

VERITAS NetBackup™ 5.1

System Administrator's Guide, Volume I

for UNIX

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.

Copyright

Copyright © 1993-2004 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS Software, the VERITAS logo, VERITAS NetBackup, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS, the VERITAS Logo, VERITAS NetBackup Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

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

ACE 5.2A: ACE(TM) is copyrighted by Douglas C. Schmidt and his research group at Washington University and University of California, Irvine, Copyright (c) 1993-2002, all rights reserved.

IBM XML for C++ (XML4C) 3.5.1: Copyright (c) 1999,2000,2001 Compaq Computer Corporation; Copyright (c) 1999,2000,2001 Hewlett-Packard Company; Copyright (c) 1999,2000,2001 IBM Corporation; Copyright (c) 1999,2000,2001 Hummingbird Communications Ltd.; Copyright (c) 1999,2000,2001 Silicon Graphics, Inc.; Copyright (c) 1999,2000,2001 Sun Microsystems, Inc.; Copyright (c) 1999,2000,2001 The Open Group; All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

JacORB 1.4.1: The licensed software is covered by the GNU Library General Public License, Version 2, June 1991.

Open SSL 0.9.6: This product includes software developed by the OpenSSL Project * for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)

TAO (ACE ORB) 1.2a: TAO(TM) is copyrighted by Douglas C. Schmidt and his research group at Washington University and University of California, Irvine, Copyright (c) 1993-2002, all rights reserved.



Contents

Preface	xxxiii
What Is In This Manual	xxxiii
Getting Help	xxxiv
▼ <i>To locate the telephone support directory on the VERITAS web site</i>	xxxv
▼ <i>To contact support using E-mail on the VERITAS web site</i>	xxxv
NetBackup Manuals	xxxv
Glossary	xxxvi
▼ <i>To access the NetBackup online glossary</i>	xxxvii
Accessibility Features	xxxvii
Using the Keyboard to Navigate in NetBackup	xxxvii
Navigating in a NetBackup Tree View	xxxviii
Using Accelerator Keys	xxxix
Using Mnemonic Keys	xxxix
Using the Keyboard in Dialogs	xxxix
Accessing Online Documentation	xli
Conventions	xli
Chapter 1. Introduction to NetBackup	1
Overview	1
NetBackup Administration Interfaces	4
NetBackup Administration Console Setup	4
Running the NetBackup-Java Interface on a UNIX System	5
▼ <i>To prepare a CDE (Common Desktop Environment) for NetBackup-Java interfaces</i>	5



▼ <i>To start the NetBackup-Java Administration Console on a NetBackup-Java capable UNIX system 6</i>	
Running the Java-Based Windows Display Console	7
▼ <i>To start the Windows Display Console</i>	7
Administering Remote Servers	8
Using the NetBackup Administration Console	9
User Backups, Archives, and Restores	10
Activity Monitor	10
NetBackup Management	10
Reports	10
Policies	10
Storage Units	12
Catalog	12
Host Properties	13
Media and Device Management	13
Access Management	13
NetBackup Configuration Wizards	13
Menus	14
File Menu	14
Edit Menu	15
▼ <i>To Find using the Advanced tab</i>	17
View Menu	17
▼ <i>To Filter using the Advanced tab</i>	21
Actions Menu	21
Help Menu	21
▼ <i>To invoke the pop-up menu on Solaris X86 systems</i>	22
Standard and User Toolbars	23
Customizing the Administration Console	23
Configuring NetBackup Without Wizards	24



Chapter 2. Managing Storage Units	27
Introduction to Storage Units	28
Viewing Storage Units and Storage Unit Groups	30
Using the Device Configuration Wizard	30
▼ <i>To use the Device Configuration Wizard</i>	31
Media Manager Storage Unit Considerations	31
Allowable Media Manager Characters	32
Before Adding a Media Manager Storage Unit	33
Disk Storage Unit Considerations	36
NDMP Storage Unit Considerations	37
Disk Staging Storage Unit Considerations	38
Disk Staging: Stage I	40
Disk Staging: Stage II	41
▼ <i>To manually initiate a disk staging storage unit relocation schedule</i>	42
Disk Staging Schedule Button	43
Name	43
Final Destination Storage Unit	43
Final Destination Volume Pool	44
Cleaning the Disk Staging Storage Unit	44
Maintaining Storage Units	45
Creating a New Storage Unit	45
▼ <i>To create a storage unit from the Actions menu</i>	45
▼ <i>To create a storage unit by copying an existing storage unit</i>	45
Changing Storage Unit Properties	46
▼ <i>To change storage unit properties</i>	46
Deleting Storage Units	46
▼ <i>To delete storage units</i>	46
Storage Unit Properties	47
Absolute Pathname to Directory	48
Density	48



Disk Staging Relocation Schedule	49
Maximum Concurrent Drives Used for Backup	49
Maximum Concurrent Jobs	49
Maximum Fragment Size	50
Maximum Multiplexing per Drive	51
Media Server	51
New Media Server (... Button)	51
NDMP Host	51
On Demand Only	51
Robot Number	52
Robot Type	52
Storage Device	52
Storage Unit Name	52
Storage Unit Type	53
Configuring Drive Availability Checking	53
Interval Between Status Checks	53
Drive Count Timeout	53
Requeuing Jobs If Required Storage Units are Unavailable	54
Creating and Changing Storage Unit Groups	54
▼ <i>To create a storage unit group</i>	54
▼ <i>To change a storage unit group</i>	55
▼ <i>To delete a storage unit group</i>	56
Chapter 3. Managing Backup Policies	57
Using the Policies Utility	58
Tree and Detail Views	58
Menus	58
Actions Menu	59
Standard and User Toolbars	60
Introduction to Backup Policies	60



General Attributes on the Attributes Tab	60
Schedules on the Schedules Tab	60
Client List on the Clients Tab	61
Backup Selections on the Selections Tab	61
Configuring Backup Policies	61
▼ <i>To create a policy using the wizard</i>	62
▼ <i>To create a policy without using the wizard</i>	62
Example Policies	62
Policy Planning Guidelines for Backups	63
Changing Policies	70
▼ <i>To add or change schedules in a policy</i>	70
▼ <i>To add or change clients in a policy</i>	71
▼ <i>To add or change backup selections in a policy</i>	71
▼ <i>To delete schedules, backup selections, or clients from a policy</i>	72
▼ <i>To copy and paste items</i>	72
What Type of Policy: Policy Attributes Tab	73
▼ <i>To set the general policy attributes</i>	73
Policy Type	74
Policy Storage Unit	76
Policy Storage Unit Example	76
Notes on Specifying a Storage Unit	77
Policy Volume Pool	77
Volume Pool Example	78
Notes on Volume Pools	78
Checkpoint Restart for Backup Jobs	79
Checkpoint Frequency	79
Checkpoint Restart Support	79
Checkpoint Restart for Restore Jobs	80
Limit Jobs Per Policy	81
Notes on Limit Jobs Per Policy	81



Job Priority	82
Active. Go Into Effect At	82
Follow NFS	83
Notes on Follow NFS	83
Disadvantages of Using Follow NFS	83
Advantages of Using Follow NFS Mounts	84
Cross Mount Points	84
Notes on Cross Mount Points	84
Cases That Can Require Separate Policies	84
How Cross Mount Points Policy Attribute Interacts With Follow NFS	85
Cross Mount Point Examples	85
Collect True Image Restore Information	87
Collect True Image Restore With Move Detection	87
What Happens During True Image Restores	88
Notes On True Image Restores and Move Detection	90
Compression	90
Advantages of Using Compression	91
Disadvantages of Using Compression	91
How Much Compression Can You Expect?	91
Encryption	92
Collect Disaster Recovery Information	93
Allow Multiple Data Streams	93
When to Use Multiple Data Streams	93
Keyword Phrase (Optional)	96
Advanced Client Options	97
Which Clients Will Be Backed Up: Clients Tab	98
▼ <i>To add a client to a policy</i>	98
▼ <i>To change a client list entry</i>	99
Installing Client Software on Trusting UNIX Clients	99
▼ <i>To install UNIX client software</i>	100



Installing Software on Secure UNIX Clients	101
Installing Software on Windows Clients	101
Configuring a Snapshot Method	101
Which Selections Will Be Backed Up: Backup Selections Tab	102
Creating the Backup Selections List for Standard Policies	105
▼ <i>To add or change backup selections for a Standard, Exchange, or Lotus Notes policy</i>	105
Creating the Backup Selections List for Database Policies	107
▼ <i>To create or change backup selections containing scripts for a database policy</i> ..	107
Creating the Backup Selections List for Oracle or DB2 Policies	108
▼ <i>To add templates or scripts to the Backup Selections List</i>	108
Verifying the Backup Selections List	109
▼ <i>To verify a backup selections list</i>	109
Rules for Backup File Paths	111
File-Path Rules for UNIX Clients	111
Notes on Backup Selection Lists for UNIX Clients	112
Symbolic Links to Files or Directories	113
Hard Links to Directories	114
Hard Links to Files	114
UNIX Raw Partitions	116
Backup and Restore of Extended Attribute Files and Named Data Streams ...	119
NetBackup Client, Media Server, and Master Server Versions	120
Ramifications of Backing Up Extended Attributes or Named Data Streams .	120
Restoring Extended Attributes or Named Data Streams	120
▼ <i>To disable the restore of extended attribute files and named data streams</i>	121
File-Path Rules for Microsoft Windows Clients	122
File Backups	122
Windows Disk-Image (Raw) Backup	124
Microsoft Windows Registry Backup	125
Hard Links to Files (NTFS volumes only)	126



File-Path Rules for NetWare NonTarget Clients	127
File-Path Rules for NetWare Target Clients	128
File-Path Rules for Clients Running Extension Products	129
Backup Selections List Directives: General Discussion	130
ALL_LOCAL_DRIVES Directive	130
SYSTEM_STATE Directive	131
Shadow Copy Components:\ Directive	131
Directives for Multiple Data Streams	132
Directives for Specific Policy Types	133
Backup Selections List Directives for Multiple Data Streams	133
NEW_STREAM Directive	134
ALL_LOCAL_DRIVES Directives	137
UNSET and UNSET_ALL Directive	138
Creating an Exclude List on a UNIX Client	139
Creating an Include List on a UNIX Client	143
When Will the Job Run: Schedules Tab	144
Creating or Editing a Schedule	144
▼ <i>To create or change schedules</i>	145
Schedule Attributes Tab	145
Name	146
Type of Backup	146
More on Incremental Backups	148
Synthetic Backups	154
Advantages of Using Synthetic Backups	154
Policy Considerations and Synthetic Backups	154
Two Types of Synthetic Backups	157
Synthetic Full Backups	157
Synthetic Cumulative Incremental Backups	158
Recommendations for Synthetic Backups	160
Notes on Synthetic Backups	161



Displaying Synthetic Backups in the Activity Monitor	163
Logs Produced During Synthetic Backups	163
Synthetic Backups and Directory and File Attributes	164
Instant Recovery Backups to Disk Only	164
Calendar Schedule Type	165
Retries Allowed After Runday	165
Frequency Schedule Type	165
Backup Frequency Determines Schedule Priority	165
Multiple Copies	166
▼ <i>To configure a schedule to create multiple copies during a backup</i>	167
Override Policy Storage Unit	169
Override Policy Volume Pool	169
Retention	169
Default Retention Periods	170
Precautions For Assigning Retention Periods	170
Mixing Retention Levels on Backup Volumes	171
Media Multiplexing	171
Final Destination Storage Unit	172
Final Destination Volume Pool	172
Start Window Tab	173
▼ <i>To create a window</i>	173
Exclude Dates Tab	175
▼ <i>To exclude a date from the policy schedule</i>	175
Calendar Schedule Tab	176
Schedule by Specific Dates	176
▼ <i>To schedule a task on specific dates</i>	176
Schedule by Recurring Week Days	177
▼ <i>To schedule a recurring weekly task</i>	177
Schedule by Recurring Days of the Month	178
▼ <i>To schedule a recurring monthly task</i>	178



How Calendar Scheduling Interacts with Daily Windows	179
Examples of Automatic-Backup Schedules	180
Example 1	180
Example 2	183
Example 3	188
Example 4	189
Example 5	190
Example 6	191
Considerations for User Schedules	193
Planning User Backup and Archive Schedules	193
Creating Separate Policies for User Schedules	193
Using a Specific Policy and User Schedule	195
Creating a Vault Policy	195
▼ <i>To create a Vault policy</i>	196
Performing Manual Backups	197
▼ <i>To perform a manual backup</i>	197
Chapter 4. Managing Catalogs and Images	199
Introduction to the Catalog Application	200
Catalog Backups	200
Where are the Catalog Files?	201
What Method Do I Use to Back Them Up?	201
What NetBackup Servers Can I Use?	201
What Types of Media Can I Use?	201
How Do I Know If a Catalog Backup Succeeded?	202
How Do I Restore The Catalog Backups?	202
Important Precautions to Observe	202
Configuring Catalog Backups	203
▼ <i>To configure the catalog backup using the Catalog Backup Wizard</i>	203
▼ <i>To configure the catalog backup manually</i>	204



Catalog Attributes Tab	205
Media Server	205
Last Media Used	205
Media 1 and Media 2 Areas	205
Catalog Schedule Tab	209
Recommendations	210
Catalog Files Tab	210
▼ <i>To add a pathname</i>	211
▼ <i>To change a pathname</i>	212
▼ <i>To delete a pathname</i>	212
Catalog Pathnames	212
Backing Up the Catalogs Manually	215
▼ <i>To perform the catalog backup manually</i>	216
Protecting Large NetBackup Catalogs	216
Layout of the NetBackup Catalogs	216
Catalog Backup and Restore Concepts	217
Multiple-Tape Catalog Backups	217
Multiple-Tape Catalog Restores	218
Setting up Multiple-Tape NetBackup Catalog Backups	218
▼ <i>To define a NetBackup policy for catalog backups</i>	219
▼ <i>To configure the NetBackup catalog backups</i>	219
Create a Shell Script to Initiate the Backups	220
How To Initiate a Multiple-Tape Catalog Backup	221
Managing the NetBackup Catalogs	222
About the Binary Catalog Format	222
Catalog Conversion Utility	222
Binary Catalog File Limitations	222
Determining Catalog Space Requirements	223
▼ <i>To estimate the disk space required for a catalog backup</i>	223
File Size Considerations	225



Compressing the Image Catalog	226
Uncompressing the Image Catalog	227
▼ <i>To uncompress client records</i>	228
Moving the NetBackup Image Catalog	228
▼ <i>To move the NetBackup image catalog</i>	228
Catalog Archiving	230
Examining the Catalog Image	230
Image Files	230
Image .f Files	230
Catalog Archiving Overview	231
Creating a Catalog Archiving Policy	232
Policy Name	232
Inactive Policy	232
Type of Backup	232
Retention Level Setting	232
Schedule	233
Files	233
Clients	233
Catalog Archiving Commands	233
Create a Catalog List with <code>bpcatlist</code>	234
Back Up the Catalog with <code>bpcatarc</code>	234
Remove the Catalog with <code>bpcatrm</code>	235
Restore the Catalog with <code>bpcatres</code>	235
Recommendations for Using Catalog Archiving	235
Using Vault with the Catalog Archiving Feature	235
Browsing Offline Catalog Archive	235
Extracting Images from the Catalog Archives	236
▼ <i>To extract images from the catalog archives based on a specific client</i>	236
Reduce Restore Times by Indexing the Image Catalog	236
Catalog Index Examples	237



Catalog Index Space Requirements	237
Disabling Catalog Indexing	238
Searching for Backup Images	238
Messages Pane	240
Verifying Backup Images	240
▼ <i>To verify backup images</i>	241
Duplicating Backup Images	241
▼ <i>To duplicate backup images</i>	243
Inline Tape Copy Jobs	245
Promoting a Copy to a Primary Copy	246
▼ <i>To promote a backup copy to a primary copy</i>	246
▼ <i>To promote many copies to a primary copy</i>	247
▼ <i>To promote a backup copy to a primary copy using bpduplicate</i>	247
Expiring Backup Images	248
▼ <i>To expire a backup image</i>	248
Importing NetBackup or Backup Exec Images	249
Importing Expired Images	249
Importing Images from Backup Exec Media	249
Host Properties for Backup Exec	250
More on vmphyinv and bephyinv	251
Differences Between Importing, Browsing and Restoring Backup Exec and Net-Backup Images	252
Importing Images	253
▼ <i>To initiate an import – Phase I</i>	253
▼ <i>To import backup images – Phase II</i>	255
Viewing Job Results	256
▼ <i>To view or delete a log file</i>	256



Chapter 5. Viewing NetBackup Reports	257
Introduction to the Reports Application	258
▼ <i>To run a report</i>	258
Menu Bar	259
Reports Window	260
Report Toolbar	260
Report Contents Pane	260
Shortcut Menus	260
Reports Settings	261
Date/Time Range	261
Client	261
Media Server	261
Job ID	261
Media ID	262
Volume Pool	262
Verbose Listing	262
Run Report	262
Stop Report	262
NetBackup Report Types	263
Status of Backups Report	263
Client Backups Report	264
Problems Report	265
All Log Entries Report	266
Media Lists Report	267
Media Contents Report	270
Images on Media Report	271
Media Logs Report	273
Media Summary Report	273
Media Written Report	274
Using the Troubleshooter Within Reports	275



▼ <i>To run Troubleshooter within Reports</i>	275
---	-----

Chapter 6. Monitoring NetBackup Activity	277
---	------------

Introduction to the Activity Monitor	278
Menu Bar	279
▼ <i>To view specific column heads in the Details pane</i>	281
▼ <i>To monitor the detailed status of selected jobs</i>	281
▼ <i>To delete completed jobs</i>	281
▼ <i>To cancel uncompleted jobs</i>	281
▼ <i>To suspend a restore or backup job</i>	282
▼ <i>To resume a suspended or incomplete job</i>	282
▼ <i>To restart a completed job</i>	282
▼ <i>To export Activity Monitor data to a text file</i>	282
▼ <i>To run Troubleshooter within the Activity Monitor</i>	283
Shortcut Menus	283
Activity Monitor Toolbar	283
Status Bar	284
Setting Activity Monitor Options	284
Jobs Tab	285
Viewing Job Details	285
Daemons Tab	289
▼ <i>To monitor NetBackup daemons</i>	290
▼ <i>To start or stop a daemon</i>	291
Processes Tab	291
▼ <i>To monitor NetBackup processes</i>	291
Media Mount Errors	292
Queued Media Mount Errors	292
Cancelled Media Mount Errors	292
Managing the Jobs Database	293
Retaining Job Information in the Database	293



Changing the Default on a Permanent Basis	293
Changing the Default Temporarily	294
bpdjobs Debug Log	296
Customizing bpdjobs Output	296
Chapter 7. Configuring Host Properties	297
Introduction to Host Properties	298
Menu Bar	299
Viewing Host Properties	300
▼ <i>To view master server, media server, or client properties</i>	300
Changing Host Properties	300
Interpreting the Initial Settings	300
Selecting Multiple Hosts	302
▼ <i>To simultaneously change the properties on multiple hosts</i>	302
Required Permissions	303
Master Server, Media Server, and Client Host Properties	304
Access Control Properties	304
VERITAS Security Services (VxSS)	304
VxSS Tab within Access Control Properties Dialog	304
VxSS Networks List	305
Add Button	305
Remove Button	306
Authentication Domain Tab within Access Control Properties Dialog	307
Add Button	307
Remove Button	308
Authorization Service Tab within Access Control Properties Dialog	308
Host Name	309
Customize the Port Number of the Authorization Service	309
Authorization Properties	310
User	310



Host	310
Domain\Group	310
Group/Domain Type	311
User must be an OS Administrator	311
Backup Exec Tape Reader Properties	312
Add Button	312
GRFS Advertised Name	312
Actual Client Name	313
Actual Path	313
Change Button	313
Remove Button	313
Bandwidth Properties	314
Bandwidth Throttle Setting for the Range of IP Addresses	314
From IP Address	314
To IP Address	314
Bandwidth	314
Bandwidth Throttle Settings List	315
Add Button	315
Remove Button	315
Busy File Properties	316
Working Directory	316
Operator's E-mail Address	316
Process Busy Files	316
File Action File List	317
Add Button	317
Add to All Button	317
Remove Button	317
Busy File Action	317
Retry Count	317
Client Attributes Properties	318



Allow Client Browse	318
Allow Client Restore	318
Clients List	318
Add Button	319
Remove Button	319
General Tab	319
Connect on Non-reserved Port	319
BPCD Connect-back	320
Maximum Data Streams	320
Browse and Restore Ability	320
Free Browse	321
Client Name Properties	322
Client Name	322
Client Settings (NetWare) Properties	323
Back Up Migrated Files	323
Uncompress Files Before Backing Up	323
Keep Status of User-directed Backups, Archives, and Restores	323
Client Settings (UNIX) Properties	324
Locked File Action	324
Reset File Access Time to the Value Before Backup	324
Megabytes of Memory to Use for File Compression	324
Do Not Compress Files Ending With	325
Add Button	325
Add to All Button	325
Remove Button	325
Client Settings (Windows) Properties	326
General Level Logging	326
TCP Level Logging	326
Wait Time Before Clearing Archive Bit	326
Use Change Journal in Incrementals	327



Incrementals Based on Timestamp	328
Incrementals Based on Archive Bit	328
Time Overlap	329
Communications Buffer	329
User Directed Timeout	329
Maximum Repetitive Error Messages for Server	330
Keep Status of User-directed Backups, Archives, and Restores	330
Perform Default Search of Backup Images for Restore	330
Encryption Properties	331
Encryption Permissions	331
Enable Encryption	331
Use Standard Encryption	332
Client Cipher	332
Use Legacy DES Encryption	332
Encryption Strength	332
Encryption Libraries	332
Encryption Key File	333
Exchange Properties	334
Mailbox for Message Level Backup and Restore	334
Enable Single Instance Backup for Message Attachments	334
Exclude Lists Properties	335
Use Case Sensitive Exclude List	335
Exclude List	335
Exceptions to the Exclude List	335
Add Buttons	336
Add to All Buttons	336
Remove Buttons	337
Shared Fields in Exclude Lists	337
Policy	337
Schedule	338



Files/Directories	338
Exclude Lists for Specific Policies or Schedules	338
▼ <i>To create an exclude or include list for a specific policy</i>	338
Syntax Rules for Exclude Lists	340
Traversing Excluded Directories	342
Firewall Properties	344
Host	344
Add Button	344
Add to All Button	344
Remove Button	344
BPCD Connect-back	344
Ports	345
Daemon Connection Port	345
Minimum Required Connections	347
▼ <i>To set up vnetd between a server and a client</i>	355
▼ <i>To set up vnetd between servers</i>	355
▼ <i>To enable logging for vnetd</i>	356
Example Setup for Using the vnetd Port	356
GDM (Global Data Manager) Properties	358
Dashboard Port Window	358
Use OS Selected Non-reserved Port	358
General Server Properties	359
Delay on Multiplexed Restores	359
Re-read Interval for Available Drives	359
Must Use Local Drive	359
Use Direct Access Recovery for NDMP Restores	360
Allow Block Incrementals	360
Use Media Host Override	360
▼ <i>To force restores to go to a specific server</i>	361
Global Attributes Properties	362



Wakeup Interval	362
Schedule Backup Attempts	362
Status Report Interval	362
Maximum Jobs per Client	363
Maximum Backup Copies	364
Compress Catalog Interval	364
Keep Logs	364
Delete Vault Logs	364
Keep True Image Restoration (TIR) Information	364
Move Restore Job From Incomplete State to Done State	365
Move Backup Job from Incomplete State to Done State	365
Administrator's E-mail Address	366
Logging Properties	367
Global Logging Level	367
BPSCHED Logging Level	368
BPBRM Logging Level	368
BPTM Logging Level	368
BPDM Logging Level	368
BPRD Logging Level	368
BPDBM Logging Level	368
Vault Logging Level	368
Lotus Notes Properties	369
Path	369
INI File	369
Executable Directory	369
Media Properties	370
Allow Media Overwrite	370
Allow Multiple Retentions Per Media	371
Allow Backups to Span Media	371
Enable SCSI Reserve/Release	372



Enable Standalone Drive Extension	372
Enable Job Logging	372
Media ID Prefix (Non-robotic)	372
Media Unmount Delay	372
Media Request Delay	373
NetWare Client Properties	373
Network Properties	374
NetBackup Client Service Port (BPCD)	374
NetBackup Request Service Port (BPRD)	374
Announce DHCP Interval	374
Open File Backup (NetWare Client) Properties	375
Enable Open File Backup During Backups	375
OTM Properties	375
Port Ranges Properties	376
Use Random Port Assignments	376
Client Port Window	376
Client Reserved Port Window	377
Server Port Window	377
Server Reserved Port Window	377
Restore Failover Properties	378
Alternate Restore Failover Machines List	378
Add Button	379
▼ <i>To add a media server to the Alternate Restore Failover Machine list</i>	379
Change Button	379
Remove Button	379
Retention Periods Properties	380
Value	380
Units	380
Retention Periods List	380
Schedules List	380



Impact Report Button	381
▼ <i>To change a retention period</i>	381
Note on Redefining Retention Periods	382
SANPoint Control (SPC) Properties	383
SPC Server Name	383
SPC WebServer Name	383
Specifying the Browser Path	384
More About SANPoint Control	384
Servers Properties	385
Master Server	385
Additional Servers	385
Media Servers	385
Timeouts Properties	386
Client Connect Timeout	386
Backup Start Notify Timeout	386
File Browse Timeout	386
Use OS Dependent Timeouts	386
Media Mount Timeout	387
Client Read Timeout	387
Backup End Notify Timeout	387
Media Server Connect Timeout	388
BPTM (Drive Count) Query Timeout	388
Timeout in Job Queue	388
Requeue Active Jobs if Required Storage Unit is Unavailable	388
Requeue Scheduled Jobs if Required Storage Unit is Unavailable	388
Universal Settings Properties	389
Restore Retries	389
Browse Timeframe for Restores	389
Last Full Backup	390
Use Specified Network Interface	390



Use Preferred Group for Enhanced Authorization	391
Allow Server File Writes	392
Accept Connections on Non-reserved Ports	392
Enable Performance Data Collection	393
Client Sends Mail	393
Server Sends Mail	393
Client Administrator's E-mail	393
UNIX Client Properties	394
UNIX Server Properties	394
NFS Access Timeout	394
VERITAS Products Properties	395
VSP (Volume Snapshot Provider) Properties	396
VSP Overview	396
Logging VSP Messages	397
Cache File Volume List	398
VSP Volume Exclude List	398
Customize Cache Size	401
Initial Cache Size	401
Maximum Cache Size	402
Busy File Wait	403
Busy File Timeout	403
Using VSP with Databases	403
Windows Client Properties	405
Windows Open File Backup Properties	406
Add and Remove Buttons	406
Enable Windows Open File Backups for this Client	407
Use VERITAS Volume Snapshot Provider (VSP)	407
Use Microsoft Volume Shadow Copy Service (VSS)	407
Individual Drive Snapshot	407
Global Drive Snapshot	408



Abort Backup on Error	408
Disable Snapshot and Continue	409
Chapter 8. Managing NetBackup	411
Powering Down and Rebooting NetBackup Servers	412
▼ <i>To power down a server</i>	412
▼ <i>To reboot a NetBackup master server</i>	412
▼ <i>To reboot a NetBackup media server</i>	413
Managing Daemons	413
Displaying Active Processes with bpps	413
Starting and Stopping NetBackup and Media Manager Daemons	414
Starting NetBackup and Media Manager Daemons	414
▼ <i>To start NetBackup and Media Manager</i>	414
Stopping NetBackup and Media Manager Daemons	414
▼ <i>To stop bprd</i>	414
▼ <i>To stop all daemons</i>	414
Starting and Stopping bpdbm	415
▼ <i>To start bpdbm separately</i>	415
▼ <i>To stop bpdbm</i>	415
Administering NetBackup Licenses	416
▼ <i>To access license keys for a NetBackup server</i>	416
▼ <i>To add a new license key</i>	417
▼ <i>To delete a license key</i>	417
▼ <i>To view the properties of one license key</i>	418
▼ <i>To export the license keys</i>	418
Using the NetBackup License Utility to Administer Licenses	419
▼ <i>To start the NetBackup License Key utility</i>	419
Administering a Remote Master Server	420
Adding a NetBackup Server to a Server List	420
▼ <i>To add a server to a UNIX server list</i>	421



▼ <i>To add a server to a Windows server list</i>	423
Choosing a Remote Server to Administer	424
▼ <i>To use the Change Server command to administer a remote server</i>	424
▼ <i>To indicate a remote system upon login</i>	425
Administering through a NetBackup Client	426
If You Cannot Access a Remote Server	426
Using the NetBackup-Java Windows Display Console	427
Authorizing NetBackup-Java Users on Windows	427
Restricting Access on Windows	428
Authorizing Users for Specific Applications	429
Managing Client Restores	431
Server-Directed Restores	431
Allowing Redirected Restores	431
How NetBackup Enforces Restore Restrictions	432
Allowing All Clients to Perform Redirected Restores	432
Allowing a Single Client to Perform Redirected Restores	433
Allowing Redirected Restores of a Specific Client's Files	433
Redirected Restore Examples	434
Restoring Files and Access Control Lists	440
Restoring Files that Possess ACLs	440
Restoring Files without Restoring ACLs	440
▼ <i>To restore files without restoring ACLs</i>	441
Setting Client List and Restore Permissions	441
Adding Clients to the NetBackup Client Database	441
Setting the List and Restore Permissions	442
Examples	444
Improving Search Times by Creating an Image List	445
Set Original atime for Files During Restores	446
Checkpoint Restart for Restore Jobs	446
Suspending and Resuming a Restore Job	446



Limitations to Checkpoint Restart for Restore Jobs	447
Restoring System State	447
Important Notes on System State	448
▼ <i>To restore the System State</i>	448
Goodies Scripts	450
Server Independent Restores	450
Supported Configurations	450
Methods for Performing Server Independent Restores	453
Method 1: Modifying the NetBackup Catalogs	453
Method 2: Overriding the Original Server	454
Method 3: Automatic Failover to Alternate Server	456
Notes on Server Independent Restores	457
Configuring NetBackup Ports	458
Server and Client Connections: General Case	461
Backups	461
Restores	464
Multiplexing	468
Multiple Data Streams	468
Configuring Ports for Backups and Restores	470
Important Note on Configuring Port Limitations	471
Configuration Example	471
Administration Client Connections	473
Configuring Ports When Using an Administration Client	475
Configuration Example	477
NetBackup-Java Console Connections	480
Configuring Ports When Using the NetBackup-Java Console	481
▼ <i>To configure the master server according to the example</i>	484
▼ <i>To configure the master server to use vnetd according to the example</i>	484
Load Balancing	485
Adjust Backup Load on Server	485



Adjust Backup Load on Server Only During Specific Time Periods	485
Adjust Backup Load on Client	485
Reduce Time To Back Up Clients	485
Give Preference To a Policy	485
Adjust Load Between Fast and Slow Networks	486
Limit the Backup Load Produced By One or More Clients	486
Maximize Use of Devices	486
Prevent Backups From Monopolizing Devices	486
Using NetBackup with Storage Migrator	486
Set a Large Enough Media Mount Timeout	487
Do Not Use the RESTORE_ORIGINAL_ETIME File	487
Do Not Use the Following Client bp.conf File Settings	487
Configuring the NetBackup-Java Console	489
NetBackup-Java Administration Console Architectural Overview	489
Authorizing NetBackup-Java Users	491
Authorization File	492
Configuring Nonroot Usage	494
Authorizing Nonroot Users for Specific Applications	494
Capabilities Authorization for jbpSA	495
Runtime Configuration Options	496
BPJAVA_PORT, VNETD_PORT	496
FORCE_IPADDR_LOOKUP	496
INITIAL_MEMORY, MAX_MEMORY	498
MEM_USE_WARNING	498
NBJAVA_CLIENT_PORT_WINDOW	499
NBJAVA_CONNECT_OPTION	499
Configuration Options Relevant to jnbSA and jbpSA	500
Logging Command Lines Used by the NetBackup Interfaces	500
Customizing jnbSA and jbpSA with bp.conf Entries	500
NetBackup-Java Performance Improvement Hints	501



How do I Run a Console Locally and Administer a Remote Server?	502
How do I Make the Console Perform Even Better?	502
Is Performance Better When Remotely Displaying Back or Running Locally? ..	503
Administrator's Quick Reference	505
Index	507





Preface

This guide describes how to configure and manage the operation of VERITAS NetBackup™ Server and VERITAS NetBackup Enterprise Server for UNIX platforms and applies to all supported platforms and operating systems. 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 `/usr/opensv/netbackup/version` file.

What Is In This Manual

- ◆ Chapter 1, “Introduction to NetBackup,” contains an overview of the product.
- ◆ Chapter 2, “Managing Storage Units,” explains how to configure NetBackup to use the storage devices in your network.
- ◆ Chapter 3, “Managing Backup Policies,” explains how to configure NetBackup policies. A policy defines the backup characteristics for a group of clients that have the same or similar backup requirements.
- ◆ Chapter 4, “Managing Catalogs and Images,” explains how to manage and back up the NetBackup internal catalogs. **Catalog** is also used to search for specific backup images in order to verify, duplicate, or import backup images.
- ◆ Chapter 5, “Viewing NetBackup Reports,” explains how to run reports in order to obtain information about NetBackup activities.
- ◆ Chapter 6, “Monitoring NetBackup Activity,” explains how to monitor and control NetBackup jobs, processes, and services.
- ◆ Chapter 7, “Configuring Host Properties,” describes the Master Server, Media Servers and Client properties and how to change the settings.
- ◆ Chapter 8, “Managing NetBackup,” contains various topics regarding managing NetBackup operations.



Getting Help

VERITAS offers you a variety of support options.

Accessing the VERITAS Technical Support Web Site

The VERITAS Support Web site allows you to:

- ◆ 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

The address for the VERITAS Technical Support Web site follows:

- ◆ <http://support.veritas.com>

Subscribing to VERITAS Email Notification Service

Subscribe to the VERITAS Email notification service to be informed of software alerts, newly published documentation, Beta programs, and other services.

Go to <http://support.veritas.com>. Select a product and click “E-mail Notifications” on the right side of the page. Your customer profile ensures you receive the latest VERITAS technical information pertaining to your specific interests.

Accessing VERITAS Telephone Support

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.

▼ **To locate the telephone support directory on the VERITAS web site**

1. Open <http://support.veritas.com> in your web browser.
2. Click the **Phone Support** icon. A page that contains VERITAS support numbers from around the world appears.

Accessing VERITAS E-mail Support

▼ **To contact support using E-mail on the VERITAS web site**

1. Open <http://support.veritas.com> in your web browser.
2. Click the **E-mail Support** icon. A brief electronic form will appear and prompt you to:
 - ◆ Select a language of your preference
 - ◆ Select a product and a platform
 - ◆ Associate your message to an existing technical support case
 - ◆ Provide additional contact and product information, and your message
3. Click **Send Message**.

Contacting VERITAS Licensing

For license information call 1-800-634-4747 option 3, fax 1-650-527-0952, or e-mail amercustomercare@veritas.com.

NetBackup Manuals

The following manuals, along with the online help, comprise the NetBackup documentation set. The manuals are provided in Adobe Portable Document Format (PDF) on the NetBackup CD-ROM.

- ◆ *NetBackup Release Notes for UNIX and Windows*

Provides important information about NetBackup Server and Enterprise Server products on UNIX- and Windows-based servers, such as the platforms and operating systems that are supported and operating notes that may not be in the NetBackup manuals or the online help.
- ◆ *NetBackup Installation Guide for UNIX*



Explains how to install NetBackup Server and Enterprise Server software on UNIX-based platforms.

◆ *VERITAS Security Services Installation Guide*

Explains install and configure the VERITAS Security Services. This manual is found on the VERITAS Security Services CD-ROM.

◆ *NetBackup Media Manager System Administrator's Guide for UNIX*

Explains how to configure and manage the storage devices and media on UNIX servers running NetBackup Server and Enterprise Server. Media Manager is part of NetBackup.

◆ *NetBackup Backup, Archive, and Restore Getting Started Guide*

Explains how to use the NetBackup Backup, Archive, and Restore interface to perform basic backup and restore operations for UNIX and Windows systems.

◆ *VERITAS Security Services Administrator's Guide*

Explains how to configure and manage core security mechanisms, including authentication, protected communications, and authorization. This manual is found on the VERITAS Security Services CD-ROM.

◆ *NetBackup Vault System Administrator's Guide for UNIX and Windows*

Describes how to configure and use logical vaults and profiles to duplicate backups, perform catalog backups, eject media, and generate reports.

◆ *NetBackup Vault Operator's Guide for UNIX and Windows*

Describes procedures for sending tapes offsite, receiving tapes on site, and running reports on offsite media and vault jobs.

◆ *NetBackup Commands for UNIX*

Describes NetBackup commands and processes that can be run from a UNIX command line.

◆ *NetBackup Troubleshooting Guide for UNIX and Windows*

Provides troubleshooting information for UNIX- and Windows-based NetBackup Server and Enterprise Server, including Media Manager.

Glossary

If you encounter unfamiliar terminology, consult the NetBackup online glossary. The glossary contains terms and definitions for NetBackup and all additional NetBackup options and agents.



The NetBackup online glossary is included in the NetBackup help file.

▼ **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.

Accessibility Features

The NetBackup interface can be used by people who are vision impaired and by people who have limited dexterity. Accessibility features include the following:

- ◆ Using the Keyboard to Navigate in NetBackup
- ◆ Accessing Online Documentation

Note Text that appears in the NetBackup interface is accessible through an application programmer's interface (API) to assistive technologies such as voice or assistive device input products and to speech output products.

Using the Keyboard to Navigate in NetBackup

You can use your keyboard to navigate in the NetBackup interface:

- ◆ Press window navigation keys to move from one window element to another. For example, press **Tab** to move from one pane to another.
- ◆ Perform common actions quickly using accelerator keys. Accelerator keys let you initiate actions without first accessing a menu. For example, press **Ctrl+n** to create a new policy.
- ◆ Press mnemonic keys to select items using only the keyboard. Mnemonic keys are indicated by an underlined letter. For example, press **Alt+h** to access the **Help** menu.
- ◆ You can also use the keyboard to select control options in a dialog.



Navigating in a NetBackup Tree View

Use the following keys or key combinations to navigate through the NetBackup Console window.

Keyboard Input	Result
Tab or F6	Moves to the next (right or down) pane in the active NetBackup window.
Shift+Tab or Shift+F6	Moves to the previous (left or up) pane in the active NetBackup window.
Ctrl+Tab or Ctrl+F6	Moves to the next (right or down) NetBackup window.
Ctrl+Shift+Tab or Ctrl+Shift+F6	Moves to the previous (left or up) NetBackup window.
Plus Sign (+) on the numeric keypad	Expands the highlighted item.
Minus Sign (-) on the numeric keypad	Collapses the highlighted item.
Asterisk (*) on the numeric keypad	Expands the entire tree below the first item in the active NetBackup window.
Up Arrow	Gives focus to the next item up in the pane.
Down Arrow	Gives focus to the next item down in the pane.
Shift+Up Arrow	Selects the next item up in the pane.
Shift+Down Arrow	Selects the next item down in the pane.
Page Up	Moves to the top item visible in a pane.
Page Down	Moves to the bottom item visible in a pane.
Home	Moves to the first item (whether visible or not) in a pane.
End	Moves to the last item (whether visible or not) in a pane.
Right Arrow	Expands the highlighted item. If the highlighted item does not contain hidden items, using the Right Arrow has the same effect as using the Down Arrow .

Keyboard Input	Result
Left Arrow	Collapses the highlighted item. If the highlighted item does not contain expanded items, using the Left Arrow has the same effect as using the Up Arrow .
Alt+Right Arrow	Moves to the next (right or down) option control in the interface.
Alt+Left Arrow	Moves to the previous (left or up) option control in the interface.
Alt+Spacebar	Displays the NetBackup window menu.

Using Accelerator Keys

Accelerator keys let you use NetBackup from the keyboard, rather than using the mouse. Accelerator keys are either a single keystroke or two or more keystrokes that can be pressed in succession (rather than holding them simultaneously). If available, accelerator keys are shown to the right of the menu item they perform.

For example, to refresh the information in the window, press **F5**.

Using Mnemonic Keys

A mnemonic key is a keyboard equivalent for a mouse click that is used to activate a component such as a menu item. To select a menu item, press the **Alt** key to initiate menu pull-down mode, then press a mnemonic key to open a menu, and another mnemonic key to select a menu item.

Mnemonics are case-insensitive. Keys can be pressed either sequentially or simultaneously.

For example, to change the Master Server, press and hold the **Alt** key then press the **f** key to pull down the File menu, and press the **c** key to invoke the **Change Server** menu option.

Using the Keyboard in Dialogs

To select or choose controls that have an underlined letter in their titles, type **Alt+underlined_letter** at any time when the dialog is active. For example, typing **Alt+O** is the same as clicking the OK button in a dialog.

To move forward (right or down) from one control to the next, press **Tab**. To reverse the direction (for example, from moving right to moving left), press **Tab** and **Shift**.



To move within a list box, groups of option controls, or groups of page tabs, press the arrow key that points the direction you want to move.

Options that are unavailable appear dimmed and cannot be selected.

The following conventions are typically used in NetBackup dialogs:

- ◆ **Command buttons (also known as push buttons)**

Command buttons initiate an immediate action. One command button in each dialog carries out the command you have chosen, using the information supplied in the dialog. This button is generally labeled **OK**. Other command buttons let you cancel the command or choose from additional options.

- ◆ **Command buttons containing an ellipsis (...)**

Command buttons containing an ellipsis (...) open another dialog so you can provide more information or confirm an action. Command buttons marked with an arrow display a menu.

- ◆ **Command buttons outlined by a dark border**

A dark border around a button initially indicates the default button. Press **Enter** or the **Spacebar** at any time to choose the button with a dark border. Press **Tab** to move the keyboard focus to the next control. When you change focus to a command button, it temporarily has the dark border. If the focus is not on a control, the dark border returns to the default command button in the pane.

- ◆ **Check boxes**

Check boxes may be selected or cleared to turn an option on or off. Check boxes can have two states (checked and unchecked) or three states (checked, unchecked, and indeterminate).

Press **Tab** to move from one checkbox to another and the **Spacebar** to change the check box to the next state. Typing the mnemonic key for a check box also moves the focus to the box and changes its state.

- ◆ **Option controls (also known as radio buttons)**

Option controls are used to select only one option from a group of options. (Option buttons may represent two or three states, as checkboxes do.) Press **Tab** to move to an option button and press the **Spacebar** to initiate the option. Type the mnemonic key for an option control to move the focus to the control and select it.

- ◆ **Page series**

A series of pages are used to fit many options into a single dialog. Each page contains separate groups of controls such as check boxes or option controls. Press **Tab** to move to the name of the page, then use right and left arrows to highlight a different page name. Press **Return**.

Accessing Online Documentation

In addition to online help, NetBackup provides copies of related NetBackup manuals in Adobe Portable Document Format (PDF) on the NetBackup CD-ROM (or as an option for downloading if the release is available from the Web). For a complete list of NetBackup documents, see the NetBackup release notes.

Conventions

The following conventions apply throughout the documentation set.

Product-Specific Conventions

The following term is used in the NetBackup 5.1 documentation to increase readability while maintaining technical accuracy.

- ◆ Microsoft Windows, Windows

Terms used to describe a specific product or operating system developed by Microsoft, Inc. Some examples you may encounter in NetBackup documentation are, Windows servers, Windows 2000, Windows Server 2003, Windows clients, Windows platforms, or Windows GUI.

When Windows or Windows servers is used in the documentation, it refers to all of the currently supported Windows operating systems. When a specific Windows product is identified in the documentation, only that particular product is valid in that instance.

For a complete list of Windows operating systems and platforms that NetBackup supports, refer to the *NetBackup Release Notes for UNIX and Windows* or go to the VERITAS support web site at <http://www.support.veritas.com>.

Typographical Conventions

Here are the typographical conventions used throughout the manuals:

Conventions

Convention	Description
GUI Font	Used to depict graphical user interface (GUI) objects, such as fields, listboxes, menu commands, and so on. For example: Enter your password in the Password field.



Conventions (continued)

Convention	Description
<i>Italics</i>	<p>Used for placeholder text, book titles, new terms, or emphasis. Replace placeholder text with your specific text. For example: Replace <i>filename</i> with the name of your file. Do <i>not</i> use file names that contain spaces.</p> <p>This font is also used to highlight NetBackup server-specific or operating system-specific differences. For example: <i>This step is only applicable for NetBackup Enterprise Server.</i></p>
Code	<p>Used to show what commands you need to type, to identify pathnames where files are located, and to distinguish system or application text that is displayed to you or that is part of a code example.</p>
Key+Key	<p>Used to show that you must hold down the first key while pressing the second key. For example: Ctrl+S means hold down the Ctrl key while you press S.</p>

You should use the appropriate conventions for your platform. For example, when specifying a path, use backslashes on Microsoft Windows and slashes on UNIX. Significant differences between the platforms are noted in the text.

Tips, notes, and cautions are used to emphasize information. The following samples describe when each is used.

Tip Used for nice-to-know information, like a shortcut.

Note Used for important information that you should know, but that shouldn't cause any damage to your data or your system if you choose to ignore it.

Caution Used for information that will prevent a problem. Ignore a caution at your own risk.

Command Usage

The following conventions are frequently used in the synopsis of command usage.

brackets []

The enclosed command line component is optional.

Vertical bar or pipe (|)

Separates optional arguments from which the user can choose. For example, when a command has the following format:

command *arg1|arg2*

In this example, the user can use either the *arg1* or *arg2* variable.

Navigating Multiple Menu Levels

When navigating multiple menu levels, a greater-than sign (>) is used to indicate a continued action.

The following example shows how the > is used to condense a series of menu selections into one step:

- ❖ Select **Start > Programs > VERITAS NetBackup > NetBackup Administration Console**.

The corresponding actions could be described in more steps as follows:

1. Click **Start** in the task bar.
2. Move your cursor to **Programs**.
3. Move your cursor to the right and highlight **VERITAS NetBackup**.
4. Move your cursor to the right. First highlight and then click **NetBackup Administration Console**.





This chapter provides an introduction to NetBackup and contains the following sections:

- ◆ “Overview” on page 1
- ◆ “NetBackup Administration Interfaces” on page 4
- ◆ “NetBackup Administration Console Setup” on page 4
- ◆ “Using the NetBackup Administration Console” on page 9
- ◆ “Configuring NetBackup Without Wizards” on page 24

Overview

NetBackup provides high-performance backups and restores for a variety of platforms, including Microsoft Windows, UNIX, and NetWare systems.

Administrators can set up schedules for automatic, unattended backups for clients anywhere in the network. These backups can be full or incremental and are managed entirely by the NetBackup master server.

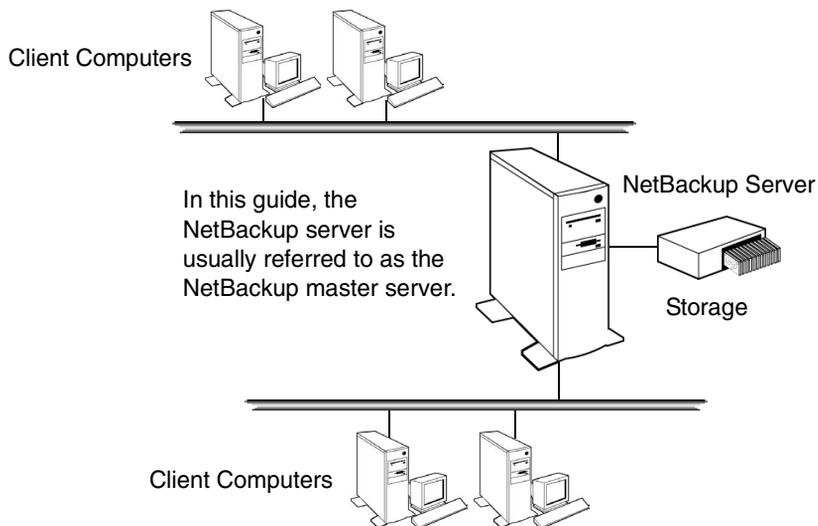
Users can start backups and restores from the computer where they are working. A user can also archive files. An archive operation backs up a file and then deletes it from the local disk if the backup is successful. Once started, user operations are managed by the NetBackup server.

NetBackup’s Media Manager software manages the media and storage devices. Robots require no intervention on the part of the administrator, operator, or the user. Standalone drives (those not in a robot) that contain appropriate media also require no intervention.



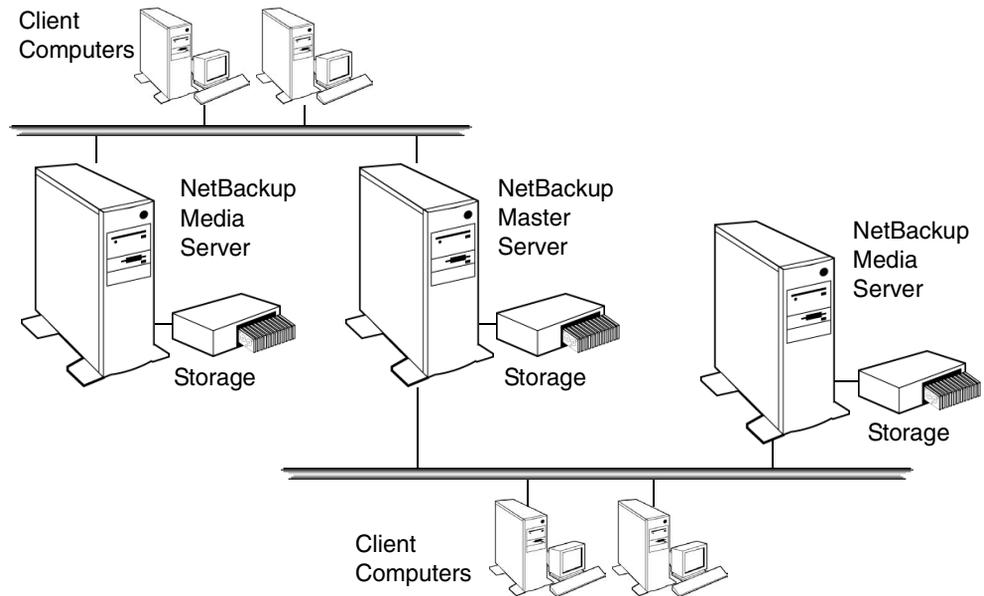
NetBackup includes both the server and client software:

- ◆ Server software is on the computer that manages the storage devices.
- ◆ Client software is on the computer whose data you want to back up. A server also has client software and can be backed up like other clients.



NetBackup servers and clients can be any one of a number of platform types as described in the data sheets and release notes for the product.

NetBackup supports both master and media servers. The master server manages the backups, archives, and restores. Media servers 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.



During a backup or archive, the client sends backup data across the network to a NetBackup server that has the type of storage specified for the client. The storage requirement is specified during NetBackup configuration (for example, 4 mm tape).

During a restore, users can browse and then select the files and directories that they want 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 4. For more details on the console, see “NetBackup-Java Administration Console Architectural Overview” on page 489.)

- ◆ 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 “Using `bpadm`” on page 183 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 the guide, *NetBackup Commands for UNIX*. 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 494.

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.

Note From NetBackup version 4.5 forward, NetBackup does not include or support the Motif interfaces: `xbpadm`, `xbpmon`, `xvadm`, and `xdevadm`. Attempting to configure NetBackup by using copies of these Motif interfaces from an earlier release will corrupt your NetBackup configuration.

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 the *NetBackup Commands for UNIX* guide.

- ◆ 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.

The startup procedures are explained below. For configuration information, see “Configuring the NetBackup-Java Console” on page 489.

Running the NetBackup-Java Interface on a UNIX System

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.

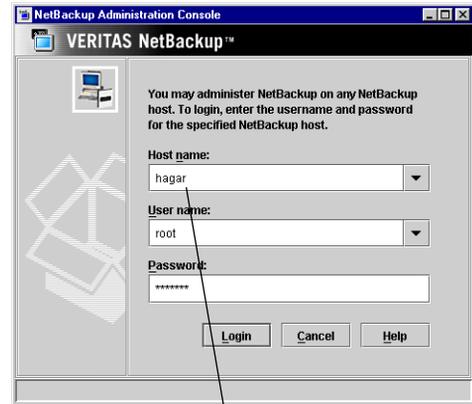
The following are general instructions for correctly setting up the focus on a CDE (Common Desktop Environment) window manager, which is the preferred window manager for NetBackup-Java applications.

- ▼ **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 scope 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

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 the *NetBackup Commands for UNIX* guide.

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.

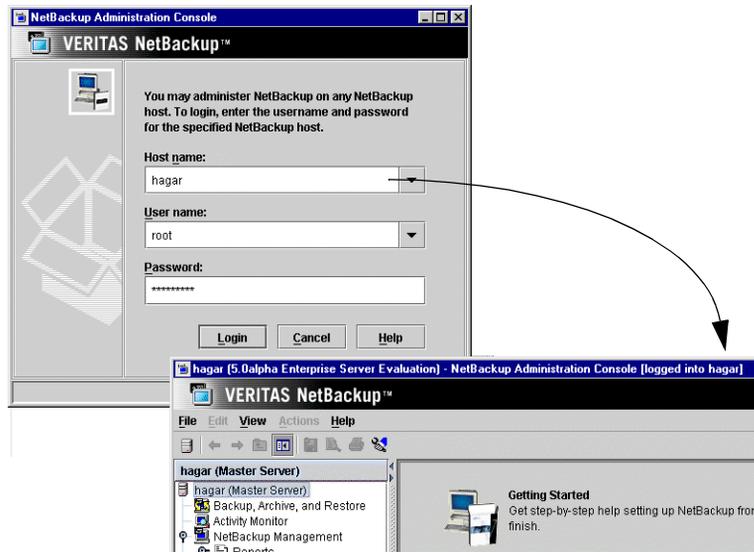
▼ To start the Windows Display Console

1. On a Windows system where the Windows Display Console is installed and configured, select **Start > Programs > VERITAS NetBackup > NetBackup-Java on host**. The *host* is the default managing NetBackup server as set during installation.
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 server's domain 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 420.

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.

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

Backup, Archive and Restore
Performs client actions for this system.

Activity Monitor
Displays information about NetBackup jobs and provides some control over the jobs.

NetBackup Management
Contains utilities for creating and viewing reports, for configuring policies, storage units, catalog backups, and a utility for configuring master server, media server, and client properties.

Media and Device Management
Contains utilities for managing the media and devices that NetBackup uses to store backups. (See the *NetBackup Media Manager System Administrator's Guide*.)

Access Management
Contents of node are viewable only by Security Administrator when VxSS components and NetBackup access control is configured.

Adjustable Split Bar

Tree view in left pane
NetBackup utilities appear as nodes.

Additional licensed utilities
Appear beneath basic NetBackup nodes.

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

The screenshot shows the NetBackup Administration Console window titled "hagar (5.0Trunk Enterprise Server Evaluation) - NetBackup Administration Console [logged into hagar]". The interface includes a menu bar (File, Edit, View, Actions, Help), a tree view on the left pane showing a hierarchy of utilities under "hagar (Master Server)", and a details pane on the right with several configuration wizards: "Getting Started", "Configure Storage Devices", "Configure Volumes", "Configure the Catalog Backup", and "Create a Backup Policy". An adjustable split bar is visible between the tree and details panes.

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 4.

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



User Backups, Archives, and Restores

To perform backups and archives for this system, and restores for this system and other clients, open the client interface by clicking on the **Backup, Archive, and Restore** button in the toolbar.

Users can back up, archive, and restore files, directories, and raw partitions that reside on their own client computer. A user can restore files at any time but can back up and archive only during the time periods that the administrator defines with the schedules. Users can also view the progress and final status of the operations they perform.

Note An archive is a special type of backup. During an archive, NetBackup first backs up the selected files then deletes them 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).

Client actions are described in the *NetBackup User's Guide for UNIX*.

Activity Monitor

Use the Activity Monitor to monitor and control NetBackup jobs, daemons, and processes. For more information see Chapter 6, “Monitoring NetBackup Activity” on page 277.

NetBackup Management

This manual describes the applications and utilities listed under **NetBackup Management** in the NetBackup Administration Console tree. The *Media Manager System Administrator's Guide* describes applications and utilities under **Media and Device Management**.

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 257.

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 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 and can be covered by more than one.

The properties of a backup policy include the following:

- ◆ General attributes that define the:
 - ◆ *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.
- ◆ 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.
- ◆ Schedules that control when backups and archives can occur for the clients.

As mentioned above, each backup policy has its own set of schedules. These schedules control when automatic backups start and also when users can start a backup or archive. Each schedule is unique with 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. Manual backups are useful if, for example, a client system is down and misses its scheduled backup.



For more information see Chapter 3, “Managing Backup Policies” on page 57.

Storage Units

Use **Storage Units** to display storage unit information and provide commands for managing NetBackup storage units.

A storage unit is a group of 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 a 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 and then defines the storage units. During a backup, NetBackup sends data to the storage unit specified by the backup policy. Media Manager then picks an available device within the storage unit.

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

Storage units simplify administration because once they are defined, the NetBackup configuration 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 27.

Catalog

Use **Catalog** to create and configure a special type of backup NetBackup requires for its own internal databases—a *catalog backup*. These databases, called catalogs, are on the NetBackup server's disk and have setup information as well as critical information on client backups. The catalog backups are set up and tracked separately from other backups to ensure recovery in case of a server crash.

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, “Managing Catalogs and Images” on page 199.

Host Properties

Use **Host Properties** to customize NetBackup configuration options. In most instances, no setting changes are necessary. However, **Host Properties** settings allow 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 297.

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 the chapter, “Access Management,” in the *NetBackup System Administrator’s Guide, Volume II*.

NetBackup Configuration Wizards

The easiest way to configure NetBackup is to use the configuration wizards. The wizard selection is visible in the Details pane on the right varies depending on what NetBackup node 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.

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 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 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 “Adjusting Time Zones in the NetBackup-Java Console” on page 131.)

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 current window. You will 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 **Edit** menu contains the following options:

New

Displays a dialog where you can specify criteria for a new item.



Change

Displays a dialog where you can specify changes to the selected item.

Delete

Deletes the selected item.

Find

Use the **Find** option (**Control+F**) to highlight rows in the Details view that meet specific criteria.

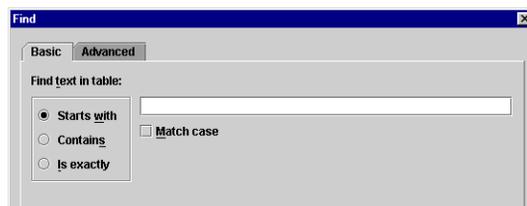
Note The Find option finds data in hidden columns. The Find option will not find data in hidden rows. That is, rows hidden due to filtering.

The Find dialog contains two tabs: Basic and Advanced.

◆ Basic Tab

In the field on the Basic tab, specify characters you wish to find in any column in the Details view.

If you're searching through a large amount of data and you know in which column the data will be found, consider using the Advanced tab to generate results more quickly. When using the Basic tab, for every row, each column is searched until a column with a match is found.



Check **Match case** to perform a case-sensitive comparison. In most cases, using **Match case** will speed up your search.

Select one of the following:

Starts with: Highlight rows containing columns where characters that start with the specified characters are found.

Contains: Highlight rows containing columns where characters composed of the specified characters are found.

Is exactly: Highlight rows containing columns where characters that match the specified characters exactly are found.

Both tabs contain the following buttons:

Find All button: Highlights all the rows in the table that meet the Find criteria.

Find Next button: Highlights the next row in the table that meets the Find criteria.

Cancel button: Click to close the dialog without performing a search.

◆ Advanced Tab

Use the Advanced Tab to specify characters you wish to find under a specific column in the table.

▼ To Find using the Advanced tab

1. Open the Find dialog by right-clicking in the Detail view and selecting **Find** or by pressing **Control+F**.
2. Define the search criteria:

- a. From the Field list, select the name of the column you wish to search.
- b. Choose a Comparison method.
- c. In the Value field, enter the value to be considered in the search criteria.

Depending on the data type, a number spinner may be available for selecting a value. If the value is case-sensitive, a **Match case** check box is available to perform a case-sensitive comparison.

- d. Check the **Add to List** button to include the criteria in the criteria list. If you wish to delete one of the search criteria, select the item and click **Remove**.
3. If multiple selection criteria were created, select whether:
 - ◆ Matches should be made only if all criteria are met (**AND**)
 - ◆ Matches should be made if at least one of the criteria is met (**OR**)
 4. Select **Find All** to highlight all the rows in the table or **Find Next** to highlight the next row in the table that meets the Find criteria.

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.

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 284.
- ◆ 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**.

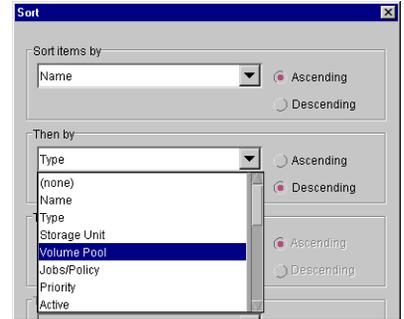
Column Layout

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 four 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	Volume	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_ipc_test	Standard	<any>	NetBackup	
wobegon_xp_stripe_r...	Standard	<any>	NetBackup	
wobegon_xp_stripe_p...	Standard	<any>	NetBackup	
wobegon_xp_stripe_tbu	FlashBackup	<any>	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.

The Filter dialog contains two tabs: Basic and Advanced.



◆ Basic Tab

In the field on the Basic tab, specify characters you wish to filter for in any column.

If you're filtering a large amount of data and you know in which column the data will be found, consider using the Advanced tab to generate results more quickly. When using the Basic tab, for every row, each column is searched until a column with a match is found.

