

VERITAS NetBackup™ DataCenter 3.4

System Administrator's Guide

UNIX

June 2000
100-001459


VERITAS

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 © 1993-2000 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo, VERITAS NetBackup, and VERITAS NetBackup BusinessServer are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

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

Printed in the USA, June 2000.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-335-8000
Fax 650-335-8050
www.veritas.com



Contents

Audience	xxx1
Organization	xxx1
Related Documents	xxxii
Conventions	xxxiii
Type Style	xxxiii
Notes and Cautions	xxxiii
Key Combinations	xxxiii
Command Usage	xxxiv
Getting Help	xxxiv
Chapter 1. Introduction to NetBackup	1
Overview	1
Media Manager	3
Storage Units	3
Volumes	3
Backup Policies (Classes)	4
User Backups, Archives, and Restores	5
NetBackup Catalog Backups	5
NetBackup Administration Interfaces	5
NetBackup Administration - Java Interface	6
Setting Up Your Window Manager for the Java Interface	7
To Start the NetBackup-Java Interface on a UNIX System	8
To Start the NetBackup-Java Interface on Windows Systems	10
NetBackup Administration - Java Window	11



Configuring NetBackup	14
Chapter 2. Managing Storage Units	17
Introduction to Storage Units	17
Using the Configure Storage Devices Wizard	18
To Start the Storage Unit Management Utility	18
The Storage Unit Management Window	18
Tree and Detail Views	19
The View if You Are Not on a Master of Masters	19
The View From a Master of Masters (requires Global Data Manager)	20
Menu Bar	21
Toolbars	22
Shortcut Menus	22
Choosing the Master Server for Storage Units	22
If You Are Not on a Master of Masters	22
If You Are On a Master of Masters (requires Global Data Manager)	23
If You Cannot Change to Another Master Server	23
Adding Media Manager Storage Units	23
Rules for Media Manager Storage Units	23
Before Adding the Storage Unit	24
Example 1	25
Example 2	26
Information Required for Storage Unit on whale	26
Information Required for Storage Units on shark	27
To Add a Media Manager Storage Unit	27
General Properties - Media Manager Storage Units	29
Storage unit type	29
On demand only	29
Storage Unit Type Properties - Media Manager Storage Units	30
NetBackup media server	30



Robot Type	30
Robot Number	30
Density	31
Maximum concurrent drives used for backup	31
Maximum fragment size	31
Maximum multiplexing per drive	31
Adding Disk Type Storage Units	32
To Add a Disk Type Storage Unit	32
General Properties - Disk Storage Units	33
Storage unit type	33
On demand only	33
Storage Unit Type Properties - Disk Storage Units	34
NetBackup media server	34
Absolute pathname to directory	34
Maximum concurrent jobs	34
Maximum fragment size	35
Changing Storage Unit Attributes	35
Deleting Storage Units	35
Automatic Drive Availability Checking	36
Interval Between Status Checks	36
Drive Count Timeout	36
Requeuing Jobs If Required Storage Units are Unavailable	37
Chapter 3. Configuring Classes	39
Introduction to NetBackup Classes	40
Parts of a Class Definition	40
General Attributes	40
Client List	40
File List	40
Schedules	41



Example Classes	41
Planning Guidelines for Classes	43
Planning Worksheet	48
Using the Backup Policy Wizard	53
To Start the Backup Policy Management Utility	53
The Backup Policy Management Window	53
Tree and Detail Views	54
The View if You Are Not On a Master of Masters	54
The View From a Master of Masters (requires Global Data Manager)	55
Menu Bar	56
Toolbar	57
Shortcut Menus	57
Choosing the Master Server for Backup Policies	57
If You Are Not On a Master of Masters	58
If You Are On a Master of Masters (requires Global Data Manager)	58
To Add a New Class	58
To Change a Class	61
Setting the General Class Attributes	63
Class type	63
Class storage unit	65
Example	66
Notes on Specifying a Storage Unit	66
Class volume pool	67
Example	67
Notes on Volume Pools	67
Limit jobs per class	68
Notes on Limit jobs per class	68
Job priority	69
Keyword phrase (optional)	69
Active	69



True image restore information	70
With move detection	70
Example of What Happens During True-Image Restores	71
Notes On True-Image Restores and Move Detection	72
Backup network drives	73
Follow NFS	75
Notes on Follow NFS	75
Disadvantages of Following NFS Mounts	76
Advantages of Following NFS Mounts	76
Cross mount points	77
Notes on Cross mount points	77
Cases That Can Require Separate Classes	77
How Cross Mount Points Interacts With Follow NFS	78
Cross Mount Point Examples	78
Compression	79
Advantages of Using Compression	79
Disadvantages of Using Compression	80
How Much Compression Can You Expect?	80
Encryption	81
Block level incremental	82
Allow multiple data streams	82
When to Use Multiple Data Streams	82
Tuning Multiple Data Streams	84
Individual file restore from raw	85
Collect disaster recovery information	86
To Add Clients to a Class	86
Installing Client Software on Trusting UNIX Clients	88
Installing Software on Secure UNIX Clients	90
Installing Software on PC Clients	90
To Add Files For Automatic Backups	90



Adding Directives to the File List	91
To Add Directives In the Same Way as Path Names	92
To Select the Directives From Templates	93
Verifying the File List	94
Rules for Backup File Paths	95
File-Path Rules for Standard UNIX and Apollo wbak Clients	95
Notes on File Lists for Standard UNIX and Apollo wbak Clients	96
Symbolic Links to Files or Directories	97
Hard Links to Directories	98
Hard Links to Files	98
UNIX Raw Partitions	100
File-Path Rules for Auspex-FastBackup UNIX Clients	102
File-Path Rules for Microsoft Windows Clients	102
File Backups - Microsoft Windows Conventions	102
File Backups - UNIX Conventions	103
Windows NT/2000 Disk-Image (raw) Backup	103
Microsoft Windows Registry Backup	104
Hard Links to Files (NTFS volumes only)	105
File-Path Rules for OS/2 Clients	107
File-Path Rules for NetWare NonTarget Clients	107
File-Path Rules for NetWare Target Clients	108
File-Path Rules for Macintosh Clients	109
File-Path Rules for Extension Clients	110
File-List Directives - General Discussion	111
ALL_LOCAL_DRIVES Directive	111
Directives for Multiple Data Streams	111
Directives for Specific Class Types	112
File-List Directives for Multiple Data Streams	112
NEW_STREAM	113
ALL_LOCAL_DRIVES	116



UNSET and UNSET_ALL	117
Excluding Files From Automatic Backups	119
Creating an Exclude List on a UNIX Client	119
Creating an Include List on a UNIX Client	122
To Add New Schedules	123
Schedule Properties	124
Name	124
Type of backup	124
Retention	126
Frequency	128
Override class storage unit	129
Override class volume pool	129
Media multiplexing	129
Schedule	130
Examples of Automatic-Backup Schedules	132
Example 1	132
Example 2	134
Example 3	136
Example 4	136
Example 5	137
Example 6	137
Considerations for User Schedules	138
Planning User Backup and Archive Schedules	139
Creating Separate Classes for User Schedules	139
To Use a Specific Class and User Schedule	141
Performing Manual Backups	141
Chapter 4. Backing Up the NetBackup Catalogs	143
Introduction to the NetBackup Catalogs	143
Where are the Catalog Files?	143



What Method Do I Use to Back Them Up?	144
What NetBackup Servers Can I Use?	144
What Types of Media Can I Use?	144
How Do I Know If a Catalog Backup Succeeded?	144
How Do I Restore The Catalog Backups?	144
Important Precautions to Observe	144
Using the Catalog Backup Wizard	145
Configuring NetBackup Catalog Backups	146
Attributes Tab	147
Media Server	147
Last Media Used	147
Media 1 and Media 2 Areas	148
Schedule Tab	151
Files Tab	153
Adding a New Pathname	153
Changing a Path	154
Deleting a Path	154
Catalog Pathnames	154
Backing Up the Catalogs Manually	156
To Perform a Manual Backup of the Catalogs	156
Protecting Large NetBackup Catalogs	157
The Layout of the NetBackup Catalogs	157
Backup and Restore Concept	158
Multiple-Tape Backups	158
Multiple-Tape Restores	158
Setting up Multiple-Tape NetBackup Catalog Backups	158
Define a NetBackup Class	159
Configure the NetBackup Catalog Backups	159
Create a Shell Script to Initiate the Backups	160
How To Initiate a Multiple-Tape Catalog Backup	161



Managing the NetBackup Catalogs	161
Determining Catalog Space Requirements	162
Compressing the Image Catalog	165
Uncompressing the Image Catalog	166
Moving the NetBackup Image Catalog	166
Chapter 5. Viewing NetBackup Reports	169
To Start the Reports Utility	170
The Reports Window	170
Report Contents	170
Menu Bar	173
Toolbar	173
Shortcut Menus (Global Data Manager only)	174
Choosing the Master Server for the Reports	174
If You Are Not on a Master of Masters	174
If You Are on a Master of Masters (requires Global Data Manager)	174
Running a Report	175
Report Settings	175
Date and Time	176
Client	176
Server	176
Media ID	176
Volume Pool	176
Verbose Listing	176
Report Descriptions	177
Backup Status Report	177
Client Backups Report	177
Problems Report	178
All Log Entries Report	179
Media Lists Report	179



Media Contents Report	181
Images on Media Report	182
Media Logs Report	183
Media Summary Report	183
Media Written Report	184
Chapter 6. Monitoring NetBackup Activity	185
To Start the Activity Monitor	186
The Activity Monitor Window	186
Jobs List	186
The View If You Are Not on a Master of Masters	186
The View From a Master of Masters (requires Global Data Manager)	188
Customizing the Jobs List Output	189
Menu Bar	192
Toolbar	193
Status Bar	193
Shortcut Menus (Global Data Manager only)	193
Refreshing the Display	194
Choosing the Master Server to Monitor	195
If You Are Not on a Master of Masters	195
If You Are on a Master of Masters (requires Global Data Manager)	195
Monitoring Jobs	196
To Monitor Jobs Without Global Data Manager	196
To Monitor Jobs With Global Data Manager	196
To Delete Done Jobs	198
To Kill Uncompleted Jobs	198
To Save Job Information	199
To Monitor Detailed Status on Backup or Archive Jobs	199
Job Details Dialog Box Description	200
Using the Troubleshooting Wizard	202



Managing the Jobs Database	203
Creating a job.conf File	203
Specifying a BPDBJOBS_OPTIONS Environment Variable	205
Using bpdjobs Command Line Options	205
bpdjobs Activity Log	206
Chapter 7. Managing NetBackup	207
Powering Down and Rebooting NetBackup Servers	208
Powering Down a NetBackup Server	208
Rebooting a NetBackup Server	208
Master Server	208
Media Server	208
Managing Daemons	208
Starting and Stopping bprd	208
Starting bprd	209
Stopping bprd	209
Starting and Stopping bpdbm	210
Displaying Active Processes With bpps	210
Managing the Restore of Client Files	210
Allowing Restores to an Alternate Client	211
How NetBackup Enforces Restore Restrictions	211
Allowing All Clients to Perform Alternate-Client Restores	212
Allowing a Single Client to Perform Alternate-Client Restores	212
Allowing Alternate-Client Restores of Specific Client's Files	213
Examples of Restoring Files to an Alternate Client	213
Setting Client List and Restore Permissions	217
Adding Clients to the NetBackup Client Database	217
Setting the List and Restore Permissions	218
Examples	220
Reduce Restore Times by Indexing the Image Catalog	221



Catalog Index Examples	222
Catalog Index Space Requirements	222
Disabling Catalog Indexing	223
Improve Search Times by Creating an Image List	223
Server-Directed Restores	223
Set Original atime for Files During Restores	223
Administering NetBackup Licenses	224
Duplicating, Verifying, and Importing Backup Images	225
Duplicating Backup Images	225
To Duplicate Backup Images	226
Duplication Results Dialog Box	230
Restoring from Duplicated Backups	231
Notes on Multiplexed Duplication	232
Verifying Backup Images	232
Verification Results Dialog Box	235
Importing Backup Images	236
Notes About Importing Backup Images	236
Import Results Dialog Box	242
Goodies Scripts	242
Server Independent Restores	243
Supported Configurations	243
Methods for Performing Server Independent Restores	246
Method 1 -- Modifying the NetBackup Catalogs	246
Method 2 -- Overriding the Original Server	247
Method 3 -- Automatic Failover to Alternate Server	248
Related Topics	249
Load Balancing	250
Adjust Backup Load on Server	250
Adjust Backup Load on Server Only During Specific Time Periods	250
Adjust Backup Load on Client	250



Reduce Time To Back Up Clients	251
Give Preference To a Class	251
Adjust Load Between Fast and Slow Networks	251
Limit the Backup Load Produced By One or More Clients	251
Maximize Use of Devices	251
Prevent Backups From Monopolizing Devices	251
Using NetBackup with Storage Migrator	251
Set a Large Enough Media Mount Timeout	252
Do Not Use the RESTORE_ORIGINAL_ETIME File	252
Do Not Use the Following Client bp.conf File Settings	253
Allowing Nonroot Users to Administer NetBackup	253
Configuring Nonroot Usage of All NetBackup Administrator Applications	253
Creating a Group Specifically for Media Manager Tape Users	255
Authorizing NetBackup-Java Users	256
Authorization File	256
Configuring Nonroot Usage	258
All NetBackup-Java Applications	258
Authorizing Nonroot Users for Specific Applications	258
Capabilities Authorization for jbpSA	259
Converting Previous auth.conf Files	260
Configuration Options for jbpSA	260
Default NetBackup Host to Log Into	260
Changing Fonts and Font Sizes on UNIX Platforms	261
NetBackup-Java Performance Improvement Hints	263
Administrator's Quick Reference	264
Chapter 8. Additional Configuration	267
Defining NetBackup Global Attributes	268
NetBackup Global Attributes	269
E-mail address for notifications	269



Maximum jobs per client	269
Wakeup interval	271
Schedule backup attempts	271
Duration to retain logs	271
Media mount timeout	271
Interval for status reports	272
Delay to compress catalog	272
How long to keep TIR information	272
Redefining Retention Periods	273
Note on Redefining Retention Periods	275
Multiplexing	275
When to Use Multiplexing	276
How to Configure Multiplexing	276
Maximum Multiplexing Per Drive for Storage Unit	277
Media Multiplexing for a Schedule	277
Other Configuration Settings to Consider	279
Demultiplexing	280
Using Multiple NetBackup Servers	280
Configuring a Master and Media Server Cluster	281
Software on Each Server	283
NetBackup Catalogs	283
NetBackup and Media Manager Daemons	283
To Add Media Servers	284
Dynamic Host Name and IP Addressing	286
Setting up Dynamic IP Addresses and Host Names	288
Configuring the NetBackup Master Server	289
Configuring a Dynamic Microsoft Windows Client	291
Configuring a Dynamic Macintosh NetBackup Client	291
Configuring a Dynamic UNIX NetBackup Client	292
Bandwidth Limiting	293



Read This First	293
How Bandwidth Limiting Works	294
Configuration	294
Rules for IP Address Ranges	294
Rules for Setting Bandwidth Values	296
Examples	296
Example 1	296
Example 2	297
Example 3	297
Busy-File Processing (UNIX Clients Only)	297
Getting Started	298
Modifying bp.conf	298
BUSY_FILE_PROCESSING	298
BUSY_FILE_DIRECTORY	298
BUSY_FILE_ACTION	298
Examples	299
Creating Action Files	300
Logs Directory	301
Modifying bpend_notify_busy	302
Open Transaction Management (Microsoft Windows Clients Only)	303
Enhanced Authentication	303
Configuration Files	304
methods.txt	304
methods_allow.txt	305
methods_deny.txt	306
names_allow.txt	306
names_deny.txt	307
Library Files	308
Authentication Commands	308
bpauthsync	308



vopie_util	308
vopied Daemon	308
vopie Files	309
hashed (public key) Files	309
unhashed (secret key) Files	309
temp File	310
To Configure Authentication	311
Example 1 - Typical Configuration	311
Example 2 - Disable Authentication for a Client	312
Example 3 - Add a Client	312
Example 4 - Restore Authentication After Client Disk Crash	313
Example 5 - Restore Authentication on NetBackup Master Server	314
Troubleshooting Authentication	315
NetBackup Authorization	315
Configuring NetBackup Authorization	316
Step 1: Add NetBackup Servers To Each Others Server Lists	316
Step 2: Enable NetBackup Authentication	316
Step 3: Create an authorize.txt file	316
Step 4: Specify Preferred Group on Administrative Console (optional)	318
Example Configuration	319
NetBackup User Authorization Process Description	320
Gaining Access to a Server	320
Gaining Access to a Client	321
Configuring Email Notifications	321
Adjust Time Zone	322
Specifying the Locale of the NetBackup installation	323
For NT/98/2000 platforms:	324
For Macintosh platforms:	324
For UNIX platforms:	324
TL lines	324



TM lines	325
NetBackup Configuration Options	325
Method for Specifying the Configuration Options	325
Syntax Rules for bp.conf Options	326
bp.conf Options for UNIX Servers	326
ALLOW_MEDIA_OVERWRITE	327
ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA	328
ALLOW_NON_RESERVED_PORTS	328
APOLLO_RESTORE_TIMEOUT	328
BPEND_TIMEOUT	329
BPSTART_TIMEOUT	329
BPTM_QUERY_TIMEOUT	329
CLIENT_CONNECT_TIMEOUT	330
CLIENT_PORT_WINDOW	330
CLIENT_READ_TIMEOUT	330
CLIENT_RESERVED_PORT_WINDOW	331
DISABLE_JOB_LOGGING	331
DISABLE_STANDALONE_DRIVE_EXTENSIONS	332
DISALLOW_BACKUPS_SPANNING_MEDIA	332
DISALLOW_CLIENT_LIST_RESTORE	332
DISALLOW_CLIENT_RESTORE	332
FAILOVER_RESTORE_MEDIA_SERVERS	333
FORCE_RESTORE_MEDIA_SERVER	333
INITIAL_BROWSE_SEARCH_LIMIT	334
KNOWN_MASTER	334
LIMIT_BANDWIDTH	335
MASTER_OF_MASTERS	336
MAX_APOLLO_RESTORE_ARG_CHARS	336
MEDIA_ID_PREFIX	337
MEDIA_UNMOUNT_DELAY	337



MEDIA_REQUEST_DELAY	337
MPX_RESTORE_DELAY	338
MUST_USE_LOCAL_DRIVE	338
QUEUE_ON_ERROR	339
RANDOM_PORTS	339
RE_READ_INTERVAL	339
REQUIRED_INTERFACE	340
SERVER	341
SERVER_PORT_WINDOW	342
SERVER_RESERVED_PORT_WINDOW	342
SLAVE_CONNECT_TIMEOUT	343
TIMEOUT_IN_QUEUE	343
VERBOSE	343
WAIT_IN_QUEUE	343
bp.conf File on UNIX Clients	344
ALLOW_NON_RESERVED_PORTS	344
BPARCHIVE_CLASS	345
BPARCHIVE_SCHED	345
BPBACKUP_CLASS	345
BPBACKUP_SCHED	346
BUSY_FILE_ACTION	346
BUSY_FILE_DIRECTORY	347
BUSY_FILE_NOTIFY_USER	347
BUSY_FILE_PROCESSING	347
CLIENT_NAME	348
CLIENT_PORT_WINDOW	348
CLIENT_READ_TIMEOUT	348
CLIENT_RESERVED_PORT_WINDOW	349
COMPRESS_SUFFIX	349
CRYPT_OPTION	349



CRYPT_STRENGTH	350
CRYPT_LIBPATH	350
CRYPT_KEYFILE	351
DISALLOW_SERVER_FILE_WRITES	351
DO_NOT_RESET_FILE_ACCESS_TIME	352
INFORMIX_HOME	352
INITIAL_BROWSE_SEARCH_LIMIT	352
KEEP_DATABASE_COMM_FILE	353
KEEP_LOGS_DAYS	353
LIST_FILES_TIMEOUT	353
LOCKED_FILE_ACTION	353
MEGABYTES_OF_MEMORY	354
NFS_ACCESS_TIMEOUT	354
RANDOM_PORTS	355
RESTORE_RETRIES	355
REQUIRED_INTERFACE	355
SERVER_PORT_WINDOW	355
SERVER	355
SYBASE_HOME	356
USE_CTIME_FOR_INCREMENTALS	356
USEMAIL	356
VERBOSE	357
Examples - UNIX Client	357
Example /usr/opensv/netbackup/bp.conf File	357
Example \$HOME/bp.conf File	357
Appendix A. NetBackup Commands	359
bp(1)	360
bpadm(1M)	362
bparchive(1)	363



bpauthsync(1M)	368
bpbackup(1)	371
bpbackupdb(1M)	377
bpclassnew(1M)	381
bpclclients(1M)	388
bpclininclude(1M)	395
bpclinfo(1M)	400
bpclsched(1M)	409
bpclschedrep(1M)	419
bpconfig(1M)	424
bpdbm(1M)	434
bpduplicate(1M)	436
bperror(1M)	442
bpexpdate(1M)	453
bpimagelist(1M)	459
bpimedia(1M)	466
bpimport(1M)	477
bplabel(1M)	482
bplic(1)	484
bpmedia(1M)	490
bpmedialist(1M)	493
bprd(1M)	505
bprecover(1M)	507
bprestore(1)	512
bpstuadd(1M)	520
bpstudel(1M)	526
bpstulist(1M)	528
bpsturep(1M)	532
bpverify(1M)	537
vopied(1M)	544



vopie_util(1M)	546
xbp(1)	550
xbpadm(1M)	552
xbpmon(1M)	554
Appendix B. Using bpadm	555
Starting bpadm	556
Defining and Managing Storage Units	557
Adding a Removable or Robotic Storage Unit	558
Adding a Disk Type Storage Unit	561
Displaying and Changing Storage Unit Configurations	563
Defining and Managing Classes	564
Adding a Class	565
Displaying and Changing Class Configurations	567
Defining and Managing the Client List for a Class	569
Adding Clients to a Class	569
Displaying Client Lists and Deleting Clients from a Class	572
Defining and Managing the File List for a Class	572
Adding to a File List	572
Displaying and Changing a File List	575
Defining and Managing Schedules for a Class	575
Adding a Schedule	576
Displaying and Modifying a Schedule	580
Defining NetBackup Global Attributes	581
Installing NetBackup Software on All Trusting Client Hosts	583
Displaying Reports	584
Backup Status Report	585
Client Backups Report	585
Problems Report	585
All Log Entries Report	586



Media Reports	586
Media List from NetBackup Media Catalog	587
Media Summary	587
Media Contents	587
Images on Media	588
Media Log Entries	588
Media Written	588
Managing bprd (NetBackup Request Daemon)	589
Redefining Retention Levels	590
Performing Manual Backups	592
Backing Up the NetBackup Databases (catalogs)	593
Cautions to Observe for Backing Up NetBackup Databases	594
Listing Database Backup Settings	595
Modify Database Backup Settings	597
Delete Database Backup Media ID	600
Performing Manual Database Backups	600
Adding Database Backup File Paths	601
Removing Database Backup File Paths	602
Appendix C. Reference Topics	603
Rules for Using Host Names in NetBackup	604
Qualifying Host Names	604
How NetBackup Uses Host Names	604
Server and Client Name on UNIX Servers and Clients	604
Host Names on Windows NT/2000 Servers and PC Clients	605
Class Configuration	605
Image Catalog	605
Error Catalog	606
Scheduler	606
Catalog Backup Information	606



How to Update NetBackup After Host Name Changes	606
Special Considerations For Domain Name Service (DNS)	607
Terminal Configuration on UNIX	608
Modifying the XKeysymDB File	608
Modifying the terminfo File	608
Changing X Resources for xbp	609
Reading Backup Images with tar	610
Overview of Auspex FastBackup Classes	613
Functional Overview	614
Auspex Client in a Standard Class	615
Auspex Client in an Auspex FastBackup Class	615
Master and Media Server Configuration	615
Configuring Auspex FastBackup Classes	616
Storage Units	616
Class Attributes	617
File List and Cache	617
Schedules	618
Restoring Files to an Auspex FastBackup Client	618
Factors Affecting Backup Time	618
Total data	619
Transfer rate	619
Compression	620
Device delays	620
Determining NetBackup Transfer Rate	620
Network-Transfer Rate	620
Network-Transfer Plus End-of-Backup-Processing Rate	620
Total-Transfer Rate	621
Examples	621
How NetBackup Builds Its Automatic-Backup Worklist	623
Building the Worklist (Queue)	624



Prioritizing the Worklist	625
Guidelines for Setting Retention Periods	626
Guidelines for Setting Backup Frequency	626
Determining Backup Media Requirements	627
Incremental Backups Overview	628
Retention Requirements	629
Backup and Restore Times	630
Determining Files Due for Backup - Windows NT/2000 Clients	631
Determining Files Due for Backup - UNIX Clients	632
Storage Management Overview	634
Storage Units	634
Media Manager	635
Retention	635
Volume Pools	635
Media Management Concepts	636
NetBackup and Media Manager Catalogs	636
Volume Database	636
Media Catalog	636
Device Catalogs	637
Media States	637
How NetBackup Selects Media in a Robot	638
Spanning Media	639
How NetBackup Uses Media in Standalone Drives	640
Media Selection Using Standalone Drive Extensions	640
Disabling Standalone Drive Extensions	641
Spanning Media	641
Keeping Standalone Drives in the Ready State	641
Media Format	642
Non-QIC Tape Format	642
QIC Tape Format	643



Fragmented Backups	643
Spanning Tapes	644
Multiplexing Format	644
Labeling Media	644
Mounting and Unmounting Media	645
Automatic Media Suspend Or Device Down	645
Appendix D. NetBackup Notify Scripts	647
backup_notify	648
backup_exit_notify	649
bpstart_notify (UNIX clients only)	649
bpstart_notify.bat (Microsoft Windows clients only)	651
bpend_notify (UNIX clients only)	654
bpend_notify.bat (Microsoft Windows clients only)	656
dbbackup_notify	659
diskfull_notify	660
restore_notify	660
session_notify	661
session_start_notify	661
userreq_notify	661
Appendix E. Global Data Manager	663
GDM Terminology	665
Installing GDM	665
Requirements	665
Procedure	665
Configuring GDM	666
Starting GDM	668
Interface Commands for GDM	670
Append Master Server	670
Ignore Master Server(s)	671



Recognize Master Server(s)	672
Appendix F. Using NetBackup With AFS	673
Installation	673
System Requirements	673
Installation on the Server	673
Installation on the Clients	673
Configuration	674
General Class Attributes	674
Client List	674
File List	674
File List Directives	674
Regular Expressions	675
Exclude and Include Lists	676
Backups and Restores	676
Backups	676
Automatic Backup	676
Manual Backup	676
Restores	676
Restore From the NetBackup for AFS Client	677
Restore From the NetBackup Master Server	677
Notes About Restores	677
Troubleshooting	678
Troubleshooting Backups	678
Troubleshooting Restores	678
Appendix G. Intelligent Disaster Recovery	679
Supported Windows NT/2000 Editions	679
Requirements for IDR	680
Overview of IDR Use	680
About the DR Files	681



Configuring NetBackup Classes for IDR	682
Preparing the IDR Bootable Media	683
Choosing the Bootable Media	683
Creating Bootable Diskettes	684
Creating a Bootable CD Image (applies only to Windows NT)	687
Updating IDR Media	690
Updating Disaster Recovery Diskettes	691
Updating a Disaster Recovery CD	692
Using drfile.exe To Create or Update a DR File	692
Recovering Your Computer	693
Step 1: Boot Your Computer	693
Step 2: Run Windows NT Setup (applies only to Windows NT)	694
Step 3: Run the Disaster Recovery Wizard	695
Preparing Emergency IDR Bootable Diskettes	698
Notes on Altering Hard Drive Partition Sizes	698
Notes on Recovering Specific Platforms	699
Recovering the Dell PowerEdge 6100/200 with RAID	699
Recovering IBM Computers	700
Recovering Compaq Computers	700
IDR's Frequently Asked Questions	700
Glossary	703
Index	727





Preface

This guide describes how to configure and manage the operation of VERITAS NetBackup DataCenter on all supported UNIX platforms and operating systems (in this guide, VERITAS NetBackup is referred to NetBackup). 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.

Audience

This guide is intended for system administrators and assumes a good working knowledge of the operating system on the platform where the product is used. In this guide, a system administrator is defined as a person with system administrator privileges and responsibilities. A client user is anyone that uses the client interfaces to back up, archive, or restore files.

Organization

- ◆ Chapter 1 is an overview of the product. Read this chapter first to get an overview of NetBackup and what it does.
- ◆ Chapter 2 explains how to configure NetBackup to use the storage devices in your network.
- ◆ Chapter 3 explains how to configure NetBackup classes. A class is a group of clients that have the same or similar backup requirements.
- ◆ Chapter 4 explains how to manage and back up the NetBackup internal databases (called catalogs).
- ◆ Chapter 5 explains how to run reports in order to obtain information about NetBackup activities.
- ◆ Chapter 6 explains how to monitor and control NetBackup jobs, processes, and services.
- ◆ Chapter 7 has topics on managing NetBackup operations.



- ◆ Chapter 8 explains how to configure less frequently changed features and parameters.
- ◆ Appendix A has man pages for commands that relate specifically to NetBackup. You can also use the `man` command to view these commands online.
- ◆ Appendix B explains the tasks that you can perform with the `bpadm` interface.
- ◆ Appendix C has useful reference information.
- ◆ Appendix D provides information about scripts that collect information and provide notification of events.
- ◆ Appendix E provides an overview of the NetBackup Global Data Manager product.
- ◆ Appendix F provides information about using NetBackup to back up AFS clients.
- ◆ Appendix G provides information about using NetBackup Intelligent Disaster Recovery for Windows NT/2000.

Following the last appendix is a glossary of terms that you will frequently encounter when using and discussing NetBackup.

