions to the Exclude List**.

From the Exclude List

Click **Remove** to remove the selected policy, schedule, or file from the **Exclude List**. The effect is that the item will be *included* in the backup.

From the Exceptions List

Click **Remove** to remove the selected policy, schedule, or file from the **Exceptions List**. The effect is that the item will be *excluded* from the backup.

Shared Fields in Exclude Lists

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

Policy

In the **Policy** field, enter the policy name that contains files and directories that you wish to exclude/except. You can also select the policy name from the drop-down menu. To exclude/except the backup of specific files or directories from all policies, select **<All Policies>**.

Schedule

In the **Schedule** field, enter the schedule name associated with files and directories that you wish to exclude/except. You can also select the schedule name from the drop-down menu. To exclude/except the backup of specific files or directories from all schedules, select **<All Schedules>**.

Files/Directories

In the **Files/Directories** field, enter the full path to the file(s) and directories that you wish to exclude/except.

Exclude Lists for Specific Policies or Schedules

▼ To create an exclude or include list for a specific policy

1. In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Clients**. Double-click on a client.
2. To add an entry to the exclude list:
 - a. Under the Exclude List, click **Add**. The Add to Exclude List dialog appears.
 - b. 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.
 - c. 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.
 - d. In the **Files/Directories** field, enter or browse to the files or directories to be excluded from the backups based on the selected policy and schedule.
 - e. Click **Add** to add the specified files and directories to the exclude list.
3. To add an exception to the exclude list:
 - a. Under the Exceptions to the Exclude List, click **Add**. The Add Exceptions to the Exclude List dialog appears.
 - b. 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.)

- c. 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.
 - d. In the **Files/Directories** field, enter or browse to the files or directories to be added back into the backups based on the selected policy and schedule.
 - e. Click **Add** to add the specified files and directories to the Exceptions to the Exclude List.
4. Click **Apply** to accept the changes. Click **OK** to accept the changes and close the host properties dialog.

Which List is Used If there is More Than One?

If there is more than one exclude or include list for a client, NetBackup uses only the most specific one. For example, assume a client has three 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.

Syntax Rules for Exclude Lists

Note VERITAS 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 before proceeding.

The following syntax rules apply to exclude lists:

- ◆ Only one pattern per line is allowed.
- ◆ The following special or wildcard characters are recognized:
 - []
 - ?
 - *
 - { }
- ◆ To use special or wildcard characters literally (that is, as nonwildcard characters), precede them with a backslash (\). For example, assume the brackets in the following are to be used literally

C:\abc\fun[ny]name

In the exclude list, precede them with a backslash as in

C:\abc\fun\[ny\]name

Note A backslash (\) acts as an escape character only when it precedes a special or wildcard character as in the above example. This means that NetBackup normally interprets a backslash literally and it is a legal character to use in pathnames.

- ◆ 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, just enter the name. For example:

```
test
rather than
C:\test
```

This 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 beginning with a pound sign (#) are ignored.

Windows Client Example Exclude List

Assume that an exclude list contains the following entries:

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

Given the example exclude list, the following files or directories would be 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 below `users` on drive `C`.
- ◆ All files or directories named `temp` that are two levels below the root directory on drive `C`.
- ◆ All files or directories named `core` at any level and on any drive.

Traversing Excluded Directories

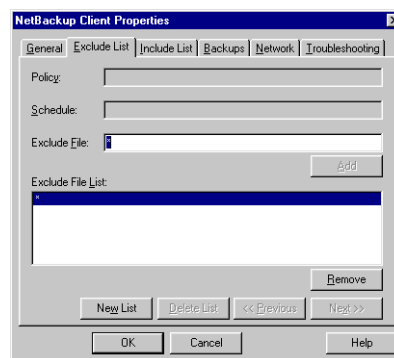
If the exclude list for a client indicates a directory for exclusion, but the client uses an include list to override the exclude list, NetBackup will traverse the excluded directories if necessary, in order to satisfy the client's include list.

Assume the following settings for a Windows client named silk:

- ◆ The backup policy backup selection list for silk indicates ALL_LOCAL_DRIVES. When a scheduled backup runs, the entire client is backed up.
The entire client would also be backed up if the backup selection list consisted of only:
/



- ◆ The exclude list on the client consists of only:
*
This indicates that all files will be excluded from the backup.



- ◆ However, since the include list on Windows client silk includes the following file:

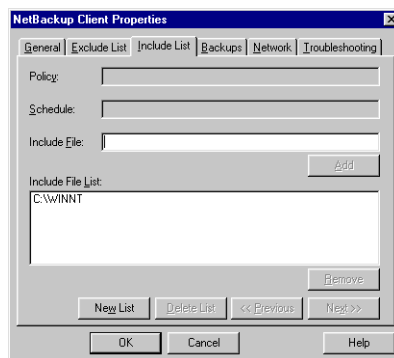
C:\WINNT

the excluded directories are traversed in order to back up C:\WINNT.

If the include list did not contain any entry, no directories would be traversed.

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

- ◆ The backup selection list for client hagar consists of the following: /
- ◆ The exclude list for UNIX client hagar consists of the following: /
- ◆ UNIX client hagar's include list consists of the following directories:
/data1
/data2



Master Server, Media Server, and Client Host Properties

/data3

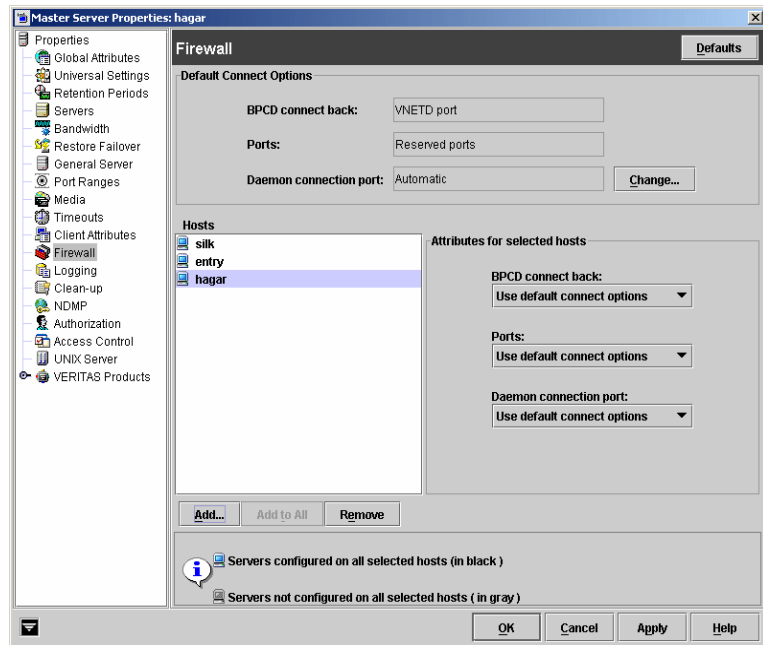
In both examples, because the include list specifies full paths and the exclude list excludes everything, NetBackup will replace the backup selection list with the client's include list.

Firewall Properties

The **Firewall** properties describe how the selected servers are connected to by other hosts.

Firewall-friendly default connect options are configured (**Default Connect Options**), but can be set up for individual servers (**Attributes for Selected Hosts**.)

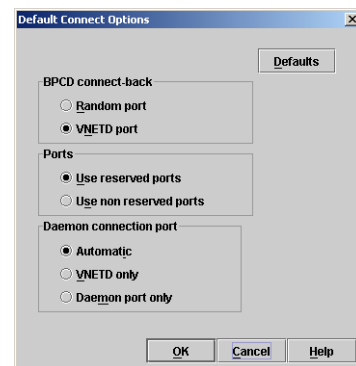
To configure port usage for clients, do so using the properties on **Client Attributes**. (See “Client Attributes Properties” on page 362.) The `bpclient` command can also be used on the master server to configure port usage for clients.



Default Connect Options

By default, the firewall settings are configured to require the fewest possible ports to be open. These properties correspond to the `DEFAULT_CONNECT_OPTIONS` entry in the `bp.conf` file.

To change any of the **Default Connect Options**, click **Change**. The **Default Connect Options** dialog appears containing the following properties:



BPCD Connect-back

The **BPCD Connect-back** property specifies how daemons are to connect back to the NetBackup Client daemon (BPCD):

- ◆ **Random Port:** NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method.

- ◆ **VNETD port:** This method requires no connect-back. The VERITAS Network Daemon (`vnetd`) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. All `bpcd` socket connections are initiated by the server.

For example, when a media server running `bpbrm` initially connects with a client running `bpcd`, the situation does not pose a firewall problem because `bpbrm` is using the well-known `bpcd` port.

Ports

Select whether the server will be connected to using a reserved or non-reserved port number:

- ◆ **Use Reserved Ports:** Connect to the server using a reserved port number.
- ◆ **Use Non-reserved Ports:** Connect to the server using a non-reserved port number. If this property is selected, also enable **Accept Connections from Non-reserved Ports** for the selected server. (See “Accept Connections on Non-reserved Ports” on page 450.) This property is located on the Universal Settings dialog under **Host Properties > Master Servers** or **Host Properties > Media Servers**.

Daemon Connection Port

The **Daemon Connection Port** setting determines which of the following methods will be used when connecting to the server:

- ◆ **Automatic**
The daemons on the server are connected to using `vnetd` if possible. If using `vnetd` is not possible, the connection is made using the daemon's traditional port number. (Automatic is the default.)
- ◆ **VNETD Only**
The daemons on the server are connected to using `vnetd` only. Select this property if your firewall rules prevent connecting to the server using the traditional port number.
- ◆ **Daemon Port Only**
The daemons on the server are connected to using only the traditional port number.

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

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.

Host List

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

Add Button

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

Add to All Button

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

Remove Button

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

Attributes for the Selected Hosts


Connect options can be configured for individual servers.

These properties correspond to the `CONNECT_OPTIONS` entry in the `bp.conf` file.

BPCD Connect-back

The **BPCD Connect-back** property specifies how daemons are to connect back to the NetBackup Client daemon (BPCD):

- ◆ **Use Default Connect Options:** Use the method specified under **Default Connect Options**. (Use **Default Connect Options** is the default for **BPCD Connect-back**.)
- ◆ **Random Port:** NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method.
- ◆ **VNETD port:** This method requires no connect-back. The VERITAS Network Daemon (`vnetd`) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. All `bpcd` socket connections are initiated by the server.




For example, when a media server running `bpbrm` initially connects with a client running `bpcd`, the situation does not pose a firewall problem because `bpbrm` is using the well-known `bpcd` port.

Ports

Select whether the server that will be connected to using a reserved or non-reserved port number:


- ◆ **Use Default Connect Options:** Use the method specified under **Default Attributes**. (Use **Default** is the default.)
- ◆ **Reserved Port:** Connect to the server using a reserved port number.
- ◆ **Non-reserved Port:** Connect to the server using a non-reserved port number. If this property is selected, also enable **Accept Connections from Non-reserved Ports** for the selected server. (See “Accept Connections on Non-reserved Ports” on page 450.) This property is located on the Universal Settings dialog under **Host Properties > Master Servers** or **Host Properties > Media Servers**.



Daemon Connection Port

The **Daemon Connection Port** setting determines which of the following methods will be used when connecting to the server:

- ◆ **Use Default Connect Options:** Use the method specified under **Default Attributes**. (Use **Default** is the default.)
- ◆ **Automatic**
The daemons on the server are connected to using `vneta` if possible. If using `vneta` is not possible, the connection is made using the daemon's traditional port number. (Automatic is the default.)
- ◆ **VNETD Only**
The daemons on the server are connected to using `vneta` only. Select this property if your firewall rules prevent connecting to the server using the traditional port number.
- ◆ **Daemon Port Only**
The daemons on the server are connected to using only the traditional port number.



Master Server, Media Server, and Client Host Properties

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.

NetBackup ports are also discussed in “Configuring NetBackup Ports” on page 530.

Note Both servers and clients must have NetBackup version 4.5 or later installed for *vnetd* to work.

▼ To set up *vnetd* between a server and a client

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

Or, add the client to the client database by running the `bpclient` command, located in `/usr/opensv/netbackup/bin/admincmd` (See “Adding Clients to the NetBackup Client Database” on page 513.)

▼ To set up *vnetd* between servers

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

Or, add a `CONNECT_OPTIONS` entry to `/usr/opensv/netbackup/bp.conf` for each server as described in “`CONNECT_OPTIONS`” on page 129 in the *NetBackup System Administrator's Guide, Volume II*.

▼ To enable logging for *vnetd*

Create a *vnetd* directory in the following location:

On Windows: `install_path\NetBackup\logs\vnetd`

On UNIX: `/usr/opensv/logs/vnetd`

Example Setup for Using the vnetd Port

The following is a sample configuration to use the `vnetd` port for `bprd`, `bpdbm`, `bpjobjd`, `bpvmd` and the robotic daemons on master and media servers and to use **Use Connect-back** `bpcd` connections:

Change in the configuration file setup:

Add the following configuration option to the `vm.conf` file on machines that may connect to `vmd` or the robotic daemons on *hostname*:

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

Where:

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

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

z is 0 for automatic connections. When selected, a `vnetd` style connection is attempted first. If that fails, a traditional connection is attempted.

1 = `vnetd`-only connections.

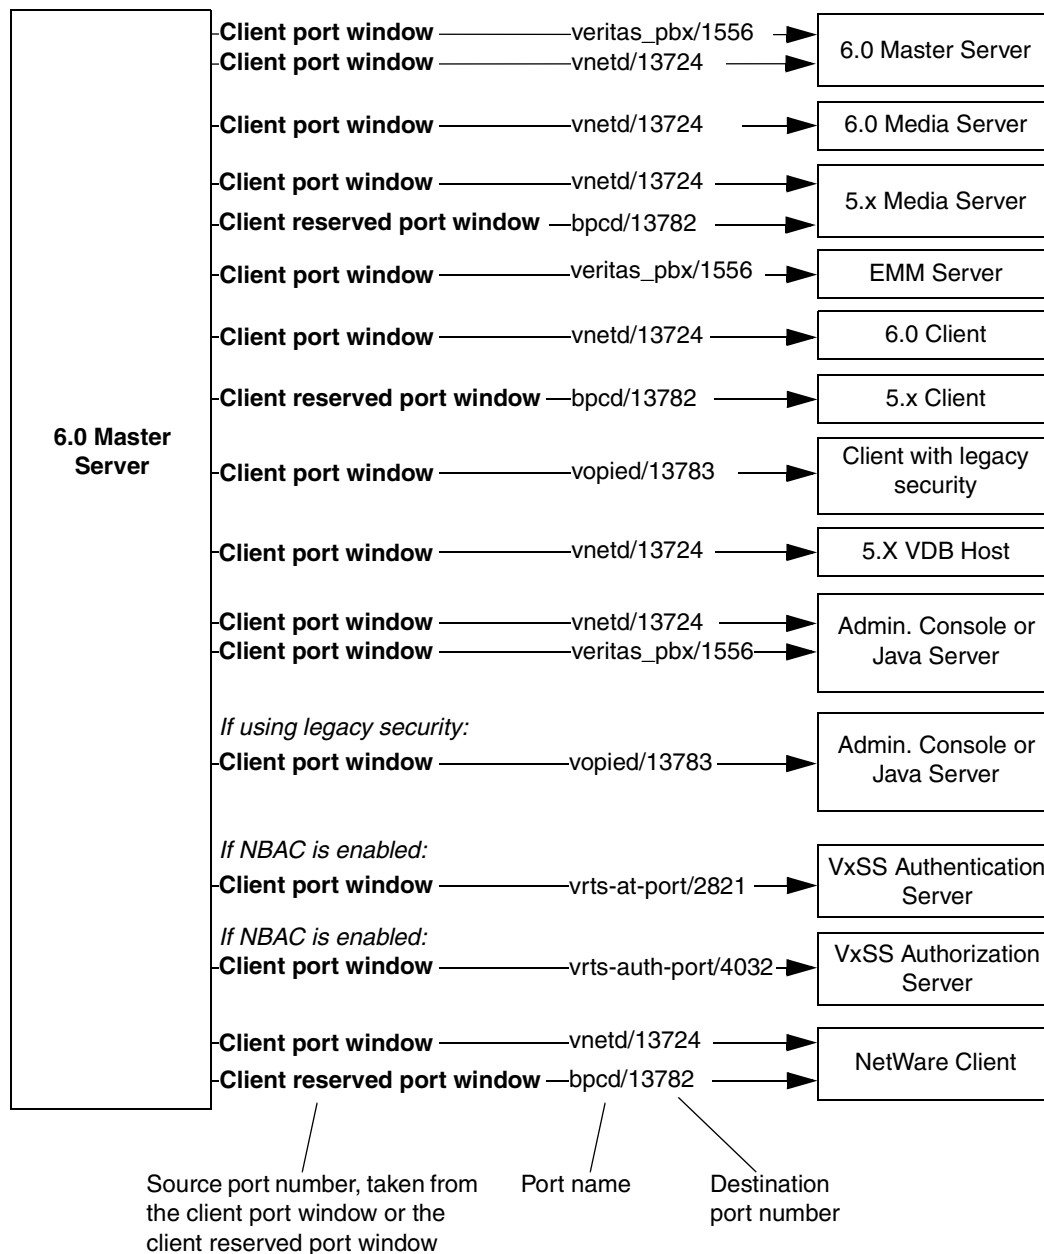
2 = Traditional connections (default)

Change in the Host Properties:

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

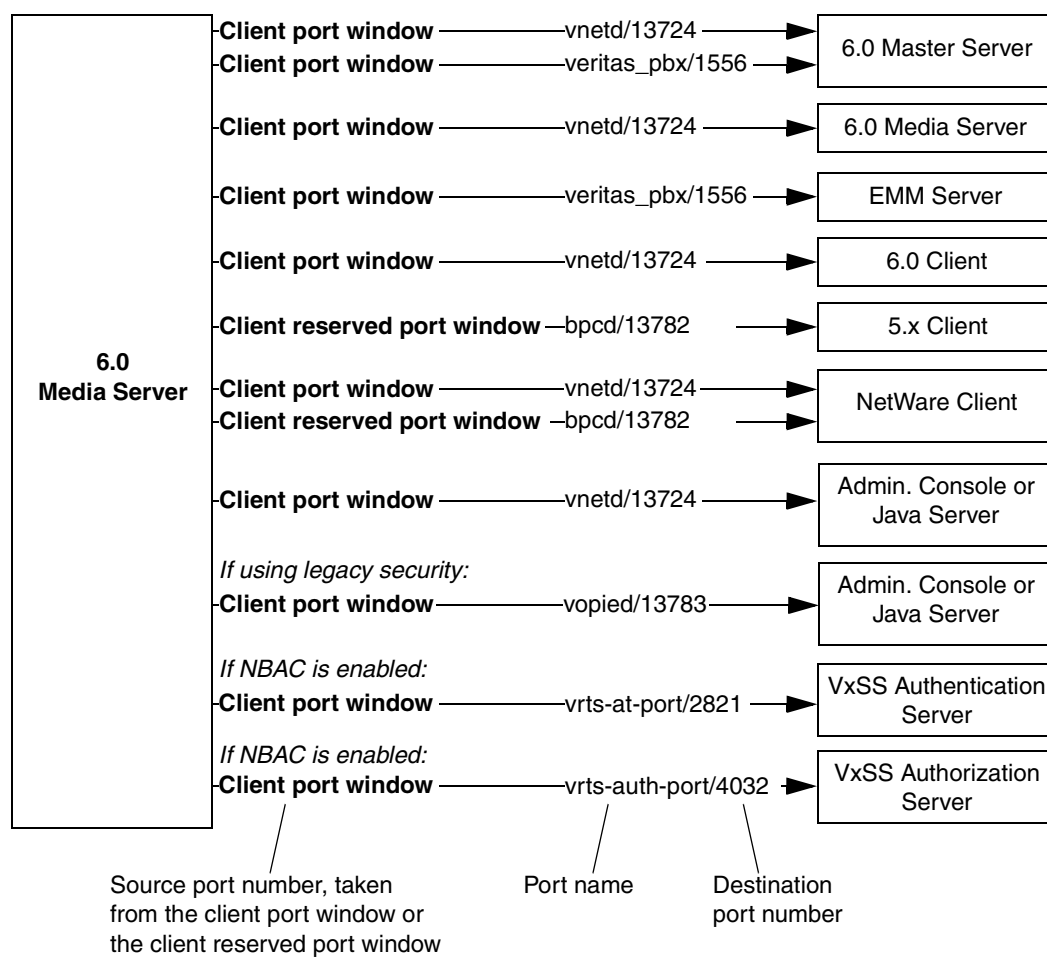
Minimum Master Server Outgoing Connections

6.0 Master Server Minimal Outgoing Port Connections



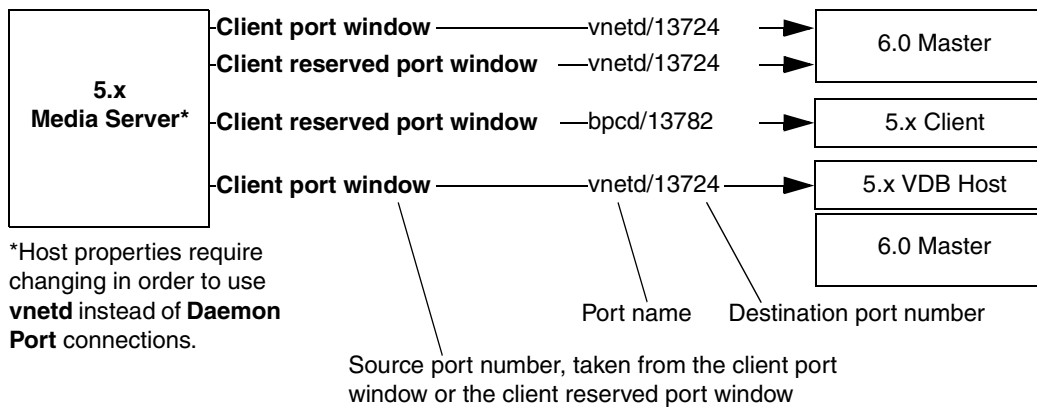
Minimum Media Server Outgoing Connections

6.0 Media Server Minimal Outgoing Port Connections

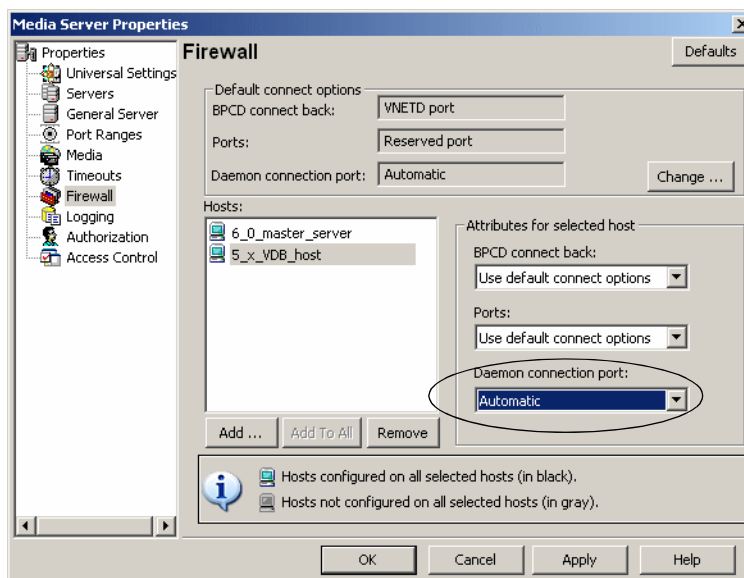


Master Server, Media Server, and Client Host Properties

5.x Media Server Minimal Outgoing Port Connections

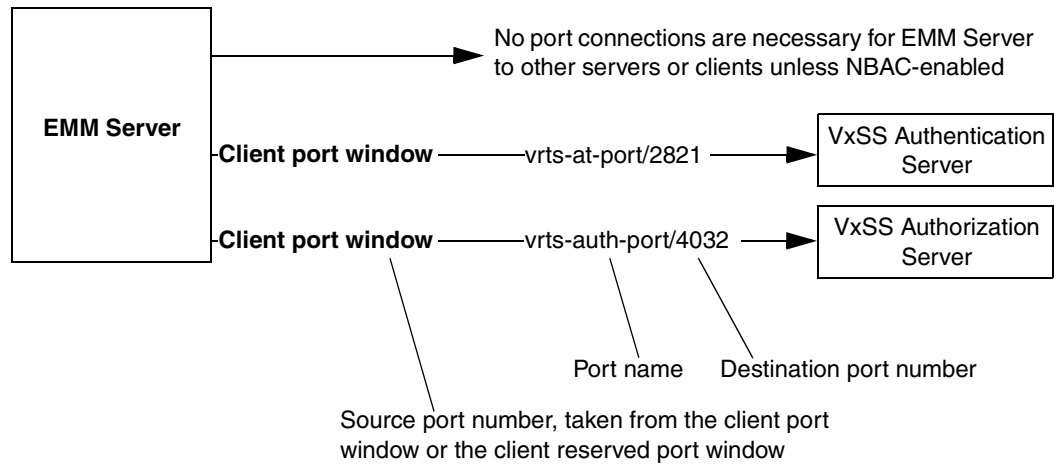


In order for a 5.x media server to use **vnetd** to communicate with a 6.0 master server or a 5.x volume database host, make the following changes in the host properties of the 5.x media server:



Minimum Enterprise Media Server Outgoing Connections

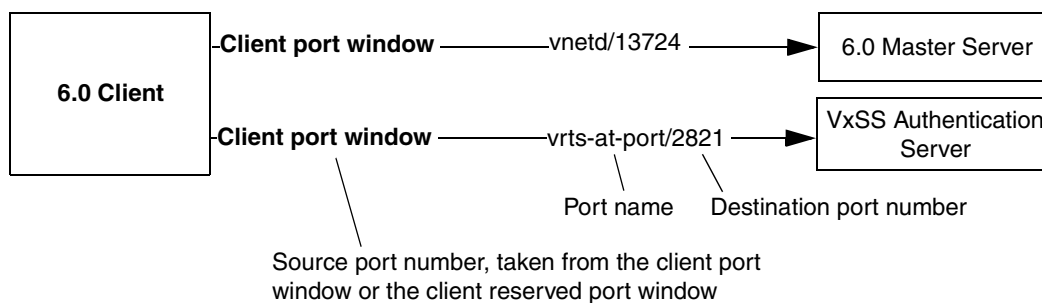
Enterprise Media Manager (EMM) Minimal Outgoing Port Connections




Master Server, Media Server, and Client Host Properties

Minimum Client Outgoing Connections

6.0 Client (Non-NetWare) Minimal Outgoing Port Connections




5.x Client (Non-NetWare) Minimal Outgoing Port Connections



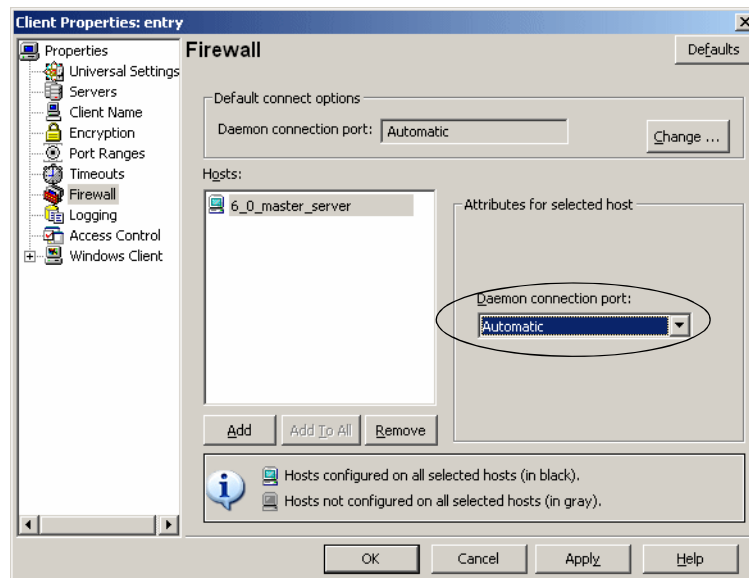
*Host properties require changing in order to use **vnetd** instead of **Daemon Port** connections. (See the next page.)

4.5 Client (Non-NetWare) Minimal Outgoing Port Connections with Mixed Versions




*Host properties require changing in order to use **vnetd** instead of **Daemon Port** connections. (See the next page.)

In order for a pre-6.0 non-NetWare client to use **vnetd** to communicate with a 6.0 master server, make the following changes to the client's Firewall host properties:




6.0 NetWare Client Minimal Outgoing Port Connections



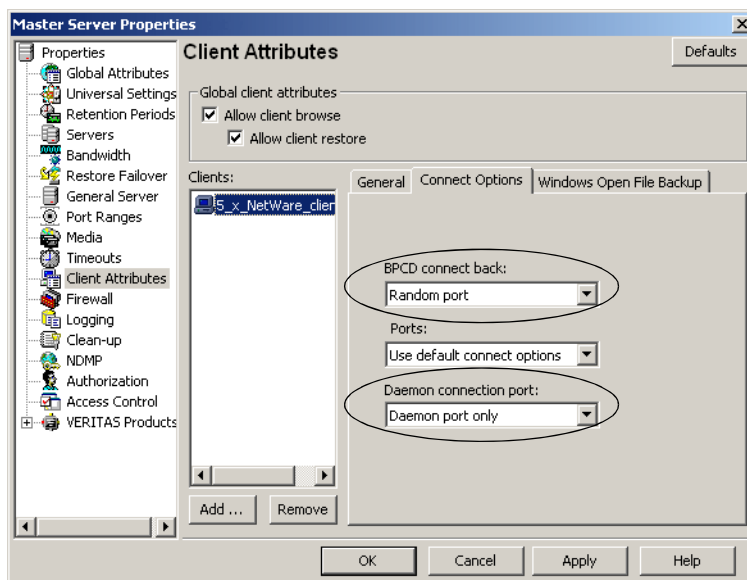
Master Server, Media Server, and Client Host Properties

5.x NetWare Client Minimal Outgoing Port Connections with Mixed Versions



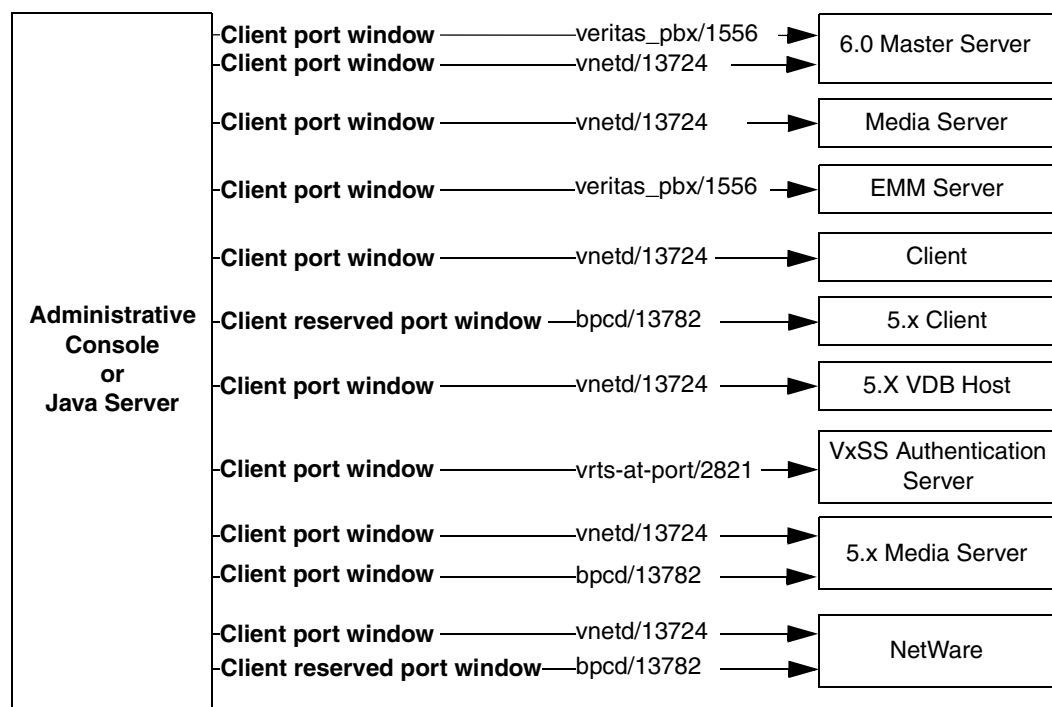
*Host properties require changing in order to use **vnetd** instead of **Demon Port** connections.

In order for a 5.x NetWare client to use **vnetd** to communicate with a 6.0 master server, make the following changes to the 6.0 master server in the **Client Attributes** host properties:




Minimum Java Server or Windows Administration Console Outgoing Connections

Java Server or Windows Administration Console Minimal Outgoing Port Connections

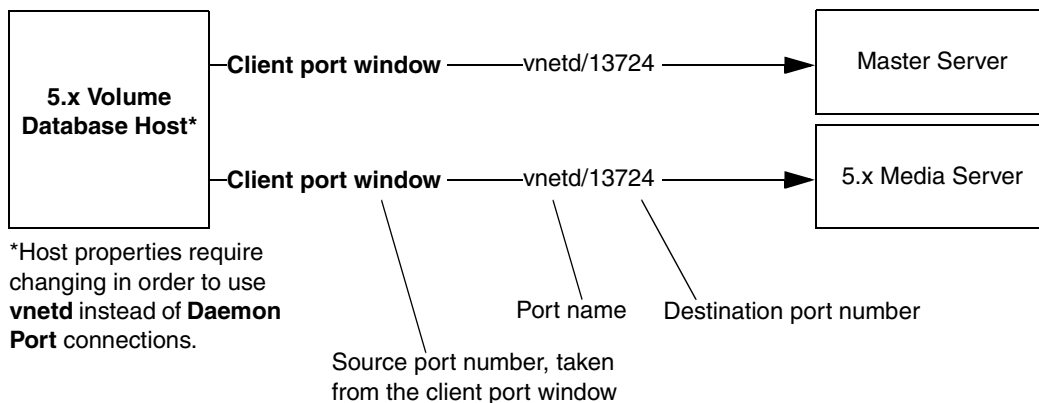


Java Console Minimal Outgoing Port Connections

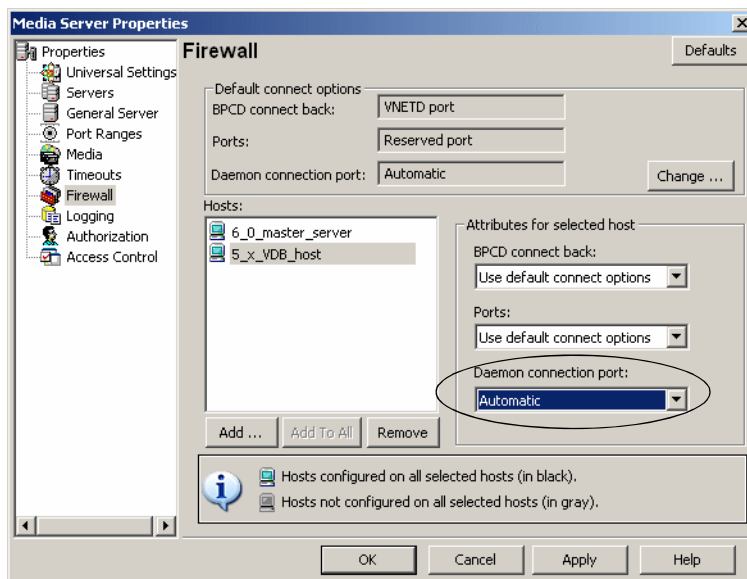


Minimum Volume Database Host Outgoing Connections

5.x Volume Database Host and Mixed Versions Minimal Outgoing Port Connections



In order for a 5.x volume database host to use **vnetd** to communicate with a 6.0 master server or a 5.x volume database host, make the following changes in the host properties of the 5.x volume database host:

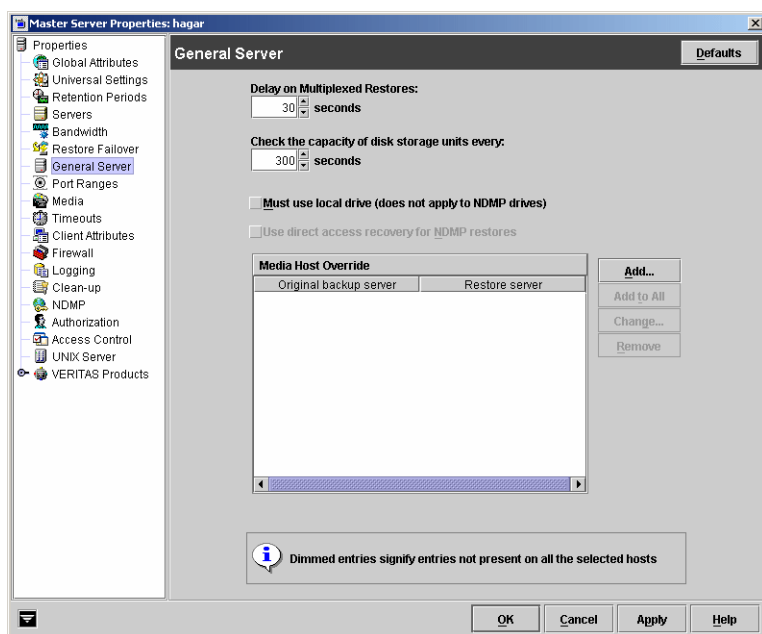


General Server Properties

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

Delay on Multiplexed Restores

The **Delay on Multiplexed Restores** property applies to multiplexed restores. It specifies how many seconds the server waits for additional restore requests of files and/or raw partitions that are in a set of multiplexed images on the same tape. All the restore requests that are received within the delay period are included in the same restore operation (one pass of the tape). Default: delay of 30 seconds.



Check the Capacity of Disk Storage Units

The **Check the Capacity of Disk Storage Units** property determines how often NetBackup checks disk storage units for available capacity. If the frequency is too often, checks are made more often than necessary and system resources are wasted. If the frequency is not often enough, too much time elapses and backup jobs are delayed. Default: 300 seconds (5 minutes).