Check **Match case** to perform a case-sensitive comparison. In many cases, using **Match case** will speed up the filtering.

Select one of the following:

Starts with: Select to show rows containing columns where characters that start with the specified characters are found.

Contains: Select to show rows containing columns where characters composed of the specified characters are found.

Is exactly: Select to show rows containing columns where characters that match the specified characters exactly are found.

Both tabs contain the following buttons:

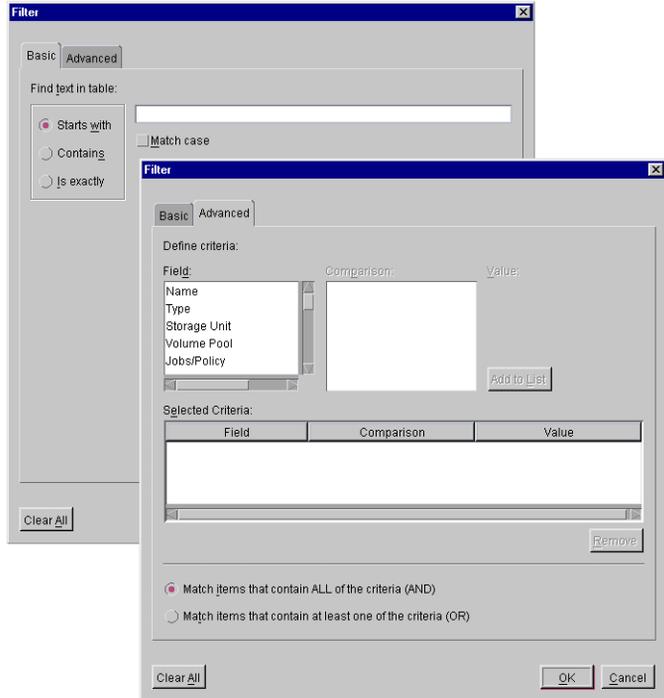
Clear All button: Click to remove all filtering criteria from the Filter dialog.

OK button: Click to apply the filtering criteria.

Cancel button: Click to close the dialog without filtering the data.

◆ Advanced Tab

Use the Advanced Tab to specify characters you wish to filter for under a specific column in the table.



▼ To Filter using the Advanced tab

1. Open the Filter dialog by right-clicking in the Detail view and selecting **Filter**.
2. Define the search criteria:
 - a. From the Field list, select the name of the column you wish to search.
 - b. Choose a Comparison method.
 - c. In the Value field, enter the value to be considered in the filtering criteria.
Depending on the data type, a number spinner may be available for selecting a value. If the value is case-sensitive, a **Match case** check box is available to perform a case-sensitive comparison.
 - d. Check the **Add to List** button to include the criteria in the criteria list. If you wish to delete one of the filtering criteria, select the item and click **Remove**.
3. If multiple selection criteria were created, select whether:
 - ◆ Matches should be made only if all criteria are met (**AND**)
 - ◆ Matches should be made if at least one of the criteria is met (**OR**)
4. Select **OK** to display only those rows in the table that meet the filtering criteria.

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.

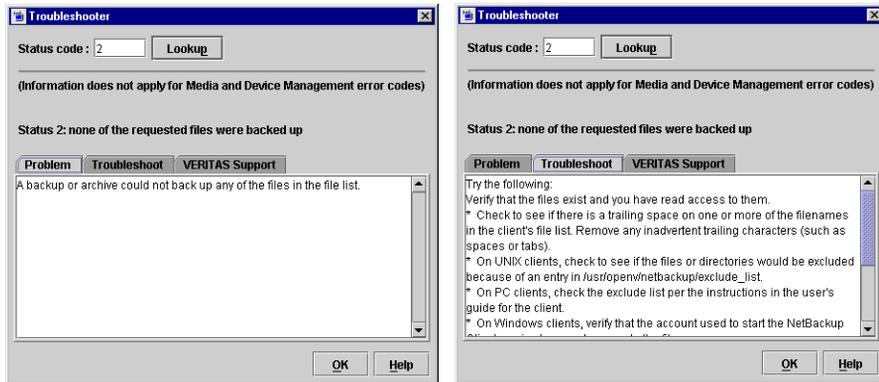
Troubleshooter

Use the **Troubleshooter** option to open the Troubleshooter dialog.

To use the Troubleshooter from the menu, 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.



The Troubleshooter is also available from the following locations:

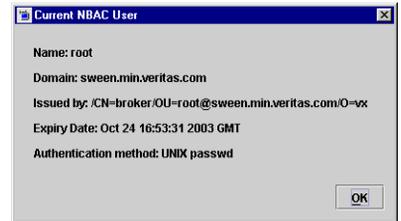
- ◆ From the Activity Monitor, on the Detailed Status tab of a job.
- ◆ From the Activity Monitor, on the right-click pop-up menu.*
- ◆ From Reports, by clicking a status code hyperlink.
- ◆ As a button on the Toolbar.

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

▼ **To invoke the pop-up menu on Solaris X86 systems**

1. Press the **Ctrl** key and hold.
2. Press the second mouse button. A menu appears.
3. Release the **Ctrl** key.
4. Select an item from the menu in either of the following ways:
 - ◆ Drag the cursor to the item and release the second mouse button, or
 - ◆ Release the second mouse button, then select the menu item with the first mouse button.

- ◆ **VERITAS Web Page:** Displays the VERITAS web page if the system has a browser configured.
- ◆ **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 you are configuring NetBackup for the first time, choose the Getting Started Wizard. This wizard steps you through the other wizards and leaves you 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. Complete instructions are not given here, but each step provides references for more information if you require it.

1. Start the NetBackup Administration Console. (See “NetBackup Administration Interfaces” on page 4.)
2. Complete the addition of storage devices. The preferred method is to use the Configure Storage Devices wizard. To perform configuration without the wizard, see the *Media Manager System Administrator’s Guide*.
3. Add the media that you will use. For instructions, see the *Media Manager System Administrator’s Guide*.
4. Ensure that the NetBackup database daemon, `bpdbm`, is running. This daemon must be running so NetBackup can update its catalogs with the new setup information.

`bprd` usually starts `bpdbm` at boot time.

To check the state of `bprd` and `bpdbm`, use the script

```
/usr/opensv/netbackup/bin/bpps
```

If necessary, start `bprd` and `bpdbm` by running the following command

```
/usr/opensv/netbackup/bin/initbprd
```

Note See “Managing Daemons” on page 413 for instructions on starting and stopping `bprd` and `bpdbm`.

5. Define the storage units. (See “Managing Storage Units” on page 27.)
6. Verify the catalog backup configuration. (See “Managing Catalogs and Images” on page 199.)
 - a. Specify the media to use.
 - b. Make any necessary changes to the backup paths. The default paths to the catalogs are added automatically.



7. Define the backup policies. (See “Managing Backup Policies” on page 57.)
8. Perform any required additional configuration as explained in “Additional Configuration” on page 99.





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

- ◆ “Introduction to Storage Units” on page 28
- ◆ “Using the Device Configuration Wizard” on page 30
- ◆ “Media Manager Storage Unit Considerations” on page 31
- ◆ “Disk Storage Unit Considerations” on page 36
- ◆ “NDMP Storage Unit Considerations” on page 37
- ◆ “Disk Staging Storage Unit Considerations” on page 38
- ◆ “Maintaining Storage Units” on page 45
- ◆ “Storage Unit Properties” on page 47
- ◆ “Configuring Drive Availability Checking” on page 53
- ◆ “Creating and Changing Storage Unit Groups” on page 54



Introduction to Storage Units

A NetBackup storage unit is a group of one or more storage devices of a specific type and density that attach to a NetBackup server. The administrator defines the storage units that are available for a backup and which storage unit to use for each policy. For example, it is possible to specify a robot as the storage unit for one policy and a standalone tape drive for another policy.

Note Storage unit names are case-sensitive.

There are four types of storage units:

◆ Media Manager storage units

A Media Manager storage unit uses tape robots, standalone tape drives, or optical disk devices, that are under 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 31.)

◆ Disk storage units

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

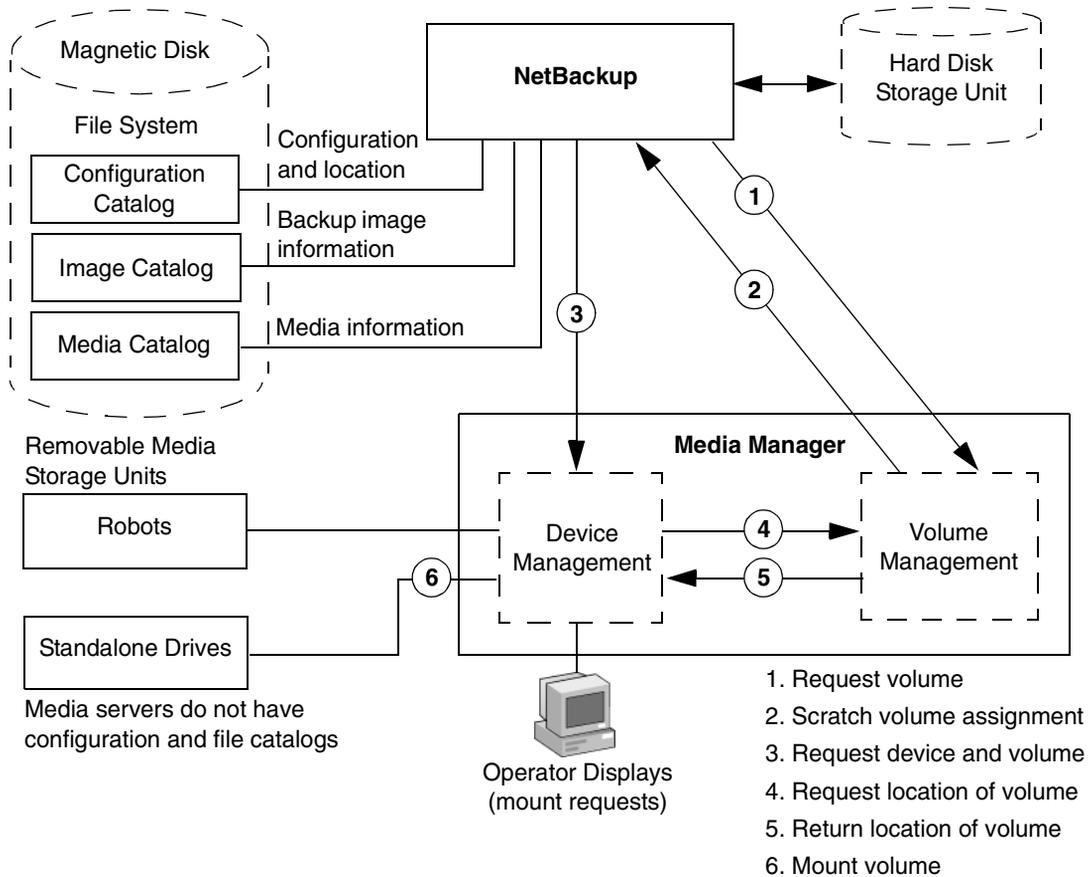
◆ NDMP storage units

NDMP storage units are controlled by Media Manager but attach to NDMP hosts and require that you have the NetBackup for NDMP option installed. (See “NDMP Storage Unit Considerations” on page 37.)

◆ Disk Staging storage units

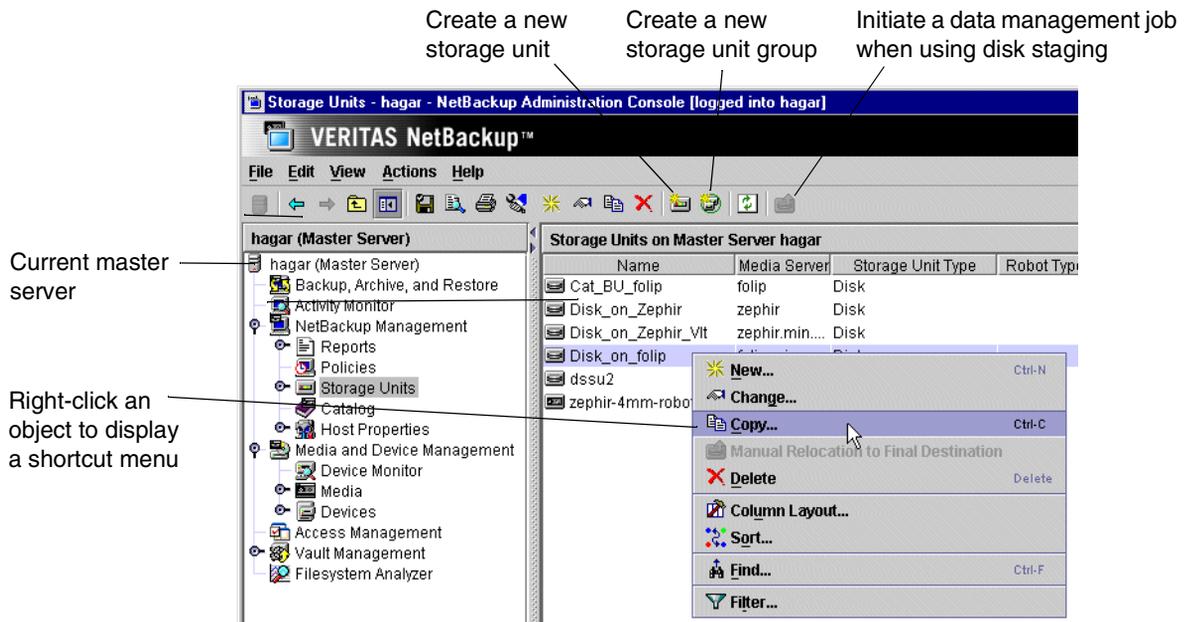
A disk staging storage unit provides the first storage location in a two-stage process called Disk Staging. In this process, client data is backed up to a disk staging storage unit, then, in the second stage, the data is relocated to another storage unit. (See “Disk Staging Storage Unit Considerations” on page 38.)

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



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 420.

Using the Device Configuration Wizard

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.

If you are modifying an existing configuration or want access to more settings, see “Maintaining Storage Units” on page 45.

▼ 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 hard disk storage unit if no devices are found.

Media Manager Storage Unit Considerations

NetBackup keeps records about the files in the backups and about the media where the records are stored. Media Manager manages the removable storage units (for example, tape drives) and tracks the location of both online and offline volumes. If the storage unit is on disk, the data goes to the file path specified during configuration of the storage unit. The operating system disk manager manages the actual reading and writing of data.

When sending a backup to a Media Manager storage unit, NetBackup looks in its media catalog for a previously used volume that is the correct density and is configured to retain backups for the desired period of time. If none of the previously used volumes are suitable, NetBackup requests a new media ID from Media Manager and then requests Media Manager to mount the volume in a device.

Note When a volume is allocated to NetBackup, other applications cannot use it until backups on the volume are no longer needed.

The request to Media Manager specifies both the media ID and device density of the volume. If a request involves a robot, the volume is then automatically mounted in a drive and assigned to the request. With a standalone drive, NetBackup attempts to use the media that is in the 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.

To restore from a Media Manager storage unit, NetBackup finds the media ID in its media catalog and requests the volume from Media Manager.

Note *Applies only to NetBackup Enterprise Server:* Media Manager is managed separately and can also be used by other applications, such as Storage Migrator.

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 Drives Used for Backup** 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.

For example, if a server has two 1/4-inch qscsi drives, add a storage unit with **Maximum Concurrent Drives Used for Backup** set to 2. Media Manager 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.

Allowable Media Manager Characters

The following set of characters can be used in user-defined names, such as storage units, volume groups, volume pool names, and media IDs that you enter when creating these entities. 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 (_)

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

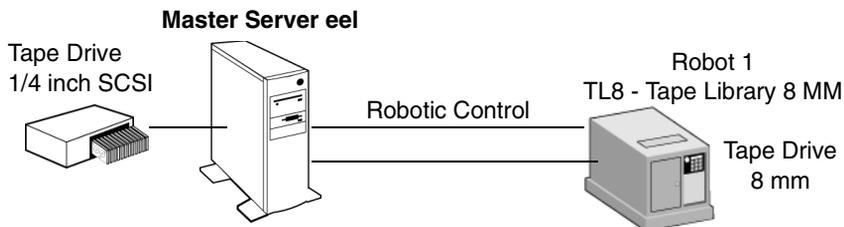
For step-by-step instructions on how to specify this information to NetBackup, see “Creating a New Storage Unit” on page 45.

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 52.)

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

- ◆ 8 mm tape drive in the robot

Storage Unit Configuration Setting	Value
Media Server	eel
Robot Type	TL8 - Tape Library 8MM
Robot Number	1
Maximum Concurrent Drives	1
Density	8mm - 8mm cartridge

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

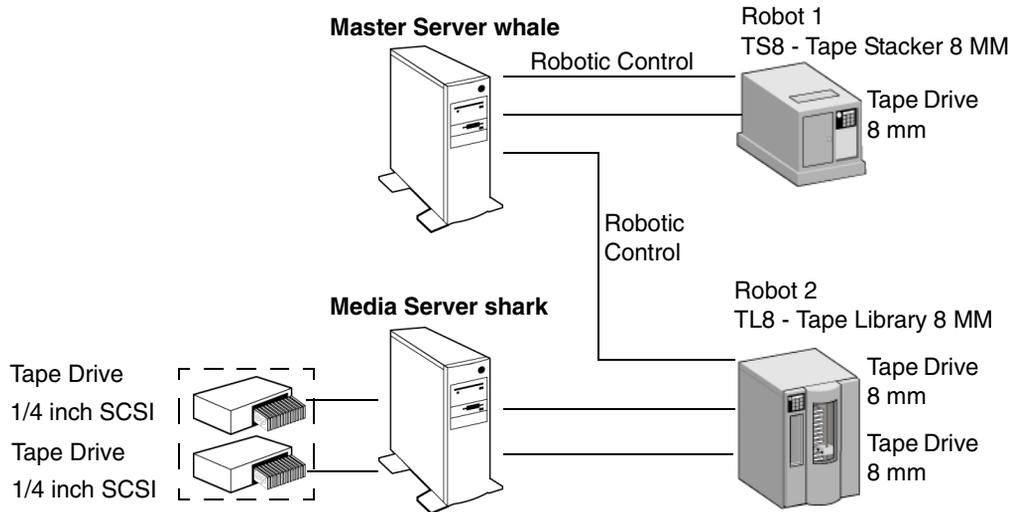
- ◆ SCSI 1/4 inch tape drive

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

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 TS8 - Tape Stacker 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	TS8 - Tape Stacker 8MM
Robot Number	1
Maximum Concurrent 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:

- ◆ 8 mm tape drives in robot 2

Storage Unit Configuration Setting	Value
Media Server	shark
Robot Type	TL8 - Tape Library 8MM
Robot Number	2
Maximum Concurrent 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	shark
Robot Type	None
Robot Number	None
Maximum Concurrent 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 consists of a directory on a hard disk that stores the backup or archive data. NetBackup permits an unlimited number of disk storage units.

A disk type storage unit is useful for testing and is useful during busy periods because it allows quick backups. However, you must be careful that it does not fill up your disk.

Before using a disk storage unit, configure the disk as explained in your operating system documentation. To calculate the approximate disk space that NetBackup requires as it creates backups, use the following formula:

$$\begin{aligned} & (\text{largest backup size} \times (\text{number of backups} + 1)) \\ & \quad + \\ & \text{Space for the restores that are concurrent with backups} \end{aligned}$$

NDMP Storage Unit Considerations

NDMP storage units are controlled by Media Manager but attach to NDMP hosts and require that you have the NetBackup for NDMP option installed. See the *NetBackup for NDMP System Administrator's Guide* for more information.



Disk Staging Storage Unit Considerations

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

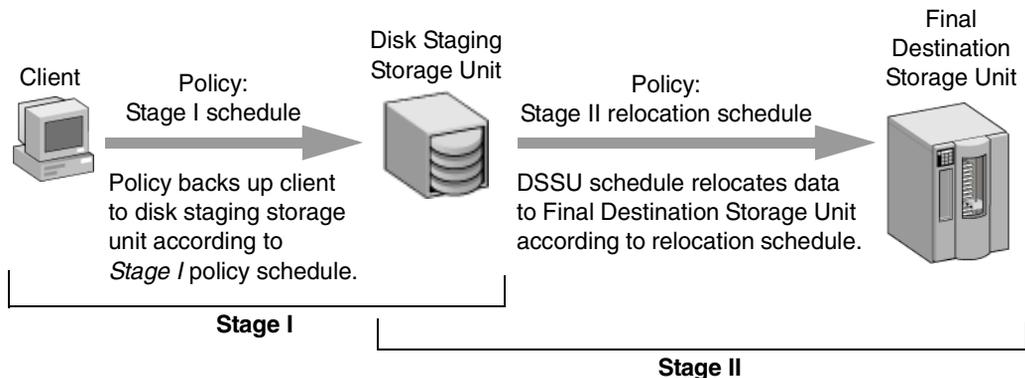
This two-stage process allows the NetBackup administrator to leverage the advantages of disk-based backups in the near term, while preserving the advantages of tape-based backups for long term.

Disk staging may be appropriate for your NetBackup environment to meet the following objectives:

- ◆ To allow backups when tape drives are scarce.
- ◆ To allow for faster restores from disk.
- ◆ To facilitate streaming to tape without image multiplexing.

Disk Staging is conducted in two separate operations:

- ◆ Stage I: A backup creates an image on the disk staging storage unit.
- ◆ Stage II: 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 image 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 staging 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.

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 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. NetBackup 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)

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 400GB and the desire is for the images to be kept for one week, the recommended average size is calculated based on the following formula:

Average Size of disk staging storage unit =
Average data per cycle * (number of cycles to keep data + 1 cycle for safety)

$$2.8\text{GB} = 400\text{MB} * (6 + 1)$$

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:

Maximum Size = Max data per cycle * (# of cycles to keep data + 1 cycle for safety)

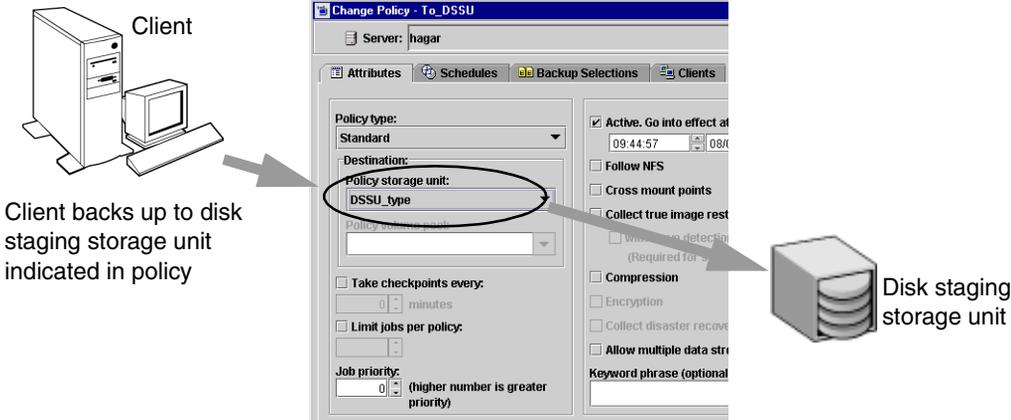
For example: $4.9\text{ GB} = 700\text{MB} * (6 + 1)$

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 storage unit. The storage unit should be one dedicated disk partition/file system for each disk staging storage unit. The schedule for Stage I is configured like any other backup.

Disk Staging: Stage I



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 **Disk Staging Schedule** button. The button is available only when **Disk Staging Storage Unit** is selected as the storage unit type.

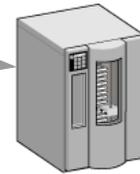
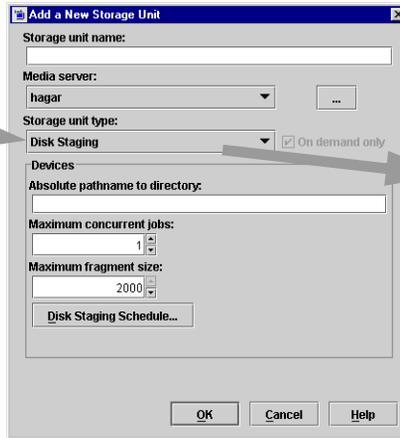


Disk Staging: Stage II

Disk staging storage unit

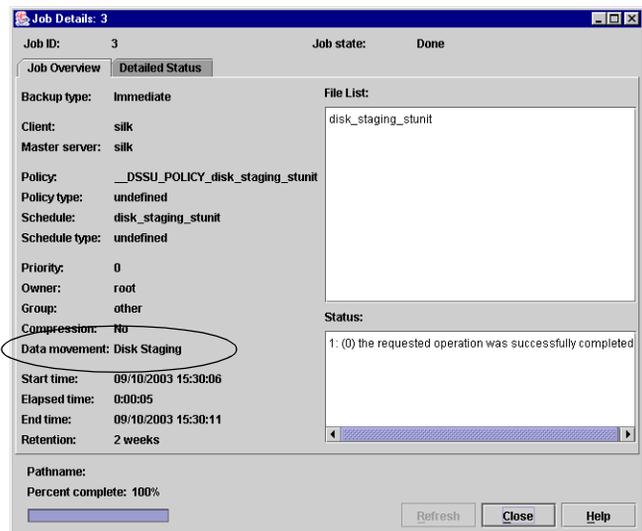


Images are relocated from the disk staging storage unit to the destination storage unit based on the relocation schedule indicated for the storage unit



Destination storage unit

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.

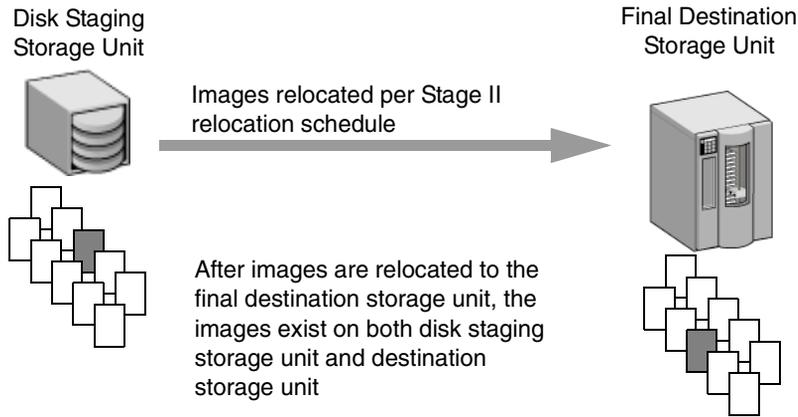
▼ To manually initiate a disk staging storage unit relocation schedule

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.



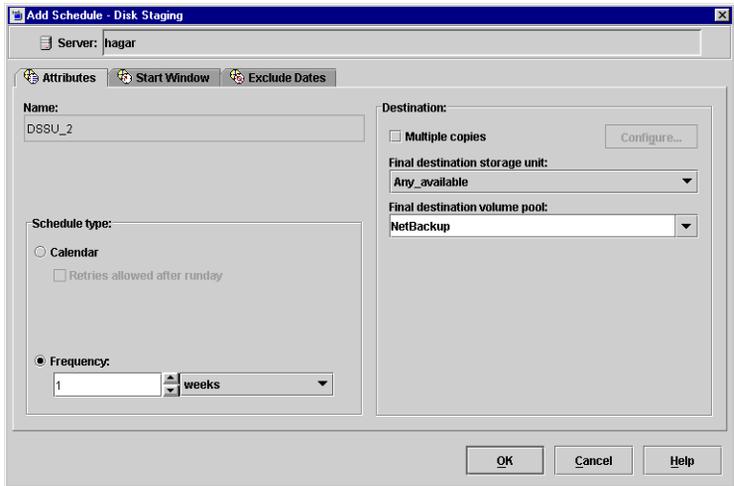
The image then exists both on the disk staging storage unit and the destination storage unit. When the disk staging storage unit becomes full, it is *cleaned* and the oldest images are deleted. (See “Cleaning the Disk Staging Storage Unit” on page 44.)



Disk Staging Schedule Button

Clicking the **Disk Staging Schedule** button brings up the **Add Schedule or Change Schedule** dialog.

This dialog is similar to the scheduling dialog seen when configuring policies. It does, however, contain some differences:



Name

The schedule name for a disk staging storage unit defaults to (and is required to be) the name of the storage unit.

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.



Note Images on the disk staging storage unit are relocated to tape in one stream only.

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 Manger storage unit (tape), or if *Any Available* is indicated for the Final Destination Storage Unit, the **Final Destination Volume Pool** is selectable.

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

Cleaning the Disk Staging Storage Unit

NetBackup detects when a disk staging storage unit is full and deletes the oldest images that have already been relocated. A backup job will not fail if there are enough images which have been duplicated, making them eligible to be cleaned. If enough space cannot be cleaned, the job will fail.

Maintaining Storage Units

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

- ◆ “Creating a New Storage Unit” on page 45
- ◆ “Changing Storage Unit Properties” on page 46
- ◆ “Deleting Storage Units” on page 46

Creating a New Storage Unit

There are two methods to create a new storage unit:

- ◆ Create a storage unit using the **Actions** menu
- ◆ Copy an existing storage unit

▼ 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 Add a New Storage Unit dialog appears.
3. Complete the fields on the Add a New Storage Unit dialog.
The options are described in “Storage Unit Properties” on page 47.
4. Click **OK** to add the storage unit to the configuration.

▼ 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 **Edit > Copy**. The Copy Storage Unit dialog appears.
4. Complete the fields on the Copy Storage Unit dialog.
The options are described in “Storage Unit Properties” on page 47.



Changing Storage Unit Properties

We suggest that you make changes only during periods when you are not expecting backup activity for policies that will be affected by the changes. This allows time for you to make adjustments before backups begin and ensures an orderly transition from one configuration to another. Regardless of your timing, NetBackup is designed to prevent serious problems or failures from occurring.

▼ To change storage unit properties

1. If your site has more than one master server, select **File > Change Server** to choose the server with the configuration that will use the storage unit.
2. In the NetBackup Administration Console, select **NetBackup Management > Storage Units**. Storage unit information appears in the Details pane.
3. Double-click the storage unit you wish to change from those listed in the Details pane.
4. Complete the fields on the Change Storage Unit dialog.

The options are described in “Storage Unit Properties” on page 47.

Deleting Storage Units

Deleting a storage unit from the NetBackup configuration does not prevent you from 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.



Storage Unit Properties

The following sections list and describe storage unit properties. Since not all properties pertain to all types of storage units, the properties are listed alphabetically.

The screenshot shows the 'Add a New Storage Unit' dialog box with the 'Storage unit type' set to 'Disk (Directory on a Hard Drive)'. The 'On demand only' checkbox is checked. The 'Devices' section includes an empty text field for the absolute path to the directory, a spinner for 'Maximum concurrent jobs' set to 1, and a spinner for 'Maximum fragment size' set to 2000 megabytes. A 'Reset to Defaults' button is located at the bottom right of the dialog, and 'OK', 'Cancel', and 'Help' buttons are at the bottom.

Properties for Disk Storage Type

The screenshot shows the 'Add a New Storage Unit' dialog box with the 'Storage unit type' set to 'Media Manager (Robot or Standalone Drive)'. The 'On demand only' checkbox is unchecked. The 'Devices' section includes a dropdown for 'Storage Device', a text field for 'Robot type' containing 'TL4', a text field for 'Density' containing '4MM', and a text field for 'Robot number' containing '0'. There are also spinners for 'Maximum concurrent drives used for backup' (set to 1) and 'Maximum multiplexing per drive' (set to 1). A disabled 'Maximum fragment size' field is also present. A 'Reset to Defaults' button is located at the bottom right of the dialog, and 'OK', 'Cancel', and 'Help' buttons are at the bottom.

Properties for Media Manager Storage Type



The screenshot shows the 'Add a New Storage Unit' dialog box with the 'NDMP' storage type selected. The 'Media server' is set to 'hagar'. The 'Storage unit type' is 'NDMP', and the 'On demand only' checkbox is unchecked. The 'Devices' section includes fields for 'NDMP host', 'Storage Device', 'Robot type' (set to 'TL4'), 'Density' (set to '4MM'), 'Robot number' (set to '0'), and 'Maximum concurrent drives used for backup' (set to '1'). A 'Reset to Defaults' button is located at the bottom right of the dialog.

Properties for NDMP Storage Type

The screenshot shows the 'Add a New Storage Unit' dialog box with the 'Disk Staging' storage type selected. The 'Media server' is set to 'hagar'. The 'Storage unit type' is 'Disk Staging', and the 'On demand only' checkbox is checked. The 'Devices' section includes a field for 'Absolute pathname to directory'. Below this are spinners for 'Maximum concurrent jobs' (set to '1') and 'Maximum fragment size' (set to '2000'). A 'Disk Staging Schedule...' button is located below these spinners. 'OK', 'Cancel', and 'Help' buttons are at the bottom.

Properties for Disk Staging Storage Unit

Absolute Pathname to Directory

Absolute Pathname to Directory specifies the absolute pathname of the file system that will store the backups. Enter the pathname directly in the field. Use any location on the disk, providing there is sufficient space available.

The following rule applies to the path specified:

In addition to the platform-specific file path separators (/ and \) and colon (:), within a drive specification on Windows, use only alphabetic (ASCII A - X, a - z), numeric (0-9), plus (+), minus (-), underscore (_), or period (.) characters. Do not use a minus as the first character.

Density

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

Disk Staging Relocation Schedule

Click the **Disk Staging Schedule** button to set up the second stage of disk staging, during which the backup image is relocated from the disk staging unit to the final destination storage unit.

The image exists on both the disk staging unit as well as the final destination storage unit until the disk staging storage unit is full. The image on the disk storage unit continues to be the primary image until it is expired or removed, after which the image on the final storage unit is the primary image.

For more information on disk staging, see “Disk Staging Storage Unit Considerations” on page 38

Maximum Concurrent Drives Used for Backup

Maximum Concurrent Drives Used for Backup specifies the number of tape drives that NetBackup can use at one time for backups in this storage unit. The number of tape drives available is limited to the maximum number of tape drives in the storage device.

Select the desired number:

- ◆ 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 hard disk storage units, **Maximum Concurrent Jobs** specifies the maximum number of backups that NetBackup can concurrently send to this disk. For example, if there are three backup jobs for this storage unit and **Maximum Concurrent Jobs** is set to two, the first two jobs start and the third one waits.

Note Specifying 0 disables the storage unit.

Maximum Concurrent Jobs corresponds to the **Maximum Concurrent Drives** setting for a Media Manager storage unit. The jobs are not multiplexed.



The number to enter depends on the available disk space and the server's ability to comfortably run multiple backup processes. The default is 1. (See "Limit Jobs Per Policy" on page 81.)

Maximum Fragment Size

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

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.

For Media Manager storage units:

The default maximum fragment size for a Media Manager storage unit is 1 terabyte. To specify a maximum fragment size other than the default, place a check in the **Maximum Fragment Size** check box, then enter a value of 50 megabytes to 1048576 megabytes (1 terabyte).

Fragmenting 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 backup and reads tar headers until finding the desired file.

For hard disk storage units:

The default maximum fragment size for a disk storage unit is 2000 megabytes. To specify a maximum fragment size other than the default, enter a value that ranges from 20 megabytes to 2000 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.

Note Changing the fragment size does not prevent restoring backups written using the former fragment size.

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.

Maximum Multiplexing per Drive

Specifies the maximum number of backups that NetBackup can multiplex onto any single drive in the storage unit.

To enable multiplexing, specify any value from 2 through 32. The default is 1, which disables multiplexing and allows only one backup job at a time per drive.

For values greater than 1, NetBackup sends concurrent, multiple backups from one or several clients to a single drive, and multiplexes the backups onto the media. See “Multiplexing” on page 100 for more information.

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 hard disk storage unit. For NDMP storage, this specifies the name of the NetBackup for NDMP server that will be backing up the NDMP host. Browse to, or enter the name that is used for that server in the NetBackup server list.

New Media Server (... Button)

Click the ... button to open the New Media Server dialog. Enter the name of a media server in the field, then click **OK**.



NDMP Host

NDMP Host specifies the NDMP host whose data will be sent to this storage unit. Enter the host manually.

On Demand Only

On Demand Only specifies whether the storage unit is available *only* 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 disk storage units and disk staging storage units, **On Demand Only** is selected by default. For all other storage types, **On Demand Only** is off by default.



Note If you make all storage units on demand only, designate a specific storage unit for each policy or schedule. Otherwise, NetBackup will be unable to find a storage unit to use.

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*.

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.

Use alphabetic (ASCII A-Z a-z), numeric (0-9), plus (+), minus (-), underscore(_), or period (.) characters. Do not use a minus as the first character or leave any spaces between characters.

Note The storage unit name is case-sensitive.

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
- ◆ **Disk Staging:** The first storage location in a two-stage process called Disk Staging.

Configuring Drive Availability Checking

NetBackup periodically checks each storage unit to determine the status of its drives and attempts to use a storage unit only if it has drives available. The following topics explain the configuration settings associated with this feature.

Interval Between Status Checks

The NetBackup host property, **Re-read Interval**, determines how often NetBackup checks storage units for available drives. (See “Re-read Interval for Available Drives” on page 359.)

Drive Count Timeout

When NetBackup checks for drive availability, it also counts the drives that are available for backups. This information is then used to prevent scheduling too many jobs for the number of drives.

The only setting associated with counting drives is the length of time that the scheduler waits for the count to complete. If you have problems with timeouts, you can extend the time that the scheduler waits by using the NetBackup host property, **BPTM (Drive Count) Query Timeout**. (See “BPTM (Drive Count) Query Timeout” on page 388.)



Requeuing Jobs If Required Storage Units are Unavailable

By default, a job fails if a required storage unit is unavailable when a job starts or, for some reason, becomes unavailable during a backup (Status code 219). You can configure NetBackup to requeue jobs for these conditions by setting the following NetBackup host properties on the Timeout dialog:

- ◆ **Requeue Active Jobs if Required Storage Unit is Unavailable**
- ◆ **Requeue Scheduled Jobs if Required Storage Unit is Unavailable**
- ◆ **Timeout in Queue** (See “Timeouts Properties” on page 386.)

Creating and Changing 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.

NetBackup uses the first storage unit listed if the storage unit is available. If the first one is not available, NetBackup attempts to use the second storage unit listed, and so on down the list.

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 76.)

▼ 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.

3. Enter a storage unit group name for the new storage unit group.

Use alphabetic (ASCII A-Z a-z), numeric (0-9), plus (+), minus (-), underscore(_), or period (.) characters.

Do not use a minus as the first character or leave any spaces between characters.

Note The storage unit group name is case-sensitive.

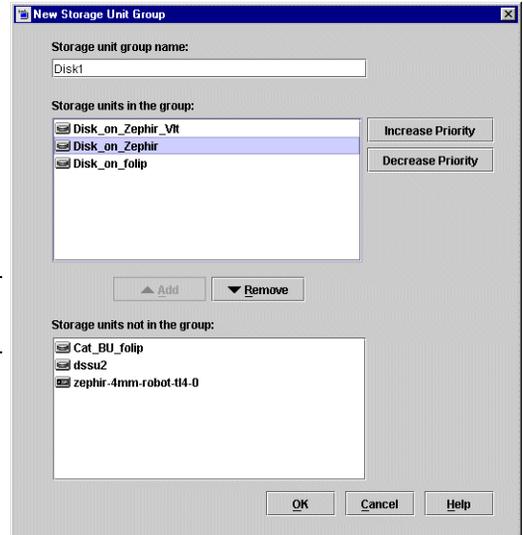
4. Add to or remove storage units from the group:

- 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**.

5. 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 priority of a storage unit, select the storage unit and click **Increase Priority** or **Decrease Priority**.



6. 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**.

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 their backup and archiving requirements, then create a policy for each group.

This chapter introduces policies, gives policy planning guidelines, and details configuration instructions:

- ◆ “Using the Policies Utility” on page 58
- ◆ “Standard and User Toolbars” on page 60
- ◆ “Introduction to Backup Policies” on page 60
- ◆ “Configuring Backup Policies” on page 61
- ◆ “Example Policies” on page 62
- ◆ “Policy Planning Guidelines for Backups” on page 63
- ◆ “Changing Policies” on page 70
- ◆ “What Type of Policy: Policy Attributes Tab” on page 73
- ◆ “Which Clients Will Be Backed Up: Clients Tab” on page 98
- ◆ “Which Selections Will Be Backed Up: Backup Selections Tab” on page 102
- ◆ “Rules for Backup File Paths” on page 111
- ◆ “When Will the Job Run: Schedules Tab” on page 144
- ◆ “Creating a Vault Policy” on page 195
- ◆ “Performing Manual Backups” on page 197



Using the Policies Utility

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

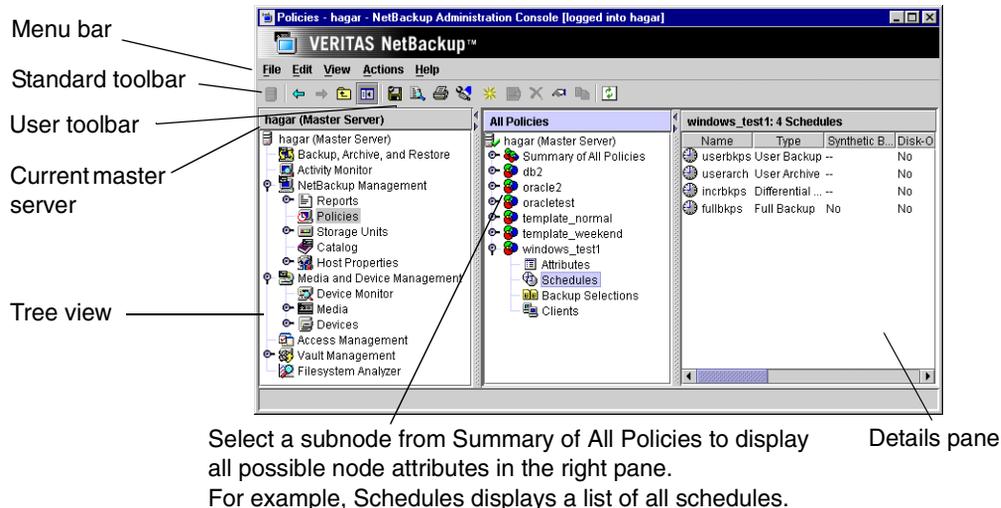
- ◆ “Tree and Detail Views” on page 58
- ◆ “Menus” on page 58
- ◆ “Standard and User Toolbars” on page 60

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

Tree and Detail Views

The center pane labeled **All Policies**, is 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.



Menus

The Menu bar consists of **File**, **Edit**, **View**, **Actions**, and **Help**. See Chapter 1 for a description of the items found on these menus.

Actions Menu

The following table describes options available on the **Actions** menu when **Policies** is selected.



Actions Menu

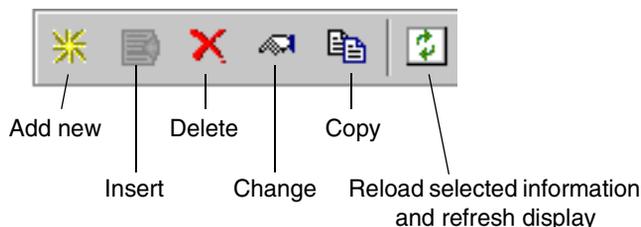
Menu	Commands
Activate	<p>Activates the policy that is selected in the Console tree. A policy must be active for NetBackup to run automatic backups or allow user backups or archives. This setting has no effect on restores.</p> <p>If the schedule is to be used for a catalog archive, the policy must <i>not</i> be active. The Active check box must be clear.</p> <p>For more information on configuring a policy for catalog archiving, see “Creating a Catalog Archiving Policy” on page 232.</p>
Deactivate	<p>Deactivates the selected policy (see Activate above).</p>
Manual Backup	<p>When a policy is selected in the Console tree, displays the Manual Backup dialog. Select a schedule and client in the Manual Backup dialog then click OK to start a manual backup.</p>
Install UNIX Client Software	<p>Allows you to install client software on a client from the NetBackup Administration Console. (See “Installing Client Software on Trusting UNIX Clients” on page 99.)</p>



Standard and User Toolbars

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

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



Introduction to Backup Policies

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

General Attributes on the Attributes Tab

The general attributes on the Add New Policy or Change Policy Attributes tab determine the basic characteristics of all the backups that NetBackup performs according to a policy. These include:

- ◆ Whether the policy is active and what date and time the policy will go into effect (so NetBackup can use it for backups).
- ◆ The type of backup policy, which primarily defines the type of clients the policy is set up to include.
- ◆ The priority that NetBackup gives to the backups for this policy relative to other policies.
- ◆ The storage unit that NetBackup uses by default when backing up clients covered by this policy. This setting can be overridden for individual schedules by specifying a storage unit for the schedule.

Schedules on the Schedules Tab

The schedules determine when the backups occur. Each schedule also includes criteria, such as how long to retain the backups.

There are two basic categories of schedules, automatic and user, and there are different types of schedules within these categories:

- ◆ *Automatic schedules* back up the backup selection list on all clients in the policy according to the timetables set up in the schedules. For example, you can set one schedule for daily incremental backups and another for weekly full backups. A cumulative incremental backup includes all files that have changed since the last full backup. A full backup includes all files in the backup selection list regardless of whether they have changed.

Note NetBackup 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).

- ◆ *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.

Client List on the Clients Tab

The client list names the computers that will be backed up according to a policy. A client must be covered by at least one backup policy and can be covered by more than one. 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 rules.

Backup Selections on the Selections Tab

The backup selections list names the files, directories, directives, scripts (used for database policies), and templates (used for Oracle and DB2 policies), that NetBackup includes in automatic backups of clients covered by a 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, as NetBackup backs up the files in directories NetBackup finds.

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. To configure a calendar-based schedule, see “Calendar Schedule Tab” on page 176.

▼ **To create a policy using the 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**.
See the *NetBackup Installation Guide* for step-by-step instructions.

▼ **To create a policy without using the wizard**

1. In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
2. Click the New button on the toolbar .
3. Type a unique name for the new policy in the **Add a New Policy** dialog.
Use alphabetic (ASCII A-Z a-z), numeric (0-9), plus (+), minus (-), underscore(_), or period (.) characters. Do not use a minus or period (.) as the first or last character. Do not leave any spaces between characters.

If you decide you’d like to use the Backup Policy Configuration Wizard to configure the policy, select **Use add policy wizard**.



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 232.

4. Click **OK**.

For help while running the wizard, click **Help** in any of the wizard screens.

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 8 mm tape.

Example Backup Policy 1

Client List	Backup selection list	Schedules		
mars	<code>/usr</code>	Daily Incrementals	Weekly Fulls	User Backups
jupiter	<code>/home</code>	Run every day between 6 pm and 6 am. Store on 8 mm tape. Keep 14 days.	Run Mondays every week between 6 pm and 6 am. Store on 8 mm tape. Keep one month.	User can run any day between 8 am and 5 pm. Store on 8 mm tape. Keep one year.
neptune				

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		
pluto	<code>/usr</code>	Daily Incrementals	Weekly Fulls	Monthly Fulls
mercury	<code>/home</code>	Run every day between 6 pm and 6 am. Store on 8 mm tape. Keep 14 days.	Run Tuesdays every week between 6 pm and 6 am. Store on 8 mm tape. Keep one month.	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. The planning worksheets in this manual may also be helpful. (See "Planning Worksheets" on page 255.)

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	mercury (file server)
WS1	mars jupiter (workstations) neptune

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.

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 8 mm 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	mercury	/ /usr /h001 /h002/projects	8 mm
S2	mercury mercury	/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	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
S1	mercury	/usr /h001 /h002/projects	high	8 mm	Daily Incr Weekly Full 4 Weeks Full
S2	mercury	/h002/devexp	high	DLT	Daily Incr Weekly Full 4 Weeks Full
S3	mercury	/	low	8 mm	Daily Incr 4 Weeks Full
S4	mercury	/h002/desdoc	high	DLT	Daily Incr Weekly Full 4 Weeks Full 12 Weeks Full
WS1	mars	/usr /people	high	8 mm	Daily Incr Weekly Full 4 Weeks Full
	jupiter	/usr /home			
	neptune	/usr /people /var			
WS2	mars	/	low	8 mm	Daily Incr 4 Weeks Full
	jupiter	/			
	neptune	/			



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 NT and Windows 2000 clients in an 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.

There are also other general attributes that are explained later in this chapter. 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 we create a new policy for `/h002/projects`. This new policy (S5) has the same requirements as S1 but we can now 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 95 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 “Multiplexing” on page 100 and “Allow Multiple Data Streams” on page 93.)

Policy	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
S1	mercury	/usr /h001	high	8 mm	Daily Incremental Cumulative Incremental 4 Weeks Full
S2	mercury	/h002/devexp	high	DLT	Daily Incremental Cumulative Incremental 4 Weeks Full
S3	mercury	/	low	8 mm	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	8 mm	Daily Incremental Weekly Full 4 Weeks Full
WS1	mars	/usr /home	high	8 mm	Daily Incremental Weekly Full 4 Weeks Full
	jupiter	/usr /home			
	neptune	/usr /home /var			

Policy	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
WS2	mars	/	low	8 mm	Daily Incremental
	jupiter	/			4 Weeks Full
	neptune	/			



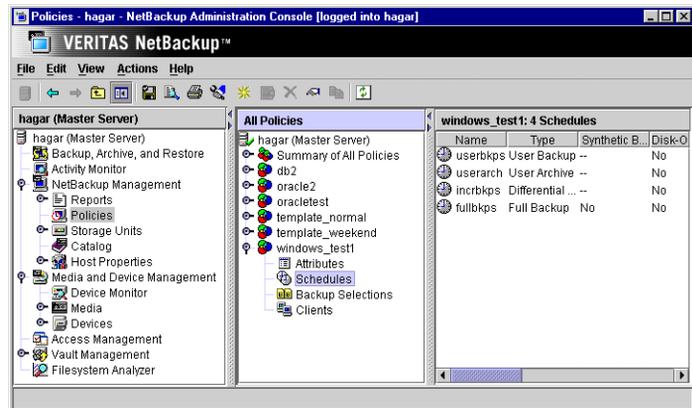
Changing Policies

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. Regardless of your timing, NetBackup is designed to prevent serious problems or failures from occurring.

▼ 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 **Edit > New**. The Add New 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 144.)

6. If this is the last schedule, click **OK**. To add more schedules, click **Add** and repeat the previous step.

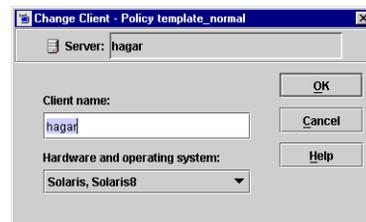
Click **Close** to cancel changes that have not been added.



▼ 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**.
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.



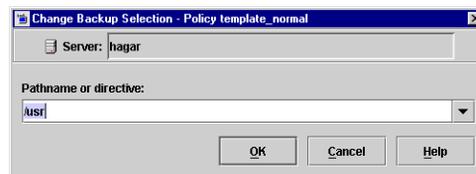
4. Complete the entries in the Add Client or Change Client dialog. (See “To add a client to a policy” on page 98.)

▼ 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 196.

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 Backup File Paths” on page 111 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

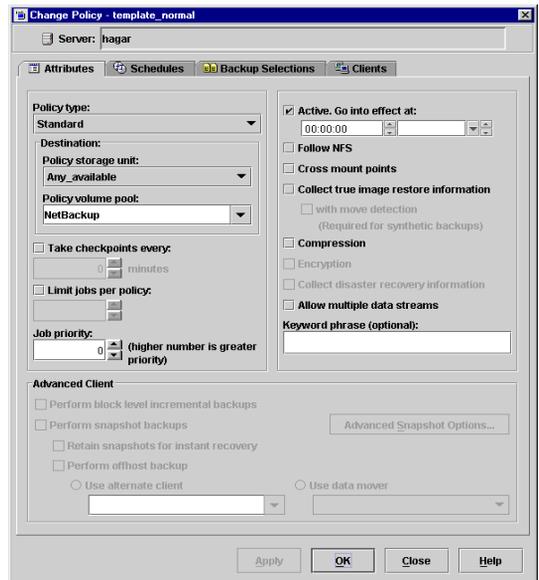
What Type of Policy: Policy Attributes Tab

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

▼ 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
 - ◆ “Which Clients Will Be Backed Up: Clients Tab” on page 98.
 - ◆ “Which Selections Will Be Backed Up: Backup Selections Tab” on page 102.
 - ◆ “When Will the Job Run: Schedules Tab” on page 144.
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.

The following sections describe the policy configuration attributes in the Attributes tab. Policy attribute 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** selection 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
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.
DataStore	A policy type reserved for use by VERITAS or its partners to provide agents for new applications or databases.
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-Windows-NT*	Use when the policy will contain only Windows 2000, NT, XP, or Windows Server 2003 clients.
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.
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.
NetWare	Use when the policy will contain only NonTarget NetBackup Novell NetWare clients (this version uses a Microsoft Windows interface).
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.
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.

Policy Type	Description
Standard*	Use when the policy will contain any combination of the following: <ul style="list-style-type: none"> ◆ NetBackup Novell NetWare clients that have the target version of NetBackup software. ◆ UNIX clients (including Mac OS X clients), except those covered by specific such as Oracle.
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.
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-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> .
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> .
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.
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.
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.
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.



Policy Type	Description
-------------	-------------

* To utilize CheckPoint Restart for backups or Checkpoint Restart for restores when backing up or restoring files, either the **Standard** or **MS-Windows-NT** policy type must be used. The creation of synthetic backups also requires the use of one of these policy types.

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

Policy Storage Unit

The **Policy Storage Unit** policy attribute specifies the default storage unit for backups of this policy. NetBackup uses the default storage unit for all schedules that do not specify another storage unit. A schedule-level storage unit (when specified) overrides the policy default. (See "Override Policy Storage Unit" on page 169.)

Select the policy storage unit from the drop-down list. You can also indicate **Any Available**. If you select **Any Available**, NetBackup tries locally-attached storage units first, and if none are found, the storage units are tried in alphabetical order. 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

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 sequence based on alphabetical order.

Policy Storage Unit Example

Assume that all schedules but one can 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.

Notes on Specifying a Storage Unit

- ◆ If your site has only one storage unit or there is no preference for storage:

- ◆ Specify *Any Available* for the policy storage unit *and*
- ◆ Do not specify a storage unit at the schedule level

However, in this instance, ensure that you 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.

- ◆ If you designate a specific storage unit and it is not available (for example, because it is down for maintenance), backups will not run for policies and schedules that require the storage unit.
- ◆ If your NetBackup configuration has several storage units and you want a policy to use *more than one but not all* of the storage units, select a storage unit group that has been configured to contain the desired storage units.

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 and volumes that are available only to the desired storage units.
 - b. For the policy, set **Policy Volume Pool** to the volume pool defined in step a.
 - 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 51.)

Policy Volume Pool

The **Policy Volume Pool** policy attribute specifies the default volume pool for backups of this policy. Select the desired volume pool name from the drop-down list. The list displays all previously-configured volume pools. Whenever a new volume is required for either a robotic or standalone drive, it is allocated to NetBackup from the requested volume pool.

A *volume pool* is a set of media used *only* by the users and hosts designated when configuring the pool. Volume pools are created and media assigned when configuring media in Media Manager type storage devices. It is not available for disk type storage devices.



A volume pool named *NetBackup* is always created by default and, unless otherwise specified in the policy, all backups go to media in the *NetBackup* pool. Other pools can be created for NetBackup to use. For example, create *Auto* and *User* volume pools, then specify that automatic backups use media from the *Auto* pool and user backups go to media in the *User* pool.

A schedule-level volume pool, when specified, overrides the policy default set here. (See “Override Policy Volume Pool” on page 169.) If there is no volume pool specified in the Attributes tab of either the policy or the schedule, NetBackup uses the *NetBackup* pool.

The volume pool concept is relevant only for storage units managed by Media Manager, and does not apply to disk storage units. For more information on volume pools, see the *NetBackup Media Manager System Administrator’s Guide*.

Volume Pool Example

Assume that you want all schedules but one to use the *backups* pool. The exception in this case is a user-archive schedule that requires the *archive* pool.

Here, 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.

Notes on Volume Pools

- ◆ This setting is optional for Media Manager type storage units and is not available for Disk type storage units.
- ◆ When configuring Media Manager, always specify the desired user and group for this Volume Pool.
- ◆ It is possible to configure a scratch pool from which NetBackup can automatically transfer volumes when another volume pool does not have media available.

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

Checkpoint Restart for Backup Jobs

The **Checkpoint Restart Every** 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 **Checkpoint Restart** 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 362.)

Policy types *MS-Windows-NT* (for Windows clients) and *Standard* (for UNIX clients) support checkpoints for backup jobs.

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 **Checkpoint Restart** 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 396.)
- ◆ **Advanced Client:** **Checkpoint Restart** is supported for use with local or alternate client backups. However, other types of Advanced Backup Methods (ABM) are not supported: Block Level Incremental Backups, Media Server Copy, Third-Party Copy Device, and Instant Recovery backups.)



- ◆ 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 40.) **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 resume, 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.
- ◆ The scheduler decides when a new job should be started instead of resuming an incomplete job. The scheduler 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 Global parameter, **Move Backup Job from Incomplete to Done State**, has passed.
 - ◆ For calendar scheduling, if another run day has arrived.

Checkpoint Restart for Restore Jobs

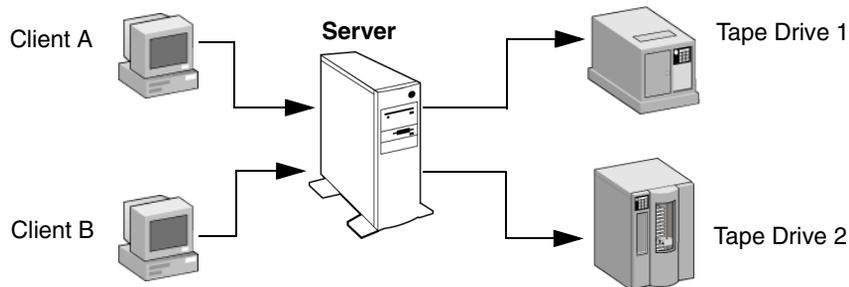
NetBackup takes checkpoints during restore jobs.

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

If the **Limit Jobs Per Policy** check box is clear (default), the maximum number of backup jobs that NetBackup will perform concurrently for this policy can be up to 999. To specify a lower limit, select the check box and specify a value from 1 to 999.

You can leave this attribute at the limit or default, except when there are enough devices that the possible number of concurrent backups will affect performance.



Client A and Client B backups can occur concurrently and to different devices

Notes on Limit Jobs Per Policy

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

- ◆ 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 magnetic disk, the storage device is defined as a file path and the available disk space determines how many paths are possible.
- ◆ 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.

- ◆ **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 its **Limit Jobs Per Policy** at 2, NetBackup can start two jobs from each policy and have a total of six running at one time.

Job Priority

The **Job Priority** policy 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, type any positive integer in the **Job Priority** text box. Higher values have higher priority. The maximum allowable priority is 99999. The default is 0.

Active. Go Into Effect At

To activate the policy, select the **Active** policy attribute check box. The policy must be active for NetBackup to run automatic-backup schedules or allow user backups or archives.

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

To deactivate a policy, remove the check from the **Active** 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 232.

Follow NFS

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

The **Follow NFS** policy 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.
- ◆ **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.

Disadvantages of Using Follow NFS

As a general rule, 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 you can also encounter problems with NFS mounts. In addition, you end up with multiple backups 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 their 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 394.)



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.

Cross Mount Points

The **Cross Mount Points** policy attribute applies only to certain policy types and NetBackup allows you to select it in only those instances.

The **Cross Mount Points** policy 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 according to 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 (/).

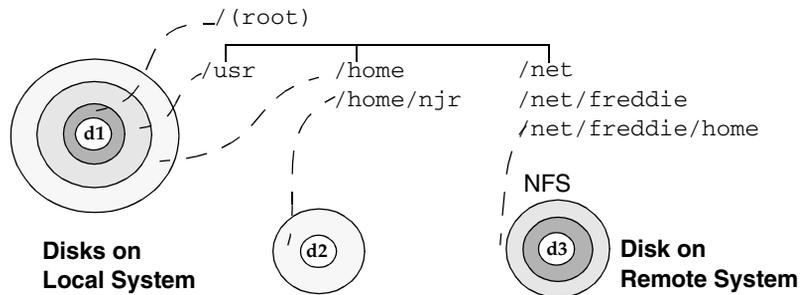
How Cross Mount Points Policy Attribute Interacts With Follow NFS

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 not back up everything, leave `/` out of the list and separately list the files and directories you want to include. The following backup selection list backs up only `/usr` and individual files under `/`:

```
/usr
/individual_files_under_root
```

Collect True Image Restore Information

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

The **Collect True Image Restore Information** policy 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.

To have true-image incremental backups include files that were moved, renamed, or newly installed in the directories, you must also select **With Move Detection**.

Note **Collect True Image Restore Information With Move Detection** must be selected if you wish to create synthetic backups. For more information on configuring synthetic backups, see “Synthetic Backups” on page 154.

Collect True Image Restore With Move Detection

The **Collect True Image Restore With Move Detection** policy attribute specifies that true-image incremental backups include files that were moved, renamed, or newly installed.

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 policy attribute must be selected if you wish 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 takes space on the client and can fail if there is not enough disk space available.

What Happens During True Image Restores

The following table shows the files backed up in the `/home/abc/doc/` directory during a series of backups between 12/01/2003 and 12/04/2003. Assume that **Collect True Image Restore Information** was selected for the policy that did the backups.