Related Documents

Other documents that will be useful when administering NetBackup on a UNIX system are:

- ◆ *NetBackup Release Notes*
Provides important information such as the platforms and operating systems that NetBackup supports and operating notes that may not be in the manuals.
- ◆ *NetBackup DataCenter Installation Guide - UNIX*
Provides instructions for installing and getting your NetBackup configuration up and running.
- ◆ *NetBackup DataCenter Media Manager System Administrator's Guide - UNIX*
Explains how to configure and manage the storage devices and media that UNIX NetBackup servers use for backups.
- ◆ *NetBackup Media Manager Device Configuration Guide*
Provides information about configuring storage devices on UNIX systems.
- ◆ *NetBackup User's Guide - UNIX*
Explains how to use NetBackup to perform backups and restores from a UNIX client.
- ◆ *NetBackup Troubleshooting Guide - UNIX*
Provides troubleshooting information for NetBackup products.

PDF copies of the above documents are provided on the CD-ROM you received with the NetBackup software. For a complete list of NetBackup documents see the *NetBackup Release Notes*.

Conventions

The following explains typographical and other conventions used in this guide.

Type Style

Table 1. Typographic Conventions

Typeface	Usage
Bold fixed width	Input. For example, type cd to change directories.
Fixed width	Paths, commands, filenames, or output. For example: The default installation directory is <code>/opt/VRTSxxx</code> .
<i>Italics</i>	Book titles, new terms, or used for emphasis. For example: <i>Do not</i> ignore cautions.
<i>Sans serif (italics)</i>	Placeholder text or variables. For example: Replace <i>filename</i> with the name of your file.
Sans serif (no italics)	Graphical user interface (GUI) objects, such as fields, menu choices, etc. For example: Enter your password in the Password field.

Notes and Cautions

Note This is a Note and is used to call attention to information that makes it easier to use the product or helps you to avoid problems.

Caution This is a Caution and is used to warn you about situations that can cause data loss.

Key Combinations

Some keyboard command sequences use two or more keys at the same time. For example, you may have to hold down the Ctrl key before you press another key. When this type of command is referenced, the keys are connected by plus signs. For example:

Press Ctrl+t



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
```

the user can use either the *arg1* or *arg2* variable.

Getting Help

For updated information about this product, including system requirements, supported platforms, supported peripherals, and a list of current patches available from Technical Support, visit our web site:

```
http://www.veritas.com/
```

VERITAS Customer Support can also be reached through electronic mail at:

```
support@veritas.com
```



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

- ◆ Overview
- ◆ Media Manager
- ◆ Storage Units
- ◆ Volumes
- ◆ Backup Policies (Classes)
- ◆ User Backups, Archives, and Restores
- ◆ NetBackup Catalog Backups
- ◆ NetBackup Administration Interfaces
- ◆ Configuring NetBackup

Overview

NetBackup provides high-performance backups and restores for a variety of computer types, including Microsoft Windows, NetWare, IBM OS/2, UNIX, and Macintosh.

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 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:

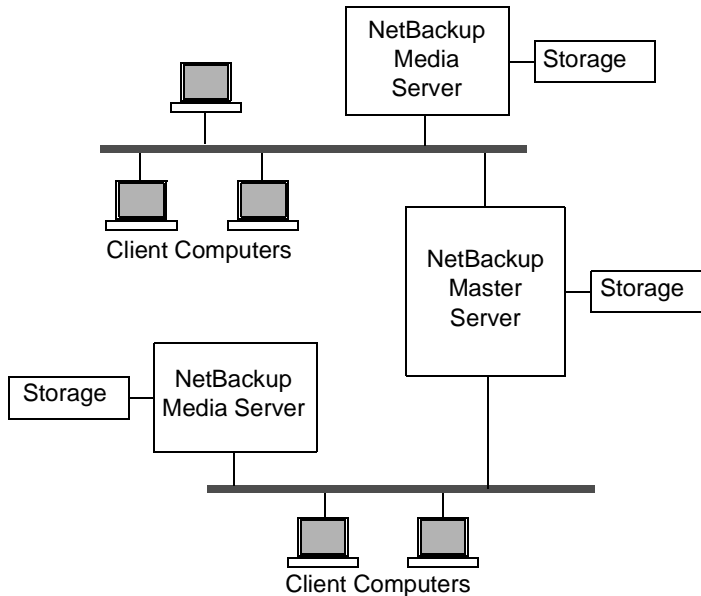
- ◆ Server software is on the computer that manages the storage devices.



- ◆ 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 computer 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 computer 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.

The remaining topics in this chapter provide an overview of other concepts and terminology that will be useful when configuring and managing NetBackup.

Media Manager

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 NetBackup administration interface contains utilities for configuring and managing Media Manager.

The Media Manager system administrator's guide has information on configuring and administering Media Manager (see "Related Documents" in the preface).

Storage Units

The devices that NetBackup uses to store backups are called 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 configuration for the class and schedule. 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.

Volumes

The removable media on which NetBackup stores data are called volumes. These media (for example, 4 mm cartridge tapes) have been assigned media IDs and other attributes so their content, location, and usage can be tracked. The attribute assignment occurs when the administrator adds media to Media Manager and can be done automatically or manually. The system administrator's guide for Media Manager explains how to add volumes.

NetBackup master servers keep a media catalog with records about the volumes where backups are stored. Media Manager controls the mounting of volumes on the devices and keeps a volume database with records that indicate where the volumes are located.



Note When the storage unit is on magnetic disk, volumes are not specified, NetBackup sends the backup to the file path specified during setup of the storage unit and records the location in the NetBackup media catalog. The operating system manages the actual writing of data.

Backup Policies (Classes)

A NetBackup class defines the backup policies for a specific group of one or more clients that have similar backup needs. For example, a class specifies when automatic backups occur for the clients in that class and when users can perform their own backups. NetBackup can have any number of classes, each of which can include one or more clients. A NetBackup client must be in at least one class and can be in more than one.

The administrator defines the properties of a class. The properties that can be assigned depends on the type of clients that the class supports and includes:

- ◆ General attributes that define things such as:
 - Priority* of backups for this class relative to other classes.
 - Storage unit* to use for backups of clients in this class.
 - Volume pool* to use for backups of clients in this class. A volume pool is a set of volumes that the administrator can assign for use by specific classes or schedules. For example, it is possible to have one volume pool for weekly backups and another for quarterly backups.
- ◆ List of client computers in the class.
- ◆ List of files to include in automatic backups of the clients. It is also possible to specify a list of files to exclude from automatic backups. The file list or exclude 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 class 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 can be unique with attributes that include:

- ◆ Type of schedule. You can specify schedules for automatic full or incremental backups or user backups or archives. There are also schedule types for database backups, such as for Microsoft Exchange (requires a separately-priced option to be installed).
- ◆ Backup window. For automatic full or incremental backup schedules, this is the time period when NetBackup can start backups of clients in the class. For user schedules, this is the time period when users can start a backup or archive of their own client.
- ◆ Frequency. How often automatic backups occur.

- ◆ Retention. How long NetBackup retains the data that is backed up by this schedule.
- ◆ Storage unit. The storage unit on which to store the data that is backed up according to this schedule. This setting, if used, overrides the storage unit specified at the class level.
- ◆ Volume pool. The volume pool to use when saving the backup data. This setting, if used, overrides the volume pool specified at the class 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.

User Backups, Archives, and Restores

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).

See the NetBackup user's guides for more information on user operations.

NetBackup Catalog Backups

NetBackup provides a special type of backup for its own internal databases. 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.

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 the workstation that is available to the administrator and personal preference.

- ◆ NetBackup Administration - Java Interface



A Java-based, graphical-user interface. This is the recommended interface and is the one referred to by most procedures and examples in this manual. See “NetBackup Administration - Java Interface” on page 6 for an introduction.

◆ **bpadm - Character-based, Menu Interface**

A character-based, menu interface that can be used from any terminal (or terminal emulation window) that has a `termcap` or `terminfo` definition. For information on `bpadm`, see Appendix B.

◆ **xnb - X Windows Interface**

An X Windows-based, graphical-user interface that can be used from any X terminal or workstation with an X server program that is compatible with release X11.R5 or later. On UNIX systems, except Solaris and HP-UX, you start this program with the `/user/opensv/netbackup/bin/xnb` command. On Solaris and HP-UX, this command is in the `/user/opensv/netbackup/bin/goodies` directory.

The `xnb` program follows OSF/Motif 1.2 conventions. For a description of Motif usage, see the *OSF/Motif User's Guide* authored by the Open Software Foundation and published by Prentice Hall (ISBN 0-14-640509-6).

◆ **Command Line**

NetBackup commands that can be entered at the system prompt or used in scripts. For details on their capabilities, see Appendix A. To view them online, use the UNIX `man` command.

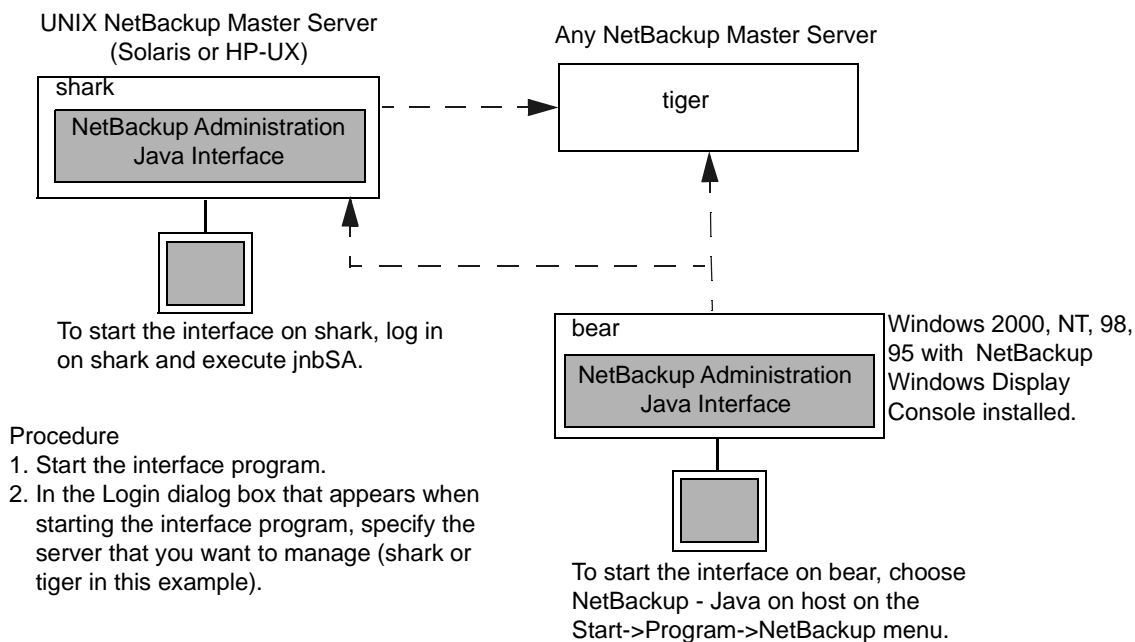
Note All NetBackup administrator programs and commands require root-user privileges by default. If it is necessary to have nonroot administrators, see “Configuring Nonroot Usage” on page 258.

NetBackup Administration - Java Interface

The NetBackup Administration - Java program (`jnbSA`) provides a graphical-user-interface through which the administrator can manage every major area of NetBackup. The interface can be started and run on the following systems (see Figure 1):

- ◆ Any supported Solaris or HP-UX system.
- ◆ Windows 2000, NT, 98, or 95 systems that have the NetBackup Java Windows Display Console installed. The Windows Display console allows you to use the Java interface for administering UNIX NetBackup servers when a Solaris or HP-UX system is not available or to directly administer a NetBackup Windows NT/2000 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.

Figure 1. NetBackup Java Interface



After starting the interface either directly on a UNIX NetBackup server or a display console, the administrator specifies the NetBackup master server they want to manage. This server can be any NetBackup master server.

The procedures for starting the interface are explained below. For configuration information, see “Authorizing NetBackup-Java Users” on page 256 and “Configuration Options for jbpSA” on page 260.

Setting Up Your Window Manager for the Java Interface

Always set your window manager so windows become active only when you click inside the windows. Do not enable auto focus, which is when windows become active if you just move the mouse pointer over them. The NetBackup-Java interfaces do not run properly with auto focus enabled. The following are general instructions for correctly setting up the focus.

CDE (Common Desktop Environment)

The following explains how to set up a CDE (Common Desktop Environment) window manager, which is the preferred window manager for NetBackup-Java applications.

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 box appears.
3. In the Style Manager-Window dialog box, click the Click In Window To Make Active button.
4. Click OK.
5. Click OK when asked to Restart the Workspace Manager.

Motif

When using the Motif window manager, set the `Mwm*keyboardFocusPolicy` X resource as follows:

```
Mwm*keyboardFocusPolicy:explicit
```

To Start the NetBackup-Java Interface on a UNIX System


1. Log in (as root) on the NetBackup client or server where you are going to start the NetBackup Administration - Java interface.

This client or server must be a Solaris or HP-UX system that is supported by NetBackup.

2. Start the interface program by executing:

```
/usr/opensv/netbackup/bin/jnbSA &
```

The login screen of the NetBackup Administration - Java window appears.



3. Type the name of the master server where you initially want to manage NetBackup (for example, shark or tiger in Figure 1). You can specify any NetBackup master server (Windows NT/2000 or UNIX).
4. Type your user name and password and click Login.

When logging into a Windows NT/2000 server, you must enter both the server's domain and the user name as follows:

`domain_name\user_name`

domainname specifies the domain of the NetBackup host (not required if the NetBackup host is not a member of a domain).

This logs you into the NetBackup Java application server program on the specified server and the screen on Figure 2 appears. The interface program continues to communicate through the server you specified for the remainder of the current session.

5. Start the desired utility by clicking its icon or selecting its command from a menu in the NetBackup Administration window (see Figure 2).




Most NetBackup-Java utilities also provide a change master server or change host command that allows you to use that specific utility to make configuration changes on a remote server. For example, Backup Policy Management has a Change Master Server command on the file menu where you can specify another server on which to configure classes.

Note The NetBackup-Java interfaces support remote X Windows display only between Solaris 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 executing `jnbSA -d tiger`. However, if *shark* were an HP system, you could display `jnbSA` only directly on *shark*.

To Start the NetBackup-Java Interface on Windows Systems

1. On a Windows NT/2000/98/95 computer where the Windows Display Console is installed and configured, click Programs on the Start menu.
2. On the Programs menu, point to VERITAS NetBackup and then click NetBackup - Java on *host* on the submenu. *host* is the default NetBackup server to manage as set during installation.

The login screen for NetBackup Administration - Java appears and displays *host* as the NetBackup server to log into. To log into another server, type its name in the box.



3. Type your user name and password and click Login. When logging into a Windows NT/2000 server, you must enter both the server's domain and the user name as follows:

`domain_name\user_name`

domainname specifies the domain of the NetBackup host (not required if the NetBackup host is not a member of a domain).

This logs you into the NetBackup-Java application server program on the specified server and the window on Figure 2 appears. The interface program continues to communicate through the server specified in the login screen for the remainder of the current session.

Note To change the default NetBackup server or to add a menu item for another server, see the instructions in the *install_directory\java\readme.txt* file on the Windows 2000, NT, 98, or 95 display console system. By default, *install_directory* is `C:\Program Files\VERITAS`.

4. Start the desired utility by clicking its icon or selecting its command from a menu in the NetBackup Administration - Java window (see the next topic).

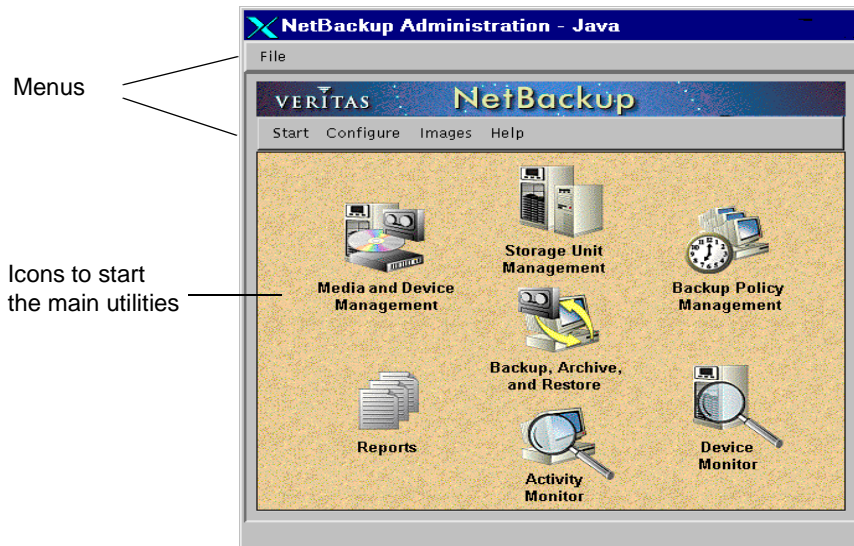
Most NetBackup-Java utilities also provide a change master server or change host command that allows you to use that specific utility to make configuration changes on a remote server. For example, Backup Policy Management has a Change Master Server command on the file menu where you can specify another server on which to configure classes.

NetBackup Administration - Java Window

In the NetBackup Administration - Java window (Figure 2), click an icon to start the associated utility. The menus contain commands to start these same utilities and to perform other functions as explained in the table below. For instructions on performing specific operations, see the other chapters in this manual and the online help.



Figure 2. NetBackup Administration - Java Window



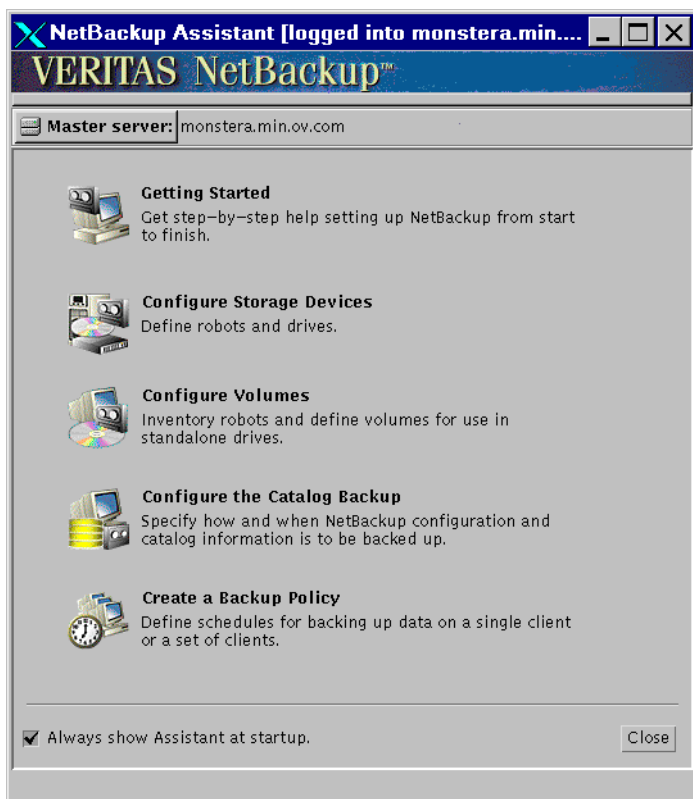
Menu	Commands
File	Exit - Closes the NetBackup Administration - Java interface, including all other NetBackup-Java windows that are open.
Start	<p>Media and Device Management - Displays the Media and Device Management window, which has commands for managing the media and devices that NetBackup uses to store its backups. For more information, see the system administrator's guide for Media Manager.</p> <p>Storage Unit Management - Displays the Storage Unit Management window, which has commands for managing NetBackup storage units.</p> <p>Backup Policy Management - Displays the Backup Policy Management window, which has commands for setting up NetBackup classes.</p> <p>Backup, Archive, and Restore - Starts the client-user interface application on the NetBackup server. This program has commands for performing backups, archives, and restores. See the online help for information on using this interface.</p>

Menu	Commands
	<p>Reports - Displays the Reports window from which you can generate reports about various NetBackup activities.</p> <p>Activity Monitor - Starts the NetBackup Activity Monitor, which allows you to monitor NetBackup jobs and also provides limited control over their execution.</p> <p>Device Monitor - Displays the Device Monitor window, from which you can monitor and control the operation of storage devices. For more information on this utility, see the system administrator's guide for Media Manager.</p> <p>Backup "NetBackup Catalog" - Starts an immediate backup of the NetBackup internal databases (called catalogs).</p> <p>Assistant - Starts the NetBackup Assistant, from which you can start the configuration wizards.</p> <p>Filesystem Analyzer - Starts the filesystem analyzer. This is useful for analyzing filesystems for use with VERITAS Storage Migrator.</p>
Configure	<p>NetBackup Catalog Backup - Displays a dialog box where you can configure backups of the NetBackup internal databases (which are called catalogs).</p> <p>NetBackup System Configuration - Displays a dialog box where you can configure NetBackup global attributes and retention periods for backups.</p> <p>Adjust Application Timezone - Displays a dialog box where you can adjust the timezone settings when the server and clients are in different timezones. For more information, see "Adjust Time Zone" on page 322.</p>
Images	<p>Import - Imports backups whose retention period is expired or are from another NetBackup server.</p> <p>Verify - Verifies that the contents of a backup matches the records in the NetBackup catalog.</p> <p>Duplicate - Duplicates NetBackup images.</p>
Help	<p>Help Topics - Provides online help information.</p> <p>About NetBackup Administration - Displays program information, version number, and copyright.</p>



Configuring NetBackup

The easiest way to configure NetBackup is to use NetBackup Assistant. This program lets you launch configuration wizards that simplify the setup process. The NetBackup Assistant starts automatically after installation is complete and you can start it again later by choosing Assistant from the Start menu in the NetBackup Administration window.



- ◆ If you are configuring NetBackup for the first time, choose the Getting Started wizard. This takes you through all the steps (including the other wizards) and leaves you with a working NetBackup configuration. You can start this wizard from NetBackup Assistant.
- ◆ When adding to an existing configuration, save time by choosing one of the following:
 - ◆ Configure Storage Devices
 - ◆ Configure Volumes
 - ◆ Configure the Catalog Backup
 - ◆ Create a Backup Policy

- ◆ To perform remote configuration with NetBackup Assistant, press the Master server button and specify the name of the remote server.
- ◆ To always use NetBackup Assistant when you start the NetBackup Administration interface, leave the Always Show Assistant at Startup box selected. The program now starts each time you start the NetBackup Administration interface. To disable this action, clear the box. You can still start NetBackup Assistant by clicking the Assistant command on the Start menu.

If you do not want to use the wizards, use the NetBackup Administration interface as explained in following steps:

1. Start the desired NetBackup Administration interface (see “NetBackup Administration Interfaces” on page 5).
2. Complete the addition of storage devices.
Using the Configure Storage Devices wizard is the preferred method. To perform configuration without the wizard, see the system administrator’s guide for Media Manager.
3. Add the media that you will use. For instructions, see the system administrator’s guide for NetBackup Media Manager.
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 executing

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

Note See “Managing Daemons” on page 208 for instructions on starting and stopping `bprd` and `bpdbm`.

5. Define the storage units as explained in Chapter 2.
6. Verify the catalog backup configuration (see Chapter 4).
 - 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 classes as explained in Chapter 3 (setting up schedules is part of defining classes).
8. Perform any required additional configuration, as explained in Chapter 8.





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

- ◆ Introduction to Storage Units
- ◆ Using the Configure Storage Devices Wizard
- ◆ To Start the Storage Unit Management Utility
- ◆ The Storage Unit Management Window
- ◆ Adding Media Manager Storage Units
- ◆ Adding Disk Type Storage Units
- ◆ Changing Storage Unit Attributes
- ◆ Deleting Storage Units
- ◆ Automatic Drive Availability Checking

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. During a backup or archive, NetBackup stores the backup data on the storage units that you have set up during configuration. The types of storage units that you can set up are as follows:

- ◆ Media Manager

A Media Manager storage unit uses robots or standalone tape drives that are under control of Media Manager. Media Manager controls the allocation and mounting of media (called volumes) in the storage devices.

- ◆ Disk

A disk type storage unit consists of a directory on a hard disk that stores the backup or archive data. The following is an example path in a Windows NT/2000 file system:
D:\NetBackup\backups. The following is an example path in a UNIX file system



/disk1/nb_storage_unit. A disk type storage unit is useful for testing and during busy periods because it allows quick backups. However, you must be careful that it does not fill up your disk.

◆ NDMP

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.

Using the Configure Storage Devices Wizard

If you are configuring storage units for the first time, the easiest way is to use the Configure Storage Devices wizard. This wizard guides you through the entire process, simplifying it by automatically choosing settings that are good for most configurations (for Media Manager storage units, the wizard starts with device configuration). If you are modifying an existing configuration or want more control over settings than the wizard provides, use the Storage Unit Management utility as explained in “Adding Media Manager Storage Units” on page 23.

To use the wizard:

1. Click Assistant on the Start menu in the NetBackup Administration window.
The NetBackup Assistant window appears.
2. In the NetBackup Assistant window, click Configure Storage Devices.

Note If the wizard does not find any devices and you want to configure a disk storage unit, you must use the Storage Management utility.

To Start the Storage Unit Management Utility

1. Start the NetBackup Administration - Java interface program. For instructions, see “NetBackup Administration Interfaces” on page 5.
The NetBackup Administration window appears.
2. In the NetBackup Administration window, click Storage Unit Management.
The Storage Unit Management window appears.

The Storage Unit Management Window

The Storage Unit Management window has tools for configuring and managing storage units. The following topics provide an overview of this window:



- ◆ Tree and Detail Views
- ◆ Menu Bar
- ◆ Toolbars
- ◆ Shortcut Menus

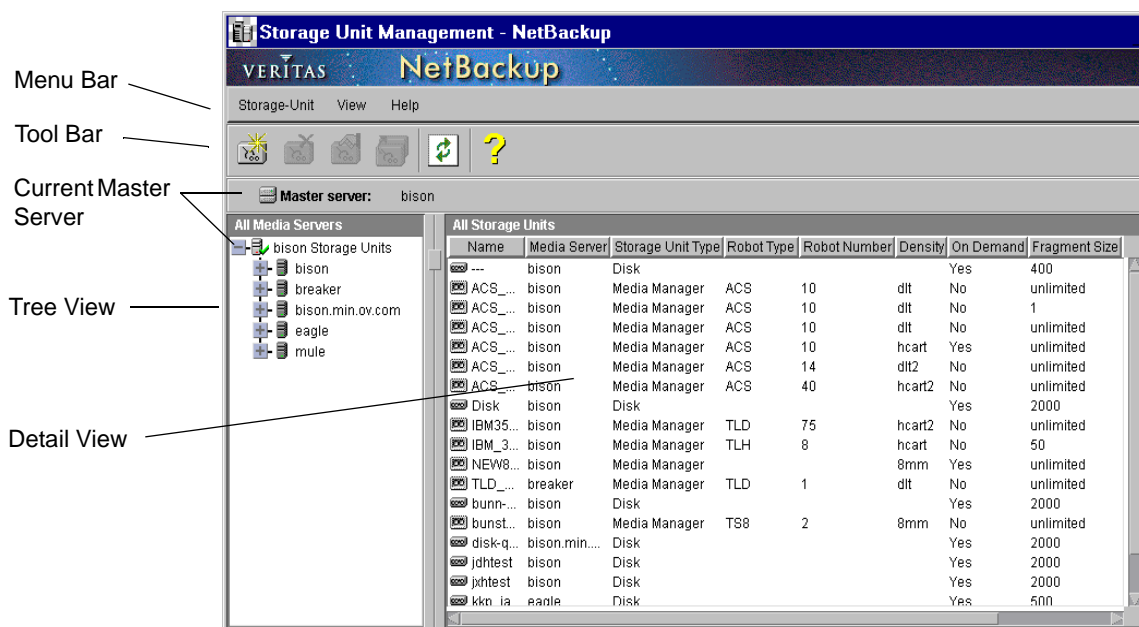
Tree and Detail Views

The Storage Unit Management window has two panes where you can select items and perform actions on them. The left pane is a hierarchical, tree view. The right pane is a detail view that shows information about whatever is selected in the left pane. The contents of these panes are different, depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters).

See Appendix E for more information on the Global Data Manager option.

The View if You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the left pane initially shows all media servers that have storage units belonging to the master server you are currently managing. The current master server appears as the first media server in the tree.



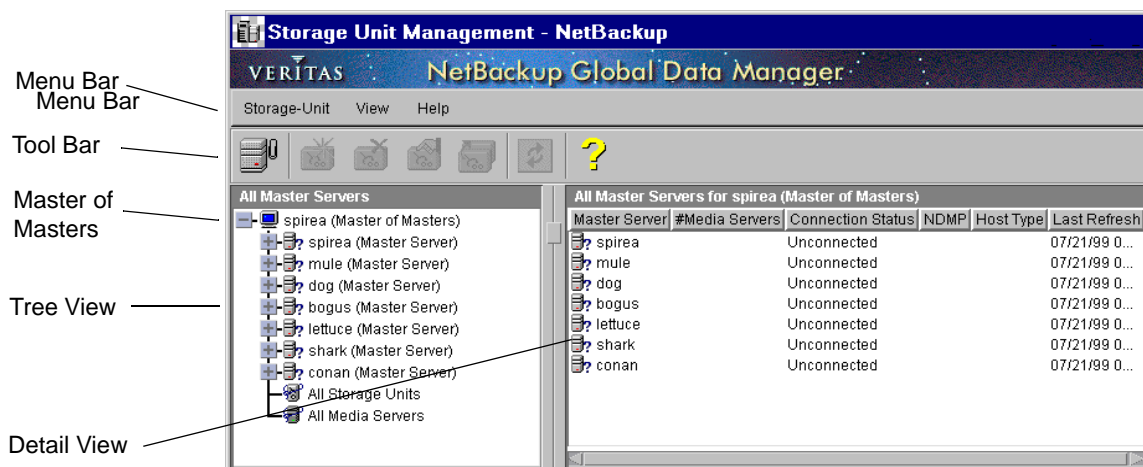
The right pane initially shows a list of all storage units.

- ◆ If you select a media server in the tree view, the right pane shows details about the storage units on that media server.
- ◆ If you select a storage unit in the tree view, the right pane shows details about only that storage unit.

To manage and view the storage unit configuration on another master server, change to that server as explained in “Choosing the Master Server for Storage Units” on page 22.

The View From a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, the left pane shows the master servers that you can currently manage, and the storage units on those servers.



The right pane shows details about the current selection in the tree.

- ◆ Selecting a master server expands the tree to show the media servers that have storage units belonging to the selected master server. The selected master server also appears in this list as a media server. To append other master servers to the tree, use the Append Master Server command on the Storage Unit menu.

Each master server also has an All storage units branch. When you select this branch, the detail view shows a list of all storage units on all media servers that are under that master server.

- ◆ The lower two branches of the tree, All storage units and All media servers, allow you to see combined lists of all the storage units and media servers that you can manage from this Master of Masters.
 - ◆ To see a list of all storage units on all master servers, select the All storage units branch of the tree.



- ◆ To see a list of all media servers on all master servers, select the All media servers branch of the tree.

Menu Bar

The following table describes the menus and commands that are available on the menu bar.

Menu	Command
Storage Unit	<p>New - Displays a dialog box where you can specify the attributes for a new storage unit.</p> <p>Delete - Deletes a selected storage unit from the configuration.</p> <p>Change - Displays a dialog box for changing the configuration of the selected storage unit(s).</p> <p>Copy - Copies the attributes of a selected storage unit and displays them in a dialog box so you can change them as necessary for a new storage unit.</p> <p>Change Server - (Applies only when Global Data Manager is not installed) Displays a dialog box where you can specify the NetBackup server that has the configuration you want to modify.</p> <p>Append Master Server - (Applies only to Global Data Manager) Displays a dialog box where you can specify a master server to add to the tree. The addition applies only to the current session and the server will not be in the tree the next time you start the administrator interface. Also, see “Append Master Server” on page 670.</p> <p>Exit - Closes this window.</p>
View	<p>Preferences - Displays a dialog box where you can specify your display preferences for the toolbar.</p> <p>Ignore Master Server(s) - (Applies only to Global Data Manager) Ignores connections to the master server(s) that you designate so the Master of Masters does not attempt connections to them. The branch of the tree for an ignored server and the associated information in the detail view appears in faded type. You cannot select an ignored server or refresh the display with updated configuration information. The only operation you can perform on an ignored server is to recognize it (see below).</p> <p>Recognize Master Server(s) - (Applies only to Global Data Manager) Recognizes an ignored master server so it can be selected and the display refreshed with updated configuration information. Also see “Recognize Master Server(s)” on page 672.</p> <p>Refresh - Updates the detail pane with new information retrieved from the master server(s).</p>



Menu	Command
Help	Help Topics - Provides online help information. About Storage Unit Management - Displays program information, version number, and copyright.

Toolbars

The buttons on the toolbar provide shortcuts for menu commands. If the toolbar is not visible or you want to change its size, click the Preferences command on the View menu and make the desired choices in the Preferences dialog box.

Shortcut Menus

Clicking the right mouse button when the pointer is over either pane in the window, displays a shortcut menu with commands that apply to what is currently selected. These commands are also on the menu bar and are explained in the “Menu Bar” topic.

Choosing the Master Server for Storage Units

If there is more than one NetBackup master server, choose the one that has the storage unit configuration you want to manage. The procedure is different depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters):

See Appendix E for an overview of the Global Data Manager option.

If You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the name of the server that you are currently managing appears on the Master Server line above the left pane of the window. To manage the configuration on another NetBackup master server:

1. On the Storage Unit menu, click Change Master Server.
2. In the Change Master Server dialog box, specify the name of the NetBackup server where you want to manage the storage unit configuration and click OK.

If you encounter problems, see “If You Cannot Change to Another Master Server.”



If You Are On a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, choose a master server by selecting it in the tree. If the server you want is not in the tree, append it by using the Append Master Server command on the Storage Unit menu. If you encounter problems, see “If You Cannot Change to Another Master Server.”

If You Cannot Change to Another Master Server

Verify the following:

1. The other server is operational.
2. The NetBackup server list on the other server has an entry for the UNIX NetBackup server that you specified in the login dialog box when starting the NetBackup Administration interface. If necessary, add the server list entry.

For example, assume you start the interface on a system named bear and specify shark in the login box. In this instance, you are running the interface on bear but managing NetBackup through shark (where you logged in). Now assume you want to change to a server named tiger but NetBackup does not permit it. After verifying that tiger is operational, try adding shark to its server list:

- ◆ If tiger is a Windows NT/2000 system, start the NetBackup administrator’s interface on tiger and make the addition on the Servers tab in the Master Server Properties or Media Server Properties dialog box (see the online help on that server for further instructions). Then, stop and restart the NetBackup Database Manager and NetBackup Request Manager services on tiger.
- ◆ If tiger is a UNIX system, add a `SERVER=host` entry below the existing entries in the `/usr/openv/netbackup/bp.conf` file on tiger. In this example, `host` is shark. Then, stop and restart the NetBackup database manager (`bpdbm`) and NetBackup request daemon (`bpird`) on tiger.

Adding Media Manager Storage Units

Rules for Media Manager Storage Units

The rules for adding Media Manager storage units are:

1. Add the storage unit to the server where the drives attach.
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 must be in the same storage unit. For example, if a robot has two drives of the same density, add only a single storage unit for the robot.
- ◆ 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.

If a robot's drives and robotic control attach to different NetBackup servers, add the storage unit on the server where the drives attach. Always specify the same robot number that is used for the robotic control on the other server (see "Information Required for Storage Units on shark" on page 27).

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 that includes both of them. 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.

Before Adding the Storage Unit

Before adding a Media Manager storage unit, set up Media Manager to recognize the devices that will compose the storage units (the system administrator's guide for Media Manager explains device configuration). As you set up the devices, record the following information from the Media Manager configuration:

For robots, record:

- ◆ NetBackup servers where the drives attach and the number of drives that attach to each server
- ◆ Robot type
- ◆ Robot number in Media Manager
- ◆ Media density for the drives in each robot


For standalone drives, record:

- ◆ Media density of each drive
- ◆ How many drives of each media density are on each NetBackup server

The following examples show the type of information required by NetBackup for various Media Manager storage unit configurations. The procedure that follows the examples provides step-by-step instructions on how to specify this information to NetBackup.

Example 1

The following figure shows a master server that has 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. For a list of vendor models corresponding to each NetBackup name, see the Platforms and Peripherals section of the NetBackup release notes that came with your software.

Each of these devices can be a storage unit and 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
Number of Drives	1
Density	8mm - 8mm cartridge

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

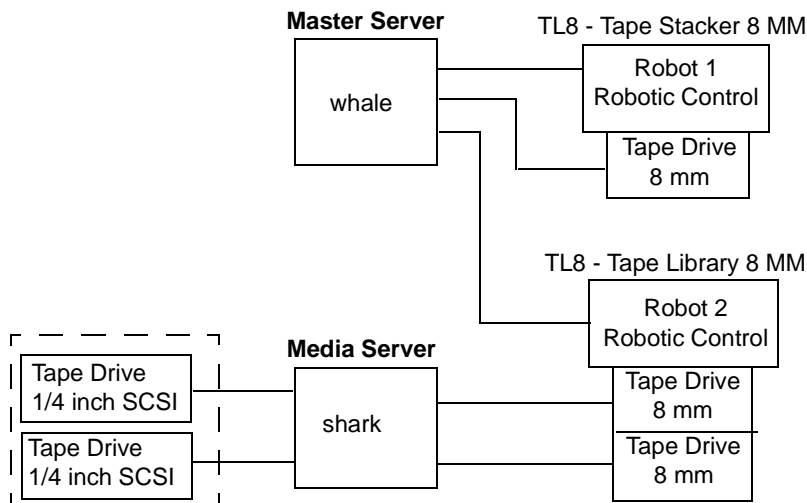
- ◆ SCSI 1/4 inch tape drive

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



Example 2

The following figure shows a master server (whale) that has a drive in a robot and a media server (shark) that has two drives in a robot and two standalone 1/4 inch SCSI tape drives.



Information Required for Storage Unit on whale

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

Storage Unit Configuration Setting	Value
Media Server	whale
Robot Type	TL8 - Tape Stacker 8MM
Robot Number	1
Number of Drives	1
Density	8mm - 8mm Cartridge

The server named whale also controls the robotics for the TL8 - Tape Library 8MM robot. However, the drives in this robot attach to shark and therefore the storage unit that contains them must also be on shark.

Information Required for Storage Units on shark

On 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
Number of Drives	2
Density	8mm - 8mm Cartridge

The robotic control for the TL8 - Tape Library 8MM is on 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
Number of 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.

To Add a Media Manager Storage Unit

1. In the NetBackup Administration window, click the Storage Unit Management icon.
The Storage Unit Management window appears.



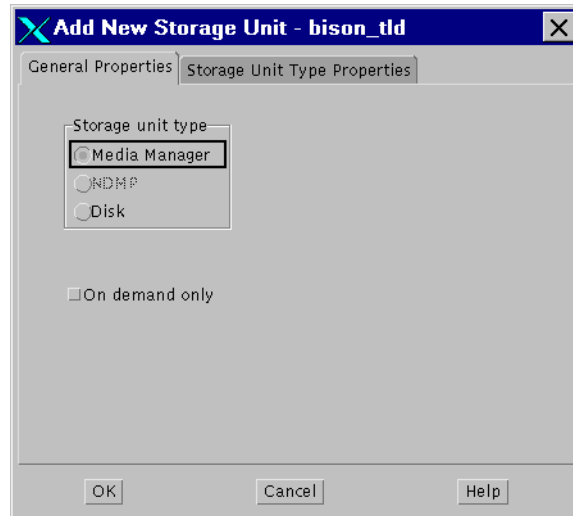
2. If your site has more than one master server, use the Change Server command on the Storage unit menu to choose the server with the configuration that will use the storage unit. See “Choosing the Master Server for Storage Units” on page 22.
3. To add the storage unit by starting with the standard defaults:
 - a. Click New on the Storage Unit menu
The Add a New Storage Unit dialog box appears.
 - b. Proceed to step 5.
4. To add the storage unit by copying and then modifying the properties of an existing storage unit, proceed as explained below.
 - a. Select the storage unit from those listed in the left or right pane of the Storage Unit Management window.
 - b. On the Storage Unit menu, click Copy.
The Copy a Storage Unit dialog box appears.
5. Type a unique name for the new storage unit.

This is the name that you use when specifying a storage unit for classes and schedules. It is best to choose a name that describes the type of storage you are defining.

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.
6. Click OK.



Another dialog box appears. Notice that the name in the title bar is the one that you just entered.



7. Complete the entries on the General Properties and Storage Unit Type Properties tabs as explained in the following topics:
 - ◆ General Properties - Media Manager Storage Units
 - ◆ Storage Unit Type Properties - Media Manager Storage Units
8. Click OK to add the storage unit to the configuration.

General Properties - Media Manager Storage Units

Storage unit type

Specifies the type of storage that this storage unit uses. Click the button for the type of storage unit you are configuring.

On demand only

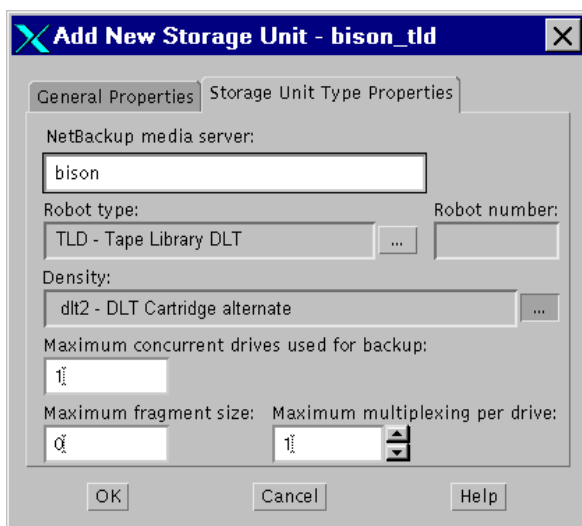
Specifies whether the storage unit is available *only* on demand (that is, only when a class or schedule is explicitly configured to use this storage unit). Clear the box (default), to make the storage unit available to any class or schedule.

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



Storage Unit Type Properties - Media Manager Storage Units

The following figure shows the Storage Unit Type Properties tab.



NetBackup media server

Specifies the name of the NetBackup server where the drives in the storage unit attach. Enter the name that is used for that server in the NetBackup server list.

Robot Type

Specifies the type of robot (if any) that the storage unit contains. Click the button to the right of the Robot Type box and select NONE-Not Robotic or one of the other robot types from the list. The list shows all the types that NetBackup supports.

The list uses NetBackup designations for robot types. For the vendor models that correspond to each robot type, go to the VERITAS support web page

www.support.veritas.com/

Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

Robot Number

For a robotic storage unit, this is the same robot number used in the Media Manager configuration. For more information on robot numbers, see the system administrator's guide for Media Manager.

Density

Specifies the media density that the storage unit will be using. Click the button to the right of the box and select from the list.

Maximum concurrent drives used for backup

Specifies the number of drives that NetBackup can use at one time for backups in this storage unit. Type the desired number in the box:

- ◆ For a storage unit that contains only standalone drives, specify a number that is less than or equal to the number of drives that are in this storage unit.
- ◆ For a robot, specify a number that is less than or equal to the number of drives that attach to the NetBackup media server for the storage unit.

Assume you have two standalone drives of the same density and you specify 1 in this box. In this instance, both drives are available to NetBackup but only one can be used for backups. This leaves the other drive available for restores and other nonbackup operations (such as importing, verifying, and duplicating backups).

Maximum fragment size

Specifies (in megabytes) the largest fragment size that NetBackup can create when storing backups. To specify a fragment size, type a value of 50 or larger. To specify unlimited fragment size, type 0 (0 is the default).

For more information, see “Fragmented Backups” on page 643.

Note If you change the fragment size, you can still restore backups that were written with the previous fragment size.

Maximum multiplexing per drive

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

- ◆ Specify any value from 1 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 37 for more information.



Adding Disk Type Storage Units

This section explains how to configure a NetBackup storage unit that resides in a directory on a hard disk. NetBackup permits an unlimited number of disk storage units.

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, multiply these numbers:

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

To Add a Disk Type Storage Unit

1. In the NetBackup Administration window, click the Storage Unit Management icon.
The Storage Unit Management window appears.
2. If your site has more than one master server, use the Change Server command on the Storage unit menu to choose the server with the configuration that will use the storage unit. See “Choosing the Master Server for Storage Units” on page 22.
3. To add the storage unit by starting with the standard defaults, click New on the Storage Unit menu.

The Add a New Storage Unit dialog box appears.

4. To add the storage unit by copying and then modifying the properties for an existing storage unit, proceed as explained below.
 - a. Select the storage unit from those listed in the left or right pane of the Storage Unit Management window.
 - b. On the Storage Unit menu, click Copy.

The Copy a Storage Unit dialog box appears.

5. Type a unique name for the new storage unit.

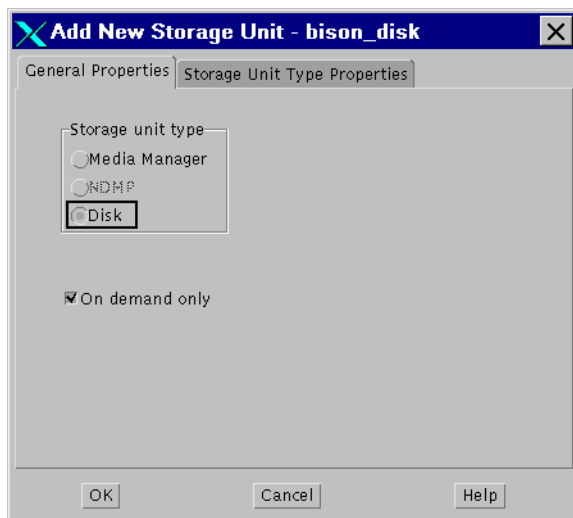
This is the name that you use when specifying a storage unit for classes and schedules. It is best to choose a name that describes the type of storage you are defining.

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.

6. Click OK.



Another dialog box opens. Notice that the name in the title bar is the one that you just entered.



7. Complete the entries on the General Properties and Storage Unit Type Properties tabs as explained in the following topics:
 - ◆ General Properties - Disk Storage Units
 - ◆ Storage Unit Type Properties - Disk Storage Units
8. Click OK to add the storage unit to the configuration.

General Properties - Disk Storage Units

Storage unit type

Specifies the type of storage that this storage unit supports. Click the button for the type of storage unit you are configuring.

On demand only

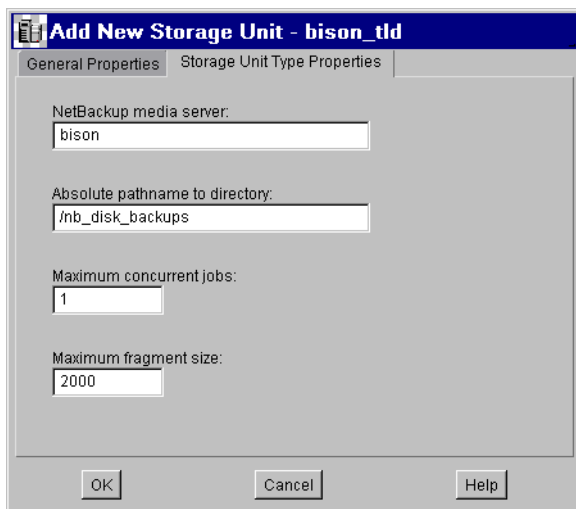
Specifies whether the storage unit is available *only* on demand (that is, only when a class or schedule is explicitly configured to use this storage unit). Clear the box, to make the storage unit available to any class or schedule. On demand only is the default for a new disk type storage unit.

Note If you make all storage units on demand, remember to specify one of those storage units for each class or schedule that you configure. Otherwise, NetBackup will be unable to find a storage unit to use.



Storage Unit Type Properties - Disk Storage Units

The following figure shows the Storage Unit Type Properties tab for disk storage units.



NetBackup media server

Specifies the name of the server that is controlling the disk. Enter the name that is used for that server in the NetBackup server list.

Absolute pathname to directory

Specifies the absolute pathname to the file system that will store the backups. You can use any location on the disk, providing there is sufficient space available.

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.

This setting equates to the Number of drives setting for a Media Manager storage unit. The jobs are not multiplexed.

The number to use here depends on the available disk space and the server's ability to comfortably execute multiple backup processes. Also, see "Limit jobs per class" on page 68. The default is 1.

Maximum fragment size

Specifies (in megabytes) the largest fragment that NetBackup can create when storing backups.

The value can range from 20 to 2000 (2000 is the default). The Maximum fragment size setting is normally used to ensure that the backup does not exceed the maximum size allowed by the file system.

For more information, see “Fragmented Backups” on page 643.

Note If you change the fragment size, you can still restore backups that were written with the previous fragment size.

Changing Storage Unit Attributes

Note We suggest that you make changes only during periods when you are not expecting backup activity for classes 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.

1. In the NetBackup Administration window, click Storage Unit Management.
The Storage Unit Management window appears.
2. If your site has more than one master server, choose the one with the configuration that has the storage unit. See “Choosing the Master Server for Storage Units” on page 22.
3. Select the storage units from those listed in the left or right pane of the Storage Unit Management window.
4. On the Storage Unit menu, click Change.
A Change Storage Unit dialog box appears. The name in the title bar is the one that you just selected.
5. After making your changes, click OK to confirm them.

Deleting Storage Units

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



1. In the NetBackup Administration window, click Storage Unit Management.
The Storage Unit Management window appears.
2. If your site has more than one master server, choose the one with the configuration that has the storage unit. See “Choosing the Master Server for Storage Units” on page 22.
3. Select the storage unit from those listed in the left or right pane of the Storage Unit Management window.
4. On the Storage Unit menu, click Delete.
The Delete a Storage Unit dialog box appears so you can confirm or cancel the deletion.
5. If any classes use the storage unit that you deleted, modify those classes to use another storage unit.

Automatic 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 `RE_READ_INTERVAL` setting in the `/usr/opensv/netbackup/bp.conf` file determines how often NetBackup checks storage units for available drives. For more information on this setting, see “`RE_READ_INTERVAL`” on page 339.

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 `BPTM_QUERY_TIMEOUT` entry in the `/usr/opensv/netbackup/bp.conf` file. For more information on this setting, see “`BPTM_QUERY_TIMEOUT`” on page 329).



Requeuing Jobs If Required Storage Units are Unavailable

By default, a job fails (status code 219) if a required storage unit is unavailable when a job starts or, for some reason, becomes unavailable during a backup. However, you can configure NetBackup to requeue jobs for either of these conditions. To configure NetBackup to requeue jobs, use the following entries in the `/usr/opensv/netbackup/bp.conf` file (for specific information on values and defaults, see “bp.conf Options for UNIX Servers” on page 326).

- ◆ `WAIT_IN_QUEUE` 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 the storage unit becomes available. A job fails if the `TIMEOUT_IN_QUEUE` time expires or its backup window closes before the storage unit becomes available.
- ◆ `QUEUE_ON_ERROR` causes jobs to enter the requeued state on startup, if the required storage unit is not available. The jobs will then run when the storage unit becomes available. If this entry is not present, the job fails with a 219 status. This entry requires that the `WAIT_IN_QUEUE` entry also exist or the job will fail immediately anyway with a 219 status if the storage unit is not available.
- ◆ `TIMEOUT_IN_QUEUE` determines how long a requeued job waits for an unavailable required storage unit.





NetBackup classes (also called backup policies) define the rules that NetBackup follows when backing up its clients. A class can contain one or more clients and every client must belong to at least one class. The best approach to setting up classes is usually to divide clients into groups according to their backup and archiving requirements and then create classes as appropriate for each group.

This chapter explains how to configure classes. The first two topics introduce classes and give guidelines for planning them. The remaining topics have detailed configuration instructions.



Introduction to NetBackup Classes

Parts of a Class Definition

There are four parts to a class definition. You set each of them up by using the Backup Policy (Classes) window as explained later in this chapter:

- ◆ General Attributes
- ◆ Client List
- ◆ File List
- ◆ Schedules

General Attributes

The general attributes determine the basic characteristics of all the backups that NetBackup performs for the class. For example:

- ◆ Whether the class is active (so NetBackup can use it for backups).
- ◆ Class type (type of clients it can include).
- ◆ Priority that NetBackup gives to jobs in this class relative to other classes.
- ◆ Storage unit for backups of clients in this class. A storage unit on the schedule, overrides the general attribute setting.

A complete list of general attributes is provided later in this chapter.

Client List

The client list names the computers that belong to the class. NetBackup backs up these computers according to the file list, schedules, and general attributes for the class. A client must belong to at least one class and can belong to more than one. Having a client in more than one class is also useful, for example, to back up different sets of files according to different rules.

File List

The file list names the files and directories that NetBackup includes in automatic backups for clients in this class. The file list does not apply to user backups or archives because in those instances the user selects the files.

NetBackup uses the same file list for all clients in the class but all the files do not have to exist on all clients. NetBackup backs up the files that it finds. Files are processed serially for each client. It is possible, however, to back up more than one client (or class) in parallel.