Must Use Local Drive

This property appears for master servers only.

If the client is also a media server and the **Must Use Local Drive** check box (for the master server) is checked, backups of the client must occur on a local drive. If the client is not a media server, this setting has no effect.

This property increases performance because backups are done locally rather than possibly being sent across the network. For example, in a SAN environment you can create a storage unit for each SAN media server, then mix the media server clients with

Master Server, Media Server, and Client Host Properties

other clients in a policy that uses ANY AVAILABLE storage unit. When a backup starts for a client that is a SAN media server, the backups go to the SAN connected drives on that server.

Use Direct Access Recovery for NDMP Restores

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

Clear the **Direct Access Recovery for NDMP Restores** 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.

Media Host Override

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


The following are some examples of when to use this capability:

- ◆ Two (or more) servers are sharing 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.

Add Button

Click **Add** to add a host to the **Media Host Override** list. The Add Media Override Settings dialog appears containing the following fields:

- ◆ **Original backup server:** Server where data was backed up originally.
- ◆ **Restore server:** Server that is to process future restore requests.



Add to All Button

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

Change Button

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

Remove Button

Select the host in the **Media Host Override** list and click **Remove** to immediately remove the host from the list.

▼ To force restores to go to a specific server

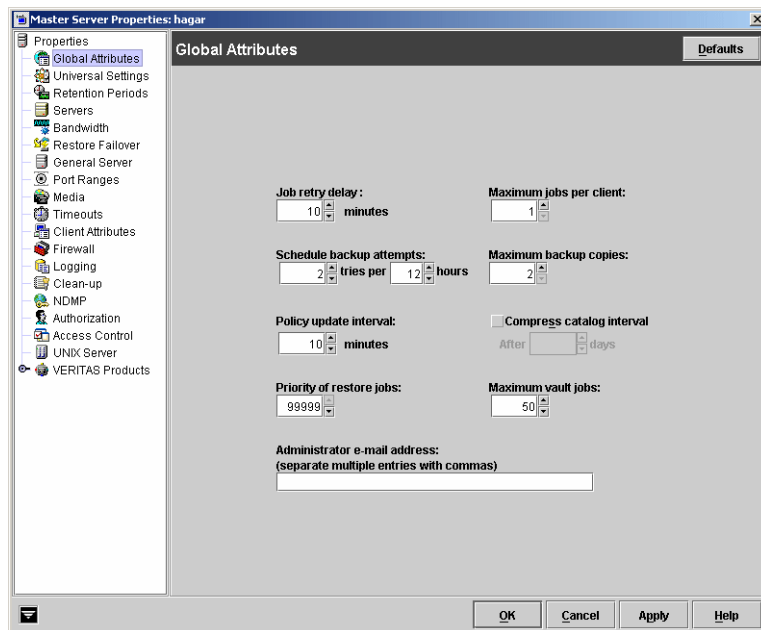
1. If necessary, physically move the media to the host that will be answering the restore requests, then update the Enterprise Media Manager database to reflect the move.
2. Modify the NetBackup configuration on the master server by adding the original backup media server and the restore server to the **Media Host Override** list in the General Server host properties.
3. Stop and restart the NetBackup Request Manager service on the master server.

This applies to all storage units on the original backup server. That is, restores for any storage unit on the server listed as the **Original backup server** will now go to the server 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.



Job Retry Delay

The **Job Retry Delay** property specifies how often NetBackup will retry a job. Default: 10 minutes. Maximum: 60 minutes; minimum: 1 minute.

Schedule Backup Attempts

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

The **Schedule Backup Attempts** property specifies the number of times that NetBackup will try to complete a scheduled backup job during the specified time period. **Schedule Backup Attempts** allows you to limit the number of tries if, for example, a client or drive is down or media is unavailable.

If the backup window closes before the retry starts, the job fails with a status code 196. Default: 2 tries in 12 hours.

Policy Update Interval

The **Policy Update Interval** property specifies the number of minutes to wait after changing a policy before that policy is processed. This allows the NetBackup administrator time to make multiple changes to the policy. Default: 10 minutes. Maximum: 1440 minutes; minimum: 1 minute.

Priority of Restore Jobs

The **Priority of Restore Jobs** property determines if all restore jobs have priority over other types of jobs when contending for drives. The higher the value, the more priority that restore jobs have. Default: 99999.

Note Multiplexed backup jobs are unaffected by the restore priority setting.


Maximum Jobs per Client

The **Maximum Jobs per Client** property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. Default: 1 job.

NetBackup can process concurrent backup jobs from different policies on the same client only if:

- ◆ There is more than one storage unit available, or,
- ◆ one of the available storage units can perform more than one backup at a time.

Files and directories that are on the same client but in different policies, can be backed up concurrently to different storage devices.



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

- ◆ Number of storage devices. NetBackup can perform concurrent backups to separate storage units or to drives within a storage unit. For example, a single Media Manager storage unit supports as many concurrent backups as it has drives. A disk storage unit is a directory on disk so the maximum number of jobs depends on system capabilities.
- ◆ Server and client speed. Too many concurrent backups on an individual client interfere with the performance of the client. The best setting depends on the hardware, operating system, and applications that are running.

Since the **Maximum Jobs per Client** property applies to all clients in all policies, consider accommodating weaker clients (ones that can handle only a small number of jobs concurrently) by using one of the following approaches:

Master Server, Media Server, and Client Host Properties

- ◆ Set the **Maximum Data Streams** property for those weaker client(s) appropriately. (This property is found under **Host Properties > Master Server > Client Attributes > General** tab.)
- ◆ Use the **Limit Jobs Per Policy** policy setting in a client-specific policy (one in which all clients share this characteristic).
- ◆ Network loading. The available bandwidth of the network affects how many backups can occur concurrently. For example, two exabyte 8500, 8mm tape drives can create up to a 900-kilobyte-per-second network load. Depending on other factors, this can be too much for a single Ethernet. If you encounter loading problems, consider backing up over multiple networks or using compression.

A special case exists when backing up a client that is on the same machine as the server. Here, network loading is not a factor because you do not use the network. Client and server loading, however, is still a factor.

Note If online, hot catalog backups are scheduled to occur concurrently with other backup types for the master server, set the **Maximum Jobs per Client** value to greater than two. This ensures that the catalog backup can proceed while the regular backup activity is occurring.

Maximum Backup Copies

The **Maximum Backup Copies** property specifies the total number of backup copies that may exist in the NetBackup catalog (2 through 10). NetBackup creates either the number of copies specified under **Multiple Copies**, or the number of copies specified as the **Maximum Backup Copies** property, whichever is smaller. For more information on creating multiple copies, see “Multiple Copies” on page 110.

Compress Catalog Interval

The **Compress Catalog Interval** property specifies the number of days that NetBackup waits after a backup before compressing the image catalog file that contains information about the backup.

Maximum Vault Jobs

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

If the limit on the number of active vault jobs is reached, subsequent vault jobs are queued and their status is shown as *Queued* in the Activity Monitor.

If a job is waiting to perform duplication or eject, its status is shown as *Active* in the Activity Monitor.

Administrator's E-mail Address

The **Administrator's E-mail Address** property specifies the address(es) where NetBackup sends notifications of scheduled backups or administrator-directed manual backups. The notification of offline, cold catalog backups includes the media ID that was used.

To send the information to more than one administrator, separate multiple e-mail addresses using a comma:

email1,email2

Note Disaster recovery information created during online, hot catalog backups is not sent to the addresses indicated here. DR information is sent to the address indicated on the Disaster Recovery tab in the catalog backup policy. See "Where Will the Catalog Data Be Located: Disaster Recovery Tab" on page 194.

Setting Up E-Mail Notifications

You may need to configure the computing environment so that notification e-mail from NetBackup functions properly.

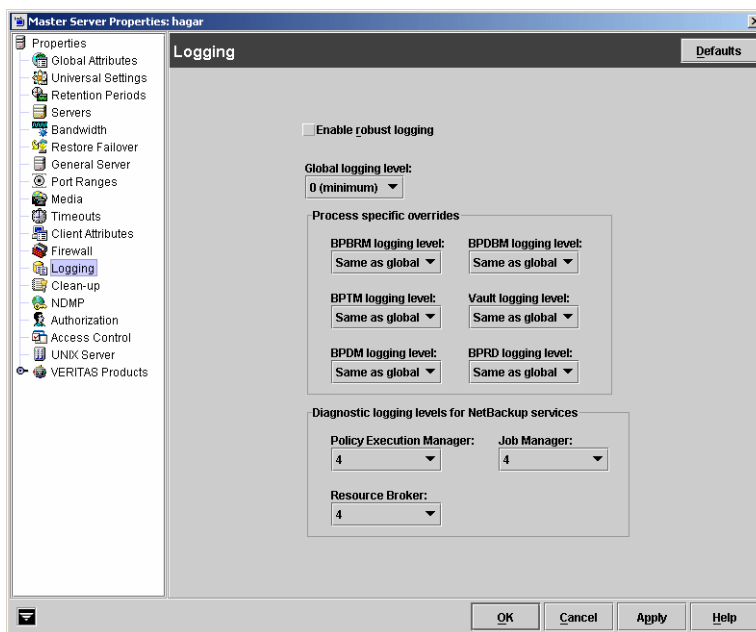
NetBackup uses the mail transfer agent `sendmail` to send e-mail notifications. If it is not installed, install it from `sendmail.org` and configure your environment accordingly so that it functions correctly.

Logging Properties

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

Types of Logging

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



Introducing Unified Logging

NetBackup introduces the use of unified logging, in which log file names and messages are created in a format that is standardized across all VERITAS products. Unified logging is used by certain NetBackup processes, primarily on the server.

The unified logs are written to `/usr/openv/logs` (UNIX) and to `install_path\NetBackup\logs` (Windows).

There is no need to create subdirectories for processes that use unified logging. This differs from logs created by processes using legacy logging.

For a list of the processes that use unified logging, and for other details on both unified and legacy logging, refer to the “Using Logs and Reports” chapter of the *NetBackup Troubleshooting Guide*.

To control the size and number of unified logs, use the `vxlogcfg` and `vxlogmgr` commands, as described in the *NetBackup Troubleshooting Guide*.

Legacy Logging

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

Create the legacy log directories in `/usr/opensv/netbackup/logs/process_name` (UNIX) and to `install_path\NetBackup\logs\process_name` (Windows).

Enable Robust Logging

A check in the **Enable Robust Logging** checkbox indicates that when a log file grows to the 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. See the *NetBackup Troubleshooting Guide* for more information about controlling the log file size.

If **Enable Robust Logging** is enabled:

The log files for bprd, bpbkar, bpbmr, bpcd, bpdbm, bptm, bpdm and are named using the following convention:

`MMDDYY_NNNNN.log`

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

If **Enable Robust Logging** is disabled:

Only one log file per day is produced:

`MMDDYY.log`

Whether Robust Logging is selected or not, the log file is pruned using `KEEP_LOGS_DAYS` and `DAYS_TO_KEEP_LOGS` settings.

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

Global Logging Level

The **Global Logging Level** property is used for debugging purposes, the logging levels control the amount of information that the NetBackup server writes to logs.

Six levels are supported. A value of 0 sets logging to minimum (default) and a value of 5 sets it to maximum.

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

Master Server, Media Server, and Client Host Properties

Some NetBackup processes allow individual control over the amount of information the process writes to logs. For those processes, it is possible to specify a different logging level other than the **Global Logging Level**.

Process Specific Overrides

The following services utilize legacy logging. All service require that you first create a log directory in `/usr/opensv/netbackup/logs/process_name`.

BPBRM Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bpbrm`: 0 (minimum) through 5 (maximum).

BPTM Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bptm`: 0 (minimum) through 5 (maximum).

BPDM Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bpdm`: 0 (minimum) through 5 (maximum).

BPRD Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bprd`: 0 (minimum) through 5 (maximum).

BPDBM Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bpdbm`: 0 (minimum) through 5 (maximum).

Vault Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bpvault`: 0 (minimum) through 5 (maximum).

Database Logging Level

If you wish to override the **Global Logging Level**, select a database logging level: 0 (minimum) through 5 (maximum). This property is configurable for clients only.

Diagnostic Logging Levels for NetBackup Services

The **Logging** properties page offers configurable diagnostic levels for services which utilize unified logging. Those services are listed below. Each service creates a log automatically in `/usr/openv/logs`.

To change the debug levels, use the `vxlogcfg` command. Please refer to the *NetBackup Troubleshooting Guide* for more information.

Policy Execution Manager

If you wish to override the Global Logging Level, select a logging level for the Policy Execution Manager (NBPEM): 0 (minimum) through 5 (maximum). Select *No Logging* to produce no log for this service at all.

The Policy Execution Manager compiles a worklist for jobs and determines when jobs are due to run. This property appears for EMM servers. This property does not appear for NetBackup hosts earlier than 6.0.

Job Manager

If you wish to override the Global Logging Level, select a logging level for the Job Manager (NBJM): 0 (minimum) through 5 (maximum). Select *No Logging* to produce no log for this service at all.

The Job Manager accepts jobs submitted by the Policy Execution Manager (NBPEM) and acquires the necessary resources. This property appears for EMM servers. This property does not appear for NetBackup hosts earlier than 6.0.

Resource Broker

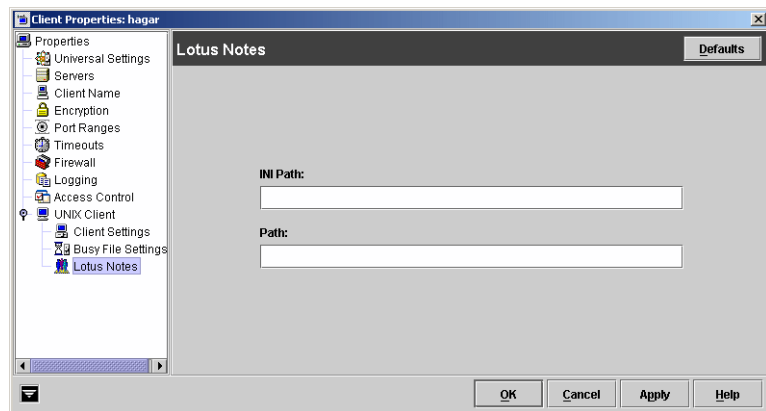
If you wish to override the Global Logging Level, select a logging level for the Resource Broker (NBRB): 0 (minimum) through 5 (maximum). Select *No Logging* to produce no log for this service at all.

The Resource Broker makes the allocations for storage units, tape drives, client reservations. This property does not appear for NetBackup hosts earlier than 6.0.

Lotus Notes Properties

The **Lotus Notes** properties apply to currently selected clients running NetBackup for Lotus Notes.

The following topics explain the settings. For more information, see the *NetBackup for Lotus Notes System Administrator's Guide*.



Path

In the **Path** field, specify the path where the Lotus Notes program files reside on the client. NetBackup must know where these files are in order to perform backup and restore operations. The value in this box overrides the one specified by the Lotus registry key, if both are defined.

INI File

In the **INI** field, specify the absolute path to the NOTES .INI file associated with the server instance to be used to back up and restore a Lotus database. Use this setting to specify the correct .INI file when backing up and restoring from Domino partitioned servers. It is not necessary to specify the .INI file for non-partitioned servers.

Media Properties

The **Media** properties apply to selected master servers and media servers. **Media** properties control how NetBackup manages media.

Allow Media Overwrite

The **Allow Media Overwrite** property overrides NetBackup's overwrite protection

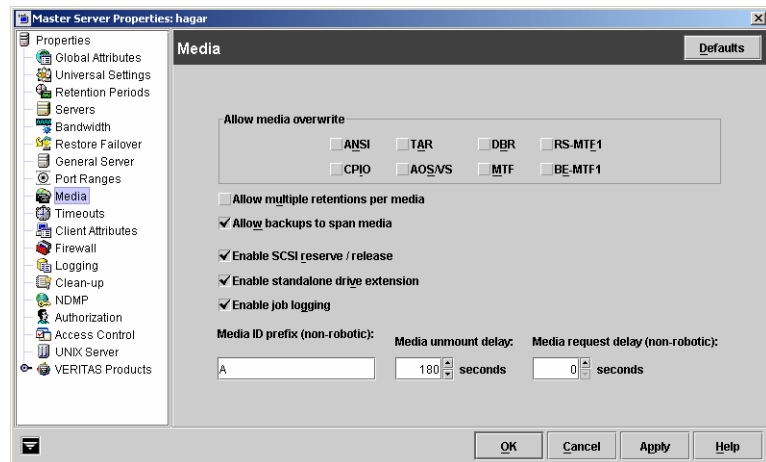
for specific media types. Normally, NetBackup will not overwrite certain media types. To disable overwrite protection, place a check in the check box of one or more of the listed media formats.

For example, place a check in the **CPIO** check box to permit NetBackup to overwrite the cpio format.

The following media formats on removable media can be selected to be overwritten:

- ◆ **ANSI:** When checked, ANSI labeled media can be overwritten.
- ◆ **AOS/VS:** When checked, AOS/VS media can be overwritten. (Data General AOS/VS backup format.)
- ◆ **CPIO:** When checked, CPIO media can be overwritten.
- ◆ **DBR:** When checked, DBR media can be overwritten. (This is a VERITAS backup format that is no longer used.)
- ◆ **RS-MTF1:** VERITAS Remote Storage MTF1 media format. When checked, VERITAS Remote Storage MTF1 media format can be overwritten.
- ◆ **TAR:** When checked, TAR media can be overwritten.
- ◆ **MTF1:** When checked, MTF1 media can be overwritten. With only MTF1 checked, all other MTF formats, apart from Backup Exec MTF (BE-MTF1) and Remote Storage MTF (RS-MTF1) media format can be overwritten.
- ◆ **BE-MTF1:** When checked, Backup Exec MTF media can be overwritten.

By default, NetBackup does not overwrite any of the above 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.



Master Server, Media Server, and Client Host Properties

If media contains one of the protected formats and you do not permit media overwriting, 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 has been previously selected for backups
 - ◆ Sets the volume's state to SUSPENDED
 - ◆ Aborts the requested backup
 - ◆ Logs an error
- ◆ If the volume is mounted for a backup of the NetBackup catalog, the backup is aborted and an error is logged that 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 that indicates the volume does not have a NetBackup format.

Allow Multiple Retentions Per Media

The **Allow Multiple Retentions per Media** setting allows NetBackup to mix retention levels on media. It applies to media in both robotic and nonrobotic drives. By default, the check box is clear and each volume can contain backups of only a single retention level.

Allow Backups to Span Media

The **Allow Backups to Span Media** property, when checked, allows backups to span more than one media. This property allows NetBackup to select another volume to begin the next fragment. The resulting backup has data fragments on more than one volume. By default, **Allow Backups to Span Media** is checked and backups are allowed to span media.

If the end of media is encountered and this property is *not* selected, the media is set to FULL and the operation terminates abnormally. This applies to both robotic and nonrobotic drives.

Enable SCSI Reserve/Release

The **Enable SCSI Reserve/Release** property ensures the use of SCSI reserve to all tape devices from this host. This feature blocks access to the device from other host systems. If unchecked, other hosts may send commands to the device that cause a loss of data.

Enable Standalone Drive Extension

Check the **Enable Standalone Drive Extension** property to allow NetBackup to use whatever labeled or unlabeled media is found in a nonrobotic drive. By default, standalone drive extensions are enabled.

Enable Job Logging

Check the **Enable Job Logging** property to allow the logging of job information used by the NetBackup Activity Monitor. By default, job logging occurs.

Media ID Prefix (Non-robotic)

The **Media ID Prefix (Non-robotic)** property specifies the media ID prefix to use in media IDs when unlabeled media is found in nonrobotic drives. The prefix must be one to three alpha-numeric characters. NetBackup appends remaining numeric characters. By default, NetBackup uses A and assigns media IDs such as A00000, A00001, and so on.

For example, if *FEB* is specified, NetBackup appends the remaining numeric characters so the assigned media IDs become FEB000, FEB001, and so on (note that this does not work with the Configure Volumes wizard).

Media Unmount Delay

Specifying a **Media Unmount Delay** property indicates that media unload is delayed for the number of seconds indicated, after the requested operation is complete. **Media Unmount Delay** applies only to user operations, including backups and restores of database agent clients, such as those running NetBackup for Oracle. The delay reduces unnecessary media unmounts and media positioning in cases where the media is requested again a short time later.

The delay can range from 0 to 1800 seconds. (Default: 180 seconds.) If you specify 0, the media unmount occurs immediately upon completion of the requested operation. Values greater than 1800 are set to 1800.

Media Request Delay

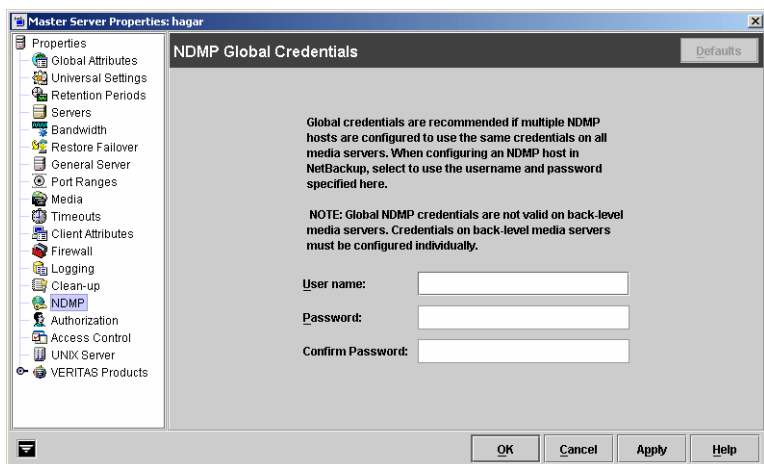
The **Media Request Delay** property specifies how long NetBackup waits for media in nonrobotic drives. This is useful if a gravity feed stacker is used on a nonrobotic drive and there is a time delay between the dismount of one media and the mounting of another. Default: 0 seconds.

During the delay period, NetBackup checks every 60 seconds to see if the drive is ready. If the drive is ready, NetBackup uses it. Otherwise, it 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 only once at the end of the delay.

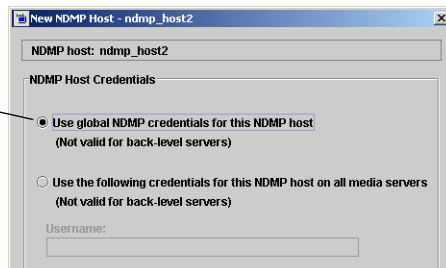
For example, assume you set the delay to 150 seconds. Here, NetBackup waits 60 seconds, checks for ready, waits 60 seconds, checks for ready, and then waits 30 seconds and checks for ready the last time. If the delay had been 50 seconds (this short a delay is not recommended), NetBackup would have checked only once, at the end of 50 seconds.

NDMP Global Credentials Properties

The credentials entered on the **NDMP Global Credentials** can apply to any NDMP host in the configuration if **Use global NDMP credentials for this NDMP host** is selected for the NDMP host.



Select on NDMP host for NDMP Global Credentials to apply



User Name

The user name under which NetBackup accesses the NDMP server. This user must have permission to run NDMP commands.

Password and Confirm Password

Enter and re-enter the password.

NetWare Client Properties

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

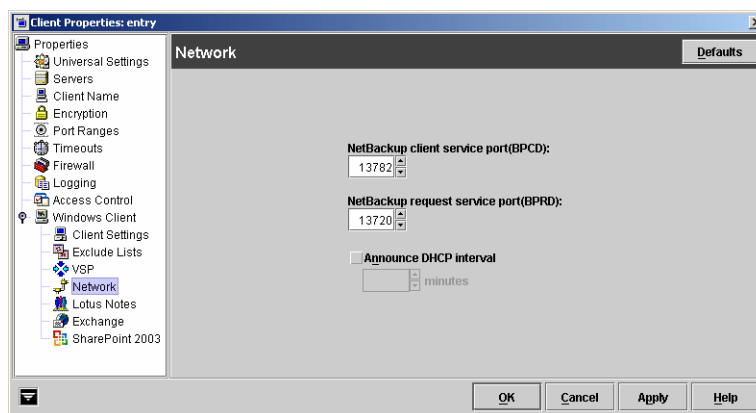
Netware Client properties include:

- ◆ “Client Settings (NetWare) Properties” on page 371
- ◆ “Open File Backup (NetWare Client) Properties” on page 430
- ◆ “OTM Properties” on page 430

Network Properties

The **Network** properties apply to currently selected Windows clients.

Under **Network** properties, set properties which define requirements for communications between clients and the master server.



NetBackup Client Service Port (BPCD)

The **NetBackup Client Service Port (BPCD)** property applies to Microsoft Windows clients and specifies the port that the NetBackup client uses to communicate with the NetBackup server. Default: 13782.

Note If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.

NetBackup Request Service Port (BPRD)

The **NetBackup Request Service Port (BPRD)** property applies to Microsoft Windows clients and specifies the port for the client to use when sending requests to the NetBackup request service (bprd process) on the NetBackup server. Default: 13720.

Note If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.

Announce DHCP Interval

The **Announce DHCP Interval** property applies to Microsoft Windows clients and specifies how many minutes the client waits before announcing that it is using a different IP address. The announcement occurs only if the specified time period has elapsed and the address has changed since the last time the client announced it.

Open File Backup (NetWare Client) Properties

The **Open File Backup** properties define Open File Backup properties on Netware clients.

Enable Open File Backup During Backups

Check the **Enable Open File Backup During Backups** check box to enable open transaction management.

OTM Properties

On Microsoft Windows and NetWare clients, previous versions of NetBackup have used Open Transaction Manager to back up files, databases, and applications that are open or active.

OTM properties do not appear for new or upgraded clients, which use Open File Backups instead. (See “Client Attributes Properties” on page 362.)

OTM host properties apply to clients at NetBackup version 3.4, 3.4.1, 4.5 GA, 4.5 MP1, MP2, MP3, MP4, and MP5. For information regarding OTM host properties, see the *NetBackup 4.5 System Administrator's Guide*.

Port Ranges Properties

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

Use Random Port Assignments

The **Use Random Port Assignments** property specifies that when NetBackup requires a port for communication with NetBackup on other computers, it will randomly choose one from those that are free in the allowed range. For example, if the range is from 1023 through 5000, it chooses randomly from the numbers in this range.

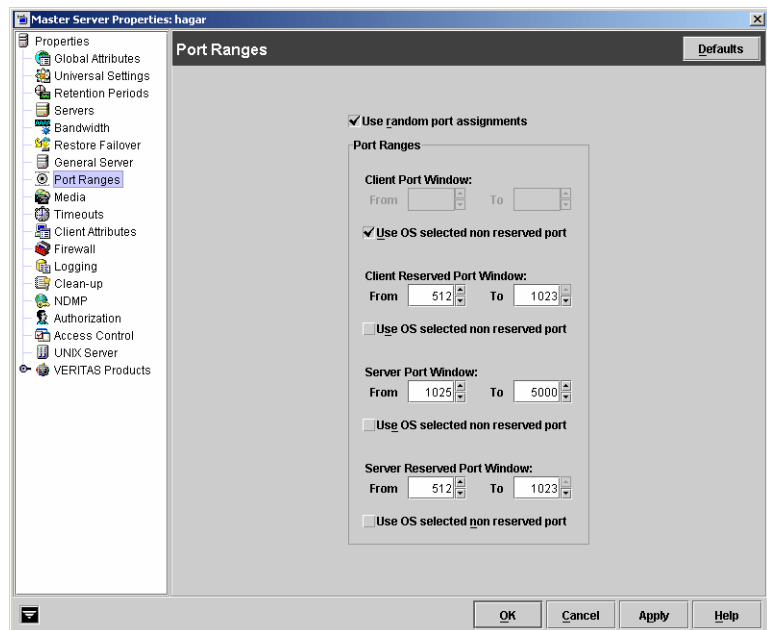
By default, **Use Random Port Assignments** is selected, and ports will be chosen randomly.

If deselected, NetBackup chooses numbers sequentially, starting with the highest number that is available in the allowed range. For example, if the range is from 1023 through 5000, NetBackup chooses 5000 (assuming port 5000 is free). If 5000 is being used, port 4999 is chosen.

Client Port Window

The **Client Port Window** property specifies the range of nonreserved ports on this computer that are used for connecting to NetBackup on other computers. This setting applies when connecting to NetBackup on a computer configured to accept nonreserved ports. (See **Allow Nonreserved Ports** on the **Universal Settings** dialog.)

If you prefer to have the operating system determine the nonreserved port to use, place a checkmark in the **Use OS selected non reserved port** check.



Client Reserved Port Window

The **Client Reserved Port Window** property specifies the range of reserved ports on this computer that are used for connecting to NetBackup on other computers. This setting applies when connecting to NetBackup on a computer configured to accept only reserved ports. (See **Allow Nonreserved Ports** on the **Universal Settings** dialog.)

Default range: 512 through 1023.

If you prefer to have the operating system determine the nonreserved port to use, place a checkmark in the **Use OS selected non reserved port** check.

Server Port Window

The **Server Port Window** property specifies the range of nonreserved ports on which this computer accepts connections from NetBackup on other computers. This setting applies when connecting to a client configured to accept only nonreserved ports. (See **Accept Connections on Non-reserved Ports** on the **Universal Settings** dialog.) **Server Port Window** does not appear when configuring a client.

Default range: 1024 through 5000.

If you prefer to have the operating system determine the nonreserved port to use, place a checkmark in the **Use OS selected non reserved port** check.

Server Reserved Port Window

The **Server Reserved Port Window** setting specifies the range of local reserved ports on which this computer accepts connections from NetBackup on other computers. This setting applies when connecting to a client configured to accept only reserved ports. (See **Allow Nonreserved Ports** on the **Universal Settings** dialog.) **Server Reserved Port Window** does not appear when configuring a client.


Default range: 512 through 1023.

If you prefer to have the operating system determine the nonreserved port to use, place a checkmark in the **Use OS selected non reserved port** check.

Restore Failover Properties

The **Restore Failover** properties apply to selected master servers.

The **Restore Failover** properties control how NetBackup performs automatic failover to another NetBackup media server in a master and media server cluster, if the regular media server is temporarily inaccessible for a restore.



The automatic failover does not require administrator intervention. By default, NetBackup does not perform automatic failover.

Examples of when to use the restore failover capability:

- ◆ Two or more media servers are sharing 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 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 media server (through `bpcd`) fails. Possible reasons for the failure are:

- ◆ The media server is down.
- ◆ The media server is up but `bpcd` is not responding (for example, if the connection is refused or access is denied).
- ◆ The media server is up and `bpcd` is all right but `bptm` is having problems (for example, if `vmd` is down or `bptm` cannot find the required tape).

Alternate Restore Failover Machines List

The **Media Server** column displays the NetBackup media servers that have failover protection for restores. The **Failover Restore Server** column displays the servers that are providing the failover protection. When automatic failover is required, NetBackup searches from top to bottom in the **Failover Restore Server** column for the failed server until it finds another server that can perform the restore.

Master Server, Media Server, and Client Host Properties


A NetBackup media server can appear only once in the **Media Server** column but can be a failover server for more than one other media server. The protected server and the failover server must both be in the same master and media server cluster.

Add Button

To include a NetBackup media server in the **Alternate Restore Failover Machines** list, click **Add**.

▼ To add or change a media server to the Alternate Restore Failover Machine list

1. To add an entry, click **Add**. The **Add Failover Servers** dialog appears.
To change an entry, click **Change**. The **Change Failover Servers** dialog appears.
2. In the **Server** field, specify the media server for which you're setting up failover protection.
3. In the **Failover Servers** field, specify the media server(s) that can be used if the server designated in the **Server** field is unavailable. Separate the names of multiple servers with a single space.
4. Click **Add** to add the name to the list. The dialog remains open for another entry. Click **Close** to close the dialog. If changing an entry, click **OK**.
5. Click **Apply** to accept the Restore Failover property changes. Click **OK** to close the host properties dialog.
6. Stop and restart the NetBackup Request daemon on the master server where you are changing the configuration.



For more information on failover, see "Method 3: Automatic Failover to Alternate Server" on page 527.

Change Button

To change an entry in the **Alternate Restore Failover Machines** list, select a media server, then click **Change**.

Remove Button

To remove a NetBackup media server from the **Alternate Restore Failover Machines** list, select the media server to be removed, then click **Remove**.

Master Server, Media Server, and Client Host Properties

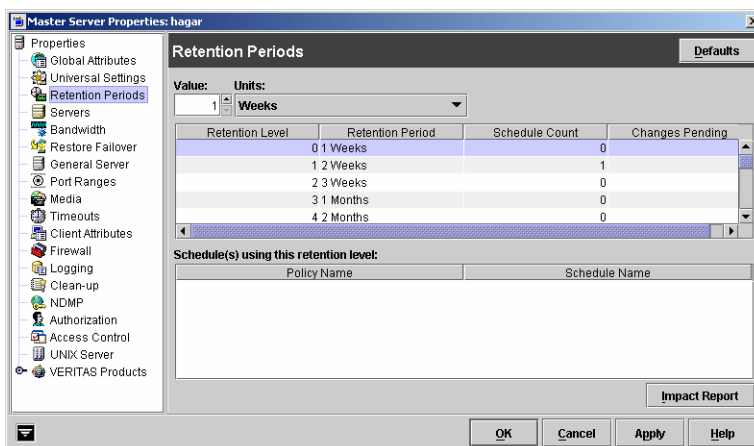


Retention Periods Properties

The **Retention Periods** properties apply to selected master servers.

When setting up a schedule, the selected retention period determines how long NetBackup retains the backups or archives created according to that schedule. There are 25 possible levels of retention from which to select. The

Retention Period properties define the length of time associated with each level.



Value

The **Value** specifies the retention level setting.

Units

The **Units** property specifies the units of time for the retention period. The list also includes the special units, **Infinite** and **Expires Immediately**.

Retention Periods List

The dialog contains a listing of the current definitions for the 25 possible levels of retention (0 through 24). 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.

Schedules List

The dialog contains a listing of the schedules that use the currently selected retention level, and the policy to which each schedule belongs.

Impact Report Button

Click **Impact Report** to display a summary of how changes will affect existing schedules. If you change a retention period, click **Impact Report**. The list displays all schedules in which the retention period is less than the frequency period (including schedules that do not use the retention periods that you have just changed.)

▼ To change a retention period

1. In the NetBackup Administration Console, expand **NetBackup Management > Host Properties > Master Servers > Double-click on master server > Retention Periods**.
2. Select the retention level that you want to change.

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

The dialog displays the names of all schedules that are using the selected retention level as well as the policy to which each schedule belongs.

3. Type the new retention period in the **Value** box.
4. Select the units of measure (*days, weeks, months, years, infinite or expires immediately*).

Note After changing either **Units** or **Value**, an asterisk (*) displays in the Changes Pending column to indicate that the period was changed. NetBackup does not change the actual configuration until **Apply** or **OK** is clicked.

5. Click **Impact Report**.

The policy impact list displays the schedules where the retention period is less than the frequency period (including schedules that do not use the retention periods that you just changed).

If any schedules are listed, correct the problem by either redefining the retention period or changing the settings for retention or frequency on the schedule.

6. To discard your changes, click **Cancel**.
7. To save your changes and update the configuration, click one of the following:
 - ◆ **Apply**: Saves changes and leaves the dialog open so you can make further changes.
 - ◆ **OK**: Saves changes since the last time you clicked **Apply**. **OK** also closes the dialog.

8. To save the changes, click **OK**.

Note on Redefining Retention Periods

NetBackup, by default, stores each backup on a volume that already has backups at the same retention level. However, NetBackup does not check the retention period defined for that level. This means that redefining the retention period for a level can result in unintentionally storing backups with different retention periods on the same volume. For example, if you change the retention period for level 3 from one month to six months, NetBackup stores future level 3 backups on the same volumes that it previously used (if they are available). That is, they are on the volumes with the level 3 backups that have a retention period of one month.

This is not a problem if the new and old retention periods are of about the same value. However, if you make a major change to a retention period (for example, from one week to infinity), it is best to suspend the volumes that were previously used for that retention level. To do this, proceed as follows:

1. Use the NetBackup Media List report to determine which volumes are currently at the level that you are going to suspend.
2. Use the `bpmedia` command to suspend the volumes.

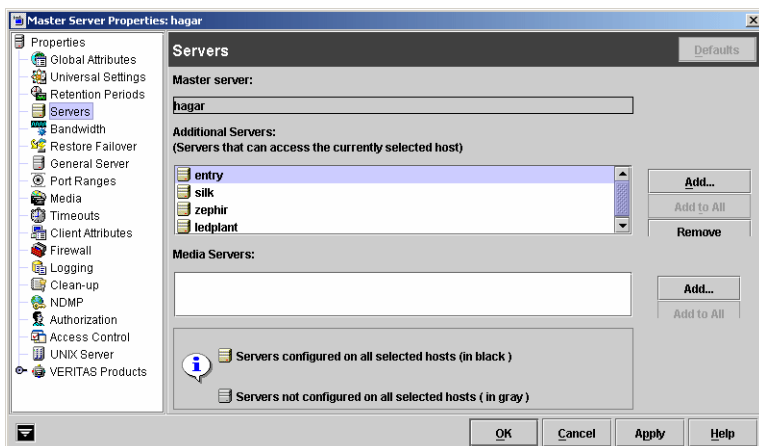
```
bpmedia -suspend -m media_ID
```


Servers Properties

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

Master Server

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



Additional Servers

Lists additional servers that can access the server specified as **Master Server**.

During installation, NetBackup sets the master server to the name of the system where the server software is being installed. NetBackup uses the master server value to validate server access to the client and to determine which server the client must connect to in order to list and restore files.

- ◆ To add a server, click **Add** and select a server.
- ◆ To delete a server, select a server from the list and click **Remove**.
- ◆ To change the master server, select another server from the list and click **Make Master**.

To configure access to a remote server, add to the server list the name of the host seeking access. For more information, see “Administering a Remote Master Server” on page 474.

Media Servers

The **Media Servers** list specifies that the listed machines are media servers only. Machines listed as media servers can back up and restore clients, but have limited administrative privileges.

- ◆ To add a new media server, click **Add** and select a server.

When adding a media server to an existing master server, run `nbsmmcmd -addhost` to add the host to the Enterprise Media Manager (EMM) database.