Day	Type of Backup	Files Backed Up in /home/abc/doc					
12/01/2003	Full	file1	file2	dirA/fileA	dirB/fileB	file3	
12/02/2003	Incremental	file1	file2	dirA/fileA	-----	-----	
12/03/2003	Incremental	file1	file2	dirA/fileA	-----	-----	
12/04/2003	User backup	file1	file2	dirA/fileA	-----	-----	dirC/fileC file4
12/04/2003	Incremental	file1	file2	-----	-----	-----	----- file4

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

Also, assume that you are going to restore the 12/04/2003 version of the /home/abc/doc/ directory.

- ◆ If you do a regular restore, the restored directory has all files and directories that ever existed in /home/abc/doc/ from 12/01/2003 (last full backup) through 12/04/2003:

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

- ◆ If you do a true-image restore of the 12/04/2003 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on 12/04/2003:

```
file1
file2
file4
```

NetBackup does not restore *any* of the files deleted prior to the 12/04/2003 incremental backup.

The restored directory does not include the dirA and dirC subdirectories, even though they were backed up on 12/04/2003 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 additional 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 364.)
- ◆ Incremental backups are slower for a policy where true-image restore information is being collected.
- ◆ If you are using the indexing feature, the INDEX files take much more space when you are collecting true-image restore information. (This warning applies only to ASCII catalogs—the binary catalog does not need INDEX files.) For more information, see “Reduce Restore Times by Indexing the Image Catalog” on page 236.
- ◆ 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 show individual files or let you select them. The NetBackup user’s guides explain this further and provide instructions for performing true-image restores.
- ◆ A true-image restore preserves files that are currently in the directory but were not present when the backup was done. In our previous example, assume you created a file named file5 after the incremental backup occurred on 12/04/2003, but before doing the restore. In this case, the contents of the directory after the restore is:

```
file1  
file2  
file4  
file5
```

Compression

The **Compression** policy attribut specifies that software compression be used for backups of this policy. Select the box to enable compression (the default is 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 that 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 therefore the network load.

Disadvantages of Using Compression

Disadvantages of compression are that it increases computing overhead on the client and also 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 8 mm) 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, thus reducing the total time to back them up.

How Much Compression Can You Expect?

The degree to which a file can be compressed depends on the types of data. A backup usually involves more than one type of data. Examples include stripped and unstripped binaries, ASCII, and repeating non-unique strings. If more of the data is favorable to compression you obtain more compression.

Note When compression is not used, it is normal to receive slightly more data at the server than is on the client (on UNIX, this is as shown by `du` or `df`) due to client disk fragmentation and file headers added by the client.

Compression Specifications

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



Compression Specifications (continued)

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 other others with the same extension) from compression by adding it under the NetBackup host UNIX Client > Client Settings dialog.																																				
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> <p>Files that are equal to or less than 512 bytes, because that is the tar block size.</p> <p>On UNIX clients, files ending with suffixes specified with the COMPRESS_SUFFIX = <i>.suffix</i> option in the <code>bp.conf</code> file.</p> <p>On UNIX clients, files with the suffixes as shown below:</p> <table border="0" style="margin-left: 40px;"> <tr> <td><code>.arc</code> or <code>.ARC</code></td> <td><code>.gz</code> or <code>GZ</code></td> <td><code>.iff</code> or <code>.IFF</code></td> <td><code>.sit.bin</code> or</td> </tr> <tr> <td><code>.arj</code> or <code>.ARJ</code></td> <td><code>.hqx</code> or <code>.HQX</code></td> <td><code>.pit</code> or <code>.PIT</code></td> <td><code>.SIT.bin</code></td> </tr> <tr> <td><code>.au</code> or <code>.AU</code></td> <td><code>.hqx.bin</code> or</td> <td><code>.pit.bin</code> or</td> <td><code>.tiff</code> or <code>.TIFF</code></td> </tr> <tr> <td><code>.cpt</code> or <code>.CPT</code></td> <td><code>.HQX.BIN</code></td> <td><code>.PIT.BIN</code></td> <td><code>.Y</code></td> </tr> <tr> <td><code>.cpt.bin</code> or</td> <td><code>.jpeg</code> or <code>.JPEG</code></td> <td><code>.scf</code> or <code>.SCF</code></td> <td><code>.zip</code> or <code>.ZIP</code></td> </tr> <tr> <td><code>.CPT.BIN</code></td> <td><code>.jpg</code> or <code>.JPG</code></td> <td><code>.sea</code> or <code>.SEA</code></td> <td><code>.zom</code> or <code>.ZOM</code></td> </tr> <tr> <td><code>.F</code></td> <td><code>.lha</code> or <code>.LHA</code></td> <td><code>.sea.bin</code> or</td> <td><code>.zoo</code> or <code>.ZOO</code></td> </tr> <tr> <td><code>.F3B</code></td> <td><code>.lzh</code></td> <td><code>.SEA.BIN</code></td> <td><code>.z</code> or <code>.Z</code></td> </tr> <tr> <td><code>.gif</code> or <code>.GIF</code></td> <td><code>.pak</code> or <code>.PAK</code></td> <td><code>.sit</code> or <code>.SIT</code></td> <td></td> </tr> </table>	<code>.arc</code> or <code>.ARC</code>	<code>.gz</code> or <code>GZ</code>	<code>.iff</code> or <code>.IFF</code>	<code>.sit.bin</code> or	<code>.arj</code> or <code>.ARJ</code>	<code>.hqx</code> or <code>.HQX</code>	<code>.pit</code> or <code>.PIT</code>	<code>.SIT.bin</code>	<code>.au</code> or <code>.AU</code>	<code>.hqx.bin</code> or	<code>.pit.bin</code> or	<code>.tiff</code> or <code>.TIFF</code>	<code>.cpt</code> or <code>.CPT</code>	<code>.HQX.BIN</code>	<code>.PIT.BIN</code>	<code>.Y</code>	<code>.cpt.bin</code> or	<code>.jpeg</code> or <code>.JPEG</code>	<code>.scf</code> or <code>.SCF</code>	<code>.zip</code> or <code>.ZIP</code>	<code>.CPT.BIN</code>	<code>.jpg</code> or <code>.JPG</code>	<code>.sea</code> or <code>.SEA</code>	<code>.zom</code> or <code>.ZOM</code>	<code>.F</code>	<code>.lha</code> or <code>.LHA</code>	<code>.sea.bin</code> or	<code>.zoo</code> or <code>.ZOO</code>	<code>.F3B</code>	<code>.lzh</code>	<code>.SEA.BIN</code>	<code>.z</code> or <code>.Z</code>	<code>.gif</code> or <code>.GIF</code>	<code>.pak</code> or <code>.PAK</code>	<code>.sit</code> or <code>.SIT</code>	
<code>.arc</code> or <code>.ARC</code>	<code>.gz</code> or <code>GZ</code>	<code>.iff</code> or <code>.IFF</code>	<code>.sit.bin</code> or																																		
<code>.arj</code> or <code>.ARJ</code>	<code>.hqx</code> or <code>.HQX</code>	<code>.pit</code> or <code>.PIT</code>	<code>.SIT.bin</code>																																		
<code>.au</code> or <code>.AU</code>	<code>.hqx.bin</code> or	<code>.pit.bin</code> or	<code>.tiff</code> or <code>.TIFF</code>																																		
<code>.cpt</code> or <code>.CPT</code>	<code>.HQX.BIN</code>	<code>.PIT.BIN</code>	<code>.Y</code>																																		
<code>.cpt.bin</code> or	<code>.jpeg</code> or <code>.JPEG</code>	<code>.scf</code> or <code>.SCF</code>	<code>.zip</code> or <code>.ZIP</code>																																		
<code>.CPT.BIN</code>	<code>.jpg</code> or <code>.JPG</code>	<code>.sea</code> or <code>.SEA</code>	<code>.zom</code> or <code>.ZOM</code>																																		
<code>.F</code>	<code>.lha</code> or <code>.LHA</code>	<code>.sea.bin</code> or	<code>.zoo</code> or <code>.ZOO</code>																																		
<code>.F3B</code>	<code>.lzh</code>	<code>.SEA.BIN</code>	<code>.z</code> or <code>.Z</code>																																		
<code>.gif</code> or <code>.GIF</code>	<code>.pak</code> or <code>.PAK</code>	<code>.sit</code> or <code>.SIT</code>																																			

Encryption

The **Encryption** policy 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

The **Collect Disaster Recovery Information** policy attribute specifies whether or not you want NetBackup to collect the information required for intelligent disaster recovery during backups of Windows clients using this policy. (See “Configuring NetBackup Policies for IDR” on page 294 in the *NetBackup System Administrator’s Guide, Volume II*.)

Allow Multiple Data Streams

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.

- ◆ How many streams (backup jobs) start for each client and how the backup selection list is divided into separate streams is determined by the directives, scripts, or templates that you specify in the backup selection list. (See “Backup Selections List Directives for Multiple Data Streams” on page 133.)
- ◆ 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 “Tuning Multiple Data Streams” on page 95.)

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

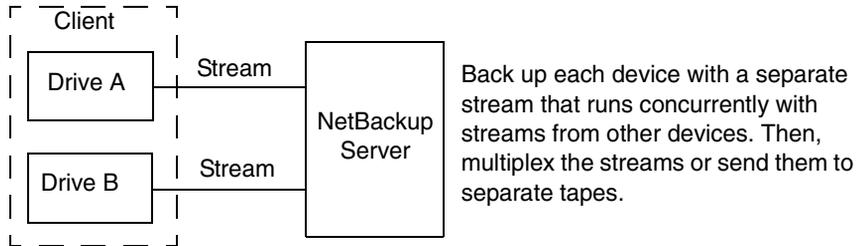
Multiple data streams can reduce the backup time for large backups. This is achieved by splitting the backup into multiple streams and then using multiplexing, multiple drives, or a combination of the two for processing the streams concurrently.

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.

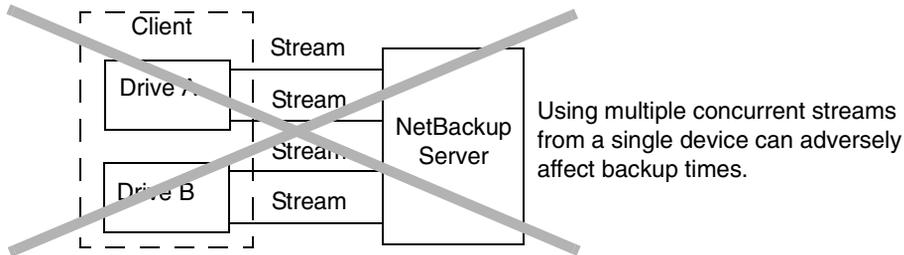


Note For best performance, use only one data stream to back up each physical device on the client. Multiple concurrent streams from a single physical device can adversely affect backup times because the heads must move back and forth between tracks containing files for the respective streams.

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** property applies to each stream. For example, if **Schedule Backup Attempts** is set to 3, NetBackup retries each stream a maximum of three times.

The Activity Monitor shows each stream as a separate job. Use the job details view to determine the files that are backed up by each of these jobs.



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.

Tuning Multiple Data Streams

The two aspects of multiple data streams that you can tune 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 133.)

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
- ◆ Number of drives that are available
- ◆ Maximum concurrent jobs settings for the policy and client

Each storage unit and each schedule has a maximum multiplex 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. MPX on storage unit 1 is set to 3 and MPX on storage unit 2 is set to 5. If MPX 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)**
- ◆ **Limit jobs per policy** (policy attribute)



- ◆ **Maximum Data Streams** (Set the number in **Host Properties > Master Servers > Client Attributes** or use the `bpclient` command `-max_jobs` option as shown below)

The maximum job settings are interdependent as follows:

- ◆ If **Maximum Data Streams** is not set, the lowest value of **Maximum Jobs Per Client** and **Limit Jobs Per Policy** is the limiting factor.
- ◆ If **Maximum Data Streams** is set, then NetBackup ignores **Maximum Jobs Per Client** and uses the lowest value of **Maximum Data Streams** and **Limit Jobs Per Policy** as the limiting factor.

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 (Optional)

The **Keyword Phrase** policy attribute specifies a key phrase that NetBackup will associate with all backups or archives for this policy. Windows and UNIX clients can then optionally list or restore only the backups that have this phrase associated with them (see the appropriate NetBackup user's guide). The user interfaces on NetBackup clients other than Windows and UNIX do not support keyword phrases.

You can use the same keyword phrase for more than one policy. This makes 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.

Windows and UNIX clients can also specify a keyword phrase for a user backup or archive. A user keyword phrase overrides the policy phrase.

Advanced Client Options

To use Advanced Client, 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*.

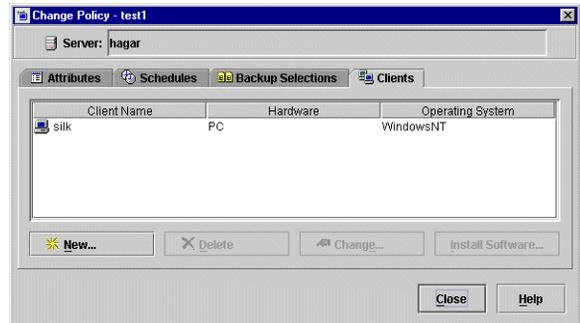


Which Clients Will Be Backed Up: Clients Tab

Add, delete, or change clients for a policy on the **Clients** tab. You can also install NetBackup software on UNIX client machines.

▼ 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/openv/netbackup/client` directory for the directories and software corresponding to the client you are trying to install. If the directories or software are not there, rerun the installation script on the server and choose the option to install client software (see the NetBackup getting started 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 or browse to find the name of the client.
4. Click the button to the right of the **Hardware and Operating System** field and select the desired entry.

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.

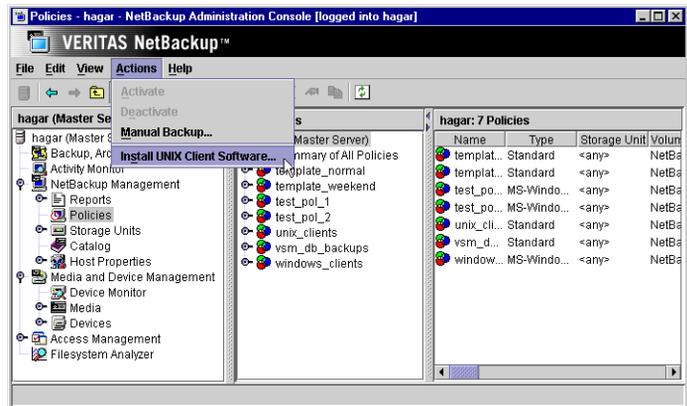


▼ To install UNIX client software

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.

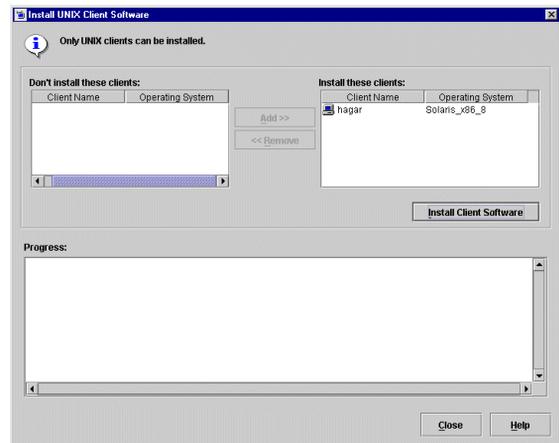
2. Select the master server name at the top of the All Policies middle pane.

3. Click **Actions > Install UNIX Client Software**. The Install UNIX Client Software dialog appears.



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/opencv/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.

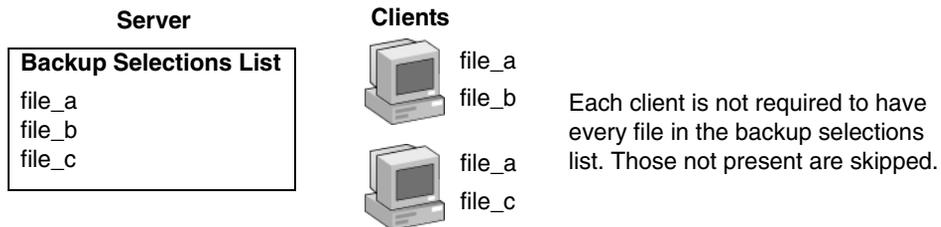
For information on configuring snapshots, see the *NetBackup Advanced Client System Administrator's Guide*.



Which Selections Will Be Backed Up: Backup Selections Tab

The backup selections list names the files, directories, directives, scripts (used for database policies), and templates (used for Oracle and DB2 policies), that NetBackup includes in automatic backups of clients covered by a policy. The policy backup selections list does not apply to user backups or archives because users select the objects to back up before starting the operation.

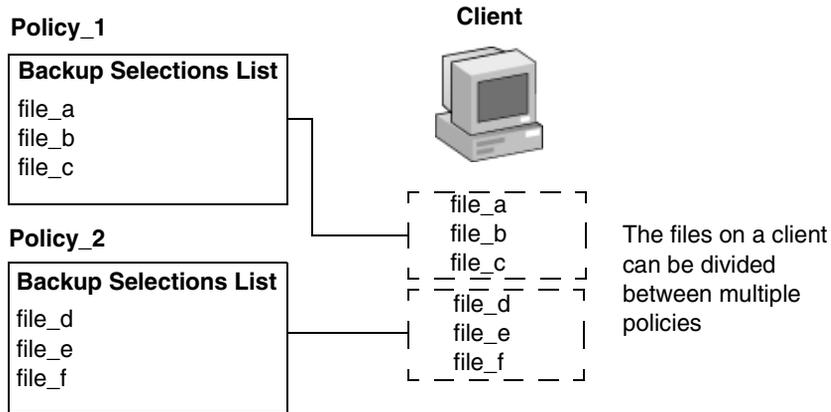
NetBackup uses the same selection list for all clients backed up according to this policy. All the files do not need to exist on all the clients, as NetBackup backs up the files that it finds. However, each client must have at least one of the files in the backup selections list or the client backup will fail with a status 71. Selection list entries are processed serially for each client, but it is possible to back up multiple clients in parallel if enough drives are available.



NetBackup processes the entries on the backup selection list one entry at a time and in the order that they appear in the backup selections. However, NetBackup backs up the files from multiple clients in parallel, assuming multiple storage devices are available and NetBackup attributes are set to allow it. (See “Setting the Number of Streams That Can Run Concurrently” on page 95.)

- ◆ 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).

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, NetBackup processes the files serially, which 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 95 for an explanation of maximum jobs settings that also applies 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 them in the same policy, schedule them at different times, or set **Maximum Jobs per Client** to 1.

Another way to reduce backup time is to use a single policy that has **Allow Multiple Data Streams** enabled and then add `NEW_STREAMS` directives to the backup selections list. For example:

```
NEW_STREAM
file_a
file_b
file_c
NEW_STREAM
file_d
file_e
file_f
```



The example 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 93.)

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 196.

Creating the Backup Selections List for Standard Policies

Standard policies and Exchange and Lotus Notes policies list pathnames and directives in the backup selection list.

For information on what directives accomplish, see “Backup Selections List Directives: General Discussion” on page 130 and “Backup Selections List Directives for Multiple Data Streams” on page 133 (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.

▼ 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:



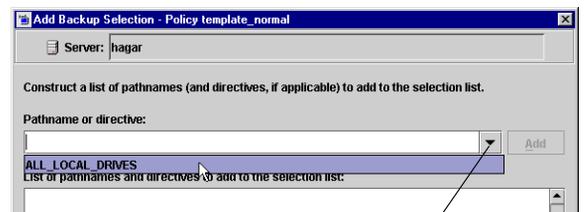
- ◆ Type the name of the path in the **Pathname or Directive** field.

If you are unfamiliar with how to specify file paths for a client:

For UNIX clients, see “File-Path Rules for UNIX Clients” on page 111.

For Windows clients, see “File-Path Rules for Microsoft Windows Clients” on page 122.

- ◆ 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.



Click to select a directive

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**.
7. To verify that the entries on the selections list are accurate, see “Verifying the Backup Selections List” on page 109.

Creating the Backup Selections List for Database Policies

Depending on the database type, the backup selection list of policies for databases contains different types of objects. For Exchange and Lotus Notes, the list contains pathnames and directories. 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 scrips and/or templates.

For information on what directives accomplish, see “Backup Selections List Directives: General Discussion” on page 130 and “Backup Selections List Directives for Multiple Data Streams” on page 133 (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.

▼ 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.



Creating the Backup Selections List for Oracle or DB2 Policies

An Oracle backup or XML export policy, or a DB2 backup policy, lists templates and/or scripts in the backup selection list. The templates and scripts listed are run during manual and automatic backups in the order in which they appear in the backup selection 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_database_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.

Verifying the Backup Selections List

After creating or modifying a backup selections list, complete the following procedure to make sure that the file-path rules 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.
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).

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:

```
02/02/03 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:

```
02/02/03 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:

```
02/02/03 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 Backup File Paths

The following topics explain the rules for specifying backup file paths for each type of NetBackup client:

- ◆ “File-Path Rules for UNIX Clients” on page 111
- ◆ “File-Path Rules for Microsoft Windows Clients” on page 122
- ◆ “File-Path Rules for NetWare NonTarget Clients” on page 127
- ◆ “File-Path Rules for NetWare Target Clients” on page 128
- ◆ “File-Path Rules for Clients Running Extension Products” on page 129
- ◆ “Backup Selections List Directives: General Discussion” on page 130
- ◆ “Backup Selections List Directives for Multiple Data Streams” on page 133

If you’re making an entry for a Vault job, see “To create a Vault policy” on page 196.

File-Path Rules for UNIX Clients

The general requirements for pathnames on UNIX clients are as follows:

- ◆ Enter one pathname per line. NetBackup supports a maximum path length of 1023 characters on UNIX clients.
- ◆ Start all pathnames with a slash (/).
- ◆ You can use the following meta or wildcard characters in policy backup selection lists:

```
*
?
[ ]
{ }
```

The following are example UNIX file specifications that use this capability:

```
/home/.[a-zA-Z0-9]*
/etc/*.conf
```

- ◆ To use meta or wildcard characters literally, precede them with a backslash (\). Assume, for example, that the brackets in the following pathname are used as literal characters:

```
/home/abc/fun[ny]name
```

In the backup selection list, precede the brackets with a backslash as in

```
/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.

The following topics provide more information on specifying UNIX file paths to back up.

Notes on Backup Selection Lists for UNIX Clients

- ◆ File paths that cross mount points or that the client mounts through NFS can affect the way that you must configure your backups. Before creating a backup selection list, familiarize yourself with the **Follow NFS** and **Cross mount points** attributes.
- ◆ You can back up operating system, kernel, and boot files with NetBackup. You cannot, however, create bootable tapes. Consult your system documentation to create a bootable tape.
- ◆ NetBackup never backs up the following:
 - ◆ 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**.
 - ◆ 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 `cd` into.
 - ◆ 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.
- ◆ Exclude specific files from backups by creating an exclusion list on the client.
- ◆ The `BUSY_FILE_ACTION` and `LOCKED_FILE_ACTION` options in the `/usr/obj/usr/lib/netbackup/bp.conf` file on the client offer alternatives for handling busy and locked files. See “NetBackup Configuration Options” on page 134.
- ◆ 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 119.)
- ◆ 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 119.)
- ◆ On Hewlett-Packard and Solaris 2.5 (and later) platforms, NetBackup backs up VxFS extent attributes.
- ◆ 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
```

Symbolic Links to Files or Directories

For symbolic (soft) links, include the file path to the source file in your 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.

Note If NetBackup restores a symbolic link as root, it changes the owner and group back to the original owner and group. When NetBackup restores a UNIX symbolic link as a nonroot user, it 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, it 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 warn you to avoid using these links.

NetBackup does not back up and restore hard-linked directories in the same manner as it does files:

- ◆ During a backup, if NetBackup encounters hard-linked directories, it backs them up multiple times, 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.

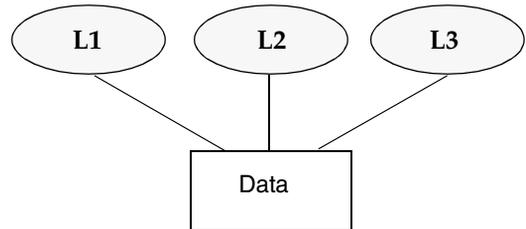
During a backup, if the backup selection list includes hard-linked files, the data is backed up only once, using the first file name reference found in the directory structure. If a second or subsequent file name reference is found, it is backed up as a link to the name of the first file. This means 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 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.

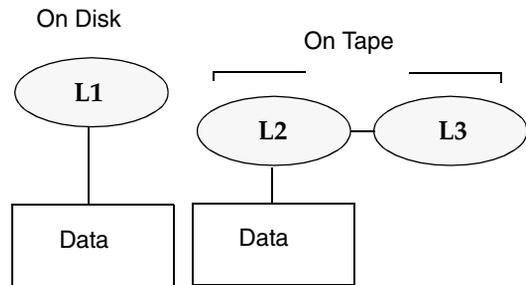
Example 1

Assume there are three hard links named L1, L2, and L3 that are pointing to the same data as shown in the figure below.

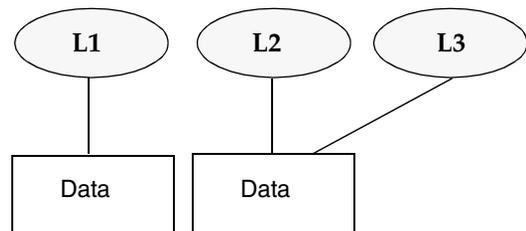
1. The three files are all hard linked to the same data.



2. L2 and L3 are backed up to tape and 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.



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. The restore therefore fails and you see 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 it 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 are 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 and then deletes the device file associated with the raw partition. However, 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 `bpemd_notify` scripts to unmount and remount the backed-up file systems.
- ◆ The **Cross Mount Points** 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 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 Solaris:

```
/devices/sbus@1,f8000000/esp@0,800000/sd@2,0:1h
```

Caution Do not specify wildcards (such as `/dev/rsd*`) in paths 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 file path 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.



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 146.)

Backup and Restore of Extended Attribute Files and Named Data Streams

NetBackup can back up and restore the following file attributes:

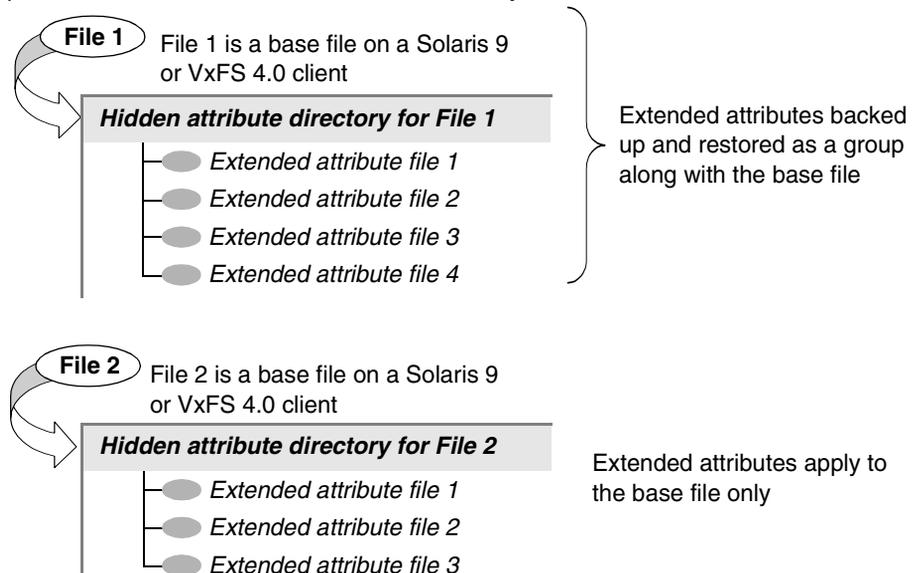
- ◆ Extended attribute files of the Solaris 9 UNIX File System (UFS) and temporary file system (TMPFS)
- ◆ Named data streams of the VxFS 4.0 file system

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.

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.

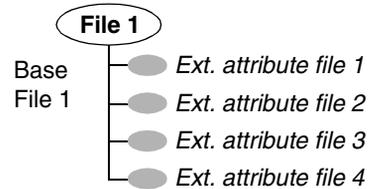
This is especially true in the case of incremental backups, during which NetBackup checks the mtime or ctime of each file individually.

Restoring Extended Attributes or Named Data Streams

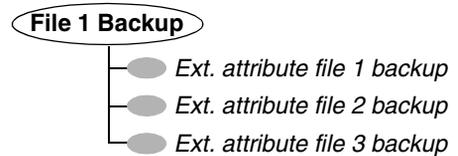
If **Overwrite Existing Files** is selected as a restore option in the Backup, Archive, and Restore client interface, and a file possessing extended attributes or named data streams is being restored, any existing attribute files or named data streams for that base file are replaced with the restored 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,

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`” on page 241 in *NetBackup System Administrator’s Guide, Volume II*.)

Note Extended attributes and named data streams cannot be compressed.



File-Path Rules for Microsoft Windows Clients

The following describes the conventions to use when specifying backups for Windows clients.

File Backups

You can use either Microsoft Windows conventions or UNIX file-path conventions, whichever you are the most comfortable with. You can also mix the two styles within the same backup selection 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:\
```

- ◆ 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 file path is 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:

```
*  
?
```

The following backs up all files ending with .doc

```
c:\Users\*.doc
```

The following backs 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
```

The drives that are not backed up include: floppy disks, CD-ROMs and drives that are located on remote systems but mounted on your system through the network.

The following is an example backup selection list that uses the Microsoft Windows conventions:

```
c:\  
d:\workfiles\  
e:\Special\status  
c:\tests\*.exe
```

Using UNIX Conventions

NetBackup permits you to use UNIX conventions in the backup selection list for Windows clients. This is convenient if your configuration has mainly UNIX clients and you are more comfortable with UNIX conventions.

The rules for the UNIX conventions are the same as explained for Microsoft Windows clients, except that you:

- ◆ 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
```



Windows Disk-Image (Raw) Backup

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.

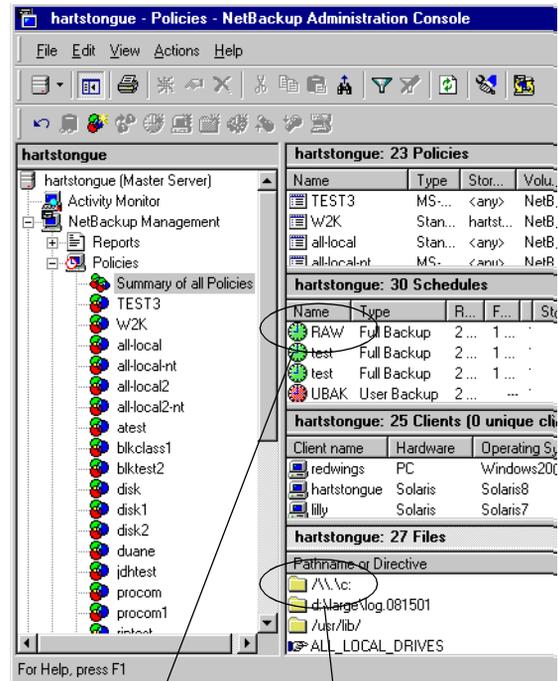
When performing a disk-image backup, be sure to select **Full Backup** for the backup type. Any other backup type will not work for backing up a disk-image.

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 would back 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:\
```



Must select Full Backup as backup type

Logical drive name in the backup selection list

Note Before starting a disk-image backup, NetBackup locks the logical drive to ensure that no changes occur during the backup. If there are open files on the logical drive, a disk-image backup is not performed.

Note 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.

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 details, see the *NetBackup User's Guide for Microsoft Windows*.

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

```
%systemroot%\system32\config
```

- ◆ On Windows 98 or 95, this directory is

```
%systemroot%
```

Where `%systemroot%` is the directory where Windows is installed.

For example, if Windows NT 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)
```

```
:\ (backs up all local drives except those that are removable)
```

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.

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.

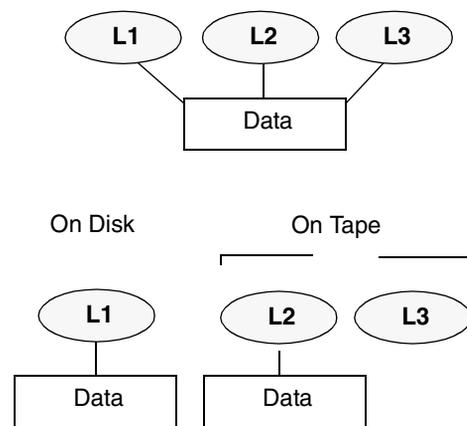
During a backup, if the backup selection list includes hard-linked files, the data is backed up only once, using the first file name reference found in the directory structure. If a second or subsequent file name reference is found, it is backed up as a link to the name of the first file. This means 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.

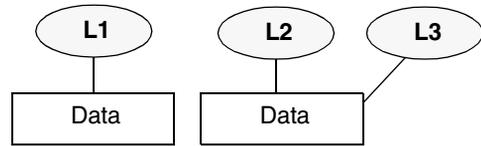
Example 1

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 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.

File-Path Rules for NetWare NonTarget Clients

For NetWare systems that are running the NonTarget version of NetBackup client software, specify the pathnames in the 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 file path is 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.
- ◆ 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*.

File-Path Rules for NetWare Target Clients

For NetWare clients that are running the target version of NetBackup client software, use the following format for the file paths:

```
/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/
```

File-Path Rules for Clients Running Extension Products

File-path 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 cause NetBackup to perform specific actions when processing the files in the list.

The directives that are available depend on the policy type and whether the **Allow Multiple Data Streams** attribute is enabled for the policy. The following is an example of a backup selections list that contains the `NEW_STREAM` directive and is for an MS-Windows-NT policy that has **Allow Multiple Data Streams** enabled:

```
NEW_STREAM
D:\Program Files
NEW_STREAM
C:\Winnt
```

The example above shows how directives appear in a backup selections list. The actions that the `NEW_STREAM` directive causes are explained in “Backup Selections List Directives for Multiple Data Streams” on page 133.

The rules for specifying backup paths in the backup selections list apply regardless of whether directives are used.

ALL_LOCAL_DRIVES Directive

Use the `ALL_LOCAL_DRIVES` directive to back up all local drives except for those 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` Directives” on page 137.)

NetBackup automatically excludes the following file system types on most platforms:

- ◆ `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)
- ◆ `cdrom` (all UNIX platforms)
- ◆ `cachefs` (AIX, Solaris, SGI, UnixWare)

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 Windows2000\XP, the list of items that get backed up can include:

- ◆ Active Directory
- ◆ COM+ Class Database
- ◆ Cluster Database
- ◆ IIS Database
- ◆ Registry
- ◆ Boot Files and Protected Files
- ◆ SYSVOL
- ◆ Certificate Server

On an NT system, 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

This directive affects Windows 2003 Server computers 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 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
 - ◆ 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 includes 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 133.

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:

- ◆ 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 above policy types 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 them 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 Directive
- ◆ ALL_LOCAL_DRIVES Directives
- ◆ UNSET and UNSET_ALL Directive



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, we 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.

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` Directives” on page 137.
- ◆ If wild cards are used, the expansion of the wild cards results in one stream per wild card 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)
Postprocess Image: immediately
Display Reports: 24 hours ago
Keep TIR Info: 1 days
Prep Interval: 12 hours

```

Example - Auto-Discover Streaming Mode

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 Directives

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 the directive causes depends on whether you also enable **Allow Multiple Data Streams** 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 136.
- ◆ 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.

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: \.

UNSET and UNSET_ALL Directive

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 133.) 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

Unsets 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.

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
```

Creating an Exclude List on a UNIX Client

If you create a `/usr/opensv/netbackup/exclude_list` file on a UNIX client, NetBackup uses the contents of the file as a list of patterns to skip during automatic full and incremental backups.



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

The following types of files typically appear in an exclude list:

- ◆ *.o files
- ◆ core files
- ◆ a.out files
- ◆ Files prefixed or suffixed by ~ (backups for editors)
- ◆ Files and directories under /tmp, /usr/tmp
- ◆ Man pages
- ◆ Software that you can restore from original installation tapes
- ◆ Automounted directories
- ◆ CD-ROM file systems
- ◆ NetBackup automatically excludes the following file system types:
 - ◆ mntfs (Solaris)
 - ◆ proc (all UNIX platforms)
 - ◆ cdrom (all UNIX platforms)
 - ◆ cacheefs (AIX, Solaris, SGI, UnixWare)

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

Check with users before excluding any files from their backups.

Syntax Rules

The following syntax rules apply to exclude lists:

- ◆ Blank lines or lines beginning with a pound sign (#) are ignored.
- ◆ 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 non-wildcard characters), precede them with a backslash (\). For example, assume the brackets in the following are to be used literally

```
/home/abc/fun[ny]name
```

In the exclude list, precede them with a backslash as in

```
/home/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.

- ◆ If you exclude all files in the backup selections list by using / or * or both symbols together (/*), NetBackup backs up only what is specified by full path names in the include list.
- ◆ 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

```
/home/testfile (with no extra space character at the end)
```

and your exclude list entry is

```
/home/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, /home/test/). If the pattern does not end in / (for example, /usr/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 without a preceding slash. For example:

```
test
```

rather than

```
/test
```

This is equivalent to prefixing the file pattern with

```
/
```

```
/*/
```

```
/**/*.
```

```
/**/*./*
```



and so on.

- ◆ 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`.

Example of an Exclude List

In this example, an exclude list contains the following entries:

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

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

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

Exclude Lists for Specific Policies or Schedules

NetBackup allows you to create an exclude list for a specific policy or a policy and schedule combination. To do this, create an `exclude_list` file with a `.policyname` or `.policyname.schedulename` suffix. The following are two examples for a policy named `wkstations` that contains a schedule named `fulls`:

```
/usr/opensv/netbackup/exclude_list.wkstations
/usr/opensv/netbackup/exclude_list.wkstations.fulls
```

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

For a given backup, NetBackup uses a single exclude list—the list containing the most specific name. For example, if there are files named:

```
exclude_list.wkstations and exclude_list.wkstations.fulls
```

NetBackup uses only:

```
exclude_list.workstations.fulls
```

Creating an Include List on a UNIX Client

To add back in files that you eliminate with the exclude list, create a `/usr/opensv/netbackup/include_list` file. The same syntax rules apply as explained previously for the exclude list.

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

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

In this case, add back in a file named `/home/jdoe/test` by creating a `/usr/opensv/netbackup/include_list` file on the client and adding the following to it:

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

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

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

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

For a given backup, NetBackup uses only one include list and that is the one with the most specific name. For example, assume there are files such as the following:

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

In such a case, NetBackup uses only the following:

```
include_list.workstations.fulls
```

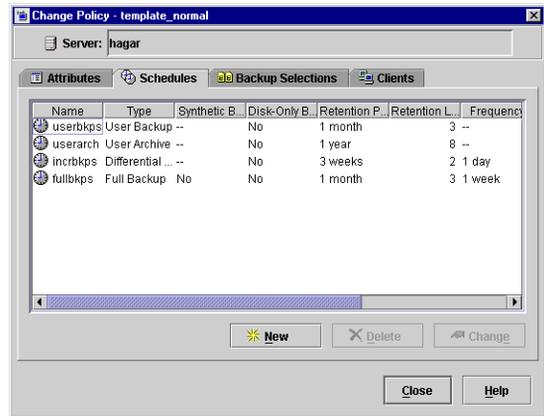


When Will the Job Run: Schedules Tab

The policy Schedules tab displays the different time schedules set up for the policy selected.

From the policy Schedules tab:

- ◆ Create a new schedule by clicking **New**.
- ◆ Delete an existing schedule by selecting the schedule and clicking **Delete**.
- ◆ Double-click an existing schedule to edit the schedule or select the schedule and click **Change**.



Creating or editing a schedule causes a second Schedule dialog to appear: the Add New Schedule or the Change Schedule dialog.

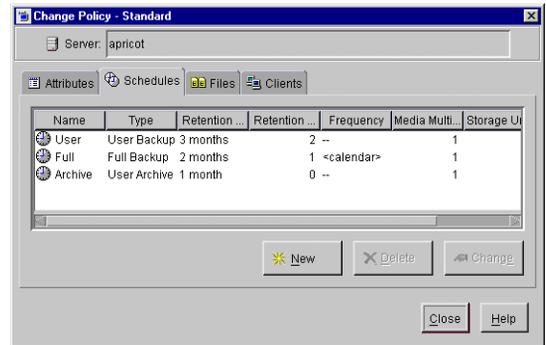
Creating or Editing a Schedule

When creating or editing a schedule, policy schedule attributes appear on four tabs in the Add New Schedule or Change Schedule dialog. The tabs allow you to schedule the days or dates on which a task will run.

- ◆ **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.

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



Note “To add or change schedules in a policy” on page 70 also provides information on changing existing policies.

Schedule Attributes Tab

The Attributes tab can be displayed by a number of methods. For example:

- ◆ By double-clicking a schedule in the Schedules tab.
- ◆ By clicking the **New** button in the Schedules tab.
- ◆ By clicking the **Disk Staging Schedule** button when configuring a relocation schedule for a disk staging storage unit.

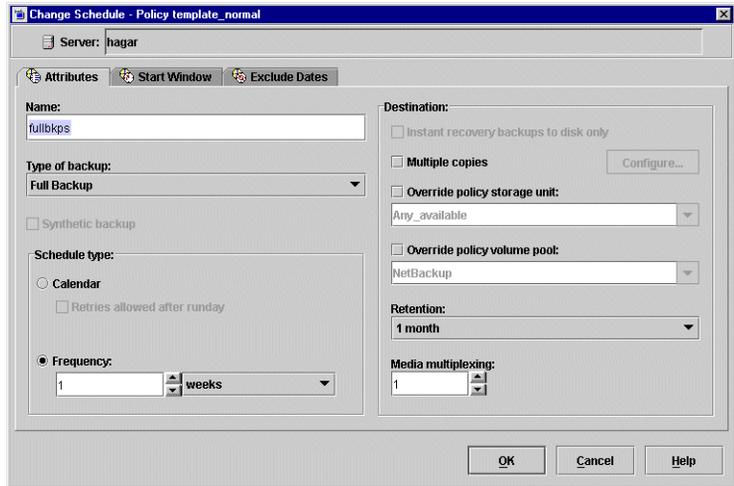
The Attributes tab contains options that define the backup type, when the backup can occur, and how long the backup image is to be kept. Other attributes such as type of storage and volume pool can also be defined.



Name

Identifies the schedule and appears on screens and messages from NetBackup. Specify a name by typing in the box. The name must be unique and can contain alphabetic (ASCII A-Z a-z), numeric (0-9), plus (+), minus (-), underscore(_), or period (.) characters. Do not use a minus or period as the first character. Do not use a period as the first or last character. Do not or leave spaces between characters.

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.



Type of Backup

The **Type of Backup** specifies the type of backup that a schedule will control. The list displays only the backup types that apply to the policy you are configuring.

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 is a list of possible selections:

◆ Full Backup

Backs up all the files that are specified in the backup selections list for the policy, regardless of when they were last modified or backed up. These backups occur automatically according to the criteria in the schedule. If you use incremental backups, you must also schedule full backups to perform a complete restore. If you're performing a raw partition backup, you must select **Full Backup**.

◆ Cumulative Incremental Backup

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

Note NetBackup 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 329.

For more information on configuring a Vault policy, see “More on Incremental Backups” on page 148.

◆ Differential Incremental Backup

Backs up all files that are specified in the backup selections list for the policy that have changed since the last successful incremental or full backup. All files are backed up if no prior backup has been done. These backups occur automatically according to the criteria in the schedule. A complete restore in this instance 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 configuring a Vault policy, see “More on Incremental Backups” on page 148.

◆ User Backup

Initiated by the user through the client interface (Backup, Archive, and Restore interface) and backs up all files that the user specifies. Users can start backups only during the times that you specify in 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 232.

◆ User Archive

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

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 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 that 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 that 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 that 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 that 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**

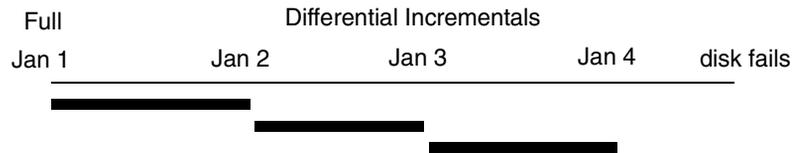
An automatic Vault session. Applies only to Vault policies. This does not run a backup, but instead runs the vault command specified in the Vault policy's backup selections list. In this way it starts an automatic, scheduled vault session or vault eject operation.

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

More on Incremental Backups

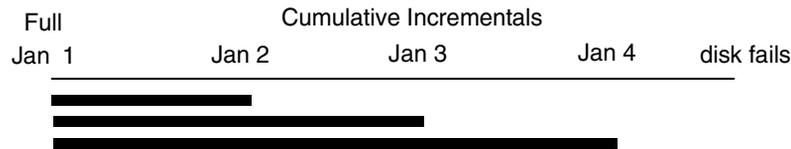
The example below shows the data included in a series of backups between January 1 and January 4. The January 1 backup is a full backup and includes all files and directories in the policy backup selections list. The subsequent backups are differential incremental

backups and 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 and all three of the incremental backups are required for the recovery.



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

A cumulative incremental backs up the data that has changed since the last full backup. The example below shows the data included in a series of 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 (incr)

The following table compares the 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.



Backup and Restore Times

The following table compares the 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

This 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

- ◆ 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 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 backup is a backup assembled from a 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.

Advantages of Using Synthetic Backups

There are a number of reasons why implementing synthetic backups may be useful in your NetBackup configuration.

Processing Takes Place on Master and Media Server(s) Instead of Client

One advantage of synthesizing a full backup lies in where the bulk of the processing takes place. 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 incremental backup has been run.

Reduce 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 containing the component images can be recycled, thereby reducing the number of tapes in use.

Use Drives More Effectively

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 87.)
- ◆ 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 155.)
- ◆ The master servers, media servers, and clients must all have NetBackup version 5.0 or later installed in order to synthesize backups.
- ◆ Generally, a tape library with multiple drives for reading and writing is needed, though not required. (See “Storage unit considerations:” on page 161.)
- ◆ NetBackup must be configured to use the binary catalog format from NetBackup version 5.0 or later. Backup images using the ASCII catalog format cannot be synthesized. (See “Notes on Synthetic Backups” on page 161.)

Schedules to Include in a Policy for Synthetic Backups

A policy for synthetic backups must contain at least three types of schedules:

- ◆ One schedule for a full, non-synthesized backup.

The traditional, full backup schedule must run successfully before the synthetic full backup schedule. The synthetic backup job will fail for a policy that contains a full synthetic backup schedule, but does not contain a traditional full backup schedule.
- ◆ 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 fail for a policy that contains full or incremental synthetic backup schedules and a traditional full backup schedule, 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.



Only one full and/or one cumulative schedule per policy is allowed to have the **Synthetic Backup** option selected. Even if multiple streams and clients are specified in the policy, only one synthetic job will run at a time for the policy. This prevents more than one job from starting at the same time and competing for the same resource.

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.

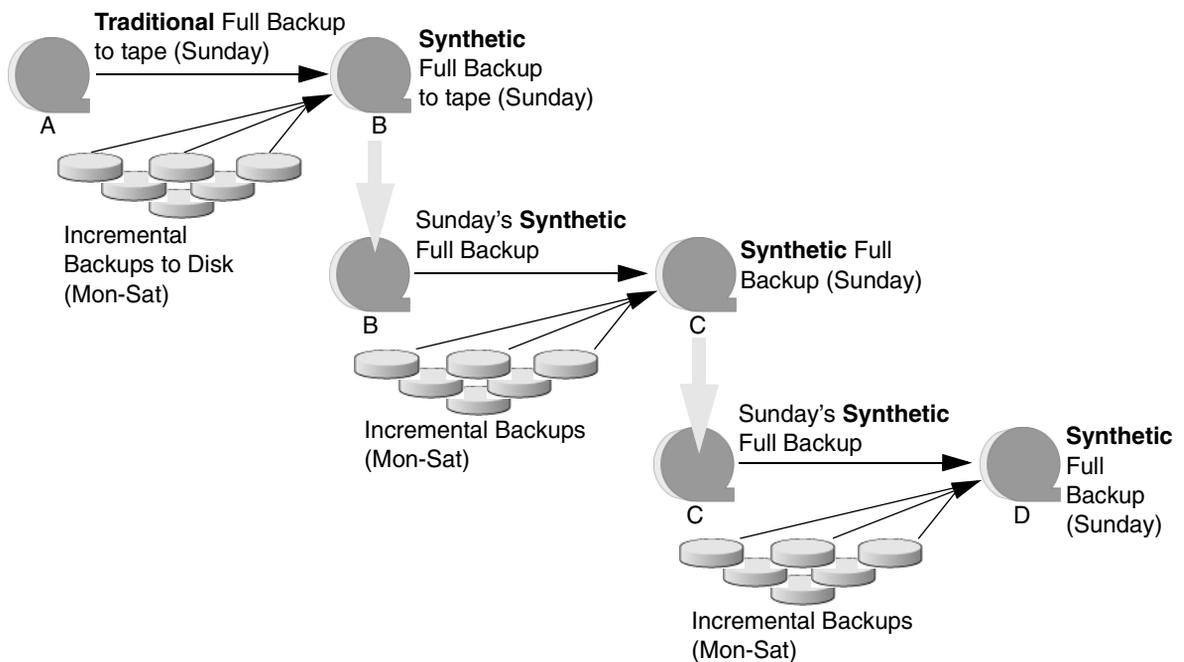
Two Types of Synthetic Backups

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 initial 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 158.

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.



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 one Sunday every 3 months.
 - ◆ 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 on the first scheduled Sunday. 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.

5. Per schedule, run a traditional full backup for the clients one Sunday every three months.

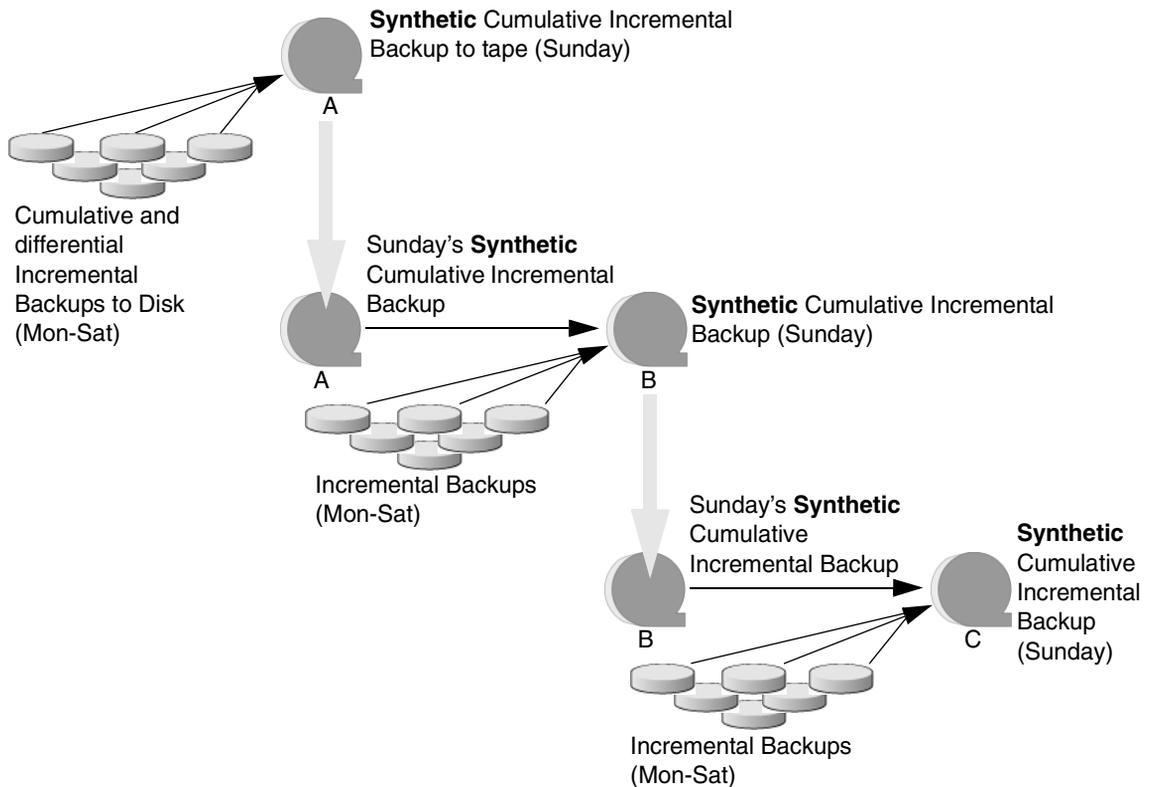
Synthetic Cumulative Incremental Backups

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 traditional full backup to run one Sunday every 3 months.
 - ◆ 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 on the first scheduled Sunday. 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 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.

5. Per schedule, run a traditional full backup for the clients one Sunday every three months.

Recommendations for Synthetic Backups

Scenario in which synthesized backups would be most beneficial:

The set of NetBackup clients to be backed up experience a moderate rate of change in their file systems every day. (Approximately 5% of the files change every day, for example.)

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 this is a concern in a NetBackup environment, to reduce a potential gap in backup coverage, consider running an incremental backup just prior to the synthetic backup.

Disk-based images are more efficient for synthesizing:

While synthesizing a backup, NetBackup processes the newest source 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.

Storage unit considerations:

- ◆ When performing tape backups, a tape outside of the tapes where full and differential images reside, is required for the formation of a synthetic image.
- ◆ If the maximum drive usage is set to utilize only a single drive (see “Maximum Concurrent Drives Used for Backup” on page 49), the generation of a synthetic image results in the maximum drive usage being ignored and multiple drives being utilized in the creation of synthetic images.
- ◆ 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**General Notes:**

- ◆ Synthetic backups are supported on the following NetBackup server platforms: Sun Solaris, Microsoft Windows, AIX, Hewlett-Packard, and Linux Red Hat.
- ◆ 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.

NetBackup must be configured to use the binary catalog format:

Synthetic backups cannot be created from catalogs in ASCII format.

Catalogs converted from ASCII to binary using `cat_convert` will not work. Synthetic backups and the associated catalog are identical to those of non-synthetic backups, except that the image time for synthetic backups is one second after the time of the latest component image, and not the time that the synthetic backup was actually created.

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 time stamp 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 time stamp. 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 169.) The two sets of catalogs have the same expiration time.

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

True Image Restore and Synthesized Backups

Since **True Image Restore with Move Detection** 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 host property, **Keep True Image Restoration (TIR) Information**. (See “Keep True Image Restoration (TIR) Information” on page 364.)

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 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 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 Change Journal host property is enabled for a client, 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 327.)

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, the scheduler (`bpsched`) starts program `bpsynth` to manage the synthetic backup process. (An administrator may also run `bpsynth` directly from the command line.) `bpsynth` plans how the synthetic backup will be built from the previous backup images.

`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`, `bpsched`, `bpcd`.

On the media server(s): `bptm` (tape images), `bpdm` (disk images), `bpcd`.



Note that several media servers may be involved if the component images are on different nodes.

Only one `bpsnth/byppcoord` pair runs at a time for any policy. This means to make a synthetic backup for several clients in a policy, the clients will be backed up one after another. Also, if the policy uses multiple streams, each stream will be synthesized one after the other.

Multiple streams and multiplexing provide parallelism that speed up the total traditional backup time. Multiplexed and multi-stream images are processed sequentially (one client or one stream at a time) during a synthetic backup. The sequential processing of these images during a synthetic backup may be acceptable if the synthetic backups take place during periods of relative inactivity on the media server. Better performance can be achieved if the component images are on disk (disk or disk-staging storage units).

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" on page 179 in the *NetBackup System Administrator's Guide, Volume II*.)

- ◆ Windows: For each Windows client, select **Perform Incrementals Based on Archive Bit.** (**NetBackup Management > Host Properties > Clients > Selected client(s) > Windows Client.**)

Instant Recovery Backups to Disk Only

The **Instant Recovery Backups to Disk Only** option is 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.

Calendar Schedule Type

Calendar-based scheduling allows administrators to specify run day options for a task. Choosing the **Calendar** schedule option causes the **Calendar Schedule** tab to appear in the Change Schedule dialog. For details on calendar-based scheduling, see “Calendar Schedule Tab” on page 176.

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.

For details on how calendar-based scheduling works with time windows, see “How Calendar Scheduling Interacts with Daily Windows” on page 179.

Retries Allowed After Runday

Select **Retries Allowed After Runday** to have the scheduler attempt to complete this schedule until the backup is successful. With this option 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 this 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.

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 the NetBackup scheduler 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 “Factors Affecting Backup Time” on page 244.

Multiple Copies

Using the **Multiple Copies** option (sometimes referred to as Inline Tape Copy or ITC), 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.

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 364.)

If you want 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 Drives Used for Backup** setting).

You can write multiple images to the following storage units:

- ◆ Media Manager storage units (except optical devices)
- ◆ Disk storage units

◆ Disk staging storage units

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: NDMP, third-party copies, or optical devices.

Also, **Multiple Copies** does not support storage units that use a QIC (quarter-inch cartridge) drive type.

▼ **To configure a schedule to create multiple copies during a backup**

1. Multiple copies can be created for a regular backup policy or for a disk staging storage unit relocation schedule:

In a policy, created for a regular backup:

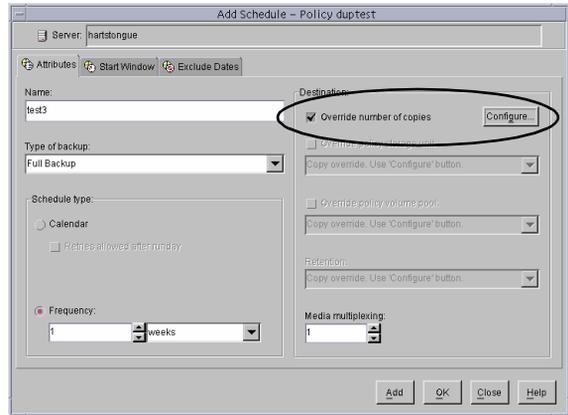
Expand **NetBackup Management > Policies**. Double-click an existing policy or select **Edit > New** to create a new policy.

For a relocation schedule, created as part of a disk staging storage unit:

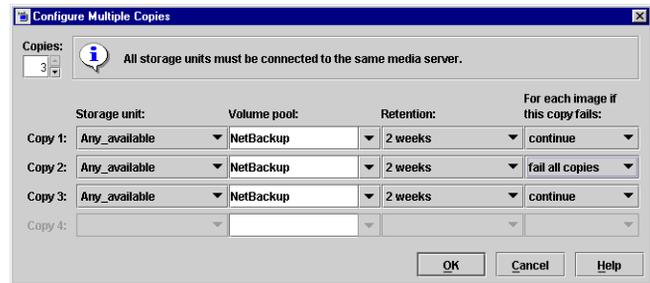
- a. Expand **NetBackup Management > Storage Units**. Double-click an existing storage unit or select **Actions > New > Storage Unit**.
 - b. Select **Disk Staging** as the storage type and configure the other storage unit selections.
 - c. Click the **Disk Staging Schedule** button.
2. Select the Schedules tab.
 3. Double-click an existing schedule or click **New** to create a new schedule.



- In the Attributes tab, select **Multiple Copies**, then click **Configure**.



- 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 364.)



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

- 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.
- Specify the volume pool where each copy will be stored.
- Select the retention level for each copy. (See “Retention” on page 169.)
- 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.

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** setting 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 76.)

Override Policy Volume Pool

The **Override Policy Volume Pool** setting 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 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.

Retention

The **Retention** setting specifies how long NetBackup retains the backups it creates according to this schedule. To set the retention period, select a **Retention** 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. For example, if you choose two weeks, you can restore the data from a backup done by this schedule for only two weeks after the backup.

For 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.

For 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.

For 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.

Default 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 380.) The default choices are shown below.

Level	Period	Level	Period
0	1 week	5	3 months
1	2 weeks	6	6 months
2	3 weeks	7	9 months
3	1 month	8	1 year
4	2 months	9 through 24	infinite

Note The levels are index numbers that correspond to the retention period (for example, the default retention period for level 0 is one week). The retention levels are shown here for reference as NetBackup uses them in some reports. NetBackup also uses the level when determining the volume to use for storing a backup. (See “Mixing Retention Levels on Backup Volumes” on page 171.)

Precautions For Assigning Retention Periods

- ◆ 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), 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 platter from the Media Manager configuration upon expiration, or Media Manager will reallocate the platter for future backups (even though WORM can be written only once).

Mixing Retention Levels on Backup Volumes

By default, NetBackup stores each backup on a volume that has existing backups at the same retention level (the period is not checked). For example, 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 than the previous backup, 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. Otherwise, for example, one small backup with an infinite retention prevents a volume from being reused, even if all other backups on the volume have expired.

If you want to mix retention levels on volumes, **Host Properties > Master Server**. Double-click a server and select **Media**. Select **Allow multiple retentions per media** or add `ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA` to the `bp.conf` file. (See “NetBackup Configuration Options” on page 134.)

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** setting 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. For more information on configuring multiplexing and the ramifications of using multiplexing, see “Multiplexing” on page 100 in the *NetBackup System Administrator’s Guide, Volume II*.

Specify a number from 1 through 32, where 1 specifies no multiplexing.

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.

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.

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.