A related item is the exclude list. You can create this list on each client to specify files that you do not want to include in automatic backups of that client. The exclude list does not apply to user backups and archives.

Schedules

Schedules control the backups for the class. The two basic categories of schedules are automatic and user and there are different types of schedules within these categories:

- ◆ Automatic schedules back up the file list on all clients in the class 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. An incremental backup includes only files that have changed since the last backup. A full backup includes all files in the file list regardless of whether they have changed.
- ◆ User schedules specify 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 but keeping a copy for future use.

Each schedule also includes criteria, such as the storage unit to use and how long to retain the backups (it is usually best to set the retention period to infinite for archives).

Example Classes

The following figure shows the clients, file list, and schedules for two example classes.

- ◆ Example class 1 specifies that files in

`/usr`

`/home`

be backed up for the clients named mars, jupiter, and neptune. This class has daily, and weekly automatic schedules and a user backup schedule. All backups go to 8-mm tape.



- ◆ Example class 2 has different scheduling requirements. One difference is that this class has monthly fulls that go to DLT tape.

Example Class 1

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

Example Class 2

Client List	File List	Schedules		
pluto	/usr	Daily Incrementals Run every day between 6 pm and 6 am Store on 8 mm tape Keep 14 days	Weekly Fulls Run Tuesdays every week between 6 pm and 6 am. Store on 8 mm tape Keep one month	Monthly Fulls Run Sundays every month between 6 pm and 6 am. Store on DLT tape Keep one year
mercury	/home			



Planning Guidelines for Classes

Classes allow you to meet the needs of a wide variety of clients in a single NetBackup configuration. Taking full advantage of classes requires careful planning before starting your configuration. The following procedure provides some planning guidelines.

1. Divide clients into groups according to the types of work they perform.

Clients used for similar tasks usually have a high level of commonality in their backup requirements. For example, most clients in an engineering department create the same types of files at similar levels of importance.

In some instances, you can create a single class for each group and this will be adequate for defining your classes. In other cases, you will have to further subdivide the clients based on their backup requirements as explained in later this procedure.

The table below is the initial grouping for our example. We assume these clients are in same work group and the initial plan is to have them all in the same class.

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 put mercury in a separate class called S1 and the workstations in a class called WS1. Later, we may find that we need more than one class to back up mercury, but we will evaluate other factors first. For now, the classes are as follows:

Class	Clients
S1	mercury (file server)
WS1	mars jupiter (workstations) neptune

3. Create classes to accommodate special storage requirements.



The storage unit and volume pool settings apply to all files that are backed up by the class. If files have special storage unit and volume pool requirements, create separate classes for them, even if other factors, such as schedules, are the same.

In our example (see below), we create a separate class (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 class that specifies a unique volume pool for those backups. Then, add the media for that volume pool as explained in the system administrator's guide for Media Manager.

Class	Clients	Files	Desired Storage
S1	mercury	/	8 mm
		/usr	
		/h001	
		/h002/projects	
S2	mercury	/h002/devexp	DLT
	mercury	/h002/desdoc	

4. Create additional classes 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 classes. For example, create different classes for night-shift and day-shift clients. In our example, we can back them all up during the same hours so no new classes are 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 class with an appropriate schedule and then assign the files and clients to that class.

In our example (see the next table), we place the root (`/`) file system on mercury in a different class (S3). The `root (/)` file system on the workstations is also placed in a separate class (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 file list, it is best if all files have similar retention requirements. Do not, for example, place files whose full backups must be retained forever in a class where full backups are retained for only a month.

In our example (see the next table), we place `/h002/desdoc` on mercury in a different class (S4). This is done because `/h002/desdoc` requires full backups every quarter and those backups must be retained for a much longer time than the other files on mercury.

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

5. Create separate classes for clients that require different general-attribute settings than other clients. Some general-attribute settings to consider are:
 - ◆ Class type. There are several types of classes and you must place the client in the correct type. For example, place Windows 2000 and NT clients in an MS-Windows-NT class.



- ◆ 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 put these clients in their own class 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 2000 clients in this class. 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-Window-NT classes).
- ◆ 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 back up time and make it unsuitable to use for all clients.
- ◆ Class priority. Use this attribute to control the order in which NetBackup starts its backups. The client in the higher priority class is backed up first.

There are also other general attributes that are explained later in this chapter. In our example, no extra classes are required because of general attribute settings.

6. Create separate classes 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 classes to reduce backup times to an acceptable level.

Compute the approximate backup time by multiplying the speed of the device by the amount of data in the backup. For example, if your backup device transfers data at 800 kilobytes per second, it takes 0.7 hours to back up 2 gigabytes.

The variable that is easiest to control here is the amount of data in the backup. NetBackup imposes no limits on backup size, but try to keep backups to less than 2 gigabytes. In addition to reducing backup time, shorter backups usually mean less time to recover files that are near the end of the backup.

In our example, it so happens that backing up `/usr`, `/h001`, and `/h002/projects` on mercury takes too much time so we create a new class for `/h002/projects`. This new class (S5) has the same requirements as S1 but we can now back up `/h002/projects` separately thus reducing backup time. The table below shows the final set of classes.

In addition to reducing the backup time for each class, putting the files in separate classes can reduce the total backup time for the server mercury. NetBackup processes files within a file list serially and in the order they appear in the file list. However, separate classes are backed up in parallel if enough drives are available and the

maximum jobs attributes are set to allow it (“Number of Streams That Can Run Concurrently” on page 84 has an explanation of maximum jobs settings that also applies to this discussion).

Multiplexing and Allow Multiple Data Streams also allow backing up classes in parallel (see “Allow multiple data streams” on page 82 and “Multiplexing” on page 275).

Note For best performance with multiple data streams, 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.

Class	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
S1	mercury	/usr /h001	high	8 mm	Daily Incr Monthly Full Monthly Full
S2	mercury	/h002/devexp	high	DLT	Daily Incr Weekly Full Monthly Full
S3	mercury	/	low	8 mm	Daily Incr Weekly Full
S4	mercury	/h002/desdoc	high	DLT	Daily Incr Weekly Full Monthly Full Quarterly Full
S5	mercury	/h002/projects	high	8 mm	Daily Incr Weekly Full Monthly Full
WS1	mars	/usr /home	high	8 mm	Daily Incr Weekly Full
	jupiter	/usr /home			Monthly Full
	neptune	/usr /home /var			



Class	Clients	Files	Frequency of Change	Desired Storage	Auto Backup Frequency
WS2	mars	/	low	8 mm	Daily Incr
	jupiter	/			Monthly Full
	neptune	/			

Planning Worksheet

The next two figures show a blank copy of a worksheet that will be useful for planning. Following the blank copy is a completed example. The previous procedure on planning and the remaining topics in this chapter provide information about the items on the worksheet.



Class Planning Worksheet (sheet 1)

Class: _____

Clients: _____

File List:

Class storage unit (if yes specify label _____)Class volume pool (if yes specify label _____)

General Attributes:

Active

Class type _____

Compression (applies to UNIX and Microsoft Windows clients only)Cross mount points (applies to UNIX and Windows 2000 clients only)Backup network drives (applies to Microsoft Windows clients only)Follow NFS (applies to UNIX clients only)Limit jobs per class (if yes specify value _____)

Job priority _____

True image recovery With move detection

Keyword phrase _____

Allow multiple data streams 

Class Planning Worksheet (sheet 2)

Schedules

	1	2
Schedule name	_____	_____
Type of backup	_____	_____
Frequency	_____	_____
Retention	_____	_____
Override Class storage unit	<input type="checkbox"/> label _____	<input type="checkbox"/> label _____
Override Class volume pool	<input type="checkbox"/> label _____	<input type="checkbox"/> label _____

Media multiplexing _____

Backup Times:

	Start	Duration	Start	Duration
Sunday	_____	_____	_____	_____
Monday	_____	_____	_____	_____
Tuesday	_____	_____	_____	_____
Wednesday	_____	_____	_____	_____
Thursday	_____	_____	_____	_____
Friday	_____	_____	_____	_____
Saturday	_____	_____	_____	_____
Sunday	_____	_____	_____	_____



Class Planning Worksheet (sheet 1)

Class: W2 on server bunny

Clients: mars (RS6000/AIX), jupiter (Solaris), neptune (HP)

File List: /usr, /home, /var

Class storage unit (if yes specify label TS_8)

Class volume pool (if yes specify label Backups)

General Attributes:

Active

Class type Standard

Compression (applies to UNIX and Microsoft Windows clients only)

Cross mount points (applies to UNIX and Windows 2000 clients only)

Backup network drives (applies to Microsoft Windows clients only)

Follow NFS (applies to UNIX clients only)

Limit jobs per class (if yes specify value _____)

Job priority 0

True image recovery **With move detection**

Keyword phrase _____

Allow multiple data streams



Class Planning Worksheet (sheet 2)

Schedules

	1	2
Schedule name	<u>W2DailyIncr</u>	<u>W2WeeklyFull</u>
Type of backup	<u>Differential Incr</u>	<u>Full</u>
Frequency	<u>1 day</u>	<u>1 week</u>
Retention	<u>1 week</u>	<u>1 month</u>

Override Class storage unit label _____ label _____

Override Class volume pool label _____ label _____

Media multiplexing 1 1

Backup Times:

	Start	Duration	Start	Duration
Sunday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Monday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Tuesday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Wednesday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Thursday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Friday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Saturday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>
Sunday	<u>22:00</u>	<u>8</u>	<u>22:00</u>	<u>8</u>



Using the Backup Policy Wizard

The easiest way to configure a class is to use the Backup Policy Configuration wizard. This wizard guides you through the setup process and simplifies it by automatically choosing default values that are good for most configurations. There are two ways to start the wizard:

- ◆ From NetBackup Assistant. Start NetBackup Assistant, by clicking Assistant on the Start menu in the NetBackup Administration window. Then, in the NetBackup Assistant window, click Create a Backup Policy to start the wizard.
- ◆ When adding a new class through the Backup Policy Management utility. Here, select the Use Add Class Wizard check box in the Add a New Class dialog box. This is explained in “To Add a New Class” on page 58.

After adding a class with the wizard, you can change any of the values by using the Backup Policy Management utility.

To Start the Backup Policy Management Utility

1. Start the NetBackup Administration - Java interface program. For instructions, see “NetBackup Administration Interfaces” on page 5.

The NetBackup Administration window appears.

2. In the NetBackup Administration window, click Backup Policy Management.

The Backup Policy Management (Classes) window appears.

The Backup Policy Management Window

The Backup Policy Management (Classes) utility has tools for configuring and managing classes. The topics below provide an overview of this window:

- ◆ Tree and Detail Views
- ◆ Menu Bar
- ◆ Toolbar
- ◆ Shortcut Menus



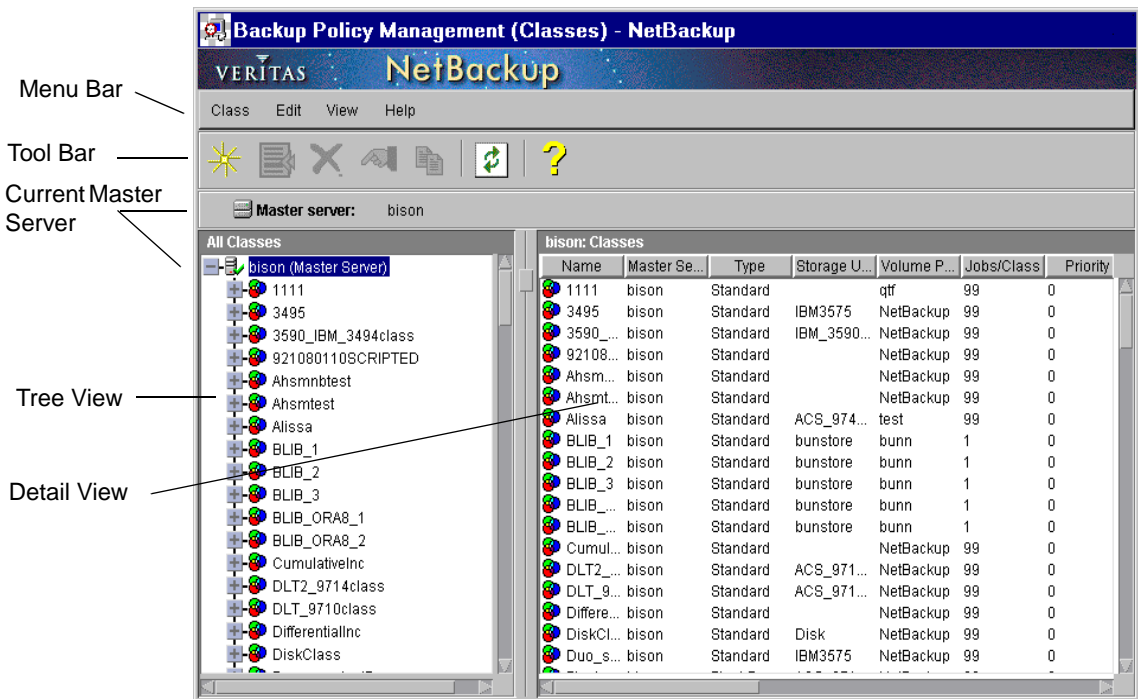
Tree and Detail Views

The window has two panes where you can select items and perform actions on them. The left pane is a hierarchical, tree view. The right pane is a detail view that shows information about whatever is selected in the left pane. The contents of these panes are different, depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters).

See Appendix E for more information on the Global Data Manager option.

The View if You Are Not On a Master of Masters

If you are not administering NetBackup from a Master of Masters, the left pane is a tree view of the classes on the master server that you are currently managing.



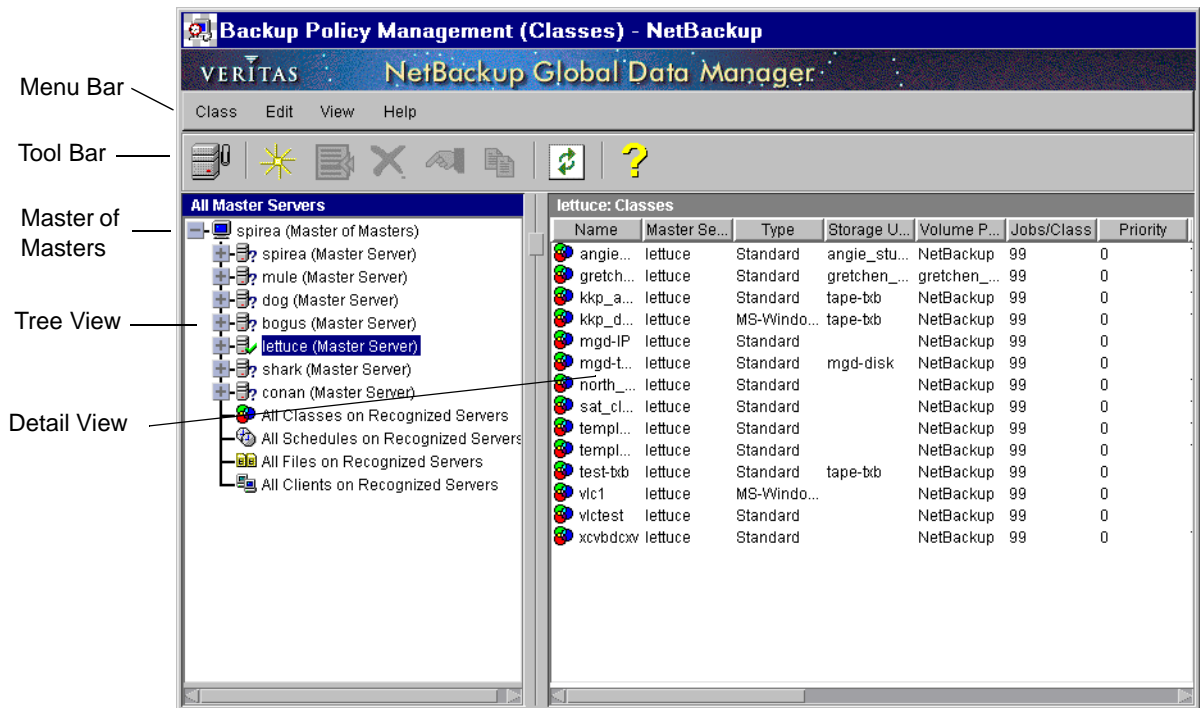
- ◆ If you expand the master server, the right pane shows information about all classes on the master server that you are currently managing.
- ◆ If you select one of the All Schedules, Files, or Clients nodes in the tree, the right pane shows a list of all the schedules, files, or clients for all classes on the master server.
- ◆ If you select a class in the tree, the right pane shows the general attributes for that class.

- ◆ If you expand the class in the tree and then select Attributes, Schedules, Clients, or Files under that class, the right pane shows details for the selected item.

To manage classes on another master server, change to that server as explained in “Choosing the Master Server for Backup Policies” on page 57.

The View From a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, the left pane is a tree view of the master servers you can currently manage. To append other master servers to the tree, use the Append Master Server command on the Class menu.



If you select the Master of Masters in the tree (default), the right pane lists all the media servers belonging to that master. If you select another master server, the right pane lists all the classes that are configured on that master.

You can expand each server in the tree to list the classes on that server. If you then expand a class and select Attributes, Schedules, Clients, or Files under that class, the right pane shows details for the selected item.

The bottom four “All ...” branches of the tree allow you to see combined lists of the classes and attributes that you can administer from this Master of Masters. For example, to see a list of all classes, select the All Classes on Recognized Servers branch of the tree. The

detail view now shows all classes on all *recognized* master servers, their general attributes, and the master server where they are configured (also see Recognize Master Server(s) and Ignore Master Server(s) in the “Menu Bar” topic).

Menu Bar

The following table describes the menus and commands that are available on the menu bar:

Menu	Commands
Class	<p>Activate - Activates the class that is selected in the left pane of the window. A class must be active for NetBackup to execute automatic backups or allow user backups or archives. This setting has no effect on restores.</p> <p>Deactivate - Deactivates the selected class (see Activate above).</p> <p>Manual Backup - Displays a dialog box where you can start a manual backup of schedules and clients for a specific class.</p> <p>Change Server - (Applies only when Global Data Manager is not installed) Displays a dialog box where you can specify the NetBackup server that has the configuration you want to modify.</p> <p>Append Master Server - (Applies only to Global Data Manager) Displays a dialog box where you can specify a master server to add to the tree. The addition applies only to the current session and the server will not be in the tree the next time you start the administrator interface. Also, see “Append Master Server” on page 670.</p> <p>Exit - closes this window.</p>
Edit	<p>Copy - Copies the selected item to a new location (applies only to classes and schedules).</p> <p>New - Displays a dialog box where you can specify criteria for a new class, schedule, client, or file.</p> <p>Insert - Applies to the file list and inserts a line ahead of the one you have selected. This is useful when adding directives.</p> <p>Change - Displays a dialog box where you can specify changes to the selected class attributes, client, file, or schedule.</p> <p>Delete - Deletes the selected class, client, file, or schedule.</p> <p>Install UNIX Client Software - Displays a dialog box where you can choose to install software on UNIX clients.</p>

Menu	Commands
View	<p>Preferences - Displays a dialog box where you can specify your display preferences for the toolbars.</p> <p>Ignore Master Server(s) - (Applies only to Global Data Manager) Ignores connections to the master server(s) that you designate so the Master of Masters does not attempt connections to them. The branch of the tree for an ignored server and the associated information in the detail view appears in faded type. You cannot select an ignored server or refresh the display with updated configuration information. The only operation you can perform on an ignored server is to recognize it (see below).</p> <p>Recognize Master Server(s) - (Applies only to Global Data Manager) Recognizes an ignored master server so it can be selected and the display refreshed with updated configuration information. Also see “Recognize Master Server(s)” on page 672.</p> <p>Refresh - Updates the information in the window with new information retrieved from the master server(s).</p>
Help	<p>Help Topics - Provides online help information.</p> <p>About Backup Policy Management - Displays program information, version number, and copyright.</p>

Toolbar

The buttons on the toolbar provide shortcuts for menu commands. If the toolbar is not visible or you want to change the size of its buttons, click the Preferences command on the View menu and make the desired choices in the Preferences dialog box.

Shortcut Menus

Clicking the right mouse button when the pointer is over a pane in the window, displays a popup menu with commands that apply to what is currently selected. These commands also appear on the menu bar (see “Menu Bar”).

Choosing the Master Server for Backup Policies

If there is more than one master server, choose the one where you want to manage the backup policy configuration.

The procedure for choosing a server to manage is different, depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters).



If You Are Not On a Master of Masters

If you are not administering NetBackup from a Master of Masters, the name of the server that you are currently managing appears on the Master Server line above the left pane of the window. To manage the configuration on another NetBackup master server:

1. On the Class menu, click Change Server (or click the button that appears when you position the pointer over the master server name below the toolbar).
2. In the dialog box, specify the name of the NetBackup server where you want to manage the class configuration and click OK.

If you encounter problems, see “If You Cannot Change to Another Master Server” on page 23.

If You Are On a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, choose a master server by selecting it in either the tree or detail view.


- ◆ The tree view in the left pane shows the master servers that you can currently manage. Click on a server to select it.
- ◆ If you select the Master of Masters at the top of the tree, the right pane shows the same list of master servers as the tree. In the right pane, double click on a server to select it.

If the server you want is not displayed, append it by using the Append Master Server command on the Class menu. If you encounter problems, see “If You Cannot Change to Another Master Server” on page 23.

To Add a New Class

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. If your site has more than one master server, choose the one where you want to add the class (see “Choosing the Master Server for Backup Policies” on page 57).
3. In the Backup Policy Management (Classes) window, click New on the Edit menu.

The Add a New Class dialog box appears.

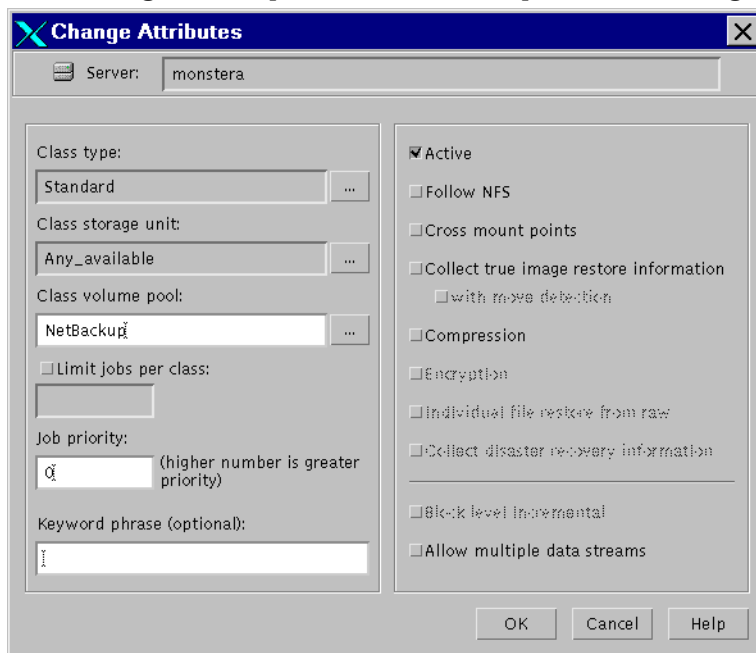


- a. In the Class Name box, type a unique name for the new class.
Use alphabetic (ASCII A-Z a-z), numeric (0-9), plus (+), minus (-), underscore(_), or period (.) characters. Do not use a minus as the first character or leave any spaces between characters.
 - b. Choose whether to use the wizard for configuring the class. The wizard guides you through the setup process and simplifies it by automatically choosing default values that are good for most configurations. If necessary, you can change the defaults later by editing the class.
 - ◆ To use the wizard, select the Use add class wizard box and click OK. The wizard starts and you create the class by following the prompts (skip the rest of this procedure).
 - ◆ If you require more control over the settings than the wizard provides, then clear the box and proceed to step 4.
4. Click OK.

A dialog box appears where you can specify the general attributes for the class.



5. In the dialog box, complete the entries as explained in “Setting the General Class



Attributes” on page 63 and click OK to close the dialog box.

6. Add schedules, clients, and files as explained below.

To Add Schedules to the Class:

- a. Under the class name in the left pane, select Schedules and then click New on the Edit menu.
- b. In the Add Schedules dialog box, complete the entries as explained in “Schedule Properties” on page 124.
- c. If this is the last schedule, click OK. To add more schedules, click Add and repeat step b. Click Close to cancel changes that you have not yet added and close the dialog box.

To Add Clients to the Class:

- a. Under the class name in the left pane, select Clients and then click New on the Edit menu.
- b. In the Add Clients dialog box, complete the entries as explained in “To Add Clients to a Class” on page 86 starting with step 6.

To Add Files to the Class:

- a. Under the class name in the left pane, select Files and then click New on the Edit menu.
- b. In the Add Files dialog box, specify the file list as explained in “To Add Files For Automatic Backups” on page 90 starting with step 5.

To Change a Class

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

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. If your site has more than one master server, choose the one that has the class you want to modify (see “Choosing the Master Server for Backup Policies” on page 57).
3. Find the class name in the left pane and click the + to the left of the name.
This expands the tree to show Attributes, Schedules, Files, and Clients below the selected class name.
4. Make your changes as explained below (or use online help in the dialog box).

To Add Schedules, Files, or Clients

- a. Under the class name in the left pane, select Schedules, Files, or Clients.
- b. On the Edit menu, click New. For instructions, see the online help for the dialog box or refer to the following:
 - ◆ For Schedules see, “To Add New Schedules” on page 123 starting with step 5.
 - ◆ For Files, see “To Add Files For Automatic Backups” on page 90 starting with step 5.
 - ◆ For Clients, see “To Add Clients to a Class” on page 86 starting with step 6.

Your additions (for example, files) are appended to the end of the existing list.



To Change Attributes, Schedules, Files, or Clients

- a. Under the class name in the left pane, select Attributes, Schedules, Files, or Clients.
- b. Select the item in the right pane and click Change on the Edit menu.
- c. In the Change dialog box, make your changes.
- d. Click OK to close the dialog box and update the configuration.

To Delete Schedules, Files, or Clients

- a. Under the class name in the left pane, select Schedules, Files, or Clients.
- b. In the right pane of the window, select the item (for example, a file).
- c. On the Edit menu, click Delete.

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.

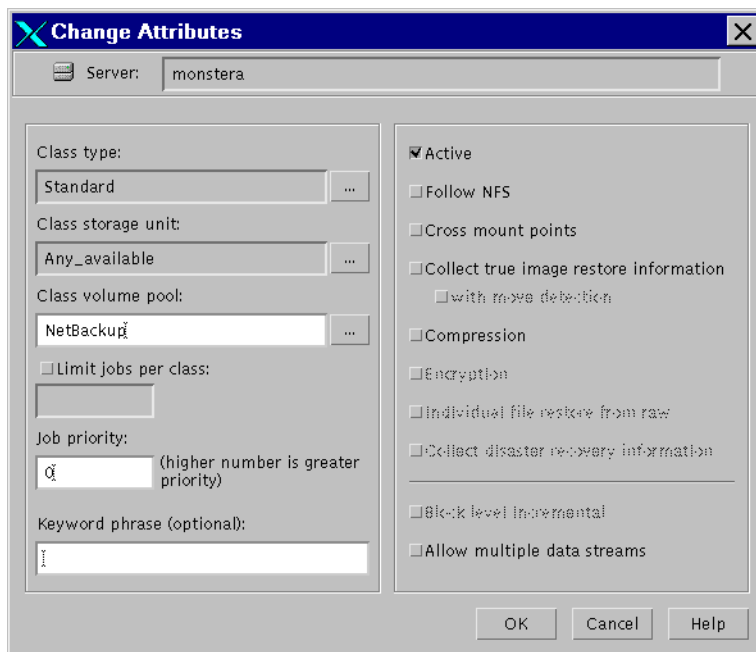
Also, deleting a file only deletes the file from the list of files designated for automatic backup. It does not delete the actual file from the disk.

To Insert Within the File List

- a. Under the class name in the left pane, select Files.
- b. In the right pane, select the point in the list where you want to insert the entry.
- c. On the Edit menu, click Insert.
The Insert File dialog box appears.
- d. Click in the Pathname or Directive box and specify the pathname for the file or directory.
 - ◆ If you are unfamiliar with how to specify file paths for your clients, read “Rules for Backup File Paths” on page 95 before proceeding.
 - ◆ If you are using directives in the file list, see “Adding Directives to the File List” on page 91.
- e. Click the Add to List button at the right of the Pathname or Directive box.
The new entry appears in the list.
- f. When you have made all your additions, click the Insert in File List button.
Your additions (files or directives) are inserted before the selected entry.

Setting the General Class Attributes

The general attributes determine the basic characteristics of all the backups that NetBackup performs for this class. For a new class, you set these attributes in the dialog box below. This dialog box appears when you first create the class or if you select the attributes later to change them.



The attributes you can specify depend on the type of class you are configuring and the options that are installed. For example, Encryption is available only with the NetBackup Encryption option. The following topics explain all the general attributes.

Class type

Determines the type of clients that can be in the class and in some cases the types of backups that can be performed on those clients. Select the type of class by clicking the button to the right of the box and then the desired class type (Table 2). If you change the class type for an existing class that has schedules that are invalid for the new class type, NetBackup prompts you and then either deletes the invalid schedules or if possible changes them to an equivalent type.



Table 2. Class Types

Class Type	Description
DB2	Use when the class will have only clients with the NetBackup for DB2 option. See the guide for that option for information on setting up this class type.
Extensible-Client	A class type reserved for use by VERITAS or its partners to provide agents for new database types.
Lotus-Notes	Use when the class will have only clients with the NetBackup for Lotus Notes option. See the guide for that option for information on setting up this class type.
MS-Windows-NT	Use when the class will have only Windows 2000 or NT clients.
MS-Exchange-Server	Use when the class will have only clients with the NetBackup for MS-Exchange option. See the guide for that option for information on setting up this class type.
MS-SQL-Server	Use when the class will have only clients with the NetBackup for MS-SQL Server option. See the guide for that option for information on setting up this class type.
NCR-Teradata	Use when the class will have only clients with the NetBackup for Teradata option. See the guide for that option for information on setting up this class type.
NetWare	Use when the class will have only nontarget NetBackup Novell NetWare clients (this version uses a Microsoft Windows interface).
NDMP	Use when the class will have only clients with the NetBackup for NDMP option. This class is available only when the NetBackup NDMP option is installed. See the guide for that option for information on setting up this class type.
OS/2	Use when the class will have only NetBackup OS/2 clients.
Standard	Use when the class will have any combination of the following: <ul style="list-style-type: none">◆ Windows 98 or 95 clients.◆ Macintosh clients.◆ NetBackup Novell NetWare clients that have the target version of NetBackup software.◆ UNIX clients, except those covered by other special classes such as Oracle.

Table 2. Class Types (continued)

Class Type	Description
Note: The following class types apply only to UNIX clients.	
AFS	Use when the class will be backing up only AFS file systems on clients. See the AFS appendix in this manual for information on setting up these classes.
Apollo wbak	Use when the class will have only Apollo clients.
Auspex-FastBackup	Use when the class will have only Auspex clients that you will be backing up by using the Auspex-FastBackup API. An Auspex FastBackup client must be on a NetBackup UNIX server.
DataTools-SQL-BackTrack	Use when the class will have only clients with the NetBackup for DataTools-SQL-BackTrack option. See the guide for that option for information on setting up this class type.
FlashBackup	Use when the class will have only NetBackup FlashBackup clients. This class is available only when the NetBackup FlashBackup option is installed. See the guide for that option for information on setting up this class type.
Informix-On-BAR	Use when the class will have only clients that are running the NetBackup for Informix option. See the guide for that option for information on setting up this class type.
Split-Mirror	Use when the class will have only clients with the NetBackup for EMC option. See the guide for that option for information on setting up this class type.
Oracle	Use when the class will have only clients with the NetBackup for Oracle option. See the guide for that option for information on setting up this class type.
SAP	Use when the class will have only clients with the NetBackup for SAP option. See the guide for that option for information on setting up this class type.
Sybase	Use when the class will have only clients with the NetBackup for Sybase option. See the guide for that option for information on setting up this class type.

Class storage unit

Specifies the default storage unit for backups of this class and NetBackup uses it for all schedules that do not specify another storage unit. A schedule-level storage unit (when specified) overrides the class default (see “Override class storage unit” on page 129).



To select the class storage unit, click the button to the right of the box to display a selection list. Then, select either a specific storage unit or Any Available.

If you select Any Available, NetBackup tries storage units in alphabetical order and uses the first one that meets the following requirements:

- ◆ Is Not “on demand only”
- ◆ Has available drives
- ◆ Has media available in the required volume pool

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 class level and specify the following on the schedules:

- ◆ For schedules that can use the Tape Stacker 8MM, clear Override class storage unit. When these schedules run, NetBackup uses a Tape Stacker 8MM.
- ◆ For the schedule that requires DLT, select Override class storage unit and select Tape Library DLT. When this schedule runs, NetBackup overrides the class 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 class 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 classes and schedules that require the storage unit.
- ◆ An Asepex FastBackup client and the storage unit must always be on a NetBackup UNIX server. See “Overview of Asepex FastBackup Classes” on page 613.
- ◆ If your NetBackup configuration has several storage units and you want a class to use *more than one but not all* of the storage units, perform the following:
 - a. When you configure volumes in Media Manager, define a volume pool and volumes that are available only to the desired storage units.
 - b. For the class, set Class volume pool to the volume pool defined in step a.

- c. For all classes, set Class storage unit to Any Available.

Class volume pool

Specifies the default volume pool for backups of this class and NetBackup uses it for all schedules that do not specify another volume pool. A schedule-level volume pool (when specified) overrides the class default (explained later under “Override class volume pool” on page 129). If you do not name a volume pool for either the class or the schedule, NetBackup uses the “NetBackup” pool.

To specify the class volume pool, click the button to the right of the box and select the desired volume pool name from the list. The list shows all previously configured volume pools. You can also type the name of the pool in the box and then configure it in Media Manager.

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 Class volume pool to *backups* When you set up the schedules for the class, set Override class volume Pool as follows:

- ◆ For schedules that use the *backups* volume pool, clear Override class volume pool.
- ◆ For the schedule that requires the *archive* volume pool, select Override class 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.



Limit jobs per class

If the Limit jobs per class check box is clear, the maximum number of backup and restore jobs that NetBackup will perform concurrently for this class can be up to a limit of 999. To specify a lower limit, select the check box and specify a value from 1 to 999 (the default is 99).

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 class

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

- ◆ Number of storage devices available. 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. 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 class high enough to support the specified level of multiplexing.

Lower values can limit multiplexing within a class if there are jobs from different schedules within that class. For example, if Limit jobs per class 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.

For more information on multiplexing, see “Multiplexing” on page 275.

- ◆ Limit jobs per class does not prevent concurrent jobs if the jobs are from different classes.

For example, if there are three classes and each has its Limit jobs per class at 2, NetBackup can start two jobs from each class and have a total of six running at one time.

Job priority

Specifies the priority that NetBackup assigns to automatic-backup jobs for this class. When a drive becomes available, NetBackup assigns it to the first client in the highest priority class. For an explanation of all the factors involved in establishing backup priority, see “Factors Affecting Backup Time” on page 618.

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.

Keyword phrase (optional)

Specifies a keyword phrase that NetBackup will associate with all backups or archives for this class. Users on Windows NT/2000 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 other NetBackup clients do not support keyword phrases.

You can use the same keyword phrase for more than one class. This makes it possible to link backups from related classes. 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.

Users on Windows NT/2000 and UNIX clients can also specify a keyword phrase for a user backup or archive. A user phrase overrides the class phrase.

Active

To activate the class, select the box. The class must be active for NetBackup to execute automatic-backup schedules or allow user backups or archives.



Leaving a class inactive is useful if there are problems and you want to suspend backups until they are resolved. It is also useful if you are creating new classes but cannot complete them right away.

True image restore information

Note The True image restore information attribute applies only to certain class types and NetBackup allows you to select it only in those instances.

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 class. 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*.

With move detection

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.

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 class. 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.

Example of 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 12/01/1995 and 12/04/1995. Assume that the True image restore information attribute was selected for the class that did the backups.

Day	Type of Backup	Files Backed Up in <code>/home/abc/doc</code>							
12/01/1995	Full	file1	file2	dirA/fileA	dirB/fileB	file3			
12/02/1995	Incremental	file1	file2	dirA/fileA	-----	-----			
12/03/1995	Incremental	file1	file2	dirA/fileA	-----	-----			
12/04/1995	User backup	file1	file2	dirA/fileA	-----	-----	dirC/fileC	file4	
12/04/1995	Incremental	file1	file2	-----	-----	-----	-----	file4	

Note: dashes (-----) mean that the file was deleted prior to this backup.

Also, assume that you are going to restore the 12/04/1995 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 12/01/1995 (last full backup) through 12/04/1995:

```
file1
file2
dirA/fileA
dirB/fileB
file3
```



dirC/fileC

file4

- ◆ If you do a true-image restore of the 12/04/1995 backup, the restored directory has only the files and directories that existed at the time of the incremental backup on 12/04/1995:

file1

file2

file4

NetBackup does not restore *any* of the files deleted prior to the 12/04/1995 incremental backup.

The restored directory does not include the `dirA` and `dirC` subdirectories, even though they were backed up on 12/04/1995 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 using the How long to keep TIR Information global attribute. See “How long to keep TIR information” on page 100.
- ◆ Incremental backups are slower for a class where true-image restore information is being collected.
- ◆ If you are using the indexing feature (see “Reduce Restore Times by Indexing the Image Catalog” on page 221), the INDEX files take much more space when you are collecting true-image restore information.
- ◆ You can perform true-image restores only on directories that were backed up by a class 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 class 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 12/04/1995, but before doing the restore. In this case, the contents of the directory after the restore is:

```
file1  
file2  
file4  
file5
```

Backup network drives

Note The Backup network drives attribute applies only to certain class types and NetBackup allows you to select it only in those instances.

Specifies that you want NetBackup to back up or archive files that are on network drives and are named in the file list (or by the user in the case of a user backup). With this attribute selected, NetBackup also backs up CD-ROM drives on Windows NT/2000 systems. Clear the box to prevent the backup or archive of network drives or backup of CD-ROM drives.

The network (shared) drives must be available to the service account that the NetBackup Client service logs into at startup. By default, the startup account is set to System. You must change this account on each Windows NT or 2000 client where you are backing up network drives.

Note Windows 95 and 98 do not have services so changing the account does not apply.

To change the account on a Windows NT client:

1. Start the Services application in the Control Panel.
2. Select the NetBackup Client service.
3. Click the Startup button and specify an account that has read permissions for the share to be backed up. The account must have write permission to perform restores.
4. Stop and start the NetBackup Client Service so the new account will take effect.

To Change the account on a Windows 2000 client:

1. From Windows Explorer or the desktop, right click on My Computer.
2. Click Manage.



3. Expand Services and Applications in the left pane.
4. Click Services.
5. Double-click the NetBackup Client service and select the Log On tab.
6. Specify an account that has read permissions for the share to be backed up and click OK. Note that the account must have write permission to perform restores.
7. Select the General tab and stop and start the service.

Example 1 - User Backup

Assume that:

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

The steps to perform are as follows:

1. On `wildrice`, the NetBackup master server, select Backup network drives for the class to be used for the backup.
2. On `buck`, the NetBackup client:
 - a. Map a drive to `\\pepper\share` (for example, `F:`)
 - b. Change the NetBackup Client Service on `buck` to either *Start Up* or *Log On* with the same account as the user that will do the backup. This user account must have read permissions for the share that is to be backed up. The account must have write permission to perform restores.
 - c. Stop and start the NetBackup Client Service so the new account takes effect.

Note The user must log directly into the NetBackup client. It does not work to log in using Terminal Services Client because this does not map drives to the shared folders and the NetBackup Client Service does not have access to the mapped drives.

3. The user can now perform the backup as follows:
 - a. Start the Backup, Archive, and Restore program on `buck`
 - b. Open a backup window and mark the mapped drive (`F:` in our example)
 - c. Start the backup.

Example 2 - Manual backup of a client from the server

Assume that:

- ◆ wildrice is the NetBackup master server.
- ◆ buck is a Windows NT NetBackup client.
- ◆ pepper is a Windows NT computer (not necessarily a NetBackup client) and has a shared folder named `share`.
- ◆ You want to schedule automatic backups for buck that will back up the folder named `share` on pepper.

The steps to perform are as follows:

1. On wildrice, the NetBackup server:
 - a. Select Backup Network Drives in the class to be used for the backup.
 - b. Add the UNC path name to the file list for the class that to be used for the backup. In our example this path is `\\pepper\share`.
2. On buck, the NetBackup client:
 - a. Change the NetBackup Client Service to either *Start Up* or *Log On* with the account for a user that has read permissions for the share that is to be backed up. The account must have write permission to perform restores from buck.
 - b. Stop and start the NetBackup Client Service so the new account takes effect.
3. A full or incremental backup of buck now includes the folder named `share` on pepper. A manual backup of the class or client (from the server) will also back up the `share` folder.

Follow NFS

Note The Follow NFS attribute applies only to certain class types and NetBackup allows you to select it in only those instances.

Specifies that you want NetBackup to back up or archive any NFS mounted files that are named in the file 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 setting of Cross mount points (explained later in this chapter).



- ◆ Follow NFS does not affect Apollo clients. These clients always behave as if the box is selected. Do not put NFS-mounted files in the file list for Apollo wbak classes unless you want to back up those files.
- ◆ 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 Following NFS Mounts

As a general rule, do not back up NetBackup clients over NFS. It is best to backup 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 class 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, it continues to retry for five seconds (by default) before assuming the file system to be unavailable. To change the five second value, use the `bp.conf` file option `NFS_ACCESS_TIMEOUT`.

Advantages of Following 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

Note The Cross mount points attribute applies only to certain class types and NetBackup allows you to select it in only those instances.

Controls whether NetBackup will cross 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 2000 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.
- ◆ Cross mount points does not affect Apollo clients. These clients always behave as if Cross mount points is selected.
- ◆ Do not use Cross mount points in classes where you use the `ALL_LOCAL_DRIVES` directive in the file list.

Cases That Can Require Separate Classes

In some cases, it is best to create separate classes according to whether you want to cross mount points. For example, to back up the root file system without also backing up files systems mounted on it, create a class where Cross mount points is not selected and the file list contains only root (/). Place other file systems in another class or classes.

To back up all the data on a client, create a class where Cross mount points is selected and the file list includes root (/).



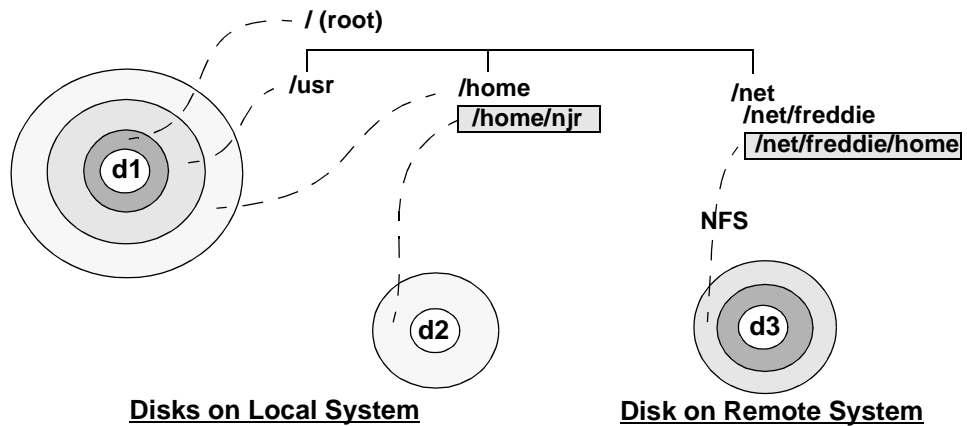
How Cross Mount Points Interacts With Follow NFS

To back up NFS mounted files, select Follow NFS. The table below summarizes the behavior of the Cross mount points and Follow NFS Attributes:

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 above concepts. 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, has a directory named /net/freddie/home that is NFS-mounted on /net/freddie.

Example 1

Assume that you clear the Cross mount points and Follow NFS attributes and have the following entries in the file list:


```
/
/usr
/home
```

In this case, NetBackup considers only the directories and files that are in the same file system as the file list entry it is processing. It does not back up `/home/njr` or `/net/freddie/home`.

Example 2

Assume that you select the Cross mount points and Follow NFS attributes and include only `/` in the file 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 file list backs up only `/usr` and individual files under `/`:

```
/usr
/individual_files_under_root
```

Compression

Note The Compression attribute applies only to certain class types and NetBackup allows you to select it in only those instances.

Specifies that software compression be used for backups of this class. Select the box to enable compression (the default is no compression).

Advantages of Using Compression

Compression reduces the size of a backup by reducing the size of files in that backup. This in turn 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). In addition, 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 are: stripped and unstripped binaries, ASCII, and repeating non-unique strings. If more of the data is favorable to compression you obtain more compression.

The following shows the specifications for 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.

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).
Worst-case compression	Files that are already compressed become slightly larger if compressed again. On UNIX clients, if you find these types of files and they have a unique file extension, exclude them by adding a <code>COMPRESS_SUFFIX = .suffix</code> option to the <code>bp.conf</code> file.
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 CPU time and as much memory as the administrator configures.

Effect on client speed	Compression uses as much of the CPU as available and affects other applications that require the CPU. 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 Kbytes, because that is the tar block size.</p> <p>On UNIX clients, files ending with suffixes specified with the <code>COMPRESS_SUFFIX = .suffix</code> option in the <code>bp.conf</code> file.</p> <p>On UNIX clients, files with the suffixes in the table below.</p>

.arc or .ARC	.pak or .PAK
.arj or .ARJ	.iff or .IFF
.au or .AU	.pit or .PIT
.cpt or .CPT	.pit.bin or .PIT.BIN
.cpt.bin or .CPT.BIN	.scf or .SCF
.F	.sea or .SEA
.F3B	.sea.bin or .SEA.BIN
.gif or .GIF	.sit or .SIT
.gz or GZ	.sit.bin or .SIT.bin
.hqx or .HQX	.tiff or .TIFF
.hqx.bin or .HQX.BIN	.Y
.jpeg or .JPEG	.zip or .ZIP
.jpg or .JPG	.zom or .ZOM
.lha or .LHA	.zoo or .ZOO
.lzh	.z or .Z

Encryption

Note Available only when the NetBackup Encryption option is installed and configured. See the *NetBackup Encryption System Administrator's Guide* for more information.

Specifies encryption for backups of clients in this class.



Block level incremental

Note Available only with VERITAS Oracle Edition. See the documentation that came with that software for more information.

Specifies block-level-incremental backups for clients in this class. Note that this option and Allow multiple data streams cannot both be selected for the same class (you can use only one or the other).

Allow multiple data streams

Specifies that, depending on directives in the file list, NetBackup can divide automatic backups for each client into multiple jobs, with each job backing up only a part of the file list. The jobs are in separate data streams and can occur concurrently.

- ◆ Whether multiple backup jobs start for each client and how the file list is divided into separate jobs is determined by the directives that you specify in the file list (see “File-List Directives for Multiple Data Streams” on page 112).
- ◆ The total number of streams and how many can run concurrently is determined by the following settings (see “Tuning Multiple Data Streams” on page 84):
 - ◆ Number of available storage units
 - ◆ Multiplexing settings
 - ◆ Maximum jobs parameters.

When to Use Multiple Data Streams


Reduce Backup Time

Multiple data streams can reduce the backup time for very 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, multiple data streams also provide a form of checkpoint restart. A single failure only terminates a single stream and NetBackup can restart the failed stream without restarting the others. This improves retry time when a backup failure occurs.

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 global attribute applies to each stream. For example, if it 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 Classes

When a configuration contains large file servers with many file systems and volumes, using multiple data streams will provide more backups with fewer classes 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.

Total Number of Streams

The file 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 “File-List Directives for Multiple Data Streams” on page 112).

Number of Streams That Can Run Concurrently

The number of streams that can run concurrently for a class 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 class 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 (global attribute)
- ◆ Limit jobs per class (class attribute)

- ◆ Maximum Jobs This Client (set with `bpclient` command `-max_jobs` option, see below)

The maximum job settings are interdependent as follows:

- ◆ If Maximum Jobs This Client is not set, the lowest value of Maximum jobs per client and Limit jobs per class is the limiting factor.
- ◆ If Maximum Jobs This Client is set, then NetBackup ignores Maximum jobs per client and uses the lowest value of Maximum Jobs This Client and Limit jobs per class as the limiting factor.

To specify a value for Maximum Jobs This Client with the `bpclient` command

1. Determine if the client is in the client database on the master server by executing (all on one line):

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

2. If the client is not in the client database, execute the following on the master server (all on one line):

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

3. If the client is in the client database, execute (all on one line):

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

Individual file restore from raw

Note The Individual file restore from raw class attribute applies only to Auspex FastBackup and FlashBackup type classes. These are available only when their associated NetBackup option is installed and configured.

The Individual file restore from raw attribute enables NetBackup to restore individual files from full or incremental raw-partition backups performed on clients in FlashBackup or Auspex FastBackup classes. NetBackup starts saving the necessary information with the next backup of the class after the attribute is selected (the attribute is always selected in a FlashBackup class).

You can restore either individual files or an entire partition.

- ◆ An entire raw partition can be restored only from a FlashBackup or Auspex FastBackup full backup. To restore entire raw partitions from a FlashBackup or Auspex FastBackup backup, you can use `jbpsA` and proceed in the same way as when restoring other raw-partition backups.



- ◆ The incrementals for these classes support only individual file restores. To restore individual files from an Auspex FastBackup or FlashBackup backup, you can use `jbpsa` and proceed in the same way as if you were restoring from a standard-file-system backup. You can select the files from either a full or an incremental backup.

Note For an Auspex FastBackup class, if you connect tape drives from a single robotic library to multiple Storage Processors (SPs) or to both a Storage Processor and the Host Processor (HP), then also perform the following steps.

1. Create a separate storage unit (same robot number) for each set of drives that connect to a different processor (HP or SP).
2. In the class for each processor, specify the storage unit that connects to that processor (you can do this either at the class or schedule level). The purpose of this step is to prevent the NetBackup scheduler from over-committing a specific SP or HP drive.

Example:

Assume you have an ACL 4/52 with two drives on SP0 and two drives on SP1.

1. Define two storage units (same robot number), each with two drives (for example, TLDSP0 and TLDSP1).
2. In the class that has that SP0 disks, specify TLDSP0 as the storage unit and in the class that has the SP1 disks specify TLDSP1 as the storage unit.

Collect disaster recovery information

Specifies that you want NetBackup to collect the information required for intelligent disaster recovery when it backs up Windows NT/2000 clients in this class. For more information, see “Configuring NetBackup Classes for IDR” on page 682.

To Add Clients to a Class

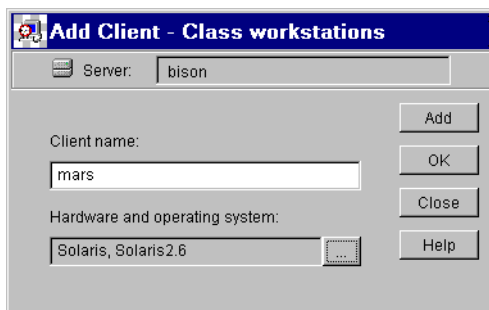
Note See “To Change a Class” on page 61 for instructions on inserting, changing, or deleting entries in an existing class.

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. Choose the master server where you want to add clients to a class (see “Choosing the Master Server for Backup Policies” on page 57).
3. Find the class name in the left pane and click the + to the left of that name.



4. This expands the tree to show Attributes, Schedules, Files, and Clients below the selected class name.
5. Under the class name, select Clients and then click New on the Edit menu.

A dialog box appears. The title bar shows the name of the class where you are adding the clients.



6. In the Client name text box, type the name of the client that you are adding.

Observe the following rules for assigning client names:

- ◆ 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.
 - ◆ If you put the client in multiple classes, use the same name in each class.
 - ◆ 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.
7. Click the Hardware and operating system list box and then select the desired entry in the list.

Add only clients with hardware and operating systems that this class supports. For example, do not add a Novell NetWare client to an MS-Windows-NT class. If you add the same client to more than one class, be sure to designate the same hardware and operating system in each of the classes.

Note If the desired hardware and operating system is not in the list, it means that the associated client software is not installed on the server. Check the `/usr/opensv/netbackup/client` directory for the directories and software corresponding to the client you are trying to install. If the directories or software are 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).



8. If this is the last client, click OK. If you are going to add more clients click Add and repeat step 6 and step 7. Click Close to cancel changes that you have not yet added and close the dialog box.

Installing Client Software on Trusting UNIX Clients

You can install client software on trusting UNIX clients by using the administration interface 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 box when starting the interface. This server must also be the master where you are currently managing backup policies and clients must be in a class on this master.

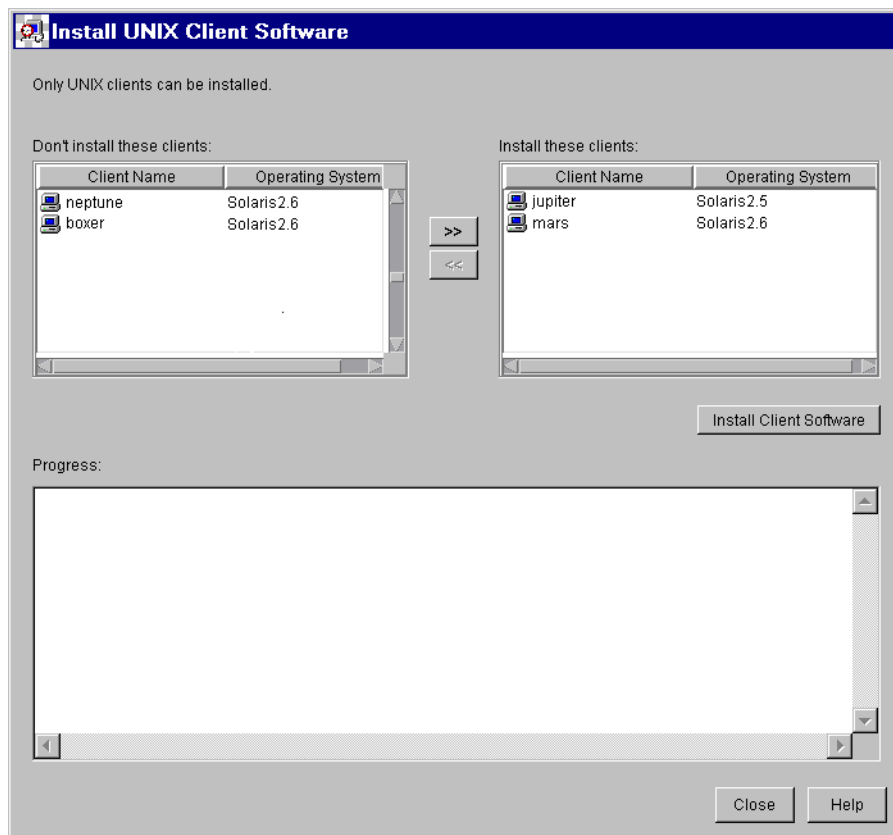
For example, assume you want to install clients that are in a class on a master server named shark. Here, you must have specified shark in the Login dialog box and therefore be managing NetBackup through the Java application server on this system. shark must also be the master server you are currently 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 box.

- ◆ 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.

The procedure for installing the UNIX client software is as follows:

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. Choose the master server where you want to add clients to a class (see “Choosing the Master Server for Backup Policies” on page 57).
This must be the UNIX system where you are running the NetBackup Java interface.
3. With the master server selected in the left pane, click Install UNIX Client Software on the Edit menu.

The Install UNIX Client Software dialog box appears.



4. In the Don't install on these clients box, select the clients you want to install and click the right arrows.

The clients are moved to the Install these clients box.

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 class. 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 and then create `/usr/opensv/netbackup` as a link to that directory prior to installing software.

6. When the install is complete, click Close.

Installing Software on Secure UNIX Clients

As defined here, a *secure* UNIX client is one that does not have 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 getting started guide that came with your NetBackup server software.

Installing Software on PC Clients

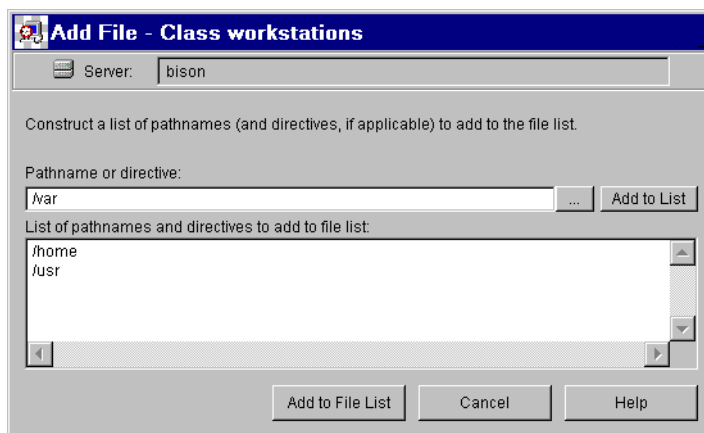
You install NetBackup PC client software by using the same CD-ROM that has your server software. For instructions, see the NetBackup getting started guide that came with your NetBackup server software or the NetBackup installation guide for PC clients.

To Add Files For Automatic Backups

Note See “To Change a Class” on page 61 for instructions on inserting, changing, or deleting entries in an existing class.

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. Choose the master server where you want to add files to a class (see “Choosing the Master Server for Backup Policies” on page 57).
3. Find the class name in the left pane and click the + that is to the left of that name.
This expands the tree to show Attributes, Schedules, Files, and Clients below the selected class name.
4. Under the class name, select Files and then click New on the Edit menu.
The Add File dialog box appears. The title bar shows the name of the class to which you are adding the files.





5. If you are unfamiliar with how to specify file paths for your clients, read “Rules for Backup File Paths” on page 95 before proceeding.
6. Click in the Pathname or directive box and specify the pathname for the file or directory.
7. Click the Add to List button at the right of the Pathname or directive box.
The new entry appears in the list.
8. Repeat step 6 and step 7 to add other pathnames.
9. If you are using directives in the file list, see “Adding Directives to the File List” on page 91.
10. When the list is complete, click the Add to File List button at the bottom of the dialog box.
11. Verify the file list as explained in “Verifying the File List” on page 94.

Adding Directives to the File List

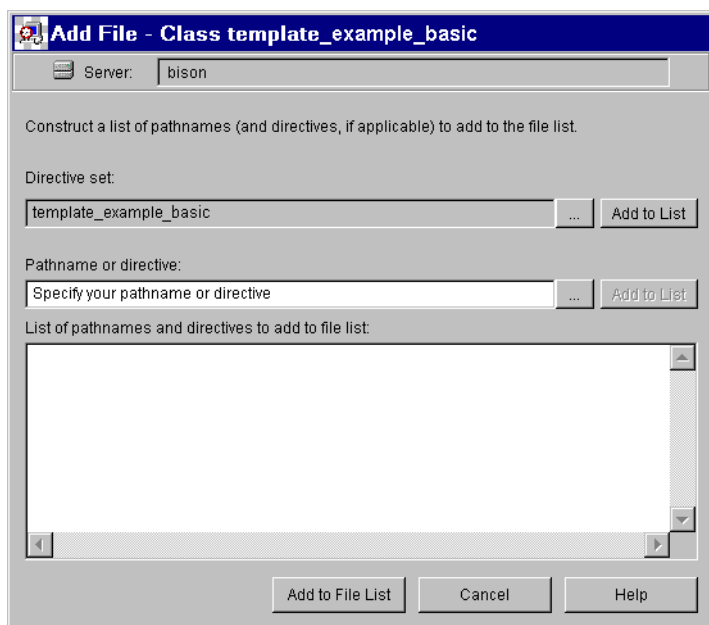
The directives that you can use for a class depends on the class type. There are two ways to add directives to the file list.

- ◆ Add them individually in the same way as other entries. This method is required if a directives template does not exist.
- ◆ Select them from a directives template supplied with a NetBackup option.



Some separately-priced NetBackup options come with one or more templates that contain directives for the class types that you can configure for that option. If a template is installed, the add and change files dialog boxes for the associated class contains a Directive set text box where you can select the template that has the directives you want to use (for example there may be basic and advanced sets).

Note If the Directive set text box is not present, it means that templates are not installed or do not apply to the class type you are configuring. The NetBackup guide for each option explains how to use any special directives that apply to that option.



The following explains both methods for adding directives. For more information on what directives do, see “File-List Directives - General Discussion” on page 111 and “File-List Directives for Multiple Data Streams” on page 112 (if Allow multiple data streams is enabled). For separately-priced options, also see the NetBackup guide that came with the option.

To Add Directives In the Same Way as Path Names

1. Click the ... button to the right of the Pathname or directive box and select from the list of directives (if any) that apply to the class type you are configuring.

Or

Type the name of the directive in the Pathname or directive box. For example, type ALL_LOCAL_DRIVES.


2. Click Add to list.

The directive is appended to the end of the list of additions. Use the keyboard to rearrange the list. For example, you can cut and paste a directive to another place in the list.


To Select the Directives From Templates

1. Click the browse button to the right of the Directive set box.

A list of one or more templates appears. For example, there can be basic and advanced versions.



2. Select a template and click OK.
3. To Add selected directives to the file list:
 - a. Click the browse button (...) to the right of the Pathname or directive box.
The list of directives appears.



- b. Select the directive from the list and click OK.
4. To Add the entire list of directives to the file list:
 - a. Click the Add to list button to the right of the Directive set box.
All the directives appear in the file list.
 - b. Use the keyboard to insert, cut, and paste as necessary to arrange the file list.



Verifying the File List

After creating or modifying a file 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 `/usr/opensv/netbackup/bin/goodies`).

This step can reveal mistakes that result in not backing up files because the files are not found. The status code for a backup does not always indicate an error of this nature because NetBackup does not require all paths in the file 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 classes, 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 file 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 file 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 Backup Problems or All Log Entries report shows a message similar to the following:

```
02/02/99 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 file 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 Backup Problems or All Log Entries report shows a message similar to the following:

```
02/02/99 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 file list names a symbolic link. NetBackup does not follow symbolic links and provides a message such as the following in the Backup Problems or All Log Entries report:

```
02/02/99 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 Standard UNIX and Apollo wbak Clients
- ◆ File-Path Rules for Auspex-FastBackup UNIX Clients
- ◆ File-Path Rules for Extension Clients
- ◆ File-Path Rules for Microsoft Windows Clients
- ◆ File-Path Rules for OS/2 Clients
- ◆ File-Path Rules for NetWare Target Clients
- ◆ File-Path Rules for NetWare NonTarget Clients
- ◆ File-Path Rules for Macintosh Clients
- ◆ File-Path Rules for Extension Clients

File-Path Rules for Standard UNIX and Apollo wbak Clients

The general requirements for pathnames on UNIX Standard and Apollo wbak 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 class file 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 file 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.

Notes on File Lists for Standard UNIX and Apollo wbak 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 file 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).

- ◆ 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/opensv/netbackup/bp.conf` file on the client offer alternatives for handling busy and locked files. See “NetBackup Configuration Options” on page 325.
- ◆ 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.
- ◆ 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 file 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 permission 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 file 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 `bpend_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.

- ◆ Apollo clients cannot back up raw partitions. Specifying a raw-partition device file in the file list of an Apollo client results in backing up the `/dev` device file.
- ◆ For disks managed by disk volume managers such as VERITAS VxVm, specify the logical partition names.
- ◆ For clients in a FlashBackup class, refer to the *NetBackup FlashBackup System Administrator's Guide* (file list and cache section) for the differences between Standard and FlashBackup classes.



Specifying UNIX Raw Partitions In the File List

To specify a UNIX raw partition in the class file 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 file 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 whole partition is backed up in both cases.

Raw-partition backups occur only if the absolute file path in the file 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 they avoid going through the file system to access 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.

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 `bpemd_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 alternate-path 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.

File-Path Rules for Auspex-FastBackup UNIX Clients

The entries in the file list for an Auspex FastBackup class are the same as for a standard UNIX raw-partition backup, except that you must add a special `CACHE` entry to the top of the list (see “File List and Cache” on page 617).

File-Path Rules for Microsoft Windows Clients

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

File Backups - 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 actual 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 NT/2000 pathnames:

```
*
```



?

The following backs up all files ending with .doc

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

The following backs up all files named log01_97, log02_97, and so on.

```
c:\system\log??_97
```

- ◆ To back up all local drives except for those that use removable media, specify:

```
:\ or *:\
```

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 file list that uses the Microsoft Windows conventions:

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

File Backups - 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 NT/2000 Disk-Image (raw) Backup

On Windows NT/2000 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.



To specify a disk-image backup, add the logical name for the drive to the class file list as shown in the following example:

```
\\.\c:
```

This example backs up drive C (use the exact form shown above).

You can include a disk-image backup in the same file list as other backups. For example:

```
\\.\c:  
d:\workfiles\  
e:\Special\status  
HKEY_LOCAL_MACHINE:\
```

When a user subsequently lists the backups for the client, the disk image appears as a file with the same name that was specified in the file list. In this example:

```
\\.\c:
```

To restore the backup, the user selects `\\.\c:` and starts the restore in the same way as for files.

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.

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 Windows NT/2000, this directory is

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

- ◆ On Windows 98 or 95, this directory is

```
%systemroot%
```

Where `%systemroot%` is the directory where Windows 2000, NT, 98, or 95 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 file 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.

Back Up Individual HKEYs (do not use for disaster recovery)

As mentioned above, do not include HKEY entries in the same class file list used to back up the entire registry. However, if you want the ability to restore individual keys within the registry, create a separate class and then specify the desired HKEYs in the file list for that class. The following is an example HKEY entry for a class file list:

```
HKEY_LOCAL_MACHINE: \
```

Remember, you cannot perform a disaster recovery by restoring HKEYs. In addition, backups and restores will also 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 file 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. 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 ID and are assigned a new file ID.



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 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 OS/2 Clients

The requirements for OS/2 clients are the same as for Microsoft Windows clients.

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 back up 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 User's Guide NonTarget Version - Novell NetWare*.
- ◆ 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 path name 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 NT clients.
- ◆ To back up all NetBackup for NetWare clients that are in this class, 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 User's Guide NonTarget Version - Novell NetWare*.

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 User's Guide Target Version - Novell NetWare*).

- ◆ 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 NT clients.

The following example backs up the targets: NETWARE, SYSTEM, and BINDERY:

```
/NETWARE/  
/SYSTEM/  
/BINDERY/
```

File-Path Rules for Macintosh Clients

The following explains how to map your Macintosh file and folder names to the names required in the class file list. Other syntax rules are also explained.

- ◆ Enter one pathname per line.
- ◆ Specify the Macintosh paths as follows:

```
/ volume/ folders ... /file
```

Where:

- ◆ *volume* is the name of the Macintosh volume (disk) where the data is located.
The first directory in the pathname is always the volume. It is a good idea to use a wildcard (*) in place of the volume because the Macintosh hard disk name is not necessarily the same on each client.
- ◆ *folders* are the names of the Macintosh folders that contain the files to be backed up. If there are multiple folders in the path, separate the names with slashes. The Macintosh folder names map to directories in a UNIX path.
- ◆ *file* is the Macintosh file name. Macintosh file names map to the file name in a UNIX path.
- ◆ 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 serves as a reminder that the file path is a directory instead of a file.

```
/Macintosh HD/Graphics/
```

If the last component is a file, include the file extension and omit the slash from the end of the name.

```
/Macintosh HD/Graphics/Checklist
```

- ◆ Upper and lower case letters in the pathname must match those in the actual pathname on the client.



```
/Macintosh HD/Graphics/
```

- ◆ Macintosh file and folder names can contain the space character.
- ◆ Wildcard usage is the same as for UNIX clients. For example:

```
/* /BackMeUp
```

This example refers to the folders or files named `BackMeUp` at the top level of all the mounted volumes on the Macintosh. To specify all mounted volumes on the Macintosh, use:

```
/*
```

- ◆ Any slash (/) character in a Macintosh file or folder name, maps to a colon (:) character in the pathname on the master server.