Master Server, Media Server, and Client Host Properties

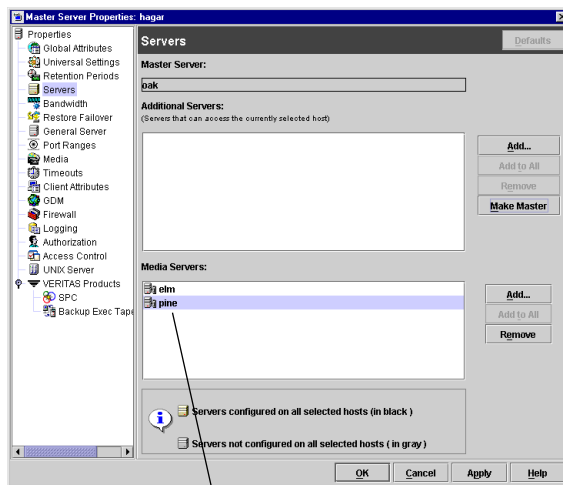
- ◆ To delete a media server, select a media server from the list and click **Remove**.

Restricting Administrative Privileges of Media Servers

The servers included in the **Media Servers** list are media servers only. (**Host Properties > Master Server or Media Servers > Servers.**)

Machines listed as media servers can back up and restore clients, but have limited administrative privileges.

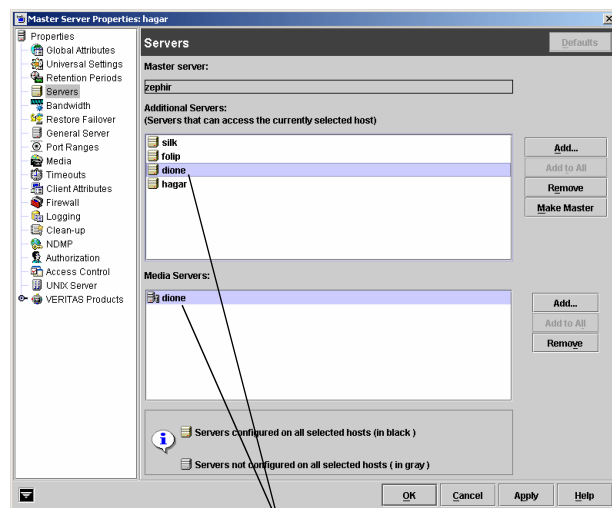
Suppose you have a configuration consisting of master server *oak* and two media servers—*elm* and *pine*. Set up *oak* as the master server and *elm* and *pine* as media servers.



Administrative scope of media servers is limited

If a machine is defined as both a master server and a media server, the master server entry takes precedence.

A consequence of listing a server as both a master and media server is that a system administrator on a media server would also be a NetBackup administrator on other master servers.



A machine listed as both an additional server and a media server has full administrative privileges

Multiple Masters Sharing One Enterprise Media Manager Host

Multiple master servers can share one EMM database located on a single host. The host containing the EMM database can be either a master server or a media server.

The **Servers** host properties must be set up to allow multiple master servers to access the EMM host. This can be set using the **Host Properties** or configured in the `bp.conf` file.

Shared EMM Database Located on a Master Server

In the following example, three master servers are sharing one EMM database located on one of the servers (*meadow*).

The `bp.conf` server entries on each master server would read as described in the following table:

Meadow	Havarti	Study
SERVER = meadow	SERVER = havarti	SERVER = study
SERVER = havarti	SERVER = meadow	SERVER = meadow
SERVER = study	CLIENT_NAME = havarti	CLIENT_NAME = study
CLIENT_NAME = meadow	EMMSERVER = meadow	EMMSERVER = meadow
EMMSERVER = meadow		

SERVER entries:

- ◆ The first `SERVER` entry must be the name of the master server.

Other master servers that need to be listed:

- ◆ In order for the NetBackup Administration Console to administer other servers, the servers must be listed. (**File > Change Server.**)
- ◆ If the EMM database is on another master server, that server needs to be listed. In the table above, *meadow* is listed on *havarti* and *study*.

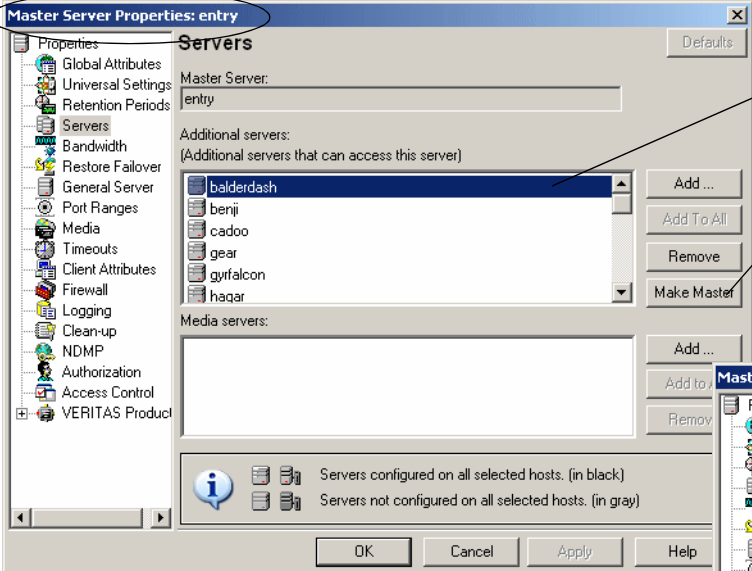
The `EMMSERVER` entry must be present on all master servers sharing the EMM host. In the table above, *meadow* is listed as the `EMMSERVER` on *havarti*, *study*, as well as on *meadow*.

Shared EMM Database Located on a Media Server

Master Server, Media Server, and Client Host Properties

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

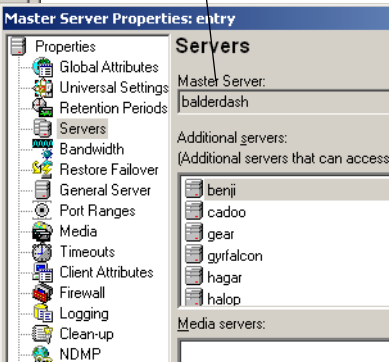
1 **Master Server Properties: entry**



To modify the master server:

Select a server from the list.
Click **Make Master**.

The server then appears as the new master.



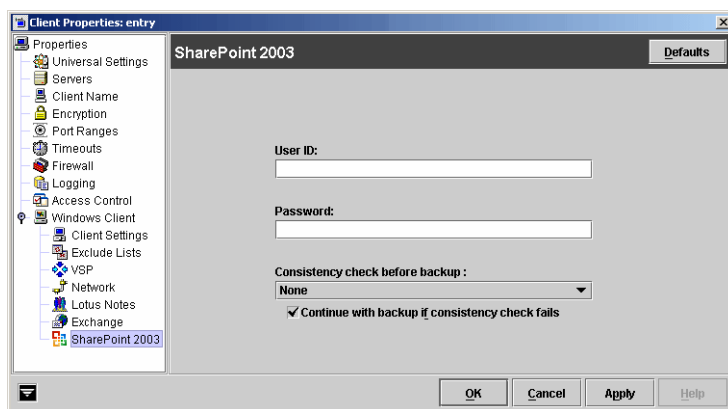
- 2** To update the EMM database, after changing the master server for a media server, run:
- ```
/usr/opensv/netbackup/bin/admincmd/nbemcmd -updatehost
```

## SharePoint 2003 Properties

The **SharePoint 2003** properties apply to currently selected Windows 2003 clients in order to protect SharePoint 2003 installations.

### User ID

Specify the user id for the account used to log on to SharePoint  
(DOMAIN\user name).



### Password

Specify the password for the account.

### Consistency Check Before Backup

Select what kind, if any, of consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation. These checks are performed for both server- and user-directed backups.

**None:** No consistency check will be performed.

**Full check, excluding indexes:** Select this 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 nonclustered index pages is not checked.

**Fullcheck, including indexes:** Include indexes in the consistency check. Any errors are logged.

**Physical check only (SQL 2000 only):** Select this to perform a low overhead check of the physical consistency of the SQL Server 2000 database. This option only checks the integrity of the physical structure of the page and record headers, and the consistency between the pages' object ID and index ID and the allocation structures.

### Continue with Backup if Consistency Check Fails

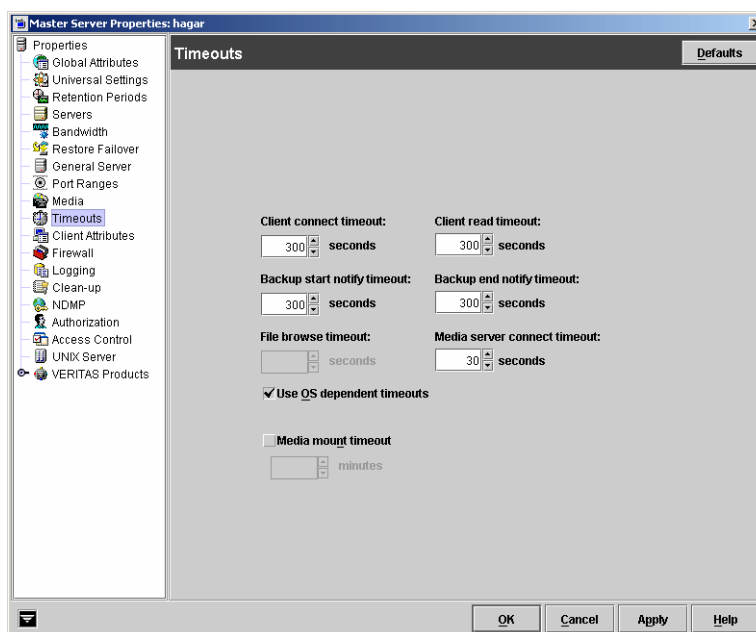
Select whether or not the backup should continue if the consistency check fails.

## Timeouts Properties

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

### Client Connect Timeout

The **Client Connect Timeout** property specifies the number of seconds the server waits before timing out when connecting to a client. Default: 300 seconds.



### Backup Start Notify Timeout

The **Backup Start Notify Timeout** property specifies the number of seconds the server waits for the `bpstart_notify` script on a client to complete. Default: 300 seconds.

**Note** If you change this timeout, verify that **Client Read Timeout** is set to the same or higher value.

### File Browse Timeout

The **File Browse Timeout** property specifies the number of seconds for the client to wait for a response from the NetBackup master server when listing files.

**Note** On a UNIX client, the value in the user's `$HOME/bp.conf` file takes precedence, if it exists, to the property here.

If **File Browse Timeout** is exceeded, the user receives a *socket read failed* error even if the server is still processing the request.

## Use OS Dependent Timeouts

The **Use OS Dependent Timeouts** property specifies that the client wait for the timeout period as determined by the operating system when listing files:

- ◆ Windows client: 300 seconds
- ◆ UNIX client: 1800 seconds

## Media Mount Timeout

The **Media Mount Timeout** property appears as a master server property only and specifies the number of minutes that NetBackup waits for the requested media to be mounted, positioned, and ready on backups, restores, and duplications.

Use this timeout to eliminate excessive waits when it is necessary to manually mount media (for example, when robotic media is out of the robot or off site). When restoring backups or archives that were written to a disk being managed by Storage Migrator on a UNIX server, the media mount timeout value is in effect during the caching of potentially migrated files. If a file is part of a large disk image that Storage Migrator has migrated to tape, there must be enough time to cache in the entire disk file.

## Client Read Timeout

The **Client Read Timeout** property specifies the number of seconds to use for the client-read timeout on a NetBackup master, remote media server, or database-extension client (such as NetBackup for Oracle). Default: 300 seconds.

The client-read timeout on a database-extension client is a special case. Clients can initially require more time to get ready than other clients because database backup utilities frequently start several backup jobs at the same time, slowing the central processing unit.

The sequence on a database-extension client is as follows:

- ◆ NetBackup on the database-extension client reads the client's client-read timeout to find the initial value. If the option is not set, the standard five minute default is used.
- ◆ When the database-extension API receives the server's value, it uses it as the client-read timeout.

---

**Note** For database-extension clients, VERITAS suggests that you set the client-read timeout to a value greater than 5 minutes. 15 minutes is adequate for many installations. For other clients, change **Client Read Timeout** only if problems are encountered.

---

## Backup End Notify Timeout

The **Backup End Notify Timeout** property specifies the number of seconds that the server waits for the `bpend_notify` script on a client to complete. Default: 300 seconds.

---

**Note** If you change this property, verify that **Client Read Timeout** is set to the same or higher value.

---

## Media Server Connect Timeout

The **Media Server Connect Timeout** property specifies the number of seconds that the master server waits before timing out when connecting to a remote media server. Default: 30 seconds.

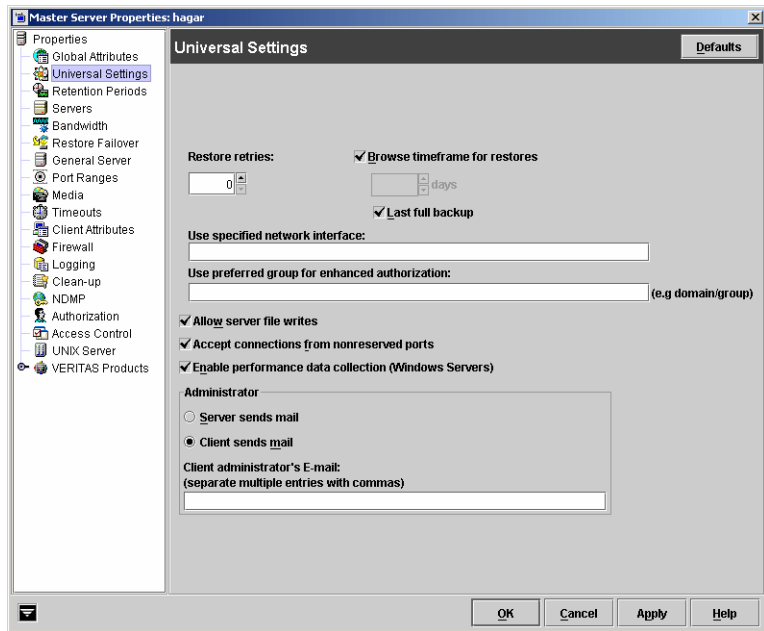


## Universal Settings Properties

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

### Restore Retries

The **Restore Retries** setting specifies the number of attempts (1 through 3) a client will try to restore after a failure. The default is 0 (client will not attempt to retry). If a job is of a type that can be checkpointed, the job will retry from the start of the last checkpointed file rather than at the beginning of the job.



Change **Restore Retries** only if problems are encountered.

If a job fails after the number of retries, the job remains in the incomplete state as determined by the **Move Restore Job From Incomplete State to Done State** property on the Global Attributes host properties page. Checkpoint Restart for restores allows a failed restore job to be resumed by a NetBackup administrator from the Activity Monitor.

### Browse Timeframe for Restores

The **Browse Timeframe for Restores** property specifies the number of days in the past that NetBackup searches for files to restore. For example, to limit the browse range to the seven days prior to the current date, clear the **Last Full Backup** checkbox, then specify 7.

This limit is specified on the master server and applies to all NetBackup clients. It can also be specified on a client and in this instance applies only to that client and can reduce the size of the search window from what you specify on the server (the client setting cannot make the window larger).

By default, NetBackup includes files from the time of the last-full backup through the latest backup for the client. If the client belongs to more than one policy, then the browse starts with the earliest of the set of last-full backups.

## Last Full Backup

The **Last Full Backup** property indicates that NetBackup should automatically include in its browse range, all backups since the last successful full backup. The **Last Full Backup** check box must be cleared in order to enter a value for the **Browse Timeframe for Restores** property.

## Use Specified Network Interface

The **Use Specified Network Interface** property specifies the network interface that NetBackup uses when connecting 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.

### Example 1 - Client with multiple network interfaces.

Assume a NetBackup client with two network interfaces:

- ◆ One network interface is for the regular network. The host name for the regular interface is fred.
- ◆ One network interface is for the backup network. The host name for the backup interface is fred\_nb.

The NetBackup client name setting on both the client and server is fred\_nb.

When client fred starts a backup, restore, or list operation, the request goes out on the fred\_nb interface and over the backup network. This assumes that fred and the network are set up to do so. If this configuration is not in place, fred can send the request out on the fred interface and over the regular network. The server receives the request from client fred\_nb with host name fred and refuses it because the host and client names do not match.

One way to solve this problem is to set up the master server to allow redirected restores for client fred. This allows the server to accept the request, but leaves NetBackup traffic on the regular network.

A better solution is to set **Use Specified Network Interface** on fred to fred\_nb. Now, all backup, restore, and list requests use the fred\_nb interface, the server receives requests from client fred\_nb with host name fred\_nb, and everything works as intended.

### Example 2 - Server with multiple network interfaces.

Assume a NetBackup server with two network interfaces:

- ◆ One network interface is for the regular network. The host name for the regular interface is barney.
- ◆ One network interface is for the backup network. The host name for the backup interface is barney\_nb.

The server list on all NetBackup servers and clients have an entry for barney\_nb.

When barney connects to a client for a backup, the request ideally goes out on the barney\_nb interface and over the backup network. This assumes that barney and the network are set up to do so. If this configuration is not in place, barney can send the request out on the barney interface and over the regular network. The client now receives the request from barney rather than barney\_nb and refuses it as coming from an invalid server.

One way to solve this problem is to add an entry for barney to the server list on the client. The client now accepts requests from barney, but NetBackup traffic continues on the regular network.

A better solution is to set **Use Specified Network Interface** on barney to barney\_nb. Now, when barney connects to a client, the connection is always through the barney\_nb interface and everything works as intended.

## Use Preferred Group for Enhanced Authorization

The **Use Preferred Group for Enhanced Authorization** setting specifies the domain group name that is passed by this computer to the server when NetBackup-user authorization is used. The default is the user's primary *domain\group*. The **Use Preferred Group for Enhanced Authorization** entry is intended specifically for use with NetBackup enhanced authorization. The entry is case sensitive and must be in the form *domain\group*. For example:

```
NTDOMAINNAME\Backup Operators
```

When **Use Preferred Group for Enhanced Authorization** is specified, Windows global groups are checked to determine if the user is a member of the specified *domain\group*:

- ◆ If the specified *domain\group* is a global group and the user is a member, then this *domain\group* value is used.
- ◆ If the specified *domain\group* is a local group or the user is not a member, then the user's primary *domain\group* is used. Note that if the domain name is an empty string or is the name of the local machine, it is considered to be local.

Some NetBackup processes also use the **Use Preferred Group for Enhanced Authorization** entry for Media Manager authorization. For more information on this, see "Media Manager Configuration File (vm.conf)" in the *NetBackup Media Manager System Administrator's Guide*.

## Master Server, Media Server, and Client Host Properties

Adding a **Use Preferred Group for Enhanced Authorization** entry in the Universal Settings dialog has the following effect on UNIX and Windows systems:

The `PREFERRED_GROUP` entry is added to the `bp.conf` file:

```
PREFERRED_GROUP = netgroup name
```

- ◆ If the `bp.conf` configuration file has a `PREFERRED_GROUP` entry, the `innnetgr()` function is used to determine if the user is in the netgroup (for further details refer to the `innnetgr` man page).
- ◆ If the `PREFERRED_GROUP` entry does not exist or the user is not a member of the netgroup, the local group name is obtained.

---

**Note** Netgroups are not supported for Sequent systems.

---

### Allow Server File Writes

The **Allow Server File Writes** setting prevents the NetBackup server from creating or modifying files on the NetBackup client. For example, checking this box would prevent server-directed restores and remote changes to the client properties.

Once **Allow Server File Writes** is applied, it can be cleared only by modifying the client configuration. Default: server writes are allowed.

### Accept Connections on Non-reserved Ports

The **Accept Connections on Non-reserved Ports** property specifies that the NetBackup client service (`bpcd`) can accept remote connections from nonprivileged ports (port numbers 1024 or greater). If this property is not specified, `bpcd` requires remote connections to come from privileged ports (port numbers less than 1024). **Accept Connections on Non-reserved Ports** is useful when NetBackup clients and servers are on opposite sides of a firewall.

When unchecked, this also means that the source ports for connections to `bpcd` use reserved ports as well.

If **Accept Connections on Non-reserved Ports** is checked (default) on a client or server, and you want to use non-reserved ports, the server connecting to the client or server must also be set up to use non-reserved ports for the client.

In addition to changing **Accept Connections on Non-reserved Ports** here, specify that the server use nonreserved ports for this client: select **Accept Connections from Non-reserved Ports** on the server properties Client attributes tab.

## Enable Performance Data Collection

The **Enable Performance Data Collection** property specifies to NetBackup to update disk and tape performance object counters. (Applies only to Windows master and media servers. The NetBackup performance counters can be viewed using the Windows Performance Monitor utility (perfmon).

## Client Sends Mail

The **Client Sends Mail** property specifies that the client send the E-mail to the address specified in the box labeled for the administrator's E-mail address. If the client cannot send E-mail, select **Server Sends Mail**.

## Server Sends Mail

The **Server Sends Mail** setting specifies that the server send the mail to the address specified in the box for the administrator's E-mail address. This is useful if the client cannot send mail.

## Client Administrator's E-mail

The **Client Administrator's E-mail** property specifies the E-mail address of the administrator on the client and is the address where NetBackup sends status on the outcome of automatic or manual backup operations for the client. By default, no E-mail is sent. To enter multiple addresses or E-mail aliases, separate entries with commas.

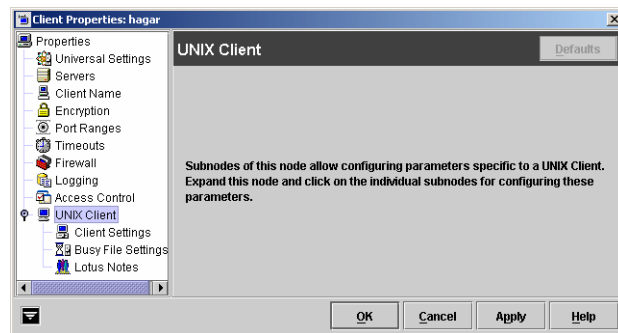
See "Setting Up E-Mail Notifications" on page 417 to ensure that the system is set up to send mail.

## UNIX Client Properties

The **UNIX Client** properties define NetBackup properties of UNIX clients.

**UNIX Client** properties include:

- ◆ “Client Settings (UNIX) Properties” on page 372
- ◆ “Busy File Properties” on page 357
- ◆ “Lotus Notes Properties” on page 422

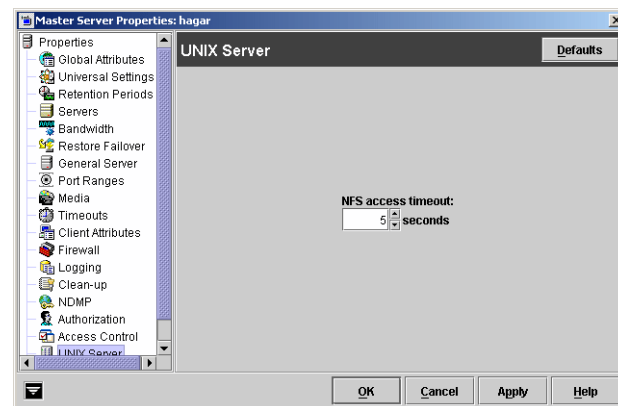


## UNIX Server Properties

The **UNIX Server** properties apply to selected UNIX master servers.

### NFS Access Timeout

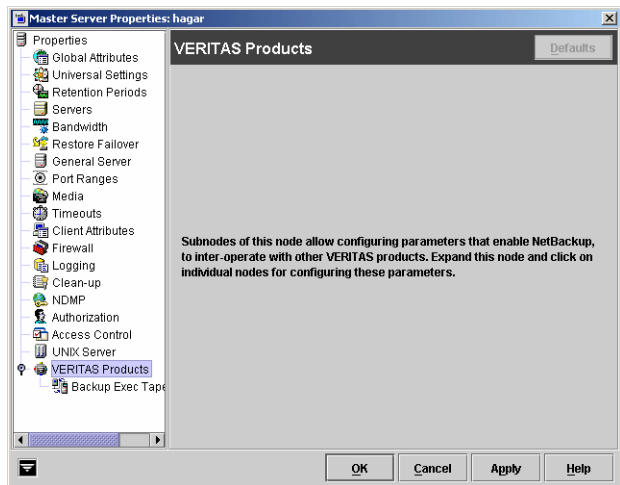
The **NFS Access Timeout** property specifies the number of seconds that the backup process waits when processing the mount table before considering an NFS file system unavailable. Default: 5 seconds.



## VERITAS Products Properties

The **VERITAS Products** properties apply to currently selected master servers.

**VERITAS Products** properties include the subnode, “Backup Exec Tape Reader Properties” on page 352.

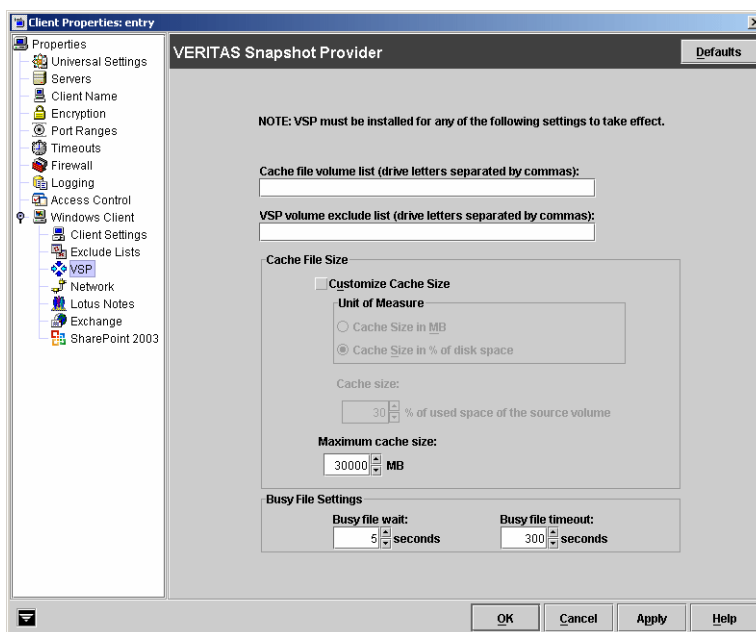


## VSP (Volume Snapshot Provider) Properties

The VSP properties apply to currently selected supported Windows clients.

In order for the properties in this dialog to affect client(s), VSP must be selected as the snapshot provider for the client(s). VSP is the default Windows 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.



To make VSP the snapshot provider for a Windows client, select **NetBackup Management** > **Master Servers** > *Select master server for the client* > **Client Attributes** > . (See “Client Attributes Properties” on page 362.)

### VSP Overview

NetBackup uses VSP to back up open and active files on Windows 2000, Windows XP, and Windows Server 2003 (32- and 64-bit) clients. To make backing up open and active files possible, VSP first captures a snapshot of each volume that needs to be backed up. After creating a snapshot of the volume, a virtual drive representing a static copy of the volume in a point-in-time is created along with a corresponding VSP cache file. NetBackup backs up files using the virtual drive instead of the actual drive. For each snapshot that is created for a volume, a VSP cache file is created to maintain the integrity of the snapshot. The original data corresponding to the changes that occur during the backup is stored in the cache file that was created along with the volume snapshot.

VSP is similar to OTM (used in previous releases) in that VSP creates volume snapshots using a caching mechanism. However, it is important to keep in mind the following considerations when using VSP:

- ◆ VSP uses a cache file for each volume that requires a snapshot, while OTM uses only one cache for all snapshots.



- ◆ Using VSP, a snapshot of a volume cannot be created if the volume already contains a VSP cache file.
- ◆ Using VSP, a cache file cannot be placed on a volume that has had a snapshot taken of it, or is in the process of having a snapshot taken. Only when the snapshot for the volume has been destroyed can it be used as a location for a VSP cache file.
- ◆ All VSP cache files are placed at the root level of a volume and are removed when its VSP snapshot has been destroyed.
- ◆ VSP cannot be used to perform hot database backups. See “Using VSP with Databases” on page 461.

### Stepping through the Backup Process with VSP

The following steps describe the sequence of events during a backup using VSP:

1. Before the backup begins, NetBackup uses VSP to create snapshots for the backup job. NetBackup waits for a quiet period to occur when no writes are being performed on the drives that contain data to be backed up. This wait is required to ensure that the file system is in a consistent state. The length of the quiet period is defined by the property. If a quiet period of sufficient length does not occur within the time specified by **Busy File Timeout**, the backup proceeds without VSP.
2. If a quiet period of sufficient length is detected, NetBackup performs the actions necessary to record the VSP snapshot.
3. The backup begins and NetBackup starts reading data from the client. If an application requests a read or write during the backup, VSP reads or writes the disk or its cache as necessary to maintain the snapshot and provide accurate data to the application.
4. Once the backup completes, NetBackup attempts to destroy the VSP volume snapshots created for the backup job while deleting the VSP cache files for the volume snapshots.

### Logging VSP Messages

VSP snapshot activity is logged in the `bpfis` and `online_util` debug logs. To enable VSP logging messages, create directories `bpfis` and `online_util` in the following location (default):

```
C:\Program Files\VERITAS\NetBackup\Logs\
```

If you wish, specify a different location during client installation.

To create detailed log information, set the **Global Logging Level** to a higher value on the master server host property Logging dialog. (**Host Properties > Master Servers > Selected master server > Logging.**) Eventually, these directories can require extensive of disk space. Delete the directories when you are finished troubleshooting and reset the **Global Logging Level** to a lower value.

## Cache File Volume List

The **Cache File Volume List** serves as a list of preferred locations for NetBackup to place VSP cache files. Volumes should be listed in this list as drive letters separated by commas and spaces. For example: C, D, E

The list indicates that all backup jobs requiring VSP snapshots will have their VSP cache files placed in a volume listed in the **Cache File Volume List**. If multiple volumes are listed in the **Cache File Volume List**, the volume with the most free disk space at the time of the backup is the preferred location for VSP cache files. If no volumes are listed, NetBackup will automatically determine the best location for VSP cache files.

NetBackup places cache files on one of these volumes unless it is determined that it is undesirable to use the volumes as cache file locations, even if it is specified by the user. The location is considered undesirable if a volume in the list is also being targeted to be snapshot. Because VSP does not allow snapshots of volumes already containing active cache files, NetBackup would not allow other VSP cache files to be placed in the volume.

Assume a backup job is backing up the C and D volumes and needs to create VSP volume snapshots for the volumes. The **Cache File Volume List** lists the C volume as a preferred location for all backups to place the VSP cache files:

1. NetBackup uses VSP to create VSP snapshots for the C and D volumes.
2. Instead of placing the VSP cache files in C because C is listed in the **Cache File Volume List**, NetBackup proceeds to place the VSP cache files for the C and D snapshots in the C and D volumes because the C volume cannot be used as a preferred location for VSP snapshots since it is having a VSP snapshot created for it.
3. All subsequent backups will not be able to use C as a VSP cache file location until the VSP volume snapshots created in step 2 have been destroyed.

## VSP Volume Exclude List

The **VSP Volume Exclude List** contains volumes that are never to be snapped by VSP during backups or never to have VSP cache files placed on the volumes. Volumes in the **VSP Volume Exclude List** are excluded from VSP activity and are backed up without snapshot protection. Volumes should be listed in this list as drive letters separated by commas and spaces. For example: C, D, E

### Ramifications of the Precedence of the Volume List over the Exclude List

The volumes in the **Cache File Volume List** have precedence over the volumes listed in the **VSP Volume Exclude List**. The **Cache File Volume List** overrides the **VSP Volume Exclude List** if both lists contain the same volume.

For example, if a user specifies `C : \` in the **Cache File Volume List** as well as the **VSP Volume Exclude List**, this means that the user wants the `C : \` volume to be the preferred location for VSP cache files, yet would not like VSP to snap or place cache files on the volume.

Because the **Cache File Volume List** takes precedence over the **VSP Volume Exclude List**, NetBackup places cache files in `C : \` even though it is listed in the **VSP Volume Exclude List**. NetBackup will not create snapshots for `C : \` until `C : \` is removed from the **VSP Volume Exclude List**.

### Using the Cache File Volume List and VSP Volume Exclude List for Multiple Simultaneous Backup Jobs or Multiple Groups of Multistreamed Jobs

NetBackup allows a scheduled backup to be broken into several backup streams that can run simultaneously to increase performance. (See “Allow Multiple Data Streams” on page 92.)

If a backup policy is configured to allow multiple data streams, a scheduled backup of a client can be divided into multiple data streams, with each file list directive in the policy forming a separate backup job (stream) that can run concurrently with other streams to help complete the scheduled backup. All backup jobs (streams) in a policy are grouped into an entity called a *stream group*. All backups that are part of a stream group have their VSP volume snapshots shared between backup jobs in the stream group.

Additionally, multiple backup jobs could also run concurrently on a single client even if a backup policy is configured not to allow multiple data streams.

For both these types of backups, it is necessary to use the **Cache File Volume List** and **VSP Volume Exclude List** to make sure that VSP snapshot creation is successful. When running these kinds of backups, it is highly recommended that a volume be listed in both the **Cache File Volume List** and the **VSP Volume Exclude List**. This volume would effectively be used as the volume for all VSP cache files. However, it will not have VSP snapshots created for it and all backups backing up the volume will not have VSP snapshots enabled for it.

### Example 1: Running Multiple Simultaneous Backups with VSP

Assume two backup jobs are run simultaneously; both jobs backing up the C and D volumes on a client that contains only volumes C and D:

1. Place either the C or D volume in the **Cache File Volume List** and the **VSP Volume Exclude List**. For this example, the D volume has been placed in the **Cache File Volume List** and the **VSP Volume Exclude List**.
2. Both backup jobs are run simultaneously, both backing up the C and D drives.
3. Snapshots of the C drive are created successfully for both backup jobs while their cache files were placed in the D drive. Since the D drive was listed in both the **Cache File Volume List** and the **VSP Volume Exclude List**, VSP cache files for the C drive for both backup jobs were placed in the D drive. Both backup jobs also backed up the D drive without VSP.

If the D drive was not listed in the **Cache File Volume List** and the **VSP Volume Exclude List**, VSP would not have been enabled for the C and D drives for both backup jobs.

### Example 2: Running Multiple Groups of Multistreamed Backups Simultaneously with VSP

Assume two multistreamed policies contain the following file lists:

Policy 1:

C:\ Dir1  
D:\ Dir2

Policy 2:

C:\ Dir3  
D:\ Dir4

When both policies are run simultaneously, two groups of multistreamed backup jobs will be run (with each group running a backup job for each file list item). Both groups of multistreamed jobs will be backing up the C and D volumes on a client that contains only volumes C and D:

1. Place either the C or D volume in the **Cache File Volume List** and the **VSP Volume Exclude List**. For this example, the D volume has been placed in the **Cache File Volume List** and the **VSP Volume Exclude List**.
2. Both policies are run simultaneously, which results in two groups of multistreamed jobs running at the same time and backing up the C and D drive contents.
3. Snapshots of the C drive are created successfully for both groups of multistreamed jobs while their cache files were placed in the D drive. Since the D drive was listed in both the **Cache File Volume List** and the **VSP Volume Exclude List**, VSP cache files for the C drive for both groups of multistreamed jobs were placed in the D drive. Both groups of multistreamed jobs also backed up the D drive without VSP.

If the D drive was not listed in the **Cache File Volume List** and the **VSP Volume Exclude List**, VSP would not have been enabled for the C and D drives for both groups of multistreamed jobs.

## Customize Cache Size

The **Customize Cache Size** property enables a number of properties that help you set specific cache size characteristics for your backup configuration. If the **Customize Cache Size** property is not enabled, NetBackup will automatically size VSP cache files for the client. By default, the **Customize Cache Size** property is disabled to allow NetBackup to automatically calculate cache file sizes for VSP snapshots. However, if the cache file sizes need to be configured manually, the **Customize Cache Size** property can be disabled and the VSP cache file sizes can be adjusted manually.

### Cache Size in % of Disk Space

The **Cache Size in % of Disk Space** property specifies that the **Cache Size** and **Maximum Cache Size** properties will use percentage of disk space as the form of measurement.

### Cache Size in Megabytes

The **Cache Size in Megabytes** property specifies that the **Cache Size** and **Maximum Cache Size** properties will use megabytes as the form of measurement.

## Cache Size

The **Cache Size** property is the initial size for the VSP cache file when creating snapshots.

## Maximum Cache Size

The **Maximum Cache Size** is the maximum for the VSP cache file to grow to when creating snapshots. The **Maximum Cache Size** is an optional configuration property and is only applicable when the VSP cache file is placed on a volume that is not being snapped. When the VSP cache file is placed on a volume that is not being snapped, the cache file size begins at 0 megabytes and can grow to a maximum size of **Maximum Cache Size**.

The **Maximum Cache Size** is calculated as a percentage of free disk space (of the cache file volume) from the value specified in **Maximum Cache Size** if the form of measurement used is percentage of disk space. The **Maximum Cache Size** is calculated in megabytes if the form of measurement selected is **Cache Size in MB**.

## Master Server, Media Server, and Client Host Properties

---

If the **Customize Cache Size** property is disabled, NetBackup automatically determines **Maximum Cache Size** for the VSP snapshot if the cache file is being placed in a volume that is not being snapped.

The following items are VSP best practices when configuring VSP cache sizes:

- ◆ Allow NetBackup to automatically determine cache file sizes for VSP snapshots by disabling the **Customize Cache Size** property. This allows NetBackup to allocate as much cache space as possible whenever creating VSP snapshots. In most cases, allowing NetBackup to automatically size cache files should avoid VSP snapshot errors from occurring. However, in some cases, VSP snapshot errors could occur (even if **Customize Cache Size** is disabled), depending on the data being backed up and the I/O activity of the client being backed up.
- ◆ If snapshot errors still occur even if **Customize Cache Size** is disabled, then increase the **Cache Size** and **Maximum Cache Size** properties to values that best fit your client's installation. The recommended setting for the **Cache Size** is 30% of used disk space of the volume that is being snapshot (the cache file size will be the value set at the **Cache Size** if the VSP cache file is placed in the same volume as being snapshot. It will be ignored otherwise). The recommended setting for **Maximum Cache Size** is 95% of free disk space of the cache file volume (the cache file will begin at 0 MB and will grow until a maximum of **Maximum Cache Size** if the cache file is placed on a different volume that is being snapshot. It will be ignored otherwise).
- ◆ Use caution when manually configuring the cache file sizes since they are used regardless of the sizes of the volumes being backed up. If enough space is not allocated, the backup job could fail with a VSP error.