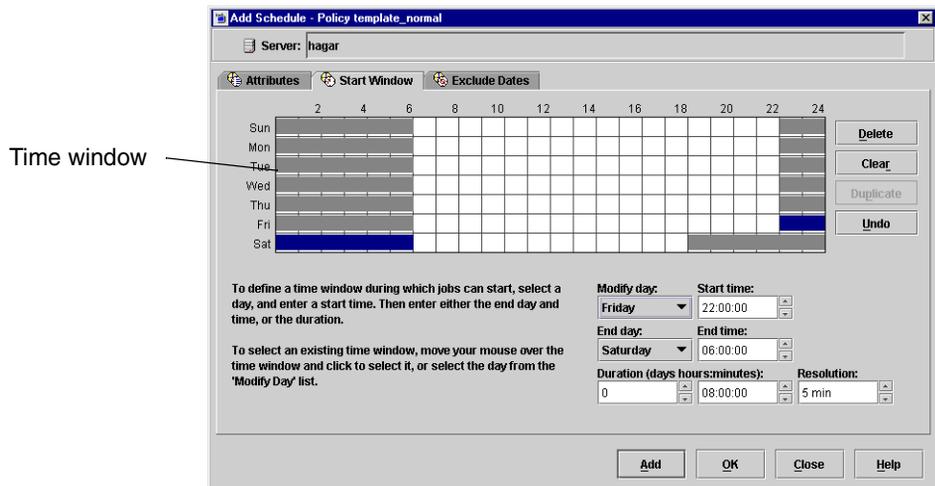
Start Window Tab

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

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:



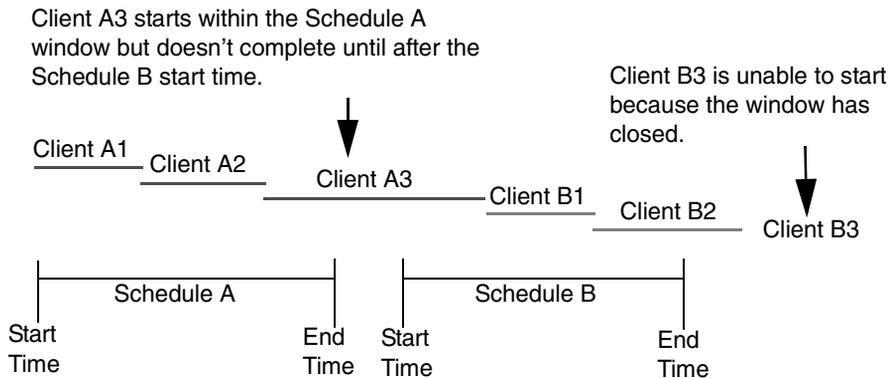
Click the arrow to the right of **End day** and select the last day in the time window. Then, click the up and down arrows to the right of **End time** to select when the time window will end.

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 the second schedule (B) begins shortly after the end time for the previous schedule (A). Both schedules have three clients with backups due.



The backup for client A3 in Schedule A does not finish until well 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 that 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. You may want to exclude, for example, the dates of holidays.

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. There are two methods to exclude a date from the schedule:
 - ◆ 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 for your task, 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.

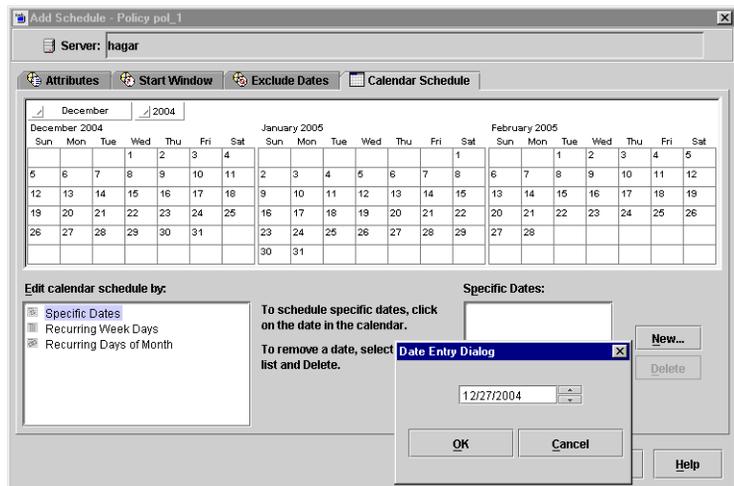
The Calendar Schedule tab displays a 3-month calendar. Use the controls at the top of the calendar to change the month or year.

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 for your task, select another tab to make changes or click **OK** to 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, to 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 for a particular day to select that day or to clear it.

Add Schedule - Policy test1
Server: silk

Attributes Start Window Exclude Dates Calendar Schedule

November 2003 2003

November 2003							December 2003							January 2004						
Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat
						1	7	8	9	10	11	12	13	4	5	6	7	8	9	10
2	3	4	5	6	7	8	14	15	16	17	18	19	20	11	12	13	14	15	16	17
9	10	11	12	13	14	15	21	22	23	24	25	26	27	18	19	20	21	22	23	24
16	17	18	19	20	21	22	28	29	30	31				25	26	27	28	29	30	31
23	24	25	26	27	28	29														
30																				

Edit calendar schedule by:

- Specific Dates
- Recurring Week Days
- Recurring Days of Month

A check indicates that the task is scheduled to run on that day of its respective week.

	Sun	Mon	Tue	Wed	Thu	Fri	Sat
1st	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				
2nd	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				
3rd	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				
4th	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				
Last	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>				

Set All Clear All

Add OK Close Help

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 for your task, select another tab to make changes or click **OK** to close the dialog.

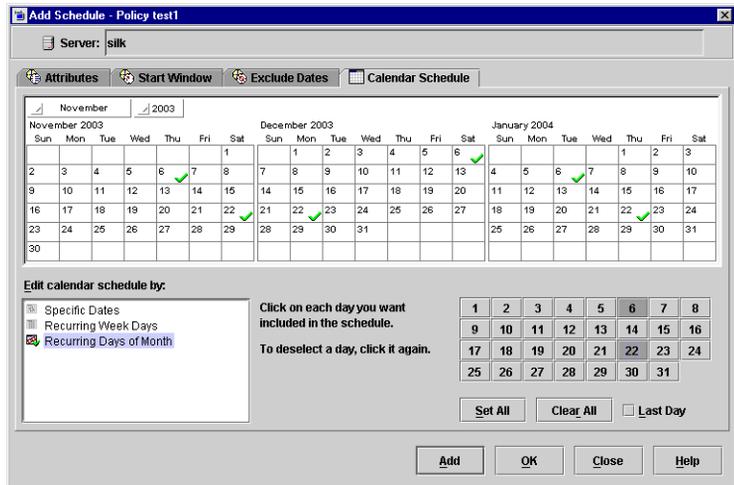


Schedule by Recurring Days of the Month

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 for your task, select another tab to make changes or click **OK** to 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.

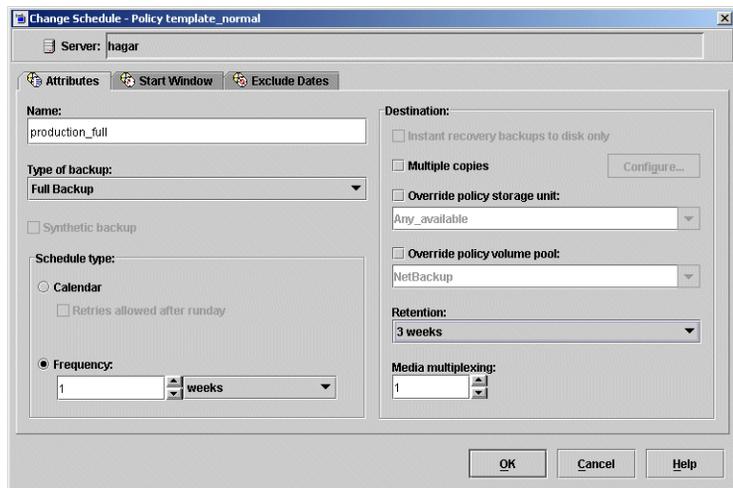
For details on how calendar-based scheduling works with backup windows, see “How Calendar Scheduling Interacts with Daily Windows” on page 179.

Example 1

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

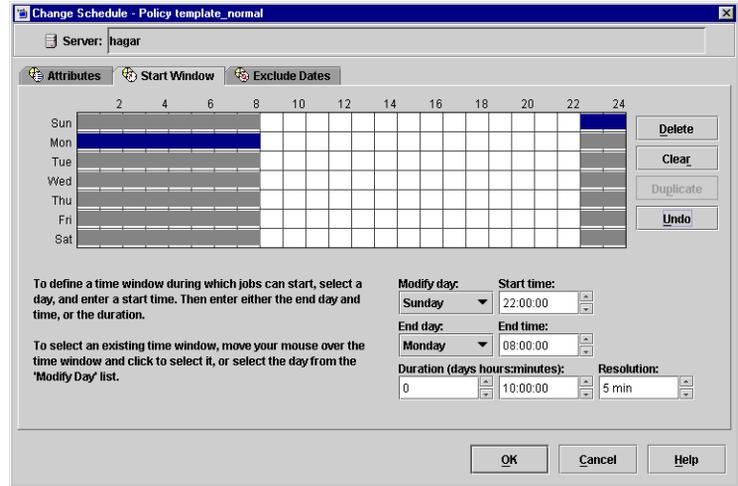
Schedule Runs Every Day (recommended method)

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.

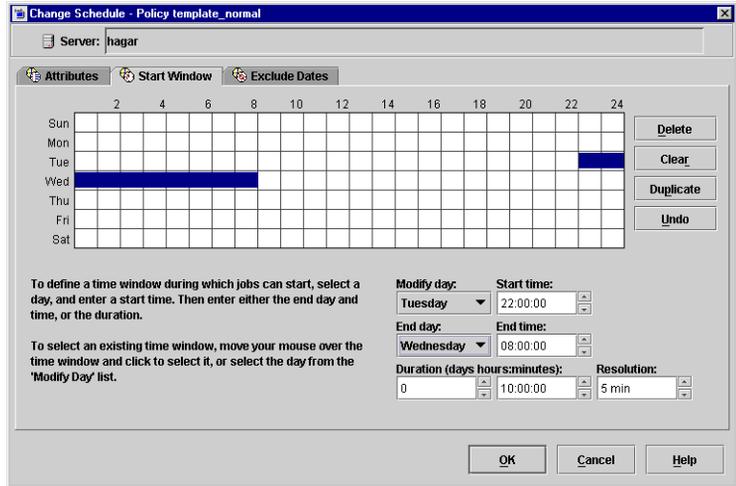
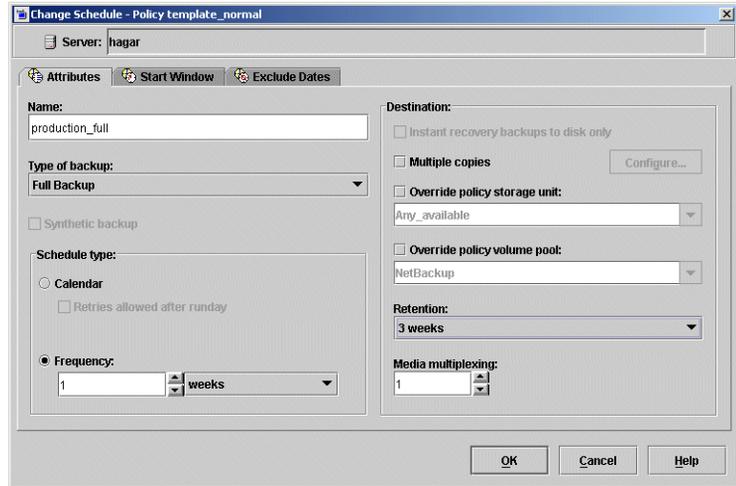


If the backup fails on Monday, NetBackup attempts it at the same time each day until it does successfully complete. NetBackup can attempt the backup on each subsequent day because the schedule allows backups to occur on any day, but only once during any seven day period. If the backup completes on Tuesday,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

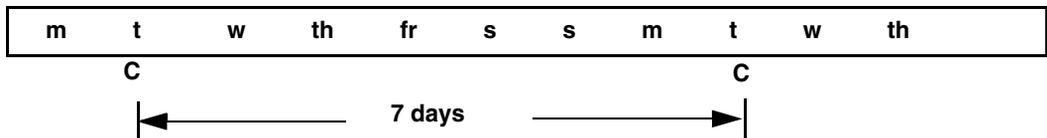


Another Method

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.



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.



If the backup fails on Tuesday, NetBackup must wait until the following Tuesday before trying again.



Example 2

The following shows a complete set of frequency schedules that have a backup window every day (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 Configure...

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: Start time: Sunday 22:00:00

End day: End time: Monday 08:00:00

Duration (days:hours:minutes): Resolution: 0 10:00:00 5 min

Add OK Close Help



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													
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: 10:00:00

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 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

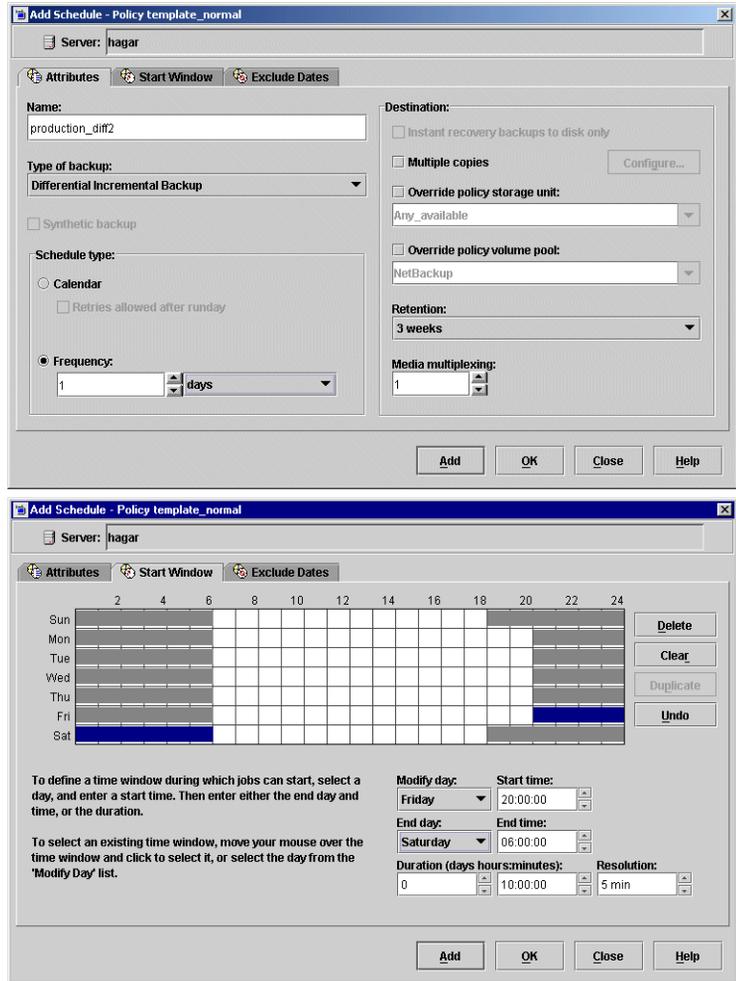
Duration (days hours:minutes): 0 Resolution: 5 min

Add OK Close Help



Example 3

The following is an example of using different backup windows, depending on the day.



Example 4

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.

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

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: Saturday Start time: 12:00:00

End day: Sunday End time: 12:00:00

Duration (days:hours:minutes): 1 Resolution: 5 min

Add OK Close Help

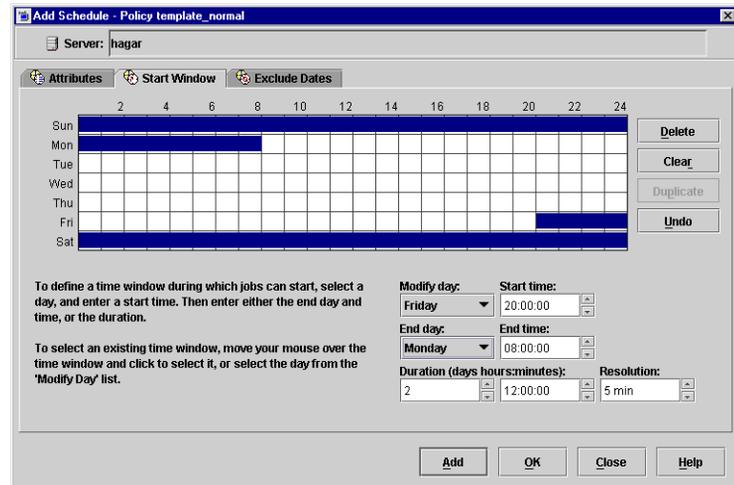
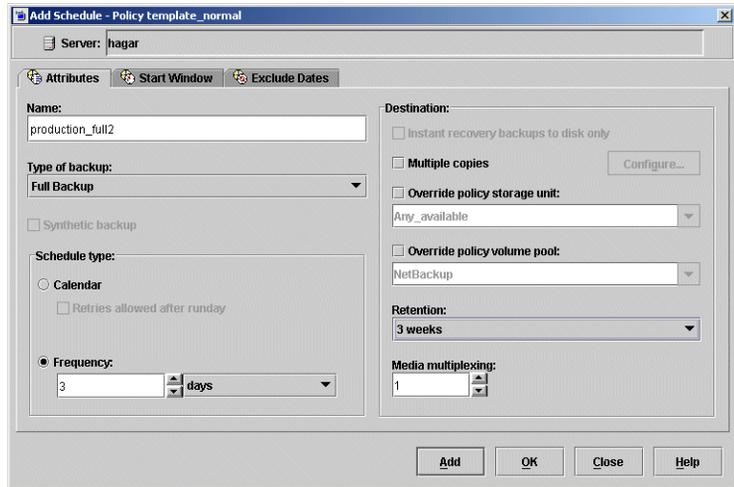


Example 5

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.



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 Configure...

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

	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: 5 min

Add OK Close Help



Considerations for User Schedules

To allow user backups and archives, you must create schedules for them. There is no requirement, however, to create a policy exclusively for user backups.

Restores can occur at any time and do not have schedules.

Note An archive is different from a backup: NetBackup first backs up the selected files, *then deletes them* 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 user backup and archive schedules, consider the following:

- ◆ Best times to perform backups. For user backups, this is the time most convenient to the users.

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.

- ◆ Storage unit. Using a different storage unit can eliminate conflicts with automatic backups.
- ◆ 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, 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.



The following table shows a set of clients in two user policies.

Policy	Client	Desired Storage	Best Backup Time	Retention
User1	mercury	8 mm tape stacker	08:00 to 16:00	Backups - 6 months
	mars			Archives - Infinite
	jupiter			
	neptune			
User2	pluto	8 mm 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 honors it 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	mercury, mars, jupiter, neptune
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.

Creating a Vault Policy

Creating a Vault policy differs from creating a regular policy in the following ways:

- ◆ You must specify *Vault* as the policy type.
- ◆ You do not specify clients for Vault policies.
- ◆ 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 window, expand **NetBackup Management > Policies**. Select Master Server at the top of the middle pane.
2. Click the New button .
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.
6. Bypass the Client tab, as clients are not specified for Vault jobs.
7. 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]
```

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*.

8. 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.
- ◆ When workstations miss their regular backups.
- ◆ Before installing new software (to preserve the old configuration).
- ◆ Preserving records before a special event such as when companies split or merge.
- ◆ 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. You can do this by creating it 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.

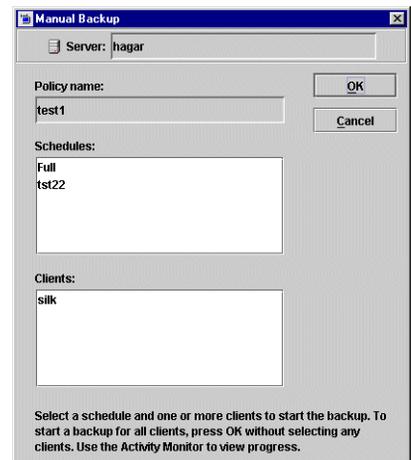
Note Not only does the policy need to be Active, but if **Go into effect** is set on the policy to a future date and time, the backup will not run.

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.





This chapter explains how to back up and manage the NetBackup catalog files. This chapter contains the following sections:

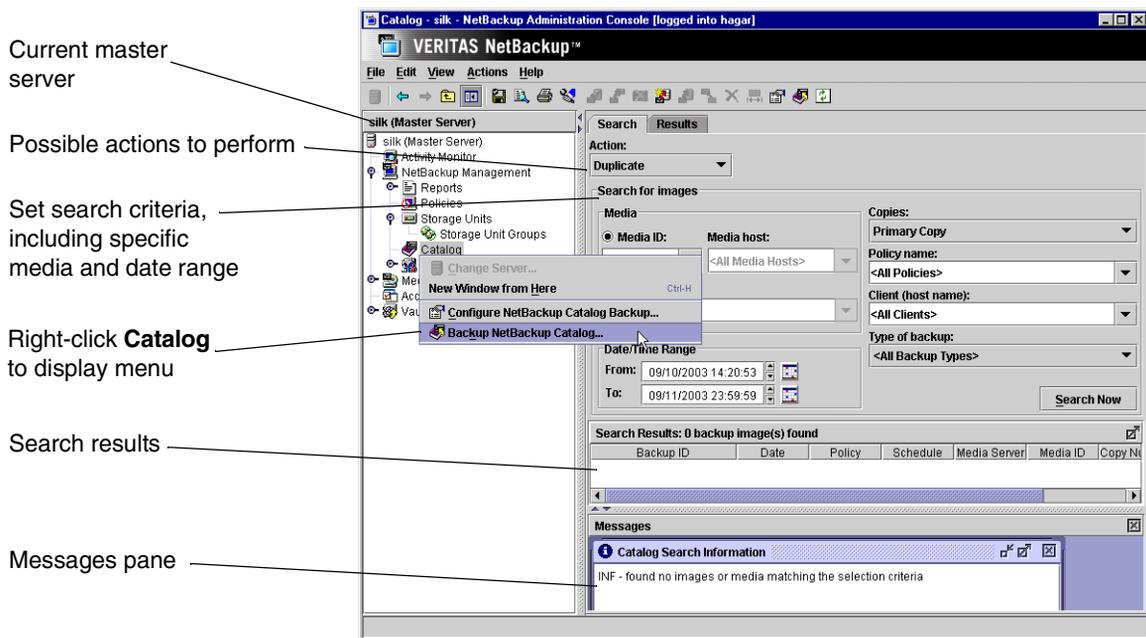
- ◆ “Introduction to the Catalog Application” on page 200
- ◆ “Catalog Backups” on page 200
- ◆ “Configuring Catalog Backups” on page 203
- ◆ “Backing Up the Catalogs Manually” on page 215
- ◆ “Protecting Large NetBackup Catalogs” on page 216
- ◆ “Managing the NetBackup Catalogs” on page 222
- ◆ “About the Binary Catalog Format” on page 222
- ◆ “Catalog Archiving” on page 230
- ◆ “Searching for Backup Images” on page 238
- ◆ “Verifying Backup Images” on page 240
- ◆ “Duplicating Backup Images” on page 241
- ◆ “Expiring Backup Images” on page 248
- ◆ “Importing NetBackup or Backup Exec Images” on page 249
- ◆ “Viewing Job Results” on page 256



Introduction to the Catalog Application

Use the **Catalog** application 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.

Catalog is also used to search for backup images 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.



Catalog Backups

NetBackup catalogs are internal databases that contain information about the NetBackup configuration and backups. Backup information includes records of the files and the media on which the files were stored. The catalogs also contain information about the media and storage devices that are under the control of Media Manager.



NetBackup requires the catalog information in order to recover any backups that have been performed. Therefore, it is extremely important to configure catalog backups 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.

NetBackup defaults to use a binary format for new catalogs on Windows, Solaris, HP_UX, Compaq Tru64 UNIX, AIX, Linux and SGI platforms. Releases prior to the first 4.5 feature pack created catalogs in ASCII format.

Existing catalogs can be upgraded to binary format using the catalog conversion utility, `cat_convert` as described in “Catalog Conversion Utility” on page 222.

Where are the Catalog Files?

The catalogs reside on disk on NetBackup servers. NetBackup chooses default locations for them during installation. If you change the default locations, you must change your catalog backup configuration accordingly. (See “Catalog Files Tab” on page 210.)

What Method Do I Use to Back Them Up?

Because the catalogs are essential to restoring files in case of a disk crash, the process for backing them up is separate and different than for standard backups. The two available methods are:

- ◆ Automatic backup according to your configuration as you defined it in “Configuring Catalog Backups” on page 203.
- ◆ Manual backup as explained in “Backing Up the Catalogs Manually” on page 215.

What NetBackup Servers Can I Use?

Applies only to NetBackup Enterprise Server:

The catalogs can be backed up to either the master server or one of its remote media servers. During the configuration process, explained later in this chapter, you specify both the media server and the media to use for the backups.

What Types of Media Can I Use?

You can use either removable media (such as a tape) that is configured under Media Manager, or a directory on a hard disk. (See “Media Type” on page 206.)



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:

- ◆ `dbbackup_notify` script.
- ◆ E-mail, if you configure this capability with the E-mail Address for NetBackup Administrator Global attribute. (See “Administrator’s E-mail Address” on page 366.)

How Do I Restore The Catalog Backups?

If it is necessary to perform a disaster recovery, restore the catalogs by using the NetBackup `bprecover` command. See the *NetBackup Troubleshooting Guide for UNIX and Windows* for recovery procedures.

Important Precautions to Observe

- ◆ Use only the methods described in this chapter to back up the catalogs. The special backup operations described here are the only ones that can track all relevant NetBackup activities and ensure consistency between the catalog files.
 - ◆ Do not use scheduling or backup methods provided by any other vendor.
 - ◆ Do not rely on user backups or regular-scheduled backups. If you use these methods and the disk fails, the catalogs as well as the backups are lost and you may not be able to recover any data.
- ◆ *Applies only to NetBackup Enterprise Server:* If you are using media servers, manually alter the NetBackup catalog configuration to include the catalogs on the media servers.
- ◆ 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 NetBackup catalog backup and the time that the disk crash occurred.
- ◆ Never manually compress the catalogs. If you compress them manually, NetBackup may not be able to read them with its standard mechanism, the `bprecover` command.
- ◆ Keep a hard-copy record of the media IDs where you store the NetBackup catalog backups, or configure the E-mail global attribute. The E-mail global attribute 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 has the catalogs.

- ◆ 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 catalog is more sensitive to the location of the catalog. Storing your 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.

Configuring Catalog Backups

The easiest way to configure NetBackup catalog backups is to use the Catalog Backup Wizard. This wizard guides you through the configuration process, simplifying it by automatically choosing settings that are good for most configurations. If you are modifying an existing configuration or want access to all available configuration settings, use the manual method. The following sections explain both the wizard and the manual method.

▼ To configure the catalog backup using the Catalog Backup Wizard

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. You can change a policy after it is created. See “Backing Up the Catalogs Manually” on page 215.



Note If you are unfamiliar with NetBackup catalog backups, read “Introduction to the Catalog Application” on page 200 before proceeding. In particular, read the precautions under “Important Precautions to Observe” on page 202.

▼ **To configure the catalog backup manually**

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

If your site has more than one master server, use **File > Change Server** to select a different server. (See “Administering a Remote Master Server” on page 420.)

2. Select **Actions > Configure NetBackup Catalog Backup**. The Catalog Backup Configuration dialog appears containing three tabs: Attributes, Schedule, Files.
3. Specify the properties on each tab of the dialog using the Catalog Attributes, Catalog Schedule, and Catalog Files tabs:
 - ◆ “Catalog Attributes Tab” on page 205
 - ◆ “Catalog Schedule Tab” on page 209
 - ◆ “Catalog Files Tab” on page 210
4. Click **OK**.

Catalog Attributes Tab

The Catalog **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 NetBackup server to which catalogs 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.

The screenshot shows the 'Catalog Backup Configuration - NetBackup' dialog box with the 'Attributes' tab selected. It contains the following fields and values:

- Media server:** collie
- Last media used:** F:\NB_Catalog
- Media 1:**
 - Media type: Disk (hard drive)
 - Volume: Media ID Density
 - Pathname (disk media type): F:\NB_Catalog
 - Last written: 11/14/2001 17:45:00
 - Allocated: 11/14/2001 17:42:13
- Media 2:**
 - Media type: None
 - Volume: Media ID Density
 - Pathname (disk media type):
 - Last written: never
 - Allocated: never

Buttons at the bottom include OK, Cancel, Apply, and Help.

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. “Catalog Files Tab” on page 210. Also, ensure that the media server was named in the `bp.conf` file on the master server at the time that you started `bpd` 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
- ◆ **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 (though this is not as safe as method 4 below). 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. If you choose to back up the catalog to disk, the destination of the catalog backup must be on a different drive.

`/usr/opensv/netbackup/db`

`/usr/opensv/volmgr/db`

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 change to a different volume.
- ◆ 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 configure the E-mail global attribute. When this is done, NetBackup sends 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 disk containing the catalogs.

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.

Caution 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/volmgr/database  
/usr/opensv/var
```

- ◆ Ensure that the disk has adequate space for the catalogs. If the disk fills up, the catalog backups will fail.

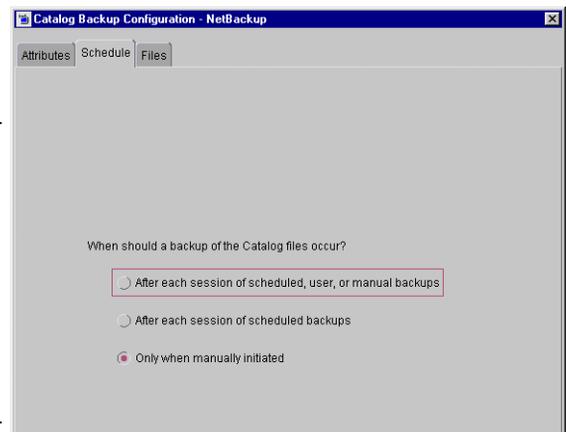
- ◆ Ensure that the path is a directory rather than a file. If the path is a file, an error occurs when the backup is done (*not* when you specify the path).
- ◆ The following rule applies to the path you specify:

In addition to the platform-specific file path separators (/ and \) and colon (:), within a drive specification on Windows, use only alphabetic (ASCII A - X, a - z), numeric (0-9), plus (+), minus (-), underscore (_), or period (.) characters. Do not use a minus as the first character.

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.

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.

Recommendations

- ◆ If you are sending your catalog backups to a robot or tape stacker, a second standalone tape drive, or to disk, choose either of the two automatic backups.
- ◆ If you must use a single standalone tape drive to back up both catalog *and* business data, choose either:
 - ◆ If you will be running only one backup session per day or night, choose:
After each session of scheduled backups
 - ◆ If you will be running multiple backup sessions in a single day or night, choose:
Only when manually initiated

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:

1. Insert the tape configured for catalog backups.
2. Manually start the backup. (See “Backing Up the Catalogs Manually” on page 215.)
3. When the backup is complete, remove the tape and store it in a safe place.

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.

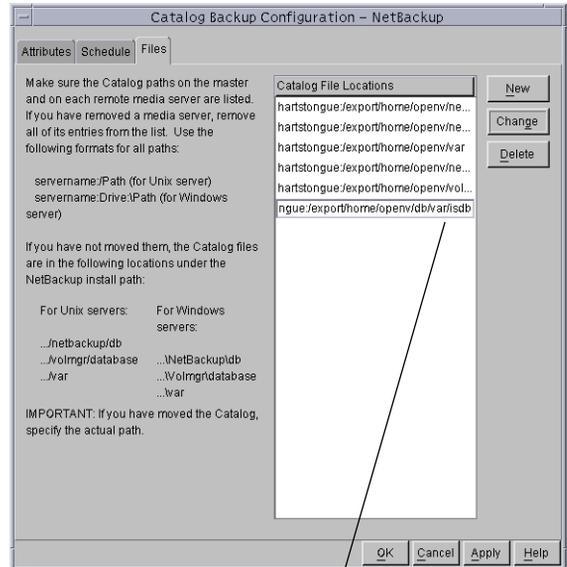
Catalog Files Tab

The Catalog **Files** tab contains the absolute pathnames to the catalog files to be backed up.

For more information on pathnames, see “Catalog Pathnames” on page 212 and “Pathnames for the NetBackup Database” on page 213.

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.

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.



Enter a directory or an individual table

Note The table names and database names in the database pathname are case-sensitive. The database catalog backups will fail if typed without regard to case.

▼ To add a pathname

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Select **Actions > Configure NetBackup Catalog Backup**.
3. Select the **Files** tab.
4. Click **New**.
5. Type the pathname in the **Catalog File Locations** list. (See "Catalog Pathnames" on page 212.)
6. Click **OK** to complete the addition.



Caution 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.

Caution Do not specify a 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.

▼ **To change a pathname**

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Select **Actions > Configure NetBackup Catalog Backup**.
3. Select the Files tab.
4. Select the pathname you wish to change and click **Change**.
5. Change the pathname and click **OK**.

▼ **To delete a pathname**

1. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
2. Select **Actions > Configure NetBackup Catalog Backup**.
3. Select the Files tab.
4. Select the pathname you wish to delete and click **Delete**.

Catalog Pathnames

Applies only to NetBackup Enterprise Server:

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.

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.

```
/usr/opensv/netbackup/db
```

The files in this directory have NetBackup scheduling information, error logs, and all information about files backed up from client workstations.

```
/usr/opensv/volmgr/database
```

The files in this directory have the information about the media and devices being used in the configuration.

```
/usr/opensv/var
```

The files in this directory contain license key and authentication information.

If you are backing up the catalogs to a media server, prefix each pathname with the name of the master server:

```
master_name:catalog_backup_path
```

For example, if the master server is named *venus*, the paths are:

```
venus:/usr/opensv/netbackup/db  
venus:/usr/opensv/volmgr/database  
venus:/usr/opensv/var
```

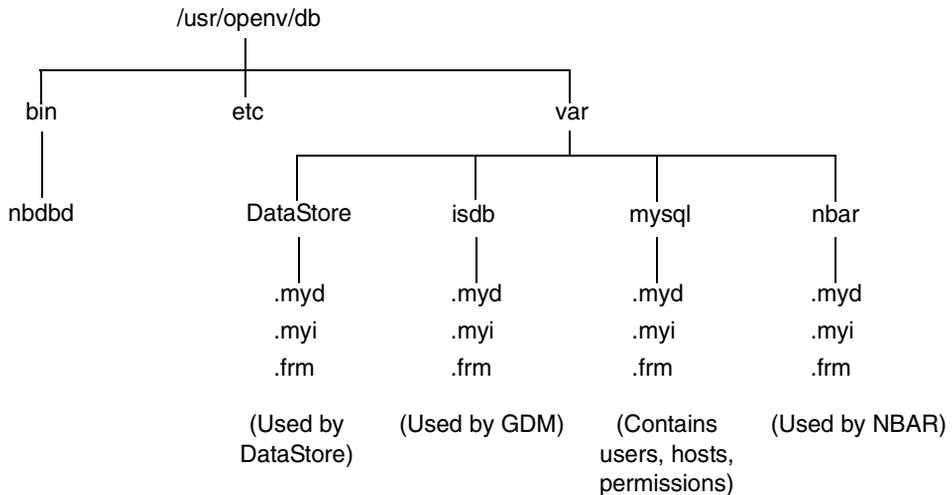
Pathnames for the NetBackup Database

The pathnames to the NetBackup database on the master server are *not* automatically added during installation and require that you add them to the list. For example:

```
/usr/opensv/db
```

Note The table names and database names in the database pathname are case-sensitive. The database catalog backups will fail if typed without regard to case.





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:

```
server_name:catalog_backup_path
```

The paths that you must add depend on whether the platform has a volume database or devices configured and on the version of NetBackup installed on the system.

- ◆ For UNIX NetBackup media servers that have a volume database or devices configured, 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 NetBackup media servers that do not have a volume database or devices configured, add the following path:

```
media_server_name:/usr/opensv/netbackup/db/media
```

The files in this directory have information about files backed up or archived from client workstations.

- ◆ For UNIX NetBackup media servers at version 3.4 or later, also include the following path:

```
media_server_name:/usr/opensv/var
```

For example, to add the paths for a UNIX NetBackup media server (3.4 or later) named elk that has a volume database or devices configured, make the following entries:

```
elk:/usr/opensv/netbackup/db/media
elk:/usr/opensv/volmgr/database
elk:/usr/opensv/var
```

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
```

For example, the paths for a Windows NetBackup server (3.4 or later) named mars are as follows (*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.

Backing Up the Catalogs Manually

A manual backup starts a backup of the catalogs immediately. Starting a manual backup is useful in the following situations:

- ◆ To perform an emergency backup. For instance, if you anticipate a problem or are moving the system and do not want to wait for the next scheduled catalog backup.
- ◆ 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.



▼ **To perform the catalog backup manually**

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

If your site has more than one master server, use **File > Change Server** to select a different server. (See “Administering a Remote Master Server” on page 420.)

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 started by the scheduler must wait until the catalog backup is complete.

Protecting Large NetBackup Catalogs

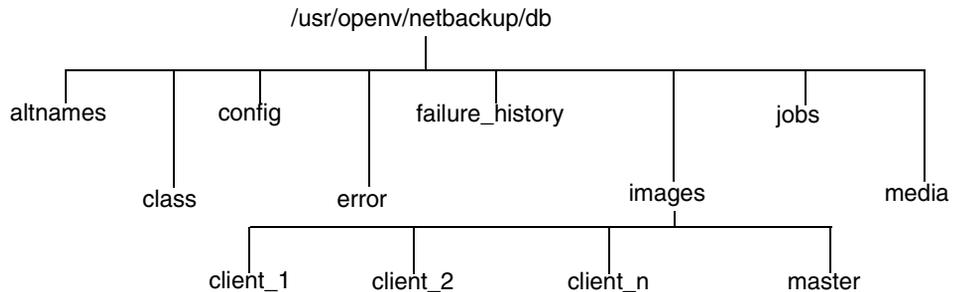
It is very important to ensure that the NetBackup catalogs on the master and media servers are backed up regularly. NetBackup provides a built-in mechanism for achieving this. However, this mechanism imposes a limit on the size of the data that can be backed up; namely, the data must all fit on a single piece of media.

You can set up a method for multiple-tape catalog backups to back up and recover the NetBackup catalog files if they become too large to fit onto a single tape.

Layout of the NetBackup Catalogs

Before implementing a solution for backing up large NetBackup catalogs across multiple tapes, it is important to understand the structure of the catalogs.

The NetBackup and Media Manager catalogs are held within subdirectories on the master server and media servers. The NetBackup catalogs reside in the directory `/usr/opensv/netbackup/db` and the Media Manager catalogs reside in `/usr/opensv/volmgr/database`. Typically, it is the NetBackup catalogs on the master server that will grow large and can fail to fit on a single tape. The diagram below shows the directory layout of the first few directory depths of the NetBackup catalogs on the master server.



The directories under `db` contain further subdirectories or files, which together make up the NetBackup catalogs. The `images` directory contains a directory sub-tree, with one subdirectory for each NetBackup client that has been backed up (including the master server and any media servers). Beneath these subdirectories are further directories and files, which hold the information about all the backup images held by NetBackup.

While most of the subdirectories in the NetBackup catalogs are relatively small, the `images` directory can grow to several tens, or even hundreds of gigabytes. (See “Determining Catalog Space Requirements” on page 223 for more information on estimating the size of the NetBackup catalogs.)

Due to its potentially large size, it is the `images` subdirectory that can become too large to fit onto a single tape and it is therefore this subdirectory that is addressed in the following sections.

Catalog Backup and Restore Concepts

The following sections present the concepts underlying multiple-tape catalog backups and restores.

Multiple-Tape Catalog Backups

The basic concept behind the protection of large NetBackup catalogs is to split the catalog-backup process into two steps:



1. Back up the majority of the data from the images subdirectory on the master server.
2. Back up a small sub-set of the images subdirectory, together with the remainder of the NetBackup and Media Manager catalog files and directories from the master server and media servers.

Since the first backup contains the majority of the data, it must be able to span tapes. This is achieved by using a normal NetBackup job to back up the data. As a result of this normal backup, an entry is placed in the images subdirectory tree for the master server. This catalog entry allows the user to browse the catalog for files during a restore operation.

The second backup must back up the portion of the images subdirectory that contains the catalog entries for the master server, together with the other parts of the NetBackup and Media Manager catalogs. Since this is a relatively small amount of data, it fits onto a single tape. It must also be possible to recover this backup without the NetBackup catalogs being available. This is achieved by using the normal NetBackup catalog-backup mechanism to perform the backups.

Multiple-Tape Catalog Restores

A restore of the NetBackup catalogs is also achieved in two steps. The first step is to use the most recent NetBackup catalog backup to recover the portion of the image catalog containing information about the backups taken from the master server, together with the other parts of the NetBackup and Media Manager catalogs on the master server and, if configured, the media servers.

Once this information has been recovered, NetBackup can be started and one of the user interfaces can be used to browse the files backed up from the master. These include the files and directories that constitute the NetBackup images catalog, which were backed up using the first step of the catalog backup described above. Using the normal restore process, these files and directories are restored, completing the operation. You must ensure the option **Overwrite Existing Files** is not selected, since this replaces the files previously recovered in stage 1.

Setting up Multiple-Tape NetBackup Catalog Backups

In order to configure NetBackup to perform multiple-tape backups of its catalogs, define a normal NetBackup policy and make changes to the NetBackup catalog-backup configuration. In addition, you must create a shell script or executable file to initiate the multiple-tape catalog backups. These steps are detailed below.

▼ To define a NetBackup policy for catalog backups

1. Use the NetBackup Administration Console to create a new policy with the following policy attributes:
 - ◆ Set the **Policy Type** to Standard if the master server is a UNIX machine or MS-Windows-NT if the master server is a Windows machine.
 - ◆ Do not choose **Cross Mount Points** if the master server is a UNIX or Windows 2000 machine.
 - ◆ Do not choose **Follow NFS** for UNIX or **Backup network drives** for Windows NT.
 - ◆ Pick a suitable storage unit and volume pool.
 - ◆ Set **Limit Jobs per Policy** to 1.
 - ◆ Do not choose **Compression**.
 - ◆ Set **Job Priority** to 0.
2. Add the master server to the client list.
3. Enter the following path in the file list:

```
/usr/opensv/netbackup/db/images
```

Note On UNIX, 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.

4. Set up schedules to meet your requirements. VERITAS recommends that the policy contains only a full backup schedule, since this will minimize tape mounting and positioning during restores.

Do not set any backup windows for the schedules that you define. This ensures that the backup policy is never initiated automatically by the NetBackup scheduler. Instead, you must initiate the backup job manually.
5. Save your changes.

▼ To configure the NetBackup catalog backups

1. In the NetBackup Administration Console, ensure that the **Media Server** setting specifies the required backup server.
2. Specify the following for **Absolute Pathname**:

```
masterserver:/usr/opensv/netbackup/db/[A-Za-hj-z]*
```



```
masterserver: /usr/opensv/netbackup/db/images/masterserver
masterserver: /usr/opensv/var
mediaserver1: /usr/opensv/netbackup/db/media
mediaserver1: /usr/opensv/volmgr/database
```

(repeat for additional media servers)

3. Change the schedule to **Only When Manually Initiated**. This stops the NetBackup catalog backups from running automatically and allows you to control when they run manually.
4. Select appropriate media types, densities, and IDs for the two catalog-backup media.
5. Save your changes.

Create a Shell Script to Initiate the Backups

It is also important that the second-stage backup of the NetBackup catalogs occurs directly after the first-stage backup. A good way to ensure this is to write a script that initiates both backups, one after the other.

Example Catalog-backup Script

```
#!/bin/sh
#
# catalog_backup script
#
# Performs a two-stage backup of the NetBackup catalogs
#
POLICY=nbu_cat_backup # Change to the name of the correct policy
SCHED=full_backup # Change to the name of the correct schedule
LOGDIR=/usr/opensv/netbackup/logs/catalog_backup
if [ -d $LOGDIR ]; then
    exec >> $LOGDIR/log.`date +%m%d%y` 2>&1
else
    exec > /dev/null 2>&1
fi
echo "Running first stage catalog backup"
/usr/opensv/netbackup/bin/bpbackup -w -i -c $POLICY -s $SCHED
EXIT_STAT=$?
if [ $EXIT_STAT -ne 0 ]; then
    echo "First stage catalog backup failed ($EXIT_STAT)"
    exit 1;
fi
echo "Running second stage catalog backup"
/usr/opensv/netbackup/bin/admincmd/bpbackupdb
EXIT_STAT=$?
```

```
if [ $EXIT_STAT -ne 0 ]; then
    echo "Second stage catalog backup failed ($EXIT_STAT)"
    exit 1;
fi
exit 0;
```

How To Initiate a Multiple-Tape Catalog Backup

Similarly to how the automatic-catalog backup works, it is important to ensure that no other NetBackup operations that modify the NetBackup catalogs are in progress while the two catalog backups are performed. Such operations include:

- ◆ Backups and archives
- ◆ Catalog compression
- ◆ TIR record expiration or retrieval (during a restore operation)
- ◆ Catalog image record expiration
- ◆ Image imports
- ◆ Image duplication

Performing the catalog backups when any of these operations are in progress can cause an inconsistent catalog backup.

Since both image import and image duplication operations must be initiated manually by the NetBackup administrator, it is relatively easy to ensure that these are not in progress during the catalog backup. However, it is more difficult to ensure that no backups or restores are running, since both the NetBackup scheduler and other users can initiate these.

More difficult still, are operations that are started automatically by NetBackup, such as catalog compression, TIR record expiration or retrieval, and image record expiration. Due to the way the NetBackup scheduler interlocks processes, do not start the two-step backup script with the `/usr/opensv/netbackup/bin/session_notify` script. We suggest using another scheduler (such as `cron` on UNIX) to start the two-step backup script or run it manually when the above operations are not occurring.



Managing the NetBackup Catalogs

The following sections explain various aspects of managing the NetBackup catalogs:

- ◆ “About the Binary Catalog Format” on page 222
- ◆ “Determining Catalog Space Requirements” on page 223
- ◆ “Compressing the Image Catalog” on page 226
- ◆ “Uncompressing the Image Catalog” on page 227
- ◆ “Moving the NetBackup Image Catalog” on page 228

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 in binary format. The binary representations of numbers, dates, and so on, takes up less disk space than the text representations.
- ◆ The catalog in binary format is much faster to browse and search, especially for large file sizes.
- ◆ The catalog in binary format supports alternate backup methods without requiring post-processing, improving catalog performance for alternate backup methods.

Catalog Conversion Utility

In order to allow users to convert from ASCII to binary, NetBackup offers a catalog format conversion utility called `cat_convert`. The utility converts NetBackup catalog `.f` files between version 3.4, 4.0V or 4.5 ASCII format and 4.5 binary format. `cat_convert` automatically detects the source catalog file format and converts it to the other format. See the *NetBackup Commands Guide for UNIX* for information on `cat_convert`.

Upon installation, NetBackup does *not* convert existing ASCII catalogs to the binary catalog format. However, any new catalogs created will be binary. You may elect to upgrade any existing NetBackup catalogs to binary format using the catalog conversion utility, `cat_convert`, described below.

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):
(2³¹) – 1 IDs = 2,147,483,647 IDs = 7FFFFFFF IDs

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):

$$\text{Files per Backup} \times \text{Backups per Retention Period} = \text{Max Files}$$



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.

Determining the space required for binary catalogs:

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:

$$(8400 \times 150) + (12,000 \times 150) = 3,060,000 \text{ bytes (or about 2988 kilobytes)}$$

Determining the space required for ASCII catalogs:

If you are unsure of the average length of a file's full pathname, use 150. (Averages from 100 to 150 are common.) Using the results from the examples in step 3 yields:

$$(8400 \times 150) + (12,000 \times 150) = 2,988 \text{ kilobytes (1024 bytes in a kilobyte)}$$

Note If you have ASCII catalogs and use catalog indexing, multiply the number in step 4 by 1.5%. For information on catalog indexing, see "Reduce Restore Times by Indexing the Image Catalog" on page 236.

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

File system limitations:

- ◆ For a FAT 32 file system, the maximum file size is 4GB.
- ◆ 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.

One file size and security-related limitation:

For UNIX systems, set the file size limit for the root user account to *unlimited* in order to support large file support.

“Example Reference Table for Catalog Requirements” on page 226 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:

$$\begin{aligned} & (\text{Files per Incremental} \times \text{Backups per Retention Period}) \\ & \quad + \\ & (\text{Files per Monthly Full Backups} \times \text{Backups per Retention Period}) \end{aligned}$$

2. Substitute values from “Example Reference Table for Catalog Requirements” on page 226:

$$1000 \text{ files} \times 30 + 64,000 \text{ files} \times 12 = 798,000 \text{ files}$$

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:

$$460.59 \text{ megabytes} (1,048,576 \text{ bytes in a megabyte})$$

Adding 15 megabytes for error logs results in a final uncompressed catalog space requirement of:



475.59 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
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

Compressing 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 the Global NetBackup attribute, **Compress Catalog Interval**. This attribute 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.

For more information, see “Global Attributes Properties” on page 362.



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 the scheduler 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** page, clear the **Compress Catalog Interval** check box.
5. Set the **Compress Catalog Interval** Global NetBackup attribute to 0.
6. Change your working directory to `/usr/opensv/netbackup/bin` and run the command:

```
admincmd/bpimage -decompress -client name
```
7. Restart the request daemon, `bprd`, by running:

```
/usr/opensv/netbackup/bin/initbprd
```
8. Perform the file restorations from the client.
9. Set the **Compress Catalog After** Global NetBackup attribute to its previous value.
The records that were uncompressed for this client will be compressed after the backup scheduler, `bpsched`, runs the next backup schedule.

Moving the NetBackup Image Catalog

If the image catalog becomes too large for the file system in which it is currently located, you can move it to one containing more space.

▼ **To move the NetBackup 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. (See “Configuring Catalog Backups” on page 203.)

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`.



Catalog Archiving

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.

Examining the Catalog Image

The NetBackup binary catalog image consists of two types of files: image files and image .f files.

Image Files

The image file is generally less than 1 kilobyte in size because it contains only backup set summary information. For example, the backup ID, the backup type, the expiration date, and fragment information.

NetBackup Catalog



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*.

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
```

Multiple File Layout

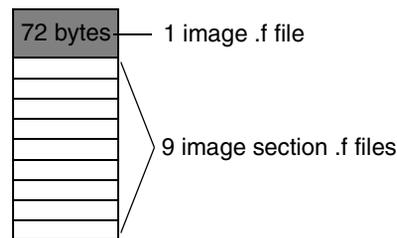
When the detailed backup file information of one catalog backup is greater than or equal to 4 megabytes, the information is stored in multiple .f files: one main image .f file plus nine section .f files.

Separating the section .f files from the image .f file and storing them 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
```

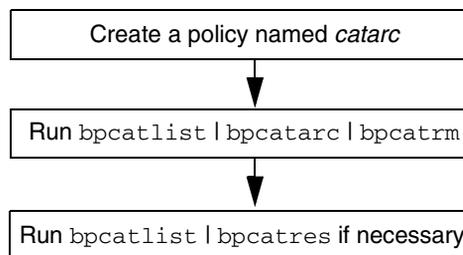
NetBackup Catalog in
Multiple File Layout



Catalog Archiving Overview

The following section describes the steps to archive a catalog. Catalog archiving operations must be performed when NetBackup is in a quiet state.

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 232.)
2. Run `bpcatlist | bpcatarc | bpcatrm` to archive the list of .f files indicated by `bpcatlist`, then remove the files from the database after archiving them. (See “Catalog Archiving Commands” on page 233.)
3. If necessary, run `bpcatlist | bpcatres` to restore the .f files to the catalog.



Creating a Catalog Archiving Policy

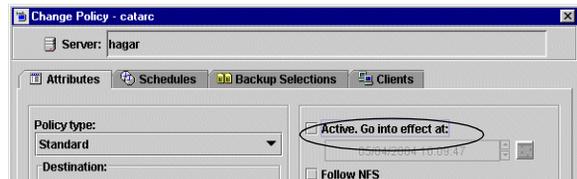
The catalog archiving feature requires the presence of a policy named *catarc* in order to have catalog archiving commands run properly. The policy can be reused for catalog archiving.

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.

Inactive Policy

The catalog archive policy must be set up as inactive. On the Attributes tab, clear the **Active** field.



Type of Backup

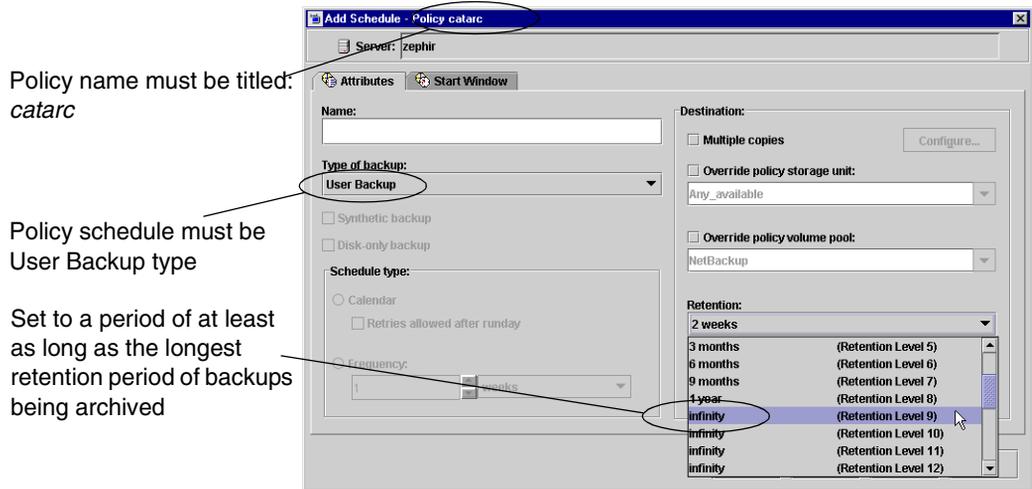
The type of backup indicated for the catalog archive policy must be *User Backup*. This is set in the Change Schedule dialog on the Attributes tab.

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.



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, 2000, the following would be entered:

```
# bpcatlist -client all -before Jan 1 2000 | bpcatarc | bpcatrm
```

`bpcatlist` is also used to provide status information. For each catalog, it lists the following information:

- ◆ Backup ID (Backupid)
- ◆ Backup date (Backup Date)
- ◆ Catalog archive ID (Catarcid). After an `.f` file is successfully backed up, a catalog archive ID is entered into the `catarcid` 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      ...Catarcid  S C Files file
alpha_0972380832 Oct 24 10:47:12 2000 ... 973187218 1 0 alpha_0972380832_UBAK.f
alpha_0972336776 Oct 23 22:32:56 2000 ... 973187218 1 0 alpha_0972336776_FULLL.f
alpha_0972327197 Oct 23 19:53:17 2000 ... 973187218 1 0 alpha_0972327197_UBAK.f
```

For detailed information on `bpcatlist`, see `bpcatlist` in the *NetBackup Commands for UNIX Guide*.

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 232.)

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 a quiet state.
- ◆ 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 380.

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 *bpclist* 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.

Reduce Restore Times by Indexing the Image Catalog

If you have large numbers of backups, reduce the total time required to restore files by creating indexes of the backed up files that are recorded in the NetBackup image catalog. NetBackup can then use the indexes to go directly to the catalog entry for a file rather than starting the search at the beginning of the catalog entries.

Note This section applies to ASCII catalogs only. Binary catalogs do not need catalog indexing.

Use the following command to generate indexes for one or all clients, and 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.
- ◆ *client_name* is the name of the client of the backups you want to index. The default is all clients.

Run this command once, to activate indexing for a client. Once activated, indexing is done automatically each night when NetBackup does its cleanup for the previous day's activities.

Catalog Index Examples

- ◆ To index client mars to index level 5 (five levels of directories), run:
- ◆ To index selected clients, run a command for each of them (you cannot use wildcards). The following indexes 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
```

Note Changing the index level affects only future index creation and does not immediately create index files.

Catalog Index Space Requirements

The index files do not require much space. Regardless of how many clients you have, indexing all clients to level 9 requires about 1.5 percent more space in the NetBackup catalog than if you do not use indexing for any clients. NetBackup does not produce index files for backups that contain less than 200 files.

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
```

Note If you are collecting true-image restore information, the INDEX files take much more space for incrementals.



Disabling Catalog Indexing

Note When using a binary catalog, disable catalog indexing.

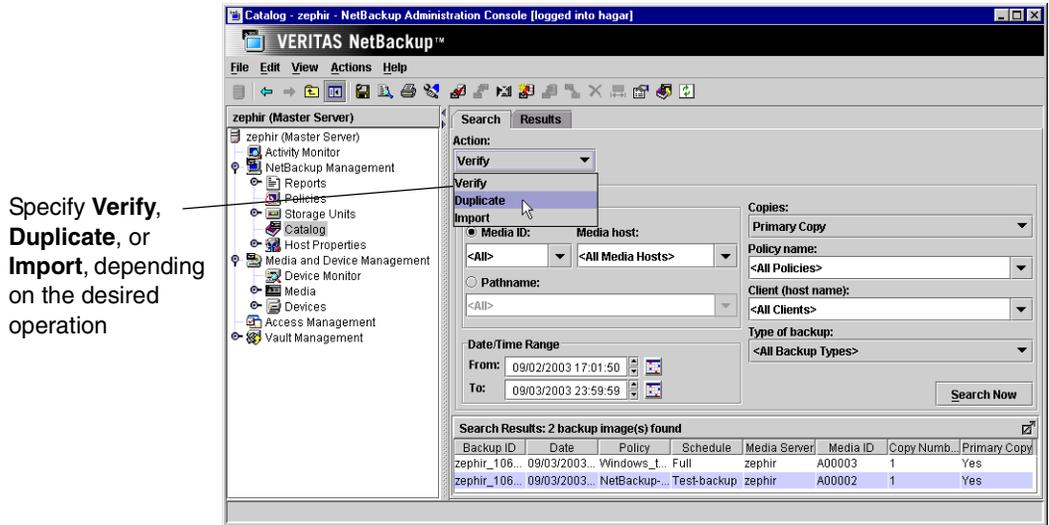
- ◆ To stop NetBackup from generating new `INDEX` files for a client, delete the `INDEXLEVEL` file. NetBackup continues to use existing `INDEX` files.
- ◆ To temporarily stop using the `INDEX` files during searches but retain existing index files, change the `INDEX` directory to `INDEX.ignore`. When 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.

Searching for Backup Images

Use **Catalog** 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 .
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 attribute setting, Interval for status reports .
Copies	The source you want to search. From the scroll-down list, select either Primary or the copy number.



Search Criteria for Backup Images (continued)

Search Criteria	Description
-----------------	-------------

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 253 it is not imported. If a backup ends on a media ID that has not been processed by first step of “To initiate an import – Phase I” on page 253, 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 appears.
To display information on each file that NetBackup verifies, select **Log all files found in image(s) verified**.
4. Click the **Results** tab, then select the verification job just created to view the job results. (See “Viewing Job Results” on page 256.)



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 362.)

NetBackup can create up to four of the copies simultaneously.

An alternative to taking time to duplicate backups is to use Inline Tape Copy. Inline Tape Copy allows you to 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 when using Inline Tape Copy).
- ◆ when the backup has expired.
- ◆ by using the NetBackup scheduler to schedule duplications automatically (unless you use a Vault policy to schedule duplication)
- ◆ of the 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.

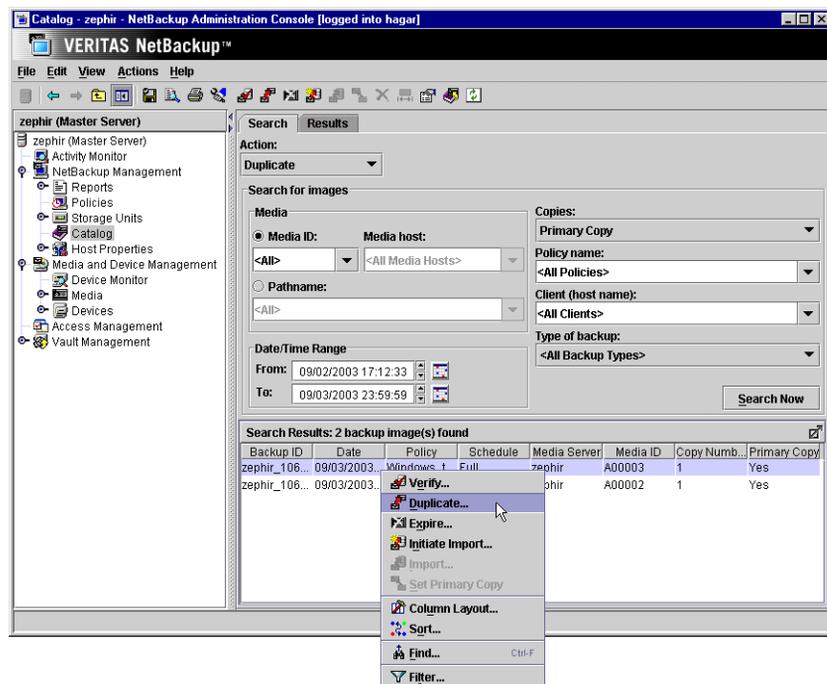
Notes on Multiplexed Duplication

- ◆ When duplicating multiplexed SQL-BackTrack backups with multiplex mode enabled, it is necessary to duplicate all the backups in the multiplexed group. This ensures that the fragment order and size is maintained in the duplicate. Otherwise, it is possible that restores from the duplicated backups will not work. A multiplexed group is a set of backups that were multiplexed together during a single multiplexing session.
- ◆ 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. (This concerns only for SQL-BackTrack backups.)

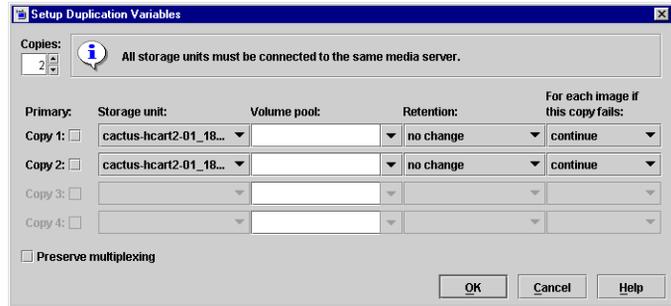
▼ To duplicate backup images

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 you wish to duplicate and select **Duplicate** from the shortcut menu. The **Setup Duplication Variables** dialog appears.