### VSP and Interaction with Virus Scanners

Virus scanners can intermittently cause deletion of VSP snapshots and their cache files to fail after NetBackup backup jobs complete successfully. If this occurs, NetBackup has the following features to allow users to exclude VSP snapshots and cache files from virus scanning activity:

- ◆ VSP cache files are created in VSP cache file directories named `NBU_VSP_Cache`. For example, `D:\ NBU_VSP_Cache`  
  
Exclude directories are created from virus scanning activity if VSP snapshot deletion failures occur. The directories are named `NBU_VSP_Cache`.
- ◆ VSP cache files are named with a `.VSP` extension. Exclude files with `.VSP` extensions from virus scanning activity if VSP snapshot deletion failures occur.

## Busy File Wait

The **Busy File Wait** property specifies in seconds how long VSP should wait for a quiet period (quiesce wait time) before creating the snapshot. A quiet period is a time during which no file write-activity occurs on the drive being snapped using VSP. Default: 5 seconds. A value less than 5 seconds for the **Busy File Wait** property is not recommended because the data backed up with this property may be corrupted.

## Busy File Timeout

The **Busy File Timeout** property specifies in seconds how long VSP should wait for a quiet period to occur. If this time expires, the backup proceeds without VSP. Default: 300 seconds.

## Using VSP with Databases

There are special considerations regarding using VSP (Volume Snapshot Provider) to back up and restore databases.

Many popular database vendors provide a formal application program interface (API) specifically designed for use with backup products. VERITAS works closely with many database vendors to ensure these interfaces are stable, efficient, and reliable when used in conjunction with NetBackup and the various NetBackup database extension features. Many of these APIs were jointly developed to ensure that data is protected and can be restored when needed. Oracle, Microsoft (SQL Server, Exchange), IBM (Lotus Notes, DB2), NCR (Teradata), Sybase and Informix are examples of database vendors that provide an API for use with backup products. VERITAS strongly recommends that the NetBackup database extension features be used when a backup API is available and when backing up a database in a hot mode is required.

### Databases with an API

Hot backups are done on active databases and only by using these formal APIs will the confidence of a backup and the ability to perform a successful restore be achieved. VERITAS does not recommend that VSP be used for hot backups of these databases.

Cold or inactive backups of these databases may be possible with VSP, but success varies with each database vendor. Customers should contact the specific database vendor to identify the recommended method for database backup where data reliability is ensured as database programs recover from a point-in-time restore differently. If the data being backed up and restored does not conform to the specification designed into the database product being used, the integrity of the database can be in question.

### Databases without an API

When using VSP to back up databases that do not have a backup and restore API, the safest method is to back up the databases when the database is inactive (cold). For databases where there is no VERITAS database extension product, shut down the database and perform a file system level or cold backup.

If the databases cannot be backed up cold and the only option is a hot backup, set **Busy File Wait** to 5 seconds. If the file system does not achieve a quiescent or inactive state, NetBackup will not perform the VSP snapshot. NetBackup does not fail the backup when a quiescent state is not achieved. Instead, NetBackup continues the backup as if VSP was not being used. The result is that NetBackup skips open, active, or locked files. The backup job ends with an exit status code 1, indicating that the backup job completed but not all files were successfully backed up.

If VSP is used to back up database environments, VERITAS strongly recommends first backing up the data and validating that the backup exited with a Status 0. Then, restore the database and confirm the integrity of the data and the functionality of the database.

Using VSP to back up active databases without using a formal API presents risk. Customers should contact the database supplier to ensure support of database backups using point-in-time technology. Also, significant back up and restore testing should be performed to assure database availability and reliability.

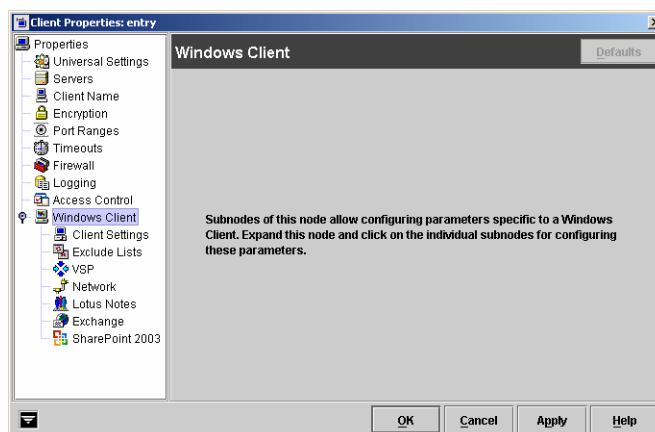



## Windows Client Properties

The **Windows Client** properties define NetBackup properties for Microsoft Windows clients.

**Windows Client** properties include:



- ◆ “Client Settings (Windows) Properties” on page 377
- ◆ “Exclude Lists Properties” on page 387
- ◆ “VSP (Volume Snapshot Provider) Properties” on page 454
- ◆ “Network Properties” on page 429
- ◆ “Lotus Notes Properties” on page 422
- ◆ “Exchange Properties” on page 385





## Master Server, Media Server, and Client Host Properties

---



## Managing NetBackup

## 8

This chapter contains topics related to the administration and management of NetBackup.

- ◆ “Powering Down and Rebooting NetBackup Servers” on page 466
- ◆ “Administering NetBackup Licenses” on page 469
- ◆ “Using the NetBackup License Utility to Administer Licenses” on page 473
- ◆ “Administering a Remote Master Server” on page 474
- ◆ “Using the NetBackup-Java Windows Display Console” on page 483
- ◆ “Configuring the NetBackup-Java Administration Console” on page 484
- ◆ “NetBackup-Java Performance Improvement Hints” on page 497
- ◆ “Administrator’s Quick Reference” on page 502
- ◆ “Managing Client Restores” on page 504
- ◆ “Goodies Scripts” on page 522
- ◆ “Server Independent Restores” on page 522
- ◆ “Configuring NetBackup Ports” on page 530
- ◆ “Load Balancing” on page 533
- ◆ “Using NetBackup with Storage Migrator” on page 535

## Powering Down and Rebooting NetBackup Servers

When closing down and restarting NetBackup servers, use the following recommended procedures.

### ▼ To power down a server

1. In the NetBackup Administration Console, click **Activity Monitor**, then select the Jobs tab to make sure no backups or restores are running.
2. Use the NetBackup Administration Console or the command line to stop the NetBackup Request daemon `bprd`. This stops additional backup and restore activity and to allows current activity to end gracefully:
  - ◆ In the NetBackup Administration Console, click **Activity Monitor**, then select the Processes tab. Right-click the request daemon (`bprd`) and select **Stop Daemon**.
  - ◆ From the command line, run:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

---

**Note** `bprdreq` does not run on a media server.

---

3. Run the system shutdown command.

---

**Note** The installation process copies the appropriate startup/shutdown script from `/usr/opensv/netbackup/bin/goodies` to `/init.d` and creates links to it from the appropriate `/rc` directory.

You can use system startup scripts to automatically start the Media Manager and NetBackup daemons when the system boots and use shutdown scripts to terminate them at system shutdown. See the *NetBackup Installation Guide* for instructions on editing the script.

---

4. Power down the server.

### ▼ To shut down all NetBackup daemons

From a command line, enter:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```

or

```
/usr/opensv/netbackup/bin/goodies/bp.kill_all
```

This script does not stop all processes on media servers, and is intended to stop daemons when there is no backup activity in progress.

▼ **To start up all NetBackup daemons**

From a command line, enter:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

▼ **To reboot a NetBackup master server**

1. Restart the system.
2. Ensure that `bprd`, `bpdbm`, and `vmd` are up by running the following script:

```
/usr/opensv/netbackup/bin/bpps -a
```

3. If necessary, start the NetBackup and Media Manager daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

▼ **To reboot a NetBackup media server**

1. Restart the system.
2. Start `ltid` if it is not already running:  
From the NetBackup Administration Console:

- a. Click **Activity Monitor**, then select the Processes tab.
- b. Right-click `ltid` and select **Start Daemon**.

From the command line, run:

```
/usr/opensv/volmgr/bin/ltid
```

## Displaying Active Processes with bpps

NetBackup provides a script called `bpps` that determines 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
```

Prevent `bpps` from displaying processes you do not want to check by adding the processes to an exclude list. Refer to comments within the script for more information.



To display both NetBackup and Media Manager options, run:



## Powering Down and Rebooting NetBackup Servers

---

```
/usr/opensv/netbackup/bin/bpps -a
```



## Administering NetBackup Licenses

The license key for each computer is initially entered when the software is installed. At some point you may need to modify the licensing, for example, when changing to a different level of NetBackup or adding separately-priced options.

**Note** When making and saving any license key updates in the NetBackup-Java Administration Console, you must restart the NetBackup Administration Console.

### ▼ To access license keys for a NetBackup server

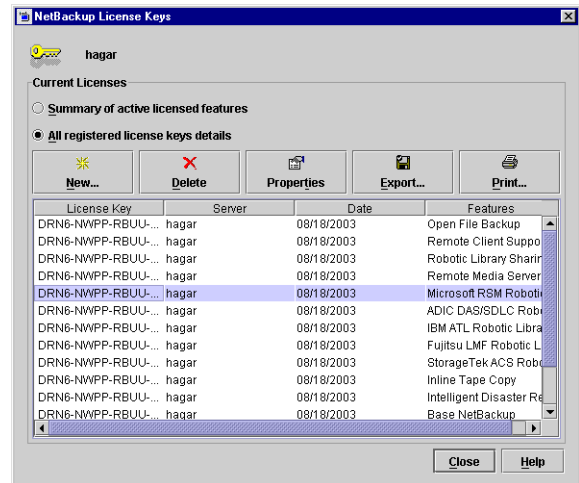
#### 1. Select a server:

- a. To view the license keys of the current server:  
In the NetBackup Administration Console, click **Help > License Keys**.
- b. To view the license keys of another server:  
Click **File > Change Server**, then select another server. Click **Help > License Keys**.

**Note** The licenses displayed are for the current server. To view the licenses for a particular master or media server, that server must be selected as the current server using **File > Change Server**.

#### 2. Choose to display either a summary listing or the details for each license key:

- ◆ Select **Summary of active licensed features** to show a summary of the active features that are licensed on this server. This view lists each feature and how many instances of it are permitted.
- ◆ Select **All registered license keys details** to show the details of the license keys registered on this server. This view lists each license key, the server where it is registered, when it was registered, and the features that it provides, and whether the feature is active or inactive.



## Administering NetBackup Licenses

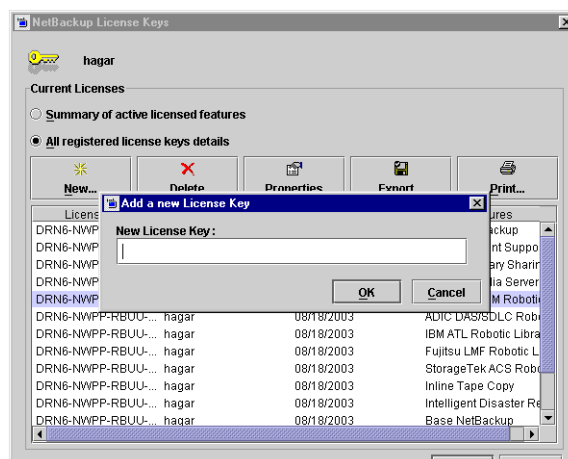
### 3. From the NetBackup License Keys dialog, you can:

- ◆ Add a new license
- ◆ Delete a license
- ◆ View the properties of one license
- ◆ Export the license listing

### ▼ To add new license keys

1. In the NetBackup License Keys dialog, click **New**.
2. In the Add a New License Key dialog, enter the license key and click **Add**. The new license key appears in the license listing.

**Note** After deleting the license keys, all the NetBackup utilities including NetBackup-Java Administration Console should be restarted.



### ▼ To print license key lists

1. In the NetBackup License Keys dialog, select the license key you wish to print. If no selection is made, all licenses are printed.

The printed information includes:

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



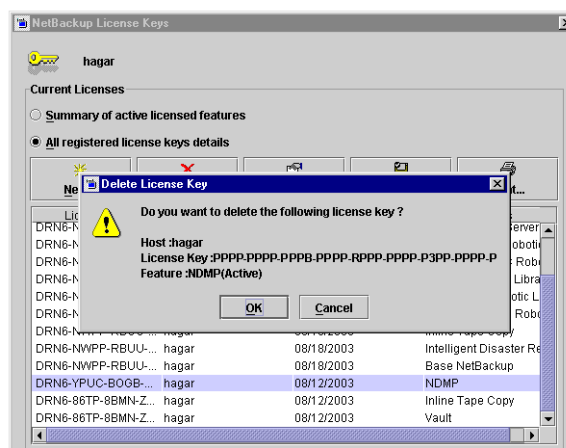
2. In the NetBackup License Keys dialog, click the **Print** button.  
The Print dialog appears.
3. Make your print selections and click **OK**.

#### ▼ To delete license keys

1. Select the license key you wish to delete from the license key list. If the key has more than one feature, all the features are listed in the dialog.
2. In the NetBackup License Keys dialog, click **Delete**. A confirmation dialog appears.
3. Click **Yes** to delete all the features associated with the key. The license key cannot be restored.


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

**Note** After deleting the license keys, all the NetBackup utilities including NetBackup-Java Administration Console should be restarted.



#### ▼ To view the properties of one license key


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



## Administering NetBackup Licenses

### ▼ To export license keys

1. In the NetBackup License Keys dialog, click **Export**. The **Export File Name** dialog appears.
2. Enter the path and file name where you'd like the key properties of all licenses to be exported.



The 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

## Using the NetBackup License Utility to Administer Licenses

### ▼ To start the NetBackup License Key utility

Run `/usr/openv/netbackup/bin/admincmd/get_license_key` command.

The License Key Utility menu appears:

```
License Key Utility

A) Add a License Key
D) Delete a License Key
F) List Active License Keys
L) List Registered License Keys
H) Help
q) Quit License Key Utility
```

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. Licenses that are expired 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.

## Administering a Remote Master Server

If your site has more than one NetBackup master server, you can configure the systems so multiple servers can be accessed from one NetBackup Administrator Console.


In order to access remote servers:

- ◆ First, make the remote server accessible to the local server. See the following section, “Adding a NetBackup Server to a Server List” on page 474
- ◆ Second, indicate the remote server you want to administer. See “Choosing a Remote Server to Administer” on page 478.

### Adding a NetBackup Server to a Server List

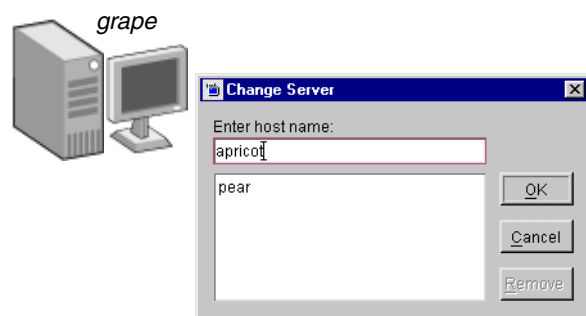
In order for a local host to administer a remote server, the name of the local host must appear in the server list of the remote server.

For example, assume server *grape* wants to remotely administer host *apricot*.



Grape selects **File > Change Server** and types *apricot* as the host name.

If *grape* is not listed on the server list of *apricot*, *grape* receives an error message after trying to change servers to *apricot*.



Assuming *apricot* is an authorized NetBackup server, the message that appears may indicate that *grape* is considered invalid because it does not appear on the server list of *apricot*.

To add *grape* to the server list of *apricot*, follow the steps in “To add a server to a UNIX server list.” For other reasons why a remote server may be inaccessible, see “If You Cannot Access a Remote Server” on page 482.

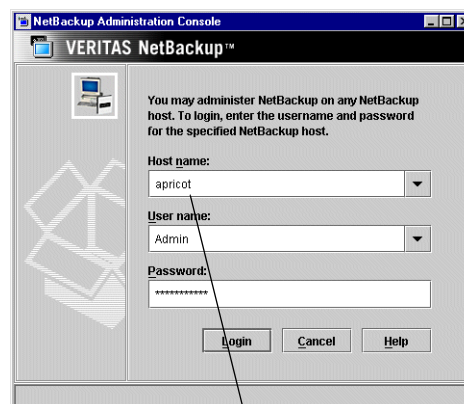


**Note** If you are logging in to a remote master server through the log in dialog, it is not necessary for the name of the local host to appear in the server list of the remote server. This method of logging in to a remote host is explained in “To indicate a remote system upon log in” on page 479.

#### ▼ To add a server to a UNIX server list

1. Access the server properties of the destination host using one of the following methods:

- ◆ Start the NetBackup-Java Administration Console (jnbSA) on the local server (*grape*). Indicate destination host *apricot* on the log in dialog. The jnbSA command is described in *NetBackup Commands for UNIX and Linux*.
- ◆ Start the NetBackup-Java Windows Display Console on a Windows system. Indicate destination host *apricot* on the log in dialog.
- ◆ Physically go to the destination host (*apricot*) and start jnbSA. Indicate *apricot* on the log in dialog.



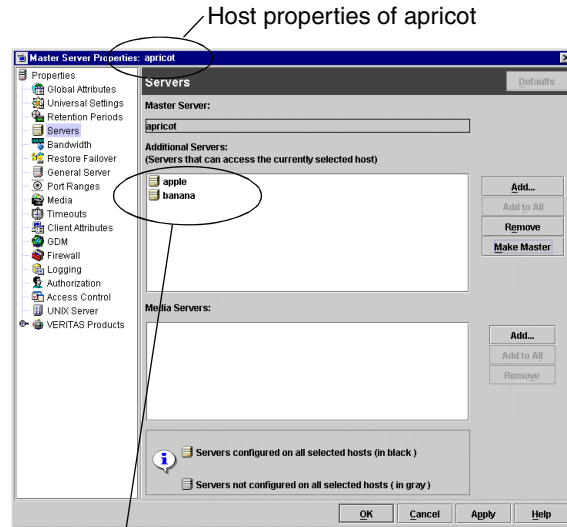
Log in to *apricot* from *grape* (provided the user name has sufficient privileges), or log in at *apricot*

2. In the NetBackup Administration Console, expand **Host Properties > Master Servers**.

## Administering a Remote Master Server


3. Double-click the server name (*apricot*) to view the properties.
4. Select **Servers** to display the server list.

The **Additional Servers** list contains, as the dialog explains, “Servers that can access the currently selected host.” Since the **Additional Servers** list does not include server *grape*, *apricot* considers *grape* to be an invalid server.

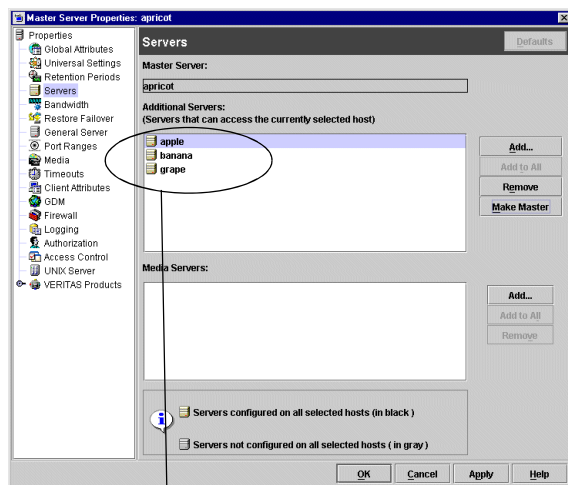


Currently, *apricot* allows remote access by two additional servers: *apple* and *banana*

5. To add a server to the server list, click **Add**. The New Server dialog appears.
6. Type the server name (*grape*) in the field and click **Add** to add the server to the list. Click **Close** to close the dialog without adding a server to the list.



7. As when changing any NetBackup property through the Host Properties dialogs, restart all daemons and utilities on the server where the change was made to ensure that the new configuration values are used. Restart the NetBackup Administration Console as well.



*Apricot* now includes *grape* among the servers to which it allows remote access

---

**Note** The `bp.conf` file on every UNIX server contains `SERVER` and possibly `MEDIA_SERVER` entries. The server list in the properties dialog represents these entries. Hosts listed as media servers have limited administrative privileges.

---

▼ **To add a server to a Windows server list**

1. Go to the destination host and start the NetBackup Administration Console.
2. Expand **Host Properties** > **Master Server**.
3. Double-click the server name to view the properties.
4. Select the **Servers** tab to display the server list. The server list contains, as the dialog explains, "Servers that can access the currently selected host."
5. To add a server to the **Additional Server List**, click **Add**.
6. Enter the server name in the **New Server** dialog. Click **Add** or **Close**.
7. Restart all services on the server where the change was made to ensure that the new configuration values are used. Restart the NetBackup Administration Console as well.

## Choosing a Remote Server to Administer

Indicate a remote server using one of the following methods:

- ◆ Select the **File > Change Server** menu command.
- ◆ Specify the remote server as hostname upon NetBackup logging in using the NetBackup-Java console.

### ▼ To use the Change Server command to administer a remote server

1. Start the NetBackup Administration Console on a NetBackup-Java capable machine:

Log in and run jnbSA:

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



2. In the NetBackup Administration Console log in screen, specify the local server to manage.

3. Click **Login**.

4. Select **Master Server** in the left pane (tree view) of the NetBackup Administration Console.

5. Select **File > Change Server**.

6. Type or select the host name and click **OK**.



**Note** When moving between Master A and Master B, if the user's identity has the necessary permissions on both machines, the user will transition to Master B without needing to set up any trust relationships as was required in NetBackup 4.5. If the user's identity on Master B, that has administrative privileges is different from the user's identity on Master A, the user would be required to reauthenticate. This can be done from the NetBackup Administration Console by using **File > Login as New User...** for Windows or closing and reopening the NetBackup-Java Administration Console.



▼ To indicate a remote system upon log in


1. Log in to the NetBackup client or server where you want to start the NetBackup Administration Console.
2. Start the NetBackup Administration Console on the local system:

- ◆ For example, to start the console on a Solaris system named *tiger*, log in on *tiger* and run the following command line:

```
/usr/opencv/java/jnbSA
```

- ◆ To start the console on a Windows system named *bear*, from the Windows desktop, select **Start > Programs > VERITAS NetBackup > NetBackup-Java Version 6.0**. (The system must have the Windows Display Console installed.)

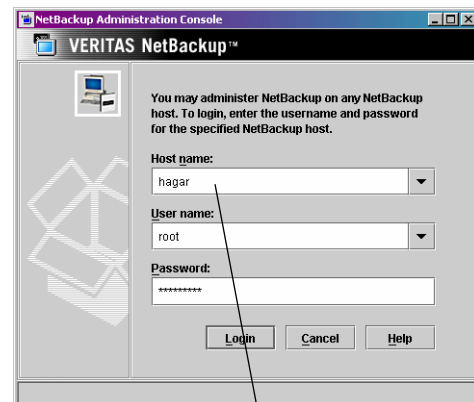
The log in screen appears.



3. In the **Host name** field, type the name of the remote NetBackup server you want to manage. In this example, *shark*.
4. Type in the user name and password for an authorized NetBackup administrator (for example: *root*), then click **Login**.

This process logs you in to the NetBackup-Java application server program on the specified server.

The NetBackup Administration Console appears. The console program continues to communicate through the server you specified for the remainder of the current session.



Type in the name of the remote server you'd like to administer

## Administering via a NetBackup Client

Even though a machine may not contain the NetBackup server software, running the NetBackup Administration Console on a client is useful in order to administer a NetBackup server remotely. You can run the NetBackup Administration Console on a client under the following conditions:

- ◆ On a Windows client if the NetBackup-Java Windows Display Console is installed.
- ◆ On a UNIX client if the client is NetBackup-Java capable.

## Using the Remote Administration Console

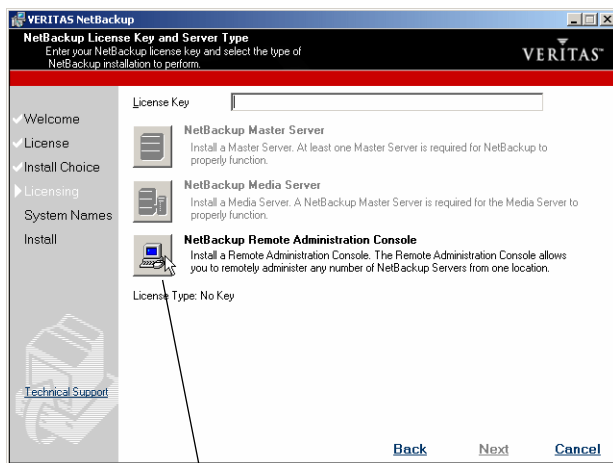
You can install the Remote Administration Console on a Windows machine to fully administer or manage any remote NetBackup server—Windows or UNIX. No license is required to install the Remote Administration Console.

Upon installation of the Remote Administration Console, the Administration Console and the client software is installed. No NetBackup master or media server software is installed. The presence of the client software enables the machine to be backed up like any other client.

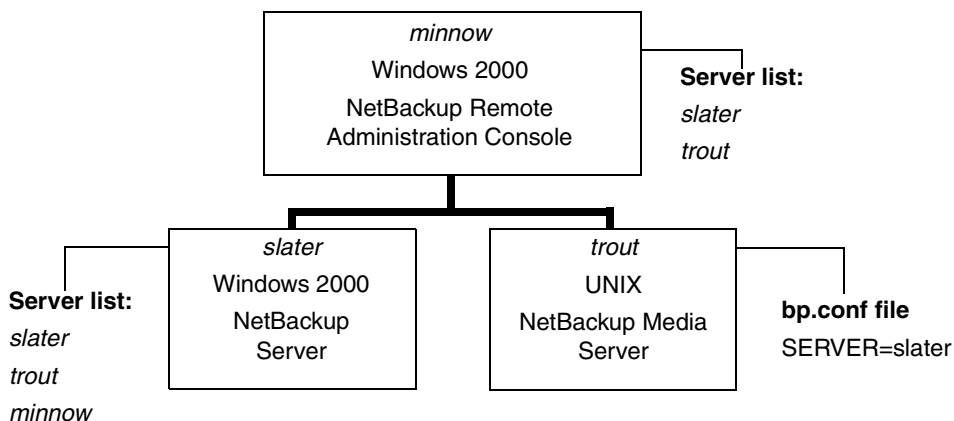
In the following figure, the administrator can use the Remote Administration Console on *minnow* to add a comment for a drive in the configuration on the server *trout*.

The administrator can manage NetBackup on *slater* or *trout* from any of the systems shown.

*minnow* must be listed in the server list on *slater* but does not need to be listed on *trout* because *trout* is a media server.



Installing the Remote Administration Console installs the administration console and client software, but no server software



After starting the Remote Administration Console, change to a NetBackup server by clicking **File > Change Server**.

## If You Cannot Access a Remote Server

In order to administer a server from another master server make sure that the following conditions are true:

- ◆ The destination server is operational.
- ◆ NetBackup daemons are running on both hosts.
- ◆ There is a valid network connection.
- ◆ The user has administrative privileges on the destination host.
- ◆ The current host is listed in the server list of the destination host “Adding a NetBackup Server to a Server List” on page 474. This is not required for a media server, client, media and device management, or device monitoring.

To ensure that the new server entry is used by all NetBackup processes that require it, stop and restart:

- ◆ The NetBackup Database Manager and NetBackup Request Manager services on the remote server if it is Windows.
- ◆ The NetBackup Database Manager (`bpdbm`) and NetBackup Request Manager (`bprd`) on the remote server if it is UNIX.
- ◆ Authentication is set up correctly, if used.
- ◆ If you have problems changing servers when configuring media or devices or monitoring devices:
  - ◆ If the remote server is Windows, verify that the NetBackup Volume Manager service is running on that server and start it if necessary.
  - ◆ If the remote server is UNIX, verify that the Media Manager Volume daemon is running on that server and start it if necessary.
- ◆ 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 *Media Manager System Administrator's Guide* for instructions.

## Using the NetBackup-Java Windows Display Console

This section describes methods to authorize NetBackup-Java users on a Windows system, to restrict access to NetBackup-Java applications on a Windows system, and to authorize users to use a subset of the NetBackup-Java administrator applications.

### Authorizing NetBackup-Java Users on Windows

To use the NetBackup-Java Windows Display Console, you must first log in to the NetBackup-Java application server that is on the NetBackup host where you want to perform NetBackup administration or user operations.

Users log in to the application server when logging in to the dialog that appears when starting the console. This is done through the Windows Display Console or by starting the NetBackup Administration Console on a UNIX system.

During log in, users provide a user name and password that is valid on the computer specified in the **NetBackup host** field of the log in dialog box.

The user name for Windows must be of the form: *domainname\username*

*domainname* specifies the domain of the NetBackup host. The domain is not required if the NetBackup host is not a member of a domain.

The NetBackup-Java application server authenticates the user name and password by using standard Windows authentication capabilities for the specified computer.

If neither NetBackup Access Control nor Enhanced Authorization and Authentication are configured for the users, by default the NetBackup-Java application server provides authorization data that allows all users that are members of the administrator group for the host's domain to use all the NetBackup-Java applications. Other users are allowed to access only Backup, Archive, and Restore.

If desired, restrict access to NetBackup-Java or some of its applications by creating an *nbjava\_install\_path\java\auth.conf* authorization file as described in "Restricting Access on Windows" on page 483.

See "Configuring the NetBackup-Java Administration Console" on page 484 for additional details.

### Restricting Access on Windows

To restrict access to one or more of the NetBackup-Java applications, create the following file on the Windows system:

```
nbjava_install_path\java\auth.conf
```

## Configuring the NetBackup-Java Administration Console

---

Add an entry in `auth.conf` for every user that will be granted access to the NetBackup-Java applications. The existence of this file, along with the entries it contains, prohibits unlisted users from accessing NetBackup-Java applications on the Windows system. The following is a sample `auth.conf` file on a Windows system:

```
mydomain\Administrator ADMIN=ALL JBP=ALL
mydomain\joe ADMIN=ALL JBP=ALL
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

See “Authorizing NetBackup-Java Users” on page 486 for additional details.

## Configuring the NetBackup-Java Administration Console

The following sections contain information about the NetBackup-Java Administration Console. This is the NetBackup-Java Windows Display Console on Windows systems.

### NetBackup-Java Administration Console Architectural Overview

The NetBackup-Java Administration Console is a distributed application consisting of two separate, major system processes:

- ◆ The NetBackup Administration Console graphical user interface
  - ◆ Available on UNIX systems by running `jnbSA`
  - ◆ Available on Windows systems by installing and running the NetBackup-Java Windows Display Console
- ◆ The application server (`bpjava` processes)

These processes may be run on two physically different NetBackup hosts. This distributed application architecture holds true for the UNIX Backup, Archive, and Restore client graphical user interface (`jbpSA`) as well.

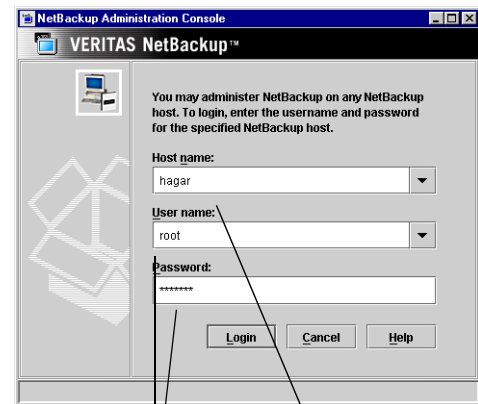
After the NetBackup-Java Administration Console interface is started, the user is required to log in to the application server on the host specified in the log in dialog.

The NetBackup-Java Administration Console can be started in one of the following methods:

- ◆ Run the `jnbSA` command on a UNIX system
- ◆ Select **Start > VERITAS NetBackup > NetBackup-Java Version 6.0** on a Windows system on which the Windows Display Console is installed

**Note** The NetBackup server or client you specify on the log in dialog must be running the same version of NetBackup as is installed on the machine running the NetBackup Administration Console.

NetBackup log in dialog:



Application server


User name and password must be valid on application server

The log in credentials of the user are authenticated by the application server on the host specified in the NetBackup Administration Console log in dialog using standard UNIX system user account data and associated APIs. This means that the provided log in credentials must be valid on the host specified in the log in dialog.

The server that is usually the object of all administrative tasks is the one specified in the NetBackup Administration Console log in dialog.

The exception to this 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 log in dialog).

Remote server



Regardless of the server being administered (a remote server or the server specified on the log in dialog), all administrative tasks performed in the NetBackup Administration Console make requests of the application server and are run on the application server host.

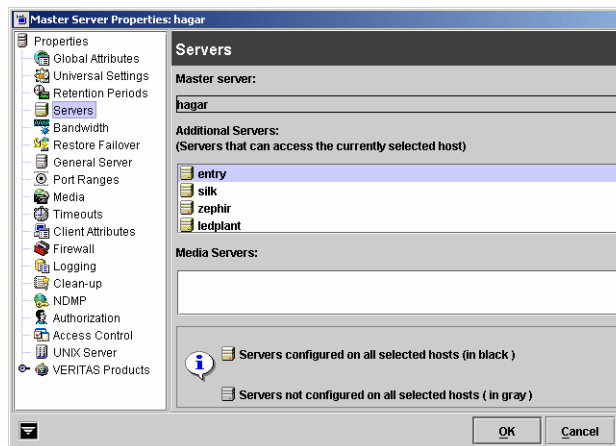
However, regardless of which NetBackup authorization method is configured, authorization for tasks in the Administration Console is specific to the server being administered. For example, if NetBackup-Java authorization capabilities are in use on Host\_A, then **Change Server** is used to change to Host\_B, the permissions as configured in the `auth.conf` on Host\_B are honored.

## Configuring the NetBackup-Java Administration Console

For successful administration of a remote server, the application server host must be included in the server list of the remote server. (See “Adding a NetBackup Server to a Server List” on page 474.)

This context (switching to a remote server from the application server) also applies to the Enhanced Authentication and Authorization capabilities (see the *NetBackup System Administrator's Guide, Volume II*). For instance, the host where the NetBackup Administration Console is running is not the host requiring access to any server host unless both the NetBackup-Java Administration Console and its application server are running on the same host.

In addition, this context (switching to a remote server from the application server) applies to configuration scenarios for administration in firewall environments with one exception: the host where the NetBackup Administration Console is running must be able to access the `vnetd` daemon on either the remote host or the host specified in the log in dialog for activity monitoring tasks.



## Authorizing NetBackup-Java Users

The *NetBackup System Administrator's Guide, Volume II* documents two types of user authorization: NetBackup Access Control and Enhanced Authorization and Authentication. If neither method is configured, you may choose to authorize users of the NetBackup-Java administration console for specific applications. In addition, you may use Enhanced Authorization and the NetBackup-Java console capabilities authorization together. (See “USE\_NBJAUTH\_WITH\_ENHAUTH” on page 496.) The following sections document how to do so.

NetBackup Access Control and enhanced authorization, when configured as described, always take precedence over the capabilities authorization of NetBackup-Java as described in “Configuring Nonroot Usage” on page 490.

When NetBackup Access Control or Enhanced Authorization is configured, but a user is not authorized as an administrator of NetBackup, the capabilities allowed to this user in the Backup, Archive, and Restore (`jbpsa`) application are those specified for the user in the `auth.conf` file resident on the host specified in the NetBackup-Java log in dialog.



Users of the NetBackup-Java interfaces must log in to the NetBackup-Java application server that is on the NetBackup host where they want to perform administrator or user operations.

The `/usr/opensv/java/auth.conf` file contains the authorization data for accessing NetBackup-Java applications. This file exists only on NetBackup-Java capable machines where the NetBackup-Java interface software is installed. The default `auth.conf` file provides the following authorizations:

- ◆ On NetBackup servers: Administration capabilities for the root user and user backup and restore capabilities for all other users.
- ◆ On NetBackup clients: User backup and restore capabilities for all users.

On all other UNIX NetBackup systems, the file does not exist but the NetBackup-Java application server provides the same default authorization. To change these defaults on other UNIX systems, you must create the `/usr/opensv/java/auth.conf` file.

To perform remote administration or user operations with `jbpSA` a user must have valid accounts on the NetBackup UNIX server or client machine.

---

**Note** Nonroot or non-administrator users can be authorized to remotely administer Windows NetBackup servers from the NetBackup-Java Console by setting up the desired authorization in the `auth.conf` file on the Windows server. The `auth.conf` file must contain entries for the UNIX usernames used on the log in dialog of the NetBackup-Java Console. The `auth.conf` file must reside in `install_path\VERITAS\java` on each Windows server you wish to provide nonroot administration capability. If no `auth.conf` file exists, or it doesn't contain an entry for the username and the host authorization between the two is set up, (that is, `SERVER` entries in the configuration of each), the user will have the same privileges to administer the remote Windows server as they have on the server specified in the log in dialog for the NetBackup-Java Console.

---

## Authorization File Characteristics

The released version of the UNIX `/usr/opensv/java/auth.conf` file is installed on all NetBackup-Java capable hosts and contains only the following entries:

```
root ADMIN=ALL JBP=ALL
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

- ◆ The first field of each entry is the user name that is granted access to the rights specified by that entry. In the released version, the first field allows root users to use all of the NetBackup-Java applications.

## Configuring the NetBackup-Java Administration Console

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 an entry containing an asterisk (\*) in the username field; users without entries cannot access any NetBackup-Java applications. Any entries that designate specific user names must precede a line that contains an asterisk in the username field.

---

**Note** The asterisk specification cannot be used to authorize all users for any administrator capabilities. Each user must be authorized via individual entries in the `auth.conf` file.

---

If you wish to deny all capabilities to a specific user, add a line indicating the user before a line starting with an asterisk. For example:

```
mydomain\ray ADMIN= JBP=
* ADMIN=JBP JBP=ENDUSER+BU+ARC
```

- ◆ The remaining fields specify the access rights.
  - ◆ The ADMIN keyword specifies the applications that the user can access. ADMIN=ALL allows access to all NetBackup-Java applications and their related administrator related capabilities. To allow the use of only specific applications, see “Authorizing Nonroot Users for Specific Applications” on page 490.
  - ◆ The JBP keyword specifies what the user can do with the Backup, Archive, and Restore client application (jbpSA). JBP=ALL allows access to all Backup, Archive, and Restore capabilities, including those for administration. To allow only a subset of those capabilities, see “Capabilities Authorization for jbpSA” on page 491.
  - ◆ 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 has an asterisk in the first field, which means that NetBackup-Java validates any user name for access to the Backup, Archive, and Restore client application (jbpSA). JBP=ENDUSER+BU+ARC allows end users to only back up, archive and restore files.

When starting the NetBackup-Java administrator applications or the Backup, Archive, and Restore application (jbpSA), you must provide a user name and password that is valid on the machine that you specify in the NetBackup host field of the log in dialog. The NetBackup-Java application server authenticates the user name and password by using the system password file data for the specified machine, so the password must be the same as used when logging in to that machine.

For example, assume you log in with:

```
username = joe
password = access
```

Here you must use the same user name and password when logging in to NetBackup-Java.

---

**Note** The NetBackup-Java log in box accepts passwords greater than eight characters. However, only the first eight are significant when logging in to a NetBackup-Java application server running on a UNIX system.

---

It is possible to log in to the NetBackup-Java application server under a different user name than the one used for logging in to the operating system. For example, if you log in to the operating system with a user name of joe, you could subsequently log in to jnbSA as root. When you exit, in this instance, some application state information (for example, table column order) is automatically saved in joe's `$HOME/.java/.userPrefs/vrts` directory and is restored the next time you log in to the operating system under account joe and initiate the NetBackup-Java application. This method of logging in is useful if there is more than one administrator because it saves the state information for each of them.

---

**Note** NetBackup-Java creates a user's `$HOME/.java/.userPrefs/vrts` directory the first time an application is exited. Only NetBackup-Java applications use the `.java/.userPrefs/vrts` directory.

---

If the user name is not valid according to the contents of the `auth.conf` file, the user sees the following error message in a popup message dialog and all applications are inaccessible.

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, you have two basic choices for types of entries in the `auth.conf` file:

- ◆ Use the released defaults to allow anyone with any valid user name to use the Backup, Archive, and Restore client application (jbpSA) and only root users to use 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 when browsing directories and files to back up or restore.

---

## Configuring Nonroot Usage

### Authorizing Nonroot Users for Specific Applications

It is possible to authorize nonroot users 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:

---

#### **auth.conf ADMIN Identifiers for Administrator Applications**

---

|            |                                                           |
|------------|-----------------------------------------------------------|
| ALL        | Indicates administration of all applications listed below |
| 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                                   |
| VLТ        | Vault Management                                          |

---

For example, to give a user named `joe` access only to the Device Monitor and Activity Monitor, add the following entry to the `auth.conf` file:

```
joe ADMIN=DM+AM
```

If necessary for a nonroot administrator to modify files used by the NetBackup-Java Administration Console, the script `/usr/opensv/java/nonroot_admin_nbjava` can be executed to change the permissions on the following files:

```
/usr/openssl/java/auth.conf
/usr/openssl/java/Debug.properties
/usr/openssl/java/nbj.conf
```

## Capabilities Authorization for jbpSA

Capabilities authorization in the Backup, Archive, and Restore interface enables certain parts of the user interface to allow one to perform certain tasks. Not all tasks can be performed successfully without some additional configuration. The following require additional configuration and are documented elsewhere:

- ◆ Redirected restores. See “Managing Client Restores” on page 504.
- ◆ User backups or archives require a policy schedule of these types and the task to be submitted within the time window of the schedule.

To authorize users for a subset of Backup, Archive, and Restore capabilities, use the following identifiers for the JBP keyword in the `auth.conf` file:

- ◆ `ENDUSER` - Allows user to perform restore tasks from true image, archive or regular backups plus redirected restores
- ◆ `BU` - Allows user to perform backup tasks
- ◆ `ARC` - Allows user to perform archive tasks (BU capability required for this)
- ◆ `RAWPART` - Allows user to perform raw partition restores
- ◆ `ALL` - Allows user to perform all of the above actions, including restoring to a different client from the one you are logging into (that is, server-directed restores). Server-directed restores can only be performed from a NetBackup master server.

The following example entry allows a user named *bill* to restore but not back up or archive files:

```
bill ADMIN=JBP JBP=ENDUSER
```

## Runtime Configuration Options

On UNIX systems, file `/usr/openssl/java/nbj.conf` contains configuration options for the NetBackup-Java Administration Console. Enter one option per line, following the same syntax rules as exist for the `bp.conf` file.

On Windows systems, the analogous file containing configuration options for the NetBackup-Java Windows Display Console is `nbjava_install_path\java\setconf.bat`.

There are set commands in this file for each of the configuration options described in the following sections. To make changes, simply change the value after the equal sign in the relevant set command.

## Configuring the NetBackup-Java Administration Console

The NetBackup-Java Administration Console configuration options consist of the following:

- ◆ BPJAVA\_PORT (see “BPJAVA\_PORT, VNETD\_PORT” on page 492)
- ◆ FORCE\_IPADDR\_LOOKUP (see “FORCE\_IPADDR\_LOOKUP” on page 492)
- ◆ INITIAL\_MEMORY (see “INITIAL\_MEMORY, MAX\_MEMORY” on page 494)
- ◆ MAX\_MEMORY (see “INITIAL\_MEMORY, MAX\_MEMORY” on page 494)
- ◆ MEM\_USE\_MEMORY (see “MEM\_USE\_WARNING” on page 495)
- ◆ NBJAVA\_CLIENT\_PORT\_WINDOW (see “NBJAVA\_CLIENT\_PORT\_WINDOW” on page 495)
- ◆ NBJAVA\_CONNECT\_OPTION (see “NBJAVA\_CONNECT\_OPTION” on page 496)
- ◆ VNETD\_PORT (see “BPJAVA\_PORT, VNETD\_PORT” on page 492)
- ◆ USE\_NBJAUTH\_WITH\_ENHAUTH (see “USE\_NBJAUTH\_WITH\_ENHAUTH” on page 496)

### BPJAVA\_PORT, VNETD\_PORT

The following ports are the configured ports for the `bpjava-msvc` and `vnetd` daemon processes. These ports are registered with the Internet Assigned Numbers Authority (IANA).

| Port        | Process and Registered Default Port Number |
|-------------|--------------------------------------------|
| bpjava-msvc | BPJAVA_PORT=13722                          |
| vnetd       | VNETD_PORT=13724                           |

VERITAS recommends that these ports are not changed. If changes are necessary, make the change on all NetBackup hosts in the relevant NetBackup cluster as described in the *NetBackup Installation Guide*. In addition, the value must be set in the corresponding `nbj.conf` (UNIX) or `setconf.bat` (Windows) configuration option.

### FORCE\_IPADDR\_LOOKUP

The `FORCE_IPADDR_LOOKUP` configuration option specifies whether NetBackup will perform 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 do not perform an IP address lookup to determine if two host name strings are indeed the same host. They will be considered the same host if the host name strings compare equally or a short name compares equally to the short name of a partially or fully qualified host name.

1 = Indicates to perform an IP address lookup if the two host name strings do not match to determine if they have the same host (default). The default is to perform an IP address lookup if necessary to resolve the comparison. The IP address lookup will not be 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 using host name aliases.

---

There are many places in the NetBackup Administration Console where comparisons of host names are done to determine if the two are indeed the same host (when using the **File > Change Server** command, for example).

The IP address lookup can be time consuming and result in slower response time. However, it is important to be accurate with the comparisons. If following the rules for host names as documented in Chapter 5 of the *NetBackup System Administrator's Guide, Volume II*, there should not be any issues as the string comparison will be accurate.

No IP address lookup should be necessary if you are always consistent in the way you specify the host name in the NetBackup Administration Console login dialog and it matches how the host names are configured in NetBackup. Host names are identified in the server list 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 will result in no IP address lookup and the hosts will be considered 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 will be considered different for any comparisons of short, partially or fully qualified host names of *eagle* and *hawk* regardless of aliasing.

- ◆ `FORCE_IPADDR_LOOKUP = 1`

## Configuring the NetBackup-Java Administration Console

---

Comparisons of the following will result in no IP address lookup and the hosts will be considered the same host.

```
eagle and eagle
eagle.abc and eagle.abc
eagle.abc.def and eagle.abc.def
```

However, in addition to all comparisons of *eagle* and *hawk*, the following will result in an IP address lookup to determine 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).

VERITAS recommends running the NetBackup-Java Administration Console, the NetBackup-Java Windows Display Console, or the NetBackup, Archive, and Restore user interface on a machine with 1 gigabyte of physical memory with 256 megabytes of memory available to the application.

INITIAL\_MEMORY specifies how much memory is allocated for the heap when the JVM starts. It is unlikely that this value will require changing as the default is sufficient for quickest initialization of jnbSA, the NetBackup-Java Windows Display Console, or jbpSA on a machine 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 the JVM uses for dynamically allocated objects and arrays. This is useful if the amount of data is large (for example, a large number of jobs in the Activity Monitor).

On UNIX systems, the maximum memory allocation can also be specified as part of the jnbSA or jbpSA command. For example:

```
jnbSA -mx 512M
```

Default = 256M (megabytes).



## MEM\_USE\_WARNING

The `MEM_USE_WARNING` configuration option specifies the percent of memory used compared to `MAX_MEMORY`, at which time a warning dialog is displayed 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 nonreserved ports on this computer that are used for connecting to the NetBackup-Java application server or the `bpjobjd` daemon (or service on Windows) from the NetBackup-Java Administration Console's Activity Monitor.

This option uses the following format:

```
NBJAVA_CLIENT_PORT_WINDOW = n m
```

Where:

- ◆ *n* indicates the first in a range of nonreserved ports used for connecting to the `bpjava` processes on the NetBackup-Java application server or the `bpjobjd` daemon or Windows service from the NetBackup-Java Administration Console's Activity Monitor.

If *n* is set to 0, the operating system determines the nonreserved port to use (default).

- ◆ *m* indicates the last in a range of nonreserved ports used for connecting to the NetBackup-Java Administration Console/NetBackup-Java Windows Display Console.

If *n* and *m* are set to 0, the operating system determines the nonreserved 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 stating that there was an invalid value during initialization.

---

**Note** Performance is somewhat reduced with the use of `NBJAVA_CLIENT_PORT_WINDOW`.

---

## Configuring the NetBackup-Java Administration Console

### NBJAVA\_CONNECT\_OPTION

The `NBJAVA_CONNECT_OPTION` configuration option specifies whether the connection to the NetBackup-Java application server is done via the `vnetd` daemon (`VNETD_PORT`) or directly via the application server's port (`BPJAVA_PORT`). The option also specifies the call-back method the server or client will use when communicating with the NetBackup-Java consoles (`jnbSA`, `jbpSA`).

The default for `NBJAVA_CONNECT_OPTION` requires only that the `vnetd` port be accessible through any firewall.

```
NBJAVA_CONNECT_OPTION = [0 | 1]
```

Where:

0 = Indicates direct connection to the application server and the traditional call-back method.

1 = Indicates connection via `vnetd` and the no call-back method (default).

### USE\_NBJAUTH\_WITH\_ENHAUTH

The `USE_NBJAUTH_WITH_ENHAUTH` configuration option specifies whether or not NetBackup-Java Capabilities Authorization will be used with NetBackup Enhanced Authentication and Authorization.

Where:

0 = Do not allow NetBackup-Java Capabilities Authorization with NetBackup Enhanced Authentication and Authorization. NetBackup Enhanced Authentication and Authorization will take precedence over the NetBackup-Java console's Capabilities Authorization. The NetBackup-Java console's Capabilities Authorization configuration file, `auth.conf`, will be ignored. (Default.)

1 = Allow NetBackup-Java Capabilities Authorization with NetBackup Enhanced Authentication and Authorization. The NetBackup-Java console's Capabilities Authorization can be used to further restrict access by users to various console components (**Policies**, for example), when NetBackup Enhanced Authentication and Authorization is configured.

## Configuration Options Relevant to `jnbSA` and `jbpSA`

There are several configuration options available to administrators when using the NetBackup-Java Administration Console or the NetBackup, Archive, and Restore user interface through the NetBackup-Java Administration Console.

## Logging Command Lines Used by the NetBackup Interfaces

You may find it helpful to see which command lines are used by the NetBackup-Java Administration Console or the NetBackup, Archive, and Restore user interface. To log the command lines used by jnbSA or jbpSA to a log file, use option `-lc`. No value is necessary. For example:

```
/usr/opensv/java/jbpSA -lc
```

**Note** jnbSA and jbpSA don't always use the command lines to retrieve or update data. The interfaces have some protocols that instruct the application server to perform tasks using NetBackup and Media Manager APIs.

## 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 allow the administrator and users to customize the following aspects of jbpSA operation

- ◆ `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 the number of days to keep job and progress log files generated by the NetBackup-Java Backup, Archive, and Restore application (jbpSA). These files are written into the  
`/usr/opensv/netbackup/logs/user_ops/_username_/jobs` and  
`/usr/opensv/netbackup/logs/user_ops/_username_/logs` directories.  
There is a directory for each user that uses the NetBackup-Java applications. The default is three days.

The number of days to keep the NetBackup-Java GUI log files contained in `/usr/opensv/netbackup/logs/user_ops/nbjlogs` is controlled by this option as well.

## NetBackup-Java Performance Improvement Hints

The most important factor to consider when faced with performance issues while using the NetBackup-Java Administration Console, the NetBackup-Java Windows Display Console, or the NetBackup Backup, Archive, and Restore user interface is the platform on which the console is running. Regardless of the platform, you have the choice of running the Administration Console from the following locations:

- ◆ Locally on your desktop host (on supported Windows and UNIX platforms), or

## NetBackup-Java Performance Improvement Hints

---

- ◆ Remotely and displaying back to your desktop host (from supported UNIX platforms)

The recommended method for using these consoles is to run the consoles locally on your desktop host. This method provides the best performance and does not exhibit font and display issues that can be present in some remote display back configuration cases.

### What it Means to be Running the Java Console Locally on a UNIX Platform

On supported UNIX platforms, you are running the console locally if you enter the `jnbSA` or `jbpSA` commands on the same host on which the console is displayed. That is, your display environment variable is set to the host on which you entered the `jnbSA` or `jbpSA` commands.

Though improvements in the Java technology have made remote X-display back potentially viable on some platforms, there continues to be problems with certain controls in the consoles. For example, incorrect combo box operations, sluggish scrolling and display problems in tables with many rows. More serious issues have also occurred. For example, consoles aborting and hanging caused by a Java Virtual Machine (JVM) failure when run in this mode on some platforms with a variety of configurations. These JVM failures have most often been seen on the AIX platform. *Therefore, VERITAS cannot recommend running the consoles in a remote X-display back configuration.*

### What it Means to be Running the Console Locally on a Windows Platform

On Windows platforms, you are running the console locally if you start the Windows Display Console by selecting **Start > VERITAS NetBackup > NetBackup-Java Version 6.0** menu item or its equivalent desktop shortcut. This Start menu item or shortcut appears if you install the optional NetBackup-Java Windows Display Console available on the main NetBackup for Windows installation screen.

### How do I Run a Console Locally and Administer a Remote Server?

The NetBackup Administration Console and the Backup, Archive, and Restore user console are distributed applications that consist of two major and separate system processes that can run on different machines. For example:

- ◆ The NetBackup Administration Console on one machine, and
- ◆ the console's application server - `bpjava` processes on another machine.

While the NetBackup Administration Console does not have to run on a NetBackup server host, the application server must run on this host in order for you to be able to administer NetBackup. Refer to “NetBackup-Java Administration Console Architectural Overview” on page 484 for more details.

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. This distributed application architecture enables direct (logically local) administration (either server or client backup/restore tasks) of all NetBackup platforms even though the consoles themselves only run on a subset of the NetBackup supported platforms.

When logging into the NetBackup-Java Administration Console, you specify a host name. This is the machine where the application server (`bpjava`) runs. For example, a NetBackup master server. All requests or updates initiated in the console are sent to its application server running on this host.

## How do I Make the Console Perform Even Better?

Performance of the NetBackup-Java applications depends on the environment where the applications are running, including available resources and network throughput. The default configuration of NetBackup-Java, specifically the `INITIAL_MEMORY` and `MAX_MEMORY` configuration options, assumes sufficient memory resources on the machine where the console is running, for example, where the `jnbSA` command is executed or the NetBackup-Java Windows Display Console is started.

Following are guidelines for improving performance:

- ◆ Consider the network communication speed and the amount of data being transferred.
- ◆ Consider the amount of work being performed on the relevant machines.

Run NetBackup-Java on a machine that has a low level of activity. For example, there can be dramatic differences in response time when other memory-intensive applications are running on the machine. (For example, Web browsers.) Multiple instances of NetBackup-Java on the same machine have the same effect.

- ◆ Run NetBackup-Java on a 1 gigabyte machine that has at least 256 MB of RAM available to the application. In some instances, the application does not even initiate due to insufficient memory. These failures can be identified by a variety of messages in the xterm window where the `jnbSA` command was executed or the application log file. Possible messages include:

```
Error occurred during initialization of VM
Could not reserve enough space for object heap
Out of Memory
```

## NetBackup-Java Performance Improvement Hints

---

For more information, refer to the conf options “INITIAL\_MEMORY, MAX\_MEMORY” on page 494.

- ◆ Consider the amount of physical memory on the relevant machines, possibly adding memory on the host being administered (the console’s application server host).
- ◆ Consider increasing the swap space to relevant machines:
  - ◆ The console host (the host where the console is started)
  - ◆ The host being administered

Increasing the amount of swap space available to the system where you are running the applications can increase performance, especially if there is a great deal of other activity on the machine. Increasing the amount of swap space can also alleviate hangs or other problems related to insufficient memory for the applications.

- ◆ Consider additional or faster CPUs to relevant machines:
  - ◆ The console host (the host where the console is started)
  - ◆ The host being administered
- ◆ Since startup of the Java virtual machine and some applications can take longer than others, leaving NetBackup-Java running (iconified) rather than exiting and restarting is beneficial.
- ◆ Consider limiting the amount of NetBackup data 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 “Managing the Jobs Database” on page 332.)

## Is Performance Better When Remotely Displaying Back or Running Locally?

Performance depends on the speed of your network, the console and application server machine resources, the workloads on the console and application server hosts, as well as the amount of NetBackup data. (Data being the number of jobs in the Activity Monitor or number of NetBackup policies.) Given these considerations, the console may perform better if started on the console’s application server host and displayed back to the desktop host. However, VERITAS is not aware of a situation where this is true and, as mentioned above, this is *not recommended* due to problems unrelated to performance issues.

Consider the following scenarios when determining what would provide the best performance for your configuration.

### Scenario 1

Assume no deficiency in either the console host's resources or the application server host's resources. Assume that the amount of NetBackup configuration data being transferred to the console host far exceeds the X-Windows pixel display data—that is, the actual console screen being sent from the remote host.

Unfortunately, the only certain method to determine this is to try it. The situation will likely be specific to your NetBackup configuration and certainly be influenced by your network capabilities and proximity of the two hosts involved.

### Scenario 2

Assume that the available resources of the application server host far exceed that of the console host.

For example, if the console host (the machine on which the console is started) has a *very* limited CPU and memory in comparison to the NetBackup master server being administered, you may see better performance by running the console on the master server and displaying back to your desktop host.

If your desktop host is a Windows machine, X-terminal emulation or remote display tools such as Exceed and VNC are required.

These scenarios address the performance aspect of this type of use of the NetBackup-Java console. There may be other reasons that require you to remotely display back to your desktop. However, as mentioned in previous sections, this is *not recommended*. Review the Release Notes for operational notes or known issues and limitations for additional issues of relevance to the NetBackup-Java Administration Console and Backup, Archive, and Restore client console.

## Administrator's Quick Reference

The following tables show information that the NetBackup administrator will frequently use. The man page appendix in this manual provides details on most of the commands displayed in this table.

| Command                        | Description                                                                  |
|--------------------------------|------------------------------------------------------------------------------|
| <b>Administrator Utilities</b> |                                                                              |
| bpadm                          | Starts character-based, menu-driven administrator's interface on the server. |
| jnbSA                          | Starts Java-based, NetBackup administrator's interface on the server.        |
| <b>Client-User Interfaces</b>  |                                                                              |
| bp                             | Starts character-based, menu-driven client-user interface.                   |
| jbpSA                          | Starts Java-based, client-user interface on the client.                      |
| <b>Daemon Control</b>          |                                                                              |
| initbprd                       | Starts bprd (request daemon).                                                |
| bprdreq -terminate             | Stops bprd (request daemon)                                                  |
| initbpdbm                      | Starts bpdbm (database manager).                                             |
| bpadm                          | Has option for starting and stopping bprd.                                   |
| jnbSA (Activity Monitor)       | Has option for starting and stopping bprd.                                   |
| <b>Monitor Processes</b>       |                                                                              |
| bpps                           | Lists active NetBackup processes.                                            |
| jnbSA (Activity Monitor)       | Lists active NetBackup processes.                                            |



| File                                       | Description                                          |
|--------------------------------------------|------------------------------------------------------|
| <code>/usr/opensv/java/auth.conf</code>    | Authorization options.                               |
| <code>/usr/opensv/netbackup/bp.conf</code> | Configuration options (server and client).           |
| <code>/usr/opensv/java/nbj.conf</code>     | Configuration options for the NetBackup-Java Console |
| <code>\$HOME/bp.conf</code>                | Configuration options for user (on client).          |

## Managing Client Restores

The topics in this section concern aspects of managing restores for NetBackup clients.

- ◆ “Server-Directed Restores” on page 504
- ◆ “Client-Redirected Restores” on page 505
- ◆ “Restoring Files and Access Control Lists” on page 512
- ◆ “Setting Client List and Restore Permissions” on page 513
- ◆ “Improving Search Times by Creating an Image List” on page 517
- ◆ “Set Original atime for Files During Restores” on page 518
- ◆ “Checkpoint Restart for Restore Jobs” on page 518
- ◆ “Restoring System State” on page 519

Find related topics in Chapter 5, “Reference Topics,” in the *NetBackup System Administrator’s Guide, Volume II*. Incorrectly specified host names are often a factor in file restore problems.

## Server-Directed Restores

NetBackup clients are configured, by default, to allow NetBackup administrators on a master server to direct restores to any client.

To prevent server-directed restores, configure the client accordingly:

- ◆ Windows clients: Open the Backup, Archive, and Restore interface on the client.  
Select **File > NetBackup Client Properties > General** tab > Clear the **Allow server-directed restores** checkbox.
- ◆ UNIX clients: Add `DISALLOW_SERVER_FILE_WRITES` to the following file on the client:

`/usr/opencv/netbackup/bp.conf`

---

**Note** On UNIX systems, redirected restores can set the UIDs or GIDs incorrectly when the UIDs or GIDs are too long. When restoring files from one platform type to another, it is possible that UIDs and GIDs on one system may be represented with more bits on the source system than on the destination. This means that if the name for the UID/GID in question is not common to both systems, the original UID/GID could be invalid on the destination system. In this case, the UID/GID would be replaced with that of the user doing the restore.

---

## Client-Redirected Restores

The Backup, Archive, and Restore client interface contains options for allowing clients to restore files that were backed up by other clients. The operation is called a *redirected restore*.

### How NetBackup Enforces Restore Restrictions


By default, NetBackup permits only the client that backs up files to restore those files. NetBackup enforces this restriction by ensuring 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 (due to the use of a gateway and token ring combination, or multiple connections), the peer name is equivalent to the client's host name. When a client connects through a gateway, the gateway can use its own peer name to make the connection.

The NetBackup client name is normally the client's short host name, such as `client1` rather than a longer form such as `client1.null.com`.

The client name is found in the following locations:

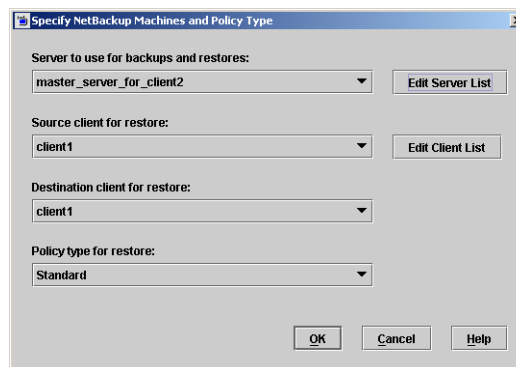
- ◆ Windows clients (including NetWare NonTarget):  
Open Backup, Archive, and Restore and select **File > Specify NetBackup Machines and Policy Type**. The client name selected as **Source Client for Restores** is the source of the backups you want to restore.
- ◆ On NetWare target clients:  
Specify the client name in the `bp.ini` file.



## Managing Client Restores

### ◆ UNIX clients:

Open Backup, Archive, and Restore and select the client name as the **Source client for restore**.



## Allowing All Clients to Perform Redirected Restores

The NetBackup administrator can allow clients to perform *redirected restores*. That is, allow all clients to restore backups belonging to other clients. This is done by placing an empty `No.Restrictions` file on the NetBackup master server where the policy that backed up the other clients resides.

**Note** The information within this section applies to restores made using the command line, not the Backup, Archive, and Restore client interface.

Create an `altnames` directory in the following location, then place the empty file inside of the directory:

```
/usr/opensv/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.

**Caution** The `/usr/opensv/netbackup/db/altnames` directory can present a potential breach of security if users permitted to select and restore files from other clients also have permission to locally create the files found in the backup.

## Allowing a Single Client to Perform Redirected Restores

The NetBackup administrator can permit a single client to restore backups belonging to other clients. This is done by creating an empty file on the NetBackup master server where the policy that backed up the other client(s) resides.

**Note** The information within this section applies to restores made using the command line, not the Backup, Archive, and Restore client interface.

Create an `altnames` directory in the following location, then place the empty file inside of the directory:

```
/usr/opensv/netbackup/db/altnames/peername
```

Where *peername* is the client that will possess restore privileges.

In this case, the requesting client (*peername*) can access files backed up by another client if the NetBackup client name setting on *peername* matches 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 backups belonging to another specific client. This is done by creating a file on the NetBackup master server of the requesting client.

---

**Note** The information within this section applies to restores made using the command line, not the Backup, Archive, and Restore client interface.

---

Create an `altnames` directory in the following location, then place the *peername* file inside of the directory:

```
/usr/opensv/netbackup/db/altnames/peername
```

Where *peername* is the client that will possess restore privileges. Add to the *peername* file the names of the client(s) whose files the requesting client wishes to restore.

The requesting client can restore files backed up by another client if:

- ◆ The names of the other clients appear in the *peername* file, and
- ◆ the NetBackup client name setting on the requesting client is changed to match the name of the client whose files the requesting client wishes to restore.

## Redirected Restore Examples

This section provides NetBackup example configurations that allow clients to restore 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 an image-catalog directory on the master server in the file below, or the requesting client must be a member of an existing NetBackup policy:

```
/usr/opensv/netbackup/db/images/client_name
```

---

**Caution** Not all file system types on all machines support the same features and you may encounter problems when restoring from one file system type to another. For example, the S51K file system on SCO machines does not support symbolic links nor does it support names greater than 14 characters long. If you restore to

## Managing Client Restores

a machine or file system that does not support all the features of the machine or file system from which you performed the restore, you may not be able to recover all the files.

In the following examples:

- ◆ *client1* is the client that is requesting the restore.
- ◆ *client2* is the client that created the backups that the requesting client wants to restore.

---

**Note** The information within this section applies to restores made using the command line, not the Backup, Archive, and Restore client interface.

---

---

**Note** You must be a root user for any of the steps that must be performed on the NetBackup server. You may also have to be a root user to make the changes on the client.

---

### Example 1: Redirected Client Restore

Assume you must restore files to *client1* that were backed up from *client2*. The *client1* and *client2* names are those specified by the NetBackup client name setting on the clients.

In the nominal case, follow these steps to perform the restore:

1. Log in as root on the NetBackup server and perform one of the following actions:
  - ◆ Edit `/usr/opensv/netbackup/db/altnames/client1` so it includes the name of *client2*. Or,
  - ◆ Run the `touch` command on the following file:  
`/usr/opensv/netbackup/db/altnames/No.Restrictions`

---

**Caution** Creating the `No.Restrictions` file allows any client to restore files from *client2*.

---

2. Log in on *client1* and change the NetBackup client name to *client2*.
3. Restore the file.
4. Undo the changes made on the server and client.

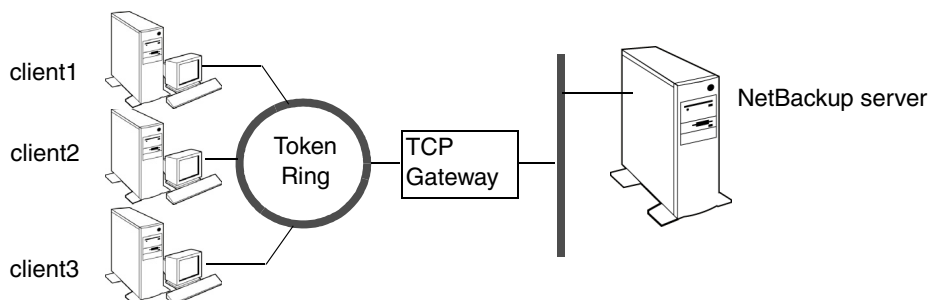
### Example 2: Redirected Client Restore using the altnames File

This example explains how the `altnames` file can provide restore capabilities to clients that do not use their own host name when connecting to the NetBackup server.

By default, the NetBackup client name of the requesting client must match the peer name 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 Ethernets or connect to the NetBackup server through a gateway. Consider the configuration in the following figure:

Example Restore from Token Ring Client



In this example, restore requests coming 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. This means that clients cannot restore even their own files.

Perform the following steps to correct the situation:

1. Determine the peer name of the gateway:
  - a. Attempt 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
```
  - b. Examine the NetBackup problems report and identify the peer name used on the request. Entries in the report will be similar to:
 

```
01/29/05 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`.

2. Run the `touch` command on the following file:

```
/usr/opensv/netbackup/db/altnames/peername
```

## Managing Client Restores

In our example, the file is:

```
/usr/opensv/netbackup/db/altnames/client1.dvlp.null.com
```

3. Edit the *peername* file to include the desired client names.

For example, if you leave the file

```
/usr/opensv/netbackup/db/altnames/client1.dvlp.null.com
```

empty, *client1*, *client2*, and *client3* can all access the backups corresponding to their NetBackup client name setting. (See “Allowing a Single Client to Perform Redirected Restores” on page 506.)

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 507.)

Note that this example requires no changes on the clients.

4. Restore the files.

### Example 3: Troubleshooting Redirected Client Restore using the altnames File

If you cannot restore files with a redirected client restore using the *altnames* file, troubleshoot the situation by performing the following steps:

1. On the NetBackup master server, add the *VERBOSE* entry to the *bp.conf* file.
2. Create the debug log directory for *bprd* by running:

```
mkdir /usr/opensv/netbackup/logs/bprd
```

3. On the NetBackup server, stop the NetBackup request daemon, *bprd*, and restart it in verbose mode by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
/usr/opensv/netbackup/bin/bprd -verbose
```

This ensures that *bprd* logs information regarding client requests.

4. On *client1*, attempt the file restore.
5. On the NetBackup server, identify the peer name connection used by *client1*.

Examine the failure as logged in the All Log Entries report or examine the *bprd* debug log:

```
/usr/opensv/netbackup/logs/bprd/log.date
```

to identify the failing name combination.



6. Perform 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
```

This allows 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
```

This 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. This allows *client1* access to the backups created on *client2* only.

7. On *client1*, change the NetBackup client name setting in the user interface to match what is specified on *client2*.

8. Restore the files to *client1*.

9. Perform 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.

10. To return the configuration to what it was before the restore:

- ◆ Delete */usr/opensv/netbackup/db/altnames/peer.or.hostname* (if existent)
- ◆ Delete */usr/opensv/netbackup/db/altnames/No.Restrictions* (if existent)
- ◆ On *client1*, restore the NetBackup client name setting to its original value.

## Restoring Files and Access Control Lists

An access control list (ACL) is a table that conveys the access rights users have to a file or directory. Each file or directory can have a security attribute which extends or restricts users' access.

### Restoring Files that Possess ACLs

By default, the NetBackup modified GNU `tar` (`/usr/opensv/netbackup/bin/tar`) restores ACLs along with file and directory data. However, there are situations when the ACLs cannot be restored to the file data:

- ◆ Where the restore is cross-platform. (Examples: restoring an AIX ACL to a Solaris client; restoring a Windows ACL to a HP client.)
- ◆ When a `tar` other than the NetBackup modified `tar` is used to restore files.

In these instances, NetBackup stores the ACL information in a series of generated files in the root directory using the following naming form:

`.SeCuRiT.y.nnnn`

These files can be deleted or can be read and the ACLs regenerated by hand.

For a list of other files that NetBackup generates due to cross-platform restores, see Chapter 5, "Reference Topics," in the *NetBackup System Administrator's Guide, Volume II*.

### Restoring Files without Restoring ACLs

The option to restore file and directory data without restoring ACLs is available to NetBackup administrators from the NetBackup client interface if the destination client and the source of the backup are both Windows systems. In order to restore files without restoring ACLs, the following conditions must be met:

- ◆ The policy that backed up the client must have been of policy type *MS-Windows-NT*.
- ◆ The restore must be performed by an administrator logged into a NetBackup server (Windows or UNIX). The option is set from the client interface running on the server. 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 Windows 2000, Windows XP, or Windows Server 2003. The option is disabled on UNIX clients.

### ▼ To restore files without restoring ACLs

1. Log in to the NetBackup server as administrator. Open the Backup, Archive, and Restore client interface.
2. From the client interface, initiate a restore.
3. After selecting the files to be restored, select **Actions > Start Restore of Marked Files**. The Restore Marked Files dialog appears.
4. Place a check in the **Restore without access-control attributes** check box.
5. Make any other selections for the restore job and click **Start Restore**.

## Setting Client List and Restore Permissions

You can specify the list and restore permissions for clients by modifying the `bp.conf` file and (or) the client database. This is explained in the following sections:

- ◆ “Setting the List and Restore Permissions” on page 514
- ◆ “Examples” on page 516

## Adding Clients to the NetBackup Client Database

**Note** The following explains how to add clients when you are using fixed IP addresses. If you are using dynamic addressing (DHCP), see Chapter 3, “Additional Configuration,” in *NetBackup System Administrator’s Guide, Volume II* for instructions on adding clients to the client database.

Before you can set list and restore permissions for a client, you must add the client to the NetBackup client catalog on the master server. The client catalog consists of directories and files in the following directory:

```
/usr/opensv/netbackup/db/client
```

You can create, update, list, and delete client entries with the `bpclient` command. The `bpclient` command is in the directory:

```
/usr/opensv/netbackup/bin/admincmd
```

### ▼ To create an entry in the client catalog

- ❖ To create a client entry, run the following command:

## Managing Client Restores

```
bpclient -add -client client_name -current_host host_name
```

The variables are described as follows:

- ◆ `-client client_name` specifies the NetBackup client name as it appears in the NetBackup configuration.
- ◆ `-current_host host_name` adds the client to the catalog with the name specified by *host\_name*. This host name must already be configured with an IP address in the name service that you are using (for example, DNS). When you run this command, NetBackup queries the name service for the IP address and updates the NetBackup client catalog.

For example:

```
cd /usr/opensv/netbackup/bin/admincmd
bpclient -add -client shark -current_host shark
```

### ▼ To delete and list entries in the client catalog

1. To delete client entries, run `bpclient -delete -client client_name`
2. To list specific client entries, run `bpclient -L -client client_name`
3. To list all client entries, run `bpclient -L -All`

## Setting the List and Restore Permissions

To set the list and restore permissions, use the `bpclient` command to change the `list_restore` settings for the desired clients. The `list_restore` setting is a part of the NetBackup client catalog entry for each client and you can modify it only with the `bpclient` command in the following directory:

```
/usr/opensv/netbackup/bin/admincmd/bpclient
```

The syntax for changing `list_restore` with the `bpclient` command is as follows (one line):

```
bpclient -client client_name -update -current_host host_name
-list_restore [0 | 1 | 2 | 3]
```

Where:

- 0 = List or restore control is not specified (default, see below)
- 1 = Allow both list and restore
- 2 = Allow list only
- 3 = Deny both list and restore

For example, to prevent both lists and restores from the client shark (one line):

```
bpclient -client shark -update -current_host shark
-list_restore 3
```

If you select 0, the standard default action is to allow both lists and restores. However, you can change this by adding `DISALLOW_CLIENT_LIST_RESTORE` and `DISALLOW_CLIENT_RESTORE` options to the `bp.conf` file on the master server.

- ◆ Adding `DISALLOW_CLIENT_LIST_RESTORE` changes the default to deny both lists and restores.
- ◆ Adding `DISALLOW_CLIENT_RESTORE` changes the default to deny restores.

If you add both the `DISALLOW_CLIENT_RESTORE` and `DISALLOW_CLIENT_LIST_RESTORE`, NetBackup behaves as though only `DISALLOW_CLIENT_LIST_RESTORE` is present.

The following table shows the combinations that are possible for setting list and restore permissions. Notice that you can use `list_restore` in combination with the `DISALLOW_CLIENT_RESTORE` and `DISALLOW_CLIENT_LIST_RESTORE` options in the `bp.conf` file. But for any specific client, a `list_restore` setting other than 0 always overrides the `bp.conf` file option.

## Managing Client Restores

| Desired Result |         | Settings                          |                         |                              |
|----------------|---------|-----------------------------------|-------------------------|------------------------------|
| List           | Restore | list_restore value                | DISALLOW_CLIENT_RESTORE | DISALLOW_CLIENT_LIST_RESTORE |
| Yes            | Yes     | 0 (list or restore not specified) | No                      | No                           |
| Yes            | No      | 0 (list or restore not specified) | Yes                     | No                           |
| No             | No      | 0 (list or restore not specified) | No                      | Yes                          |
| No             | No      | 0 (list or restore not specified) | Yes                     | Yes                          |
| Yes            | Yes     | 1 (allow both)                    | No                      | No                           |
| Yes            | Yes     | 1 (allow both)                    | Yes                     | No                           |
| Yes            | Yes     | 1 (allow both)                    | No                      | Yes                          |
| Yes            | Yes     | 1 (allow both)                    | Yes                     | Yes                          |
| Yes            | No      | 2 (allow list only)               | No                      | No                           |
| Yes            | No      | 2 (allow list only)               | Yes                     | No                           |
| Yes            | No      | 2 (allow list only)               | No                      | Yes                          |
| Yes            | No      | 2 (allow list only)               | Yes                     | Yes                          |
| No             | No      | 3 (deny both)                     | No                      | No                           |
| No             | No      | 3 (deny both)                     | Yes                     | No                           |
| No             | No      | 3 (deny both)                     | No                      | Yes                          |
| No             | No      | 3 (deny both)                     | Yes                     | Yes                          |

**Note** In the `DISALLOW_CLIENT_RESTORE` and `DISALLOW_CLIENT_LIST_RESTORE` columns, *Yes* means it is in the `bp.conf` file. *No* means that it is not in the `bp.conf` file.

## Examples

The following examples show several approaches to limiting list and restore privileges for your clients. Each of these examples assume there are three clients: shark, eel, and whale.

### Example 1: Prevent lists and restores on all clients

1. Add `DISALLOW_CLIENT_LIST_RESTORE` to the `bp.conf` file.
2. Leave the `list_restore` setting at 0 (default) for these clients.

**Example 2: Prevent restores but allow lists on all clients except one**

Prevent restores but allow lists on all clients except shark. Prevent both lists and restores on shark.

1. Add `DISALLOW_CLIENT_RESTORE` to the `bp.conf` file.
2. Use `bpclient` to set `list_restore` to 3 for shark. Leave the `list_restore` setting at 0 (default) on the other clients.

**Example 3: Prevent lists and restores for all clients except one**

Prevent lists and restores for all clients except eel. Allow eel to both list and restore files.

1. Add `DISALLOW_CLIENT_LIST_RESTORE` to the `bp.conf` file.
2. Use `bpclient` to set `list_restore` to 1 for eel. Leave the `list_restore` setting at 0 (default) on the other clients.

**Example 4: Allow lists and restores on all clients except one**

Allow lists and restores on all clients except whale. Allow users on whale to list but not restore files.

1. Remove `DISALLOW_CLIENT_LIST_RESTORE` and `DISALLOW_CLIENT_RESTORE` from the `bp.conf` file. (if they exist).
2. Use `bpclient` to set `list_restore` to 2 for whale. Leave the `list_restore` setting at 0 (default) on the other clients.

## Improving Search Times by Creating an Image List

To improve search performance when you have many small backup images, run the following command (one line) as root on the master server:

```
/usr/opensv/netbackup/bin/admincmd/bpimage -create_image_list
-client name
```

The *name* is the name of the client that has many small backup images.

This creates the following files in the

`/usr/opensv/netbackup/db/images/clientname` directory:

IMAGE\_LIST: List of images for this client

IMAGE\_INFO: Information about the images for this client

## Managing Client Restores

---

**IMAGE\_FILES:** The file information for small images

Do not edit these files because they contain offsets and byte counts that are used for seeking to and reading the image information.

These files take 35 to 40% more space in the client directory and if you use them, verify that there is adequate space. Also, they improve search performance only when there are thousands of small backup images for a client.

## Set Original atime for Files During Restores

During a restore NetBackup by default sets the `atime` for each file to the current time. If you want NetBackup to set the `atime` for each restored file to the value it had when it was backed up, create the following special file on the client.