- 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.



- 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.)

- 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 Inline Tape Copy does not support the following storage types: NDMP, third-party copies, or optical devices.
Also, Inline Tape Copy does not support storage units that use a QIC (quarter-inch cartridge) drive type.

- Specify the volume pool where each copy will be stored.

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.

- 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 `bpxpdate` command to change the expiration date of the duplicate.
 - ◆ If a retention period is indicated, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2002 and its retention period is one week, the new copy's expiration date is November 21, 2002.
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 or disk staging storage unit. However, if the source is a multiplexed tape and the destination is a disk storage unit or disk staging 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 256.)

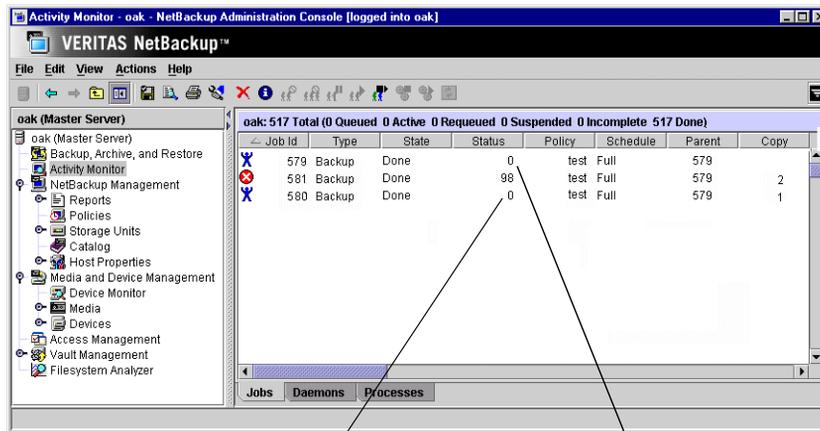
Inline Tape Copy Jobs

When using Inline Tape Copy to create simultaneous copies, 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

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 238.) 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 the guide, *NetBackup Commands for UNIX*.

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 271.)

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. See the *NetBackup User's Guide* for instructions.

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 reimporting.

▼ 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 239. 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.
5. Select **Yes** to proceed with expiring the image or **No**.



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 (7.0 and later). 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.

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

- ◆ You cannot import data from a disk image.
- ◆ 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.
- ◆ 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

In order to import Backup Exec media, the following conditions must be satisfied:



Master Server

The master server on which the Backup Exec images will be created must be at NetBackup 5.0 or later.

Media Server

The media server on which the Backup Exec media will be mounted must be at NetBackup 5.0 or later. Importing Backup Exec images from media that is not compatible with the devices supported by NetBackup 5.0 or later media servers is not supported.

It is not possible to restore Backup Exec media with a NetBackup 5.0 master server (or later) using the BE Tape Reader on a NetBackup 4.5 media server. In order to use a NetBackup 4.5 media server as a BE Tape Reader, it must be used with a NetBackup 4.5 master.

Clients

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 312.)

Consideration for 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.
- ◆ To restore media written by Backup Exec using an RSM-controlled standalone drive or robot, the media must be put in a compatible non-RSM standalone drive or a library.
- ◆ 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.

More on `vmphyinv` and `bephyinv`

Whether you run `vmphyinv` or `bephyinv` to inventory a robot is determined by what kind of robot and media server you have.

`vmphyinv`

Use `vmphyinv` on Backup Exec media to physically inventory the media contents of a robotic library or standalone drive and update the volume database. `vmphyinv` runs on all supported media server platforms. `vmphyinv` must be used when inventorying Backup Exec media on UNIX platforms, and before performing Phase I imports of fresh Backup Exec media.

`bephyinv`

Use `bephyinv` to inventory Backup Exec media in API robots on Windows media servers to make fresh Backup Exec media known to Backup Exec and NetBackup. `bephyinv` works with all robot types but runs only on Windows media server platforms.

If Backup Exec media needs to be used in API robots on UNIX media server platforms, the ADAMM GUID field of the Media Manager volume record needs to be manually updated using the unsupported options `vmchange -g` or `vmadd -a`.

To find out the media GUID of the Backup Exec media, mount the media in the drive and wait for the drive to become ready. Once the drive becomes ready, dump the `ltid` shared memory tables (`ltid -tables` option) and get the media GUID from the `appl_guid` field of the drive status table.



bephyinv may not be supported in future releases. Instead, vmphyinv will be used. vmphyinv cannot, however, be used with the `-auto_correct` option provided by bephyinv. Also, vmphyinv cannot be used for taking an inventory of API robots.

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:

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.

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 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.

Importing Images

▼ 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. Add the media IDs that have the backups to Media Manager on the server where you are going to import the backups.
2. In the NetBackup Administration Console, expand **NetBackup Management > Catalog**.
3. Select **Actions > Initiate Import**. A dialog appears titled Step 1: Read Catalog from Media.
 - ◆ 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.
 - ◆ In the **Media ID (to import)** field, type the Media ID of the volume that contains the backups you are importing.
 - ◆ If you're importing Backup Exec images, check **Password**, then enter the password.

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. If the password contains non-ASCII characters, the media can be imported only by using a NetBackup Windows interface or by using the `bpimport` command.

Click **OK**. The Confirm Initiate Import dialog appears.



4. Click **OK** to start the process of reading the catalog information from the source volume.
5. 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.*

2. 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 **Import** to search for imported images

Select the date range that includes the images to import

Images eligible for importing appear as a result

3. Select the image(s) you wish to import and Select **Actions > Import**. The Confirm Import dialog appears.
4. To view the log, click the **Results** tab, then select the import job log just created.



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.

▼ 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.

NetBackup provides reports for verifying, managing, and troubleshooting NetBackup operations. NetBackup reports show status or problem information for NetBackup servers or clients. The Troubleshooter is available to help analyze the cause of errors that can appear in a NetBackup report.

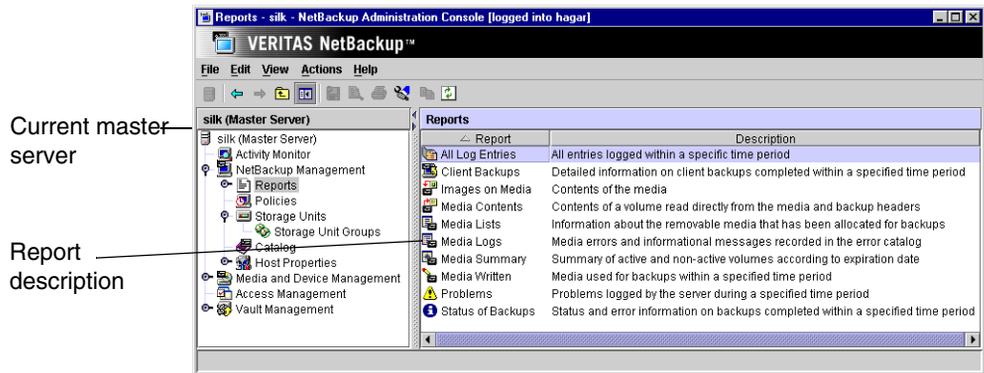
The following topics are discussed in this chapter:

- ◆ “Introduction to the Reports Application” on page 258
- ◆ “Reports Window” on page 260
- ◆ “NetBackup Report Types” on page 263
- ◆ “Using the Troubleshooter Within Reports” on page 275



Introduction to the Reports Application

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 263.



▼ To run a report

1. In the NetBackup Administration Console, expand **NetBackup Management > Reports**. A list of possible reports 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 420.)

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 263.

Menu Bar

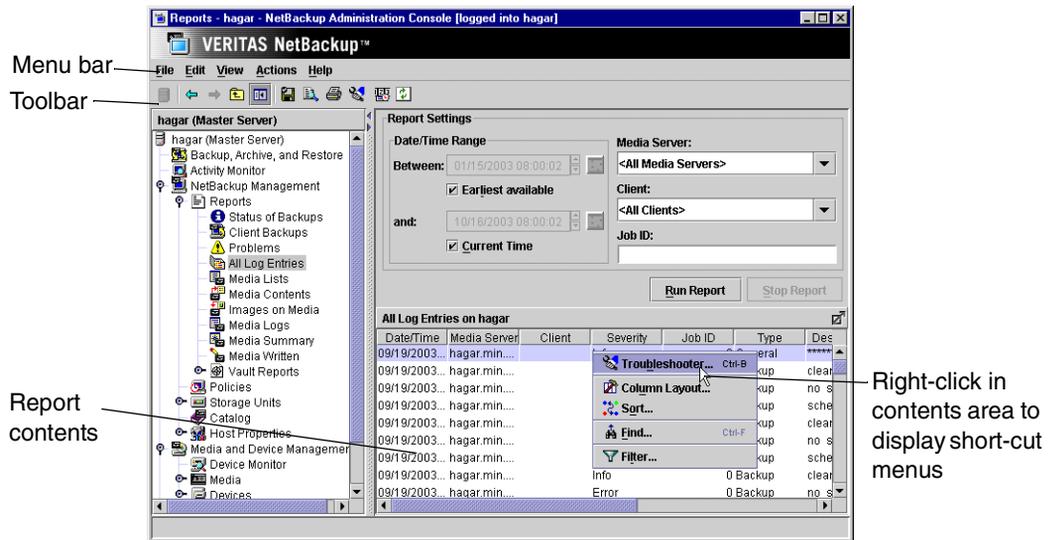
The Menu bar consists of the following menu items:

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 , Back , Forward , Up One Level , Options , Refresh , Column Layout , Sort , and Filter are described in the section, “View Menu” on page 17.
Actions	<i>Applies to NetBackup Enterprise Server only:</i> View SANPoint Control: Launches the browser to display the SANPoint Control reports page. From there, the user can view NetBackup-related SANPoint Control logs. The logs are not viewed from the NetBackup Administration Console. Check the logs to help diagnose SANPoint problems, for example, look for intermittent drive failures and network errors.
Help	Options Help Topics , Troubleshooter , License Keys , Current NBAC User , and About NetBackup Administration Console are described in the section, “Help Menu” on page 21.



Reports Window

The Reports window contains a number of methods to make it easier for you 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 10.

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.

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 cover. 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 Global host property, **Keep Logs For**, determines the period of time for which the information is available.

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 its media servers appear in the report.

Job ID

Specify the Job ID for which you want the report.



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, click the hyperlink to open the Troubleshooter.



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	<i>Yes</i> , if 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 its 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 <i>0</i> 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	<i>Yes</i> 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, Standard, MS-Windows-NT, and so on).
Primary Copy	Primary copy shows which copy (1 or 2) NetBackup uses to satisfy restore requests.
Retention Period	Retention period for the backups on this volume. An asterisk after the retention period number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown was the first level assigned. (See “Retention” on page 169.)
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	Yes 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 266.)



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 Log 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 it performs the backup.
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, click the hyperlink to open the Troubleshooter.
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 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 catalog backups, keep a hard copy record or configure the E-mail global attribute. The E-mail global attribute 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 has 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.
Media Server	Media server that controlled the backup.



Media Lists Report (continued)

Column	Meaning
Number of Restores	Number of times this volume has been used for restores.
Partner ID	Media ID that is assigned when the volume is added to Media Manager.
Retention Period	How long the backups will be considered valid.
Sector Size	Optical sector size in bytes.
Status	<p>The messages that commonly appear under Status are the following:</p> <p>Suspended: The volume cannot be used for further backups until retention periods for all backups on it have expired. At that time, the suspended volume is deleted from the NetBackup media catalog and unassigned from NetBackup. (The <code>bpmmedia</code> command can also be used to manually suspend or unsuspend volumes.)</p> <p>A suspended volume is available for restores. If the backups have expired, the backups first require importing.</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>Full: The volume 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.</p> <p>A full volume is unavailable for future backups until the retention period expires for all backups that are on it. At that time, the volume is deleted from the NetBackup media catalog and unassigned from NetBackup.</p> <p>Expired: All backups have expired.</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>
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 MPX.

Media Lists Report (continued)

Column	Meaning
Volume Pool	A number that corresponds to the volume pool for the media. 0 = None 1 = NetBackup For other numbers, find the media ID in the volume list on the volume database host for the media ID. This list shows the volume pool for each volume.



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 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 Period	Period of time that NetBackup retains the backup. An asterisk after the retention period number means that the volume can have backups with different retention levels and that the number shown was the first level assigned. (See "Retention" on page 169.)



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 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 (Yes or No).
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	Y if 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	Y if 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. IDX (Index file) if the fragment contains true image restore information or is for an individual-file-restore-from-raw backup. For TIR backups, this, displays as TIR.



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 has the backup image. For disk, it is a pathname.
Media Server	Server with the database that has this information.
Media Type	Type of storage and can be removable (Rmed) or disk (Disk).
Multiplexed	Y if 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, Standard, MS-Windows-NT, and so on).
Remainder	Bytes written beyond kilobytes filed. Size of fragment is exactly: Kilobytes*1024 + Remainder.
Retention Period	Retention period for the backups on this volume. An asterisk after the retention period number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown was the first level assigned. (See "Retention" on page 169.)
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 266.)

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 its 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 267.)

The only expired volumes that appear in this report are those that are FROZEN. NetBackup deletes other expired volumes from its media catalog when it runs backups. An expired volume with other status can show up only if you run the report between the time the volume expires and the next backup is done.



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 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 volume database with the records for this volume.
Retention Period	How long the backups will be considered valid.
Times Written	Number of times this media was written.

Using the Troubleshooter Within Reports

You can 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.

You can also open the Troubleshooter at any time (**Help > Troubleshooter**), enter a status code, then click **Lookup**.





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 278
- ◆ “Jobs Tab” on page 285
- ◆ “Daemons Tab” on page 289
- ◆ “Processes Tab” on page 291
- ◆ “Media Mount Errors” on page 292
- ◆ “Managing the Jobs Database” on page 293



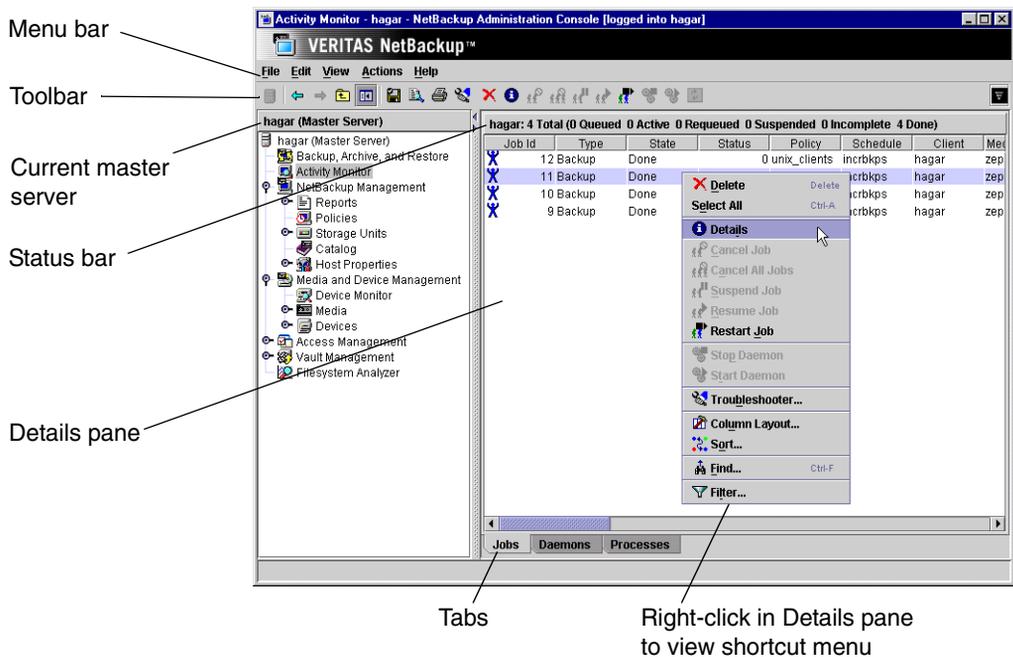
Introduction to the Activity Monitor

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 `bpjobjd` daemon continuously communicates and pushes NetBackup job activity to the Activity Monitor.

When receiving job activity data from the `bpjobjd` 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:



Menu Bar

The 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	Options Delete , 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 , Column Layout , Sort , and Filter are described in the section, "View Menu" on page 17. The Options option allows you to configure certain aspects of the Activity Monitor. (See "Setting Activity Monitor Options" on page 284.) The Filter option 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.



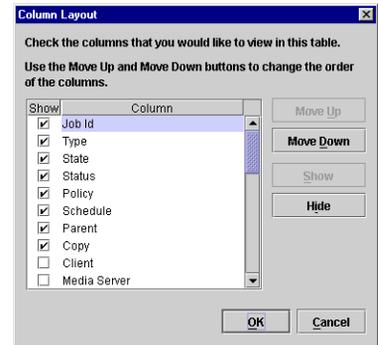
Option	Description
Actions	<p>The Actions menu contains the following options:</p> <p>Details: Displays detailed information about the job, daemon, or process you select in the list.</p> <p>Cancel Job: Cancels uncompleted jobs that you have selected in the Jobs list. A cancelled checkpointed backup or restore job cannot be resumed from the last checkpoint. If this is desired, use Suspend Job instead.</p> <p>Cancel All Jobs: Cancels all uncompleted backup jobs.</p> <p>Suspend Job: Suspends an Active, Queued or Requeued 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 365 and “Move Backup Job from Incomplete State to Done State” on page 365.)</p> <p>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 <i>Any available</i> 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.</p> <p>A job may be in a 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 365 and “Move Backup Job from Incomplete State to Done State” on page 365.)</p> <p>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.</p> <p>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 79.)When a job is restarted, a new job ID is assigned to the restarted job.</p> <p>Note The job must be marked as Done to restart it. To restart an Incomplete or Suspended job, cancel the job to force it into the Done state.</p> <p>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:</p> <pre>timestamp Job manually restarted as new_jobID.</pre> <p>Stop Daemon: Stops daemons that you have selected in the Daemons list.</p> <p>Start Daemon: Starts daemons that you have selected in the Daemons list.</p>
Help	<p>Options Help Topics, Troubleshooter, License Keys, and About NetBackup Administration Console are described in the section, “Help Menu” on page 21.</p>

The following sections describe common Activity Monitor operations using the menu options:



▼ To view specific column heads in the Details pane

1. Open the Activity Monitor.
2. Click **View > Columns Layout**. The **Column Layout** dialog appears, showing the current settings.
3. Select a column heading you wish to show or hide.
 - ◆ Select the **Show** button or the **Hide** button to display or hide the column.
4. Select a column heading you wish to move up or down in the order of appearance. Select 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.

To cancel all uncompleted jobs in the jobs list, click **Actions > Cancel All Jobs**.



▼ **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**

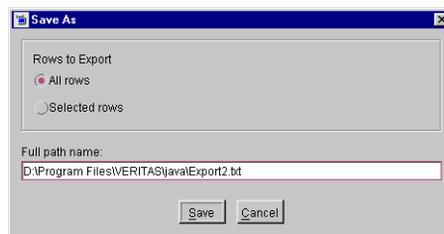
1. Open the Activity Monitor and select the **Jobs** tab.
2. Select the suspended or incomplete job you want to resume.
3. Select **Actions > Resume Job**. All selected jobs are resumed.

▼ **To 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.

▼ **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 the information you want to export:
 - ◆ All jobs visible in the Jobs tab, or
 - ◆ Only the rows that are currently selected in the Activity Monitor.
4. Enter the full pathname to the file where you want the job data to be written, then click **Save**.



▼ 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. Open the Activity Monitor and select the **Jobs** tab.
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**.
 - ◆ Right-click on the job and select **Troubleshooter**.
 - ◆ Open the job details for a job, click the Detailed Status tab. Then click **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 look up the troubleshooting information. You can open the Troubleshooter at any time and enter a status code.

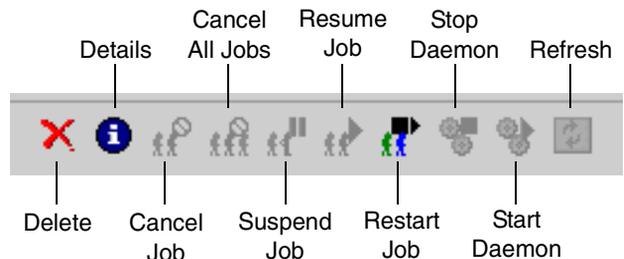
Shortcut Menus

Right-clicking in the Activity Monitor Details pane produces different shortcut menus, depending on which tab is displayed.

Activity Monitor Toolbar

The buttons on the toolbars 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.



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.
- ◆ The number of jobs in each of the job states: Active, Queued, Requeued, Suspended, Incomplete, and Done.

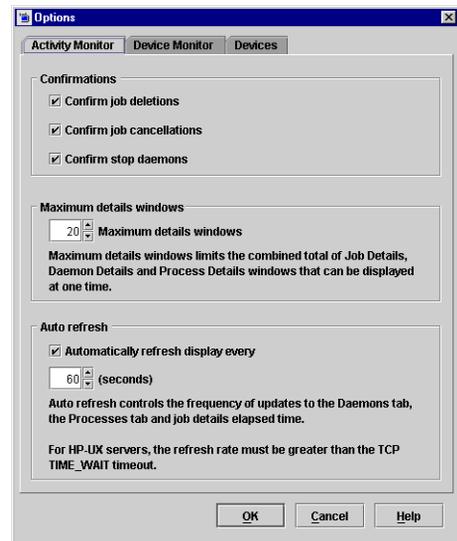
The numbers always reflect the actual number of jobs, even when filtering is used.

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.

Jobs Tab

The **Jobs** tab displays all jobs that are in process or have been completed for the master server currently selected. The Activity Monitor contains detailed information for the following types of jobs:

- ◆ Backup
- ◆ Archive
- ◆ DB Backup (backup of the catalog database)
- ◆ Duplicate
- ◆ Import
- ◆ Restore
- ◆ Verify
- ◆ Vault
- ◆ Label
- ◆ Erase

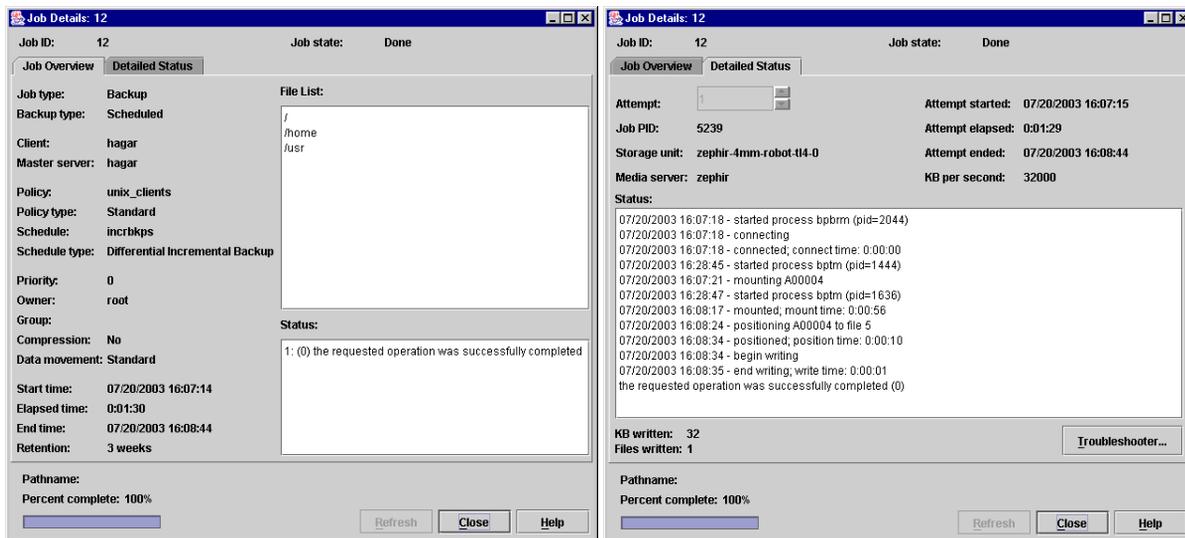
Viewing Job Details

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:

- ◆ Job Overview tab, containing general information about the entire job.
- ◆ Detailed Status tab, containing specific information about job attempts.



The Job Details dialog contains detailed job information about the selected job.



The following table describes fields that appear in the Jobs tab and/or the Job Details tabs:

Field	Description
Active elapsed	The time since the most recent attempt became active.
Active start	The time when the most recent attempt became active.
Attempt	The selector for the attempt number. Used to display detailed information for the selected attempt if NetBackup tried more than once. This field is disabled if there is only one attempt. (Detailed Status tab)
Attempt elapsed	The elapsed time for this attempt. (Detailed Status tab)
Attempt ended	The time when this attempt ended. (Detailed Status tab)
Attempt started	The time when this attempt began. (Detailed Status tab)
Client	The name of the client associated with the job.
Data movement	Distinguishes between synthetic backups (<i>Synthetic</i>), standard backups (<i>Standard</i>), and disk staging (<i>Disk Staging</i>).
Elapsed time	The amount of time that has elapsed since the job was initially queued.

Field	Description
End time	The date and time that the operation was completed.
Files	The number of files written for the last backup of the policy and schedule. (Detailed Status tab)
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 the backup is multiplexed, all jobs associated with the same multiplexed storage unit have the same PID. (Detailed Status tab)
Job state	<p><i>Queued:</i> Jobs in the NetBackup scheduler queue. A queued restore job is one for which NetBackup is still determining which files are needed.</p> <p><i>Active:</i> Currently active jobs.</p> <p><i>Requeued:</i> Jobs that are placed back in the scheduler queue as retries because the previous attempt was unsuccessful.</p> <p><i>Incomplete:</i> 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. 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. (See “Move Restore Job From Incomplete State to Done State” on page 365 and “Move Backup Job from Incomplete State to Done State” on page 365.)</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> <p><i>Done:</i> Completed jobs.</p>
KB per second	The data transfer rate in kilobytes per second. (Detailed Status tab)
KB per second	The average data transfer rate in kilobytes per second over the length of the current attempt. (Detailed Status tab)
Kilobytes	The number of kilobytes that have been written.
Master	The master server on which the job is run.
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.
Media server	The NetBackup server controlling the media. (Detailed Status tab)
Operation	For Active jobs, this indicates the operation that is currently being performed.



Field	Description
Owner	The owner of the job.
Parent	When Vault is used, each child job refers to the parent job by this number. The Parent Job ID for a parent job is 0.
Pathname	For Active jobs, this is 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.
% 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> .
Profile	The name of the profile that defines the processing to be done by a Vault job. Multiple profiles can be configured for the Vault. (Appears for Vault jobs only.)
Robot	The robot used for a Vault job. (Appears for Vault jobs only.)
Schedule	The name of the schedule that NetBackup is using to back up the client.
Schedule type	The type of schedule controlling the backup. For example, Full or Cumulative-Incremental. <i>Unknown schedule type</i> refers to a schedule associated with a disk staging storage unit.
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.
State	The current state of the job.
Start time	The date and time that the first attempt was initially queued.
Status	<ul style="list-style-type: none">On the Job Overview tab: 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.On the Job Details tab: Events that have occurred up to this point. For example, this box contains entries for when the client connects to the server and when the server begins writing data. When the job is complete, the last line shows the completion status.

Field	Description
Storage unit	The name of the storage unit that the job is using. (Detailed Status tab)
Type	The type of backup: scheduled, user-directed, immediate (manual backup), or archive.
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 master server you are monitoring.

More About Services:

Standalone daemons: These NetBackup daemons are always running and listening to accept connections.

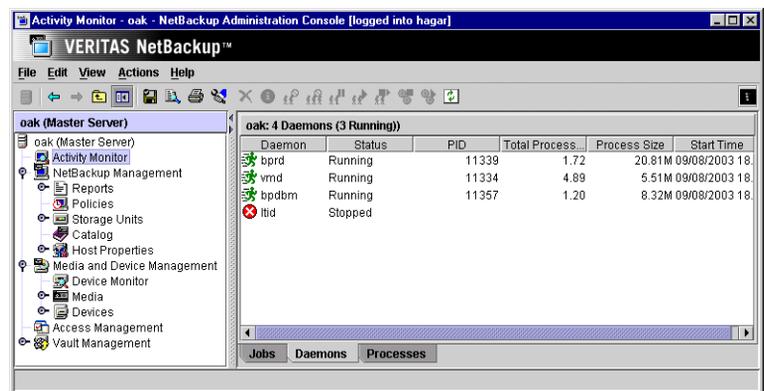
Examples include `bpdbm`, `bprd`, `bpjobd`, `vmd`, the robotic daemons, `nbdbd`, and `visd`.

The Global Data Manager product's Persistent Storage daemon, (`nbdbd`) runs only if the GDM server or GDM-managed server license is added to the NetBackup server. `visd` (the Global Data Manager product's Information Server daemon), requires `nbdbd` to be running before `visd` will start. `visd` is active only if a GDM server or managed server license is present.

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.

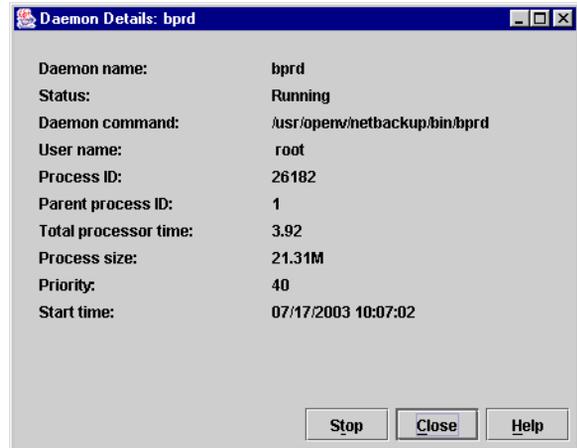
inetd daemons: NetBackup daemons that are usually launched by way of `inetd` (`1m`) or `bpinetd`. Examples include `bpcd`, `bpjava-msvc`, `vopied`, and `vnetd`.



Note After restarting daemons in the Activity Monitor or by using a command, we recommend exiting all instances of the NetBackup-Java Administration Console, then restarting the console using the `jnbSA` command. (The `jnbSA` command is described in the *NetBackup Commands for UNIX* guide.)

▼ 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.



The following table describes the fields that appear in the Daemons tab and the Daemon Details dialog:

Field	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> .

Field	Description
Total processor time	The processor time used by the daemon in seconds.
User name	The user name under which the daemon was started.

▼ 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**. Or, right-click the daemon and select **Start Daemon** or **Stop Daemon** from the shortcut menu.

Processes Tab

The **Processes** tab displays the NetBackup processes running on the selected master server.

▼ 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.



The following table describes the fields that appear in the Processes tab and the Process Details dialog:

Field	Description
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 name	The name of 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 processor 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 either becomes 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 greater likelihood that the backup is not delayed.

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 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 293.

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*.

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/openv/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*.

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:

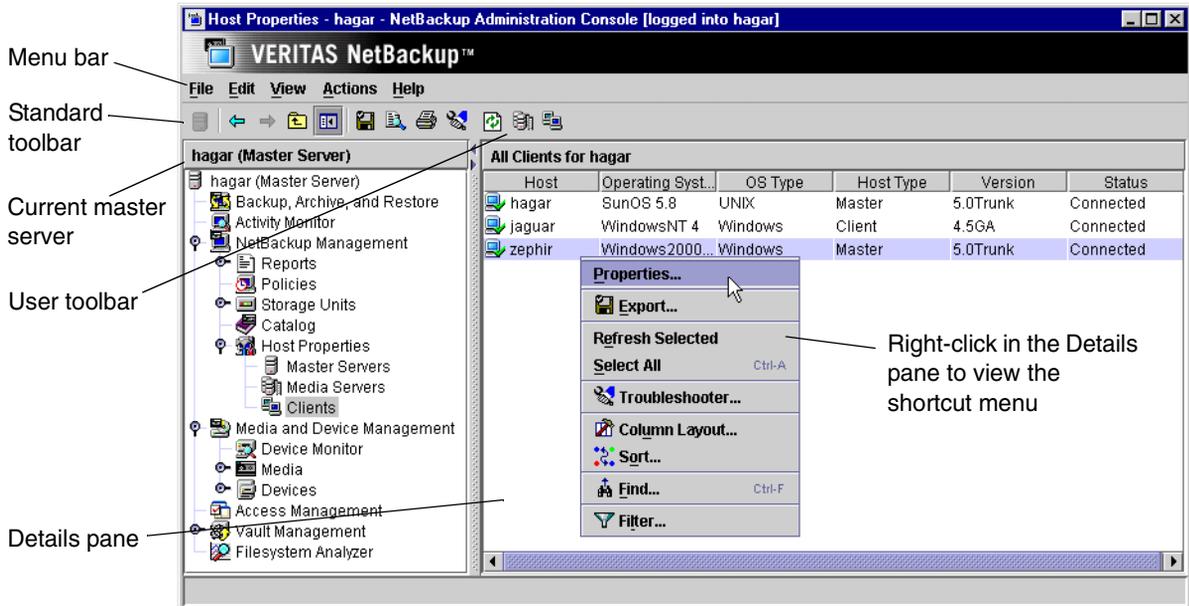
- ◆ “Introduction to Host Properties” on page 298
- ◆ “Menu Bar” on page 299
- ◆ “Viewing Host Properties” on page 300
- ◆ “Changing Host Properties” on page 300
- ◆ “Required Permissions” on page 303
- ◆ “Master Server, Media Server, and Client Host Properties” on page 304



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.

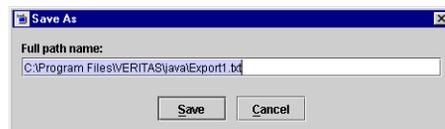
The Host Properties nodes in the Administration Console tree and the Details pane contain the following information:



Menu Bar

The 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 17.
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 21.



Viewing Host Properties

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

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.



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.

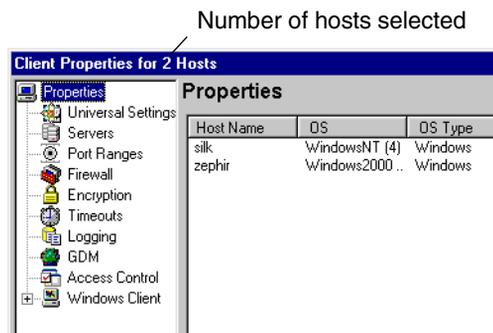
▼ To simultaneously change the properties on multiple hosts

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 420.)

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 389.)



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

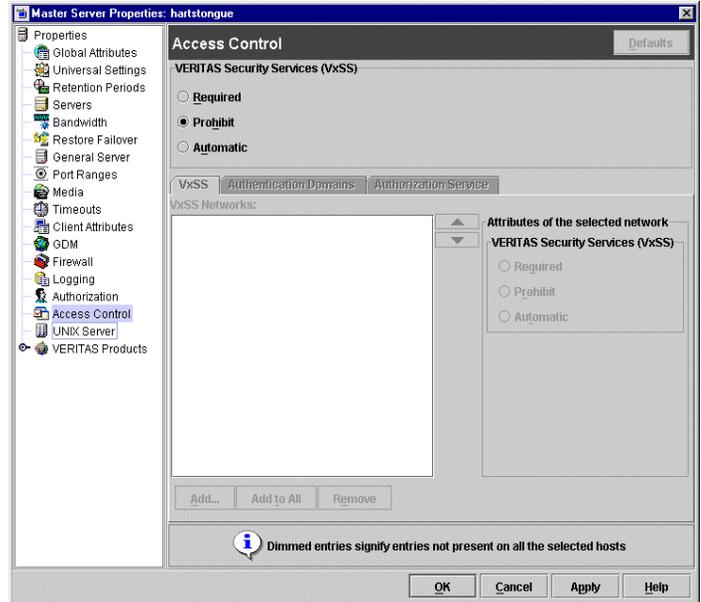
The VxSS tab contains a list of networks that are allowed or (not allowed) to use VxSS with the local system.

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**.

Add Button

To add a network to the **VxSS Network** list, click **Add**. The **Add VxSS Network** dialog displays, containing the following properties:

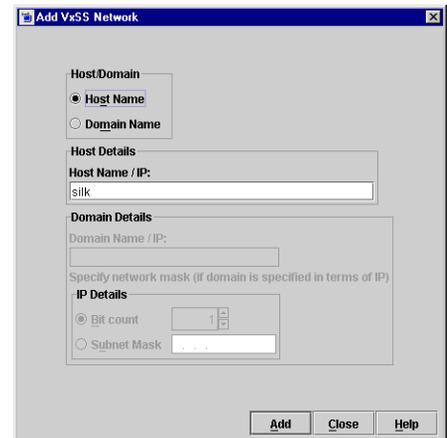
Host/Domain

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.

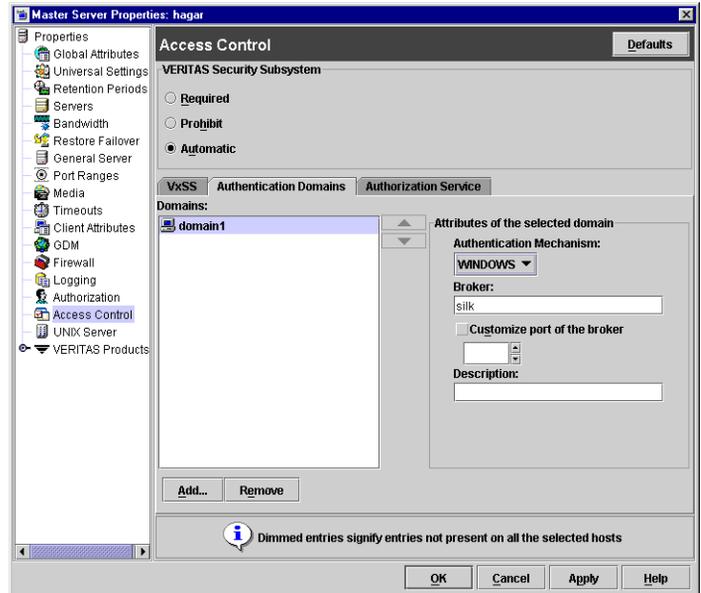
Remove Button

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.



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:

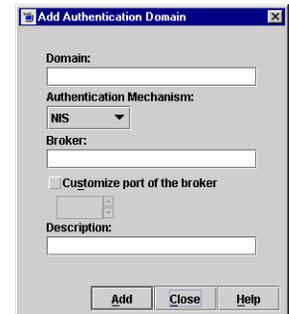
NIS: The Network Information Service, version 1.

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.

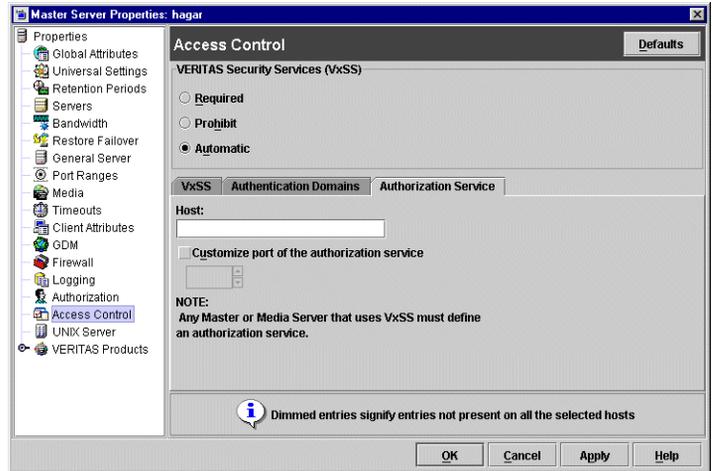
Note If configuring this tab for a media server using Access Control, you must define the host that will perform authorization.

Host Name

Enter the host name or IP address of the authorization service.

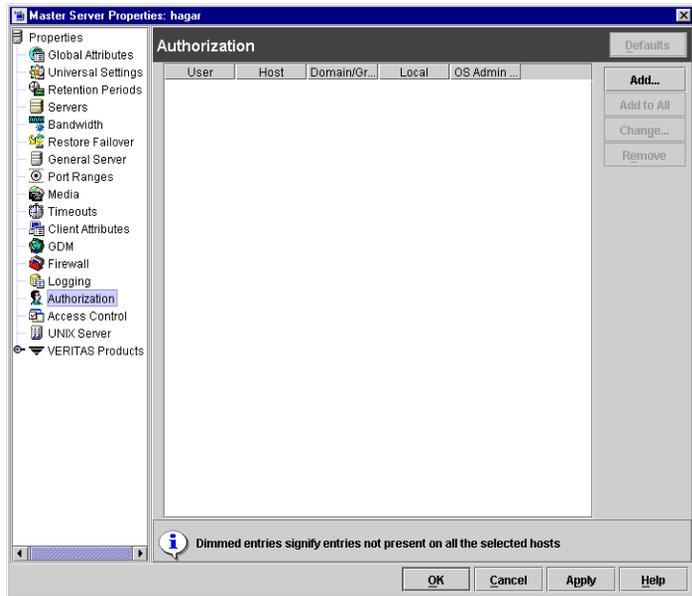
Customize the Port Number of the Authorization Service

To use a non-standard port number, select **Customize the Port Number** and enter the port number of the authorization service.

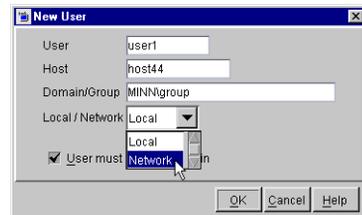


Authorization Properties

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.

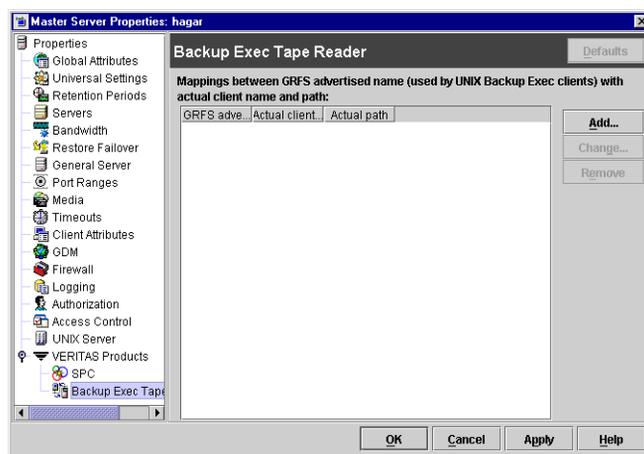
For configuration information, see “Enhanced Authentication and Authorization” on page 71 in *NetBackup System Administrator’s Guide, Volume II*.



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 249.)



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.



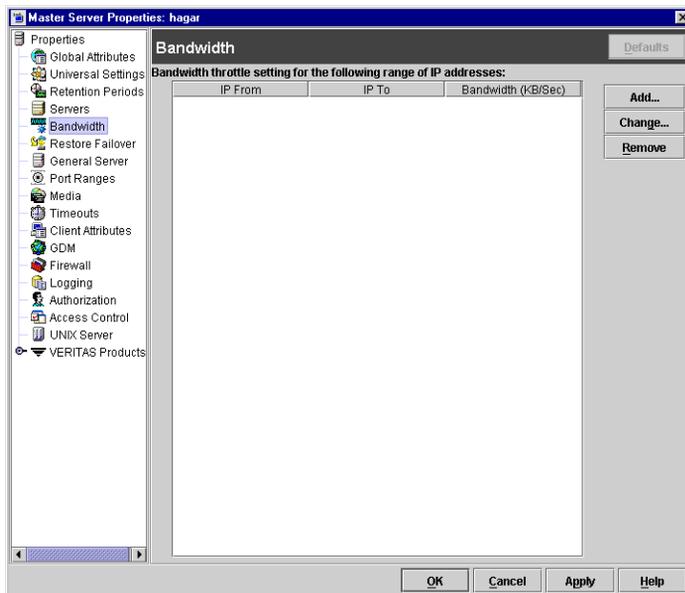
Bandwidth Properties

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. By default, the bandwidth is not limited.

The limiting occurs on the client side of the backup connection and applies only to backups. Restores are unaffected.

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.

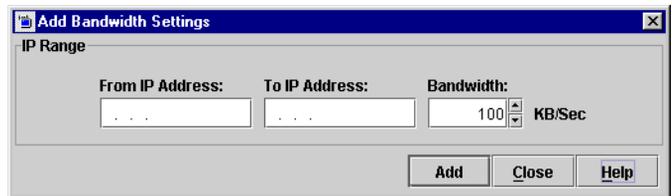
For example, a value of 200 indicates 200 kilobytes per second.

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.



The screenshot shows a dialog box titled "Add Bandwidth Settings" with a close button (X) in the top right corner. Below the title bar, the text "IP Range" is displayed. The dialog contains three input fields: "From IP Address:" with a text box containing "...", "To IP Address:" with a text box containing "...", and "Bandwidth:" with a text box containing "100" and a dropdown arrow. To the right of the "Bandwidth:" field is the unit "KB/Sec". At the bottom of the dialog, there are three buttons: "Add", "Close", and "Help".

Remove Button

Click the **Remove** button to remove a selected entry from the bandwidth table.



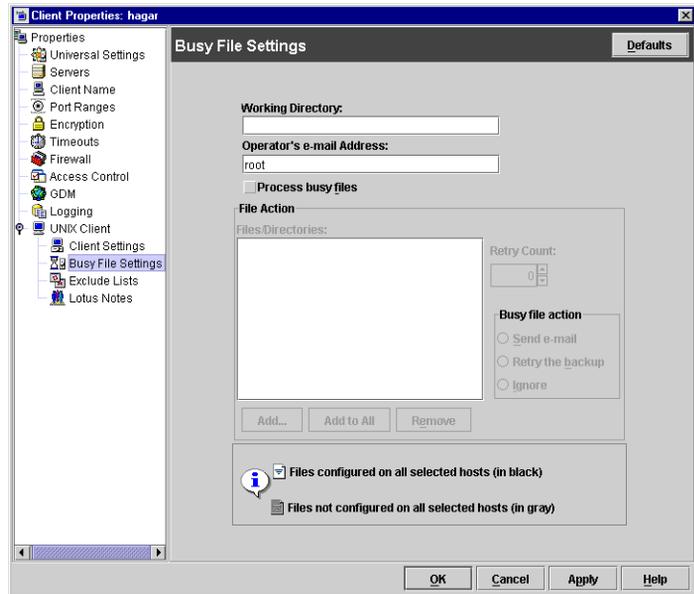
Busy File Properties

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/obj/np/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 122 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.



Client Attributes Properties

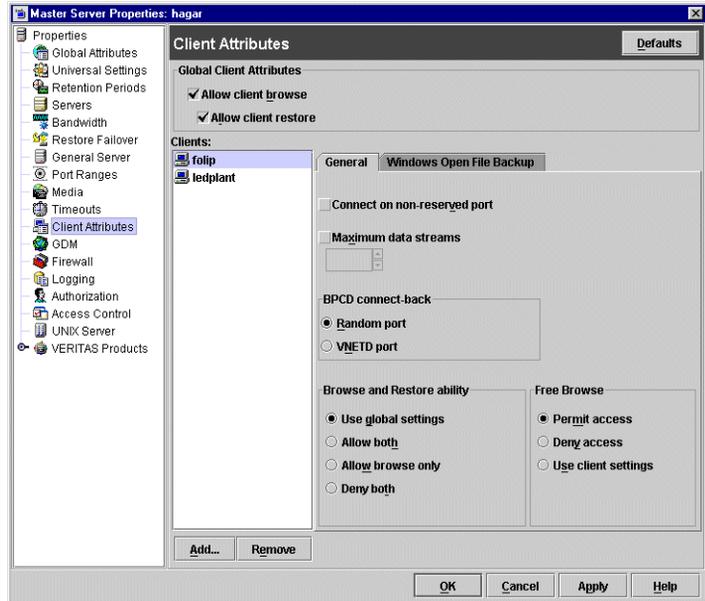
Client Attributes

properties apply to clients of currently selected master servers. **Client Attributes** contain two subtabs:

- ◆ **General Tab** (described below)
- ◆ **Windows Open File Backup** (described in “Windows Open File Backup Properties” on page 406)

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/opencv/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/opensv/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 322.)

Note If you are using dynamic addressing (DHCP), use the `bpclient` command to add clients to the client database. (See “Dynamic Host Name and IP Addressing” on page 113 in the *NetBackup System Administrator’s Guide, Volume II* for instructions.)

Add Button

Click the **Add** button 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 Properties” on page 406.

Connect on Non-reserved Port

The **Connect on Non-reserved Port** property specifies that the currently selected server use a non-reserved port when connecting to the clients selected in the General tab of **Client Attributes**. To enable **Connect on Non-reserved Port**:

1. Select the desired client in the Clients list box on the **General** tab in the **Client Attributes** dialog.
2. Select the **Connect on Non-reserved Port** check box.
3. Enable **Accept Connections on Non-reserved Ports** for each of the selected clients. This client property is found under **Host Properties > Clients > Universal Settings**. (See “Accept Connections on Non-reserved Ports” on page 392.)



BPCD Connect-back

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

- ◆ By using a **Random Port**: NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method.
- ◆ By using the **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.

Maximum Data Streams

The **Maximum Data Streams** property specifies that there is a maximum number of concurrent data streams allowed for the clients selected in the General tab of the **Client Attributes** dialog.

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 318 and “Allow Client Restore” on page 318), 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.

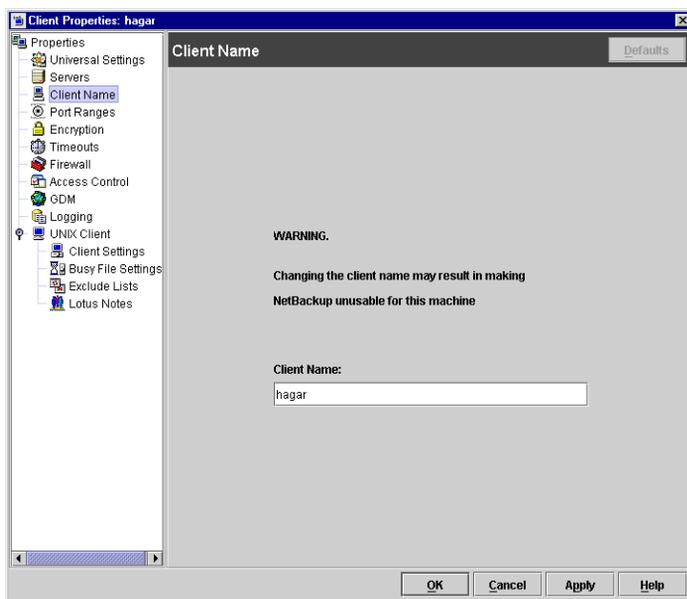


Client Name Properties

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 318.)

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.



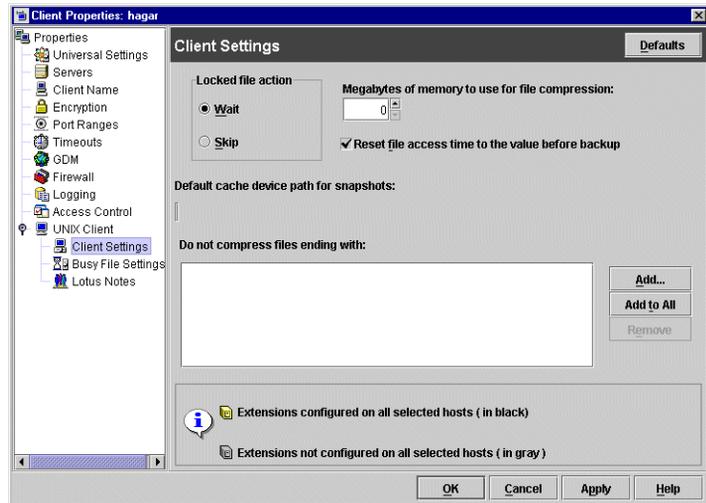
Client Settings (UNIX) Properties

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. A message is logged if waiting was necessary.
- ◆ **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.



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.

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.

Add to All Button

Use the **Add to All Lists** 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 Lists**.

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

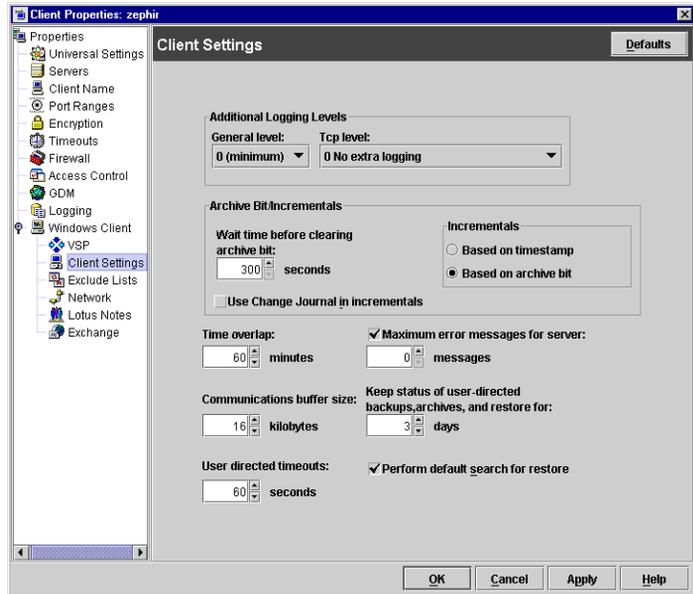
The **General Level Logging** property enables `bpineted`, `bbpkar`, `tar`, and `nbwin` logging. Scroll to the desired level of logging. The higher the level, the more information is written. Default: 0.

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 affect while using the following options or products:
 - ◆ True Image Restore (TIR) (See “Collect True Image Restore Information” on page 87.)
 - ◆ True Image Restore with Move Detection (See “Collect True Image Restore With Move Detection” on page 87.)
 - ◆ Synthetic backups (See “Synthetic Backups” on page 154.)
 - ◆ 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 NetBackup 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 Repetitive Error Messages for Server

The **Maximum Repetitive 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 of Backup Images for Restore

The **Perform Default Search of Backup Images 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.

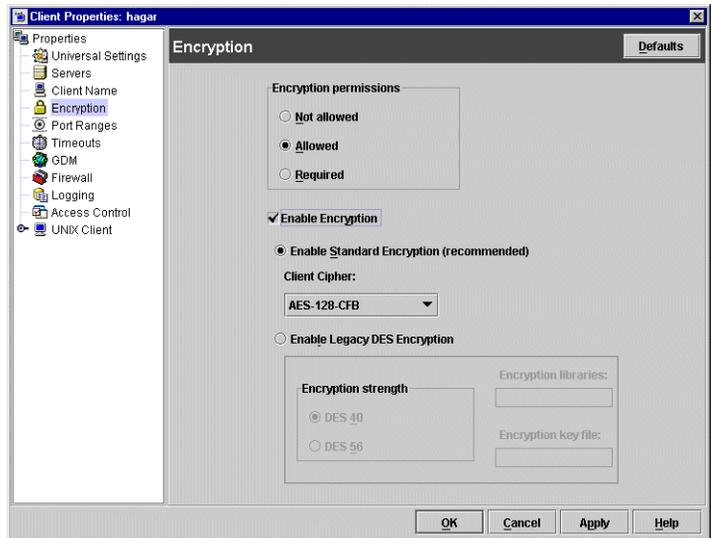
Clear the **Perform Default Search** check box to disable the initial search. With the property disabled, the NetBackup Restore window does not display any files or folders upon opening. Clicking a backup image, or selecting a range of backup images, starts a search. Default: option is enabled.

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 errors.
- ◆ **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 errors.

Enable Encryption

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



Use Standard Encryption

The **Use 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, **Use 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

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

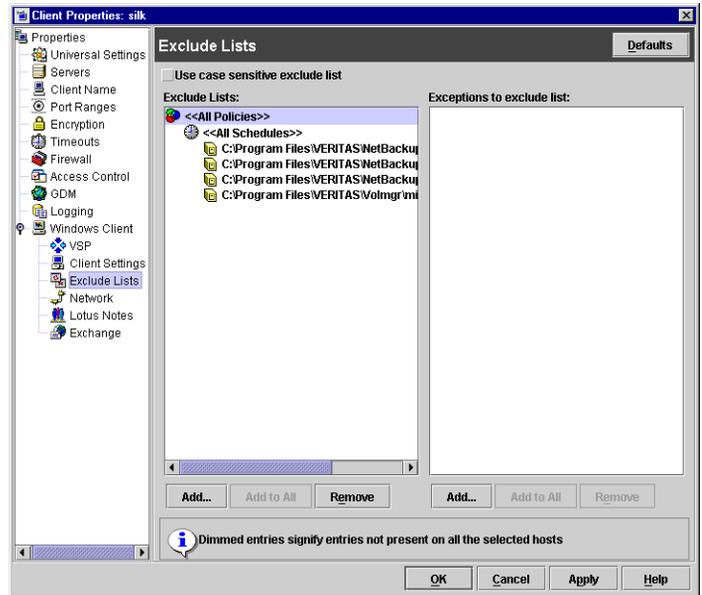
- ◆ A mailbox alias

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 "Creating an Exclude List on a UNIX Client" on page 139.

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.

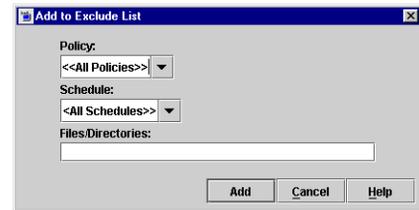
Add Buttons

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 affect 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 affect 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.)



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.

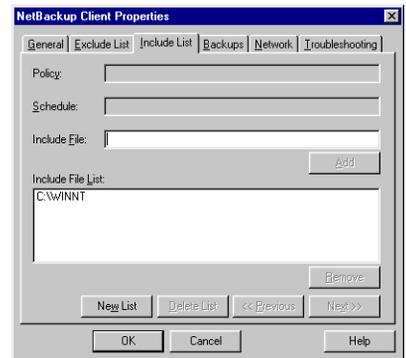
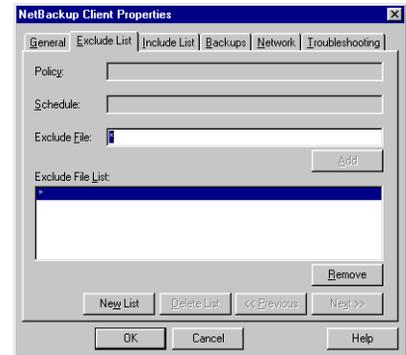
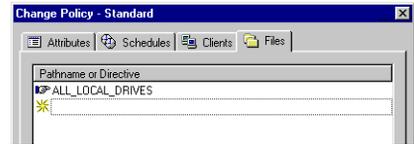
If the include list did not contain any entry, no directories would be traversed.

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



/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

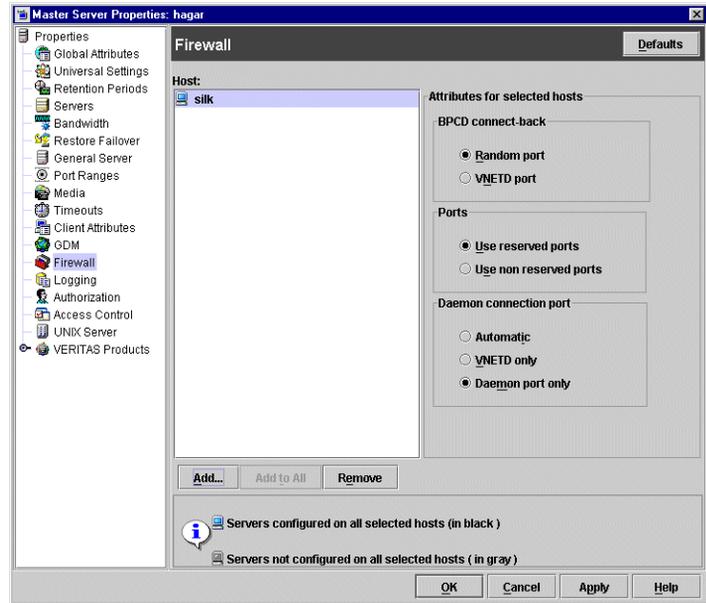
The **Firewall** properties apply to selected master servers, media servers, and clients.

Host

The Host list displays the configuration for the host(s) for that are currently selected.

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(s) 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.

BPCD Connect-back

The **BPCD Connect-back** property specifies how daemons are to connect back to BPCD (the NetBackup Client daemon):

- ◆ By using a **Random Port**: NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method.
- ◆ By using the **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.

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**

Select **Use Reserved Ports** to connect to the server using a reserved port number.

- ◆ **Use Non Reserved Ports**

Select **Use Non Reserved Ports** to connect to the server using a non-reserved port number. If using, also enable **Accept Connections from Non-reserved Ports** for the selected server. (See “Accept Connections on Non-reserved Ports” on page 392.) This property is located on the 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 will be connected to using `vnetd` if possible. If using `vnetd` is not possible, the connection will be made using the daemon’s traditional port number. (Automatic is the default.)

- ◆ **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 traditional port number, check this option.