```
/usr/opensv/netbackup/RESTORE_ORIGINAL_ETIME
```

---

**Note** If you are using VERITAS Storage Migrator, do not create the `RESTORE_ORIGINAL_ETIME` file. If you do, it is possible that restored files will be immediately migrated because of their older `atime`.

---

## Checkpoint Restart for Restore Jobs

Checkpoint Restart for restore jobs saves time by providing the mechanism for NetBackup to automatically resume a failed restore job from the start of the file last checkpointed rather than from the beginning of the entire restore job. The checkpoints are taken once every minute during a restore job.

Checkpoint Restart for restore jobs is enabled by default, requiring no additional configuration. However, there are two host properties that impact Checkpoint Restart for restore jobs:

- ◆ Master server host property **Clean-up > Move Restore Job from Incomplete State to Done State** (see “Clean-up Properties” on page 359).
- ◆ Master server host property **Universal > Restore Retries** (see “Universal Settings Properties” on page 447).

## Suspending and Resuming a Restore Job

A NetBackup administrator can choose to suspend a checkpointed restore job and resume the job at a later time.



For example, while running a restore job for several hours, the administrator may receive a request for a second restore of a higher priority that requires the resources being used by the first job. The administrator can suspend the first job, start the second restore job and let it complete. Then, resume the first job from the Activity Monitor and let the job complete.

**Note** If a checkpointed restore that has no end date is suspended, then resumed, and a new backup occurred prior to initiating the resume, the files from that new backup will be included in the restore.  
For example, a user makes a restore request of a directory, then that restore is suspended. The request is resumed the next day, after another backup of the directory has been performed. The files that are restored are from the latest backup.

For more on suspending restore jobs and resuming incomplete jobs, see “Activity Monitor Menu Bar” on page 311.

### Limitations to Checkpoint Restart for Restore Jobs

Limitations to Checkpoint Restart for restore jobs include the following:

- ◆ The restore restarts at the beginning of the last checkpointed file only, not within the file.
- ◆ Checkpoint Restart for restore jobs works only on files backed up using Standard or MS-Windows-NT policy types.

**Note** Although NetWare clients use the Standard policy type, Checkpoint Restart for restores is not supported on NetWare clients.

- ◆ Third Party Copy and Media Server Copy images that use Standard policy types are supported, but cannot be suspended or resumed if the backup image has changed blocks. Flashbackup is not supported.

## Restoring System State

On all hosts running Windows 2000 or later, the System State includes the registry, the COM+ Class Registration database, and boot and system files. For Windows 2000 servers, the Certificate Services database is included if the server is operating as a certificate server. If the server is a domain controller, the data also includes the Active Directory services database and the SYSVOL directory.

## Managing Client Restores

---

**Note** If you are restoring a Windows server from a complete system failure, the best recovery procedure depends on many hardware and software variables pertaining to your server and its environment. A complete Windows recovery procedure is beyond the scope of this manual; you may need to contact Microsoft or refer to your Microsoft documentation.

---

### Important Notes on System State

Before restoring the System State, please read the following notes carefully.

- ◆ The System State should be restored in its entirety: restoring selected files is not recommended.
- ◆ Although incremental backups of the System State can be configured, NetBackup always performs a full backup. Therefore, only the most recent backup of the System State must be restored.
- ◆ For Windows 2000 systems, Service Pack 2 is required.
- ◆ Do not redirect a System State restore. System State is computer-specific and restoring it to an alternate computer can result in an unusable system.
- ◆ Do not cancel a System State restore operation. Canceling this operation could leave the system unusable.
- ◆ When restoring the System State to a domain controller, the Active Directory must not be running. Refer to the following procedure for directions on restoring the Active Directory.


#### ▼ To restore the System State

1. If you want to restore the Active Directory, or if the system to which you are restoring is a Windows domain controller, restart the system and press F8 during the boot process. Otherwise, begin with step 4 below.

F8 brings up a startup options menu.

2. From the startup options, select **Directory Services Restore Mode** and continue the boot process.
3. Make sure the **NetBackup Client Service** has started. (Select **Control Panel > Administrative Tools > Services** to check.)

4. Start the Backup, Archive, and Restore client interface. Click **Select for Restore**, and place a checkmark next to **System State** (as shown in the Windows Backup, Archive, and Restore console to the right).
5. From the **Actions** menu, choose **Start Restore of Marked Files**.
6. From the **Restore Marked Files** dialog, select **Restore everything to its original location** and **Overwrite the existing file**.



**Caution** Do not redirect the System State restore to a different host. System State is computer-specific: restoring it to a different computer can result in an unusable system.

7. Click **Start Restore**.
8. If you have more than one domain controller in the network and you want Active Directory replicated to the other domain controllers, you must 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 have 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.

For more information about authoritative restore and the `ntdsutil` utility, please refer to your Microsoft documentation.

9. Reboot your system before performing subsequent restore operations.

If this is a domain controller and you have booted into **Directory Services Restore Mode**, reboot into normal mode when the restore is complete.

## Goodies Scripts

The `/usr/opensv/netbackup/bin/goodies` directory contains sample shell scripts that you can modify. You can use some of them in conjunction with the `cron` utility to create periodic mailings of information relating to NetBackup. They can also serve as examples of how to use NetBackup commands in scripts. If you use the example scripts, ensure that they are executable by *other*. Do this by running `chmod 755 script_name`, where `script_name` is the name of the script.

---

**Note** The scripts in the `goodies` directory are not officially supported but are intended as examples that you can customize according to your needs.

---

## Server Independent Restores

This section explains how to restore files by using a NetBackup server other than the one that was used to write the backup. This is called a *server independent restore* and allows easier access to data for restores in master and media server clusters and provides better failover and disaster recovery capabilities.


NetBackup has a master and media server architecture that allows storage devices to be located on multiple servers (either separate storage devices or a shared robot). For successfully completed backups, the NetBackup image catalog stored on the master server contains an entry that defines the server (master or media server) to which each backup was written. In addition, information specific to the backup media is held within both the master server image catalog (in the attribute file for each backup) and in the master server Media Manager database.

Because NetBackup tracks the server that was used to create a backup, restoring data through a device on another server is more involved than other restores but can be accomplished by using the methods described in this section. These methods do not require you to expire and import backup images; although, that can be useful in some instances. (See “Notes on Server Independent Restores” on page 528.)

## Supported Configurations

The next two figures show configurations where NetBackup supports server independent restores. All of these methods require that the server used for the restore be in the same cluster as the server that did the original backup and also share the same Enterprise Media Manager database.

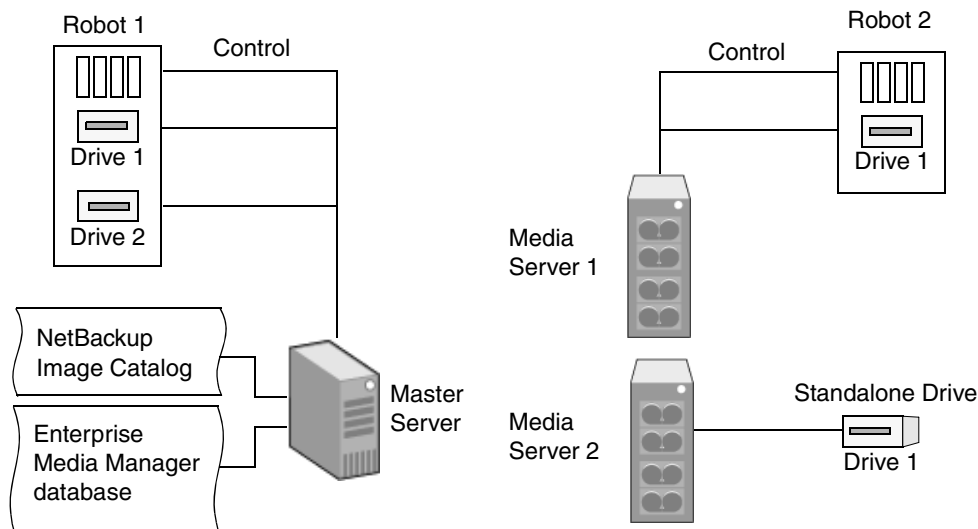
## NetBackup Servers Sharing Robotic Peripherals



In the figure above, the following assumptions are made:

- ◆ 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.

## NetBackup Servers with Separate Non-shared Peripherals



**Note:** Media servers 1 and/or 2 may be offsite.

In the figure above, the following assumptions are made:

- ◆ The media is made physically accessible through an available NetBackup server and the Enterprise Media Manager database is updated to reflect this move.

## Server Independent Restores

---

- ◆ 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.

## Methods for Performing Server Independent Restores

The method that NetBackup administrators can use to perform server independent restores depends on the configuration and situation, and can include one or more of the following:

- ◆ “Method 1: Modifying the NetBackup Catalogs” on page 524
- ◆ “Method 2: Overriding the Original Server” on page 526
- ◆ “Method 3: Automatic Failover to Alternate Server” on page 527

### Method 1: Modifying the NetBackup Catalogs

This method changes the contents of NetBackup catalogs and thus requires administrator intervention. Use this method only when the server reassignment is permanent. Some examples of when to use this method:

- ◆ Media is moved to an offsite location, where a media server exists.
- ◆ A robot has been moved from one server to another.
- ◆ Two (or more) servers are sharing a robot, each has connected drives. One of the servers will soon 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 plenty of empty slots exist on another server’s robot.

The actual steps used in the process vary depending on whether the original server is still available.

▼ **To modify catalogs when the server that wrote the media is available**

1. If necessary, physically move the media. Then, update the Enterprise Media Manager database by using move volume options in the Media Manager administration utilities.
2. Update the NetBackup image catalog on the master server and 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. The `admincmd` command above must be entered on one line.

- ◆ As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpmedia -movedb -m media_id -newserver hostname
-oldserver hostname
```

- ◆ As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -m media_id
-newserver hostname -oldserver hostname
```

▼ **To modify catalogs when the server that wrote the media is not available**

1. If necessary, physically move the media and update the Enterprise Media Manager database by using the move volume options in the Media and Device Management window.
2. Update only the NetBackup image catalog on the master server. Use the following commands from the NetBackup master server. The `admincmd` command above must be entered on one line.

- ◆ As root on a UNIX NetBackup server:

```
cd /usr/opensv/netbackup/bin/admincmd
bpimage -id media_id -newserver hostname
-oldserver hostname
```

- ◆ As administrator on a Windows NetBackup server:

```
cd install_path\NetBackup\bin\admincmd
bpimage.exe -id media_id -newserver hostname
```

## Server Independent Restores

`-oldserver hostname`

### Method 2: Overriding the Original Server

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:

- ◆ Two (or more) servers are sharing 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 server has been removed from the NetBackup configuration, and is no longer available.

#### ▼ To enable overriding of the original server for restores

1. In the NetBackup Administration console, open the **General Server** host properties dialog. (See “General Server Properties” on page 411.)
2. Add an entry in the **Media Host Override** list, naming the original backup server and the restore server. Click **OK**.

#### ▼ To manually override the original server for restores

To revert to the original configuration for future restores, delete the changes made in step 2.

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:
  - ◆ Using the NetBackup Administration Console:
 

Open the **General Server** host properties dialog of the master server.

Add an entry in the **Media Host Override** list, naming the original backup server and the restore server. Click **OK**.
  - ◆ By modifying the `bp.conf` file on a UNIX NetBackup server:
 

As `root` add the following entry to the `/usr/opensv/netbackup/bp.conf` file:

```
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.

3. Stop and restart the NetBackup Request daemon on the master server.

---

**Note** The override applies to all storage units on the original server. This means restores for any storage unit on *fromhost* will go to *tohost*.

---

### Method 3: Automatic Failover to Alternate Server

NetBackup allows the administrator to configure automatic restore failover to an alternate server, if the original server is temporarily inaccessible. Once configured, this method does not require administrator intervention. (See “Restore Failover Properties” on page 433.)

Some examples of when to use this method are:

- ◆ Two or more servers are sharing a robot and each has 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:

- ◆ Original server is down.
- ◆ Original server is up but `bpcd` on that server is not responding (for example, if the connection is refused or access is denied).
- ◆ Original server is up and `bpcd` is fine, but `bptm` is having problems (for example, if `bptm` cannot find the required tape).

---

**Note** The failover uses only failover hosts that are listed in the NetBackup configuration (see the following procedure). By default, no servers are listed so NetBackup does not perform the automatic failover.

---

#### ▼ To enable automatic failover to an alternate server

1. Modify the NetBackup configuration on the master server:
  - ◆ Using the NetBackup Administration Console:  
Open the **Restore Failover** host properties dialog of the master server.

## Server Independent Restores

Add an entry in the **Alternate Restore Failover Machines** list, naming the media server and failover restore server(s).

- ◆ By modifying the `bp.conf` file on a UNIX NetBackup server:

As `root`, add the following entry to the `/usr/opensv/netbackup/bp.conf` file:

```
FAILOVER_RESTORE_MEDIA_SERVERS =
 failed_host host1 host2 ... hostN
```

where:

`failed_host` is the server that is not operational.

`host1 ... hostN` are the servers that provide failover capabilities.

When automatic failover is necessary for a given server, NetBackup searches through the relevant `FAILOVER_RESTORE_MEDIA_SERVERS` list from left to right to determine the first server eligible to perform the restore.

---

**Note** 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.

## Notes on Server Independent Restores

### Expiring and importing media

Even with the above server independent restore capabilities, there are still instances when it is necessary to expire media and then import it.

### Identifying *media spanning groups*

A server independent restore operation can involve media IDs with backup images that span media. For any of these media IDs, it can be necessary to identify the rest of the media IDs that contain fragments of the same spanned images. The group of related media, in this instance, is called a *media spanning group*.

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
```

---

## Server Independent Restores

To display all media in all spanning groups, omit `-mediaid media_id` from the command.

## Configuring NetBackup Ports

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

- ◆ Registered ports are registered with the Internet Assigned Numbers Authority (IANA) and are permanently assigned to specific NetBackup services. For example, the port for the NetBackup client daemon (`bpcd`) is 13782. These ports are specified in a system configuration file:

```
/etc/services
```

Media Manager services include tape library control daemons, which accept connections from daemons on other servers that are sharing the same library. See the `services` file on the media server to determine the ports required for a specific library.

- ◆ In NetBackup 6.0 installations, it is not likely that dynamically-allocated ports will be as much of a concern as in previous releases.

Dynamically-allocated ports are assigned, as needed, from configurable ranges on NetBackup clients and servers. In addition to the range of numbers, you can configure the following for dynamically-allocated ports:

- ◆ Whether NetBackup selects a port number at random from the allowed range or starts at the top of the range and uses the first one available.
- ◆ Whether connections to `bpcd` on a client use reserved or non-reserved ports.

## Port Descriptions

The following five daemons figure most prominently in NetBackup 6.0 installations concerning firewalls. If additional port configuration is necessary due to firewalls, these ports would most likely be affected:

- ◆ `vnetd` (port 13724)  
VERITAS Network Daemon allows all socket communication to take place while connecting to a single port. Legacy NetBackup services introduced before NetBackup 6.0 use the `vnetd` port number.
- ◆ `veritas_pbx` (port 1556)  
VERITAS Private Branch Exchange allows all socket communication to take place while connecting through a single port. NetBackup services introduced in NetBackup 6.0 use the `veritas_pbx` port number.

*If using NetBackup Access Control (NBAC):*

- ◆ `vrts-auth-port` (port 4032)

The VERITAS Authorization Service is one of the VERITAS Security Services (VxSS). It verifies that an identity has permission to perform a specific task.

- ◆ `vrts-at-port` (port 2821)

The VERITAS Authentication Service is one of the VERITAS Security Services (VxSS). It validates, identifies, and forms the basis for authorization and access.

*If using an NDMP server:*

- ◆ `ndmp` (port 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.

The following table lists all ports used by NetBackup and Media Manager:

NetBackup and Media Manager Ports

| Daemon                   | Port  | Description                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|-------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>acsd</code>        | 13702 | The Automated Cartridge System (ACS) daemon is a robotic daemons.                                                                                                                                                                                                                                                                                                                                                                       |
| <code>bpcd</code>        | 13782 | The NetBackup Client daemon.<br>On UNIX clients, <code>bpcd</code> can only be run in standalone mode.<br>On Windows, <code>bpcd</code> always runs under the supervision of <code>bpinstd.exe</code> . There is a NetBackup-specific configuration parameter for <code>bpcd</code> . If the port number is changed within the NetBackup configuration, the software causes the port number in the services file to be updated as well. |
| <code>bpdbm</code>       | 13721 | The NetBackup database manager daemon.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <code>bpjava-msvc</code> | 13722 | The NetBackup-Java application server authentication service program.                                                                                                                                                                                                                                                                                                                                                                   |
| <code>bpjobd</code>      | 13723 | The NetBackup Jobs Database Management daemon.                                                                                                                                                                                                                                                                                                                                                                                          |
| <code>bprd</code>        | 13720 | The NetBackup Request daemon.<br>On Windows, there is a NetBackup specific configuration parameter for <code>bprd</code> . If the port number is changed within the NetBackup configuration, the software causes the port number in the services file to be updated as well.                                                                                                                                                            |
| <code>migrd</code>       | 13699 | The VSM request daemon (database request management) for Storage Migrator. <code>migrd</code> handles communication for VSM-Java and commands.                                                                                                                                                                                                                                                                                          |

## Configuring NetBackup Ports

---

|                             |       |                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------------|-------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ndmp</code>           | 10000 | (See description above table.)                                                                                                                                                                                                                                                                                                                                    |
| <code>odld</code>           | 13706 | The Optical Disk Library (ODL) daemon is a robotic daemon.                                                                                                                                                                                                                                                                                                        |
| <code>tl4d</code>           | 13713 | The Tape Library 4MM (TL4) daemon is a robotic daemon.                                                                                                                                                                                                                                                                                                            |
| <code>tl8cd</code>          | 13705 | The Tape Library 8MM (TL8) control daemon is a robotic daemon.                                                                                                                                                                                                                                                                                                    |
| <code>tlbcd</code>          | 13711 | The Tape Library DLT (TLD) control daemon is a robotic daemon.                                                                                                                                                                                                                                                                                                    |
| <code>tlhcd</code>          | 13717 | The Tape Library Half-inch (TLH) control daemon is a robotic daemon.                                                                                                                                                                                                                                                                                              |
| <code>tlmd</code>           | 13716 | The Tape Library Multimedia (TLM) daemon is a robotic daemons.                                                                                                                                                                                                                                                                                                    |
| <code>tshd</code>           | 13715 | The Tape Stacker Half-inch (TSH) daemon is a robotic daemon.                                                                                                                                                                                                                                                                                                      |
| <code>visd</code>           | 9284  | <p>The VERITAS Information Server Daemon. <code>visd</code> requires <code>nbdabd</code> to be running before it will start.</p> <p>Note: Do not use port number 65535 for either <code>visd</code> or the Dashboards, as this port number causes problems for both <code>visd</code> and the Dashboards.</p>                                                     |
| <code>vmd</code>            | 13701 | <p>The Media Manager volume daemon. <code>vmd</code> logs an error message using <code>syslogd</code> on UNIX or the Event Viewer on Windows, if the port that it binds to is in use. If this occurs, it may be necessary to override the services file.</p>                                                                                                      |
| <code>vnetd</code>          | 13724 | (See description above table.)                                                                                                                                                                                                                                                                                                                                    |
| <code>vopied</code>         | 13783 | <p>The daemon that provides VERITAS One-time Password user authentication. <code>vopied</code> is used to authenticate user names, hosts names, and group/domain names.</p> <p>On UNIX clients, <code>vopied</code> can only be run in standalone mode.</p> <p>On Windows, <code>vopied</code> always runs under the supervision of <code>bpinetd.exe</code>.</p> |
| <code>vrts-auth-port</code> | 4032  | (See description above table.)                                                                                                                                                                                                                                                                                                                                    |
| <code>vrts-at-port</code>   | 2821  | (See description above table.)                                                                                                                                                                                                                                                                                                                                    |
| <code>veritas_pbx</code>    | 1556  | (See description above table.)                                                                                                                                                                                                                                                                                                                                    |

## Load Balancing

NetBackup provides ways to balance loads between servers, clients, policies, and devices. These features are explained in the following topics. When making changes, remember that these settings are interactive, and compensating for one problem can cause another. The best approach to configuring these attributes is to use the defaults unless you anticipate or encounter a problem.

### Adjust Backup Load on Server

Change the **Limit Jobs Per Policy** attribute for one or more of the policies that the server is backing up. For example, decreasing **Limit Jobs Per Policy** reduces the load on a server on a specific network segment. Reconfiguring policies or schedules to use storage units on other servers also reduces the load. Another possibility is to use NetBackup's bandwidth limiting on one or more clients.

### Adjust Backup Load on Server Only During Specific Time Periods

Reconfigure schedules that run during those time periods, so they use storage units on servers that can handle the load (assuming you are using media servers).

### Adjust Backup Load on Client

Change the **Maximum Jobs Per Client** global attribute. For example, increasing **Maximum Jobs Per Client** increases the number of concurrent jobs that any one client can process and therefore increases the load.

### Reduce Time To Back Up Clients

Increase the number of jobs that clients can perform concurrently, or use multiplexing. Another possibility is to increase the number of jobs that the server can perform concurrently for the policy or policies that are backing up the clients.

### Give Preference To a Policy

Increase the **Limit Jobs Per Policy** attribute for the preferred policy relative to other policies. Or, increase the priority for the policy.

### Adjust Load Between Fast and Slow Networks

Increase the **Limit Jobs Per Policy** and **Maximum Jobs Per Client** for policies and clients in a faster network and decrease these numbers for slower networks. Another solution is to use NetBackup's bandwidth limiting.

## Load Balancing

---

### **Limit the Backup Load Produced By One or More Clients**

Use NetBackup's bandwidth limiting to reduce the bandwidth used by the clients.

### **Maximize Use of Devices**

Use multiplexing. Also, 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 Devices**

Limit the number of devices that NetBackup can use concurrently for each policy or the number of drives per storage unit. Another approach is to not put some devices under Media Manager control.

You can also place some drives in a down state or limit the number used concurrently in a specific storage unit. For example, if there are four drives in a robot, allow only two to be used concurrently.



## Using NetBackup with Storage Migrator

VERITAS recommends Storage Migrator for UNIX (VSM) as a storage migration solution. NetBackup can back up files from a disk type storage unit that has file systems managed by VSM.

When a file is migrated, the data is copied to secondary storage. The data can then be deleted (or purged) from the disk storage unit because it is copied elsewhere. The files can then be recalled (or cached) from secondary storage should the files be needed locally.

NetBackup backs up files that have been purged by Storage Migrator in the following order:

- ◆ For user backups by nonroot users, NetBackup first caches the files, then backs up the files.
- ◆ For scheduled backups and user backups by a root user, NetBackup backs up only the migration information for the files. Because the data is already resident on secondary storage, NetBackup neither backs up the data nor caches it.

**Caution** Because NetBackup does not set the Storage Migrator obsolescence date for a file, you must ensure that your migrated copies are retained at least as long as your backups. Restores will not be possible unless you ensure the copies are retained.

When NetBackup restores files that have been purged, Storage Migrator considers the restored files to be purged, with a file slice value of zero. If files have been selected for migration and not yet copied to secondary storage, NetBackup backs them up.

A `bparchive` back up and remove operation always caches a purge file.

## Set a Large Enough Media Mount Timeout

When NetBackup restores files to a disk storage unit managed by Storage Migrator, the **Media Mount Timeout** host property is in effect during the caching of the (potentially) migrated backups. This property is located under **Host Properties > Master Server > Timeouts > Media Mount Timeout**.

If the file being restored is part of a large backup that was migrated to tape, the **Media Mount Timeout** must provide enough time to cache in the entire disk file.

## Do Not Use the RESTORE\_ORIGINAL\_ETIME File

Do not create the `/usr/opensv/netbackup/RESTORE_ORIGINAL_ETIME` on any clients that are running Storage Migrator or restored files may be immediately migrated because of the older `etime`. (Also see “Set Original `etime` for Files During Restores” on page 518.)

## Using NetBackup with Storage Migrator

---

**Note** If using another migration product, ensure that it provides adequate and full recoverability of the disk-resident data and fully transparent access to these disk files at the application level.

---

### Do Not Use the Following Client `bp.conf` File Settings

Ensure that the `bp.conf` file on a client using Storage Migrator does not have entries for either of the following:

- ◆ `DO_NOT_RESET_FILE_ACCESS_TIME`
- ◆ `USE_CTIME_FOR_INCREMENTALS`

These entries cause the `atime` for files to be updated each time the files are backed up. This makes it appear as if the files are frequently used and stops Storage Migrator from selecting the files for migration.

## NetBackup Relational Database

## A

This appendix contains information concerning the proper installation and operation of the Sybase Adaptive Server Anywhere (ASA) relational database management system. Generally, the implementation of Sybase ASA in the NetBackup catalog is transparent. However, this appendix is for administrators concerned with the following aspects of the RDBMS used in NetBackup:

- ◆ “Installation Overview” on page 538
- ◆ “Post-installation Tasks” on page 548
- ◆ “Backup and Recovery Procedures” on page 553
- ◆ “Database Unloading Tool” on page 558
- ◆ “Moving the NetBackup Database from One Host to Another” on page 560

NetBackup installs Sybase ASA 9.0.1 during the master server installation, as a private, non-shared server for the NetBackup database (NBDB). The NetBackup database, NBDB, contains the Enterprise Media Manager (EMM) data as well as other NetBackup data used by NetBackup services.

The same installation of Sybase ASA 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.

## Installation Overview

By default, the NetBackup relational database (NBDB) is installed on the master server, which is also the default location for the Enterprise Media Manager (EMM) server. Since the primary usage of NBDB is by EMM, the NetBackup database always resides on the same machine as the Enterprise Media Manager.

For performance reasons, the EMM server and the relational database can be moved to another server.

The following steps are performed automatically during installation, but can be performed independently after installation.

1. As part of the NetBackup master server installation, the Sybase ASA 9.0.1 server is created. The server parameters are set in the `server.conf` file:

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

The contents of the `server.conf` file are described in “server.conf” on page 540.

2. The following entry sets the database location. It is added to the `bp.conf` file:

```
VxDBMS_NB_DATA = /usr/opensv/db/data
```

For more about this entry, see “NetBackup Configuration Entry” on page 545.

3. The VxDBMS configuration file for NetBackup is created. This file requires the read/write permissions of root:

```
/usr/opensv/db/data/vxdbms.conf
```

4. The NetBackup database is created:

```
/usr/opensv/db/data/NBDB.db
```

5. DBA password is set for the NetBackup database in `vxdbms.conf`:

```
VxDBMS_NB_PASSWORD = encrypted_password
```

6. A minimum of four additional database files are created with contiguous space pre-allocated:

The NetBackup system database file (mentioned in step 4):

```
/usr/opensv/db/data/NBDB.db
```

The EMM database files:

```
/usr/opensv/db/data/EMM_DATA.db
```

```
/usr/opensv/db/data/EMM_INDEX.db
```

The NetBackup transaction log, necessary for recovering the database:

```
/usr/opensv/db/data/NBDB.log
```

7. The Sybase ASA accounts and schema for each of the NetBackup components (for example, EMM\_MAIN) making use of the NetBackup database are created.
8. The EMM data is initialized by running:

```
/usr/opensv/volmgr/bin/tpext
```

## NetBackup Master Server Installation

Sybase ASA is installed in the following directories:

- ◆ /usr/opensv/var/global

The files located in /global can be shared within a cluster.

- ◆ /usr/opensv/db

The contents of each directory are examined in the following sections.

### 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. NBDB must always be on the same host as the EMM server. Also, 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 549.)

---

**Note** If Bare Metal Restore is installed, BMRDB must be located on the master server.

---

### server.conf

---

**Caution** VERITAS strongly recommends that this file *not be edited* without assistance from Technical Support. Editing this file may result in NetBackup not starting.

---

/usr/opensv/var/global/server.conf is read when the ASA daemon is started. The ASA daemon gets all configuration information from this file:

```
-n VERITAS_NB_server_name
-x tcpip(LocalOnly=YES;ServerPort=13785) -gp 4096 -ct+ -gd DBA -gk DBA -gl
DBA -ti 0 -c 25M -ch 500M -cl 25M -gn 10 -o /usr/opensv/db//log/server.log -ud
-n VERITAS_NB_server_name
```

Where *server\_name* indicates the name of the Sybase ASA 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.

---

**Caution** If this name is changed, the Enterprise Media Manager will be unable to connect to the database.

---

```
-x tcpip(LocalOnly=YES;ServerPort=13785)
```

Indicates what kind of connections are allowed in addition to shared memory: local TCP/IP connections using port 13785.

`-gp 4096`

Indicates maximum page size (in bytes) for the database. This parameter is given during database creation.

`-ct+`

Indicates that character set translation is used. UTF8 encoding is used.

`-gd DBA -gk DBA -gl DBA`

Indicates that the DBA user is the account used to start, stop, load, and unload data.

`-ti 0`

Indicates the client idle time allowed before shutdown. By default, no idle time is allowed, preventing the database from shutting down.

`-c 25M`

Indicates the initial memory reserved for caching database pages and other server information. (May be changed for performance reasons.)

`-ch 500M`

Indicates the maximum cache size, as a limit to automatic cache growth (May be changed for performance reasons.)

`-cl 25M`

Indicates minimum cache size, as a limit to automatic cache resizing. (May be changed for performance reasons.)

`-gn 10`

Indicates the number of requests the database server can handle at one time. This parameter limits the number of threads upon startup. (May be changed for performance reasons.)

`-o /usr/opensv/db/log/server.log`

Indicates location of server output messages (includes start/stop events, checkpoints, error conditions, cache changing size). This log is not managed, but growth is slow.

`-ud`

Indicates that the server should run as a daemon.

## **databases.conf**

The `/usr/opensv/var/global/databases.conf` configuration file contains the locations of the main database files and the database names for automatic startup when the ASA daemon is started. For example, if NBDB and BMRDB are both located on the master server in the default locations, `databases.conf` contains:

## Installation Overview

---

```
"/usr/opensv/db/data/NBDB.db" -n NBDB
"/usr/opensv/db/data/BMRDB.db" -n BMRDB
```

### **vxdbms\_env.csh, vxdbms\_env.sh**

These scripts set up the ASA environment and are used by other scripts and commands:

- ◆ /usr/opensv/db/vxdbms\_env.csh
- ◆ /usr/opensv/db/vxdbms\_env.sh

### **/bin**

/usr/opensv/db/bin contains all ASA commands and NetBackup-specific commands:

- ◆ 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/stop individual databases.
- ◆ nbdb\_backup  
Use to make an online or an offline backup of the ASA database files to a file system directory.
- ◆ nbdb\_move  
Use to change the location of the ASA database files from the default location.
- ◆ nbdb\_ping  
Displays the status of the ASA database.
- ◆ nbdb\_restore  
Use to recover from an online or an offline 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 or BMRDB database schema and data.
- ◆ nbdbms\_start\_server  
Use to start and stop the ASA daemon.
- ◆ nbdb\_upgrade  
Used internally to upgrade the NetBackup and BMR databases.



---

**Note** The commands listed above are described fully in *NetBackup Command for UNIX* and in the online help.

---

## /charsets

/usr/opensv/db/charsets contains ASA-specific information.

## /data

/usr/opensv/db/data is the default location of the database, NBDB:

- ◆ NBDB.db  
Main NetBackup database file; considered a *dbspace*.
- ◆ EMM\_DATA.db  
An additional *dbspace* that contains EMM data.
- ◆ EMM\_INDEX.db  
Enhances EMM database performance.
- ◆ NBDB.log  
The transaction log for the NetBackup database, necessary for recovery.  
NBDB.log is automatically truncated after a successful full or incremental online, hot or offline, cold catalog backup of the ASA database.
- ◆ vxdbms.conf

Contains configuration information specific to the Sybase ASA installation:

```
VXDBMS_NB_SERVER = VERITAS_NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
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_PASSWORD = encrypted_password
```

The encrypted password used to log into both the DBA accounts for NBDB and BMRDB, and other data accounts is stored in vxdbms.conf.

---

**Note** The password is set to a default upon installation (*nbusql*). VERITAS recommends changing the password after installation. (See “Changing the Database Password” on page 548.)

---

## Installation Overview

---

- ◆ If BMR is installed, the directory will also contain: `BMRDB.db`, `BMRDB.log` (transaction log for BMR), `BMR_DATA.db`, `BMR_INDEX.db`

### **/lib**

`/usr/opensv/db/lib` contains all the ASA shared libraries, including ODBC libraries, used to connect to NBDB and BMRDB.

### **/log**

`/usr/opensv/db/log` contains the ASA server log file `server.log` which contains Sybase logging only

### **/res**

`/usr/opensv/db/res` contains ASA-specific information.

### **/scripts**

`/usr/opensv/db/scripts` contains ASA SQL scripts used in creating the database, and NetBackup SQL scripts used to create the EMM and other schemas.

---

**Caution** The scripts located in `/usr/opensv/db/scripts` should not be edited.

---

### **/staging**

`/usr/opensv/db/staging` is used as a temporary staging area during online, hot catalog backup and recovery.

### **/tix**

`/usr/opensv/db/tix` contains ASA-specific information.

## NetBackup Configuration Entry

The `bp.conf` entry, `VXDBMS_NB_DATA`, is a required entry and is created upon installation. The entry indicates the pathname to the directory where `NBDB.db`, `BMRDB.db`, and the `vxdbms.conf` file are located.

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

```
VXDBMS_NB_DATA = /usr/opensv/db/data
```

## Sybase ASA Server Management

Upon startup, the Sybase ASA 9.0.1 server uses the ASA (Adaptive Server Anywhere) daemon to set the server parameters in the `server.conf` file. Then, the ASA daemon starts the databases indicated in the `databases.conf` file.

### ▼ To start and stop the ASA daemon

Use one of the following methods.

- ◆ Select NB\_dbsrv in the Activity Monitor in the NetBackup Administration Console.
- ◆ From the command line:
  - ◆ `/usr/opensv/db/bin/goodies/netbackup stop | start`  
The ASA daemon is included in the `stop` or `start` command, which starts and stops all NetBackup daemons.
  - ◆ `/usr/opensv/db/bin/nbdbms_start_server`  
`nbdbms_start_server` without any option indicated, starts the ASA server.
  - ◆ `/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:  
Adaptive Server Anywhere Server Ping Utility Version  
9.0.1.1965  
Ping server successful.
  - ◆ `/usr/opensv/db/bin/nbdbms_start_server -h`  
Use `-h` to display usage information about the `nbdbms_start_server`.

### ▼ To start/stop individual databases

The individual databases can be started or stopped, while leaving the ASA daemon to continue running:

- ◆ `nbdb_admin [-start | -stop]`  
Starts or stops NBDB without shutting down the ASA server.  
To see whether the database is up, enter:  
`nbdb_ping`
- ◆ `nbdb_admin [-start | -stop BMRDB]`  
Starts or stops BMRDB without shutting down the ASA server.

To see whether the BMRDB database is up, enter:

```
nbdb_ping -dbn BMRDB
```

## Clustering

Sybase ASA is supported in a clustered environment. Sybase ASA failover is included with the NetBackup server failover solution. The software is installed on all machines in the cluster, but, the database files are created on a shared disk.

To facilitate this, database and configuration files are installed on a shared drive:

Configuration files are stored in `/usr/opensv/var/global`.

## Post-installation Tasks

All tasks are optional and performed at the command line:

- ◆ Change the database password (described below)
- ◆ Move NBDB and BMRDB database files (possibly for performance tuning) (see “Moving NBDB Database Files After Installation” on page 549)
- ◆ Add a mirrored transaction log (see “Adding a Mirrored Transaction Log” on page 550)
- ◆ Recreate NBDB (“Creating the NBDB Database” on page 551)

## Changing the Database Password

The DBA and application password may be changed at any time. The password is encrypted using AES-128-CFB and stored in the `vxdbms.conf` file. The permissions on `vxdbms.conf` allow only root to read or write to the file.

---

**Note** VERITAS recommends changing the password after installation.

---

The default password set during installation is *nbu\$sql*. This password is used for NBDB and BMRDB and for all DBA and application accounts. (For example, *EMM\_MAIN*.)

### ▼ To change the database password

1. Log on to the server as `root`.
2. Run the following command:

```
/usr/opensv/db/bin/nbdb_admin -dba new_password
```

The `vxdbms.conf` file is updated with the new, encrypted string.

## Moving NBDB Database Files After Installation

The `nbdb_move` command allows the administrator to change the location of the database files or split the database files into multiple directories. Doing so could improve performance in the case of large databases. This command moves both NBDB and BMRDB, if present.

`nbdb_move` is located in the following directory:

```
/usr/opensv/db/bin/nbdb_move
```

The `nbdb_move` command can be run at any time because it does not drop, then recreate the database. Thus, all data is preserved.

---

**Note** VERITAS recommends backing up NBDB and BMRDB using the catalog backup method of choice both before and after running `nbdb_move`.

---

### ▼ To move the NBDB and BMRDB database files

1. Perform a catalog backup.

2. Shut down all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```

3. Start the ASA daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

4. Move the existing data, index, and transaction log files:

```
/usr/opensv/db/bin/nbdb_move -data data_directory
-index index_directory -tlog log_directory
```

Or, if a mirrored transaction log is being used:

```
/usr/opensv/db/bin/nbdb_move -data data_directory
-index index_directory -tlog log_directory
-mlog log_mirror_directory
```

5. Start all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

6. Perform a catalog backup using the configured method.

## Adding a Mirrored Transaction Log

The transaction logs, `NBDB.log` and `BMRDB.log`, are critical files used to recover the ASA databases.

For extra protection, a mirrored transaction log can be used. This mirrored log should be created in a different directory from the original log.

### ▼ To create a mirrored transaction log

1. Perform a catalog backup.

2. Shut down all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```

3. Start the ASA daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

4. To create the mirrored transaction log only, enter:

```
/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, enter:

```
/usr/opensv/db/bin/nbdb_move -data data_directory
-index index_directory -tlog log_directory
-mlog log_mirror_directory
```

5. Start up all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

6. Perform a catalog backup.



## Creating the NBDB Database

The NBDB database is created automatically during NetBackup installation. However, it may be necessary during certain catalog recovery situations to create it manually using the `create_nbdb` command. `create_nbdb` is located in the following directory:

```
/usr/opensv/db/bin/create_nbdb
```

---

**Caution** Recreating the database manually is not recommended in most situations.

---



---

**Caution** If NBDB.db database already exists, running `create_nbdb` will overwrite it. If you want to move the database, move it using the `nldb_move` command.

---

### ▼ To manually create the NBDB database

1. Shut down all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```

2. Start the ASA daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

3. Run the following command:

```
/usr/opensv/db/bin/create_nbdb
```

4. Start up all NetBackup daemons:

```
/usr/opensv/NetBackup/bin/goodies/netbackup start
```

5. The new NBDB database is empty and does not contain the EMM data that is loaded during a normal installation.

Before repopulating this data, make sure that you have the most current support for new devices. New devices are added approximately every two months.

---

**Note** This procedure differs from releases prior to NetBackup 6.0.

---

- a. Obtain the `external_types.txt` mapping file from `support.VERITAS.com`.

- b. Place `external_types.txt` in `/usr/opensv/var/global`

This replaces the current `external_types.txt` file.

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.

## Post-installation Tasks

---

`/usr/opensv/volmgr/bin/tpext`

During regular installation, `tpext` is run automatically.

---

**Caution** If using the `create_nbdb` command to manually create a database, the `tpext` utility must also be run. `tpext` loads EMM data into the database.

---

### Additional `create_nbdb` Options

Besides using the `create_nbdb` command to create the NBDB database, it can also be used to perform the following actions. In each command, `VERITAS_NB_server_name` matches the name in `server.conf`. (See “`server.conf`” on page 540.)

- ◆ Drop the existing NBDB database and recreate it in the default location:

```
create_nbdb -drop
```

The `-drop` option instructs NetBackup to drop the existing NBDB database. The location of the current NBDB data directory is retrieved automatically from the `bp.conf` file.

- ◆ Drop the existing NBDB database and do not recreate:

```
create_nbdb -drop_only
```

- ◆ Drop the existing NBDB database and recreate it in the directories as specified:

```
create_nbdb -drop -data data_directory -index index_directory
-tlog log_directory [-mlog log_mirror_directory]
```

If the NBDB database files have been moved from the default location using `nbdb_move`, use this command to recreate them in the same location by specifying `current_data_directory`.

---

**Caution** If the location of `NBDB.db` is changed from the default, and `BMRDB.db` also exists on the server, `BMRDB.db` must also be recreated since its files must reside in the same location as the NetBackup database files.

---

## Backup and Recovery Procedures

There are two methods for backing up a NetBackup catalog:

- ◆ Online, hot catalog backup (recommended method)

The online, hot catalog is considered an *online, hot* method because it can be performed while regular backup activity is taking place.

Other benefits of the online, hot catalog method include:

- ◆ Runs per a policy and is virtually transparent to the customer; the policy can be set up using either the Catalog Backup wizard or the Policy wizard.

Either wizard automatically includes all the necessary catalog files, including the database files (NBDB and BMRDB), and any catalog configuration files (`vxdbms.conf`, `server.conf`, `databases.conf`.)

- ◆ Allows the administrator to recover either the entire catalog or pieces of the catalog. (For example, the databases separately from the image catalog.)
- ◆ Offers the incremental backup. For Sybase ASA, this means just a backup of the transaction log. Transaction logs are managed automatically, truncated after each successful backup.

For more information, see “Online, Hot Catalog Backup Method” on page 218.

- ◆ Offline, cold catalog backup

This type of catalog backup is considered an *offline, cold* backup because it should not be run when regular backup activity is taking place. For Sybase ASA, the databases (NBDB and BMRDB) are shut down during the catalog backup. For more information, see “Offline, Cold Catalog Backup Method” on page 228.

The default for environments upgrading to NetBackup 6.0, is to remain with the offline, cold catalog backup method. A hot catalog backup would need to be configured.

## Using the Online, Hot Catalog Backup Method

Normally, during a hot, online catalog backup there is one parent job and two or more child jobs. Logging for these jobs appears in the `dbm log`.

An overview of the hot catalog backup process consists of the following steps:

1. A temporary copy of database files is made to a staging directory:

`/usr/opensv/db/staging`

Once the copy is made, NetBackup can back up the catalog files.

2. A child job backs up files in a single stream:

- ◆ Configuration files (`server.conf`, `database.conf`, `vxdbms.conf`)

- ◆ Database files:

- ◆ `NBDB.db`
- ◆ `NBDB.log`
- ◆ `EMM_DATA.db`
- ◆ `EMM_INDEX.db`

If BMR has been installed:

- ◆ `BMRDB.db`
- ◆ `BMRDB.log`
- ◆ `BMR_DATA.db`
- ◆ `BMR_INDEX.db`

3. A second child job begins the image catalog backup.

---

**Note** The backup of any 5.x media server displays as a separate job.

---

---

**Note** If BMR is installed and a remote EMM server is being used, the backup of the EMM server displays as a separate job.

---

4. Transaction logs are truncated after a successful full or incremental backup.

If the transaction logs are manually changed or deleted, there could be a hole in the recovery.

The child job for the backup of the relational database files will normally be run on the master server because this is the default location for `NBDB` and the required location for `BMRDB`.

If NBDB has been moved to a media server, the child job will run on the media server. In this case, there will be additional logging for the job in the admin log on the media server.

If NBDB has been moved to a media server, and BMRDB is installed on the master server, there will be two child jobs for the relational database backup portion of the online, hot catalog backup: one on the media server for NBDB and one on the master server for BMRDB.

## Using the Offline, Cold Catalog Backup Method

The NetBackup relational database files are included during an offline, cold catalog backup.

An overview of the cold catalog backup process consists of the following steps:

1. The ASA databases are queried for the location of the database files associated with the database. In the case of NBDB, the following files are dynamically located:

- ◆ NBDB.db
- ◆ NBDB.log
- ◆ EMM\_DATA.db
- ◆ EMM\_INDEX.db

In the case of BMRDB, the following files are dynamically located:

- ◆ BMRDB.db
- ◆ BMRDB.log
- ◆ BMR\_DATA.db
- ◆ BMR\_INDEX.db

If these files have been moved (using `nbdm_move`) from the default location (`/usr/opensv/db/data`), the locations will be determined automatically.

2. The databases, NBDB, and, if it exists, BMRDB, are shut down. The Sybase ASA daemon continues to run.
3. The relational database files identified in step 1 as well as the image catalog files, are backed up.
4. If the backup is successful, the transaction logs are truncated and the databases are restarted. If the backup was not successful, the databases are restarted without truncating the transaction logs.

## Backup and Recovery Procedures

---

If NBDB has been moved to a media server, the offline, cold catalog backup will include the database files on the media server and will shutdown/startup the database remotely. Additional logs will appear in the admin log on the media server.

## Transaction Log Management

The transaction log for the NetBackup database, necessary for recovering the database, is automatically truncated after a successful catalog backup (either online, hot or offline, cold). The transaction log, NBDB.log, is located by default in:

```
/usr/opensv/db/data/NBDB.log
```

The transaction log continues to grow until it is truncated, so it is crucial that either online, hot or offline, cold catalog backups are being run frequently enough so that the transaction log doesn't grow to the point of filling up the file system on which it is located.

In addition to the default transaction log, a mirrored transaction log can be created for additional protection of NBDB using:

```
/usr/opensv/db/bin/nbdb_move -mlog mirrored_log_directory
```

The log is called:

```
mirrored_log_directory/NBDB.m.log
```

The directory for the mirrored log should not be the same as the directory for the default transaction log, and, ideally, would 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:

```
/usr/opensv/db/data/BMRDB.log
```

with an optional mirrored log in:

```
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.

---

**Caution** If a catalog backup is not being run, the logs won't be truncated. Managing the truncation in this manner is critical to recovery of the database.

---

## Catalog Recovery

The method used to recover the catalog in a disaster recovery situation depends on the method used to back up the catalog.

Recovery scenarios include:

- ◆ 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 the command line, `bprecover`, the relational database files and configuration files can be restored and recovered from either an online, hot or offline, cold catalog backup.

Catalog recovery scenarios and procedures are discussed in the *NetBackup Troubleshooting Guide*.

## Additional Command Lines for Backup and Recovery of the Relational Databases

The recommended method for protection of the relational databases is via the catalog backup and recovery interfaces.

In addition, a temporary backup of the NBDB and BMRDB databases can be made to a directory. This backup could be used for extra protection before performing database administration activities such as moving or reorganizing the database files.

### **nbdb\_backup**

Use `nbdb_backup` to make either an online or an offline copy of the NBDB and BMRDB database files in a directory. The transaction log won't be truncated using `nbdb_backup`. Transaction logs are managed only by using the catalog backup.

```
/usr/opensv/db/bin/nbdb_backup [-dbn database_name] [-online |
-offline] destination_directory
```

`-dbn database_name` only backs up the specified database (NBDB or BMRDB).

`-offline` shuts down the database and access to the database. Connections to the database at this time will be refused. The ASA daemon does not shut down.

---

**Caution** The transaction logs are not truncated using `nbdb_backup`. A catalog backup must be run in order to truncate the logs.

---

### **nbdb\_restore**

Use `nbdb_restore` to recover from a database backup that was made using `nbdb_backup`.

```
/usr/opensv/db/bin/nbdb_restore -recover source_directory
```

## Database Unloading Tool

---

Logging for these commands is in the admin directory.

## Database Unloading Tool

The `nbdb_unload` command line utility can be used to dump the entire NetBackup or Bare Metal Restore databases, or individual tables (one `.dat` file is created for each table), or schema. The utility can be used to create a copy of the ASA database which may be requested in some customer support situations.

There should be no active connections to the database when running `nbdb_unload`. See the following section, "Terminating Database Connections."

A `reload.sql` script is generated as part of running `nbdb_unload`. The script contains all the code required to recreate the database. This script and associated files can be used by VERITAS Technical Support to assist in troubleshooting a support case.

```
/usr/opensv/db/bin/nbdb_unload [-dbn database_name] [-t
table_list] [-s] destination_directory
```

Where:

- ◆ `-dbn database_name`

*database\_name* is NBDB (default) or BMRDB.

- ◆ `-t table_list`

Must give the owner of the table, then the table name. Using EMM, for example: all tables are owned by the account `EMM_MAIN`.

```
nbdb_unload -t EMM_MAIN.EMM_Device, EMM_MAIN.EMM_Density
```

- ◆ `-s`

Schema only is dumped; no data.

- ◆ `destination_directory`

Specify the location where dump is created.

## Terminating Database Connections

To eliminate concurrency problems, terminate all active connections to the database by shutting down NetBackup before running `nbdb_unload`.

### ▼ To terminate connections

5. Shut down all NetBackup daemons:



```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```

**6.** Start the ASA daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

**7.** Start only the database server by using

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

**8.** Run `nbdb_unload`, indicating the desired outputs (database name, table lists, or schema only) and the destination directory.

**9.** Shut down the database server by using `/usr/opensv/netbackup/bin/nbdbms_start_stop stop`.

**10.** Stop the ASA daemon:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop stop
```

**11.** Start up all NetBackup daemons:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

---

**Caution** VERITAS does not recommend using `reload.sql` to make a copy of the relational databases in a production environment. Either `nbdb_backup` should be used to make a physical copy, or `nbdb_move` should be used to relocate the database files.

---

## Moving the NetBackup Database from One Host to Another

The NetBackup database, NBDB, must always reside on the same host as the EMM server. If NBDB is moved, the EMM server must also be moved. The Bare Metal Restore database, BMRDB, must always reside on the master server. So, if NBDB and EMM server are moved to a media server from a master server, BMRDB must remain on the master server.

Use the following procedure to move the NetBackup database (NBDB) from host A to host B.

1. Perform a catalog backup.
2. If NetBackup is currently installed on B:
  - a. Shut down all NetBackup daemons on B:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```
  - b. Change the EMMServer entry in the `bp.conf` file from A to B on B.
  - c. Start the Sybase ASA server on B:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```
  - d. Create NBDB and associated files in the default location (`/usr/opensv/db/data`) on B:

```
/usr/opensv/db/bin/create_nbdb
```

Or, if NetBackup has not been installed on B:

Install NetBackup on B, identifying B as the EMM server during installation.

3. Set the database password on host B to match the password on A if the password has been changed from the default:

```
/usr/opensv/db/bin/nbdb_admin -dba password
```
4. Shut down NetBackup on A, B, and on all master and media servers that are using host A as the EMM server:

```
/usr/opensv/netbackup/bin/goodies/netbackup stop
```
5. Copy the following catalog files from their location on A to the desired final location on B:

---

**Note** The desired final location on B does not need to be the same as the original location on A.

---

---

**Note** Do not copy `vxdbms.conf`.

---

NBDB.db  
 EMM\_DATA.db  
 EMM\_INDEX.db  
 NBDB.log  
 NBDB.m.log (optional)

If the database files on both A and B are in the default location  
 (/usr/opensv/db/data) and server A is also a UNIX server, go to step 11.

6. Change `databases.conf` on A and B so that the databases don't start automatically when the server is started:

```
/usr/opensv/db/bin/nbdb_admin -auto_start NONE
```

7. Start the Sybase ASA server on B:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

8. To update the catalog with the location of the database files on B:

```
nbdb_move -data dataDirectoryB -index indexDirectoryB

-tlog tlogDirectoryB [-mlog mlogDirectoryB] -config_only
```

9. Stop the Sybase ASA server on B:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop stop
```

10. On B, for any database files that were copied to non-default locations in step 5, go to the default directory on B and delete the appropriate database files.

EMM\_DATA.db  
 EMM\_INDEX.db  
 NBDB.db  
 NBDB.log  
 NBDB.m.log (optional)

11. Change the `EMMSERVER` entry in the `bp.conf` file from A to B on all master and media servers that were using A as the EMM server.

12. On A, delete the database and configuration files:

EMM\_DATA.db  
 EMM\_INDEX.db  
 NBDB.db  
 NBDB.log  
 NBDB.m.log (optional)

---

## Moving the NetBackup Database from One Host to Another

### 13. On A:

- ◆ If BMRDB does not exist on A, delete the following configuration files:

```
dataDirectoryA/vxdbms.conf
/usr/opensv/var/global/databases.conf
/usr/opensv/var/global/server.conf
```

Remove the VXDBMS\_NB\_DATA entry from the bp.conf file.

- ◆ If BMRDB exists on A, execute the following command on A so that BMRDB starts automatically when the server is started:

```
/usr/opensv/db/bin/nbdb_admin -auto_start BMRDB
```

### 14. Start NetBackup on B:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

### 15. Start NetBackup on all master and media servers that are now using B as the EMM server:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

### 16. Perform a full catalog backup.

## NearStore Storage Unit Considerations

## B

NetBackup 6.0 introduces the Network Appliance (NetApp) NearStore™ disk storage unit.

This appendix contains information for NearStore administrators beyond the storage unit information that appears in “Managing Storage Units” on page 23. This appendix addresses the following topics:

- ◆ “Required Software, Hardware, and Licenses” on page 564
- ◆ “Advantages of the NearStore Storage Unit” on page 566
- ◆ “NearStore Configuration” on page 568
- ◆ “Viewing the Backup Image” on page 572
- ◆ “Disk Consumption” on page 574
- ◆ “Logging Information” on page 575
- ◆ “Troubleshooting” on page 575

Essentially, NetBackup writes client backup data to NearStore disk in tar format. After the tar image is complete, a snapshot is taken of the tar image and the data is converted into a WAFL qtree on the NearStore.

Useful terms to understand in this section include:

*WAFL* (Write Anywhere File Layout): The file system used in all Network Appliance storage servers. WAFL supports snapshot creation.

*qtree* (quota tree): A subdirectory in a NearStore volume that acts as a virtual subvolume with special attributes, primarily quotas and permissions.

*Snapshot*: A read-only, point-in-time copy of the entire volume. A snapshot captures file modifications without duplicating file contents.

## Required Software, Hardware, and Licenses

In order to configure a NearStore storage unit, the following hardware and software (and accompanying licenses) must be in place and configured:

- ◆ NetBackup Enterprise 6.0 on the master server and media server, with the Disk Optimization Option license installed.
- ◆ NetBackup client software 4.5 or later.
- ◆ A Network Appliance NearStore appliance with the following software installed:
  - ◆ Network Appliance Data ONTAP 7.1 or later
  - ◆ A SnapVault™ secondary license (enabled)

---

**Note** The credentials for the NearStore can be created using the `tpconfig` command. A NetBackup NDMP license is *not* required to create a NearStore storage unit. However, NDMP should be enabled on the NearStore since this enables the NearStore credentials to be entered using the NetBackup Administration Console. Specifically, credentials can be added via **Add NDMP Host** under **Media and Device Management > Devices**.

---

## NearStore and SnapVault Topics

### NearStore Storage Units and SnapVault Storage Units Cannot Share Volumes

NearStore and SnapVault storage units cannot share the same volume. To prepare a volume to support NearStore storage units:

- ◆ Disable any incremental backups to the secondary qtrees that were originally scheduled for the volume.
- ◆ Release existing SnapVault relationships. See the *Data ONTAP System Administrator's Guide* for instructions on running the `snapvault -stop` command to stop all backups and delete the qtrees and configurations on a volume.
- ◆ Disable any existing WAFL snapshots on the volume. This includes the default WAFL snapshot schedule that is automatically configured when the volume is created.

### Cleaning Up Configured qtrees

Qtree entries in the SnapVault configuration database are not deleted when a volume is destroyed. Make sure to delete the qtree entry in the configured database:

Log in to the filer.

```
r200> snapvault status -c /* lists config'd Qtrees */
r200> snapvault stop -f /vol/volume_name/Qtrees_name
```

Or, as an alternative, run:

```
rsh r200 snapvault status -c | grep /volume_name/ |
awk '{ print $1; }' | while read QT; do
rsh r200 snapvault stop -f $QT; done
```

## NearStore SnapVault Snapshot Schedules

To display the currently configured SnapVault snapshot schedule, enter the `snapvault snap sched` command and view the basenames of all snapshot sets for all SnapVault volumes on the filer. To accomplish this, run:

```
r200> snap sched MYVOLUME 0 0 0
```

---

**Note** Volumes configured for NetBackup NearStore storage units do not support the `snapvault snap sched` command. Any attempt to run the `snapvault snap sched` command on an NetBackup NearStore volume will fail.

---

To turn off the SnapVault schedule for a set of snapshots and stop the snapshot process for the SnapVault secondary storage system, enter the following command:

```
r200> snapvault snap unsched VOLUME_NAME
```

---

**Note** This command does not end the SnapVault relationships between the secondary system qtrees and their platform source drives or directories; this must be accomplished by running:

```
r200> snapvault stop -f /vol/VOLUME_NAME/QTREE_NAME
```

for each qtree configured or existing on the volume to be used as a NearStore storage unit.

---

## Advantages of the NearStore Storage Unit

### ◆ Backup data reduction

NearStore avoids duplicating disk blocks by comparing blocks in the most recent backup with the preceding backup image. Incremental backups do not consume disk space unless blocks in the new backup image differ from blocks in the active file system. As a result, multiple backups to the same volume store only uncommon blocks, and blocks that are common continue to share.

To view storage savings, use the Data ONTAP `df -s` command:

```
r200> df -s
Filesystem used shared saved %saved
/vol/vol0/ 1335000 0 0 0%
/vol/flexsle/ 96 0 0 0%
/vol/sim/ 292 0 0 0%
/vol/p3/ 21304124 14637384 21731976 50%
```

The `df` command is described in the *Data ONTAP System Administrator's Guide*.

### ◆ Tar image retained for interoperability

The NetBackup tar image is preserved in a WAFL qtree in order to support NetBackup administrative functions such as file restoration, duplication, staging, import from disk, and twinning to another storage unit.

### ◆ No administrator necessary to perform user restores

Client backups to a NearStore storage unit are translated by the NetBackup media server into WAFL qtree images. This feature is available with Data ONTAP 7.2 and future NetBackup releases.

### ◆ Extensive support

There are virtually no restrictions using a NearStore storage unit:

- ◆ Policy types, attributes, and schedule types are supported.
- ◆ All supported NetBackup clients can be backed up to, and restored from, a NearStore storage unit.
- ◆ NearStore storage units can be used for synthetic backups.
- ◆ One of the best applications for a NearStore storage unit is that of a disk staging storage unit or the target of a disk staging storage unit.
- ◆ NearStore allows twinning to other storage units—tape or disk.
- ◆ NearStore allows duplication, expiration, and verification of images.
- ◆ NearStore images can be imported.
- ◆ NearStore storage units can be used for NetBackup catalog backups and restores.



---

### Advantages of the NearStore Storage Unit

- ◆ NearStore can serve as a source or a target for Vault.
- ◆ **Restrictions in current release**
  - ◆ The NearStore storage unit does not support the checkpoint restart feature. This restriction will be removed in future NetBackup releases.
  - ◆ The NearStore storage unit does not support backups based on the following policy configurations:
    - ◆ Multistreamed backups using NetBackup for Microsoft SQL Server.
    - ◆ Multistreamed backups using NetBackup for Sybase.
    - ◆ Multistreamed NetBackup for Oracle Proxy Block Level Incremental (BLI) backups.

## NearStore Configuration

To make the NearStore storage unit available for backups, add and enable the secondary SnapVault license:

1. Add the secondary SnapVault license:

```
r200> license add sv_secondary_license
```

2. Enable SnapVault:

```
r200> options snapvault.enable on
```

3. Grant access to media servers authorized to access the NearStore by entering the following command:

```
r200> options snapvault.access host=nbu_media_server1,
nbu_media_server2...
```

## NearStore Authentication

A NearStore user name and password must be configured in NetBackup before any backups are run. NearStore authentication information is stored in the NetBackup Enterprise Media Manager (EMM) database. This allows for global authentication or authentication on a per-media server basis.

### ▼ To authenticate the NetBackup media server

1. Make sure that the SnapVault license is enabled on the media server.
2. Use the `tpconfig` command to add a NearStore user ID and password into the EMM global database:

- ◆ Authenticates only the media server where this is run:

```
tpconfig -add -snap_vault_filer -nh hostname -user_id userID
[-password password]
```

- ◆ Authenticates all media servers:

```
tpconfig -add -snap_vault_filer -nh hostname -filer_user_id
userID
```

### ▼ To create a root NearStore user name and password

```
tpconfig -add -snap_vault_filer -nh hostname -user_id root_ID
[-password root_password]
```

---

**Note** It is important to avoid using `root` as an NDMP password because `root` is not encrypted, and could compromise the integrity of your storage system.

---

▼ **To create a non-root NearStore user name and password**

Administrators can create user accounts on the NearStore in order to perform backups and restores. The following procedure describes a method to send a user's encrypted password over the network.

1. Log onto the Nearstore.
2. To create a new user, enter the following command:  
`useradmin user add userID -g groupID`  
 Enter a password when prompted.
3. To receive the encrypted version of the password, enter:

```
ndmpd password userID
```

where *userID* is the name of the user just added.

4. Record the password.
5. Log out of the NearStore.
6. Enter the following `tpconfig` command:

```
./tpconfig -add -user_id username -nh nearstore_name
-snap_vault_filer -password encrypted_password
```

The encrypted password is passed across the network.

▼ **To verify that the NearStore credentials have been entered into the NetBackup EMM database**

Once the `tpconfig` command is run, ensure that the media server is authenticated by running the following command:

```
nbemmcmd -listhosts -list_snap_vault_filers -machinename
media_server_name
```

For example:

```
C:\Program Files\VERITAS\NetBackup\bin\admincmd>
nbemmcmd -listhosts -list_snap_vault_filers -machinename entry
NBEMMCMD, Version:6.0CA(20050628)
The following hosts were found:
```

## NearStore Configuration

```
ndmp mmnetapp.xxx.yyy.com
Command completed successfully.
```

## NearStore Disk Storage Unit Properties

Settings specific to NearStore are available when configuring a NearStore disk storage unit. The following sections describe each. Storage units are also discussed in “Managing Storage Units” on page 23.

**Note** NearStore storage units cannot be included in storage unit groups.

### ◆ Storage Unit Type

For a NearStore storage unit, select *Disk* as the **Storage unit type**.

### ◆ On Demand Only

A NearStore storage unit can only be used on demand. **On demand only** cannot be deselected.

### ◆ Disk Type

When configuring disk storage units, there are three disk types available to choose from: Basic Disk, SnapVault, and NearStore.

### ◆ NearStore Server

The **NearStore server** drop-down list contains all NearStore hosts configured for the selected media server and available to NetBackup.

### ◆ Absolute Pathname to Volume

The **Absolute pathname to volume** drop-down list contains all the volumes in the selected NearStore that are capable of serving as NetBackup storage units. The list displays only flexible volumes. For example, WORM volumes are filtered out.

### ◆ Properties button

Click the **Properties** button to display the capacity of the selected volume and the available storage.

### ◆ High Water Mark, Low Water Mark

To avoid filling up the NearStore and potentially causing problems, use the **High water mark** setting to control the amount of data on the NearStore.

◆ **Enable Block Sharing**

The **Enable block sharing** setting allows data blocks that have not changed from one backup to the next to be shared. Sharing data blocks can significantly save disk space in the storage unit.

◆ **Temporary Staging Area**

Using a NearStore storage unit for disk staging is a recommended use because of the quick speed of NearStore servers.

◆ **Reduce Fragment Size**

NearStore uses the **Reduce Fragment Size** setting differently than other storage units. NetBackup writes to a NearStore by laying out the data in one large image, and not dividing the data into fragments.

## Viewing the Backup Image

Use the `bpstsinfo` command to look for images on the NearStore, as well as to look at server and logical storage unit (LSU) attributes. The `bpstsinfo` command is described in *NetBackup Commands for UNIX and Linux*. The command options are listed below:

`bpstsinfo`

Only one of the following must be specified:

`-serverinfo`, `-lsuinfo`, `-imageinfo`, or `-deleteimage`.

`-serverinfo`

Print information about the server.

`-lsuinfo`

Print information about the logical storage unit.

`-imageinfo`

Print information about the image.

`-deleteimage`

Delete the specified image. Requires `-servername`, `-serverprefix`, `-lsuname`, `-imagename`, and `-imagedate` arguments.

`-servername server_name`

Optional argument. *server\_name* is the hostname of the storage server.

If `-servername` is not specified, the hostname of the local host is used.

Can be used with `-serverinfo`, `-lsuinfo`, or `-imageinfo`.

`-serverprefix server_prefix`

Optional filtering argument. By default, all server prefixes are used.

Specify *server\_prefix* to limit to one prefix.

Can be used with `-serverinfo`, `-lsuinfo`, or `-imageinfo`.

`-lsuname lsu_name [-lsuname lsu_name ...]`

Optional filtering argument. By default, all logical storage units are used.

Specify *lsu\_name* to limit to one logical storage unit for each `-lsuname` supplied.

Can be used with `-lsuinfo` or `-imageinfo`.

`-imagename image_name`

Optional filtering argument. By default, all images are used.

Specify an *image\_name* to limit to only matching images.

Can be used with `-imageinfo`.

`-imagedate image_date`

Used to specify a single image. Acceptable formats:

03/08/2005 09:41:22  
1110296416

Can be used with `-imageinfo` only.

Cannot be used with `-imagedatestart` or `-imagedateend`.

`-imagedatestart image_date`

Optional filtering argument. By default, all images are used.

Specify an image date to limit to images equal to or newer than the *image\_date*.

Acceptable formats:

03/08/2005 09:41:22  
1110296416

Can be used with `-imageinfo`. Cannot be used with `-imagedate` option.

`-imagedateend image_date`

Optional filtering argument. By default, all images are used.

Specify an image date to limit to images equal to or newer than the *image\_date*.

Acceptable formats:

03/08/2005 09:41:22  
1110296416

Can be used with `-imageinfo`. Cannot be used with `-imagedate` option.

`-imagetype image_type`

Optional filtering argument. By default, both full and incremental images are used.

Specify *image\_type* of `STS_FULL_ONLY` or `STS_INCR_ONLY` to limit to only a specific image type.

Can be used with `-imageinfo`.

`-remote server_name [-remote remote_server ...]`

Optional argument to query remotely for disk information about each `-remote` supplied. Can be used with `-serverinfo`, `-lsuinfo`, `-imageinfo`, or `-deleteimage`.

## Disk Consumption

Snapshot creation consumes a large amount of disk space. NetBackup prepares for this space requirement by reserving 20% of the disk space on the volume to be used exclusively for the snapshot, and not for the active file system.

If the snapshots exceed the reserved amount, space is consumed as needed from the active file system. The active file system cannot, however, consume disk space reserved for snapshots.

### File System Full Conditions

Whenever snapshots consume more than 100% of the reserved space, the active file system is in danger of becoming full. Under these conditions, backups fail until administrative action is taken.

Administrative action could include:

- ◆ Expiring images through NetBackup. This can be accomplished through the Catalog interface.
- ◆ Lowering the retention level for images so that images are expired sooner.


### End of Media Detection on Disk Staging Storage Units

In order to permit End of Media detection on NearStore disk staging storage units, backup performance is diminished when a NearStore volume is within 4 gigabytes of being full.

### Multi-using the NearStore

The volume properties of the NearStore storage unit display a value for the storage available on the volume. However, the value doesn't reflect any multi-use situations in which an administrator has allotted part of the volume for another use.

**Caution** VERITAS strongly recommends using the volume for NetBackup only.





## Logging Information

Logging occurs in the following locations log files:

- ◆ By default (if not twinning): `/usr/opensv/netbackup/logs/bpdm`
- ◆ For twinning logs only: `/usr/opensv/netbackup/logs/bptm`

All other logging is similar to a standard backup, producing, for example, progress logs.

Logs contain more information about the interaction with NearStore.

On the NearStore, the root volume contains a NetBackup-specific log file that details the protocol between NetBackup and the NearStore.

ONTAP debug logs are found in the following location:

```
/vol/vol10/etc/logs/nbu_snapvault
```


## Troubleshooting

- ◆ Make sure that the permissions on the disk storage unit are set so that data can be written to the volume. If permissions are Read Only, NetBackup cannot write to it.
- ◆ Make sure that the SnapVault license has been added and is turned on:
 

```
license add sv_secondary_license
```
- ◆ Make sure the `tpconfig` command is used to add the NearStore user ID and password into the EMM global database. (See “To authenticate the NetBackup media server” on page 568.)
- ◆ Check the storage unit configuration to make sure that NearStore is selected as the storage unit type.
- ◆ If jobs are failing to write to the NearStore, make sure that the space reserved for snapshots on the NearStore is not completely full. When the reserved space is full, NetBackup uses the active file system space as needed.
- ◆ In the case of a disk full condition on the NearStore, make sure that there are no WAFL snapshots consuming disk space unnecessarily.
- ◆ If the maximum number of transfers allowed to a single NearStore is exceeded, the Ontap kernel reports the following error:



```
inf Wed Jul 6 07:28:27 CDT [10.80.106.36:58645] Maximum active
transfers reached.
```

The maximum number of concurrent backup and/or restore connections is 128.