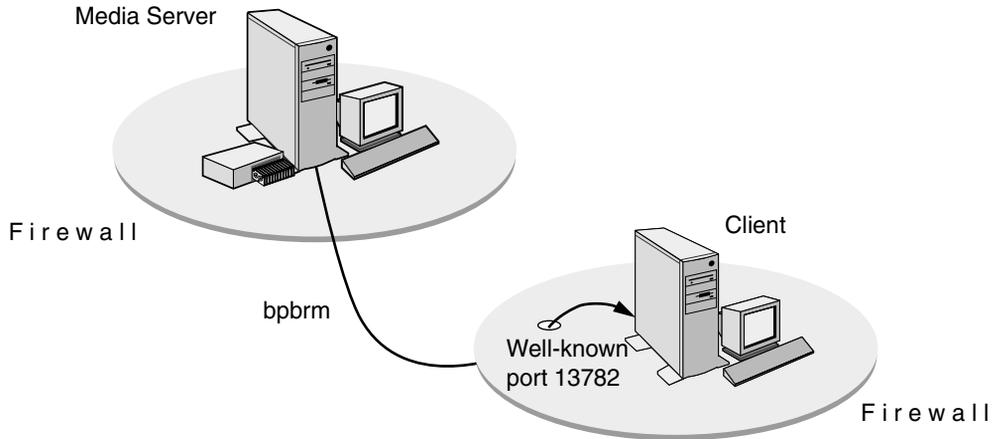
- ◆ **Daemon Port Only**

The daemons on the server will be connected to using only the traditional port number. If the server is running NetBackup software earlier than NetBackup 4.5 with the feature pack 3, use this option.

NetBackup ports are also discussed in “Configuring NetBackup Ports” on page 458.

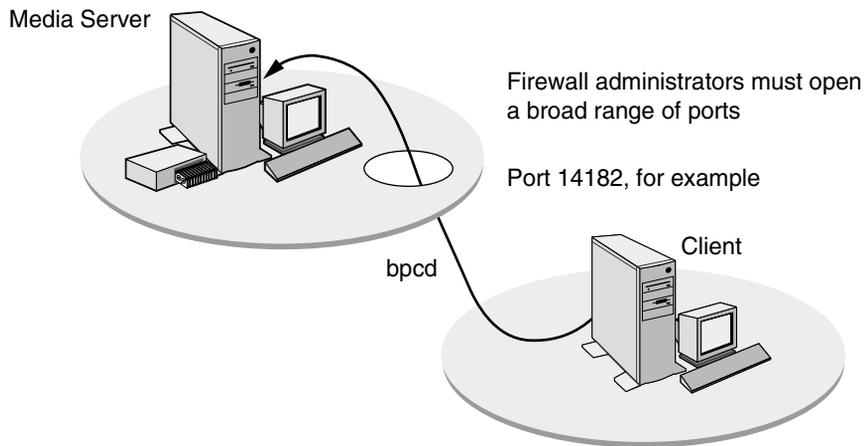


Media Server Running `bpbrm` and Client Running `bpcd`



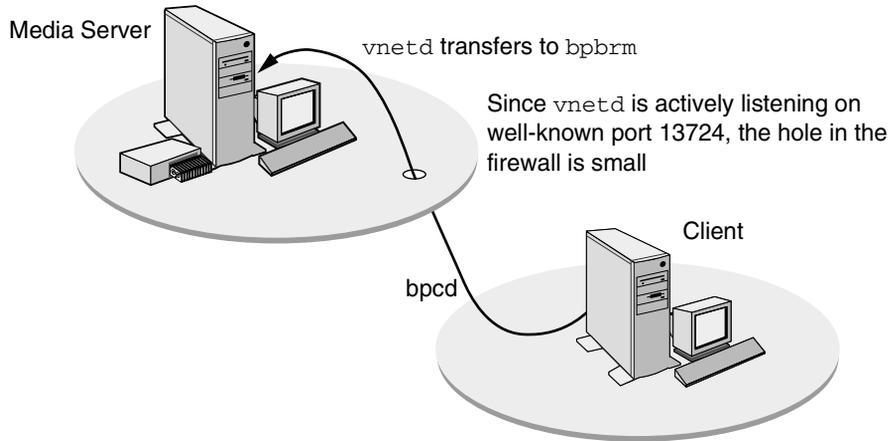
Using the traditional call-back method, `bpbrm` is not actively listening on a port for a specific connection from `bpcd`. Because `bpcd` could connect back to the media server on one of many ports, firewall administrators must make more ports available on the firewall to accommodate that communication.

Traditional Call-back Method



The no call-back method uses `vnetd`. `vnetd` uses the well-known port 13724 to actively listen for a specific connection from another NetBackup process. If a firewall is in place, administrators need only to leave port 13724 open. `vnetd` will transfer a socket from itself to another process on the same machine.

No Call-back Method Using vnetd



Note Both servers and clients must have NetBackup version 4.5 or later installed for vnetd to work.

Minimum Required Connections

The following table lists the minimum ports that must be open for NetBackup to operate with a firewall in place.

If a computer is performing more than one responsibility (for example, a master server is also performing media server operations), that computer can be classified as a media server and a master server. Therefore, when looking at the tables, examine all the tables for all the operations that will be performed by that computer.

Between a media server and a client:

Media Server		Client
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound

Between a master server and a client (user backup or restore):

Master Server		Client
Outbound >	vopied >	Inbound (if authentication)



Outbound >	bpcd >	Inbound (if progress logging or DHCP)
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Inbound	< bprd	< Outbound

Between a master server and a client (multi-streamed scheduled backup):

Master Server		Client
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound

Between a master server and a media server:

Master Server		Media Server
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Inbound	< bpjobd	< Outbound
Inbound	< bpdbm	< Outbound
Inbound	< bprd	< Outbound

Between a media server and a media server:

Media Server		Media Server
Outbound >	bpcd >	Inbound

Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Outbound	robotic daemons >	Inbound

Between a media server and the volume database host:

Media Server		Volume Database Host
Outbound >	vnetd >	Inbound
Outbound	vmd >	Inbound

Between a media server and the global device database host:

Media Server		Global Device Database Host
Outbound >	vnetd >	Inbound
Outbound	vmd >	Inbound

Between a media server and the SSO device allocation host:

Media Server		SSO Scan Host
Outbound >	vnetd >	Inbound
Outbound	vmd >	Inbound

Note The SSO device allocation host is the host that is serving as the volume database host for the robot with shared drives.

Between a media server and the SSO scan host:

Media Server		SSO Scan Host
Outbound >	vnetd >	Inbound
Outbound	vmd >	Inbound



Between a media server and an NDMP server:

Media Server		NDMP Server
Outbound >	ndmp >	Inbound

Between a NDMP tape/data server and an NDMP tape/data server:

NDMP Tape/Data Server		NDMP Tape/Data Server
Outbound >	ndmp >	Inbound

Note In the preceding tables, an NDMP server refers to either a physical NDMP host or the Remote NDMP functionality. The Remote NDMP functionality resides on a NetBackup media server but it is not being considered as being part of the NetBackup media server.

Between a Windows System Administration Console and a client:

Windows Console		Client
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Inbound	< vopied	< Outbound

Between a Windows System Administration Console and a media server:

Windows Console		Media Server
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Inbound	< vopied	< Outbound (if authentication/authorization)
Outbound >	vmd >	Inbound
Outbound >	Robotic daemons>	Inbound

Between a Windows System Administration Console and a master server:

Windows Console		Master Server
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Inbound	< vopied	< Outbound (if authentication/authorization)
Outbound >	bprd >	Inbound
Outbound >	bpdbm >	Inbound
Outbound >	bpjobd >	Inbound



Between the NetBackup-Java Console and a NetBackup-Java application server:

Java Console		Java Server
Outbound >	bpjava-msvcd >	Inbound
Outbound >	vnetd >	Inbound
Outbound >	bpjobd >	Inbound

Between the NetBackup-Java Console Activity Monitor and a master server:

Java Console Activity Monitor		Master Server
Outbound >	vnetd >	Inbound
Outbound >	bpjobd >	Inbound

Between a NetBackup-Java application server and a client:

NetBackup-Java Application Server		Client
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Inbound	< vopied	< Outbound (If authorization/authentication)



Between a NetBackup-Java application server and a media server:

NetBackup-Java Application Server		Media Server
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Inbound	< vopied	< Outbound (If authorization/authentication)
Outbound >	vmd >	Inbound
Outbound >	Robotic daemons >	Inbound

Between a NetBackup-Java application server and a master server:

NetBackup-Java Application Server		Master Server
Outbound >	bpcd >	Inbound
Inbound	< callback	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound
Inbound	< vopied	< Outbound (If authorization/authentication)
Outbound >	bprd >	Inbound
Outbound >	bpdbm >	Inbound
Outbound >	bpjobd >	Inbound



Between a GDM Dashboard and a GDM server:

GDM Dashboard		GDM Server
Outbound >	visd >	Inbound
Inbound	< callback	< Outbound

Between a GDM server and a GDM-managed server:

GDM Server		GDM-Managed Server
Outbound >	visd >	Inbound
Inbound	< visd	< Outbound

Note The range of port numbers for the source ports for GDM connections cannot be configured. The allocation of the source port numbers for GDM connections is left up to the operating system.

Between a NetBackup Advanced Reporter (NBAR) browser and a NBAR server:

NBAR Browser		NBAR Server
Outbound >	Web Server >	Inbound

Between a NBAR server and a GDM-managed server:

NBAR Server		GDM-Managed Server
Outbound >	arbdb >	Inbound
Inbound	< bpcd	< Outbound
Inbound	< bprd	< Outbound
Inbound	< arbdb	< Outbound
Inbound	< vnetd	< Outbound
Outbound >	vnetd >	Inbound

Note The range of port numbers for the source ports for NBAR connections to the web server and arbdb cannot be configured. The allocation of the source port numbers for these connections is left up to the operating system.

▼ 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 441.)

▼ 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 148 in the *NetBackup System Administrator’s Guide, Volume II*.

▼ **To enable logging for vnetd**

Create a `vnetd` directory in the following location, then restart `vnetd`:

On Windows: `install_path\NetBackup\logs\vnetd`

On UNIX: `/usr/opensv/netbackup/logs/vnetd`

Example Setup for Using the vnetd Port

The following is a sample configuration to use the `vnetd` port for `bprd`, `bpdbm`, `bpjobd`, `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**.

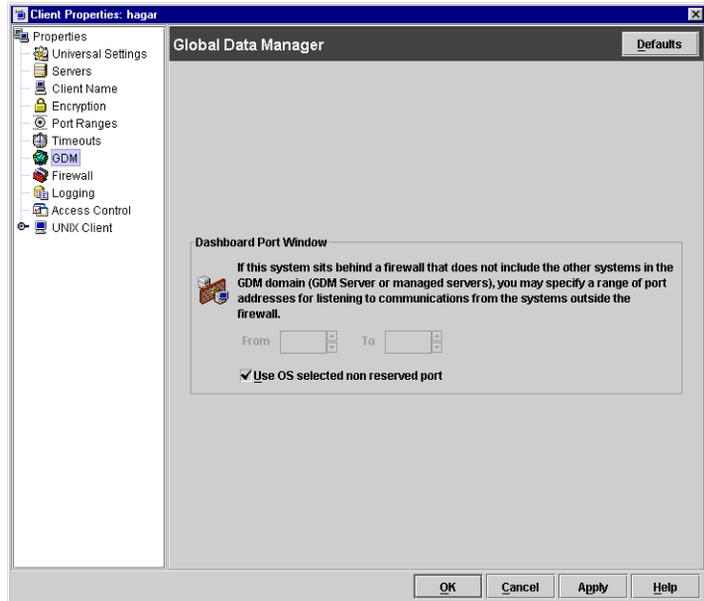


GDM (Global Data Manager) Properties

The **Global Data Manager** provides remote monitoring and remote administration of NetBackup master servers that have been grouped into a GDM domain. Additional configuration information about GDM can be found in the *Global Data Manager Administrator's Guide*.

Dashboard Port Window

The **Dashboard Port Window** property specifies the range of port addresses that can be used for listening to communications from systems outside of the firewall around the selected host(s).



Use OS Selected Non-reserved Port

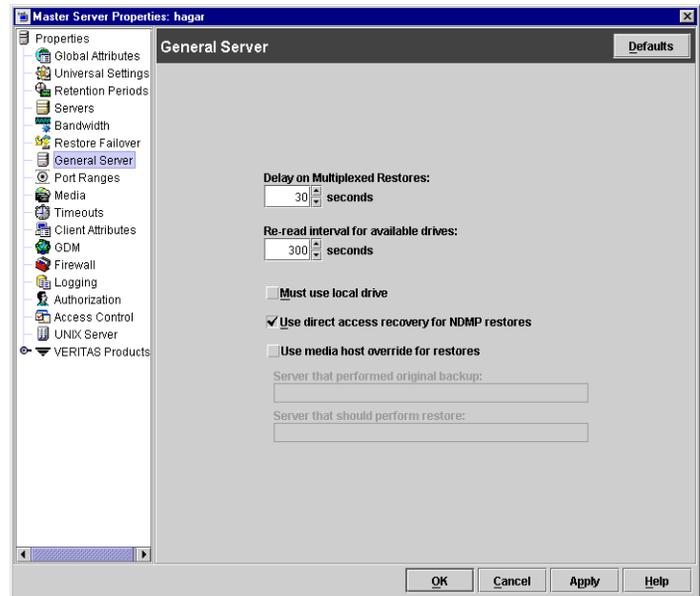
Select **Use OS Selected Non-reserved Port** if you want the operating system to determine which non-reserved port to use.

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 and 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.



Re-read Interval for Available Drives

The **Re-read Interval for Available Drives** property determines how often NetBackup checks storage units for available drives. If this value is too high, too much time elapses between drives becoming available and NetBackup discovering their availability, thus delaying backup jobs. If it is too low, checks are made more often than necessary thus wasting system resources. Default: 300 seconds (5 minutes).

Must Use Local Drive

This property appears for master servers only.

If the client is also a master server and the **Must Use Local Drive** check box is checked, backups must occur on a local drive. If the client is not a master server, this setting has no effect.

Note Although this property appears as a setting for media servers as well, the media server must also be a master server in order for the property to take effect.



This setting 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 and then mix the media-server clients with other clients in a policy that uses ANY AVAILABLE storage unit. When a backup starts for a client that is a SAN media server, the backups go to the SAN connected drives on that server.

Use Direct Access Recovery for NDMP Restores

By default, NetBackup for NDMP is configured to use Direct Access Recovery (DAR) during NDMP restores. DAR can 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.

Allow Block Incrementals

The **Allow Block Incrementals** property specifies that block level incrementals are allowed. The NetBackup Block Level Incrementals option must be installed to use this setting.

Use Media Host Override

The **Use Media Host Override** property forces restores to go to a specific server, 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.

▼ **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 and update the Media Manager volume database to reflect the move.
2. Modify the NetBackup configuration on the master server by specifying the original media server in the **Server than performed original backups** box and the new media server in the **Server that should perform restore** box.
3. Stop and restart the NetBackup Request Manager service on the master server.

This applies to all storage units on the original media server. That is, restores for any storage unit on the **Server than performed original backup** host will now go to the **Server that should perform restore** host.

To revert to the original configuration for future restores, clear the check box.

Server than performed original backup: Specifies the media server that performed the original backup and to which a restore request will normally go.

Server that should perform restore: Specifies the media server to which the restore request is being forced to go.

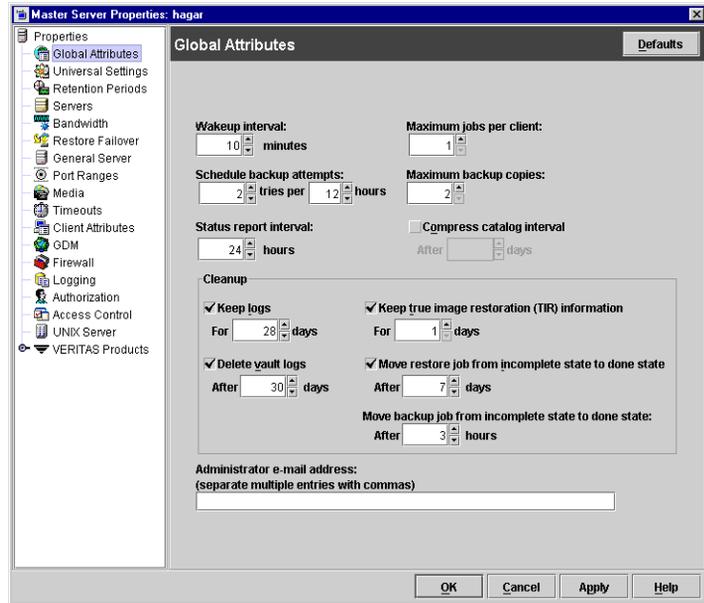


Global Attributes Properties

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.

Wakeup Interval

The **Wakeup Interval** property specifies how often the scheduler checks schedules for backups that are due. Long wakeup intervals can cause the scheduler to start too late in a backup window to complete all the backups for a schedule. Minimum setting: 1 minute. Default: 10 minutes.



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.

Retries do not occur until all backups on the worklist have been tried at least once within the backup window. If the backup window closes before the retry starts, the job fails with a status code 196. Default: 2 tries in 12 hours.

Status Report Interval

The **Status Report Interval** property specifies the default time period during which NetBackup accumulates information to put into a report. For example, a setting of 8 hours provides a report covering the previous 8 hour period. Minimum setting: 1 hour. Default: 24 hours.

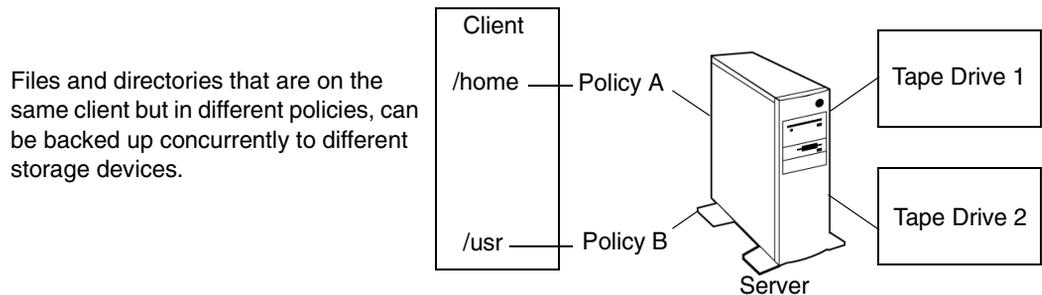


Maximum Jobs per Client

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.



You can specify any number of concurrent jobs within the following constraints. Default: 1 job:

- ◆ 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 actual number that you can use depends on the hardware, operating system, and applications that are running.

Because **Maximum Jobs per Client** applies to all clients in all policies, set it to accommodate the client that can handle the lowest number of concurrent jobs.

- ◆ Network loading. The available bandwidth of the network affects how many backups can occur concurrently. For example, two exabyte 8500, 8 mm 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.



Maximum Backup Copies

The **Maximum Backup Copies** property specifies that the total number of backup copies that can be created is from 1 to 10 copies.

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. NetBackup uses NTFS file compression and the catalog must be in an NTFS partition for compression to occur.

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:

- ◆ 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 Global 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 389 and “Checkpoint Restart for Restore Jobs” on page 446.)

Move Backup Job from Incomplete State to Done State

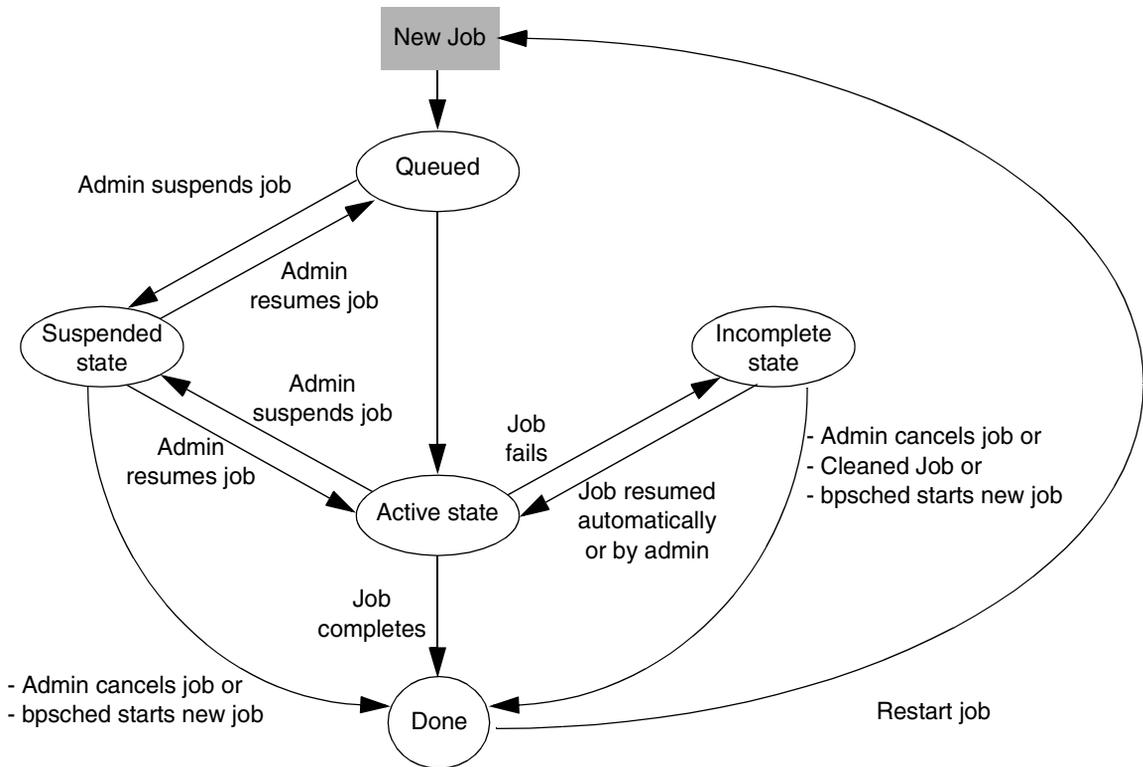
The **Move Backup Job From Incomplete State to Done State** property indicates the maximum number of hours that a failed backup job can remain in an incomplete state before Activity Monitor shows the job as done. Minimum setting: 1 hour. Maximum setting: 72 hours. Default: 3 hours.

The following figure depicts the different states for a checkpointed backup job:

When an active job errors, 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.





Administrator's E-mail Address

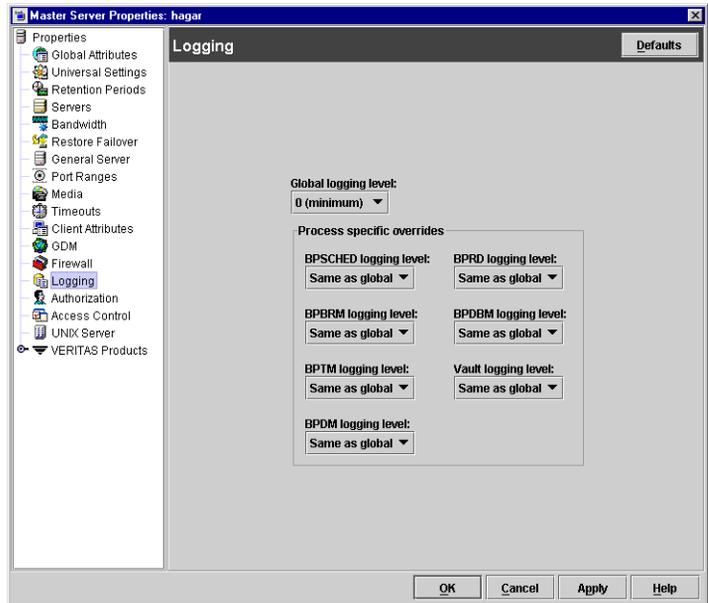
The **Administrator's E-mail Address** property specifies the address(es) where NetBackup sends notifications of scheduled backups, administrator-directed manual backups, or NetBackup catalog backups. The notification of catalog backups includes the media ID that was used.

Multiple addresses can be indicated, but need to be separated by commas. Default: no address.

On Windows NetBackup servers, it may be necessary to configure the `install_path\NetBackup\bin\nbmail.cmd` script in addition to specifying the above address. This is necessary because on Windows servers, NetBackup performs the notification by passing the specified E-mail address, subject and message to the script. The script then uses the mailing program that you specified in the script to send E-mail to the user. See the comments in the script for configuration instructions. Default: `nbmail.cmd` does not send E-mail.

Logging Properties

The **Logging** properties apply to currently selected master servers, media servers, and clients. The available properties differ between a server and a client.



Indicating a logging level does not enable logging.

You must first create a log directory for every process in the following location:

```
/usr/opensv/netbackup/logs/process_name
```

Global Logging Level

The **Global Logging Level** setting 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.

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**.



BPSCHED Logging Level

If you wish to override the **Global Logging Level**, select a logging level for `bpsched`: 0 (minimum) through 5 (maximum).

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).

Lotus Notes Properties

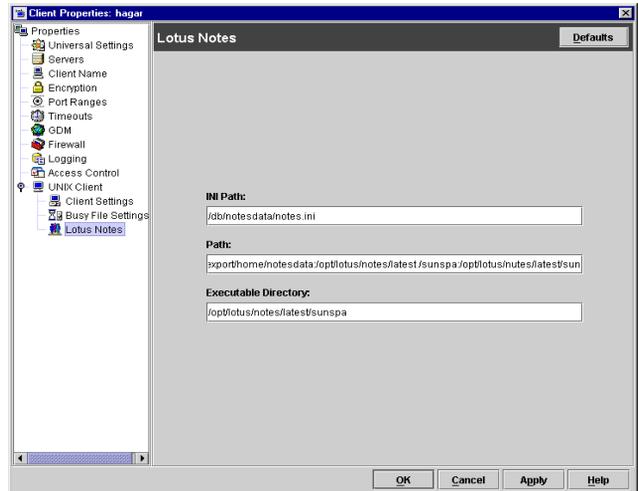
The **Lotus Notes** properties apply to currently selected clients running NetBackup for Lotus Notes.

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.

Executable Directory

The **Executable Directory** property is available on UNIX clients only. To specify an alternate Lotus program files directory other than specified in the **Path** directory, specify the directory in the **Executable Directory** box. For example:

```
/opt/lotus/notes/latest/sunspa
```

NetBackup for Lotus Notes for UNIX will now look for the Lotus program files in the `LOTUS_NOTES_EXECDIR` directory, rather than in the `LOTUS_NOTES_PATH` directory.

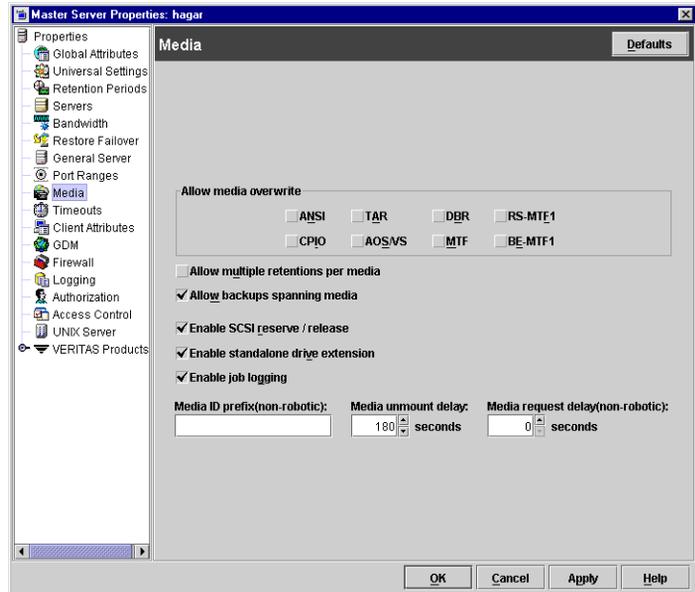


Media Properties

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.

Note You must allow overwriting the MTF1 format if you are using RSM robots. This is necessary because Free Media Labels are in MTF1 format.

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.

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 allows backups to span more than one media. If the end of media is encountered and this option is not selected, the media is set to FULL and the operation terminates abnormally (applies to both robotic and nonrobotic drives). By default, **Allow Backups to Span Media** is checked and backups can span media.



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.

NetWare Client Properties

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

Netware Client properties include:

- ◆ “Client Settings (NetWare) Properties” on page 323
- ◆ “Open File Backup (NetWare Client) Properties” on page 375
- ◆ “OTM Properties” on page 375



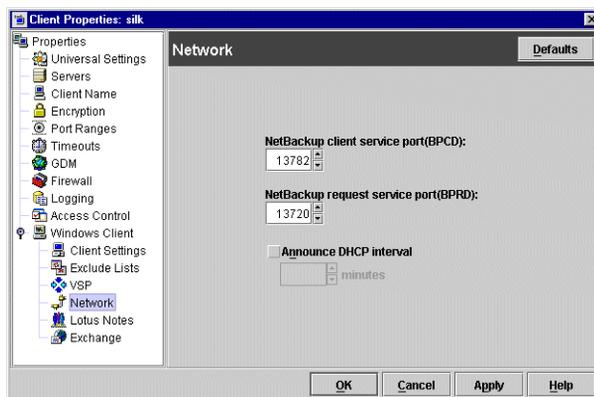
Network Properties

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 318.)

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.

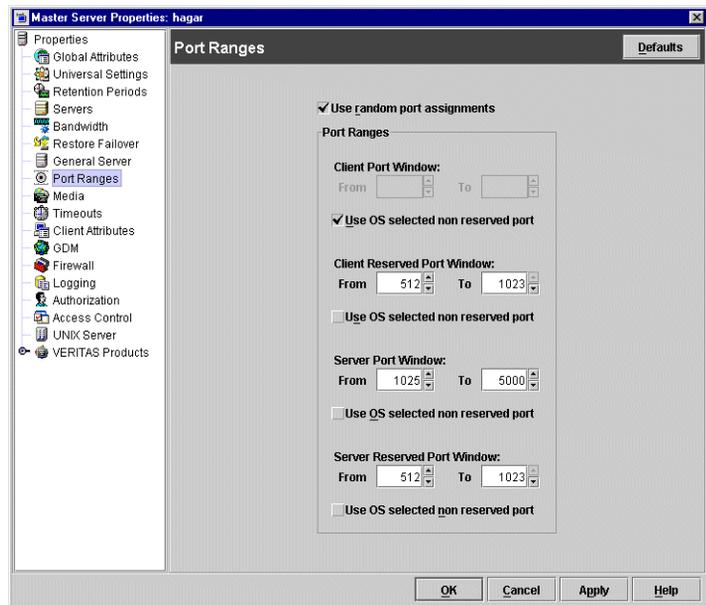
By default, **Use Random Port Assignments** is selected, and ports will be chosen randomly.

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.

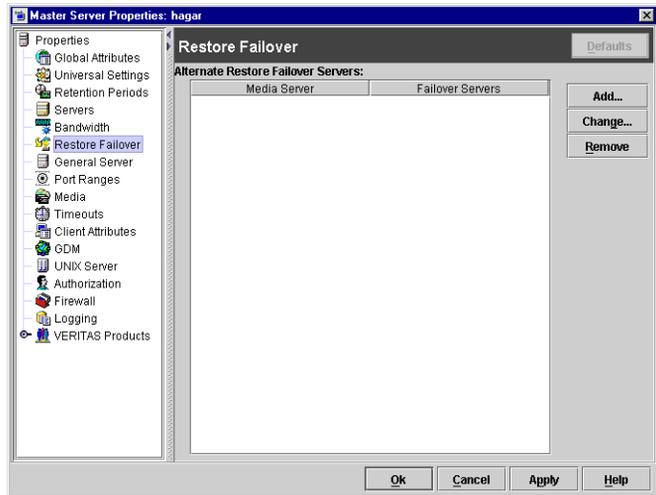


Restore Failover Properties

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.

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 a media server to the Alternate Restore Failover Machine list

4. In the **Server** field, specify the media server for which you're setting up failover protection.
5. 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.
6. Click **Add**.
7. Click **Apply** to accept the changes. Click **OK** to accept the changes and close the host properties dialog.
8. 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 456.

Change Button

To change an entry in the **Alternate Restore Failover Machines** list, select a media server, then click **Change**.

Remove Button

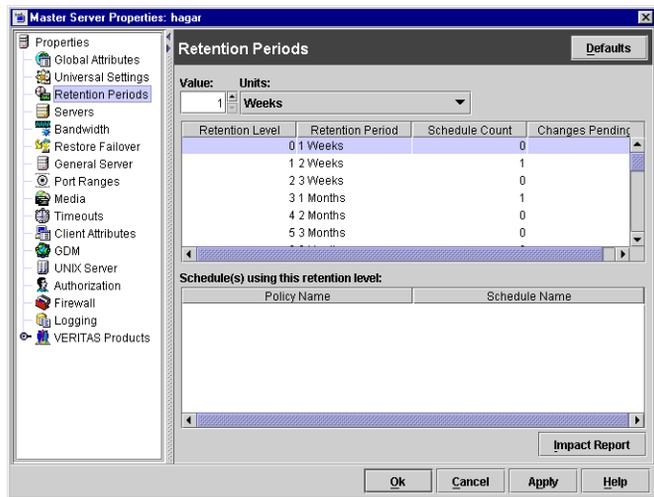
To remove a NetBackup media server from the **Alternate Restore Failover Machines** list, select the media server to be removed, then click **Remove**.



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

Specifies the retention level setting.

Units

Specifies the units of time for the retention period. The list also includes the special units, **Infinite** and **Expires Immediately**.

Retention Periods List

Lists 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

Lists the schedules that use the currently selected retention level, and the policy to which each schedule belongs.

Impact Report Button

Displays 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
```

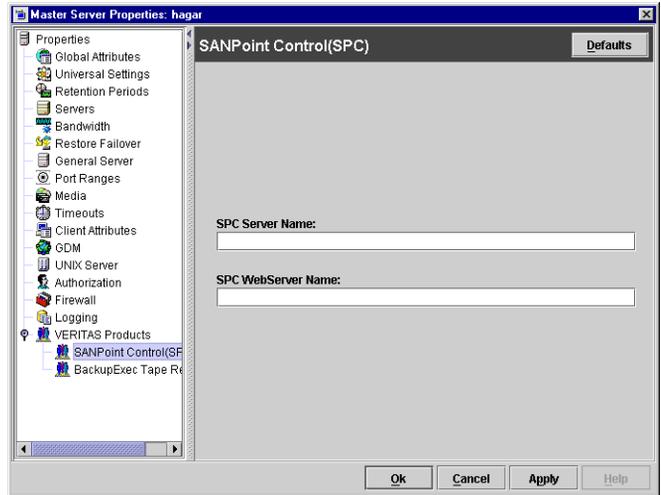


SANPoint Control (SPC) Properties

The SANPoint properties apply to selected NetBackup Enterprise master servers.

VERITAS SANPoint Control is a standalone product used to assist in troubleshooting NetBackup problems that may be related to the storage area network (SAN).

A trial version of SANPoint Control is available for evaluation to NetBackup customers. (See “More About SANPoint Control” on page 384.)



SPC Server Name

The **SPC Server Name** specifies the host where the master SANPoint Control server is installed.

If a server name alone is indicated, the SPC server port defaults to 8181. A port can also be indicated in the following format:

server_name:port_number

Where *server_name* is the name of the server, and *port_number* is the port that the SPC server will use.

For example, `ida.minnesota.com:9292`

SANPoint version 3.6 or later must be installed and configured correctly in order to view the console. A SANPoint Control-supported web browser must be installed on the local system in order to run the console. SANPoint Control supports Netscape 6.2 or later.

SPC WebServer Name

The **SPC WebServer Name** specifies the host where the SANPoint Control WebServer is installed. The SANPoint Control Web Console allows you to view your SAN from any host system that has a supported browser and network access to the appropriate SANPoint Control host.

If a web server name alone is indicated, the SPC webserver port defaults to 8181. A port can also be indicated in the following format:



webservice_name:port_number

Where *webservice_name* is the name of the web server, and *port_number* is the port that the SPC webservice will use.

For example, `silk.california.com:9696`

Specifying the Browser Path

If you are running NetBackup Administration Console (jnbSA), let the console know the path to the web browser by adding a `BROWSER_PATH=`*path_to_browser* entry to the `nbu.conf` file. The default location for this file is: `usr/openssl/java/nbj.conf`

You might enter, for example:

```
...  
BROWSER_PATH=/usr/bin/netscape
```

More About SANPoint Control

The SANPoint Control Console provides a central interface from which you can view the SAN, create and modify policies, administer zoning and storage provisioning, generate reports, and launch third-party SAN management tools.

NetBackup also provides a NetBackup Edition of SANPoint Control that gives users SANPoint Control view-only privileges. To configure SANPoint Control, the full product needs to be installed.

VERITAS recommends that the SANPoint Control server be on a machine other than a NetBackup master server.

While it is not necessary for the SANPoint Control Server and the SANPoint Web Server to reside on the same host, both must have network connectivity with one another.

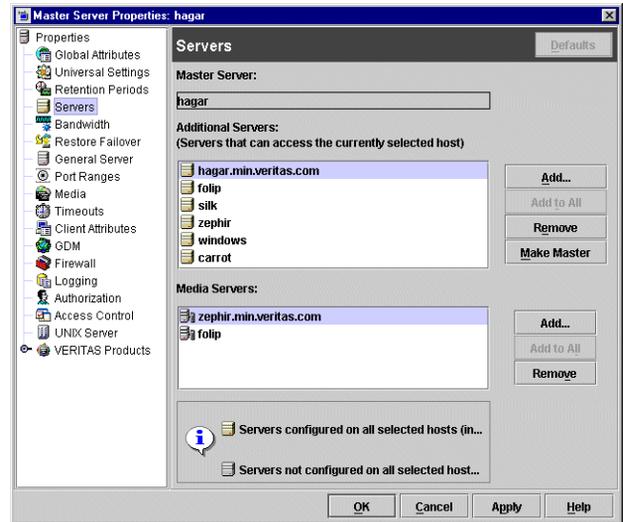
SANPoint Control can push SANPoint Control agents to remote machines so that the SANPoint Control has access to the machines.

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 420.

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.
- ◆ To delete a media server, select a media server from the list and click **Remove**.



Timeouts Properties

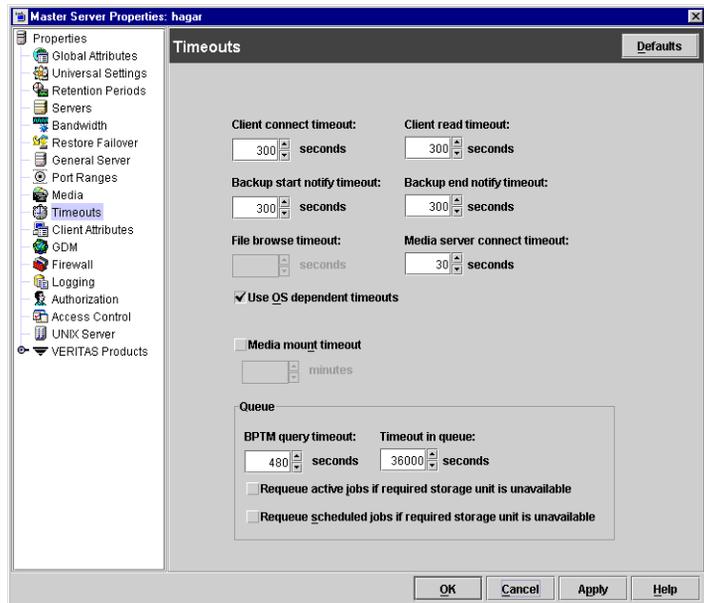
The **Timeouts** properties apply to selected master servers, media servers, and clients.

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 After** property specifies the number of minutes that NetBackup waits for the requested media to be mounted, positioned, and ready on backups and restores.

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.

BPTM (Drive Count) Query Timeout

The **BPTM (Drive Count) Query Timeout** property determines how long the scheduler waits for a drive-count query to bptm to complete. For timeout problems, increase this property to extend the time that the scheduler waits. Default: 480 seconds (8 minutes).

Timeout in Job Queue

The **Timeout in Job Queue** property determines how long a requeued job will wait for a required storage unit if the storage unit is currently unavailable. Default: 36000 seconds (10 hours.)

Requeue Active Jobs if Required Storage Unit is Unavailable

The **Requeue Active Jobs if Required Storage Unit is Unavailable** property causes active jobs to enter the requeued state if the required storage unit becomes unavailable (for example, a drive goes down). The jobs will then run when NetBackup can use the storage unit again.

A job fails if the **Requeue Active Jobs if Required Storage Unit is Unavailable** time expires or the backup window for the job closes before the storage unit becomes available. By default, this option is not selected and the job is not requeued.

Requeue Scheduled Jobs if Required Storage Unit is Unavailable

The **Requeue Scheduled Jobs if Required Storage Unit is Unavailable** property causes jobs to enter the requeued state on startup, if the required storage unit is not available. The jobs will then run when NetBackup can use the storage unit again.

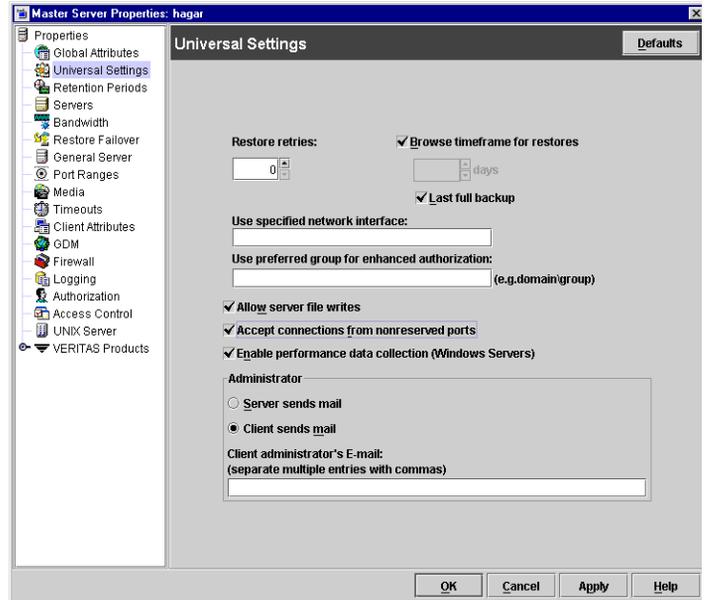
By default, the option is not selected and jobs fail with a status code 219 if the storage unit is not available.

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*.



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 (default), 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 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 to only to Windows master and media servers. The NetBackup performance counters can be viewed using the Windows 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.

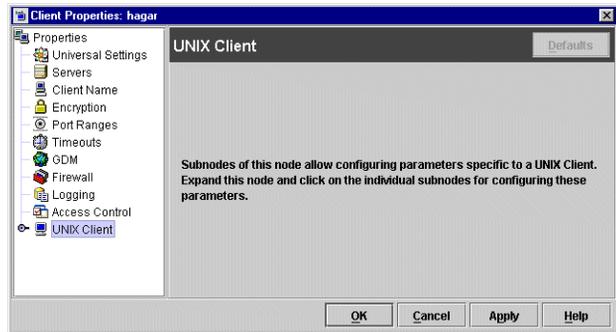


UNIX Client Properties

The **UNIX Client** properties define NetBackup properties of UNIX clients.

UNIX Client properties include:

- ◆ “Client Settings (UNIX) Properties” on page 324
- ◆ “Busy File Properties” on page 316
- ◆ “Lotus Notes Properties” on page 369

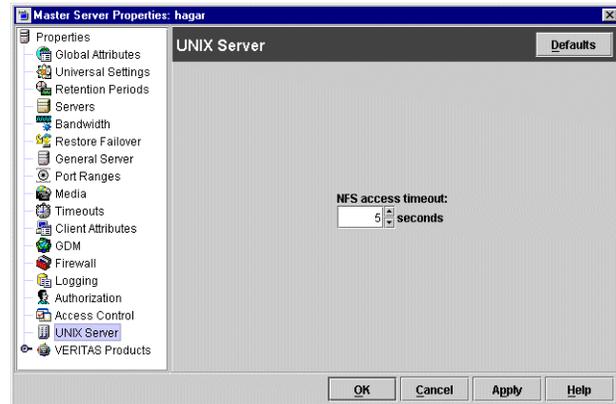


UNIX Server Properties

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 following subnodes:

- ◆ “Backup Exec Tape Reader Properties” on page 312
- ◆ “SANPoint Control (SPC) Properties” on page 383



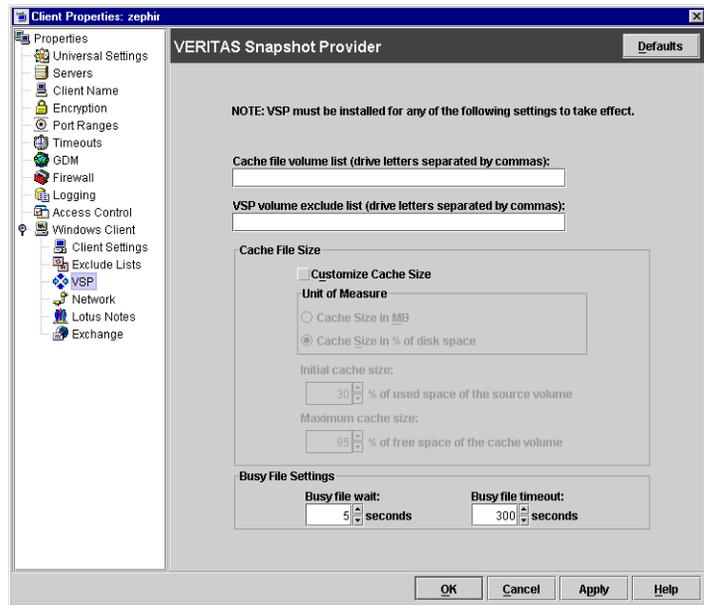
VSP (Volume Snapshot Provider) Properties

The VSP properties apply to currently selected Windows NT, Windows 2000, Windows XP, and Windows Server 2003 client(s).

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 318.)



VSP Overview

NetBackup uses VSP to back up open and active files on Windows NT, 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 403.

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 Multi-streamed 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 93.)

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 Multi-Streamed Backups Simultaneously with VSP

Assume two multi-streamed 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 multi-streamed backup jobs will be run (with each group running a backup job for each file list item). Both groups of multi-streamed 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 multi-streamed 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 multi-streamed 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 multi-streamed jobs were placed in the D drive. Both groups of multi-streamed 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 multi-streamed 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 **Initial 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 **Initial Cache Size** and **Maximum Cache Size** properties will use megabytes as the form of measurement.

Initial Cache Size

The **Initial Cache Size** property is the initial size for the VSP cache file when creating snapshots.

If the cache file is placed on a volume *not* being snapped, the **Initial Cache Size** property is ignored. If the cache file is placed on a volume being snapped, the size for a VSP cache file is at the **Initial Cache Size** property. The VSP cache file is sized differently, depending on the value used for the **Initial Cache Size** property and the **Customize Cache Size** property:

1. If **Customize Cache Size** is disabled, NetBackup automatically determines the initial cache size for the VSP snapshot if the cache file is being placed in the volume that is being snapped.
2. If **Customize Cache Size** property is enabled, NetBackup uses the user specified size in the **Initial Cache Size** property. If the VSP cache file is placed in the same volume as the volume being snapped, the cache file size will be as follows:



- ◆ A percentage of used disk space of the source volume based on the **Initial Cache Size** property if the form of measurement selected is **Cache Size in % of Disk Space**.
- ◆ The value in megabytes specified in the **Initial Cache Size** property if the form of measurement selected is **Cache Size in MB**.

If the cache file is not placed in the same volume that is being snapped, the **Initial Cache Size** property is ignored.

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**.

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 **Initial Cache Size** and **Maximum Cache Size** properties to values that best fit your client's installation. The recommended setting for the **Initial 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 **Initial 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.

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.

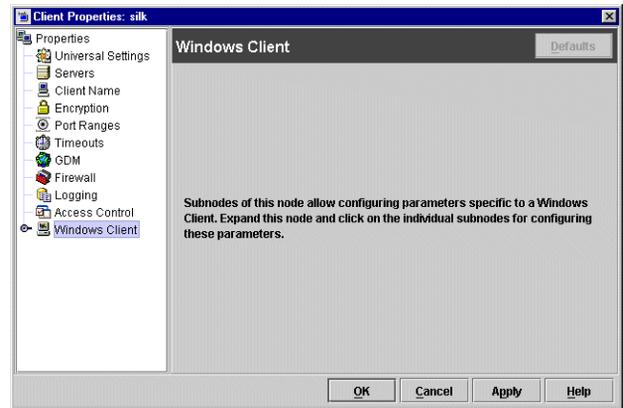
Using VSP to back up active databases without using a formal API presents risk. Customers should contact the database supplier to ensure support of database backups using point-in-time technology. Also, significant back up and restore testing should be performed to assure database availability and reliability.

Windows Client Properties

The **Windows Client** properties define NetBackup properties for Microsoft Windows clients.

Windows Client properties include:

- ◆ “Client Settings (Windows) Properties” on page 326
- ◆ “Exclude Lists Properties” on page 335
- ◆ “VSP (Volume Snapshot Provider) Properties” on page 396
- ◆ “Network Properties” on page 374
- ◆ “Lotus Notes Properties” on page 369
- ◆ “Exchange Properties” on page 334



Windows Open File Backup Properties

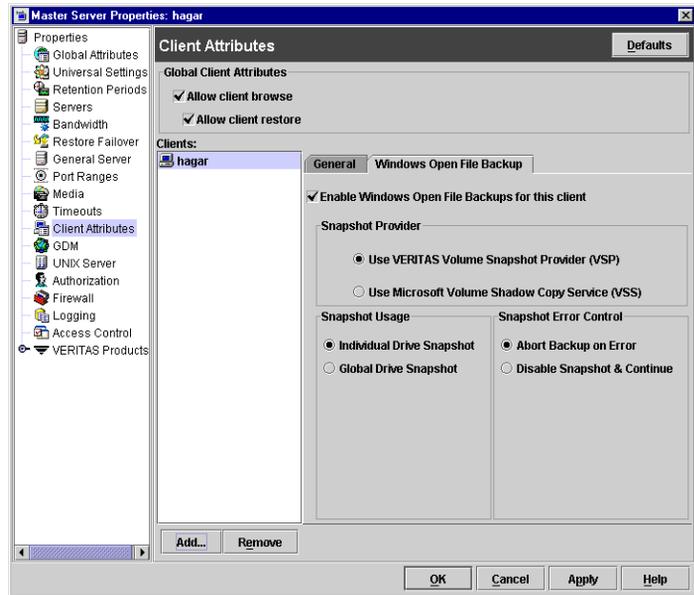
Windows Open File Backup properties apply to selected Windows master servers. The properties appear as a tab on the dialog.

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 396.)

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.





This chapter contains topics related to the administration and management of NetBackup.

- ◆ “Powering Down and Rebooting NetBackup Servers” on page 412
- ◆ “Managing Daemons” on page 413
- ◆ “Administering NetBackup Licenses” on page 416
- ◆ “Using the NetBackup License Utility to Administer Licenses” on page 419
- ◆ “Administering a Remote Master Server” on page 420
- ◆ “Using the NetBackup-Java Windows Display Console” on page 427
- ◆ “Managing Client Restores” on page 431
- ◆ “Goodies Scripts” on page 450
- ◆ “Server Independent Restores” on page 450
- ◆ “Configuring NetBackup Ports” on page 458
- ◆ “Load Balancing” on page 485
- ◆ “Using NetBackup with Storage Migrator” on page 486
- ◆ “Configuring the NetBackup-Java Console” on page 489
- ◆ “NetBackup-Java Performance Improvement Hints” on page 501
- ◆ “Administrator’s Quick Reference” on page 505



Powering Down and Rebooting NetBackup Servers

When closing down and restarting NetBackup servers, use the recommended procedures.

▼ To power down a server

1. Look in the NetBackup Administration Console or use the command line to see that no backups or restores are running:

- ◆ In the NetBackup Administration Console, click **Activity Monitor**, then select the Jobs tab to view jobs currently running.
- ◆ From the command line, run `bpps`, which displays active NetBackup processes, for example, `bpsched` for backup processes and `bprd` for restore processes:

```
/usr/opensv/netbackup/bin/bpps
```

(For more information, see “Displaying Active Processes with `bpps`” on page 413.)

2. Use the NetBackup Administration Console or the command line to stop the NetBackup request daemon:

- ◆ 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.

Note During installation, you should have installed the NetBackup startup and shutdown scripts in the appropriate `init.d` and `rc` directories. The scripts run automatically during system shutdown and system startup.

3. Run system shutdown command.
4. Power down the server.

▼ 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
```

Managing Daemons

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 itself for more information.

To display both NetBackup and Media Manager options, run:

```
/usr/opensv/netbackup/bin/bpps -a
```



Starting and Stopping NetBackup and Media Manager Daemons

The NetBackup Request Manager daemon, `bprd`, starts the scheduler and the NetBackup Database Manager, `bpdbm`, in addition to controlling other functions.

To enable `bprd` debug logging, create the `/usr/opensv/netbackup/logs/bprd` directory before starting `bprd`.

The Media Manager device daemon, `ltid`, starts the Volume Manager daemon, `vmd`, and the Automatic Volume Recognition daemon, `avrd`.

Starting NetBackup and Media Manager Daemons

Before the daemons are started, the networks and network daemons must be fully operational.

▼ To start NetBackup and Media Manager

To start NetBackup and Media Manager, run:

```
/usr/opensv/netbackup/bin/goodies/netbackup start
```

This command starts `ltid`, `vmd`, `avrd`, `bprd`, `bpdbm` and `visd`, if applicable.

Stopping NetBackup and Media Manager Daemons

To stop additional backup and restore activity and to allow current activity to gracefully end, stop `bprd`:

▼ To stop bprd

To stop `bprd`, run:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

If the daemon has started any activities, this command allows the activities to complete. With `bprd` stopped, NetBackup cannot perform any backup, archive, or restore operations. Stopping `bprd` does not stop `bpdbm`.

▼ To stop all daemons

To stop all daemons, run:

```
/usr/opensv/netbackup/bin/goodies/bp.kill_all
```

This script is intended to stop all daemons when no backup or restore is in progress.

Starting and Stopping bpdbm

The NetBackup database daemon, `bpdbm`, must be running during all administrative operations. Normally, this daemon is started by the request daemon, `bprd`.

▼ To start bpdbm separately

To start `bpdbm` separately, run:

```
/usr/opensv/netbackup/bin/initbpdbm
```

▼ To stop bpdbm

To stop `bpdbm`, run:

```
bpdbm -terminate
```

For more information, see the `bpdbm(1M)` man page or NetBackup Commands for UNIX.



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:

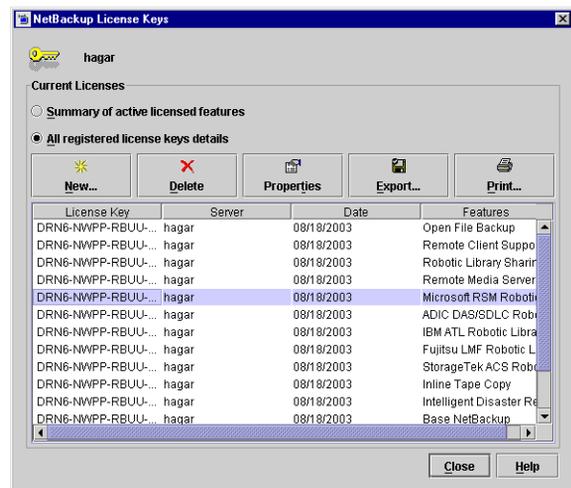
Click **File > Change Server**, then select another server. Click **Help > License Keys**.

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.

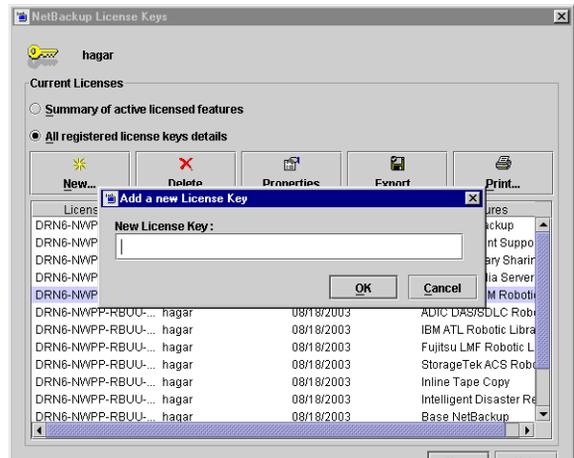


3. From the NetBackup License Keys dialog, you can:
 - ◆ Add a new license
 - ◆ Delete a license
 - ◆ View the properties of one license
 - ◆ Export the license listing

▼ To add a new license key

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 delete a license key

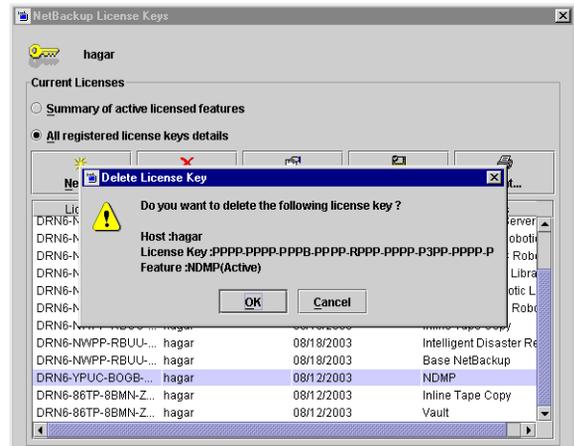
1. Select the license key you wish to delete from the license key list. If the key has more than one feature, all the features are listed in the dialog.
2. In the NetBackup License Keys dialog, click **Delete**. A confirmation dialog appears.



3. Click **Yes** to delete all the features associated with the key. The license key cannot be restored.

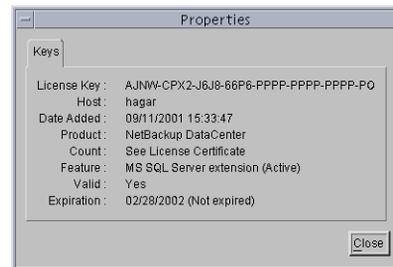
If the key appears more than once in the list, deleting one instance also deletes all other instances of the key from the list.

Note After deleting the license keys, all the NetBackup utilities including NetBackup-Java Administration Console should be restarted.



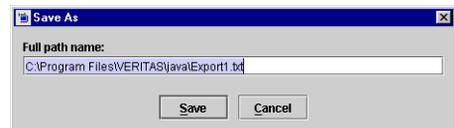
▼ **To view the properties of one license key**

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



▼ **To export the 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/opensv/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:

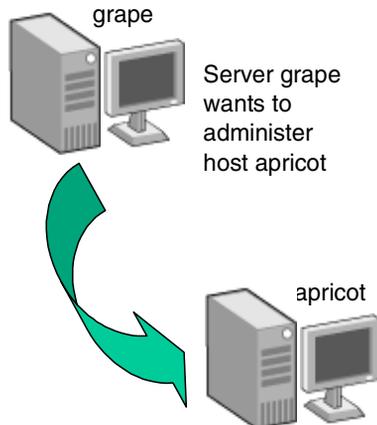
- ◆ First, make the remote server accessible to the local server. See the following section, “Adding a NetBackup Server to a Server List” on page 420
- ◆ Second, indicate the remote server you want to administer. See “Choosing a Remote Server to Administer” on page 424.

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.

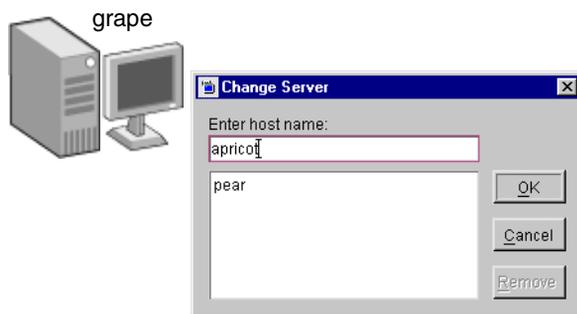
If you are logging into a remote master server through the login dialog, it is not necessary for the name of the local host to appear in the server list of the remote server. This method of logging on to a remote host is explained in “To indicate a remote system upon login” on page 425.

For example, assume UNIX server *grape* wants to remotely administer UNIX 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*.

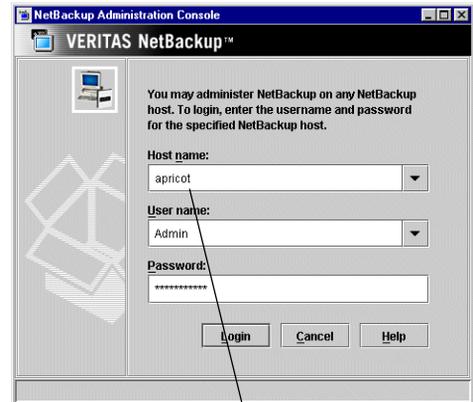
To add *grape* to the server list of *apricot*, follow the steps in “To add a server to a UNIX server list.” For other reasons why a remote server may be inaccessible, see “If You Cannot Access a Remote Server” on page 426.



▼ 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 Administration Console (jnbSA) on the local server (*grape*). Indicate destination host *apricot* on the login dialog. The *jnbSA* command is described in the *NetBackup Commands for UNIX* guide.
- ◆ Start the Windows Display Console on a Windows machine. Indicate destination host *apricot* on the login dialog.
- ◆ Physically go to the destination host (*apricot*) and start *jnbSA*. Indicate *apricot* on the login dialog.



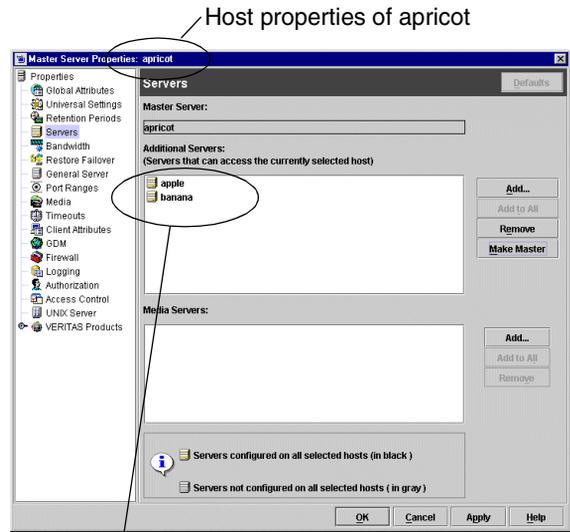
Log in to *apricot* from *grape* (provided the user name has sufficient privileges), or log in at *apricot*

2. In the NetBackup Administration Console, expand **Host Properties** > **Master Servers**.



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.