## Troubleshooting

---



# Index

---

## Symbols

- .f files in catalog 214
- .SeCuRiT.y.nnnn files 512

## A

- Absolute Pathname to Directory storage unit setting 39
- Absolute Pathname to Volume storage unit setting 39, 570
- Accept Connections on Non-reserved Ports property 450
- Access Control
  - authorizing users 483
  - host properties
    - Authentication Domain tab 347
    - Authorization Service tab 349
    - VERITAS Security Subsystem (VxSS) 344
    - VxSS Networks List 345
    - VxSS tab 345
  - NetBackup 486
- access control lists (ACLs) 169, 512
- accessibility features xxxvii
- ACL (see access control lists)
- Activity Monitor
  - bpdbjobs command 334
  - BPDBJOBS\_OPTIONS environmental variable 334
  - cancelling uncompleted jobs 313
  - deleting completed jobs 313
  - detailed job status 313, 327
  - disabling job logging 425
  - monitoring jobs 313
  - restarting a completed job 313
  - resuming suspended jobs 314
  - saving job data to a file 314
  - set column heads to view 312
  - suspending a job 314
  - using the Troubleshooter 314
- Actual Client property (Backup Exec Tape Reader) 353
- Actual Path property (Backup Exec Tape Reader) 353
- Adaptive Server Anywhere - VERITAS\_NB 325
- Administer E-mail Address property 417
- administering remote systems 479
- administrator
  - defined streaming mode 185
  - nonroot 490
- Administrator's E-mail Address property 299
- Advanced Client 71, 75, 96, 110, 180, 318, 321
- AFS policy type 70
- All Log Entries report 298, 359
- ALL\_LOCAL\_DRIVES directive 181
- Allow Backups to Span Media property 424
- Allow Client Browse property 362
- Allow Client Restore property 362
- Allow Media Overwrite property 423
- Allow Multiple Data Streams
  - directives 184
  - set policy attribute 92
  - when to use 92
- Allow Multiple Retentions per Media property 117, 424
- Allow Server File Writes property 343, 450
- alternate client restores, allowing 505
- Alternate Restore Failover Machines host properties 433
- Announce DHCP interval property 429
- ANSI format, allow overwrite 423
- AOS/VS format, allow overwrite 423
- application backups 100
- archive bit 99, 199, 377, 380
- atime 518, 535
- attributes for a policy 87

- auth.conf file
  - capabilities identifiers 491
  - description 487
  - entries for specific applications 490
  - overview 487
- Authorization
  - host properties
    - DomainGroup 350
    - Group/Domain 351
    - Host 350
    - User 350
    - User must be an OS Administrator 351
  - preferred group 449
- auto-discover streaming mode 186
- automatic
  - backups 100
  - cumulative incremental backups 100
  - differential incremental backups 101
  - full backups 101
  - Vault policy type 101
- automounted directories 81
- B**
  - backlevel administration
    - administration consoles 8
  - Backup End Notify Timeout property 446
  - Backup Exec QIC media, importing and restoring 271
  - Backup Exec Tape Reader
    - host properties
      - Actual Client 353
      - Actual Path 353
      - GRFS Advertised Name 352
  - Backup Migrated Files property 371
  - Backup Start Notify Timeout property 444
  - Backup Status report 359
  - backups
    - activating policy 78
    - application 100
    - archive 100
    - automatic 100
      - cumulative incremental 99, 100
      - differential incremental 99, 101
      - full 101
      - Vault 101
    - balancing load 533
    - Client Backups report 296
    - clients using Storage Migrator 535
    - deactivating policy 78
    - duplicating 260
    - frequency
      - effect on priority 110
      - guidelines for setting 109
      - setting 108
    - full 99
    - import 266
    - network drives 79
    - NFS mounted files 69, 81
    - offsite storage 115
    - policy management window 60
    - raw partitions on Windows 76, 164
    - registry on Windows clients 165
    - selections list, verifying 160
    - types of 98
    - user directed
      - schedules 139
      - type of backup 100
    - verifying 257
    - windows
      - duration, examples 120
      - specifying 119
  - Bandwidth
    - host properties
      - Bandwidth 355
      - Bandwidth Throttle Setting for the Range of IP Addresses 354
      - Bandwidth 355
      - Bandwidth Throttle Settings List 355
      - From IP Address 355
      - To IP Address 355
    - Bandwidth Throttle Settings List
      - property 355
  - Bare Metal Restore (BMR) 87, 216, 318, 325, 379, 537
  - Basic Disk storage units 24, 40
  - batch file example for setting bpdjobs
    - environmental variable 335
  - BE-MTF1 format, allow overwrite 423
  - Block Level Incremental Backups 75
  - BMRD (NetBackup Bare Metal Restore Master Server) 325
  - BMRDB.db
    - configuration entry 545
    - in catalog 216
    - relocating 540
  - bp.conf file

- configuring to use ctime 106
    - customizing jnbSA and jbpSA 497
    - entries for Activity Monitor 333
    - indicating database location 538
    - indicating files not to be compressed 86
    - NetBackup-Java Administration Console configuration entries 491
    - obtaining list of entries 340
    - when master servers share EMM database 441
  - BPARCHIVE\_POLICY 141
  - BPARCHIVE\_SCHED 141
  - bpbackup command 279
  - BPBACKUP\_POLICY 141
  - BPBACKUP\_SCHED 141
  - bpbackupdb command 279
  - BPBRM Logging property 420
  - bpcatarc command 252
  - bpcatlist command 251
  - bpcatres command 252
  - bpcatrm command 252
  - BPCD Connect-back property 396, 398
  - BPCD port setting on client 429
  - bpchangeprimary command 259
  - bpclient
    - add clients to catalog 513
    - delete clients from catalog 514
    - list clients in catalog 514
    - preventing lists and restores 514
  - BPCOMPATD (NetBackup Compatibility Service) 325
  - bpconfig command 187
  - bpcoord log 208
  - bpdbjobs
    - adding entries to bp.conf file 333
    - batch file example 335
    - command 334
    - debug log 335
  - BPDBJOBS\_OPTIONS environmental variable 334
  - BPDBM
    - Logging property 420
    - NetBackup Database Manager, description 325
  - BPDM Logging property 420
  - bpend 446
  - bpexptime 263
  - bpfis directory for VSP logging messages 455
  - BPJAVA\_PORT 492
  - bplabel 323
  - bpps script 467
  - BPRD
    - Logging property 420
    - NetBackup Request Manager, description 326
    - port setting on client 429
  - bpstart 444
  - bpstinfo command 572
  - bpsynth log 208
  - BPTM
    - Logging property 420
  - bpverify 324
  - Browse and Restore Ability property 364
  - buffer size 380
  - Busy Action property 358
  - Busy File
    - host properties
      - Busy File Action 358
      - File Action File List 358
      - Operator's E-mail Address 357
      - Process Busy Files 357
      - Retry Count 358
      - Working Directory 357
    - processing
      - Windows clients 430
- C**
- cachefs file systems, excluding from backup 191
  - calendar scheduling
    - how it interacts with daily windows 125
    - using 122
  - cancelling uncompleted jobs 313
  - catalog archiving
    - bpcatarc command 252
    - bpcatlist command 251
    - bpcatres command 252
    - bpcatrm command 252
    - deactivating policy for 78
  - catalog backups
    - compressing image catalog 286
    - configuration 228
    - disk path 240
    - file paths
      - media server 245
      - Windows master 244
    - last media used 237

- manual backup 278
- media
  - ID 239
  - server 237
  - type 237
- offline, cold 318
- offline, cold method
- online, hot
  - in Jobs tab 323
- online, hot method 218
  - in Jobs tab 323
  - parent and child jobs 318
  - schedules for 42, 72, 221, 226, 416
  - volume pool 73
- overview 212
- policy type 193
- setting schedules 241
- space required 275
- uncompressing 287
- CatalogBackup volume pool 73, 301
- catalogs
  - archiving 248
    - bpcatarc 252
    - bpcatlist 251
    - bpcatres 252
    - bpcatrm 252
    - catarc policy 249
    - deactivate policy 249
    - extracting images 253
    - overview 248
    - retention level setting 250
    - type of backup indicated 250
  - image files 214
  - moving client images 282
  - multiple file layout 215
  - single file layout 214
- catarc schedule 62
- cautions
  - retention time 139
- CDE (Common Desktop Environment) 4
- cdrom file system, excluding from
  - backup 191
- change journal 379
  - and synthetic backups 208
  - determining if enabling is useful 378
  - using in incremental backups 378
- Check the Capacity of Disk Storage Units
  - property 37, 411
- Checkpoint Restart
  - and synthetic backups 208
  - Move Job From Incomplete State to Done State property 360
  - Move Restore Job from Incomplete State to Done State 518
  - Restore Retries 518
  - resuming a restore job 518
  - suspending a restore job 518
- cipher types for NetBackup Encryption 383
- Clean-up
  - host properties
    - Delete Vault Logs 359
    - Keep Logs 359
    - Keep True Image Restoration Information 359
    - Move Backup Job From Incomplete State to Done State 360
    - Move Backup Job from Incomplete to Done State 76
    - Move Job From Incomplete State to Done State 360
- client
  - database 362
  - exclude and include lists 394
  - name 505
  - Client Administrator's E-mail property 451
  - Client Attributes
    - host properties
      - Allow Client Browse 362
      - Allow Client Restore 362
      - Browse and Restore Ability 364
      - Clients List 362
      - Free Browse 364
      - Maximum Data Streams 363
  - Client Backups report 296
  - Client Cipher property 383
  - Client Connect Timeout property 444
  - Client Name property 370
  - Client Port Window property 431
  - Client Read Timeout property 372, 444, 445, 446
  - Client Reserved Port Window property 432
  - Client Sends Mail setting 451
- clients
  - BPCD port 429
  - BPRD port 429
  - changing in a policy 150
  - choosing policy type 69
  - deleting from policy 67

- DHCP Interval property 429
  - exclude file list 191, 387
  - installing 150, 152
  - maximum jobs 415
  - moving image catalog 282
  - peername 505
  - secure 152
  - setting host names 149
  - trusting clients 151
- Clients List property 362
- clustering 341, 547
- cold catalog backups (see catalog backups)
- Collect Disaster Recovery Information for Bare Metal Restore 87
- Collect True Image Restoration (TIR) with Move Detection property 200
- collecting disaster recovery information 82, 87
- column heads, selecting to view 312
- Communications Buffer property 380
- Compress Catalog Interval property 286, 416
- compression, by software
  - advantages 85
  - disadvantages 85
  - specifications 85
- concurrent jobs
  - on client 415
  - per policy 77
- CONNECT\_OPTIONS 398
- Consistency Check Before Backup host property 443
- copies, third-party 263
- Copy On Write snapshots 164
- copy, primary 263
- cpio format, allow overwrite 423
- critical policies, identifying 194
- cross mount points
  - effect with UNIX raw partitions 82
  - examples 84
  - policy attribute 172
  - separate policies for 83
  - setting 82
- ctime 176
- cumulative incremental backups 99, 102, 380
- Current NBAC User 19
- D**
- Daemon Connection Port property 397, 399
- Daemon Port Only property (for selection of ports) 399
- daemons, checking processes 467
- Daily windows setting 125
- database-extension clients, adding file paths for 180
- databases, NetBackup (see catalog backups)
- DataStore
  - policy type 69
  - volume pool 73, 301
- datetime stamp 105
- DB2 policy type 70, 101
- DBR format, allow overwrite 423
- Default Cache Device Path for Snapshots property 375
- DEFAULT\_CONNECT\_OPTIONS 396
- Delay on Multiplexed Restores property 411
- Delete Vault Logs property 359
- deleting
  - storage unit groups 57
- detailed job status 313, 327
- Device Configuration Wizard 27
- Device Monitor 332
- devpts file system, excluding from backup 191
- DHCP setting on client 429
- differential incremental backups 99, 380
- Direct Access Recovery (DAR) 412
- disaster recovery
  - collect information for 82, 87
  - file, sending 194, 222, 247
  - information 417
  - sending e-mails 194, 222, 247
- Disaster Recovery tab 193
- disk consumption 574
- Disk Optimization Option 24, 36, 40, 564
- disk staging
  - creating a storage unit 48
  - Final Destination Storage Unit 53
  - Final Destination Volume Pool 53
  - Priority of Duplication Jobs 113
  - relocation schedule 98
  - schedule 45
  - schedule button 52
  - schedule name 52
  - storage units
    - size recommendations 49
    - storage unit selection within a storage unit group 56

- using Checkpoint Restart 76
  - Use Alternate Read Server 53, 113
- disk storage units 45
- Disk Type storage unit setting 570
- disk-image backups 76, 164
- Do Not Compress Files Ending With
  - property 375
- DO\_NOT\_RESET\_FILE\_ACCESS\_TIME 53  
6
- DQTS (Enhanced Device Qualification Tool Suite) 323
- duplicate backups
  - becoming a primary copy 263
  - creating 260
  - restoring from 258
- duration of backup window, examples 120

## E

- E-mail
  - address for administrator of this client 451
  - disaster recovery 194
  - send from client 451
  - send from server 451
- Enable
  - Job Logging property 425
  - Open File Backup During Backups property 430
  - SCSI Reserve/Release property 425
  - Single Instance Backup for Message Attachments property 386
  - Standalone Drive Extensions property 425
- Enable Block Sharing storage unit setting 40, 571
- Enable Encryption property 382
- Enable Multiplexing unit setting 40
- Enable Performance Data Collection property 451
- Enable Standard Encryption property 383
- Encryption
  - host properties
    - Client Cipher 383
    - Enable Encryption 382
    - Enable Standard Encryption 383
    - Encryption Key File 384
    - Encryption Libraries 383
    - Encryption Permissions 382
    - Encryption Strength 383

- Use Legacy DES Encryption 383
  - in Client Backups Report 296
  - in Images on Media Report 303
  - policy attribute 87
  - use with synthetic backups 206
- Enhanced
  - Authentication 483, 486
  - Authorization 449, 483, 486
- Enhanced Device Qualification Tool Suite (DQTS) 323
- Enterprise Media Manager 325
- Enterprise Media Manager (EMM) 30, 306, 413, 439, 522, 523, 524, 525, 526, 537
- errors, media mount 332
- escape character
  - backslash 392
  - on UNIX 168
- Exceptions to the Exclude List host
  - property 388
- Exchange
  - host properties
    - Enable Single Instance Backup for Message Attachments 386
    - Mailbox for Message Level Backup and Restore 385
- exclude file lists
  - on client 394
- exclude files list
  - overview 191
  - Windows example 393
- Exclude List
  - host properties
    - Exceptions to the Exclude List 388
    - Use Case Sensitive Exclude List 387
- excluding files and directories from
  - backup 387
- export license key 472
- extended attribute files
  - disabling the restore of 177
  - Solaris 9 169
- external\_types.txt 551

## F

- fail all copies, multiple copies 112
- failover
  - media server to alternate media server(s) 433
- File Browse Timeout property 444
- File Change Log, using in VxFS 4.1 373



- file lists
  - disk image on Windows 164
  - extension clients 180
  - links on UNIX 170
  - NetWare clients
    - NonTarget 178
    - Target 179
  - raw partitions 164, 172
  - standard clients 168
  - UNIX files not backed up 163, 169, 190
  - Windows clients 162
- file systems 82
- files
  - .SeCuRiT.y.nnnn 512
  - /.rhosts 151
  - catalog space requirements 275
  - excluding from backup 387
  - for catalog backup 243
  - goodies scripts 522
  - linked, UNIX 169
  - NFS mounted 69, 81
  - No.restrictions 506
  - NOTES.INI 422
  - peername 506
  - restoring to alternate client 507
  - restrictions on restores 505
  - version xxxv
- filters
  - applying job 312
- Final Destination Storage Unit 53
- Final Destination Volume Pool 53
- Firewalls
  - host properties
    - BPCD Connect-back 398
    - Daemon Connection Port 397, 399
    - Default Connect Options 396
    - Hosts list 398
    - Ports 397, 399
  - using vnetd with 397, 398
- FlashBackup 71, 169, 172, 173
- Follow NFS 172
- follow NFS mounts
  - advantages of 81
  - disadvantages of 82
  - notes on use
    - with cross mount points 81
    - with raw partitions 81
    - with cross mount points 83
  - Follow NFS setting 81

- FORCE\_IPADDR\_LOOKUP 492
- Free Browse property 364
- freeze media 300
- From IP Address property 355
- full backups 99, 101, 201

## G

- General Level Logging property 377
- General Server
  - host properties
    - Check the Capacity of Disk Storage Units 37, 411
    - Delay on Multiplexed Restores 411
    - Media Host Override 412
    - Must Use Local Drive 411
    - Use Direct Access Recovery for NDMP Restores 412
- generic jobs 320
- Global Attributes
  - host properties
    - Administrator's E-mail Address 417
    - Compress Catalog Interval 416
    - Job Retry Delay 414
    - Maximum Backup Copies 416
    - Maximum Jobs per Client 139, 415
    - Maximum Vault Jobs 416
    - Policy Update Interval 414
    - Priority of Restore Jobs 415
    - Schedule Backup Attempts 414
- Global Logging Level property 419
- goodies directory 522
- GRFS Advertised Name property 352

## H

- hard links
  - NTFS volumes 166
  - UNIX directories 170
- High Water Mark storage unit setting 41, 570
- HKEYS, backing up 165
- host properties
  - changing in a clustered environment 341
  - permission to change 343
- hot catalog backups (see catalog backups)

## I

- IDX (index file) 303
- images
  - changing primary copy 258
  - duplicating 260

- import 266
- moving client catalog 282
- on Media report 303
- restoring from duplicate 258
- verifying 257
- Import backup images 266
- include
  - files list 191
  - list, on client 394
- Incrementals Based on
  - Archive Bit property 380
  - Timestamp property 379
- INETD (NetBackup Client Service) 325
- Informix policy type 71
- INI file, for Lotus Notes 422
- Initial Browse Search Limit property 447
- INITIAL\_BROWSE\_SEARCH\_LIMIT 497
- INITIAL\_MEMORY 494, 499
- Inline Tape Copy option 110, 260, 263, 264
- installing client software
  - on PC clients 152
  - on secure clients 152
  - on trusting clients 150
- Instant Recovery
  - Advanced Backup Method 75
  - Backups to Disk Only setting 110
- Intelligent Disaster Recovery (IDR) 82, 87
- Internet Assigned Numbers Authority (IANA) 492, 530

## J

- Java
  - auth.conf file 487
  - authorizing users 487
  - directory 489
  - jbp.conf file 496
  - jbpSA configuration options 496
  - jnb.conf file 496
  - jnbSA configuration options 496
  - performance improvement hints 499
- Java interface 3
- Java Virtual Machine (JVM) 494
- jnbSA 3
- Job Manager Logging property 421
- Job Retry Delay property 414
- jobs 314
  - concurrent per disk storage unit 42
  - filters, specifying 312
  - maximum

- per client 415
- per policy 77
- priority for policy 78
- Jobs (see Activity Monitor)
- JVM (Java Virtual Machine) 494

## K

- Keep Logs property 293, 359
- Keep Status of User-directed Backups, Archives, and Restores property 371, 372, 381
- Keep True Image Restoration Information property 359
- KEEP\_LOGS\_DAYS 497
- keyword phrase 95

## L

- last media used, catalog backups 237
- license keys
  - accessing 469
  - deleting 471
  - export 472
  - printing 470
  - using the NetBackup License Key utility 473
  - viewing the properties of one key 471
- Limit Jobs per Policy setting 77, 139
- links
  - UNIX hard-linked directories 170
  - UNIX symbolic 169
- load balancing 533
- Locked File Action property 372
- logging
  - bpcord 208
  - bpsynth 208
  - legacy 419
  - unified 418
- Logging enabled for debug 419
- logical storage unit (LSU) attributes 572
- logs
  - deleting after a set time 359
- Lotus Notes
  - host properties
    - INI File 422
    - Path 422
  - policy type 70
  - properties 422
- Low Water Mark storage unit setting 41, 570
- ltid (Media Manager Device) 325

**M**

- Mac OS X 70
- mail notifications
  - administrator E-mail address 451
  - Disaster Recovery attachment, sending 194
  - E-mail address for administrator 417
  - Windows nbmail.cmd script 417
- Mailbox for Message Level Backup and Restore property 385
- manual backups
  - NetBackup catalogs 278
  - policy for 196
- mapping file (external\_types.txt) 551
- master servers, rebooting 467
- MAX\_MEMORY 494, 499
- maximum
  - jobs per client 415
  - jobs per policy 77
- Maximum Backup Copies property 416
- Maximum Concurrent Jobs storage unit setting 42
- Maximum Concurrent Write Drives storage unit setting 41, 111
- Maximum Data Streams property 363
- Maximum Error Messages for Server property 381
- Maximum Jobs per Client property 415
- Maximum Vault Jobs property 416
- Media
  - host properties
    - Allow Backups To Span Media 424
    - Allow Media Overwrite 423
    - Allow Multiple Retentions Per Media 424
    - Enable Job Logging 425
    - Enable SCSI Reserve/Release 425
    - Enable Standalone Drive Extensions 425
    - Media ID Prefix (Non-robotic) 425
    - Media Request Delay 426
    - Media Unmount Delay 425
- media
  - 1 and media 2, catalog backup 237
  - active 305
  - freeze 300
  - ID for catalog backup 239
  - last used for catalog backup 237
  - nonactive 305
  - type for catalog backup 237
  - unfreeze 300
  - unsuspend 301
- Media Contents report 302
- Media Host Override property 412
- Media ID Prefix (Non-robotic) property 425
- Media List report 299
- Media Log Entries report 305, 359
- Media Manager Device daemon (ltid) 325
- media mount
  - errors
    - cancelled 332
    - queued 332
    - timeout for Storage Migrator 535
- Media Mount Timeout property 445
- Media Request Delay property 426
- Media Server Connect Timeout property 446
- Media Server Copy Advanced Backup Method 75
- Media Server storage unit setting 42
- media servers
  - adding a media server to the Alternate Restore Failover Machine list 434
  - rebooting 467
  - Restore Failover host properties 433
- Media Summary report 305
- Media Unmount Delay property 425
- Media Written report 306
- Megabytes of Memory property 373
- MEM\_USE\_WARNING 495
- Microsoft Volume Shadow Copy Service (VSS) 367
- mntfs file system, excluding from backup 191
- monitoring
  - NetBackup processes 331
- monitoring NetBackup processes 331
- monthly backups, scheduling 124
- mount points 82
- Move Backup Job From Incomplete State to Done State property 360
- Move Backup Job from Incomplete to Done State property 76
- move detection 88
- Move Job From Incomplete State to Done State
  - property 360
- Move Job From Incomplete State to Done

- State property 360
- Move Restore Job from Incomplete State to Done State
  - interaction with Checkpoint Restart 518
- Move Restore Job From Incomplete State to Done State property 360
- MS-Exchange policy type 70
- MS-SharePoint policy type 71
- MS-SQL-Server policy type 70
- MS-Windows-NT policy type 70
- MTF1 format, allow overwrite 423
- mtime 176
- multiple copies
  - fail all copies 112
  - parent and child jobs 319
  - setting 110
  - using Checkpoint Restart 75
- multiple data streams
  - allowing 92, 94
  - parent and child jobs 319
  - tuning 94
- multiple file layout for NetBackup
  - catalogs 215
- multiplexing (MPX)
  - and synthetic backups 204
  - set for schedule 117
  - use with Enable Block Sharing 40
- multistreaming and synthetic backups 204
- Must Use Local Drive property 411

## N

- named data streams, disabling the restore of 177
- naming conventions 29
- NAS filers 38
- NBDB.db
  - configuration entry 545
  - in catalog 216
  - installation overview 538
  - relocating 540
- NBEMM (NetBackup Enterprise Media Manager) 325
- nbemmcmd command 442
- nbj.conf 491
- NBJAVA\_CLIENT\_PORT\_WINDOW 495
- NBJAVA\_CONNECT\_OPTION 496
- NBJM (NetBackup Job Manager) 325, 327, 421
- nbmail.cmd script 417

- NBNOS (NetBackup Notification Service) 326
- NBPEM (NetBackup Policy Execution Manager) 326, 421
- NBRB (NetBackup Resource Broker) 326, 421
- NBSL (NetBackup Service Layer) 326
- NBSVCMON (NetBackup Monitor Service) 327
- NBU-Catalog policy type 193
- NBVAULT (NetBackup Vault Manager) 326
- NCR-Teradata policy type 70
- NDMP 38, 45, 70, 263, 412
- NDMP Host storage unit setting 43
- NDMP license 564
- NDMP storage units 24
- NearStore
  - storage unit
    - properties 570
  - storage units 564
    - authenticating media servers 568
    - SnapVault schedules 565
- NearStore Server storage unit setting 570
- NearStore storage units 24, 40, 43, 56, 570
  - disk consumption 574
  - logging information 575
- NetBackup
  - client service 429
  - request service port (BPRD) 429
- NetBackup Access Control (NBAC)
  - authorizing NetBackup-Java users 486
  - Current NBAC User 19
- NetBackup Client Service (INETD) 325
- NetBackup Compatibility Service (BPCOMPATD) 325
- NetBackup Database Manager (BPDBM) 325
- NetBackup for MS-Exchange 180
- NetBackup Job Manager (NBJM) 325, 421
- NetBackup Monitor Service (NBSVCMON) 327
- NetBackup Notification Service (NBNOS) 326
- NetBackup Operations Manager 326
- NetBackup Policy Execution Manager (NBPEM) 326, 421
- NetBackup Request Manager (BPRD) 326
- NetBackup Request Service Port (BPRD)
  - property 429

- NetBackup Resource Broker (NBRB) 326, 421
- NetBackup Service Layer (NBSL) 326
- NetBackup Vault Manager (NBVAULT) 326
- NetBackup Volume Manager (VMD) 326
- NetBackup volume pool 73, 301
- NETBACKUP\_RELATIONAL\_DATABASE\_FILES directive 231, 243
- NetBackup-Java, set up for 4
- NetWare Client
  - host properties
    - Backup Migrated Files 371
    - Keep Status of User-directed Backups, Archives, and Restores 371
    - Uncompress Files Before Backing Up 371
- NetWare NonTarget clients 387
- NetWare policy type 70
- Network
  - host properties
    - Announce DHCP interval 429
    - NetBackup Client Service Port (BPCD) 429
    - NetBackup Request Service Port (BPRD) 429
  - mask for VxSS host or domain 346
- Network Appliance (NetApp) 563
- Network Attached Storage (NAS) 24
- network drives, backing up 79
- NEW\_STREAM, file list directive 185
- nonactive media 305
- none of the files in the file list exist (NetBackup status message) 154
- None volume pool 73
- non-reserved ports 450
- nonroot administration for specific applications 490
- number of drives, setting for storage units 41

## O

- obsolescence date 535
- On Demand Only storage unit setting 43, 55, 570
- online\_util directory for VSP logging messages 455
- Open File Backup properties 430
- Open Transaction Manager (OTM) properties 430

- Operator's Email Address property 357
- optical devices 111, 263
- Oracle policy type 70
- OTM (see Open Transaction Manager)
- Override Policy
  - Storage Unit setting 114
  - Volume Pool setting 114
- Overwrite Existing Files 176

## P

- parent jobs 92, 318, 320
  - in Activity Monitor Jobs tab 318
  - Limit Jobs per Policy setting 77
  - parent\_end\_notify script 318
  - parent\_start\_notify script 318
- parent\_end\_notify script 318
- parent\_start\_notify script 318
- path setting (Lotus Notes) 422
- pathname
  - catalog backup to disk 240
  - rules for policy file list 162
- PBX (VERITAS Private Branch Exchange) 530, 532
- PC NetLink files 169
- peername
  - files 506
  - of client 505
- Perform Default Search for Restore property 381
- Perform Incrementals Based on Archive Bit 104
- Perform Snapshot Backups 110
- performance
  - improving Java applications 499
- permission to change NetBackup properties 343
- planning
  - storage units 32
  - user schedules 139
- policies
  - activating 78
  - changing properties 62, 65, 66, 67
  - configuration wizard 61
  - creating policy for Vault 195
  - deactivating 78
  - example 142
  - overview 59
  - planning 142
  - setting priority 78

- storage unit 71
- user 140
- user schedules 139
- volume pool setting 73
- Policy Execution Manager
  - Logging property 421
- policy type
  - AFS 70
  - DataStore 69
  - DB2 70
  - FlashBackup 71
  - FlashBackup Windows 71
  - Informix 71
  - Lotus-Notes 70
  - MS-Exchange 70
  - MS-SharePoint 71
  - MS-SQL-Server 70
  - MS-Windows-NT 70
  - NBU-Catalog 70
  - NCR-Teradata 70
  - NDMP 70
  - NetWare 70
  - Oracle 70
  - SAP 71
  - Split-Mirror 71
  - SQL-BackTrack 71
  - Standard 70
  - Sybase 71
  - Vault 70
  - Vault Catalog Backup 101
- Policy Update Interval property 256, 414
- Port Ranges
  - host properties
    - Client Port Window 431
    - Client Reserved Port Window 432
    - Server Port Window 432
    - Server Reserved Port Window 432
    - Use OS selected non reserved port 431, 432
    - Use Random Port Assignments 431
- ports
  - allowing operating system to select non reserved port 431, 432
  - non-reserved 450
- power down NetBackup servers 466
- preprocess interval 186
- primary copy
  - becoming a 263
  - changing 258

- definition 263
- promoting to 259
- print
  - job list information 314
  - license key 470
- Priority of Duplication Jobs 113
- Priority of Relocation Jobs Started from this Schedule 52
- Priority of Restore Jobs property 415
- priority, for a policy 78
- Private Branch Exchange (veritas\_pbx) 530, 532
- Problems report 297, 359
- proc file system, excluding from
  - backups 191
- Process Busy Files property 357
- processes
  - monitoring 331
  - show active 467
- properties
  - changing on multiple hosts 342
  - overview 340
  - viewing 340

## Q

- Quiescent wait time 461

## R

- random ports, setting on server 431
- raw partitions 82
  - backing up 76, 99, 164
  - backups on UNIX 172, 173
  - Follow NFS policy attribute 81
  - restoring 164
- rebooting NetBackup servers 467
- redirected restores 173
- Reduce Fragment Size storage unit
  - setting 43, 571
- registry, backup/restore 165
- relocation schedule 98, 108, 112, 117, 118
- Remote Administration Console 481
- remote systems, administering 479
- reports
  - All Log Entries report 298
  - Client Backups report 296
  - description of utility in Administration Console 290
  - Images on Media report 303
  - Media Contents report 302
  - Media List report 299

- Media Log Entries report 305
- Media Summary report 305
- Media Written report 306
- Problems report 297
- running a report 291
- settings for building a report 293
- Status of Backups report 295
- using the Troubleshooter 307
- Reset File Access Time property 372
- Resource Broker Logging property 421
- restarting jobs 313
- Restore Failover
  - host properties
    - Alternate Restore Failover Machines list 433
- Restore job
  - resuming 518
  - suspending 518
- Restore Retries
  - interaction with Checkpoint Restart 518
  - property 447
- restores
  - alternate client 506
  - directed from the server 504
  - raw partition 164
  - reducing search time 283
  - registry on Windows clients 165
  - server independent 522
  - setting client permissions 513
  - symbolic links on UNIX 169
  - System State 519
  - to alternate clients 505
- restoring files to alternate hosts 433
- resuming suspended jobs 314
- Retain Snapshots for Instant Recovery 110
- retention levels
  - default 116
  - for archiving catalogs 250
- retention periods
  - caution for setting 139
  - changing 437
  - guidelines for setting 115
  - mixing on media 117
  - precautions for setting 116
  - redefining 436
  - setting 115
  - user schedule 139
- Retries Allowed After Runday policy
  - setting 108

- Retry Count property 358
- retry restores, setting 447
- Rmed media type 304
- Robust Logging 419
- RS-MTF1 format, allow overwrite 423

## S

- SAP policy type 71
- Schedule Backup Attempts property 93, 414
- schedules
  - adding to policy 97
  - backups on specific dates 122
  - catalog backup 241
  - examples of automatic 126
  - frequency 108
  - how calendar scheduling interacts with daily windows 125
  - monthly backups 124
  - naming 98
  - not combining calendar-based and frequency-based 97, 108, 109
  - overview 63
  - priority 110
  - retention level defaults 116
  - retention period
    - guidelines 115
  - retention periods
    - setting 115
  - setting backup times 119
  - specify multiplexing 117
  - storage unit 114
  - type of backup 98
  - user backup or archive 139
  - volume pool 114
  - weekly backups 123
- scratch volume pool 74
- scripts 318
  - bpdbjobs example 335
  - bps 467
  - goodies 522
- SeCuRiTty.nnnn files 512
- sendmail 417
- server
  - administration, backlevel 8
  - directed restore 504
  - host properties 439
  - independent restores 433, 522
  - power down 466
  - rebooting 466

- server list definition 439
  - Server Port Window property 432
  - Server Reserved Port Window property 432
  - Server Sends Mail property 451
  - setconf.bat file 491
  - Shadow Copy
    - Components directive 165, 182
    - using 367
  - SharePoint 2003 host properties 443
  - shut down NetBackup daemons 466
  - single file layout for NetBackup catalogs 214
  - single file restore program, FlashBackup 169
  - Single-Instance Storage (SIS) 76, 386
  - SnapVault
    - storage units 564
  - SnapVault storage units 24, 40, 43, 46, 56
    - license on 568
  - Solaris
    - 9 extended attributes 169
  - Source Copy Number 255
  - Split-Mirror policy type 71
  - SQL-BackTrack policy type 71
  - Staging Schedule storage unit setting 45
  - Standard policy type 70
  - start up NetBackup daemons 467
  - status 41 372
  - status codes, NetBackup
    - 71 154
  - Status of Backups report 295
  - Storage Device storage unit setting 45
  - Storage Migrator 44, 73, 535
  - storage unit groups 55, 56, 57
  - Storage Unit Name setting 45
  - storage unit selection within a storage unit
    - group 56
  - Storage Unit Type setting 45, 570
  - storage units
    - adding Media Manager type 28
    - adding NDMP type 38
    - any available 72
    - Basic Disk 24, 40
    - changing server to manage 474
    - creating 27
    - creating a disk staging unit 48
    - disk type, definition 24
    - example Media Manager type 32
    - for policy 71
    - for schedule 114
    - management window 26
    - Media Manager type, definition 24
    - naming conventions 29
    - NDMP 24, 38
    - NearStore 24, 40, 43, 56, 570
    - optical devices 111, 263
    - QIC drive type 111, 263
    - rules for Media Manager type 31
    - SnapVault 24, 40, 43, 56
    - streaming (see Allow Multiple Data Streams setting)
    - subnets
      - and bandwidth limiting 354
    - Sun PC NetLink 169
    - suspending a job 314
    - Sybase Adaptive Server Anywhere (ASA)
      - default password 543
      - in NetBackup installation 216
      - management of 546
      - starting/stopping ASA service 546
      - use in NetBackup 537
    - Sybase policy type 71
    - symbolic links
      - included in backup selection list 160
      - UNIX 169
    - synthetic backups
      - and encryption 206
      - component images 201
      - cumulative incremental 202
      - full 201
      - logs produced during 208
      - no NetBackup change journal
      - support 379
      - recommendations for running 204
      - schedules 107
      - using Checkpoint Restart 76
  - System State
    - backups 76
    - directive 181
    - restoring 519
- T**
- tar format, allow overwrite 423
  - TCP Level Logging property 377
  - temporary staging area 36, 571
  - Temporary Staging Area storage unit
    - setting 46, 113
  - third-party copies 263
  - Third-Party Copy Device Advanced Backup Method 75



- Time Overlap property 380
- Timeouts
  - host properties
    - Backup End Notify Timeout 446
    - Backup Start Notify Timeout 444
    - Client Connect Timeout 444
    - Client Read Timeout 445
    - File Browse Timeout 444
    - Media Mount Timeout 445
    - Media Server Connect Timeout 446
    - Use OS Dependent Timeouts 445
- tmpfs file system, excluding from
  - backup 191
- To IP Address property 355
- tpclean 323
- tpext utility 551
- tpformat 323
- tpreq 323
- Transfer Throttle storage unit setting 46
- traversing directories to back up a file 394
- Troubleshooter
  - using in Activity Monitor 314
  - using in Reports application 307
- True Image Restoration (TIR)
  - configuration 88
  - Error code 136 208
  - length of time to keep information 359
  - move detection 88
  - no NetBackup change journal
  - support 379
  - pruning information 207
  - with move detection 379
  - with Move Detection property 207
- U**
  - uncompress
    - client records 287
    - NetBackup catalogs 287
  - Uncompress Files Before Backing Up
    - property 371
  - unfreeze media 300
  - unified logging 418
  - Universal
    - host properties
      - Allow Server File Writes 450
      - Use Preferred Group for Enhanced Authorization 449
  - Universal Settings properties 447
  - UNIX Client
    - host properties
      - Add to All 376
      - Do Not Compress Files Ending With 375
      - Do Not Reset File Access Time 372
      - Keep Status of User-directed Backups, Archives, and Restores 372
      - Megabytes of Memory 373
      - primary node in tree 452
  - UNSET, file list directive 189
  - UNSET\_ALL, file list directive 190
  - unsuspend media 301
  - usbdevfs file system, excluding from
    - backup 191
  - Use Alternate Read Server 53, 113
  - Use Case Sensitive Exclude List host
    - property 387
  - Use Change Journal in Incrementals
    - property 378
  - Use Direct Access Recovery for NDMP
    - Restores property 412
  - Use Legacy DES Encryption property 383
  - Use Non Reserved Ports property 397, 399
  - Use OS Dependent Timeouts property 445
  - Use Preferred Group for Enhanced Authorization property 449
  - Use Random Port Assignments
    - properties 431
  - Use Reserved Ports property 397, 399
  - Use Specified Network Interface
    - property 448
  - Use VxFS File Change Log for Incremental Backups property 373
  - USE\_CTIME\_FOR\_INCREMENTALS 536
  - USE\_NBJAUTH\_WITH\_ENHAUTH 496
  - user
    - archive backups 100
    - backups 100
    - schedules
      - planning 139
  - User Directed Timeout property 381
- V**
  - Vault
    - backup type 101
    - catalog archiving 253
    - designating duplicate as the primary 258
    - disaster recovery data 194
    - Logging property 420

- Maximum Vault Jobs host property 416
- parent and child jobs 319
- policy
  - creating 195
  - type 70
- Vault Recovery Report 194
- verifying backup
  - images 257
  - selections list 160
- VERITAS Authentication Service (VRTSat) 327
- VERITAS Authorization Service (VRTSaz) 327
- VERITAS Private Branch Exchange (NBJM) 327
- VERITAS Private Branch Exchange (veritas\_pbx) 530, 532
- VERITAS Products properties 453
- veritas\_pbx (VERITAS Private Branch Exchange) 530, 532
- version file xxxv
- view properties of a license key 471
- VMD (NetBackup Volume Manager) 326
- vmphyinv 323
- vnetd
  - Only property (for selection of ports) 399
  - VERITAS Network Daemon 397, 398
- VNETD\_PORT 492
- Volume Manager (VxVM) 172
- volume pools
  - CatalogBackup 73, 301
  - DataStore 73, 301
  - for schedule 114
  - indicating one for use by a policy 73
  - NetBackup 73, 301
  - None 73
  - policy 73
  - scratch 74
- Volume Shadow Copy 181, 182
- Volume Shadow Copy Service (VSS) 367
- Volume Snapshot Provider (VSP)
  - backups using Checkpoint Restart 75
  - directory for logging messages 455
  - properties 454
  - using with databases
- volumes
  - allocation 73
  - assignments 73
  - scratch 73

- VRTSat (VERITAS Authentication Service) 327
- VRTSaz (VERITAS Authorization Service) 327
- VSP (see VERITAS Volume Snapshot Provider)
- VxFS 4.0, named data streams 174
- VxFS 4.1, File Change Log 373
- vxlogcfg command 418
- vxlogmgr command 418
- VxSS
  - domain, indicating network mask 346
  - Networks List property 345

## W

- WAFL qtree 563
  - cleaning up 564
- Wait Time Before Clearing Archive Bit property 377
- weekly backups, scheduling 123
- wildcard characters
  - escaping backslash 392
  - escaping on UNIX 168
  - in exclude files lists 392
- UNIX
  - file paths 168
- Windows clients 163
- Windows Client
  - host properties
    - Communications Buffer 380
    - General Level Logging 377
    - Incrementals Based on Archive Bit 380
    - Incrementals Based on Timestamp 379
    - Keep Status of User-directed Backups, Archives, and Restores 381
    - Maximum Error Messages for Server 381
    - Perform Default Search for Restore 381
    - TCP Level Logging 377
    - Timeout Overlap 380
    - Use Change Journal in Incrementals 378
    - User Directed Timeout 381
    - Wait Time Before Clearing Archive Bit 377
- Windows Disk-Image (raw) backups 76, 164

Windows Display Console 6, 479, 483  
Windows Open File Backups  
    host properties  
        Abort Backup on Error 368  
        Disable Snapshot and Contine 369  
        Enable Windows Open File Backups  
        for this client 367  
        Global Drive Snapshot 368  
        Individual Drive Snapshot 367  
        Use Microsoft Volume Shadow Copy  
        Service (VSS) 367

    Use VERITAS Volume Snapshot  
    Provider (VSP) 367  
Windows Terminal Services 8  
wizards  
    backup policy 61  
    catalog backup 228  
    Device Configuration 27  
Working Directory property 357  
WORM media  
    retention period caution 116