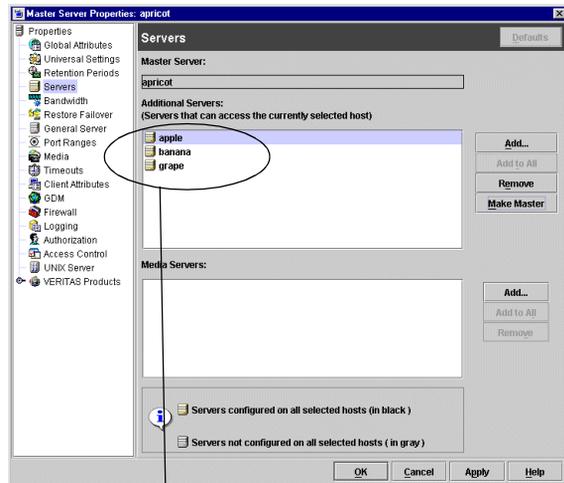


Currently, apricot allows remote access by two additional servers: apple and banana

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 these machines."
5. To add a server to the server list, type the server name in the field labeled **Add to All Lists**.
6. Click the + button next to the **Add to All Lists** field. The server name appears in the server list.
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.



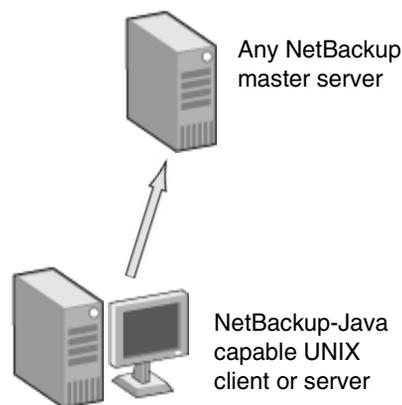
Choosing a Remote Server to Administer

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 login 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/obj/obj_rpm/java/jnbSA`
2. In the NetBackup Administration Console login 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 login

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
```

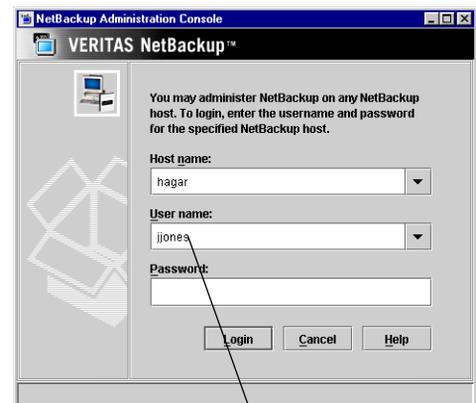
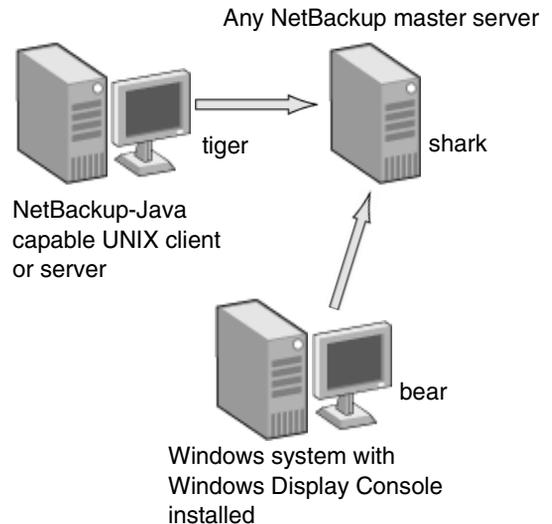
- ◆ To start the console on a Windows system named *bear*, from the Windows desktop, select **Start > Programs > VERITAS NetBackup > NetBackup-Java** on *bear*. (The system must have the Windows Display Console installed.)

The login screen appears.

3. In the **Hostname** 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 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.



Type in the name of the remote server you'd like to administer



Administering through 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 Windows Display Console is installed.
- ◆ On a UNIX client if the client is NetBackup-Java capable.

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 420. 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 (`bpird`) on the remote server if it is UNIX.
(See “Managing Daemons” on page 413 for information on starting and stopping daemons.)
- ◆ 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

Authorizing NetBackup-Java Users on Windows

To use the NetBackup-Java Windows Display Console, you must first log into 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 into 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 login, users provide a user name and password that is valid on the computer specified in the **NetBackup host** field of the login 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 "Restricting Access on Windows" on page 428.



Restricting Access on Windows

To restrict access to one or more of the NetBackup-Java applications, create the following file on the Windows system:

```
nbackup_install_path\java\auth.conf
```

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
```

The `auth.conf` file possesses the following characteristics:

- ◆ The first field of each entry is the user name that is granted access to the rights specified by that entry. 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.

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.

As in the example, any entries that designate specific user names must precede a line that contains an asterisk in the username field.

- ◆ 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 administrator-related capabilities. To restrict use to specific applications, see “Authorizing Users for Specific Applications” on page 429.
 - ◆ The `JBP` keyword specifies what the user can do with the Backup, Archive, and Restore application. `JBP=ALL` allows access to all Backup, Archive, and Restore capabilities, including those for administration. To allow only a subset of those capabilities, see “Authorizing Users for Specific Applications” on page 429.
 - ◆ 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 third line of this example has an asterisk in the first field, which means that NetBackup-Java validates any user name for access to Backup, Archive, and Restore.

- ◆ `JBP=ENDUSER+BU+ARC` allows end users to only back up, archive and restore files.

In the example above, only a user logged in as `mydomain\Administrator` or `mydomain\joe` could administer NetBackup. All other users would have access to only Backup, Archive, and Restore.

Note To use the NetBackup-Java administrator application on Windows (that is, any application except Backup, Archive, and Restore), a user must also be a member of the administrator group in the host computer's domain.

Authorizing Users for Specific Applications

To authorize users for a subset of the NetBackup-Java administrator applications, use the `ADMIN` keyword with identifiers in the `auth.conf` file. For example, to give a user (`joe`) access to device and activity monitoring only, use the following format:

```
mydomain\joe ADMIN=DM+AM
```

auth.conf ADMIN Identifiers for Administrator Applications

AM	Activity Monitor
BPM	Backup Policy Manager
CAT	Catalog
DM	Device Monitor
HPD	Host Properties
JBP	Backup, Archive, and Restore
MM	Media Manager
REP	Reports
SUM	Storage Unit Manager
VLT	Vault Management

To authorize users for a subset of Backup, Archive, and Restore capabilities, use the `JBP` keyword with identifiers in the `auth.conf` file. For example, to give a user (`jan`) permission to restore but not back up or archive files, use the following format:



```
mydomain\jan ADMIN=JBP JBP=ENDUSER
```

auth.conf JBP Identifiers for Backup, Archive, and Restore Capabilities

ENDUSER	Authorizes restore capabilities; from true image, archive or regular backups plus alternate client restores.
BU	Authorizes backup tasks.
ARC	Authorizes archive tasks (BU capability required).
RAWPART	Authorizes raw partition restores.
ALL	<p>Authorizes all of the above, including restoring to a different client from the one logged into (that is, server-directed restores). This normally requires it be run from the Administrator account or an account that is a member of the Administrator group. Alternate-client restores also require changes to the <code>altnames</code> file on the NetBackup master server.</p> <p>When authorized for ALL, the user can view a list of media IDs required for the files marked for restore through the Preview Media Required button at the bottom of the Restore Files tab.</p>



Managing Client Restores

The topics in this section concern aspects of managing restores for NetBackup clients.

- ◆ “Server-Directed Restores” on page 431
- ◆ “Allowing Redirected Restores” on page 431
- ◆ “Restoring Files and Access Control Lists” on page 440
- ◆ “Improving Search Times by Creating an Image List” on page 445
- ◆ “Checkpoint Restart for Restore Jobs” on page 446
- ◆ “Set Original atime for Files During Restores” on page 446
- ◆ “Restoring System State” on page 447

A related topic is “Rules for Using Host Names in NetBackup” on page 234 in *NetBackup System Administrator’s Guide, Volume II*. Incorrectly specified host names are often a factor in file restore problems.

Server-Directed Restores

An administrator can use the Backup, Archive, and Restore client interface on the NetBackup master server to direct restores to any client, providing NetBackup on the client is configured to permit them. See the *NetBackup User’s Guide for UNIX* or *NetBackup User’s Guide for Windows* for more information.

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.

Allowing Redirected Restores

The Backup, Archive, and Restore client interface contains options for restoring files that were backed up by other clients. The operation is called a *redirected restore*.

A client can restore files belonging to other clients only with the necessary configuration on the NetBackup master server. Create the following directory on the master server:

```
/usr/opensv/netbackup/db/altnames
```



Add files to it as explained in this section. To undo the changes, remove the `altnames` directory and its files.

Caution The `/usr/openv/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.

How NetBackup Enforces Restore Restrictions

By default, NetBackup permits restores only by the client that backs up the files. NetBackup enforces this restriction by ensuring that the name specified by the NetBackup client name setting on the requesting client matches the peer name used in the connection to the NetBackup server.

The NetBackup client name is normally the client's short host name, such as `mercury` rather than a longer form such as `mercury.null.com`.

- ◆ On Microsoft Windows clients (includes NetWare NonTarget), specify the client name in the Specify NetBackup Machines and Policy Type dialog. To display this dialog, start the Backup, Archive, and Restore client interface and select click **Actions > Specify NetBackup Machines and Policy Type**.
- ◆ On UNIX clients, specify the client name in the user interface.
- ◆ On NetWare target clients, specify the client name in the `bp.ini` file.

Peer name is the name that the client uses when it connects to the NetBackup server during the file restore request. Unless clients share an IP address due to the use of a gateway and token ring combination, or have 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.

Allowing All Clients to Perform Redirected Restores

The administrator can allow all clients to restore backups belonging to other clients by creating the following empty file on the NetBackup master server:

```
/usr/openv/netbackup/db/altnames/No.Restrictions
```

When this file exists on the master server, clients can access backups belonging to other clients if the NetBackup client name setting on the requesting client matches the name of the client for which the backup was created. The peer name of the requesting client does not have to match the NetBackup client name setting.

Example

Assume UNIX client `freddie` wants to restore a file that was backed up by client `oscar`:

1. The administrator creates the following file on the NetBackup master server:

```
/usr/opensv/netbackup/db/altnames/No.Restrictions
```
2. The user starts the Backup, Archive, and Restore application (jbpSA) and specifies **freddie** in the login dialog.
3. The user changes the NetBackup source client name setting in the Backup, Archive, and Restore user interface to **oscar**.
4. Client freddie restores the file backed up by client oscar.

Allowing a Single Client to Perform Redirected Restores

The administrator can give a single client permission to restore backups belonging to other clients by creating an empty file on the NetBackup master server:

```
/usr/opensv/netbackup/db/altnames/peername
```

Where *peername* is the client that is to possess restore privileges.

In this case, the client named *peername* can access files backed up by another client if the NetBackup client name setting on the client named *peername* matches the name of the other client.

Example

Assume UNIX client freddie wants to restore files that were backed up by client oscar:

1. The administrator creates the following file on the NetBackup master server:

```
/usr/opensv/netbackup/db/altnames/freddie
```
2. The user starts the Backup, Archive, and Restore application (jbpSA) and specifies **freddie** in the login dialog.
3. The user changes the NetBackup source client name setting in the Backup, Archive, and Restore user interface to **oscar**.
4. Client freddie restores the files backed up by client oscar.

Allowing Redirected Restores of a Specific Client's Files

The administrator can give a single client permission to restore backups belonging to specific other clients. First, create the following file on the NetBackup master server:

```
/usr/opensv/netbackup/db/altnames/peername
```



Then, add the client names to *peername*.

The client(s) named *peername* can restore files backed up by another client if:

- ◆ The name of the other client appears in the *peername* file, and
- ◆ The NetBackup client name setting on the client named *peername* is changed to match the client name in the *peername* file.

Example

Assume UNIX client *freddie* wants to restore files backed up by client *oscar*:

1. The administrator creates the following file on the NetBackup master server:

```
/usr/openv/netbackup/db/altnames/freddie
```
2. The user starts the Backup, Archive, and Restore application (*jbpsA*) and specifies **freddie** in the login dialog.
3. The administrator enters the name *oscar* on a separate line in the *freddie* file.
4. The user changes the NetBackup source client name setting in the Backup, Archive, and Restore client interface to *oscar*.
5. Client *freddie* restores the files backed up by client *oscar*.

Redirected Restore Examples

This section provides examples of configuring NetBackup to allow clients to restore files that were backed up by other clients. These example methods can be required when a client connects through a gateway or has multiple Ethernet connections. In all cases, the client you are restoring to must have an *image-catalog* directory on the master server in

```
/usr/openv/netbackup/db/images/client_name
```

or be a member of an existing NetBackup policy.

Caution Not all file system types on all machines support the same features and you may run into 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 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:

- ◆ *restore_to_client* is the client that is requesting the restore.
- ◆ *backed_up_client* is the client that created the backups that the requesting client wants to restore.

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

Assume you must restore files to *restore_to_client* that were backed up from *backed_up_client*. The *restore_to_client* and *backed_up_client* 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 either:
 - ◆ Edit `/usr/opensv/netbackup/db/altnames/restore_to_client` so it includes the name of *backed_up_client*. Or,
 - ◆ Run the `touch` command on the following file:
`/usr/opensv/netbackup/db/altnames/No.Restrictions`
2. Log in on *restore_to_client* and change the NetBackup client name on the client to *backed_up_client*.
3. Restore the file.
4. Undo the changes made on the server and client.

Example 2

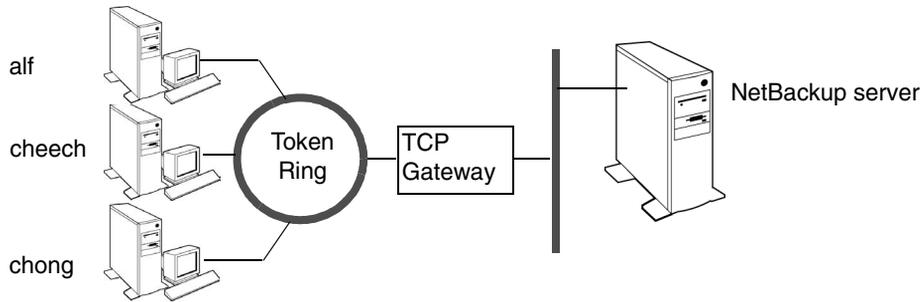
This example explains how the `altnames` file can provide restore capabilities to clients that do not use their own host name when connecting to the NetBackup server.

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 network, restore requests coming from clients alf, cheech, and chong 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.

To correct this situation proceed as follows:

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:

```
01/29/03 08:25:03 bpserver - request from invalid
server or client bilbo.dvlp.null.com
```

In this example, the peer name is `bilbo.dvlp.null.com`.

2. After determining the peer name, create the following file on the NetBackup master server:

```
/usr/opensv/netbackup/db/altnames/peername
```

In our example, the file is:

```
/usr/opensv/netbackup/db/altnames/bilbo.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/bilbo.dvlp.null.com` empty, clients `alf`, `cheech`, and `chong` can all access the backups corresponding to their NetBackup client name setting. (See “Allowing a Single Client to Perform Redirected Restores” on page 433.)

If you add the names `cheech` and `chong` to the file, you give these two clients access to NetBackup file restores, but exclude `alf`. (See “Allowing Redirected Restores of a Specific Client’s Files” on page 433.)

Note that this example requires no changes on the clients.

4. Restore the files.



Example 3

If you cannot restore files with a redirected client restore using the `altnames` file, perform troubleshooting using 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 `restore_to_client`, attempt the file restore.
5. On the NetBackup server, identify the peer-name connection used by `restore_to_client`.

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 lets any `restore_to_client` access any `backed_up_client` backups by changing its NetBackup client name setting to specify the `backed_up_client` client.

- ◆ Run the `touch` command on the `/usr/opensv/netbackup/db/altnames/peername` file. This lets `restore_to_client` access any `backed_up_client` backups by changing its NetBackup client name setting to specify the `backed_up_client` client.
- ◆ Add the `backed_up_client` name to the `/usr/opensv/netbackup/db/altnames/peername` file. This lets `restore_to_client` access only the backups created on `backed_up_client`.

7. On *restore_to_client*, change the NetBackup client name setting in the user interface to match what is specified on *backed_up_client*.
8. Restore the files from *restore_to_client*.
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 its contents.
10. To undo the changes:
 - ◆ Delete `/usr/opensv/netbackup/db/altnames/peer.or.hostname` (if you created it)
 - ◆ Delete `/usr/opensv/netbackup/db/altnames/No.Restrictions` (if you created it)
 - ◆ On *restore_to_client*, 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 "Possible Files Generated By `tar`" on page 244 in *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 NT 4.0 or later, including Windows 2000, Windows XP, and Windows Server 2003. The option is disabled on UNIX clients.

▼ To restore files without restoring ACLs

1. Log into 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 topics:

- ◆ “Setting the List and Restore Permissions” on page 442
- ◆ “Examples” on page 444

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 “Dynamic Host Name and IP Addressing” on page 113 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 a client entry, run:

```
bpclient -add -client client_name -current_host host_name
```

Where:



- ◆ `-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
```

You can also delete and list client entries:

- ◆ To delete a client entry, run: `bpclient -delete -client client_name`
- ◆ To list a client entry, run: `bpclient -L -client client_name`
- ◆ 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.



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



`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/opencv/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 configurable settings that impact Checkpoint Restart for restore jobs:

- ◆ **Move Restore Job from Incomplete State to Done State:** This is a host property of the master server. (See “Global Attributes Properties” on page 362.)
- ◆ **Restore Retries:** A host property of each client. (See “Universal Settings Properties” on page 389.)

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 “Menu Bar” on page 279.

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.



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

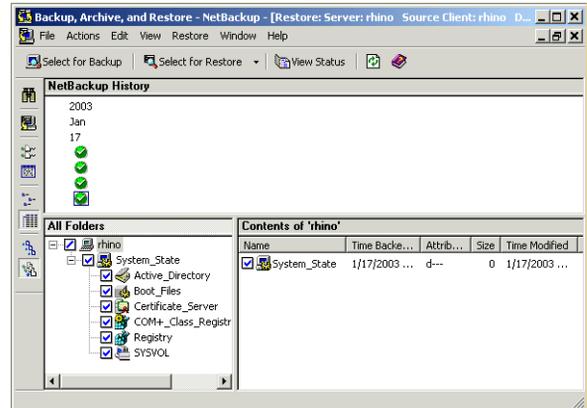
Before restoring the System State, please read the following notes carefully.

- ◆ 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

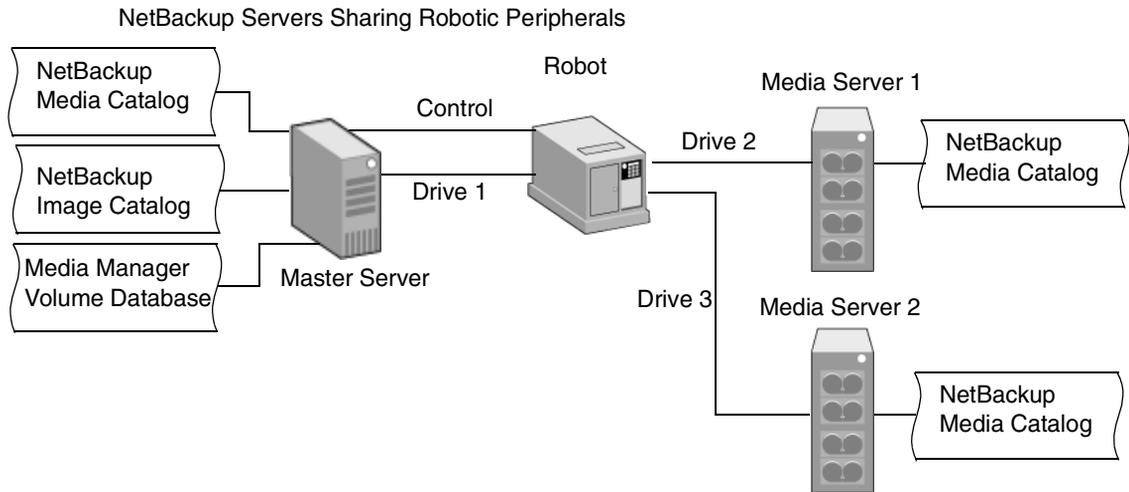
This section explains how to restore files by using a NetBackup server other than the one that was used to write the backup. This is called a server independent restore and allows easier access to data for restores in master and media server clusters and also provides better failover and disaster recovery capabilities.

NetBackup has a master and media server architecture that allows storage devices to be located on multiple servers (can be 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 media catalog on the master or media server that was used during the backup.

Due to the existence of the media catalog on each server where backups are written, 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 457.)

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 volume database.

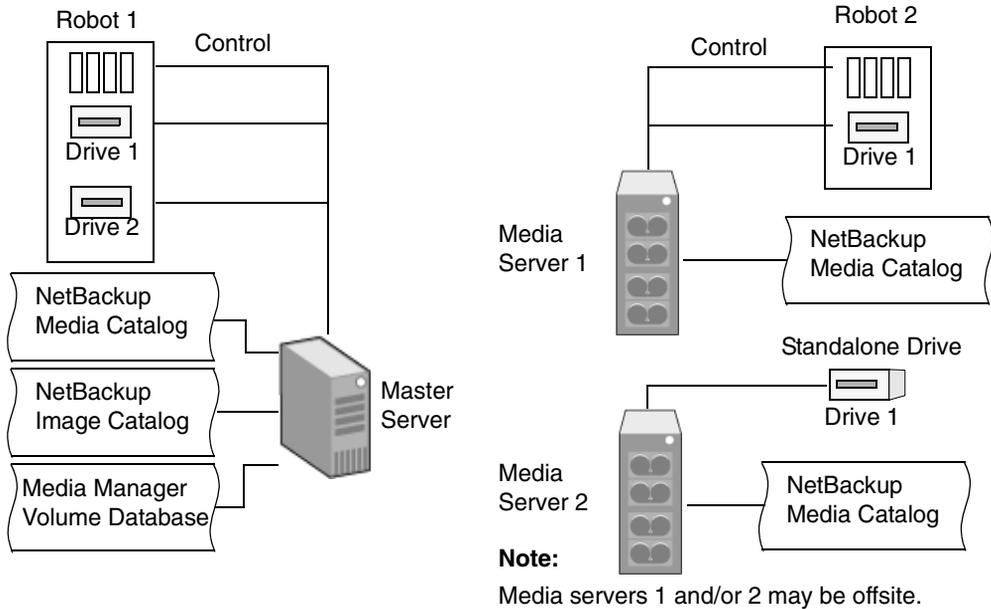


In the above figure, the following assumptions are made:

- ◆ A single, shared Media Manager volume 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



In the above figure, the following assumptions are made:

- ◆ The media is made physically accessible through an available NetBackup server and the Media Manager volume database is updated to reflect this move.
- ◆ A single, shared Media Manager volume 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 453
- ◆ “Method 2: Overriding the Original Server” on page 454
- ◆ “Method 3: Automatic Failover to Alternate Server” on page 456

Method 1: Modifying the NetBackup Catalogs

This method changes the contents of NetBackup catalogs and thus requires administrator intervention. It is best to 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.

If the Server that Originally Wrote the Media Is Available

1. If necessary, physically move the media. Then, update the Media Manager volume 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 commands, which can be run from any one of the NetBackup servers:

UNIX NetBackup server (as root user):

```
cd /usr/opensv/netbackup/bin/admincmd
```



```
bpmedia -movedb -m media_id -newserver hostname
      -oldserver hostname
```

(the `admincmd` command above must be on one line)

Windows NetBackup server (as administrator, from the MSDOS prompt):

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -m media_id
      -newserver hostname -oldserver hostname
```

(the `admincmd` command above must be on one line)

If the Host that Originally Wrote the Media Is Not Available

1. If necessary, physically move the media and update the Media Manager volume 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:

On a UNIX NetBackup server (as root user):

```
cd /usr/opensv/netbackup/bin/admincmd
bpimage -id media_id -newserver hostname
      -oldserver hostname
```

(the `admincmd` command above must be on one line)

On a Windows NetBackup server (as administrator, from the MSDOS prompt):

```
cd install_path\NetBackup\bin\admincmd
bpimage.exe -id media_id -newserver hostname
      -oldserver hostname
```

(the `admincmd` command above must be on one line)

To revert to the original configuration for future restores, perform the same steps again, switching the host names on the commands.

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.

This method requires administrator intervention in the General Server host properties dialog (see “General Server Properties” on page 359.):



Place a check in the **Use Media Host Override for Restores** check box.

- ◆ Specify the original media server in the **Server that Performed Original Backups** field.
- ◆ Specify the media server that should restore the files in the **Server that Should Perform the Restore** field.

Some 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 override the original server for a restore:

1. If necessary, physically move the media and update the Media Manager volume database to reflect the move.
2. Modify the NetBackup configuration on the master server:

On a UNIX NetBackup server:

As root user, add the following entry to the `/usr/opensv/netbackup/bp.conf` file:

```
FORCE_RESTORE_MEDIA_SERVER = fromhost tohost
```

where *fromhost* is the server that wrote the original backup and *tohost* is the server to use for the restore.

On a Windows NetBackup server, this is set through the NetBackup administration interface.

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*.

To revert to the original configuration for future restores, simply delete the changes made in step 2 above.



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 378.)

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 ok but `bptm` is having problems (for example, if `vmd` is down or `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 the automatic failover to an alternate server:

1. Modify the NetBackup configuration on the master server:

On a UNIX NetBackup server:

As root user, 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.

On a Windows NetBackup server, this is specified through the NetBackup Administration interface on the master server.

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
```

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 service (`bpcd`) is 13782. These ports are specified in a system configuration file.

On UNIX systems: `/etc/services`

On Windows systems:

`%systemroot%\system32\drivers\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.

- ◆ Dynamically-allocated ports are assigned, as needed, from ranges that you can specify 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 nonreserved ports.

These settings are useful in environments that use firewalls for security and are explained later in this section.

NetBackup and Media Manager Ports

Daemon/ Service	Port	Description
ACSD	13702	The Automated Cartridge System (ACS) daemon is a robotic daemons.
ARDBD	3306	The database server used for the NetBackup Advanced Reporter database.
BPCD	13782	The NetBackup Client daemon. On UNIX clients, BPCD can only be run in standalone mode. On Windows, BPCD always runs under the supervision of BPINETD.EXE. There is a NetBackup-specific configuration parameter for BPCD. 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.
BPDBM	13721	The NetBackup database manager daemon.
BPJAVA-MS VC	13722	The NetBackup-Java application server authentication service program.
BPJOB	13723	The NetBackup Jobs Database Management daemon.
BPRD	13720	The NetBackup Request daemon. On Windows, there is a NetBackup specific configuration parameter for BPRD. 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.
LMFCD	13718	The Library Management Facility (LMF) control daemon is a robotic daemons.
MIGRD	13699	The VSM request daemon (database request management) for Storage Migrator. MIGRD handles communication for VSM-Java and commands.
NBDBD	13784	The Persistent Storage daemon. NBDBD only runs if the GDM server or GDM-managed server license is added to the NetBackup server. visd requires nbdbd to be running before it will start.
NDMP	10000	The acronym for Network Data Management Protocol. NDMP Servers are designed to adhere to this protocol and listen on port 10000 (by default) for NDMP Clients to connect to them.
ODLD	13706	The Optical Disk Library (ODL) daemon is a robotic daemon.
RSMD	13719	The Removable Storage Manager (RSM) daemon. This process is one of the robotic daemons.



TL4D	13713	The Tape Library 4MM (TL4) daemon is a robotic daemon.
TL8CD	13705	The Tape Library 8MM (TL8) control daemon is a robotic daemon.
TLDCD	13711	The Tape Library DLT (TLD) control daemon is a robotic daemon.
TLHCD	13717	The Tape Library Half-inch (TLH) control daemon is a robotic daemon.
TLMD	13716	The Tape Library Multimedia (TLM) daemon is a robotic daemons.
TS8D	13709	The Tape Stacker 8MM (TS8) daemon is a robotic daemon.
TSDD	13714	The Tape Stacker DLT (TSD) daemon is a robotic daemon.
TSHD	13715	The Tape Stacker Half-inch (TSH) daemon is a robotic daemon.
VISD	9284	<p>The VERITAS Information Server Daemon. visd requires nbdbd to be running before it will start.</p> <p>Note: Do not use port number 65535 for either visd or the Dashboards, as this port number causes problems for both visd and the Dashboards.</p>
VMD	13701	<p>The Media Manager volume daemon. VMD logs an error message using syslogd 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>
VNETD	13724	The VERITAS Network daemon.
VOPIED	13783	<p>The daemon that provides VERITAS One-time Password user authentication. VOPIED is used to authenticate user names, hosts names, and group/domain names.</p> <p>On UNIX clients, VOPIED can only be run in standalone mode.</p> <p>On Windows, VOPIED always runs under the supervision of BPINETD.EXE.</p>
web server	8885	<p>A Web Server is used with NetBackup Advanced Reporter to provide a database connection and render reports.</p> <p>Port 8885 is used for NetBackup for UNIX. The same port is used for NetBackup version 5.0 or later for Windows.</p>



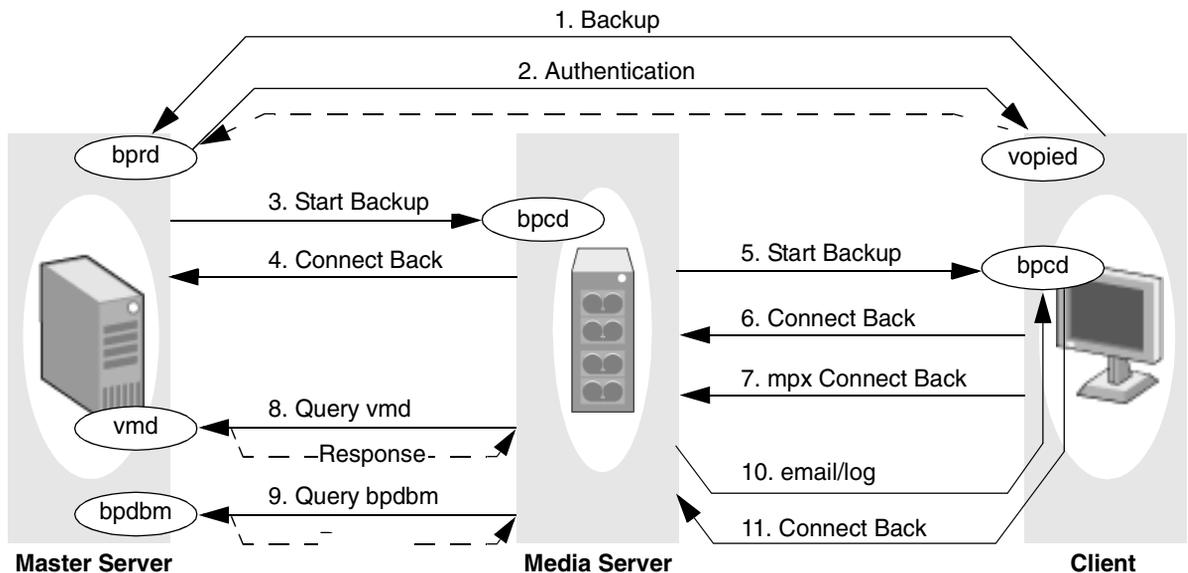
Server and Client Connections: General Case

This section explains the ports that NetBackup uses for connections between clients and servers when using the standard interfaces. The connections for alternative interfaces, such as the Windows Display Console are explained later in the chapter.

Backups

A backup can be started by either the scheduler on the master server or a request from a client. The following figure shows the connections that occur for a client-requested backup. A scheduled backup works the same way, except there is no client request.

Backup Port Connections



The table “Connections for Backups” describes each connection and defines the ports that NetBackup uses:

- ◆ For registered ports, the table shows the port number (for example, 13720).
- ◆ For dynamically-allocated ports, the table indicates *Reserved* (port numbers less than 1024) or *Nonreserved* (port numbers greater than 1024). Some `bpcd` connections are shown as *Reserved* or *NonReserved*, depending on the `allow non reserved port` setting.

In addition to the ports in the “Backup Port Connections” figure and “Connections for Backups” table, the master and media server can have connections to robotic control daemons (`t18cd` and so on). See the `services` file on the computers that share the tape library for those port numbers.



Connections for Backups

Description	Connect From ¹		Connect To ¹	
	Host	Port	Host	Port
1. Backup request to <code>bprd</code> on master server.	Client	Nonreserved	Master Server	13720 (<code>bprd</code>)
2. If NetBackup authentication is being used, and the request is from a nonprivileged user, <code>bprd</code> requests authentication from <code>vopied</code> on the client. Further communication between <code>vopied</code> and <code>bprd</code> is over the same connection.	Master Server	Nonreserved	Client	13783 (<code>vopied</code>)
3. The master server sends a backup request to <code>bpcd</code> on the media server. As part of the <code>bpcd</code> protocol, the master server also sends a port number for connecting back (see next step). The master server then listens on that port.	Master Server	Reserved	Media Server	13782 (<code>bpcd</code>)
4. Connect back to master server. Each backup has its own connect back.	Media Server	Reserved	Master Server	Reserved port specified during request in 3 above.
5. Backup request to <code>bpcd</code> on client. Again, as part of the <code>bpcd</code> protocol, the media server sends a port for connecting back.	Media Server	Reserved or Nonreserved (whichever <code>bpcd</code> on the client is configured to accept)	Client	13782 (<code>bpcd</code>)
6. Connect back to media server. Each backup job has its own connect back.	Client	Reserved or Nonreserved (whichever is used during request in 5 above)	Media Server	Reserved or NonReserved port specified during request in 5 above.



Connections for Backups (continued)

Description	Connect From ¹		Connect To ¹	
	Host	Port	Host	Port
7. If multiplexing (MPX) is used, each backup job requires an additional client-to-server connection.	Client	Reserved or Nonreserved (whichever is used during request in 5 above)	Media Server	Reserved or NonReserved port specified during request in 5 above.
8. Queries to vmd. During the backup, the media server sends queries to vmd on the master server. Responses are over the same connection.	Media Server	Nonreserved	Master Server	13701 (vmd)
9. Queries to bpdbm. During the backup, the media server sends queries to bpdbm on the master server. Responses are over the same connection.	Media Server	Nonreserved	Master Server	13721 (bpdbm)
10. Email notifications or log entries to bpcd on client. These connections also use the bpcd protocol, where the client connects back on a port specified by the server.	Media Server	Reserved or Nonreserved (whichever bpcd on the client is configured to accept)	Client	13782 (bpcd)
11. Connect back to media server.	Client	Reserved or Nonreserved (whichever is used during request in 10 above)	Media Server	Reserved or NonReserved port specified during request in 10 above.

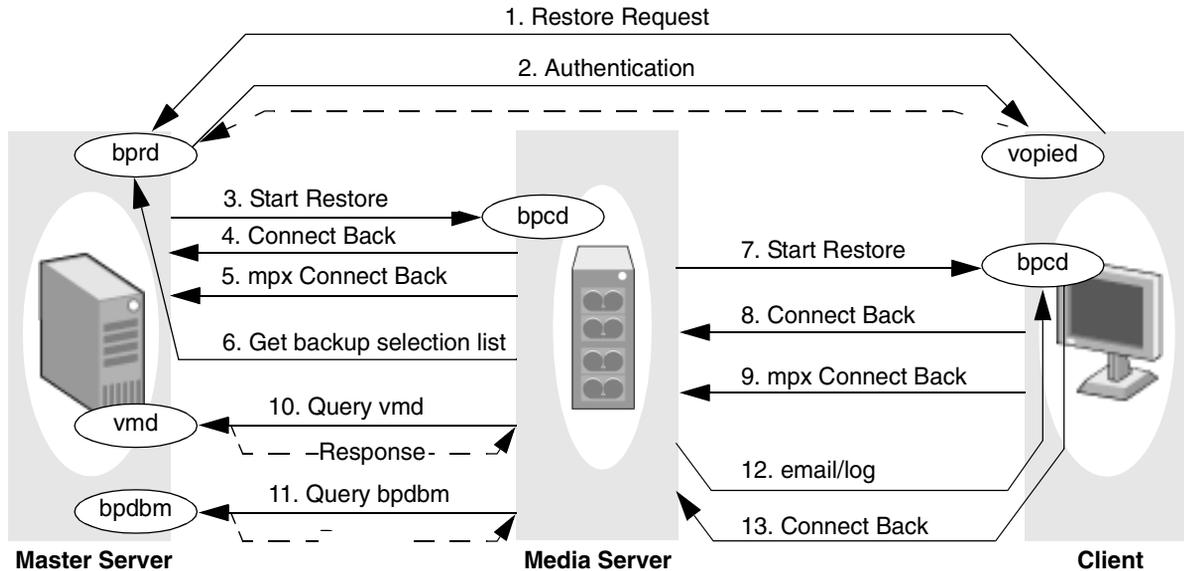
1. For configuration instructions, see "Configuring Ports for Backups and Restores" on page 470.



Restores

A restore can be started by a restore request from a client. The following figure shows the connections that occur for a restore.

Restore Port Connections



The following table, “Connections for Restores,” describes each connection and defines the ports that NetBackup uses:

- ◆ For registered ports, the table shows the port number (for example, 13720).
- ◆ For dynamically-allocated ports, the table indicates *Reserved* or *Nonreserved*. Some `bpcd` connections are shown as *Reserved* or *NonReserved*, depending on the `allow non reserved port` setting.

In addition to the ports in the “Restore Port Connections” figure and “Connections for Restores” table, the master and media server can have connections to robotic control daemons (`t18cd` and so on). See the `services` file on the computers that share the tape library for those port numbers.

Connections for Restores

Description	Connect From ¹		Connect To ¹	
	Host	Port	Host	Port
1. Restore request to <code>bprd</code> on master server.	Client	Nonreserved	Master Server	13720 (<code>bprd</code>)
2. If NetBackup authentication is being used, and the request is from a nonprivileged user, <code>bprd</code> requests authentication from <code>vopied</code> on the client. Further communication between <code>vopied</code> and <code>bprd</code> is over the same connection.	Master Server	Nonreserved	Client	13783 (<code>vopied</code>)
3. The master server sends a restore request to <code>bpcd</code> on the media server. As part of the <code>bpcd</code> protocol, the master server also sends <code>bpcd</code> a port number for connecting back (see next step). The master server then listens on that port.	Master Server	Reserved	Media Server	13782 (<code>bpcd</code>)
4. Connect back to master server. Each restore has its own connect back.	Media Server	Reserved	Master Server	Reserved port specified during request in 3 above.
5. If multiplexing (MPX) is used, each backup job requires an additional connect back to the master server.	Media Server	Reserved	Master Server	Reserved port specified during request in 3 above.
6. Get backup selection list.	Media Server	Reserved	Master Server	Reserved port specified during initial request in 1 above.
7. Restore request to <code>bpcd</code> on client. Again, as part of the <code>bpcd</code> protocol, the media server sends <code>bpcd</code> ports for connecting back.	Media Server	Reserved or Nonreserved (whichever <code>bpcd</code> on the client is configured to accept)	Client	13782 (<code>bpcd</code>)



Connections for Restores (continued)

Description	Connect From ¹		Connect To ¹	
	Host	Port	Host	Port
8. Connect back to media server. Each restore job has its own connect back.	Client	Reserved or Nonreserved (whichever is used during request in 6 above)	Media Server	Reserved or NonReserved port specified during request in 6 above.
9. If multiplexing (MPX) is used, each restore job requires an additional client-to-server connection.	Client	Reserved or Nonreserved (whichever is used during request in 6 above)	Media Server	Reserved or NonReserved port specified during request in 6 above.
10. Queries to <code>vmc</code> . During the backup, the media server sends queries to <code>vmc</code> on the master server. Responses are over the same connection.	Media Server	Nonreserved	Master Server	13701 (<code>vmc</code>)
11. Queries to <code>bpcbm</code> . During the backup, the media server sends queries to <code>bpcbm</code> on the master server. Responses are over the same connection.	Media Server	Nonreserved	Master Server	13721 (<code>bpcbm</code>)
12. Email notifications or log entries to <code>bpcd</code> on client. These connections also use the <code>bpcd</code> protocol, where the client connects back on a port specified by the server.	Media Server	Reserved or Nonreserved (whichever <code>bpcd</code> on the client is configured to accept)	Client	13782 (<code>bpcd</code>)
13. Connect back to media server.	Client	Reserved or Nonreserved (whichever is used during request in 11 above)	Media Server	Reserved or NonReserved port specified during request in 11 above.



Connections for Restores (continued)

Description	Connect From ¹		Connect To ¹	
	Host	Port	Host	Port

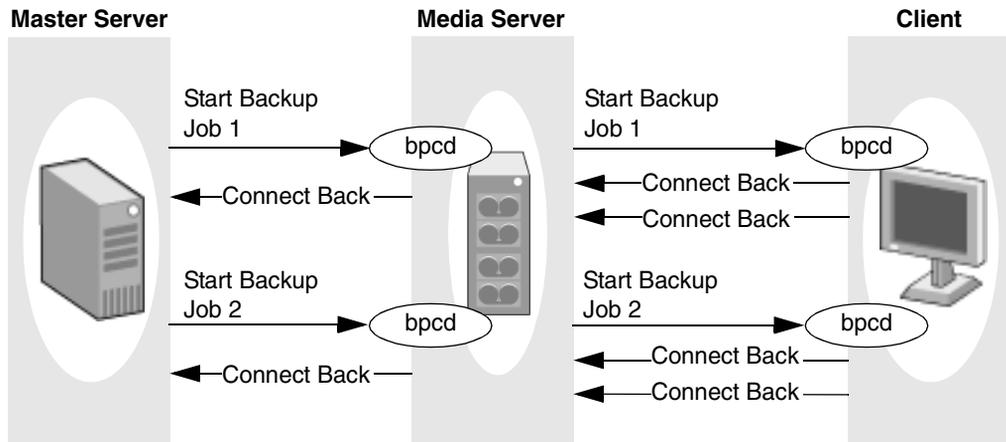
1. For configuration instructions, see “Configuring Ports for Backups and Restores” on page 470.



Multiplexing

As was mentioned in earlier discussions, multiplexing requires an extra connect back for each job. With this exception, the connections are the same as for other backups and restores. The following figure shows the connections when multiplexing results in two backup jobs.

Connections for Multiplexed Backups



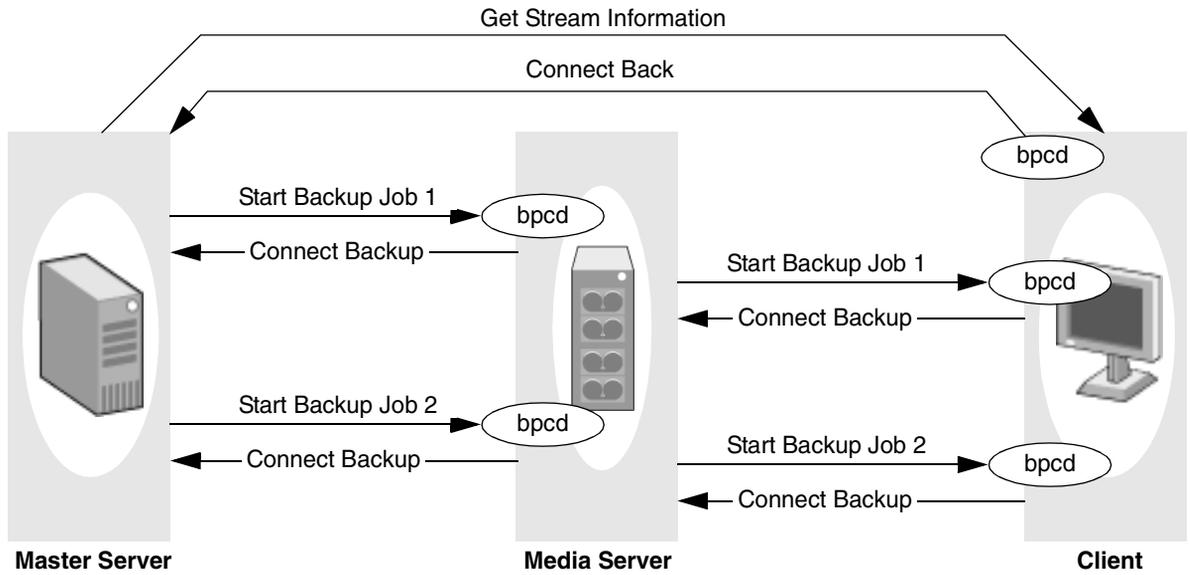
Multiple Data Streams

The connections for multiple jobs (streams) started as a result of using **Allow multiple data streams** are the same as for other backups. That is, NetBackup creates a separate set of connections for each job.

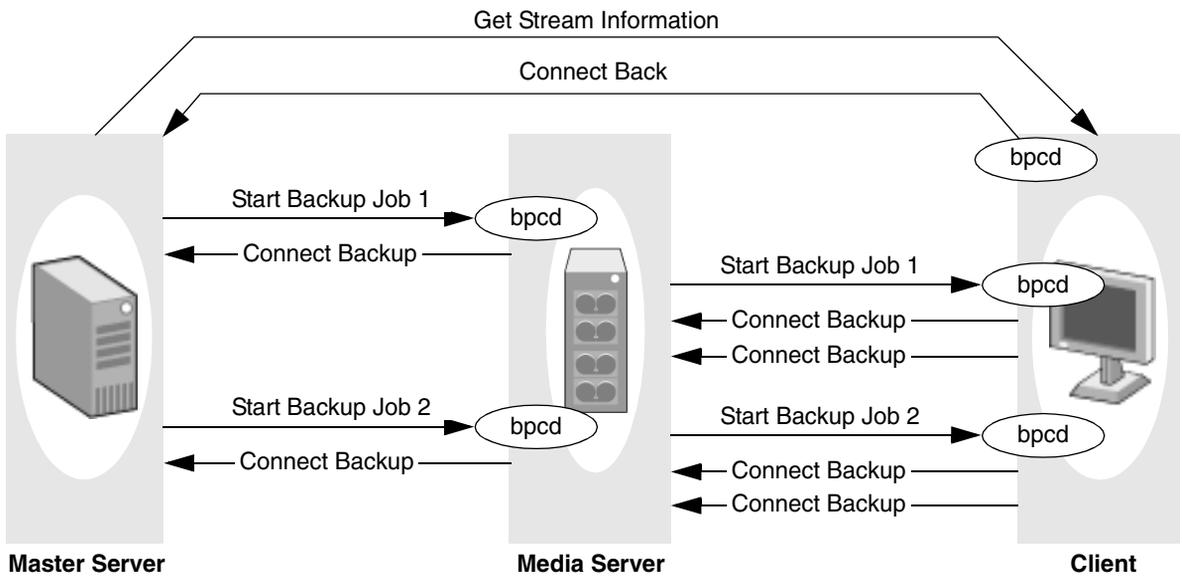
If auto-discover streaming mode is enabled (for example, by using the `ALL_LOCAL_DRIVES` directive), the master server opens an additional connection to `bpcd` on each client in order to obtain the required number of streams. This connection is to either a reserved or nonreserved port, depending on what the client is configured to accept for `bpcd`. The connect back is to the port specified during the `bpcd` connection and also is reserved or nonreserved according to what was used for `bpcd`.

The next two figures show backup connections both with and without multiplexing. Multiplexing results in an additional connect back for each job in the same way as when **Allow multiple data streams** is not used.

Connections for Multiple Data Streams - Without Multiplexing



Connections for Multiple Data Streams - With Multiplexing



Configuring Ports for Backups and Restores

The following explains the NetBackup configuration settings for ports. All settings are in the `/usr/opensv/netbackup/bp.conf` file on the respective host. For more information, see “NetBackup Configuration Options” on page 134 in the *NetBackup System Administrator’s Guide, Volume II*. Registered port numbers (for example, 13782 for `bpcd`) are not configurable with these settings and VERITAS recommends that you do not attempt to change the registered port numbers.

Ports from which NetBackup originates connections to other hosts

These are the *Connect From* ports in tables “Connections for Backups” and “Connections for Restores.” You can set the following on each host:

- ◆ Range of reserved ports from which NetBackup can originate connections. Use `CLIENT_RESERVED_PORT_WINDOW` in the `bp.conf` file.
- ◆ Range of nonreserved ports from which NetBackup can originate connections. Use `CLIENT_PORT_WINDOW` in the `bp.conf` file.
- ◆ Random or nonrandom port selection. By default, NetBackup chooses a port at random from those available in the allowed range. To have NetBackup start at the top of the allowed range and choose the first available port, set `RANDOM_PORTS=NO` in the `bp.conf` file.

Ports where NetBackup listens for connections from other hosts

These are the *Connect To* ports in tables “Connections for Backups” and “Connections for Restores.” You can set the following on each host:

- ◆ `bpcd` to accept connections from nonreserved ports (the default is to not accept these connections). To permit connections from nonreserved ports, add `ALLOW_NON_RESERVED_PORTS` to the `bp.conf` file on that host.

Note For clients, use the `bpclient` command on the master server to specify nonreserved port usage for the client. (See “Connect on Non-reserved Port” on page 319 or “`ALLOW_NON_RESERVED_PORTS`” on page 137 in the *NetBackup System Administrator’s Guide, Volume II*.)

- ◆ Range of reserved ports where NetBackup can listen for connections to this host. Use `SERVER_RESERVED_PORT_WINDOW`.
- ◆ Range of nonreserved ports where NetBackup can listen for connections to this host. Use `SERVER_PORT_WINDOW` in the `bp.conf` file.

- ◆ Random or nonrandom port selection. By default, NetBackup chooses a port at random from those available in the allowed range. To have NetBackup start at the top of the allowed range and choose the first available port, set `RANDOM_PORTS=NO` in the `bp.conf` file.

Important Note on Configuring Port Limitations

Any port limitations configured on a NetBackup host apply to connections with *all* other NetBackup hosts, not just those on the other side of the firewall. *Therefore, leave enough ports available to allow the necessary connections.* The main factors to consider are the number of clients and whether multiplexing is used. If NetBackup runs out of ports, backups and restores cannot occur.

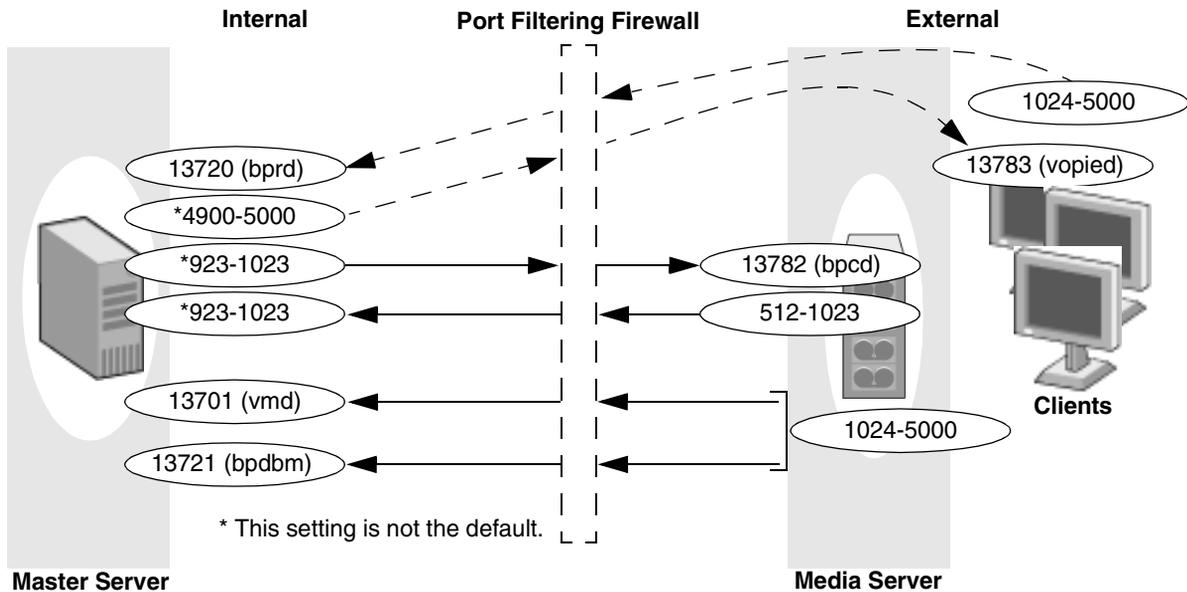
Configuration Example

The example network in the next figure “Master to Media Server and Clients Example,” shows a master server in a private (internal) network that is inside a firewall. The clients and media server are outside the firewall. To meet the port requirements shown in this figure, you must configure NetBackup to:

- ◆ Limit external connections to NetBackup in the private network by allowing the master server to accept reserved connections only on ports 923 through 1023 (the default is 512 through 1023).
- ◆ Limit NetBackup connections out of the private network by allowing the master server to:
 - ◆ Use only ports 4900 through 5000 for nonreserved-port connections to the clients (the default is 1024 through 5000).
 - ◆ Use only ports 923 through 1023 for reserved-port connections to `bpcd` on the media server (the default is 512 through 1023).



Master to Media Server and Clients Example



To configure NetBackup, perform the following on the master server (no changes are required on the media server or clients):

1. Add `CLIENT_RESERVED_PORT_WINDOW=923 1023` to the `bp.conf` file.

This specifies the reserved ports that the master server can use to originate connections, including those to `bpcd` on the media server.

2. Add `CLIENT_PORT_WINDOW=4900 5000` to the `bp.conf` file.

This specifies the nonreserved ports that the master server can use to originate connections, including those to `vopied` on the client.

3. Add `SERVER_RESERVED_PORT_WINDOW=923 1023` to the `bp.conf` file.

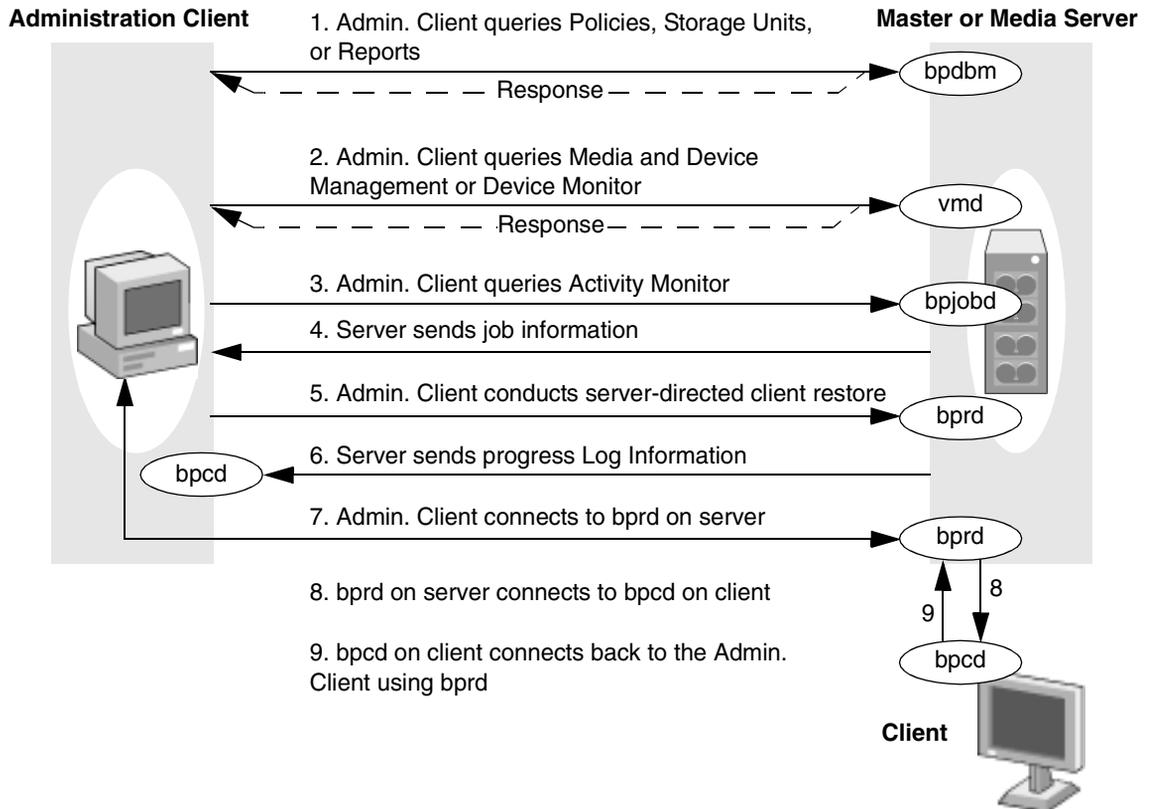
This specifies the reserved ports where the master server can elect to listen for connections.

The master server chooses from this range when it specifies the ports where it will listen for call back connections from `bpcd`.

Administration Client Connections

An Administration Client is a Windows system with the NetBackup Administration Console installed. You can use this client to perform all the administrative tasks that can be performed from a NetBackup server. The following figure shows the connections between an Administration Client and a NetBackup server. The accompanying table, "Connections for Administration Client," describes each connection and defines the ports that NetBackup uses.

Administration Client Connections



Connections for Administration Client

Description	Connect From		Connect To	
	Host	Port	Host	Port
1. Request backup policy, storage unit, or report information. The response is over the same connection. Each application requires its own connection.	Admin. Client	Non-reserved	Master Server	13721 (bpdbm)
2. Request information from Media and Device Management or Device Monitor. The response is over the same connection. Each application requires its own connection	Admin. Client	Nonreserved	Master or Media Server	13701 (vmd)
3. Request to Activity Monitor for job information.	Admin. Client	Reserved or Nonreserved (whichever the server is configured to accept for bpjobjd)	Master Server	13723 (bpjobjd)
4. Connect back to pass job information to Administration Client.	Master Server	Reserved or Nonreserved (whichever is used during request in 3 above)	Admin. Client	Reserved or nonreserved port specified during request in 3 above.
5. Request server-directed restore of a client.	Admin. Client	Nonreserved	Master Server	13720 (bprd)
6. Send progress log information to the Administration Client.	Master or Media Server	Reserved or Nonreserved (whichever the client is configured to accept)	Admin. Client	13782 (bpcd)
7. Configuring the NetBackup client: Administration Client connects to bprd on the server.	Admin. Client	Reserved or Nonreserved (whichever the client is configured to accept)	Master Server	13720 (bprd)



Connections for Administration Client (continued)

Description	Connect From		Connect To	
	Host	Port	Host	Port
8. bprd on the server connects to bpcd on the client.	Master Server	Reserved or Nonreserved (whichever the client is configured to accept)	Client	13782 (bpcd)
9. bpcd on the client connects back to the Administration Client using bprd.	Client	Reserved or Nonreserved (whichever is used during request in 8 above)	Master Server	Reserved port specified during request in 8 above.

Configuring Ports When Using an Administration Client

The following section explains the NetBackup configuration settings for ports. All settings in the following are in the `/usr/opensv/netbackup/bp.conf` file on the respective host. For more information, see “NetBackup Configuration Options” on page 134 in the *NetBackup System Administrator’s Guide, Volume II*. Registered port numbers (for example, 13782 for bpcd) are not configurable with these settings and VERITAS recommends that you do not attempt to change the registered port numbers.

Ports from which NetBackup originates connections to other hosts

These are the *Connect From* ports in the “Connections for Administration Client” table. You can set the following on each host:

- ◆ Range of reserved ports from which NetBackup can originate connections:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Set the range for the reserved ports in the **Client Reserved Port Window** fields.

Or, add `CLIENT_RESERVED_PORT_WINDOW` to the `/usr/opensv/netbackup/bp.conf` file as described in “`CLIENT_RESERVED_PORT_WINDOW`” on page 147 in the *NetBackup System Administrator’s Guide, Volume II*.

- ◆ Range of nonreserved ports from which NetBackup can originate connections:



Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Set the range for the nonreserved ports in the **Client Port Window** fields.

Or, add `CLIENT_PORT_WINDOW` to the `/usr/opensv/netbackup/bp.conf` file as described in “`CLIENT_PORT_WINDOW`” on page 146 in the *NetBackup System Administrator’s Guide, Volume II*.

- ◆ Random or nonrandom port selection. By default, NetBackup chooses a port at random from those available in the allowed range. To have NetBackup start at the top of the allowed range and choose the first available port:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Clear the check for **Use Random Port Assignment**.

Or, set `RANDOM_PORTS=NO` in the `/usr/opensv/netbackup/bp.conf` file as described in “`RANDOM_PORTS`” on page 157 in the *NetBackup System Administrator’s Guide, Volume II*.

Ports where NetBackup listens for connections from other hosts

These are the *Connect To* ports in the table “Connections for Administration Client.” You can set the following on each host:

- ◆ `bpcd` to accept connections from nonreserved ports (the default is to *not* accept these connections). To permit connections from nonreserved ports:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Universal Settings**. Check **Allow Non-reserved Ports**.

Or, add `ALLOW_NON_RESERVED_PORTS` to the `/usr/opensv/netbackup/bp.conf` file on that host as described in “`ALLOW_NON_RESERVED_PORTS`” on page 137 in the *NetBackup System Administrator’s Guide, Volume II*.

Note To specify nonreserved port usage for the client:

On the master, expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Client Attributes**. Check **Connect on Non-reserved Port**. Or, see “`ALLOW_NON_RESERVED_PORTS`” on page 137 in the *NetBackup System Administrator’s Guide, Volume II*.

- ◆ Range of reserved ports where NetBackup can listen for connections to this host:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Set the range for the reserved ports in the **Server Reserved Port Window** fields.

Or, add `SERVER_RESERVED_PORT_WINDOW` to the `/usr/opensv/netbackup/bp.conf` file on that host as described in “`SERVER_RESERVED_PORT_WINDOW`” on page 160 in the *NetBackup System Administrator’s Guide, Volume II*.

- ◆ Range of nonreserved ports where NetBackup can listen for connections to this host:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Set the range for the nonreserved ports in the **Server Port Window** fields.

Or, add `SERVER_PORT_WINDOW` to the `/usr/opensv/netbackup/bp.conf` file on that host as described in “`SERVER_PORT_WINDOW`” on page 160 in the *NetBackup System Administrator’s Guide, Volume II*.

- ◆ Random or non-random port selection. By default, NetBackup chooses a port at random from those available in the allowed range.

To have NetBackup start at the top of the allowed range and choose the first available port:

Expand **NetBackup Management > Host Properties > Master Server**. Double-click the host, then select **Port Ranges**. Clear the check in **Use Random Port Assignment**.

Or, set `RANDOM_PORTS=NO` in the `/usr/opensv/netbackup/bp.conf` file on that host as described in “`RANDOM_PORTS`” on page 157 in the *NetBackup System Administrator’s Guide, Volume II*.

Configuration Example

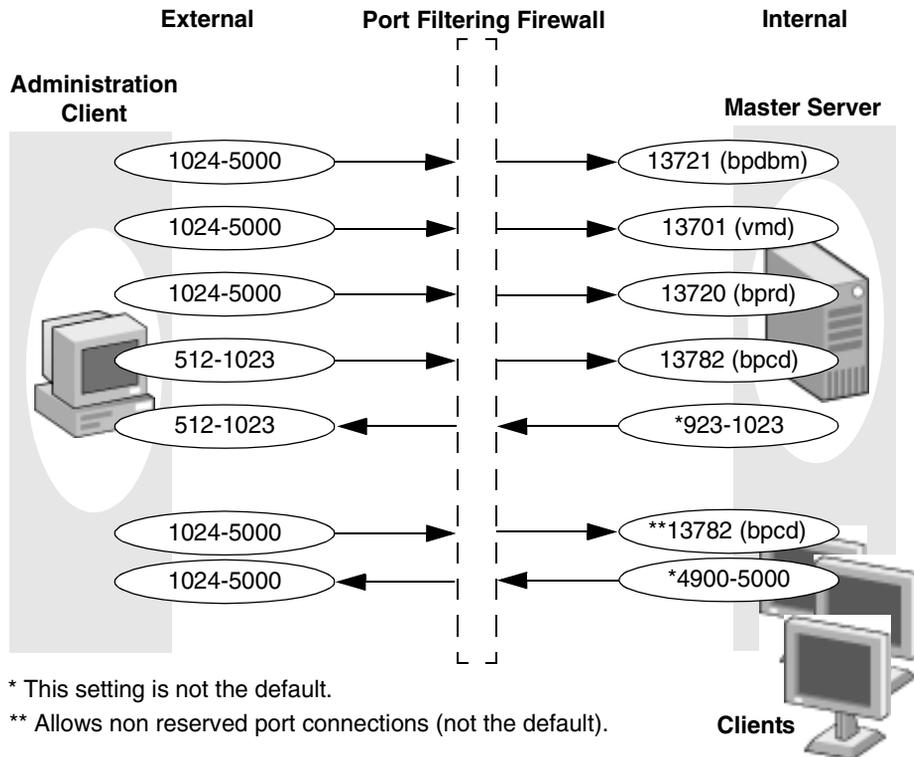
The example network in the next figure, “Master to Media Server and Clients Example,” shows a master server in a private (internal) network that is inside a firewall. You are going to use the administration client to manage the master server from outside the firewall. To meet the port requirements shown in this figure, you must configure NetBackup to:

- ◆ Limit external connections to NetBackup in the private network by allowing nonreserved port connections to `bpcd` on the master server and the clients.
- ◆ Limit NetBackup connections out of the private network by:
 - ◆ Allowing the master server to use only ports 923 through 1023 for reserved-port connections to the administration client (the default is 512 through 1023).
 - ◆ Allowing the clients to use only ports 4900 through 5000 for nonreserved-port connections to the administration client (the default is 1024 through 5000).



Note Any port limitations you configure on a NetBackup host apply to connections with *all* other NetBackup hosts, not just those on the other side of the firewall. Therefore, leave enough ports available to allow the necessary connections. The main factors to consider are the number of clients and whether multiplexing is used. If NetBackup runs out of ports, backups and restores cannot occur.

Master to Media Server and Clients Example



To configure NetBackup, perform the following steps on the master server. No configuration is required on the administration client.

1. Add `CLIENT_RESERVED_PORT_WINDOW=923 1023` to the `bp.conf` file.
This specifies the reserved ports that the master server can use to originate connections, including those to the administration client.
2. Specify that the master server can accept connections on its `bpcd` from nonreserved ports by adding `ALLOW_NONRESERVED_PORTS` to the `bp.conf` file.
3. Specify that the clients can accept connections to their `bpcd` from nonreserved ports by running the following command:



```
cd /usr/opensv/netbackup/bin/admincmd
./bpclient -client client_name -add -connect_nr_port 1
```

Where *client_name* is the name of the client (run the command for each client).

4. On the clients:

a. Add `ALLOW_NONRESERVED_PORTS` to the `bp.conf` file.

b. Add `CLIENT_PORT_WINDOW=4900 5000` to the `bp.conf` file.

This specifies the nonreserved ports the client can use to originate connections.



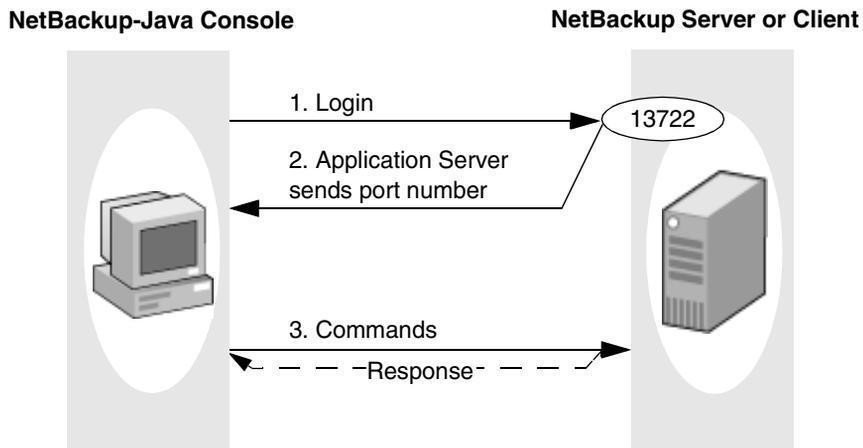
NetBackup-Java Console Connections

The NetBackup-Java administration console runs on a UNIX or Windows system where the NetBackup-Java interface software is installed. Refer to “NetBackup-Java Administration Console Architectural Overview” on page 489 for information relevant to understanding this topic.

Use the administration console on a UNIX or Windows platform to perform all the NetBackup administrative tasks. The NetBackup-Java Console can also perform backups or restores of UNIX clients.

The following figure shows the connections between a NetBackup-Java Console and the NetBackup-Java Application Server on a NetBackup server or client. The accompanying table, “Connections for the NetBackup-Java Console,” provides a brief description of each connection and defines the ports that the NetBackup Console and its application server uses.

NetBackup-Java Console Connections



Connections for the NetBackup-Java Console

Description	Connect From		Connect To	
	Host	Port	Host	Port
1. Log in to the NetBackup-Java Application Server on the NetBackup server or client.	Where the NetBackup-Java Console was started	Chosen by local host or (A)	NetBackup server or client	13722 (bpjava-msvc)

Connections for the NetBackup-Java Console (continued)

Description	Connect From		Connect To	
	Host	Port	Host	Port
<p>2. The NetBackup-Java Application Server on the server or client responds by sending the port number where the display console must connect in order to send commands. This response is over the same connection. The server then listens on that port for commands.</p> <p>A unique port (C) is specified for each user that is logged in or (B).</p>	NetBackup server or client	13722 (bpjava-msvc)	Where the NetBackup-Java Console was started	Chosen by local host or (A)
<p>3. Send commands (for example, to start Backup Policy Management). Responses from the server or client are over the same connection.</p> <p>A unique port (C) and connection is established for each user that is logged in. When a connection is established, it is used for all further commands and responses for that user or (B).</p>	Where the NetBackup-Java Console was started	Chosen by local host or (A)	NetBackup server or client	Nonreserved port specified in the login response (see step 2) or (B)
<p>4. Request to Activity Monitor for job information.</p>	Where the NetBackup-Java Console was started	Chosen by local host or (A)	NetBackup server	13724 (vnet.d) to ultimately connect to bpjobd

A. One of the ports in the range specified by the NetBackup-Java Console configuration option, `NBJAVA_CLIENT_PORT_WINDOW`. (See “`NBJAVA_CLIENT_PORT_WINDOW`” on page 499.)

B. If the NetBackup-Java Console configuration option, `NBJAVA_CONNECT_OPTION` is set to 1, no additional ports will be used. (See “`NBJAVA_CONNECT_OPTION`” on page 499.)

C. This unique port can be restricted to a configured range of ports using the `SERVER_PORT_WINDOW` option on the server or client. (See “`SERVER_PORT_WINDOW`” on page 160 in the *NetBackup System Administrator's Guide, Volume II*.)

Configuring Ports When Using the NetBackup-Java Console

On UNIX systems, all settings are in the `/usr/opensv/netbackup/bp.conf` file on the respective server, or in the `/usr/opensv/java/nbj.conf` file.



On Windows systems, all settings are made with the Host Properties dialog in the NetBackup Administration Console, or in the `install_path\Veritas\java\<hostname>.vrtsnbuj` files.

Port configurations may include:

- ◆ Range of nonreserved ports where the server or client can listen for connections. These are the *Connect To* ports in the preceding table, “Connections for the NetBackup-Java Console.” Use `SERVER_PORT_WINDOW` in the `bp.conf` file for this setting.

The server or client selects a port from this range to listen for commands from the display console. Note that the highest port available in the allowed range is always used internally by the NetBackup-Java Application Server on the server or client.

- ◆ Random or nonrandom port selection. By default, the NetBackup-Java Application Server chooses a port at random from those available in the allowed range. To have the NetBackup-Java Application Server start at the top of the allowed range and choose the first available port, set `RANDOM_PORTS=NO` in the `bp.conf` file.

Note All the above settings are in the `/usr/opensv/netbackup/bp.conf` file on the server or client. For more information, see “NetBackup Configuration Options” on page 134 in the *NetBackup System Administrator’s Guide, Volume II*. Registered port numbers (for example, 13782 for `bpcd`) are not configurable with these settings and VERITAS recommends that you do not attempt to change the registered port numbers.

On all NetBackup-Java capable platforms you can configure the following:

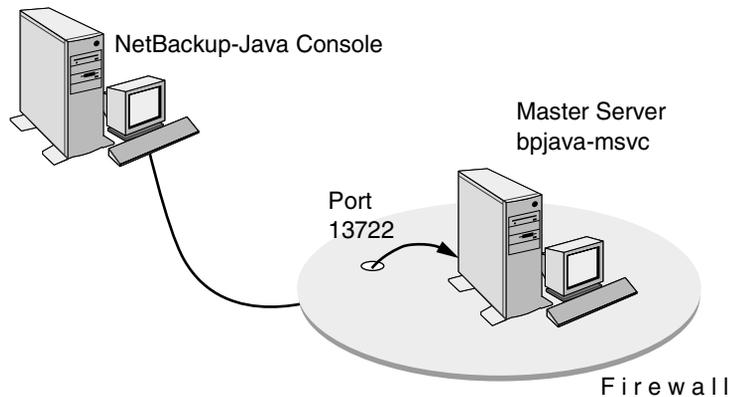
- ◆ The range of nonreserved ports on the console available for connecting to the NetBackup-Java Console configuration. (See “`NBJAVA_CLIENT_PORT_WINDOW`” on page 499.)
- ◆ Configure the NetBackup-Java Console to *not* use a unique port for every user using the console via the NetBackup-Java Console configuration option. (See “`NBJAVA_CONNECT_OPTION`” on page 499.) This requires access to the `vnetd` daemon on its port.

Note These options cannot be configured using the NetBackup-Java Console **Host Properties** dialog. The `nbj.conf` file or the `hostname.vrtsnbuj` file on the relevant host must be edited.

Configuration Example

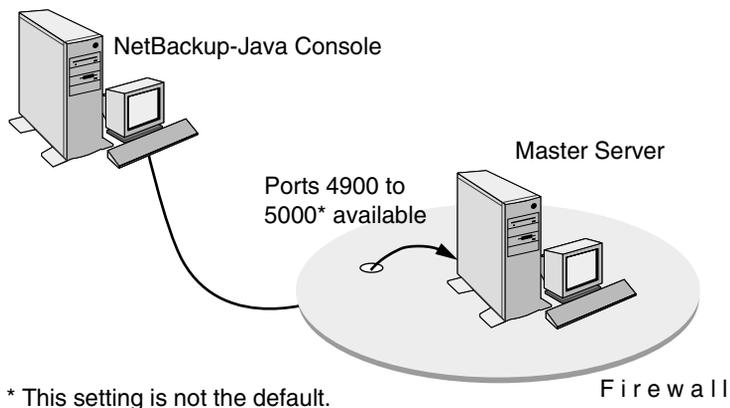
This example concerns using the NetBackup-Java Console to manage a master server that is inside a firewall.

NetBackup-Java Console to Server Example



The port requirements in this example are as follows:

- ◆ Limit external connections to the master server by allowing the master server to accept nonreserved-port connections only on ports 4900 through 5000 (the default is 1024 through 5000).
- ◆ Ports are to be selected by using the first one available, starting at the top of the allowed range.



Note Any port limitations you configure on a master server apply to connections with all other master servers, not just those on the other side of the firewall. Therefore, leave enough ports available to allow the necessary connections. The main factors to consider are the number of clients and whether multiplexing is used. If NetBackup runs out of ports, backups and restores cannot occur.

In order to effect the configuration of the master server, perform the following steps in the **Host Properties** dialog using the NetBackup-Java Console:



▼ **To configure the master server according to the example**

1. In the NetBackup Administration Console on the master server, expand **NetBackup Management > Host Properties > Master Servers**.
2. In the Details pane, double-click the host you wish to configure.
3. Select **Port Ranges**:

The **Port Ranges** settings specify the range of nonreserved ports from which the master selects a port to listen for command connections. Note that the highest port available in the range is always used internally by NetBackup-Java (in this example, the highest port that can be available is 5000).

4. Clear the check box for **Use Random Port Assignments**.

With this option unselected, NetBackup uses the first port available, starting at the top of the allowed range. In the example, the highest port that can be available is 4999 because 5000 is claimed by NetBackup-Java.

Note On a NetBackup UNIX client, add the following to the

```
/usr/opensv/netbackup/bp.conf file:  
SERVER_PORT_WINDOW = 4900 5000  
RANDOM_PORTS = NO
```

▼ **To configure the master server to use vneta according to the example**

1. In `/usr/opensv/java/nbj.conf`, indicate that the NetBackup-Java Console should use the no call-back method when communicating with other NetBackup machines:

```
NBJAVA_CONNECT_OPTION=1
```

Setting `NBJAVA_CONNECT_OPTION` to 1 means that the NetBackup-Java Console will use only one port, the `vneta` port, for communication with its application server.

2. If desired, specify a range of nonreserved outgoing ports on which the NetBackup-Java Console requires to connect to its application server. For example:

```
NBJAVA_CLIENT_PORT_WINDOW=5700 5900
```

The minimum range size for successful operation of the NetBackup-Java Console is 120.

Note Performance is somewhat reduced with the use of `NBJAVA_CONNECT_OPTION` or `NBJAVA_CLIENT_PORT_WINDOW`.

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.

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

If you require a storage migration product, VERITAS recommends that you use Storage Migrator for UNIX (VSM). NetBackup can back up files from a disk type storage unit that has file systems managed by VSM.

When a file is migrated, its 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 they be needed locally.

NetBackup proceeds as follows when backing up a file that has been purged by Storage Migrator:

- ◆ For user backups by nonroot users, NetBackup first caches the files and then backs them up.

- ◆ 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 following value is in effect during the caching of the (potentially) migrated backups: **Host Properties > Master Server > Global NetBackup Attributes > Enforce Media Mount Timeout after**.

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 446.)

Note If you use 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 they are backed up. This makes it appear as if the files are frequently used and stops Storage Migrator from selecting them for migration.

that replaces any file in `/usr/opensv/netbackup/bin/admincmd` or files `bpbackup`, `bplist` or `bprestore` in `/usr/opensv/netbackup/bin`.

Configuring the NetBackup-Java Console

NetBackup-Java Administration Console Architectural Overview

The NetBackup-Java Administration Console is a distributed application consisting of two major (and separate) system processes:

- ◆ The NetBackup Administration Console graphical user interface (`jnbSA`)
- ◆ The application server (`bpjava` processes).

These processes may be running on two physically different NetBackup server hosts. This distributed application architecture holds true for the Backup, Archive, and Restore client graphical user interface (`jbpSA`) as well.

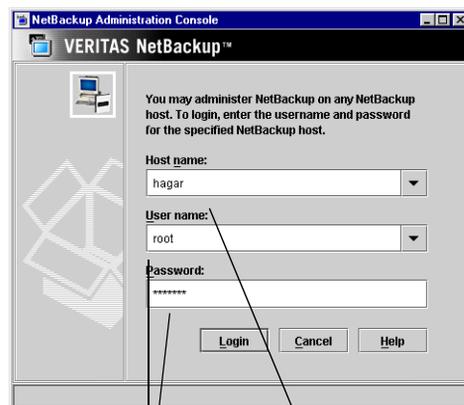
After the NetBackup Administration Console interface is started using the `jnbSA` command, the user is required to log in to the application server on the host specified in the login dialog.

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 you start the NetBackup-Java console.

The login credentials of the user are authenticated by the application server on the host specified in the NetBackup Administration Console login dialog using standard UNIX system user account data and associated APIs. This means that the provided login credentials must be valid on the host specified in the login dialog.

The server that is usually the object of all administrative tasks is the one specified in the NetBackup Administration Console login dialog.

NetBackup login dialog:



Application server
User name and password must be valid on application server



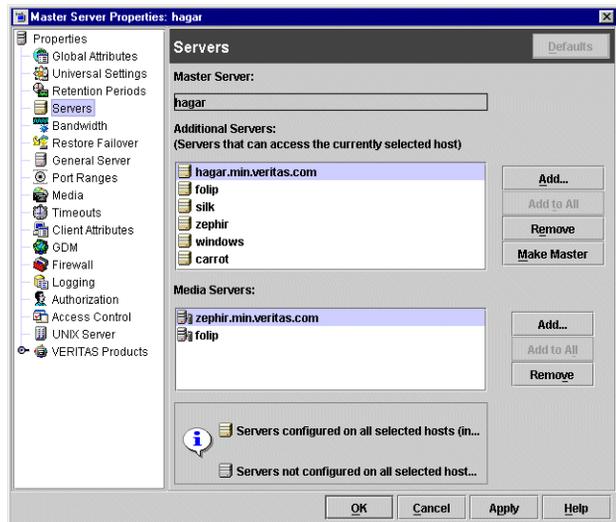
The exception to this is the use of the **File > Change Server** capability in the NetBackup Administration Console. The **Change Server** capability allows administration of a remote server (a server other than the one specified in the NetBackup Administration Console login dialog).



Regardless of which server is being administered (a remote server or the server specified on the login dialog), all administrative tasks performed in the NetBackup Administration Console make requests of the application server and are run on the application server host.

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 420.)

This context (switching to a remote server from the application server) also applies to the Enhanced Authentication and Authorization capabilities (see Chapter 9). 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 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 login dialog for activity monitoring tasks. For additional configuration information concerning this, see the information on NetBackup-Java console connections in “NetBackup-Java Console Connections” on page 480.

Authorizing NetBackup-Java Users

The *NetBackup System Administrator's Guide, Volume II* documents two types of user authorization: NetBackup Access Control (new in NetBackup 5.0), and Enhanced Authorization and Authentication. If either method is configured, you may choose to authorize users of the NetBackup-Java administration console for specific applications. 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 494.

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 login 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 login 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 login dialog for the NetBackup-Java Console.

Authorization File

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.

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.

- ◆ 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 494.
 - ◆ 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 495.
 - ◆ 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 login 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 login dialog box will accept passwords greater than eight characters. However, only the first eight are significant when logging into 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 {0}.
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
BPM	Backup Policy Management
BAR or JBP	Backup, Archive, and Restore
CAT	Catalog
DM	Device Monitor
HPD	Host Properties
MM	Media Management
REP	Reports
SUM	Storage Unit Management
VLT	Vault Management

For example, to give a user 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, script `/usr/opensv/java/nonroot_admin_nbjava` can be executed to change the permissions on the following files:

```
/usr/opensv/java/auth.conf
/usr/opensv/java/Debug.properties
/usr/opensv/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 431.
- ◆ 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.

In addition, when authorized for `ALL`, the user can view a list of media IDs required for the files marked for restore through the **Preview Media Required** button at the bottom of the **Restore Files** tab in `jbpSA`. This capability exists only 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

File `/usr/openssl/java/nbj.conf` contains configuration options for the NetBackup-Java console.

Use the following syntax rules when creating entries in `nbj.conf`:

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

BPJAVA_PORT, VNETD_PORT

These are the configured ports for the `bpjava-msvc` and `vnetd` daemon processes. These ports are registered with IANA and it is not recommended they be changed.

Port	Process	Registered Default Port Number
<code>bpjava-msvc</code>	<code>BPJAVA_PORT</code>	13722
<code>vnetd</code>	<code>VNETD_PORT</code>	13724

If the ports for these process do need to be changed, 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` option.

FORCE_IPADDR_LOOKUP

Specifies whether NetBackup will perform an IP address lookup to determine if two host name strings are indeed the same host.

This option is found in `/usr/openssl/java/nbj.conf` on NetBackup servers in 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 “Rules for Using Host Names in NetBackup” on page 234, *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 (how it appears 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`

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 options allow configuration of memory usage for the Java Virtual Machine (JVM) and are found in `/usr/opensv/java/nbj.conf`.

We recommend running the NetBackup-Java Console (`jnbSA`) or Backup, Archive and Restore client application (`jbpsA`) 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` or `jbpsA` on a machine with the recommended amount of memory. It can also be specified on 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). It can also be specified on the `jnbSA` or `jbpsA` command. For example:

```
jnbSA -mx 512M
```

Default = 256M (megabytes).

MEM_USE_WARNING

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 is found in `/usr/opensv/java/nbj.conf`.

NBJAVA_CLIENT_PORT_WINDOW

Specifies the range of nonreserved ports on this computer that are used for connecting to the NetBackup-Java application server or the `bpjobd` daemon (or service on Windows) from the NetBackup-Java Administration Console's Activity Monitor.

This option is found in `/usr/opensv/java/nbj.conf` on NetBackup servers in 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 NetBackup-Java application server (NetBackup Administration Console/`jnbSA`) or the `bpjobd` daemon (or service on Windows) 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 Administration Console/`jnbSA`.
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 `njb.conf` entry for three concurrent users might look as follows:

```
NBJAVA_CLIENT_PORT_WINDOW = 5000 5360
```

If the range is not set wide enough, `jnbSA` will exit 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.
```

NBJAVA_CONNECT_OPTION

Specifies the call-back method the server or client will use when communicating with the NetBackup-Java consoles (`jnbSA`, `jbpSA`).

This option is found in `/usr/opensv/java/nbj.conf` on NetBackup servers in the following format:

```
NBJAVA_CONNECT_OPTION = [ 0 | 1 ]
```

Where:

0 = Indicates the traditional call-back method (default).



1 = Indicates the `vnetd` no call-back method.

Note Performance is somewhat reduced with the use of `NBJAVA_CONNECT_OPTION`.

For more information, refer to the relevant topics in “Multiplexing” on page 468.

Configuration Options Relevant to `jnbSA` and `jbpSA`

There are several configuration options available to administrators when using the NetBackup Administration Console and the Backup, Archive, and Restore client interface through the 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 Administration Console or by the Backup, Archive, and Restore client 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.

For more information on the `bp.conf` file, see “NetBackup Configuration Options” on page 134, *NetBackup System Administrator’s Guide, Volume II*.

NetBackup-Java Performance Improvement Hints

The most important factor to consider when faced with performance issues while using the NetBackup-Java Administration Console or the NetBackup, 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 NetBackup-Java Administration Console from the following locations:

- ◆ Locally on your desktop host (on supported Windows and UNIX platforms), or
- ◆ Remotely and displaying back to your desktop host (from supported UNIX platforms)

The recommended method for using the NetBackup-Java Administration Console or the NetBackup, Archive, and Restore user interface, 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 on host** 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 489 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` `nbj.conf` options, assumes sufficient memory resources on the machine you execute the `jnbSA` or `jbpSA` commands.

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
```

For more information, refer to the conf options “INITIAL_MEMORY, MAX_MEMORY” on page 498.

- ◆ 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 `jnbSA` is was 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 `jnbSA` is was 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 293.)

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 `jnbSA` command was 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 `jnbSA` command 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. However, only since NetBackup Feature Pack 4.5_3_F will you possibly see adequate performance with these types of tools on a Windows machine because previous versions of NetBackup used a version of Java that performed *very* poorly in this type of configuration.

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
Administrator Utilities	
<code>bpadm</code>	Starts character-based, menu-driven administrator's interface on the server.
<code>jnbSA</code>	Starts Java-based, NetBackup administrator's interface on the server.
Client-User Interfaces	
<code>bp</code>	Starts character-based, menu-driven client-user interface.
<code>jbpSA</code>	Starts Java-based, client-user interface on the client.
Daemon Control	
<code>initbprd</code>	Starts <code>bprd</code> (request daemon).
<code>bprdreq -terminate</code>	Stops <code>bprd</code> (request daemon)
<code>initbpdbm</code>	Starts <code>bpdbm</code> (database manager).
<code>bpadm</code>	Has option for starting and stopping <code>bprd</code> .
<code>jnbSA (Activity Monitor)</code>	Has option for starting and stopping <code>bprd</code> .
Monitor Processes	
<code>bpps</code>	Lists active NetBackup processes.
<code>jnbSA (Activity Monitor)</code>	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).

Index

Symbols

- \$HOME 322
- .f files in catalog 230
- .SeCuRiT.y.nnnn files 440

A

- Accept Connections on Non-reserved Ports
 - property 392
- Access Control
 - authorizing users 427
 - host properties
 - Authentication Domain tab 307
 - Authorization Service tab 308
 - VERITAS Security Subsystem (VxSS) 304
 - VxSS Networks List 305
 - VxSS tab 304
 - NetBackup 491
- access control lists (ACLs) 112, 440
- accessibility features xxxvii
- ACL (see access control lists)
- active
 - job 287
- Activity Monitor
 - bpdbjobs command 294
 - BPDBJOBS_OPTIONS environmental variable 295
 - cancel uncompleted jobs 281
 - column heads, selecting to view 281
 - delete completed jobs 281
 - detailed job status 281, 290
 - disabling job logging 372
 - monitoring jobs 281
 - resume restore jobs 282
 - saving job data to a file 282
 - set column heads to view 281
 - suspend restore jobs 282
 - using the Troubleshooter 283
- Actual
 - Client property 313
 - Path property 313
- adding
 - a media server to the Alternate Restore Failover Machine list 379
 - catalog backup file paths 210
 - clients to policy 98
 - licenses 416
 - new license key 417
 - pathname 211
 - schedules 167
 - storage unit
 - disk type 36
 - Media Manager type 42, 45
 - NDMP type 37
- Administer E-mail Address property 366
- administering
 - a remote system 425
 - NetBackup
 - using the Java interface 9
- administrator
 - defined streaming mode 134
 - nonroot 494
- Advanced Client 164
 - backups using Checkpoint Restart 79
 - FlashBackup 75
 - in the Policy Attributes tab 97
- AFS policy type 75
- AIX cache file system 130, 140
- All Log Entries report 266, 364
- Allow Backups to Span Media property 371
- Allow Block Incrementals property 360
- Allow Client
 - Browse property 318
 - Restore property 318
- Allow Media Overwrite property 370
- Allow Multiple Data Streams
 - directives 133



- set policy attribute 93
 - when to use 93
 - Allow Multiple Retentions per Media
 - property 371
 - Allow Server File Writes property 303, 392
 - alternate client restores, allowing 432
 - Alternate Restore Failover Machines host
 - properties 378
 - altnames file 431
 - Announce DHCP interval property 374
 - ANSI format, allow overwrite 370
 - AOS/VS format, allow overwrite 370
 - application backups 147
 - archive bit 155, 326, 328, 329
 - atime 446, 487
 - attributes for a policy 93
 - auth.conf file
 - capabilities identifiers 495
 - description 492
 - entries for specific applications 494
 - overview 491
 - Authorization
 - host properties 310
 - DomainGroup 310
 - Group/Domain 311
 - Host 310
 - User 310
 - User must be an OS
 - Administrator 311
 - preferred group 391
 - auto-discover streaming mode 135
 - automatic
 - backups 148
 - cumulative incremental
 - backups 148
 - differential incremental backups 148
 - full backups 148
 - incremental backups 148
 - property (for selection of ports) 345
 - Vault policy type 148
 - automounted directories 83
- B**
- Backup End Notify Timeout property 387
 - Backup Exec QIC media, importing and
 - restoring 252
 - Backup Exec Tape Reader
 - host properties
 - Actual Client 313
 - Actual Path 313
 - GRFS Advertised Name 312
 - Backup Migrated Files property 323
 - Backup Start Notify Timeout property 386
 - Backup Status report 364
 - backups
 - activating policy 82
 - application 147
 - archive 147
 - automatic 148
 - cumulative incremental 146, 148
 - differential incremental 147, 148
 - full 148
 - incremental 148
 - Vault 148
 - automatic, introduction to 1
 - balancing load 485
 - best times for user directed 193
 - clients using Storage Migrator 486
 - duplicating 241
 - frequency
 - effect on priority 165
 - setting 165
 - full 146
 - import 249
 - NFS mounted files 73, 83
 - policy management window 58
 - raw partitions on Windows 80, 124
 - registry on Windows clients 125
 - selections list, verifying 109
 - Status of Backups report 263
 - types of 146
 - user directed
 - overview 1
 - schedules 193
 - type of backup 147
 - verifying 240
 - windows
 - duration, examples 174
 - specifying 173
 - Bandwidth
 - host properties
 - Bandwidth 314
 - Bandwidth Throttle Setting for the
 - Range of IP Addresses 314
 - Bandwidth 314
 - Bandwidth Throttle Settings List 315
 - From IP Address 314
 - To IP Address 314



-
- Bandwidth Throttle Settings List
 - property 315
 - Bare Metal Restore (BMR)
 - no NetBackup change journal support 328
 - batch file example for setting bpdjobs
 - environmental variable 295
 - BE-MTF1 format, allow overwrite 370
 - binary catalog format 161
 - Block Level Incremental Backups 79
 - bp.conf
 - entries for Activity Monitor 293
 - BPBRM Logging property 368
 - bpcatarc catalog archiving command 234
 - bpcatlist catalog archiving command 234
 - bpcatres catalog archiving command 235
 - bpcatrm catalog archiving command 235
 - BPCD Connect-back property 320, 344
 - BPCD port setting on client 374
 - bpchangeprimary command 247
 - bpclient
 - add clients to catalog 441
 - delete clients from catalog 442
 - list clients in catalog 442
 - preventing lists and restores 442
 - bpconfig command 136
 - bpcoord log 163
 - bpdjobs
 - adding entries to bp.conf file
 - Activity Monitor
 - adding bp.conf entries for 293
 - batch file example 295
 - command 294
 - debug log 296
 - BPDBJOBS_OPTIONS environmental variable 295
 - bpdbm
 - running without bprd 415
 - starting automatically 415
 - stopping bpdbm 415
 - BPDBM Logging property 368
 - BPDM Logging property 368
 - bpend 387
 - bpexpdate 245
 - bpfis directory for VSP logging
 - messages 397
 - BPJAVA_PORT 496
 - bpps script 413
 - BPRD
 - Logging property 368
 - managing 414
 - port setting on client 374
 - terminating 414
 - BPSCHED Logging property 368
 - bpstart 386
 - bpsynth log 163
 - BPTM
 - (Drive Count) Query Timeout property 388
 - Logging property 368
 - query 388
 - Browse and Restore Ability property 320
 - buffer size 329
 - Busy Action property 317
 - Busy File
 - host properties
 - Busy File Action 317
 - File Action File List 317
 - Operator's E-mail Address 316
 - Process Busy Files 316
 - Retry Count 317
 - Working Directory 316
 - processing
 - Windows clients 375
- C**
- calendar scheduling
 - how it interacts with daily windows 179
 - using 176
 - cancelling uncompleted jobs 281
 - catalog backups
 - adding a pathname 211
 - automatic 201
 - caution for compressing 227
 - changing a pathname 212
 - compressing image catalog 226
 - configuration 203
 - deleting a pathname 212
 - disk path 208
 - file paths
 - adding 210
 - media server 214
 - Windows master 213
 - indexing 236
 - last media used 205
 - manual backup 215



- media
 - ID 207
 - server 205
 - type 206
- overview 200
- recovery 202
- setting schedules 209
- space required 223
- uncompressing 227
- catalogs
 - archiving 230
 - bpcatarc 234
 - bpcatlist 234
 - bpcatres 235
 - bpcatrm 235
 - catarc policy 232
 - extracting images 236
 - inactive policy 232
 - overview 231
 - retention level setting 232
 - type of backup indicated 232
 - image files 230
 - managing 222
 - moving client images 228
 - multiple file layout 231
 - single file layout 230
- catarc schedule 62
- cautions
 - alternate client restores 432
 - alternate path restore 434
 - database compression 227
 - retention time 170, 193
 - wildcards in UNIX raw backups 117
 - worm retention 171
- CDE (Common Desktop Environment) 5
- cdrom file system 130, 140
- change journal 328
 - and synthetic backups 163
 - determining if enabling is useful 327
 - using in incremental backups 327
- changing
 - catalog backup attributes 204
 - clients in a policy 99
 - licenses 416
 - pathname 212
 - policy properties 62, 70, 71, 72
 - server
 - for configuring storage units 8
- Checkpoint Restart
 - and synthetic backups 163
 - Move Job From Incomplete State to Done State 365
 - Move Restore Job from Incomplete State to Done State 446
 - Restore Retries 446
 - resuming a restore job 446
 - suspending a restore job 446
- cipher types for NetBackup Encryption 332
- clearing sort information in Policies
 - application 19
- client
 - database 318
 - exclude and include lists 342
 - name 432
 - speed 363
- Client Administrator's E-mail property 393
- Client Attributes
 - host properties
 - Allow Client Browse 318
 - Allow Client Restore 318
 - Browse and Restore Ability 320
 - Clients List 318
 - Connect on Non-reserved Port 319
 - Free Browse 321
 - Maximum Data Streams 320
 - No Connect Back 320
- Client Backups report 264
- Client Connect Timeout property 386
- Client Name host property 322
- Client Port Window property 376
- Client Read Timeout property 386, 387, 388
- Client Reserved Port Window property 377
- Client Sends Mail setting 393
- clients
 - adding to policy 98
 - BPCD port 374
 - BPRD port 374
 - changing in a policy 99
 - choosing policy type 74
 - definition 2
 - deleting from policy 72
 - DHCP Interval property 374
 - exclude file list
 - Windows 335
 - exclude files list
 - UNIX 139
 - include files list 143
 - installing 99, 101



- maximum jobs 363
 - NetBackup
 - moving image catalog 228
 - secure 101
 - setting host names 98
 - software 2
 - trusting clients 99
 - Clients List property 318
 - clustered enviroment, changing host
 - properties 301
 - Collect True Image Restoration (TIR) with
 - Move Detection property 156
 - collecting disaster recovery information 84, 93
 - Communications Buffer property 329
 - Compress Catalog Interval property 226, 364
 - compression
 - advantages 91
 - disadvantages 91
 - specifications 91
 - concurrent jobs
 - on client 363
 - per policy 81
 - Connect on Non-reserved Port property 319
 - copies, third-party 167, 244
 - copy, primary 244
 - cpio format, allow overwrite 370
 - cross mount points
 - effect with UNIX raw partitions 84
 - examples 85
 - separate policies for 84
 - setting 84
 - ctime 120
 - cumulative incremental backups 146, 149
 - not combining with differential incremental backups 61, 147, 329
 - Current NBAC User 23
 - customizing NetBackup Administration
 - Console 23
- D**
- Daemon Connection Port property 345
- Daemon Port Only property (for selection of ports) 345
- daemons
 - bpdbm
 - starting automatically 415
 - starting with bprd 414
 - bprd
 - managing 414
 - terminating 414
 - checking processes 413
 - stopping 414
 - Daily windows setting 179
 - Dashboard Port Window property 358
 - database-extension clients
 - add file paths for 129
 - databases, NetBackup (see catalog backup)
 - DataStore policy type 74
 - datetime stamp 151
 - DB2 policy type 74
 - DBR format, allow overwrite 370
 - Delay on Multiplexed Restores property 359
 - Delete Vault Logs property 364
 - deleting
 - backup selections from a policy 72
 - clients from a policy 72
 - clients from client catalog 442
 - completed jobs 281
 - license keys 417
 - pathname 212
 - policies 72
 - schedules 72
 - storage unit groups 56
 - detailed job status 281, 290
 - Device Monitor 292
 - DHCP setting on client 374
 - differential incremental backups 147
 - not combining with cumulative incremental backups 61, 147, 329
 - Direct Access Recovery (DAR) 360
 - disaster recovery
 - catalogs 202
 - collect information for 84, 93
 - disk staging
 - schedule 49
 - storage units 38
 - size recommendations 39
 - using Checkpoint Restart 80
 - disk storage units 36, 53
 - disk-image backups 80, 124
 - Do Not Compress Files Ending With
 - property 325
 - DO_NOT_RESET_FILE_ACCESS_TIME 48
 - 7
 - done job 287
 - duplicate backups



- becoming a primary copy 244
- creating 241
- restoring from 246
- duration of backup window, examples 174

E

- E-mail
 - address for administrator of this client 393
 - send from client 393
 - send from server 393
- Enable
 - Job Logging property 372
 - Open File Backup During Backups property 375
 - Performance Data Collection property 393
 - SCSI Reserve/Release property 372
 - Single Instance Backup for Message Attachments property 334
 - Standalone Drive Extensions property 372
- Encryption
 - host properties
 - Encryption Keyfile 333
 - Encryption Libraries 332
 - Encryption Permissions 331
 - Encryption Strength 332
 - in Client Backups Report 264
 - in Images on Media Report 271
 - policy attribute 92
 - use with synthetic backups 161
- Encryption Key File property 333
- Encryption Libraries property 332
- Encryption Permissions property 331
- Encryption Strength property 332
- Encryption, NetBackup 332
- Enhanced
 - Authentication 427, 491
 - Authorization 391, 427, 491
- escape character
 - backslash 340
 - on UNIX 112, 141
- Exceptions to the Exclude List host property 335
- Exchange
 - host properties
 - Enable Single Instance Backup for Message Attachments 334

- Mailbox for Message Level Backup and Restore 334
- exclude file lists
 - on client 342
- exclude files list 139
 - Windows example 341
- Exclude List
 - host properties
 - Exceptions to the Exclude List 335
 - Use Case Sensitive Exclude List 335
- exclude lists
 - creating 139
 - example 142
 - for specific policies and schedules 142
 - syntax rules 140
 - wildcards in 140
- excluding files and directories from backup 335
- Executable Directory property 369
- export license key 418
- extended attribute files
 - disabling the restore of 121
 - Solaris 9 113

F

- failover
 - media server to alternate media server(s) 378
- File Browse Timeout property 386
- file lists
 - disk image on Windows 124
 - extension clients 129
 - links on UNIX 114
 - NetWare clients
 - NonTarget 127
 - Target 128
 - raw partitions 116, 124
 - standard clients 111
 - UNIX files not backed up 112
 - Windows clients 122
- file systems 84
- files
 - .SeCuRiT.y.nnnn 440
 - /.rhosts 99
 - catalog space requirements 223
 - excluding from backup 335
 - for catalog backup 210
 - goodies scripts 450
 - linked, UNIX 113



- NFS mounted 73, 83
- No.restrictions 432
- NOTES.INI 369
- peername 433
- restoring to alternate client 434
- restrictions on restores 432
- version xxxiii
- filters, applying job 281
- Firewalls
 - configuring port limitations 471
 - host properties
 - BPCD Connect-back 344
 - Daemon Connection Port 345
 - Hosts list 344
 - Ports 345
 - minimum ports required 347
 - using vnetd with 320, 344
- FlashBackup
 - policy type 75
 - single file restore program (sfr) 113
- follow NFS mounts
 - advantages of 84
 - disadvantages of 83
 - notes on use
 - with cross mount points 83
 - with raw partitions 83
 - with cross mount points 85
- Follow NFS setting 83
- fragment, Media Manager storage unit 50
- Free Browse property 321
- freeze media 268
- From IP Address property 314
- full backups 146, 148, 157

G

- GDM (Global Data Manager)
 - host properties
 - Dashboard Port Window 358
 - Use OS Selected Non-reserved Port 358
- General Level Logging property 326
- General Server
 - host properties
 - Allow Block Incrementals 360
 - Delay on Multiplexed Restores 359
 - Must Use Local Drive 359
 - Re-read Interval for Available Drives 359
 - Use Direct Access Recovery for

- NDMP Restores 360
 - Use Media Host Override for Restores 360
- Global Attributes
 - host properties
 - Administrator E-mail Address 366
 - Compress Catalog 364
 - Delete Vault Logs 364
 - Keep Logs 364
 - Keep True Image Restoration Information 364
 - Maximum Backup Copies 364
 - Maximum Jobs per Client 363
 - Move Job From Incomplete State to Done State 365
 - Schedule Backup Attempts 362
 - Status Report Interval 362
 - Wakeup Interval 362
- Global Attributes host properties
 - Maximum Jobs per Client 193
- Global Data Manager 289
- Global Logging Level property 367
- Glossary. *See* NetBackup Help.
- goodies directory 450
- GRFS Advertised Name property 312

H

- hard links
 - NTFS volumes 126
 - UNIX directories 114
- HKEYS, backing up 125
- host properties
 - changing in a clustered environment 301
 - permission to change 303

I

- IANA 496
- IDX (index file) 271
- images
 - changing primary copy 246
 - duplicating 241
 - import 249
 - moving client catalog 228
 - on Media report 271
 - restoring from duplicate 246
 - verifying 240
- Import backup images 249
- include
 - files list 143
 - list, on client 342



- incomplete job 287
- Incrementals Based on
 - Archive Bit property 328
 - Timestamp property 328
- index_client command 236
- indexing, image catalog 236
- Information Server daemon 289
- Informix policy type 75
- INI file, for Lotus Notes 369
- initpbdbm 415
- Initial Browse Search Limit property 389
- INITIAL_MEMORY 498
- Inline Tape Copy option 166, 241
- installing client software
 - on PC clients 101
 - on secure clients 101
 - on trusting clients 99
- Instant Recovery
 - Advanced Backup Method 79
 - Backups to Disk Only setting 164
- Intelligent Disaster Recovery
 - collect information for 84, 93
- Internet Assigned Numbers Authority (IANA) 458

J

- Java
 - auth.conf file 492
 - authorizing users 491
 - directory 493
 - jbp.conf file 500
 - jbpSA configuration options 500
 - jnb.conf file 500
 - jnbSA configuration options 500
 - performance improvement hints 502
- Java interface 4, 9
- Java Virtual Machine (JVM) 498
- jnbSA 4, 9
- Jobs
 - (see Activity monitor)
- jobs
 - concurrent per disk storage unit 49
 - done 287
 - filters
 - specify 281
 - incomplete 287
 - maximum
 - per client 363
 - per policy 81

- priority for policy 82
 - re-queued 287
- JVM (Java Virtual Machine) 498

K

- Keep Logs For property 261
- Keep Logs property 364
- Keep Status of User-directed Backups, Archives, and Restores property 323, 330
- Keep True Image Restoration Information property 364
- keyword phrase 96

L

- last media used, catalog backups 205
- license keys 416
 - accessing 416
 - adding 417
 - deleting 417
 - export 418
 - using the NetBackup License Key utility 419
 - viewing the properties of one key 418
- limit fragment size 50
- Limit Jobs per Policy setting 81, 193
- links
 - UNIX hard-linked directories 114
 - UNIX symbolic 113
- load balancing 485
- Locked File Action property 324
- Logging enabled for debug 367
- logs
 - bpcoord 163
 - bpsynth 163
 - deleting after a set time 364
 - progress, for user operations 10
- Lotus Notes
 - host properties
 - Executable Directory 369
 - INI File 369
 - Path 369
 - policy type 74
 - properties 369

M

- Mac OS X 75
- mail notifications
 - administrator E-mail address 393
 - E-mail address for administrator 366



-
- Windows nbmail.cmd script 366
 - Mailbox for Message Level Backup and Restore property 334
 - manual backups
 - NetBackup catalogs 215
 - policy for 197
 - master servers, rebooting 413
 - MAX_MEMORY 498
 - maximum
 - jobs per client
 - specifying 363
 - jobs per policy 81
 - Maximum Backup Copies property 364
 - Maximum Concurrent Drives Used for Backup setting 49
 - Maximum Concurrent Jobs
 - disk storage unit 49
 - Maximum Data Streams property 320
 - Maximum Jobs per Client property 363
 - Maximum Repetitive Error Messages for Server property 330
 - Media
 - host properties
 - Allow Backups To Span Media 371
 - Allow Media Overwrite 370
 - Allow Multiple Retentions Per Media 371
 - Enable Job Logging 372
 - Enable SCSI Reserve/Release 372
 - Enable Standalone Drive Extensions 372
 - Media ID Prefix (Non-robotic) 372
 - Media Request Delay 373
 - Media Unmount Delay 372
 - media
 - 1 and media 2, catalog backup 205
 - active 273
 - freeze 268
 - ID for catalog backup 207
 - last used for catalog backup 205
 - nonactive 273
 - type for catalog backup 206
 - unfreeze 268
 - unsuspend 268
 - Media Contents report 270
 - Media ID Prefix (Non-robotic) property 372
 - Media List report 267
 - Media Log Entries report 273, 364
 - Media Manager overview 31
 - media mount
 - errors 292
 - timeout for Storage Migrator 487
 - Media Mount Timeout property 387
 - Media Request Delay property 373
 - Media Server Connect Timeout
 - property 388
 - Media Server Copy Advanced Backup Method 79
 - media servers
 - adding a media server to the Alternate Restore Failover Machine list 379
 - rebooting 413
 - Restore Failover host properties 378
 - Media Summary report 273
 - Media Unmount Delay property 372
 - Media Written report 274
 - Megabytes of Memory property 324
 - MEM_USE_WARNING 498
 - Microsoft Volume Shadow Copy Service (VSS) 407
 - mntfs file system 130, 140
 - monitoring NetBackup processes 291
 - monthly backups, scheduling 178
 - mount points 84
 - Move Backup Job from Incomplete to Done
 - State property 80
 - move detection 87
 - Move Job From Incomplete State to Done
 - State property 365
 - Move Restore Job from Incomplete State to Done State
 - interaction with Checkpoint Restart 446
 - MS-Exchange policy type 74
 - MS-SharePoint policy type 75
 - MS-SQL-Server policy type 74
 - MS-Windows-NT policy type 74
 - MTF1 format, allow overwrite 370
 - mtime 120
 - multiple copies
 - setting 166
 - using Checkpoint Restart 79
 - multiple data streams
 - allowing 93, 95
 - tuning 95
 - multiple file layout for NetBackup
 - catalogs 231
 - multiplexing (MPX)
 - and synthetic backups 160



set for schedule 171
multistreaming and synthetic backups 160
Must Use Local Drive property 359

N

named data streams
 disabling the restore of 121
nbbdbd daemon 289
NBJAVA_CLIENT_PORT_WINDOW 499
NBJAVA_CONNECT_OPTION 499
nbmail.cmd script 366
NCR-Teradata policy type 74
NDMP 28, 37, 53, 74, 167, 244, 360
NetBackup
 administration
 using the Java interface 9
 client service 374
 configuring properties 13
 request service port (BPRD) 374
NetBackup Access Control (NBAC)
 authorizing NetBackup-Java users 491
 Current NBAC User 23
NetBackup Request Service Port (BPRD)
 property 374
NetBackup-Java, set up for 5
NetWare Client
 host properties
 Backup Migrated Files 323
 Keep Status of User-directed
 Backups, Archives, and Restores 323
 Uncompress Files Before Backing
 Up 323
NetWare NonTarget clients 335
NetWare policy type 74
Network
 host properties
 Announce DHCP interval 374
 NetBackup Client Service Port
 (BPCD) 374
 NetBackup Request Service Port
 (BPRD) 374
 loading 363
 mask for VxSS host or domain 306
NEW_STREAM, file list directive 134
nonactive media 273
non-reserved ports 392, 461
nonroot administration
 specific applications 494
number of drives, setting for storage

units 49

O

obsolescence date 487
On Demand Only storage unit setting 51
online_util directory for VSP logging
 messages 397
Open File Backup properties 375
Open Transaction Manager (OTM)
 properties 375
Operator's Email Address property 316
optical devices 167, 244
Oracle policy type 74, 75
OTM (see Open Transaction Manager)
Override Policy
 Storage Unit setting 169
 Volume Pool setting 169

P

path setting (Lotus Notes) 369
pathname
 catalog backup to disk 208
 rules for policy file list 111
PC NetLink files 112
peername, files 433
Perform Default Search of Backup Images
 for Restore property 330
Perform Snapshot Backups 164
performance
 improvement, Java applications 502
performance, reducing search time 236
permission to change NetBackup
 properties 303
Persistent Storage daemon 289
planning
 storage units 33
 user schedules 193
policies
 activating 82
 configuration wizard 61
 creating policy for Vault 196
 example 62
 overview 10, 57
 planning 63
 setting priority 82
 storage unit 76
 user 194
 user schedules 193
 volume pool setting 77
policy type



-
- AFS 75
 - DataStore 74
 - DB2 74
 - FlashBackup 75
 - FlashBackup Windows 75
 - Informix 75
 - Lotus-Notes 74
 - MS-Exchange 74
 - MS-SharePoint 75
 - MS-SQL-Server 74
 - MS-Windows-NT 74
 - NCR-Teradata 74
 - NDMP 74
 - NetWare 74
 - Oracle 74
 - SAP 75
 - Split-Mirror 75
 - SQL-BackTrack 75
 - Standard 75
 - Sybase 75
 - Vault 75
 - Port Ranges
 - host properties
 - Client Port Window 376
 - Client Reserved Port Window 377
 - Server Port Window 377
 - Server Reserved Port Window 377
 - Use OS selected non reserved port 376, 377
 - Use Random Port Assignments 376
 - ports
 - allowing operating system to select non reserved port 376, 377
 - configuring limitations 471
 - minimum required by NetBackup by firewall 347
 - non-reserved 392, 461
 - reserved 461
 - power down NetBackup servers 412
 - preprocess interval 136
 - primary copy
 - becoming a 244
 - changing 246
 - definition 244
 - promoting to 247
 - priority, for a policy 82
 - Problems report 265, 364
 - proc file system 130, 140
 - Process Busy Files property 316
 - processes
 - monitor 291
 - monitoring 291
 - show active 413
 - progress logs, client 10
 - properties, overview 300
 - Q**
 - Quiescent wait time 403
 - R**
 - random ports, setting on server 376
 - raw partition backups
 - on UNIX 116
 - relative speed on UNIX 117
 - when to use on UNIX 117
 - raw partitions
 - backing up 80, 124
 - restoring 124
 - rebooting
 - NetBackup servers 413
 - redirected restores 117
 - registry, back up on Windows clients 125
 - registry, restore on Windows clients 125
 - relocation schedule 167
 - remote systems, administering 425
 - reports
 - All Log Entries report 266
 - Client Backups report 264
 - description of console 258
 - Images on Media report 271
 - Media Contents report 270
 - Media List report 267
 - Media Log Entries report 273
 - Media Summary report 273
 - Media Written report 274
 - Problems report 265
 - running a report 258
 - settings for building a report 261
 - Status of Backups report 263
 - using the Troubleshooter 275
 - Requeue Active Jobs if Required Storage Space is Unavailable property 388
 - Requeue Active Jobs property 388
 - Requeue Scheduled Jobs if Required Storage Space is Unavailable property 388
 - Requeue Scheduled Jobs property 388
 - re-queued job 287
 - Re-read Interval for Available Drives property 359



- reserved ports 461
- Reset File Access Time property 324
- restart
 - completed jobs 282
- Restore Failover
 - host properties
 - Alternate Restore Failover Machines list 378
- Restore job
 - resuming 446
 - suspending 446
- Restore Retries
 - interaction with Checkpoint Restart 446
 - property 389
- restores
 - catalog backups 202
 - caution for alternate client 432
 - caution for alternate path 434
 - directed from the server 431
 - overview 3, 10
 - raw partition 124
 - reducing search time 236
 - registry on Windows clients 125
 - server independent 450
 - setting client permissions 441
 - symbolic links on UNIX 113
 - System State 447
 - to alternate clients 432
- restoring files to alternate hosts 378
- resume
 - restore jobs 282
- Retain Snapshots for Instant Recovery 164
- retention levels
 - default 170
 - for archiving catalogs 232
- retention periods
 - caution for setting 193
 - changing 381
 - mixing on media 171
 - precautions for setting 170
 - redefining 380
 - setting 169
 - user schedule 193
- Retries Allowed After Runday policy
 - setting 165
- Retry Count property 317
- retry restores, setting 389
- right-clicking using Solaris X86 22
- Rmed media type 272

- RS-MTF1 format, allow overwrite 370

S

- SANPoint Control (SPC)
 - host properties 383
 - SPC Server Name 383
 - SPC WebServer Name 383
- SAP policy type 75
- Schedule Backup Attempts property 94, 362
- schedules
 - adding to policy 145
 - backups on specific dates 176
 - catalog backup 209
 - examples of automatic 180
 - frequency 165
 - how calendar scheduling interacts with daily windows 179
 - monthly backups 178
 - naming 146
 - overview 10, 60
 - priority 165
 - retention level defaults 170
 - retention periods
 - setting 169
 - setting backup times 173
 - specify multiplexing 171
 - storage unit 169
 - type of backup 146
 - user backup or archive 193
 - volume pool 169
 - weekly backups 177
- scratch volume pool 78
- scripts
 - bpdbjobs example 295
 - bps 413
 - goodies 450
 - initbdbm 415
- SeCuRiT.y.nnnn files 440
- server
 - directed restore 431
 - independent restores 378, 450
 - managing storage units 8
 - power down 412
 - properties 385
 - rebooting 412
 - software 2
 - speed 363
- server list definition 385
- Server Port Window property 377



-
- Server Reserved Port Window property 377
 - Server Sends Mail property 393
 - servers
 - NetBackup
 - using Storage Migrator 486
 - Servers tab 385
 - SGI cachefs file system 130, 140
 - Shadow Copy 407
 - single file layout for NetBackup catalogs 230
 - single file restore program, FlashBackup 113
 - Single-Instance Storage (SIS) 80, 334
 - Solaris
 - 9 extended attributes 113
 - file systems 130, 140
 - X86, enabling right-click 22
 - sort column data 19
 - Source Copy Number 239
 - speed, of server and client 363
 - Split-Mirror policy type 75
 - SQL-BackTrack policy type 75
 - Standard policy type 75
 - Status of Backups report 263
 - Status Report Interval property 362
 - Storage Migrator 50, 486
 - storage unit groups
 - changing 55
 - creating 54
 - deleting 56
 - storage units
 - adding Media Manager type 42, 45
 - adding NDMP type 37
 - changing server to manage 8, 420
 - disk type, definition 28
 - example Media Manager type 33
 - for policy 76
 - for schedule 169
 - management window 30
 - Media Manager type, definition 28
 - NDMP 28, 37
 - next available 76
 - optical devices 167, 244
 - overview 12
 - QIC drive type 167, 244
 - rules for Media Manager type 31
 - streaming (see Allow Multiple Data Streams setting)
 - Sun PC NetLink 112
 - suspend
 - restore jobs 282
 - Sybase policy type 75
 - symbolic links
 - included in backup selection list 109
 - UNIX 113
 - synthetic backups
 - and encryption 161
 - component images 157
 - cumulative incremental 158
 - full 157
 - logs produced during 163
 - no NetBackup change journal support 328
 - recommendations for running 160
 - schedules 154
 - using Checkpoint Restart 80
 - System State
 - backups 80
 - restoring 447
 - System_State directive 131
- T**
- tar format, allow overwrite 370
 - TCP Level Logging property 326
 - third-party copies 167, 244
 - Third-Party Copy Device Advanced Backup Method 79
 - Time Overlap property 329
 - Timeout in Job Queue property 388
 - Timeouts
 - client read 387
 - host properties
 - Backup End Notify Timeout 387
 - Backup Start Notify Timeout 386
 - BPTM (Drive Count) Query Timeout 388
 - Client Connect Timeout 386
 - File Browse Timeout 386
 - Media Mount Timeout 387
 - Media Server Connect Timeout 388
 - Requeue Active Jobs if Required Storage Space is Unavailable 388
 - Requeue Scheduled Jobs if Required Storage Space is Unavailable 388
 - timeouts
 - host properties
 - Timeout in Job Queue 388
 - To IP Address property 314
 - traversing directories to back up a file 342
 - Troubleshooter



- using in Activity Monitor 283
- using in Reports application 275
- True Image Restoration (TIR)
 - configuration 87
 - Error code 136, TIR info was pruned from the image file
 - error codes 163
 - length of time to keep information 364
 - move detection 87
 - no NetBackup change journal
 - support 328
 - pruning information 162
 - with move detection 328
 - with Move Detection property 162
- type of backup
 - for archiving catalogs 232
 - setting 146

U

- uncompress
 - client records 228
 - NetBackup catalogs 227
- Uncompress Files Before Backing Up
 - property 323
- unfreeze media 268
- Universal
 - host properties
 - Allow Server File Writes 392
 - Use Preferred Group for Enhanced Authorization 391
- Universal Settings properties 389
- UNIX Client
 - host properties
 - Add to All Lists 325
 - Do Not Compress Files Ending With 325
 - Do Not Reset File Access Time 324
 - Megabytes of Memory 324
 - primary node in tree 394
- UnixWare cache file system 130, 140
- UNSET, file list directive 139
- UNSET_ALL, file list directive 139
- unsuspend media 268
- Use Case Sensitive Exclude List host
 - property 335
- Use Change Journal in Incrementals
 - property 327
- Use Direct Access Recovery for NDMP
 - Restores property 360

- Use Media Host Override for Restores
 - property 360
- Use Non Reserved Ports property 345
- Use OS Selected Non-reserved Port
 - property 358
- Use Preferred Group for Enhanced Authorization property 391
- Use Random Port Assignments
 - properties 376
- Use Reserved Ports property 345
- Use Specified Network Interface
 - property 390
- USE_CTIME_FOR_INCREMENTALS 487
- user
 - archive backups 147
 - backups 147
 - backups, archives, restores 10
 - schedules
 - planning 193
- User Directed Timeout property 329

V

- Vault 287
 - backup type 148
 - catalog archiving 235
 - designating duplicate as the primary 246
 - Logging property 368
 - policy
 - creating 196
 - type 75
- verifying backup
 - images 240
 - selections list 109
- VERITAS Products properties 395
- version file xxxiii
- view properties of a license key 418
- visd daemon 289
- vnetd
 - Only property (for selection of ports) 345
 - VERITAS Network Daemon 320, 344
- VNETD_PORT 496
- Volume Manager (VxVM) 116
- volume pools
 - for schedule 169
 - indicating one for use by a policy 77
 - policy 77
 - scratch 78
- Volume Shadow Copy 131
- Volume Shadow Copy Service (VSS) 407



- Volume Snapshot Provider (VSP)
 - backups using Checkpoint Restart 79
 - directory for logging messages 397
 - properties 396
 - using with databases 403
- volumes
 - allocation 77
 - assignments 77
 - scratch 77
- VSP (see VERITAS Volume Snapshot Provider) 403
- VxFS 4.0, named data streams 119
- VxSS
 - domain, indicating network mask 306
 - Networks List property 305

W

- Wait Time Before Clearing Archive Bit property 326
- Wakeup Interval property 362
- weekly backups, scheduling 177
- wildcard characters
 - escaping backslash 340
 - escaping on UNIX 112
 - in exclude files lists 340
 - in exclude lists 140
 - UNIX
 - escape character 141
 - file paths 111
 - Windows clients 122
- Windows Client
 - host properties
 - Communications Buffer 329
 - General Level Logging 326
 - Incrementals Based on Archive Bit 328
 - Incrementals Based on
 - Timestamp 328
 - Keep Status of User-directed Backups, Archives, and Restores 330
 - Maximum Repetitive Error Messages for Server 330
 - Perform Default Search of Backup Images for Restore 330
 - TCP Level Logging 326
 - Timeout Overlap 329
 - Use Change Journal in Incrementals 327
 - User Directed Timeout 329
 - Wait Time Before Clearing Archive Bit 326
- Windows Disk-Image (raw) backups 80, 124
- Windows Display Console 7, 425, 427, 480
- Windows Open File Backups
 - host properties
 - Abort Backup on Error 408
 - Disable Snapshot and Contine 409
 - Enable Windows Open File Backups for this client 407
 - Global Drive Snapshot 408
 - Individual Drive Snapshot 407
 - Use Microsoft Volume Shadow Copy Service (VSS) 407
 - Use VERITAS Volume Snapshot Provider (VSP) 407
- wizards
 - backup policy 61
 - catalog backup 203
 - Device Configuration 31
- Working Directory property 316
- WORM media
 - retention period caution 171