For example, assume there is a file named `Notes 95/09/30`, in the `My Stuff` folder, on the hard disk named `Macintosh HD`. To include this file in automatic backups, specify the following in the file list for the class:

```
/Macintosh HD/My Stuff/Notes 95:09:30
```

Example 1

The following backs up a Macintosh file named `Bldg Layout`, that is in the folder named `New Stuff`, which is in the folder named `Graphics` on the hard disk named `Macintosh HD`.

The following entry adds the `Bldg Layout` file to your backups:

```
/Macintosh HD/Graphics/New Stuff/Bldg Layout
```

`Macintosh HD` is the volume, `Graphics` and `New Stuff` are folders, and `Bldg Layout` is a file.

Example 2

The following backs up a folder named `My Stuff` on the hard disk named `Macintosh HD`.

```
/Macintosh HD/My Stuff/
```

File-Path Rules for Extension Clients

File-path rules for NetBackup extension clients, such as NetBackup for MS-Exchange, are covered in the NetBackup guide for the separately-priced extension product.



File-List Directives - General Discussion

The file list for a class 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 class type and whether the Allow multiple data streams attribute is enabled for the class. The following is an example of a file list that contains the `NEW_STREAM` directive and is for an MS-Windows-NT class that has Allow multiple data streams enabled:

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

The purpose of the above example is to show how directives look in a file list. The actions that the `NEW_STREAM` directive causes is explained later, under “File-List Directives for Multiple Data Streams” on page 112.

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

There are three categories of directives as follows:

- ◆ `ALL_LOCAL_DRIVES` Directive
- ◆ Directives for Multiple Data Streams
- ◆ Directives for Specific Class Types

ALL_LOCAL_DRIVES Directive

The `ALL_LOCAL_DRIVES` directive applies to the following class types:

- ◆ Standard (except for NetWare target clients)
- ◆ MS-Windows-NT
- ◆ NetWare (nontarget clients only)
- ◆ OS/2

An exception to the above is that you cannot use `ALL_LOCAL_DRIVES` for NetWare and OS/2 class types if you are also using Allow multiple data streams. Also see “`ALL_LOCAL_DRIVES`” on page 116.

Directives for Multiple Data Streams

If the Allow multiple data streams general attribute is set for a class, you can use the following directives in the file list:



- ◆ NEW_STREAM
- ◆ ALL_LOCAL_DRIVES
- ◆ UNSET
- ◆ UNSET_ALL

The rules for using these directives are explained in “File-List Directives for Multiple Data Streams” on page 112.

Directives for Specific Class Types

Some directives apply only to specific class types and can appear only in file lists for those classes. NetBackup passes class-specific directives to the clients along with the file list. The clients then perform the appropriate action according to the directive. The class types that currently have their own file-list directives are:

- ◆ AFS
- ◆ Auspex
- ◆ FlashBackup
- ◆ NDMP
- ◆ Split-Mirror
- ◆ Lotus-Notes
- ◆ MS-Exchange-Server

For example, the following directives can appear only in the file list for an AFS class:

```
CREATE_BACKUP_VOLUMES
```

```
SKIP_SMALL_VOLUMES
```

Except for AFS, the above class types can be used when their associated separately-priced option is installed. For information about AFS directives, see . For information on the other classes and their file-list directives, see the NetBackup guide for the option.

Caution Include class-specific directives only in file lists for the classes that support them or errors can occur.

File-List Directives for Multiple Data Streams

If the Allow multiple data streams general attribute is set for the class, the following directives can be used in the file list to control the way that NetBackup creates backup streams:

- ◆ NEW_STREAM

- ◆ ALL_LOCAL_DRIVES
- ◆ UNSET and UNSET_ALL

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

NEW_STREAM

The `NEW_STREAM` directive is recognized only if Allow multiple data streams is set for the class. If this directive is used in a file 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 file 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 file 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 file 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 file list is terminated by the end of the file 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 file list below:

```
NEW_STREAM
/usr
/lib
NEW_STREAM
/home
```



/bin

This file 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 file list entry as part of an existing stream, its first backup is according to the next schedule that is due for the class. 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 file 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 file list *and* the list contains either the `ALL_LOCAL_DRIVES` directive or wild cards. In this mode, the file list is sent to the client, which preprocesses the list and splits the backup into streams as follows:

- ◆ If the file list contains the `ALL_LOCAL_DRIVES` directive, NetBackup backs up the entire client but splits each drive volume (Windows NT) or file system (UNIX) into its own backup stream (see “`ALL_LOCAL_DRIVES`” on page 116).
- ◆ If wild cards are used, the expansion of the wild cards results in one stream per wild card expansion.

If the file 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 file list becomes a separate stream.

Auto-discover streaming mode applies to:

- ◆ Standard and MS-Windows-NT class types, except for Macintosh and 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 file list before the backup begins. The first backup of the class 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 (on UNIX, you can also use the global attributes setting in `bpadm`). 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 backup 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 (preprocessing occurs prior to every backup) by specifying `-prep 0`.

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
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 file list has the following entries:

```
/usr
/lib
/home/*
```

For this file 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 not files then NetBackup produces a separate stream for each of them: `/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

The `ALL_LOCAL_DRIVES` directive applies only to Standard (except for NetWare target clients), MS-Windows-NT, NetWare, and OS/2 classes where the clients are running NetBackup 3.2 or later software. If used, this directive must be the only entry in the file list for the class; 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 class.

- ◆ If Allow multiple data streams is enabled, the `ALL_LOCAL_DRIVES` directive is valid only if the class type is Standard (except for Macintosh and NetWare clients) or MS-Windows-NT. In this instance, NetBackup backs up the entire client and splits the data from each drive (Windows NT) or file system (UNIX) into its own backup stream. NetBackup periodically runs preprocessing on the client to make necessary changes to the streams.
- ◆ 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 classes 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 NT system with two drive volumes, `C:\` and `D:\`. The file list contains:

```
ALL_LOCAL_DRIVES
```

For this file 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 NT system with two drive volumes, `C:\` and `D:\`. The file 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

All class-specific directives (See “Directives for Specific Class Types” on page 112) that are passed to a client in a stream are passed in all subsequent streams. The `UNSET` and `UNSET_ALL` directives change this behavior. These directives are recognized only if Allow multiple data streams is set for the class.



UNSET

Unsets a class-specific directive so it is not passed with any additional streams. The directive that was unset can be defined again later in the file list and included in the current and later streams.

UNSET_ALL

UNSET_ALL has the same effect as UNSET but unsets all class-specific directives that have been defined up to this point in the file list.

Example

Assume you have a file list as shown below. In this file 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
/var
```



Excluding Files From Automatic Backups

On most NetBackup clients, you can exclude specific files from automatic backups by specifying them in an exclude list on the client. You can also create an include list to add back in some of the files by using an include list. The include list is useful, for example, if you want to exclude an entire directory except for one file.

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

The method for specifying files in the exclude and include lists depends on the type of client that you are configuring.

- ◆ On Microsoft Windows clients, specify exclude and include lists in the NetBackup Configuration dialog box in the user interface on the client. To open this dialog box, start the user interface on the client and click Configure on the Actions menu. Then, go to the Exclude List or Include list tab. For further instructions, see the NetBackup user's guide for the client.
- ◆ On NetWare target clients, you specify the exclude and include lists when adding the targets. See the NetBackup user's guide for the client.
- ◆ Macintosh, and OS/2 clients do not support exclude and include lists.
- ◆ On UNIX clients, you create the exclude and include lists in the following files on the client:

```
/usr/opensv/netbackup/exclude_list
```

```
/usr/opensv/netbackup/include_list
```

The following topics explain the rules for creating these lists on UNIX clients.

Creating an Exclude List on a UNIX Client

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

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

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 nonwildcard 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 file list by using / or *, 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

If an exclude list has the following entries:

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

the following are excluded from automatic backups:

- ◆ The file or directory named `/home/does/john`.
- ◆ The directory `/home/does/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 Classes or Schedules

NetBackup allows you to create an exclude list for a specific class or a class and schedule combination. To do this, create an `exclude_list` file with a `.classname` or `.classname.schedulename` suffix. The following are two examples for a class named `wkstations` that has 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 class named `wkstations`. The second file affects backups only when the schedule is named `fulls`.

For a given backup, NetBackup uses only one exclude list and that is the one with the most specific name. For example, if there are files named

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

NetBackup uses only

```
exclude_list.wkstations.fulls
```

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.

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 class or class and schedule combination, use a `.classname` or `.classname.schedulename` suffix. The following are two examples of include list names for a class named `wkstations` that has 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 class 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, if there are files named

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

NetBackup uses only

```
include_list.workstations.fulls
```

To Add New Schedules

Note See “To Change a Class” on page 61 for instructions on inserting, changing, or deleting schedules in an existing class.

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policy Management (Classes) window appears.
2. Choose the master server where you want to add the schedule (see “Choosing the Master Server for Backup Policies” on page 57).
3. Find the class name in the left pane and click the + to the left of that name.
This expands the tree to show Attributes, Schedules, Files, and Clients below the selected class name.
4. Double-click Schedules under the class name.



A dialog box appears. The title bar shows the name of the class to which you are adding the schedules.

The screenshot shows a dialog box titled "Add Schedule - Class workstations". The "Server:" field is set to "bison". The "Name:" field contains "weekly_full". The "Frequency:" is set to "1" and "weeks". The "Media multiplexing:" is set to "1". The "Type of backup:" is set to "Full Backup". The "Retention:" is set to "1 month". There are checkboxes for "Override class storage unit" and "Override class volume pool", both of which are unchecked. Below these fields is a "Schedule" section with a grid for days of the week (Sun-Sat) and time slots for "Start time:", "Duration:", and "Ends:". The "Ends:" row is highlighted in red. At the bottom are buttons for "Add", "OK", "Close", and "Help".

5. Specify the properties for the schedule as explained in “Schedule Properties” on page 124.
6. If this is the last schedule, click OK. To add more schedules, click Add and repeat step 5. Click Close to cancel changes that you have not yet added and close the dialog box.

Schedule Properties

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 as the first character or leave any spaces between characters.

Type of backup

Specifies the type of backup that this schedule will control. Click the button at the right of the box and select from the list. The selection list shows only the backup types that apply to the class you are configuring. A complete list of the selections that can appear is as follows:

- ◆ Full Backup

Backs up all the files that are specified in the file list for the class, 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.

◆ Cumulative-Incremental Backup

Backs up all files that are specified in the file list and 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.

◆ Differential-Incremental Backup

Backs up all files that are specified in the file list for the class and 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 and all the differential incrementals that have occurred since the last full backup.

◆ User Backup

Initiated by the user through the interface on the client and backs up all files that the user specifies. Users can start backups only during the times that you specify in the schedule.

◆ 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.

◆ Backup Policy

A backup type that applies to all database extension 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-extension clients, except NetBackup for Informix and Oracle. For more information on configuring schedules for this type of backup, see the NetBackup guide that came with the 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*.

Retention

Specifies how long NetBackup retains the backups that it creates according to this schedule. To set the retention period, click the Retention button at the right of the box and select from the list that is displayed. When the retention period expires, NetBackup deletes information about them, thus 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 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

The default choices are shown below. To change the defaults, choose NetBackup System Configuration from the Configure menu in the NetBackup administration window and then select the Retention Periods tab.

- 1 week (0)
- 2 weeks (1)
- 3 weeks (2)
- 1 month (3)
- 2 months (4)
- 3 months (5)
- 6 months (6)
- 9 months (7)
- 1 year (8)
- infinite (9)

Note The numbers in parentheses are the retention levels, which are index numbers corresponding to the retention period (for example, the default retention period for level 0 is one week). The retention levels are shown here for reference because 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 128).

Precautions For Assigning Retention Periods

- ◆ Be certain to assign a long enough retention period, because NetBackup stops tracking backups when the retention period expires, thus making them difficult or impossible to recover.
- ◆ Within a class, 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, set `ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA`. See “NetBackup Configuration Options” on page 325 for more information.

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.

Frequency

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

Specifies how much time can elapse between successful automatic backups for clients on this schedule. For example, assume that you set up a schedule for a full backup with a frequency of one week. If NetBackup successfully completes a full backup for all clients on Monday, it does not attempt another backup for this schedule until the following Monday.

To set the frequency, click in the Frequency box and either type a number or use the arrows to select a value. Click the box that is on the right under Frequency to select units of hours, days, or weeks.

Backup Frequency Determines Schedule Priority

If more than one automatic schedule is due for a client within a class, 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 full and incremental schedules have the same backup frequency and are both due for the same client, jobs from the full get precedence.

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 618.

Override class storage unit

Specifies whether to use the class storage unit or another one for this schedule.

Note There are special requirements for Auspex-FastBackup storage units (see “Auspex Client in an Auspex FastBackup Class” on page 615).

- ◆ To override the class storage unit, select the check box. Choose the storage unit by clicking the box below the check box and then selecting from the list of previously configured storage units. If the list is empty, it means that you have not yet configured any storage units.
- ◆ To use the class storage unit, do not select the check box. NetBackup will use the class storage unit that you specified with the Class storage unit general attribute. If you did not specify a class storage unit, NetBackup uses any available storage unit.

Override class volume pool

Specifies whether to use the class volume pool or another one for this schedule.

- ◆ To override volume pool specified by the Class volume pool general attribute, select the box. Choose the volume pool by clicking the box below the check box and selecting from the list of previously configured volume pools. You can also type a pool name in the box and then configure it in Media Manager.
- ◆ To use the class volume pool, do not select the box. NetBackup will use the volume pool that you specified with the Class volume pool general attribute. If you did not specify a class volume pool, NetBackup uses “NetBackup” as the default.

Media multiplexing

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



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 (also see “Multiplexing” on page 275).

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

Schedule

Provides controls for the setting time periods during which NetBackup can start backups or archives according to this schedule. These time periods are referred to as backup windows. You can create backup windows as necessary to satisfy your backup requirements. For example, you can have a different window open each day or keep the backup window open all week.

Specify the start times and durations for the backup windows in the Start and Duration boxes. The end day and time is automatically calculated and appears below the Start and Duration boxes. Table 3 describes the schedule settings and controls.

Table 3. Schedule Settings and Controls

Start	<p>Time when the backup window opens and NetBackup can start performing backups according to this schedule. Specify the time in hours, minutes, and seconds using the 24-hour clock. For example, 1 am is 01:00:00 and 11 pm is 23:00:00. 12 pm is 00:00:00.</p> <p>You can specify Start Time in any of the following forms:</p> <p><i>hours</i></p> <p><i>hours:minutes</i></p> <p><i>hours:minutes:seconds</i></p> <p>If you enter just the hours or hours and minutes, NetBackup completes the entry. For example, entering 22 for Start sets the start time to 22:00:00.</p>
-------	---

Table 3. Schedule Settings and Controls

Duration	<p>Time period (in hours) during which the backup window is open and NetBackup can start backups according to this schedule. Assume, for example, that the Start Time is 1 am and the Duration is five hours. Here, NetBackup can start client backups for this schedule only between 1 am and 6 am.</p> <p>If a backup on another schedule is in process or a required device is not available until after 6 am, the backups on this schedule do not start until the next time this schedule is run.</p> <p>However, once started, a backup runs to completion, even if the duration expires while the backup is running. See “Example of Duration” on page 132.</p> <p>Specify a large enough value to allow all clients in the class to complete their backups. Also, allow extra time in case the schedule starts late due to factors such as devices not being available.</p>
Ends	<p>Time when the duration time period has elapsed and the backup window closes. NetBackup automatically completes the Ends value after you have set the Start and Duration and chosen Duplicate, Add, or OK.</p> <p>If you create an overlapping window, NetBackup does not accept the schedule.</p>
Duplicate	<p>Click Duplicate to fill blank fields with the value that is to the left of blank fields. For example, assume that all fields are blank and you type 8 in the leftmost Duration field (Sunday). In this case, clicking Duplicate copies 8 into all the blank fields.</p>
Clear	<p>Clear all start times and durations.</p>

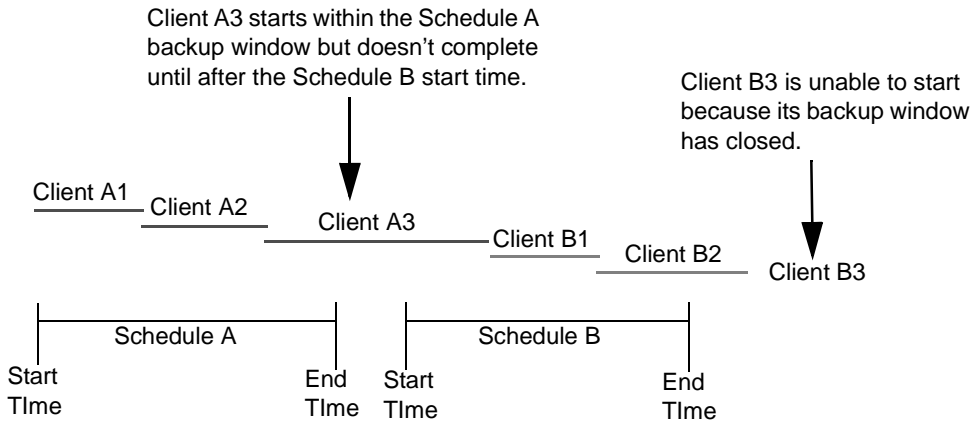
Notes on Start Time and Duration

- ◆ NetBackup does not allow overlaps in start times and duration within a single schedule definition. For example, if you set one start time for 23:00:00 on Saturday with a duration of four hours and then attempt to set another start time for 01:00:00 on Sunday, you get a message indicating that the times overlap. NetBackup does not accept the schedule until you correct the problem.
- ◆ If you set Duration to zero or leave it blank, the backups on that schedule do not execute automatically. You can, however, start them with a manual backup. This allows you to create schedules to use exclusively for manual backups.



Example of Duration

The figure below shows the effect of schedule duration on two full-backup schedules, where the start time for the second schedule is shortly after the end time for the previous schedule. 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 duration period has started and does not leave enough time for all the Schedule B backups. Client B3 must wait until the next time that NetBackup runs Schedule B.

Examples of Automatic-Backup Schedules

You can have backups occur automatically on every day of the week or only on certain days. You can also specify a different start time and duration 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.

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:

Type of Backup	Full
Frequency	7 days
Start Time	22:00
Duration	8 hours
Days of Week	All

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 lets backups 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 schedule that lets backups occur only on specific days. Full backups occur only on Tuesdays and every seven days.

Type of Backup	Full
Frequency	7 days
Start Time	22:00
Duration	8 hours
Days of Week	Tuesday

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 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:

Type of Backup	Differential Incremental
Frequency	1 day
Start Time	22:00
Duration	8 hours
Days of Week	All

Weekly Full Backups:

Type of Backup	Full
Frequency	7 days
Start Time	22:00
Duration	8 hours
Days of Week	All

Monthly Full Backups:

Type of Backup	Full
Frequency	4 weeks
Start Time	22:00
Duration	8 hours
Days of Week	All



Quarterly Backups:

Type of Backup	Full
Frequency	12 weeks
Start Time	22:00
Duration	8 hours
Days of Week	All

Example 3

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

Type of Backup	Differential Incremental
Frequency	1 day

Day	Start Time	Duration
Sunday	18:00	12
Monday	22:00	8
Tuesday	22:00	8
Wednesday	22:00	8
Thursday	22:00	8
Friday	22:00	8
Saturday	22:00	12

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.

Type of Backup	Differential Incremental
Frequency	12 hours



Day	Start Time	Duration
Sunday	00:00	168
Monday		
Tuesday		
Wednesday		
Thursday		
Friday		
Saturday		

Example 5

The following example lets full backups 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.

Type of Backup	Full
Frequency	3 days

Day	Start Time	Duration
Sunday		
Monday		
Tuesday		
Wednesday		
Thursday		
Friday	20:00	60
Saturday		

Example 6

The following is an example where a full backup runs every Sunday and cumulative incrementals run on all other days. Each of the cumulative incremental backups contain all files that have changed since the last full backup. This puts more files in each



incremental than are present for a differential but it makes restores easier. If a restore is required on Saturday, the Sunday tape and the Saturday tape are needed to do the restore. If this were a differential incremental, then all tapes Sunday through Saturday would be needed.

Type of Backup	Full
Frequency	7 days

Day	Start Time	Duration
Sunday	22:00	8
Monday		
Tuesday		
Wednesday		
Thursday		
Friday		
Saturday		

Type of Backup	Cumulative Incremental
Frequency	1 day

Day	Start Time	Duration
Sunday	22:00	8
Monday	22:00	8
Tuesday	22:00	8
Wednesday	22:00	8
Thursday	22:00	8
Friday	22:00	8
Saturday	18:00	12

Considerations for User Schedules

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



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

Note An archive is a special case of a backup. NetBackup first backs up the selected files and 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. With user backups, the best time slot is during hours that are convenient to the users.

If possible, do not permit user backups and archives when automatic backups are occurring. If an automatic backup is executing when a user submits a backup or archive, NetBackup queues the user job. 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 their 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 you do not configure a long enough retention period 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 Classes for User Schedules

If you decide to create separate classes for user backups or archives, the considerations are similar to those for automatic backups. One difference is that you do not need a file list because users select the files before starting the operation.



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

Class	Client	Desired Storage	Best Backup Time	Retention
User1	mercury mars jupiter neptune	8-mm tape stacker	08:00 to 16:00	Backups - 6 months Archives - Infinite
User2	pluto	8-mm tape stacker	12:00 to 20:00	Backups - 6 months Archives - Infinite

- ◆ All clients in class User1 have common requirements for user backups and archives.
- ◆ The class 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 class 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 class 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 class 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

To Use a Specific Class and User Schedule

To use a specific class and (or) schedule for user backups or archives, perform the following on the client:

- ◆ On Microsoft Windows clients, specify the class and schedule on the Backups tab in the NetBackup Configuration dialog box. To open this dialog box, start the user interface on the client and click Configure on the Actions menu.
- ◆ On NetWare target clients, specify the class and schedule with `backup_class` and `backup_sched` entries in the `bp.ini` file (see the NetBackup user's guide for the client).
- ◆ On UNIX and Macintosh clients, specify the class and schedule with `BPARCHIVE_CLASS`, `BPARCHIVE_SCHED`, `BPBACKUP_CLASS`, or `BPBACKUP_SCHED` options in the `bp.conf` file.

Performing Manual Backups

You can perform immediate manual backups of selected automatic backup schedules and clients within a class. A manual backup is useful for situations such as:

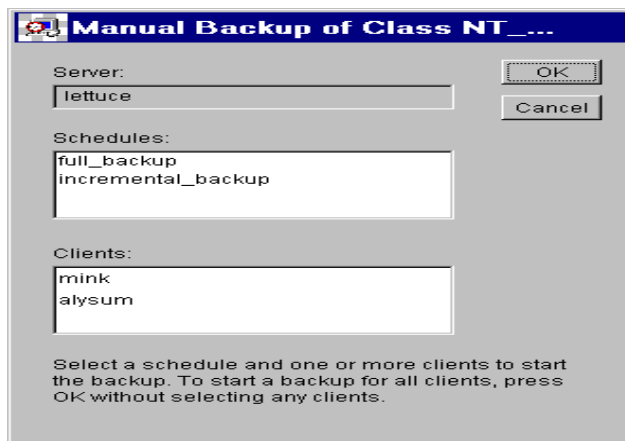
- ◆ Testing a configuration.
- ◆ When workstations miss their regular backups.
- ◆ Before installing new software (to preserve the old configuration).
- ◆ Preserving records before a special event such as when companies split or merge.
- ◆ Quarterly or yearly financial information.
- ◆ In some cases, it may be useful to create a class 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 executes automatically).

To perform a manual backup:

1. In the NetBackup Administration window, click Backup Policy Management.
The Backup Policies (Classes) window appears.
2. Find the class name in the left pane of the Backup Policies (Classes) window and click that name to select it.
3. On the Class menu, click Manual Backup.



The Manual Backup dialog box appears. Notice that the name in the dialog box is that of the class you selected.



4. In the Manual Backup dialog box, select the schedule and the clients that you want to back up.

If you do not select any schedules, NetBackup uses the schedule with the highest retention level. If you do not select any clients, NetBackup backs up all clients.

User schedules do not appear in the schedules list and cannot be manually backed up because they do not have a file list (the user selects the files).

5. Click OK to start the backup.

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

- ◆ Introduction to the NetBackup Catalogs
- ◆ Using the Catalog Backup Wizard
- ◆ Configuring NetBackup Catalog Backups
- ◆ Backing Up the Catalogs Manually
- ◆ Protecting Large NetBackup Catalogs
- ◆ Managing the NetBackup Catalogs

Introduction to the NetBackup Catalogs

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

NetBackup requires the catalog information to recover any backups that have been performed. Therefore, it is extremely important to configure catalog backups before starting to use NetBackup and to schedule them to occur on a regular basis thereafter. Otherwise, you risk losing your regular backups if there is a problem with the disk that contains the catalogs.

Where are the Catalog Files?

The catalogs reside on disk on NetBackup servers. NetBackup chooses default locations for them during installation (see “Files Tab” on page 153). If you change the default location, you must change your catalog backup configuration accordingly.



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 the configuration that you choose in “Configuring NetBackup Catalog Backups” on page 146.
- ◆ Manual backup as explained under “Backing Up the Catalogs Manually” on page 156.

What NetBackup Servers Can I Use?

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. For more information, see “Media Type” on page 148.

How Do I Know If a Catalog Backup Succeeded?

The All Log Entries, Problems, and Media Log Entries 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 global attribute.

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 - UNIX* 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.
- ◆ 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 keep a record on separate storage.
- ◆ 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.

Using the Catalog Backup Wizard

If you are configuring catalog backups for the first time, the easiest way 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 procedure explained in “Configuring NetBackup Catalog Backups” on page 146.

To use the wizard:

1. Click Assistant on the Start menu in the NetBackup Administration window.
The NetBackup Assistant window appears.
2. In the NetBackup Assistant window, click Configure the Catalog Backup.

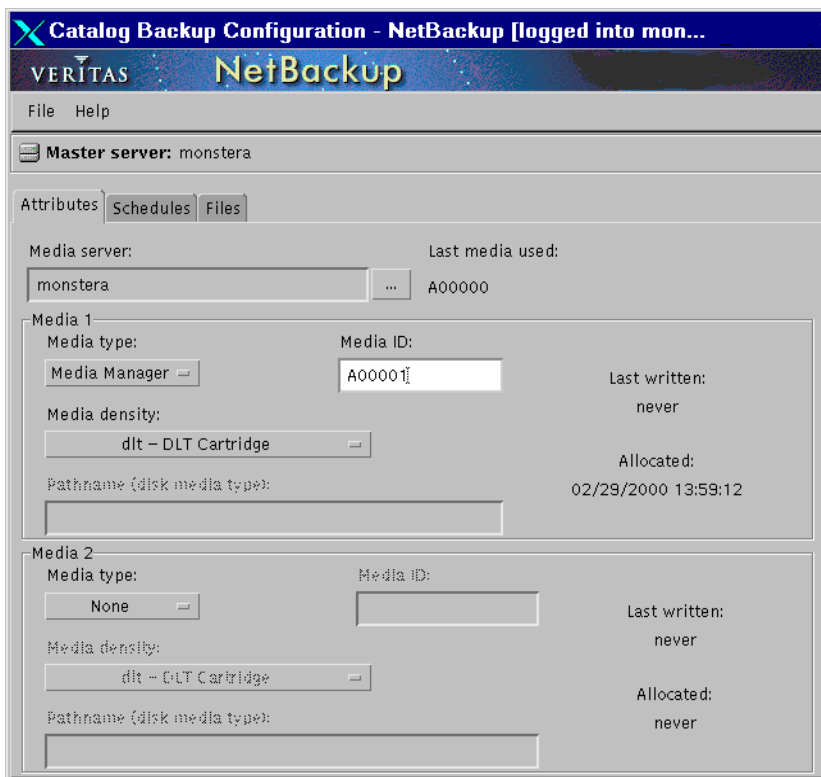


Configuring NetBackup Catalog Backups

Note If you are unfamiliar with NetBackup catalog backups, read “Introduction to the NetBackup Catalogs” on page 143 before proceeding. In particular, read the precautions under “Important Precautions to Observe” on page 144.

1. In the NetBackup Administration window, click NetBackup Catalog Backup on the Configure menu.

The Catalog Backup Configuration window appears.



2. If there is more than one NetBackup master server, check the Master server line below the menu bar and ensure that it names the master server where you want to change the configuration. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.
 - b. In the dialog box, specify the name of the desired server.

Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), click the down arrow on the Select New Master Server button in the dialog box and select from the list. The list shows the master servers that have a `KNOWN_MASTER` entry in the `bp.conf` file on the Master of Masters system.

- c. Click OK.

If you cannot change to the specified server, see “If You Cannot Change to Another Master Server” on page 23 for troubleshooting information.

3. Specify the properties on each tab of the dialog box as explained in the topics below:
 - ◆ “Attributes Tab” on page 147
 - ◆ “Schedule Tab” on page 151
 - ◆ “Files Tab” on page 153
4. On the File menu:
 - a. Click Save to save the changes.
 - b. Click Exit to close the dialog box (Exit does not save the changes).

Attributes Tab

Media Server

Specifies the name of the NetBackup server to which you are sending the catalog backups. This defaults to the master server where you are running the administration interface. To choose a server, click the button to the right of the box and select from the list. The list shows all servers that have a `SERVER` entry in the `bp.conf` file on the master server where you are changing the configuration.

If you are backing up the databases to a media server, modify the NetBackup catalog backup paths on the media server’s master server as explained under “Files Tab” on page 153. Also, ensure that the media server was named in the `bp.conf` file on the master server at the time that you started `bprd` and `bpdbm`.

Last Media Used

Shows the media ID (for Media Manager media) or absolute pathname (for disk) that has 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

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 them.

Media Type

Specifies the media type. Click the box to open the following list of choices:

- ◆ None - No media is assigned
- ◆ Media Manager - A volume that is in a robot or drive under control of Media Manager
- ◆ Disk - A directory on a disk drive

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 Media Manager for the media type and use this automated device to store your 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 Media Manager for the media type and use your 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 your catalog backups to a hard drive (though this is not as safe as method 4 below). The hard drive that you use for your catalog backup must be different than the hard drive where your catalogs reside. By default, the catalogs are stored in the following locations, so if you use this method, the destination of your 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 and then move a full set of that media to offsite storage on a regular basis. A backup written only to disk will share risk with 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 your catalogs and your 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 your 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 Media Manager for the media type. Here, you must back up your 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

Specifies the media ID if you choose Media Manager (removable) type of media. Type the media ID in the box.

The volume you use must be configured in Media Manager in the same manner as other NetBackup volumes. This means the media ID must appear in the Media and Device Management window. The volume must also meet the following requirements.

- ◆ It must be in the NetBackup volume pool. To verify this, look at the Media and Device Management window and ensure that the Pool column for the media ID shows NetBackup.
- ◆ It cannot be currently assigned to NetBackup for backups because NetBackup does not mix catalog backups and regular backups on the same media. To find an available volume, look at the Media and Device Management window and find one where the Time Assigned column is empty and the Status column is 0. After you specify a volume for catalog backups, a time appears in the Time Assigned column and the Status column changes to 1. If one of these columns does not appear, right-click in the pane and choose Show All Columns from the shortcut menu.

The Last Written field in the Media 1 and Media 2 areas shows when the volume specified in the Media ID box 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 is not actually altered until the next time it is used for a backup.

The Allocated field in the Media 1 and Media 2 areas shows 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 value in the Media Type box 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 them in its catalogs as it does with other backup media. If it did, and the disk with the catalogs crashed, the record would be lost anyway along 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 one that has 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.

Media Density

Specifies the media density for Media Manager type media. Click the Density box to list the densities that Media Manager supports and select the desired density.

Caution A NetBackup catalog backup *will not* span a tape volume. All the backup data must fit on one tape. Therefore, it is *extremely* important to choose a media density that can hold all the data to be backed up.

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 box. For example:

`/nb/dbbackup`

The path can be either of the following:

- ◆ 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 one, 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 its 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:

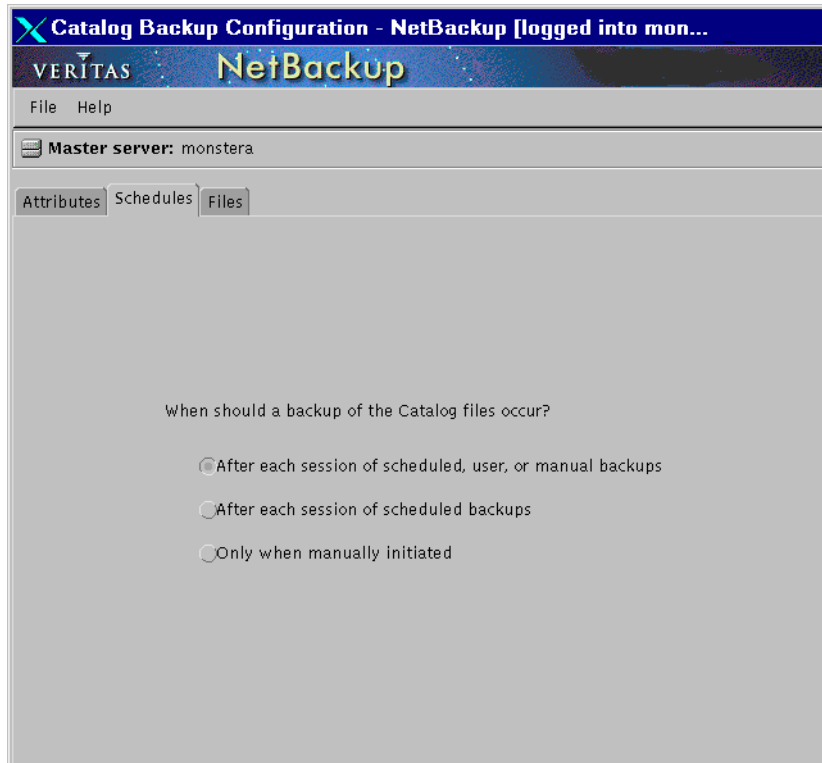
```
/usr/opensv/netbackup/db
```

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

- ◆ 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).

Schedule Tab

On the Schedule tab, select when you want to back up the catalogs.



Caution It is imperative 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.

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:

- ◆ 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.
or
- ◆ 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.

If you must use a single standalone tape drive to back up both your catalog and your business data, choose one of the methods below. Both of these methods require you to swap tapes because NetBackup will not place catalog and regular backups on the same tape.

- ◆ 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
Does not automatically back up the catalogs. To back up the catalogs in this instance, you must use the Back up “NetBackup Catalog” command on the Start menu of the NetBackup Administration window.
If you back up your catalogs manually, be certain to do so once a day or after every series of backups.

The general procedure for catalog backups when you have only one standalone drive is:

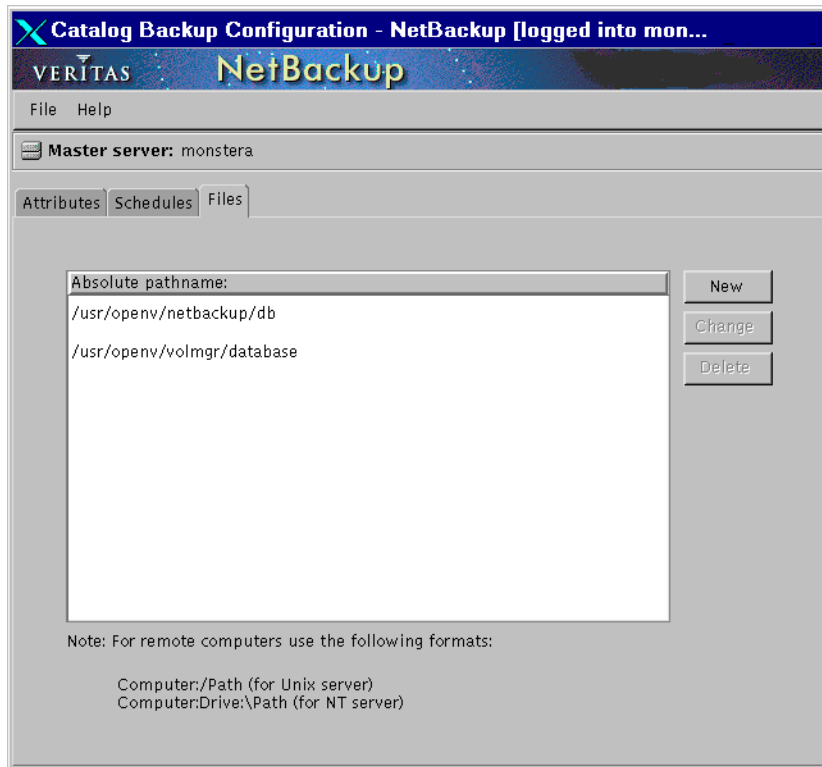
1. Insert the tape that you have configured for catalog backups.
2. Manually start the backup (for instructions, see “Backing Up the Catalogs Manually” on page 156).
3. When the backup is complete, remove the tape and store it in a safe place.

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.



Files Tab

On the Files tab, specify the absolute pathnames for the catalog files you are backing up.



Adding a New Pathname

1. Click New
2. Type the pathname in the Absolute pathname box. For the proper format, see “Catalog Pathnames” on page 154.
3. Press the Enter key to complete the addition.

Caution Ensure that you do not include any invalid paths in your list of catalog files to back up. For example, if you move the catalog files, delete the old paths and add the new paths to the catalog-backup configuration as explained in this chapter. If NetBackup cannot find or follow a path, the entire catalog backup fails.

Do not specify a link as the final component in a UNIX path. If the final



component is a link, it is not followed and the entire catalog backup fails. You can include links at other points in the path. If any other part of a listed path is a symbolic link, NetBackup saves the *actual path* during the backup.

Changing a Path

1. Double-click on the path in the Absolute pathname list.
2. Change the pathname.
3. Press the Enter key to complete the change.

Deleting a Path

1. Click on the path in the Absolute pathname list.
2. Click Delete

Catalog Pathnames

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.

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 bunny, the paths are:

```
bunny: /usr/opensv/netbackup/db
```

```
bunny: /usr/opensv/volmgr/database
```

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.

- ◆ 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 only 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 example, to add the paths for a UNIX NetBackup media server 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
```

Paths For Windows NT/2000 NetBackup Media Servers

If you are backing up catalogs that are on Windows NT/2000 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 NT/2000 NetBackup server 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
```

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 NT/2000 NetBackup server.

Backing Up the Catalogs Manually

A manual backup starts a backup of the catalogs immediately. Some instances where this is useful are:

- ◆ An emergency backup. Examples are 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 instance, 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 a Manual Backup of the Catalogs

1. In the NetBackup Administration window, click Back up “NetBackup Catalog” on the Start menu.

The Backup NetBackup Catalog dialog box appears.

2. If there is more than one NetBackup master server, check the Master server line below the menu bar and ensure that it names the master server where you want to back up the catalogs. To change master servers, see step 2 under “Configuring NetBackup Catalog Backups” on page 146.
3. Click OK to start the backup.

The backup is saved to the least recently used of Media 1 and Media 2.

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

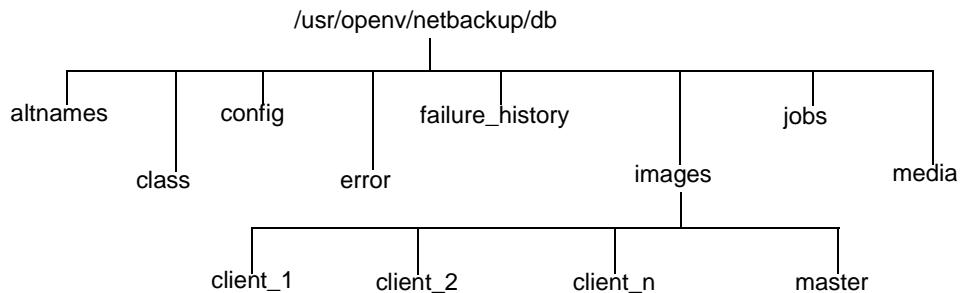
It is very important to ensure that the NetBackup catalogs on the master and media servers are backed up regularly. NetBackup provides a built-in mechanism for achieving this. However, this mechanism imposes a limit on the size of the data that can be backed up; namely, the data must all fit on a single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

This section describes a method that can be used to back up and recover the NetBackup catalog files if they become too large to fit onto a single tape.

The 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 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 sub-directories 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 will be relatively small, the `images` directory can grow to several tens, or even hundreds of gigabytes. The section “Determining Catalog Space Requirements” gives more information on how to estimate 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 here.



Backup and Restore Concept

Multiple-Tape Backups

The basic concept behind the protection of large NetBackup catalogs is to split the catalog-backup process into two steps. The first step backs up the majority of the data from the images subdirectory on the master server, while the second step backs 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 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 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 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 will 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, you must define both a normal NetBackup class and make changes to the NetBackup catalog-backup configuration. In addition, you must create a shell script or executable file that will initiate the multiple-tape catalog backups. These steps are detailed below.



Define a NetBackup Class

1. Use the NetBackup Administration interface to create a new class. Set the class attributes as follows:
 - ◆ Set the Class type to Standard if the master server is a UNIX machine or MS-Windows-NT if the master server is a Windows NT/2000 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 class to 1.
 - ◆ Do not choose Compression.
 - ◆ Set Job priority to 0.
2. Add the master server to the client list
3. Put 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 NT/2000.

4. Set up schedules to meet your requirements. VERITAS recommends that the class contain 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 class is never initiated automatically by the NetBackup scheduler. Instead, you must initiate the backup job manually.
5. Save your changes.

Configure the NetBackup Catalog Backups

Use the NetBackup Administration interface to set up the configuration of the NetBackup catalog backups. Make the following changes:

1. 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 other 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. The following is an example catalog-backup script:

```
#!/bin/sh
#
# catalog_backup script
#
# Performs a two-stage backup of the NetBackup catalogs
#

CLASS=nbu_cat_backup    # Change to the name of the correct class
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 $CLASS -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

Similar 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 execute it manually when the above operations are not occurring.

Managing the NetBackup Catalogs

This section explains the following aspects of managing the NetBackup catalogs:



- ◆ Determining Catalog Space Requirements
- ◆ Compressing the Image Catalog
- ◆ Uncompressing the Image Catalog
- ◆ Moving the NetBackup Image Catalog

Also see “Reduce Restore Times by Indexing the Image Catalog” on page 221.

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.

The following procedure explains the steps for estimating disk usage. An example follows the procedure.

1. Estimate the maximum number of files that each schedule for each class backs up during a single backup of all its clients.

For example, Table 4 shows that a full backup for class S1 includes 64,000 files.

2. Determine the frequency and retention period of your full and incremental backups for each class.
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 that 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. Here, 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, etc.):

Files per Backup x Backups per Retention Period = 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 classes that collect true-image-restore information, an incremental backup collects catalog information on all files (as if it were a full backup). This changes the above calculation for the incremental from $1200 \times 7 = 8400$ to $3000 \times 7 = 21,000$. After adding 12,000 for the fulls, the total for the two schedules is 33,000 rather than 20,400.

4. Obtain the number of bytes by multiplying the number of files by the average length of the file's full pathnames and file information. If you are unsure of the length, use 150 (averages from 100 to 150 are common). Using the results from the examples in step 3, yields:

$$(8400 \times 150) + (12,000 \times 150) = 2988 \text{ kilobytes (1024 bytes in a kilobyte)}$$

Note If you use database indexing, multiply the number in step 4 by 1.5% (see "Reduce Restore Times by Indexing the Image Catalog" on page 221).

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.

EXAMPLE

Table 4 shows backup schedules, retention times, and number of files for a group of example classes. By substituting the information from this table into the formula from step 3 above, we can calculate the maximum number of files for each class. step 1 and step 2 below demonstrate this for class S1:



1. Applying the formula to class S1:

Max Files equals:

(Files per Incremental x Backups per Retention Period)

+

(Files per Monthly Full Backups x Backups per Retention Period)

2. Substituting values from Table 4:

1000 files x 30 + 64,000 files x 12 = 798,000 files

Performing steps 1 and 2 for each class and adding the results together shows that the total files for all classes is:

4,829,600 files

Multiplying the total number of files by the bytes in the average path length and statistics (150 for this example) shows that the total amount of disk space required for file paths is:

690.88 megabytes (1,048,576 bytes in a megabyte)

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

705.88 megabytes

Table 4. Example Reference Table for Catalog Requirements

Class	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



Table 4. Example Reference Table for Catalog Requirements (continued)

Class	Schedule	Backup Type	Retention	Number of Files
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.

You control image-catalog compression with the `Delay` to compress catalog global attribute. 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, `Delay` to compress catalog is set to 0 and image compression is not enabled.

Caution Do not *manually* use the server's `compress` or `uncompress` commands to compress or uncompress the image-catalog files. This practice causes inconsistent image-catalog entries and can produce 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 the execution of 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 percent 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 client's `bp.conf` file.

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). To uncompress the records, perform the following steps as the root user on your server:

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 executing:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```
3. Verify that `bpdbm` is executing by using:

```
/usr/opensv/netbackup/bin/bpps
```
4. Set the Delay to compress database global attribute to 0.
5. Change your working directory to `/usr/opensv/netbackup/bin` and execute the command.

```
admincmd/bpimage -decompress -client name
```

6. Restart the request daemon, `bprd`, by executing:

```
/usr/opensv/netbackup/bin/initbprd
```
7. Perform the file restorations from the client.
8. Set the Delay to compress database global 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 that has more space. The steps are as follows:



1. Check that no backups are in progress by executing the `bpps` command:

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

2. Stop `bprd` by executing:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

3. Stop `bpdbm` by executing.

```
/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 NetBackup Catalog Backups” on page 146).

Caution Be certain to add the actual path for the image catalog not the link name. Otherwise, NetBackup will not back up the new location. In this example, the actual pathname is `/disk3/netbackup/db/images`.





Viewing NetBackup Reports

5

NetBackup provides reports for verifying, managing, and troubleshooting NetBackup operations. These reports show status or problem information for NetBackup servers or clients. The following topics are covered here:

- ◆ To Start the Reports Utility
- ◆ The Reports Window
- ◆ Running a Report
- ◆ Report Settings
- ◆ Report Descriptions



To Start the Reports Utility

1. Start the NetBackup Administration - Java interface program. For instructions, see “NetBackup Administration Interfaces” on page 5.
The NetBackup Administration window appears.
2. In the NetBackup Administration window, click Reports.
The Reports window appears.

The Reports Window

The Reports window has commands for running reports. The following topics provide an overview of this window:

- ◆ Report Contents
- ◆ Menu Bar
- ◆ Toolbar
- ◆ Shortcut Menus (Global Data Manager only)

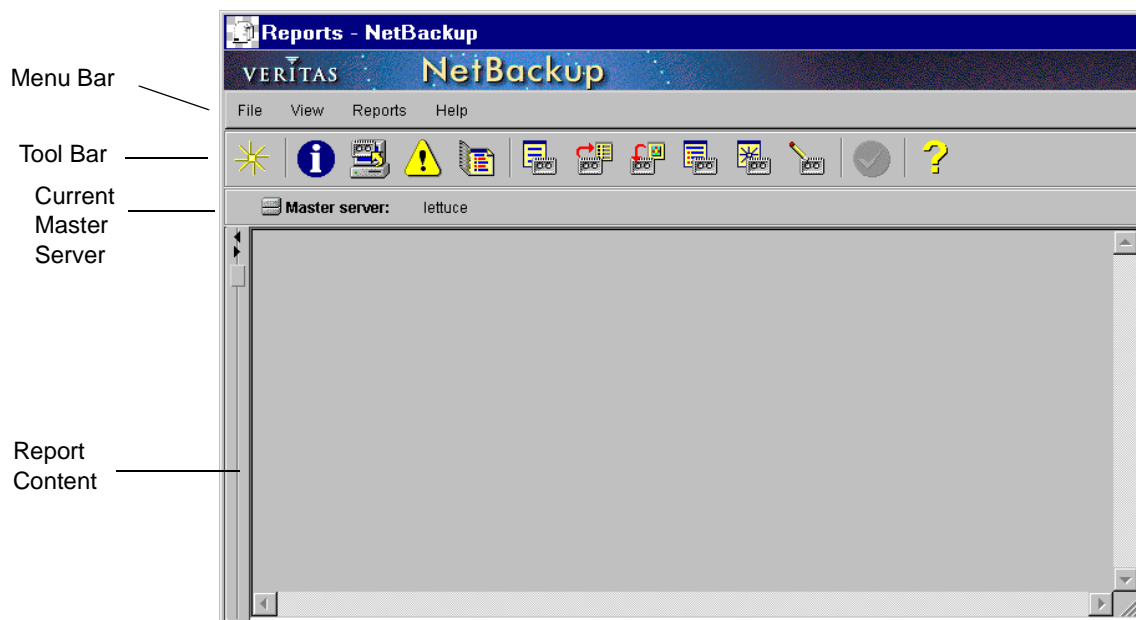
Report Contents

The content of the Reports window is different, depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters). See Appendix E for more information on the Global Data Manager option.

The next two topics explain what appears in each case.

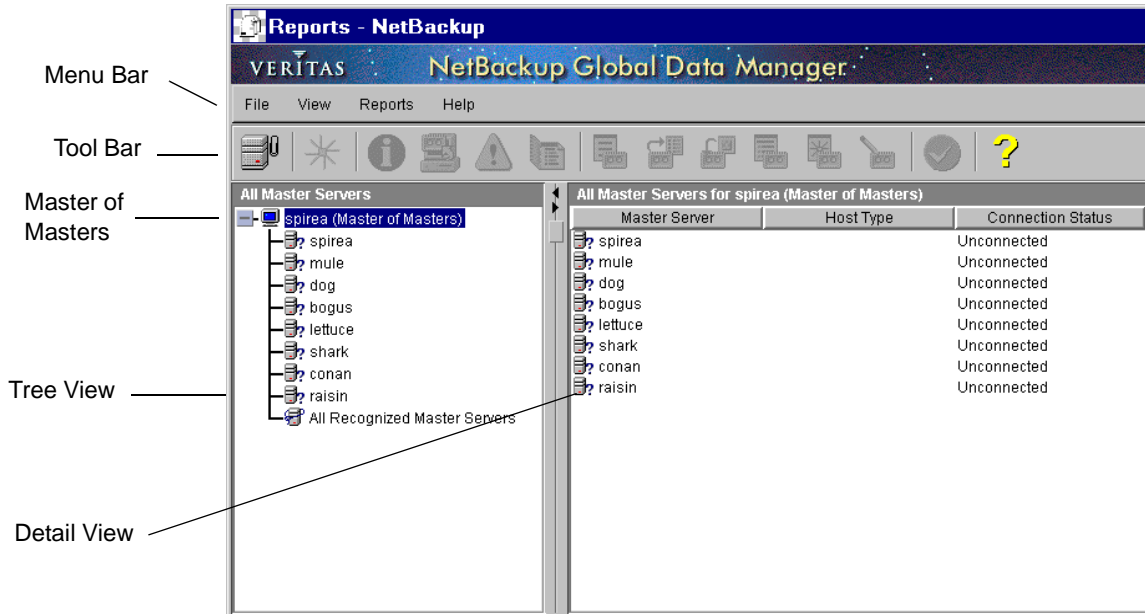
The View if You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the window shows report information for the master server that is currently selected. To select another master server, see “Choosing the Master Server for the Reports.”



The View if you are on a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, the window has two panes. The left pane has a tree view that shows the master servers you can select to run reports. The right pane shows details about whatever is selected in the tree or report information if you run a report.



Selecting the master of masters in the tree displays the list of masters in the right pane. You can then select one or more master servers and perform actions on them. To append other master servers to the tree, use the Append Master Server command on the File menu.

If you select the bottom branch of the tree, All Recognized Master Servers, before running a report, the report contains information from all *recognized* master servers (for more on recognized servers, see “Recognize Master Server(s)” on page 672).



Menu Bar

Table 5 describes the menus and commands available on the menu bar:

Table 5. Report Window Menus

Menu	Commands
File	<p>Append Master Server - (Applies only to Global Data Manager) Displays a dialog box where you can specify a master server to add to the tree. The addition applies only to the current session and the server will not be in the tree the next time you start the administrator interface. Also, see “Append Master Server” on page 670.</p> <p>New Report - Displays a dialog box where you can enter specifications for the report that you want to run.</p> <p>Exit - Closes this window.</p>
View	<p>Preferences - Displays a dialog box where you can specify your display preferences for the toolbar.</p> <p>Ignore Master Server(s) - (Applies only to Global Data Manager) Ignores connections to the master server(s) that you designate so the Master of Masters does not attempt connections to them. The branch of the tree for an ignored server and the associated information in the detail view appears in faded type. You cannot select an ignored server or refresh the display with updated configuration information. The only operation you can perform on an ignored server is to recognize it (see below).</p> <p>Recognize Master Server(s) - (Applies only to Global Data Manager) Recognizes an ignored master server so it can be selected and the display refreshed with updated configuration information. Also see “Recognize Master Server(s)” on page 672.</p> <p>Refresh - Updates the detail pane with new information retrieved from the master server(s).</p>
Reports	<p>Contains commands for generating each type of report. Clicking a command opens a dialog box where you can enter the specifications for the report.</p>
Help	<p>Help Topics - Provides online help information about the Job window.</p> <p>About NetBackup Reports - Displays program information, version number, and copyright.</p>

Toolbar

The toolbar appears at the top of the window and provides shortcut controls for running reports. If the toolbar is not visible, click the Preferences command on the View menu.



Shortcut Menus (Global Data Manager only)

If you are administering NetBackup from a Master of Masters, clicking the right mouse button when the pointer is over either pane, displays a shortcut menu with commands that apply to the current selection. The shortcut commands also appear on the menu bar and are explained in the “Menu Bar.” topic. The exception is Redisplay Data, which redisplayes the information that was last shown for the currently selected server(s).

Choosing the Master Server for the Reports

If there is more than one NetBackup master server, choose the one where you want to run reports.

The procedure is different depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, a Master of Masters).

If You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the name of the server that you are currently monitoring appears above the reports pane. To monitor activity on another NetBackup master server, proceed as follows:

1. On the File menu, click Change Master Server.
2. In the dialog box, specify the name of the NetBackup master server that you want to monitor.
3. Click OK.

The name of the master server that you specified now appears above the reports pane.

If You Are on a Master of Masters (requires Global Data Manager)

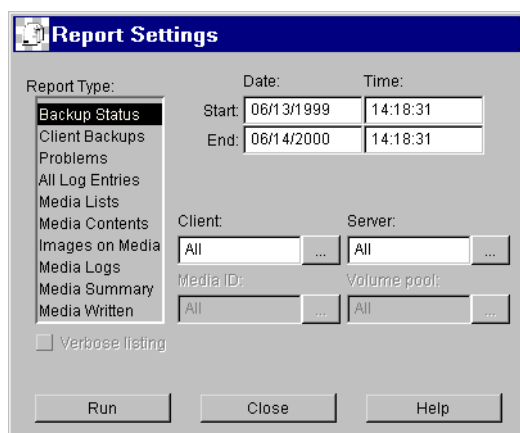
If you are administering NetBackup from a Master of Masters, you can run a report for one master server or a report that combines information from multiple master servers:

- ◆ To run a report for only one master server, select that server in the left pane and run the report. The results appear in the right pane. If the server was previously selected during this session, reselecting it redisplayes the report data.
- ◆ To create a report that includes information from more than one master server, select the Master of Masters in the left pane. Then, select the desired master servers in the right pane and run the report. The combined results appear in the right pane.
- ◆ To create a report that includes information from all recognized master servers, select the bottom branch of the tree, All Recognized Master Servers, and then run the report. The combined results appear in the right pane.

If the server you want is not in the tree, append it by using the Append Master Server command on the File menu. If you encounter problems, see “If You Cannot Change to Another Master Server” on page 23.

Running a Report

1. In the NetBackup Administration window, click Reports.
The Reports window appears.
2. If there is more than one master server, use the Change Master Server command on the File menu to choose the master server (or servers) where you want to run the reports (see “Choosing the Master Server for the Reports”).
3. On the File menu, click New Report.
The Report Settings dialog box appears.



4. In the Report Settings dialog box, specify the criteria for the report as explained in the next topic, “Report Settings.”
5. Click Run.

Refer to “Report Descriptions” on page 177 for an explanation of what appears in each type of report.

Report Settings

You can specify the following criteria for building your report. Not all of these settings are available for every report type. When running a report that includes information for more than one server, the only available options are Verbose Listing, Date, and Time.



Date and Time

Specify the time period that you want the report to cover.

- ◆ Default start time is one day before the report was run.
- ◆ Default end time is the time the report is run.

The two main factors that determine the time period for which information is available are the Duration to retain logs global attribute and the retention period as set on the class schedule that controlled the backups.

Client

Click the Client box and select either All or the individual client to which the report will apply.

Server

Click the Server box and select either All or the name of the individual server to which the report will apply. The master server that is currently selected and its media servers appear in the report.

Media ID

For media types of reports, specify the media ID or All. The Media Contents report requires a specific ID.

Volume Pool

For a media summary report, specify the volume pool name or All.

Verbose Listing

For the media summary report, select Verbose Listing to have NetBackup provide more details.

Report Descriptions

Backup Status Report

The Backup Status 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 report.

Table 6. Backup Status Report

Column	Meaning
Status	Completion status. If the status code is 0, the operation succeeded. If status is not 0, the second line of the report for that client is an error message.
Client	Name of the client for which the backup was performed.
Class	Name of the class that was used to back up the client.
Sched	Name of the schedule that was used to back up the client.
Server	Master server that controlled the backup.
Time Completed	Time that the backup began.

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 report.

Table 7. Client Backups Report

Field	Meaning
Client	Name of the client for which the backup was performed.
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Class	Name of the class that was used to back up the client.
Class Type	Type of class (for example, Standard, MS-Windows-NT, and so on).
Sched Label	Name of the schedule that was used for the backup.
Schedule Type	Type of schedule used for the backup (for example, full or incremental).



Table 7. Client Backups Report (continued)

Field	Meaning
Retention Level	Retention level indicates how long NetBackup will keep a record of this backup (for more information, see “Retention” on page 126).
Backup Time	Date and time that the backup began.
Elapsed Time	How much time the backup required.
Expiration Time	Date and time at which NetBackup will expire its record of this backup.
Compressed	Yes indicates that the backup was compressed.
Encrypted	Y, if the backup is encrypted. Encryption and decryption is possible only with the NetBackup Encryption option.
Kilobytes	Number of kilobytes in the backup.
Number of Files	Number of files in the backup.
Primary Copy	Primary copy shows which copy (1 or 2) NetBackup uses to satisfy restore requests.
Image Type	0 = Regular image (scheduled or user-directed backup) 1 = Pre-imported image (phase 1 completed) 2 = Imported image
Keyword	Keyword that the user associates with this image at the time of the backup.
Ext Security Info	This field is reserved for future use and always shows No.
File Restore Raw	Individual file restore from raw. It is set by the corresponding class attribute if it applies.
File System Only	This field is reserved for future use and always shows No.
Object Descriptor	This field is reserved for future use and is always empty.
Multiplexed	Yes indicates that the backup was multiplexed.
TIR Available	Yes indicates that NetBackup is collecting true image restore information for this class and it is possible to perform a true image restore of the directories in the backup.

Problems Report

The Problems report lists the problems that the server has logged during the specified time period. This information is a subset of the information you get from the All Log Entries report. The following table explains the columns in the report.



Table 8. Problems Report Entries

Column	Meaning
Time	Date when the event occurred.
Server/Client	NetBackup server and client involved in the event. A dash in the client field means that the event did not involve a client.
Text	Error or informational message.

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 Entries 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.

Table 9. All Log Entries Report

Column	Meaning
Time	Date when the event occurred.
Server/Client	NetBackup server and client involved in the event. A dash in the client field means that the event did not involve a client.
Text	Error or informational message.

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, use the methods explained in “Introduction to the NetBackup Catalogs” on page 143.

The following table explains the columns in the report.



Table 10. Media Lists Report

Column	Meaning
server host	NetBackup server where the volumes reside. It is possible to have more than one if the master server has media servers and ALL was selected for the server.
id	Media ID that is assigned when the volume is added to Media Manager.
rl	Retention level for the backups on this volume. An asterisk after the retention level number means that the volume can have multiple retention levels. When there are multiple retention levels, the number shown was the first level assigned.
images	Total number of backups on the volume.
vimages	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.
allocated	Date and time that Media Manager allocated the volume.
expiration	Date and time that the volume expires.
last updated	Last time the volume was used for backups.
last restore	Last time a restore was done from this volume.
den	Density of the device that produced the backups on this volume.
kbytes	Total number of kilobytes on this volume.
restores	Number of times this volume has been used for restores.
Status	The messages that commonly appear here are: SUSPENDED: The volume will not be used for further backups until retention periods for all backups on it have expired. You can still restore from the media. FROZEN: The volume is no longer used for backups. However, the information about it is frozen permanently and you can still restore from the volume (expired backups must be imported first). FULL: The volume is full and no more backups are written to it. EXPIRED: All backups have expired. IMPORTED: The backup was imported.

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. The following table explains the columns in the report.

Note The Media Contents report does not apply to disk type storage units or NetBackup catalog backups. Also, on a Master of Masters, you cannot run this report on multiple master servers.

Table 11. Media Contents Report

Column	Meaning
Media Id	Media ID that is assigned when the volume is added to Media Manager.
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.
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Creation date	Date that NetBackup created the backup.
Expiration date	Date and time that the backup expires.
Retention level	Period of time that NetBackup retains the backup. An asterisk after the retention level number means that the volume can have backups with different retention levels and that the number shown was the first level assigned.
Copy number	Shows the copy number (1 or 2).
Fragment Number	Greater than 1 only if the backup is split across multiple volumes or if the storage unit maximum fragment size is exceeded.
Fragment Size	Shows the number of bytes in the fragment and applies only to optical disk.
Data Start byte offset	Byte address where the fragment starts and applies only to optical disk.
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.
Allocated	Date and time that Media Manager allocated the volume.



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. The following table explains the columns in the report.

Note The Images on Media report does not show information for media used for NetBackup catalog backups.

Table 12. Images on Media Report

Column	Meaning
Backup ID	Identifier that NetBackup assigns when it performs the backup.
Class	NetBackup class for which the backup was created.
Type	Type of backup (full, differential incremental, cumulative incremental, or user-directed).
RL	Retention level for the schedule.
Files	Number of files in the backup.
C	Y if the backup is compressed.
E	Y if the backup is encrypted. Encryption and decryption is possible only with the NetBackup Encryption option.
T	Type of image: R = Regular image (scheduled or user-directed backup) P = Pre-imported image (phase 1 completed) I = Imported image
PC	Primary copy shows which copy (1 or 2) NetBackup uses to satisfy restore requests.
Expires	Date and time of the expiration of the first copy to expire. See description of the Expires field (for fragment 1) described later in this table.
Copy	Greater than 1 only if there are multiple copies.
Frag	Fragment number. IDX (Index file) if the fragment contains true image restore information or is for an individual-file-restore-from-raw backup.
KB	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.
Type	Type of storage and can be removable (Rmed) or disk (Disk).



Table 12. Images on Media Report (continued)

Column	Meaning
Density	Density of the device that produced the backup.
FNum	File number on the media.
Off	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.
Host	Server with the database that has this information.
DWO	Device where the backup was written. This is the drive index configured in Media Manager.
MPX	Y if the copy is multiplexed. Valid only on fragment 1 of the copy.
Expires	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 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. The following table explains the columns in the report.

Table 13. Media Logs Report

Column	Meaning
Time	Date when the event occurred.
Server/Client	NetBackup server and client involved in the event. A dash in the client field means that the event did not involve a client.
Text	Error or informational message.

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 (see “Media Lists Report” on page 179 for more information on these states). Other volumes are considered active.

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 show volumes used for duplication if the original was created prior to the specified time period.

The following table explains the columns in the report.

Table 14. Media Written Report

Column	Meaning
Media ID	Media ID that is assigned when the volume is added to Media Manager.
Last Written	Date when the media was last written.
Server	Server that has the volume database with the records for this volume.

The NetBackup Activity Monitor lets administrators:

- ◆ Monitor the progress of restores, archives, and regular backups. NetBackup catalog backup jobs are not monitored.
- ◆ Delete completed (Done) jobs and kill uncompleted jobs.
- ◆ Save job information.
- ◆ Troubleshoot failed jobs.

The following topics describe the Activity monitor and how to use it:

- ◆ To Start the Activity Monitor
- ◆ The Activity Monitor Window
- ◆ Choosing the Master Server to Monitor
- ◆ Monitoring Jobs
- ◆ Using the Troubleshooting Wizard
- ◆ Managing the Jobs Database



To Start the Activity Monitor

1. Start the NetBackup Administration - Java interface program. For instructions, see “NetBackup Administration Interfaces” on page 5.

The NetBackup Administration window appears.

2. In the NetBackup Administration window, click the Activity Monitor icon.

The Activity Monitor window appears.

The Activity Monitor Window

The following topics provide an overview of the Activity Monitor window:

- ◆ Jobs List
- ◆ Menu Bar
- ◆ Toolbar
- ◆ Status Bar
- ◆ Shortcut Menus (Global Data Manager only)
- ◆ Refreshing the Display

Jobs List

The Activity Monitor window has two panes where you can select items and perform actions on them. The contents of these panes are different, depending on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters).

See Appendix E for more information on the Global Data Manager option.

The View If You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the pane shows the list of jobs for the master server. The list includes all jobs that are in process or have been completed. Catalog backup jobs are not included.

To monitor jobs on another master server, choose that server as explained in “Choosing the Master Server to Monitor” on page 195.

The screenshot shows the NetBackup Activity Monitor window. The title bar reads "Activity Monitor - NetBackup [logged into chives]". The window features a menu bar (File, Edit, View, Help), a tool bar with icons for save, refresh, and help, and a section for the current master server (chives). Below this is a table of backup jobs. The status bar at the bottom shows summary statistics: Queued:0, Requeued:0, Active:0, Done:7, Total:7, and Master Server:chives.

Master	Job Id	Type	State	Status	Class	Schedule	Client	MediaSrv	Start Time
chives	7	Backup	Done	71	nick2	nick_sched	chives	chives	03/01/20..
chives	6	Backup	Done	0	nick	Full	chives	chives	03/01/20..
chives	5	Backup	Done	0	nick	Full	chives	chives	03/01/20..
chives	4	Backup	Done	71	nick2	nick_sched	chives	chives	03/01/20..
chives	3	Backup	Done	219	nick2	nick_sched	chives		03/01/20..
chives	2	Backup	Done	219	nick	Full	chives		03/01/20..
chives	1	Backup	Done	150	nick	Full	chives	chives	03/01/20..

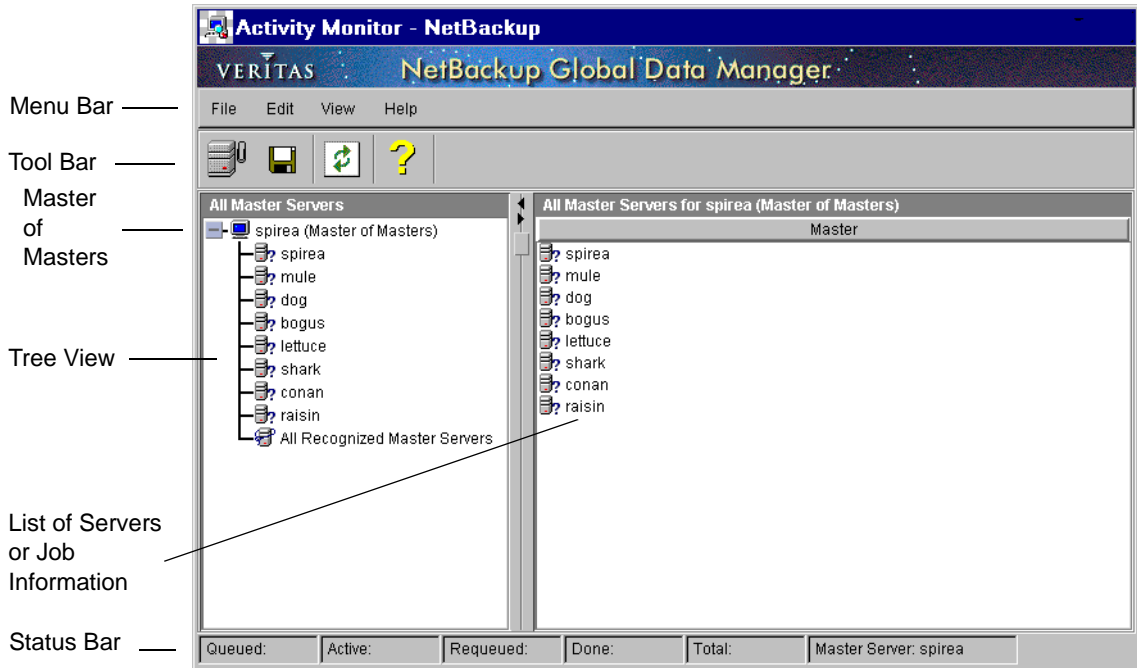
Labels on the left side of the image point to the following components:

- Menu Bar
- Tool Bar
- Current Master Server
- Jobs List
- Status Bar



The View From a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, the left pane is a hierarchical tree view that shows the master servers that you can currently monitor. To append other master servers to the tree, use the Append Master Server command on the File menu. The shows either a list of servers or job information, depending on what you select in the tree.



- ◆ Selecting the master of masters in the tree displays the list of masters in the right pane. You can then select one or more master servers and perform actions on them.
- ◆ Selecting a master server in the tree results in the following:
 - ◆ If the server has not been previously selected, NetBackup retrieves job information from that server and displays it in the right pane. Catalog backup jobs are not included.
 - ◆ If the server has been previously selected, the previously retrieved data reappears in the right pane.
 - ◆ If the server is ignored, information in the right pane is unchanged.
- ◆ Selecting multiple servers in the right pane and running the Activity Monitor creates a combined list of jobs for all the selected servers (see “Monitoring Jobs” for instructions on running the Activity Monitor).

- ◆ Selecting the bottom branch of the tree, All Recognized Master Servers, and running the Activity Monitor, creates a combined listing of all jobs on all *recognized* master servers (for more on recognized servers, see “Recognize Master Server(s)” on page 672). New information is retrieved for each server.

Customizing the Jobs List Output

You can customize the Activity Monitor output to display only the jobs and job data that you want to see. This is accomplished by using the Filters command on the View menu. Many filter combinations are available. For example, you could set the filter so the following appears:

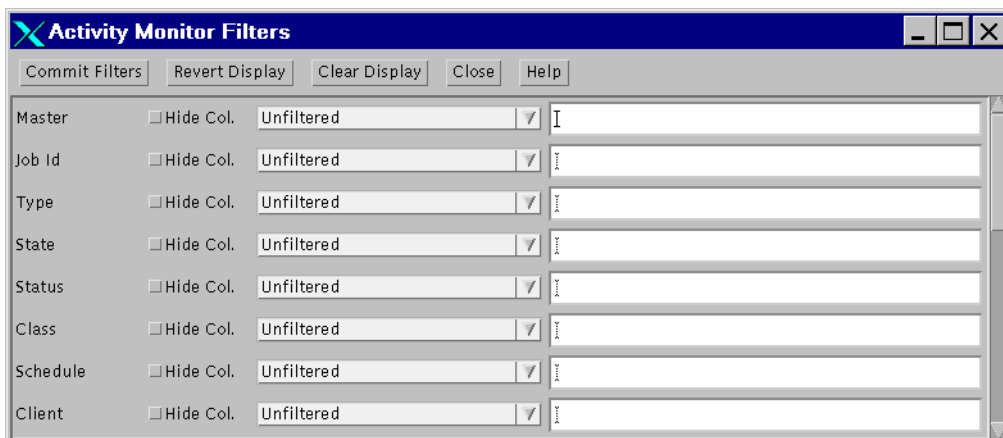
- ◆ Jobs that started before or after a certain date and time
- ◆ Jobs that are in either the active or queued state
- ◆ Jobs that had status (completion) codes within a certain range
- ◆ Specific data for each job, such as: jobid, state, status code, and kilobytes transferred

Filters remain in effect across NetBackup-Java sessions. If you exit NetBackup Java normally with filters set, those filters will be in effect the next time you start a NetBackup Java session on the same host and using the same userid. To disable all filtering, at the start of your next NB Java session, display the Activity Monitor Filters window, then click Clear Display and click Commit Filters before exiting the NetBackup-Java session.

To Specify Job Filter Settings

1. Click Filters on the View menu.

The Activity Monitor Filters window appears. The settings in the window are those currently in effect.



2. Specify the filter settings (see “Filter Settings” later in this topic).



3. Click Commit Filters to apply the changes.
4. Click Close to close the window.

Filter Settings

The five buttons across the top of the Activity Monitor Filters window function as follows:

- ◆ Commit Filters

Applies the current filter settings to the data in the job list and any data subsequently retrieved. You must click Commit Filters to apply changes that you make.

- ◆ Revert Display

Sets the filter settings display to those that are currently in effect (that is, the settings most recently applied with the Commit Filters button). This is useful for discarding changes that you have not applied (by clicking Commit Filters).

- ◆ Clear Display

Disable all filtering so the jobs list displays all data about all jobs. You must click Commit Filters to apply the clear.

- ◆ Close

Closes the Activity Monitor Filters window but does not automatically apply changes. To apply changes, you must click Commit Filters before closing the window.

- ◆ Help

Displays online help for the window.

The rows below the buttons contain filter settings for each type of job data that can appear in the job list.

The leftmost column in each row identifies the type of job data that this row filters and corresponds to one of the column identifiers in the job list. For example, Master, Job Id, and Type).

The other columns in each row contain the filter settings as follows:

- ◆ Hide button

Clicking this button to the on position hides the corresponding column in the job list.

- ◆ Filter type

Specifies the type of filter to use for this job data (see below for details).

- ◆ Filter value

Specifies the value to use for the filter.

To set the filter type, click the button to the right of the field and select from the list. The following filter types are available for all types of job data:

- ◆ Unfiltered
Leaves this type of job data unfiltered.
- ◆ Field exactly matching
The job data *must* exactly match the string specified by the filter value in order for the job to be displayed.
- ◆ Field not exactly matching
The job data *must not* exactly match the string specified by the filter value in order for the job to be displayed.
- ◆ Field contains
The job data *must* contain the string specified by the filter value in order for the job to be displayed.
- ◆ Field does not contain
The job data *must not* contain the string specified by the filter value in order for the job to be displayed.
- ◆ Regular Expression
Specifies that the job data must match the criteria for the Perl-type regular expression that appears in the filter value field. For more information on regular expressions see the following (all on one line):

[www.perl.com/CPAN-local/doc/manual/html/pod/perlre.html#Version 8 Regular Expressions](http://www.perl.com/CPAN-local/doc/manual/html/pod/perlre.html#Version_8_Regular_Expressions)

In addition, some rows have other filter types that are suitable for the type of job data represented by that row. Some examples are:

- ◆ Numeric value that a job's corresponding data value must be *greater than* in order for that job to be displayed. For example:
In the Kilobytes row selecting the filter type
greater than ###
shows only jobs for which the number of kilobytes transferred is greater than the number you enter in the filter value field of that row.
- ◆ Date and time which a job's corresponding data value must be *less than* in order for that job to be displayed. For example:
In the Start Time row selecting the filter type
before MM/DD/YYYY HH:MM:SS



shows only jobs that started before the date and time you enter in the filter value field of that row.

Menu Bar

Table 15 describes the menus and commands that are available on the menu bar:

Table 15. Activity Monitor Menus

Menu	Commands
File	<p>Append Master Server - (Applies only to Global Data Manager) Displays a dialog box where you can specify a master server to add to the tree. The addition applies only to the current session and the server will not be in the tree the next time you start the administrator interface. Also, see “Append Master Server” on page 670.</p> <p>Change Master Server - Displays a dialog box where you can specify another host that you want to monitor. This command does not apply to Global Data Manager.</p> <p>Save - Saves the contents of the job window to the last file that you saved to during this session. If you have not done a previous save, then NetBackup prompts you for a file name.</p> <p>Save As - Saves the contents of the job window to a file with the path and name that you specify.</p> <p>Redisplay Data - (Applies only to Global Data Manager) Redisplays job information that was last shown for servers that are currently selected in the right pane. This does not obtain new data from the servers.</p> <p>Run Activity Monitor - (Applies only to Global Data Manager) Runs the activity monitor on the currently selected servers. This obtains new data from the servers.</p> <p>Exit - Closes this window.</p>
Edit	<p>Delete - Deletes completed (Done) jobs that you have selected in the Jobs list. If you have selected any uncompleted jobs (Queued, Re-Queued, or Active), the Delete command is not available.</p> <p>Kill Job - Kills uncompleted jobs that you have selected in the Jobs list. If you have selected any Done jobs, the Kill command is not available.</p> <p>Kill All Backups - Kills all uncompleted backup jobs.</p> <p>Select All - Selects the entire listing.</p>



Table 15. Activity Monitor Menus (continued)

Menu	Commands
View	<p>Job Details - Displays detailed information about backup and archive jobs you have selected in the list.</p> <p>Toolbar - Shows or hides the toolbar.</p> <p>Refresh - Updates the job information in the right pane by retrieving new information from the master server(s) where the information originated. If the right pane shows the list of servers, the Refresh command is not available.</p> <p>Preferences - Displays a dialog box where you can select preferences for displaying the toolbars. This dialog box also has a tab where you can set the automatic refresh rate for the jobs list.</p> <p>Filters - Displays a dialog box where you can set the Activity Monitor to show only the job data you want to see..</p>
Help	<p>Help Topics - Provides online help information.</p> <p>Troubleshooter - Starts the troubleshooting wizard for jobs that you have selected in the jobs list.</p> <p>About Activity Monitor - Displays program information, version number, and copyright.</p>

Toolbar

The buttons on the toolbar provide shortcuts for menu commands. If the toolbar is not visible or you want to change its size, click the Preferences command on the View menu and make the desired choices in the Preferences dialog box.

Status Bar

The status bar appears at the bottom of the window and shows the number of jobs being monitored and the state they are in. The right side of the status bar shows the master server where you are currently monitoring jobs (if filtering is used, the numbers reflect only the unfiltered jobs).

Shortcut Menus (Global Data Manager only)

If you are administering NetBackup from a Master of Masters, clicking the right mouse button when the pointer is over either pane in the window, displays a shortcut menu with commands that apply to what is selected in that pane. Most of the shortcut commands are also on the menu bar and are explained in the “Menu Bar” topic. The exceptions are Recognize Master Server(s) and Ignore Master Server(s), which are explained in “Interface Commands for GDM” on page 670.



Refreshing the Display

You can manually refresh the Activity Monitor display and also set it to refresh automatically at predetermined intervals. For a manual refresh, click the Refresh command on the View menu. To set the automatic refresh rate, click Preferences on the View menu and specify the Refresh Rate on the Server Access tab.

During either manual or automatic refreshes, NetBackup normally updates only the data that has changed since the last refresh. For example, it obtains new data for a job that is in progress but not for a completed job (which does not change). This greatly reduces the refresh time.

If you run the Activity Monitor by choosing Run Activity Monitor on the File menu, NetBackup updates all its data. This operation takes longer to complete.

Automatic Refresh Behavior

The following provides additional information on automatic refreshes and will be useful in interpreting the behavior of the display.

There are two parts to the automatic refresh cycle:

- ◆ The first part is not visible to the user (there is no clock icon) and can be interrupted. For example, if you attempt to append a host the refresh stops and the Append Host dialog box appears.
- ◆ The second part shows the clock icon and cannot be interrupted. If the system is busy, the clock icon may not appear immediately but it is still too late to interrupt the refresh. If you happen to attempt an operation (for example, to append a host) between the time that the second part of the cycle begins and the clock icon appears, the following message appears:

```
"Currently unable to perform operation due to autorefresh in progress"
```

Interrupting either an automatic or manual refresh leaves the display reliable but can place the Activity Monitor in a state where it is uncertain which data requires updating. If this occurs, the Activity Monitor updates all its data at the next refresh. This will take much longer than the usual refresh where the server sends only changed data.

Refresh Behavior After Killing or Deleting Jobs

Because killing and deleting jobs takes longer than a refresh, the Activity Monitor does not automatically refresh the data after you initiate a kill or delete. If you initiate a kill or delete:

- ◆ The deleted jobs are immediately removed from the display. However, if you refresh all the data (for example, by running the Activity Monitor) before the deletion is complete, the jobs can reappear. This occurs because a deletion can take longer than a refresh. The jobs are removed again by the next refresh after the deletion has completed.
- ◆ The killed jobs continue to be displayed and following message appears:

```
"The request to kill the job(s) has been sent but the action is not yet complete. The results of this action will not be reflected in the Activity Monitor display until the action has completed and a Refresh occurs."
```

Choosing the Master Server to Monitor

If there is more than one NetBackup master server, choose the one where you want to monitor activity.

The steps required to choose a master server depend on whether you are administering NetBackup from a server that has the Global Data Manager option (that is, from a Master of Masters server).

If You Are Not on a Master of Masters

If you are not administering NetBackup from a Master of Masters, the name of the server that you are currently monitoring appears on the Master Server line in the area above the Jobs list. To monitor activity on another NetBackup master server, proceed as follows:

1. On the File menu, click Change Master Server.
2. In the dialog box, specify the name of the NetBackup master server that you want to monitor.
3. Click OK.

The name of the master server that you specified now appears in the status bar at the bottom of the window.

If You Are on a Master of Masters (requires Global Data Manager)

If you are administering NetBackup from a Master of Masters, you can select one or more master servers. See “To Monitor Jobs With Global Data Manager” on page 196.

If the server you want is not in the tree, append it by using the Append Master Server command on the File menu.



Monitoring Jobs

The procedure for monitoring jobs depends on whether you are administering NetBackup from a server that has the Global Data Manager option.

To Monitor Jobs Without Global Data Manager

1. In the NetBackup Administration window, click the Activity Monitor icon.
The Activity Monitor window appears.
2. Interpret the jobs list as explained in Table 16.
3. Perform other required actions, such as deleting or killing jobs. See the topics that follow this procedure.
4. If there is more than one master server, use the Change Master Server command on the File menu to choose another master server on which to monitor jobs see “Choosing the Master Server to Monitor” on page 195.

To Monitor Jobs With Global Data Manager

1. In the NetBackup Administration window, click the Activity Monitor icon.
The tree in the left pane shows the list of master servers. The right pane shows the same list or details.
2. To list the jobs for a single master server, click on it in the left pane:
 - ◆ If the server has not been previously selected, NetBackup retrieves job information from that server and displays it in the right pane.
 - ◆ If the server has been previously selected, the previously retrieved data reappears in the right pane.
 - ◆ If the server is ignored, no information appears in the right pane.
3. To obtain a combined List of jobs for multiple master servers:
 - a. Click the master of masters in the left pane so the list of servers appears in the right pane.
 - b. Select the servers in the right pane.
 - c. Position the pointer in the right pane, press the right mouse button, and choose Run Activity Monitor from the shortcut menu.
4. To obtain a combined list of jobs for all recognized servers, click All Recognized Master Servers in the left pane.
5. Interpret the jobs list as explained in Table 16.

6. Perform other required actions, such as deleting or killing jobs. See the topics that follow this procedure.

Table 16. Jobs List

Column	Meaning
Master	Master server where the job was run.
Job ID	Identifier that NetBackup assigns to each job. The identifier is unique on the server where the job was run.
Type	Backup, Archive, or Restore.
State	Queued - Jobs that are in the NetBackup scheduler queue. A queued restore job is one for which NetBackup is still determining which files are needed. Active - Currently active jobs. Re-Queued - Jobs that are back in the scheduler queue as retries because the previous attempt was unsuccessful. Done - Completed jobs.
Status	NetBackup status code that indicates the completion status. There is no value until the job is done. A status of zero (0) means that the job completed successfully. Any other completion value for status indicates a problem. To get the text that describes the status, double-click on the job.
Class	Name of the class that NetBackup is using to back up the client.
Schedule	Name of the schedule that NetBackup is using to back up the client.
Client	Name of the client for which the backup, archive, or restore is being performed.
Media Server	NetBackup server controlling the media.
Start Time	Date and time that the first try was initially queued.
Elapsed Time	Amount of time that has elapsed since the job was initially queued.
St. Unit	Name of the storage unit that NetBackup is using.
Ended	Date and time that the operation was completed.
Try	For active jobs, this indicates the number of the current attempt. For Done jobs it indicates the total number of attempts.
Operation	For active jobs, this indicates the operation that is currently being performed.



Table 16. Jobs List (continued)

Column	Meaning
Kilobytes	Number of kilobytes that have been written. The job monitor updates this column every 10 megabytes for an active job (every 30 megabytes for an Auspex FastBackup job).
Files	Number of files that have been written.
Pathname	For Active jobs, this is 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. Instead, each of the first 25 files appear, and after that only every 25th file until 500 files have been backed up. After 500 files, every 500th file appears.
% Comp (Estimated)	Percentage of the job that is complete. For backups, it is based on the size of the previous backup for the same class, client, schedule, and retention level. 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 restores or archives, the estimate is based on other factors.
Job PID	Process ID. If the backup is multiplexed, all jobs associated with the same multiplexed storage unit have the same PID.
Owner	Owner of the job.
KB Per second	Average kilobytes per second over the length of the current attempt.

To Delete Done Jobs

1. In the Jobs list, select the completed jobs that you want to delete.
2. On the Edit menu, click Delete.

To Kill Uncompleted Jobs

1. In the Jobs list, select the uncompleted jobs that you want to kill.
An uncompleted job is one that is in the Queued, Re-Queued, or Active state.
2. To kill selected jobs, click Kill Job on the Edit menu.

Note To kill all uncompleted backup jobs in the jobs list, click Kill All Backups on the Edit menu.

To Save Job Information

1. On the File menu, click Save (or Save As). The Save dialog box appears.

Save: writes the contents of the Jobs window to the last file that you saved to during this session. If you have not done a previous save, then NetBackup prompts you for a file name.

Save As: prompts you for the name of the file to which NetBackup should save the contents of the job window.

To save in an existing file, you must use the Overwrite file button.

2. Type the fully qualified path name for the file as it will exist on the master server where you are currently logged in.
3. Click OK.

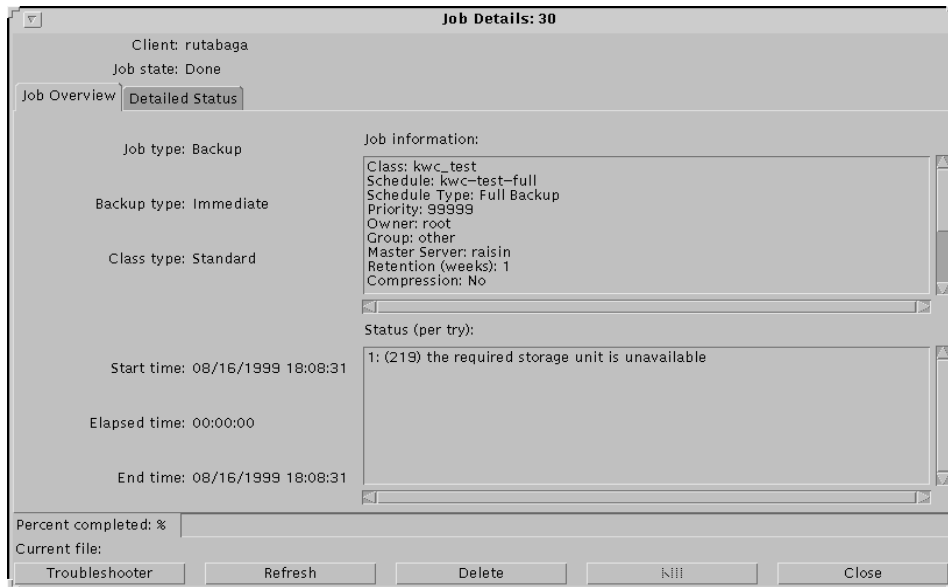
NetBackup saves the job information in a tab delimited format that most spreadsheets can use.

To Monitor Detailed Status on Backup or Archive Jobs

1. In the job list, select the backup or archive jobs for which you want detailed status. Detailed status is not available for restore jobs or for queued or requeued backup jobs.
2. On the View menu, click Job Details (or double-click on the job).

A Job Details dialog box appears for each job that you selected.





Job Details Dialog Box Description

The Job Details dialog box has detailed job information about a selected backup or archive job.

- ◆ The top of the dialog box shows the name of the client and the job state. When the job is complete, the state changes to Done. The bottom of the dialog box shows what percentage of the job is done and the current file being written.
- ◆ The Job Overview tab provides general information about the entire job. See Table 17.
- ◆ The Detailed Status tab provides detailed information about the job. See Table 18.

In addition to viewing detailed status, you can perform the following actions on a job:

- ◆ Delete a Done job by clicking Delete.
- ◆ Kill an uncompleted backup or archive job by clicking Kill. An uncompleted job is one that is in the Queued, Re-Queued, or Active state.
- ◆ For jobs that return a NetBackup status code, you can click the Troubleshooter button to start the troubleshooting wizard and obtain explanations and corrective actions.

Table 17. Job Details - Overview Tab

Column	Meaning
Job Type	See Table 16.
Backup type	Scheduled, User-directed, immediate (manual backup).
Class type	Type of class to which the client belongs. For example, MS-Windows-NT or Standard.
Start time	Date and time that the first try was initially queued.
Elapsed time	Amount of time that has elapsed since the job was initially queued.
End Time	Date and time that the operation was completed.
Class	See Table 16.
Schedule	See Table 16.
Schedule type	Type of schedule controlling the backup. For example, Full or Cumulative-Incremental
Priority	Priority for the class as specified by the Priority class attribute.
Owner	See Table 16.
Group	Group to which the job owner belongs.
Master server	Master server for the class and client that are being backed up.
Retention	Retention level assigned to the backup. This is specified on the schedule.
Compression	Yes if NetBackup is using software compression. Otherwise, No.
File List	List of files that are being backed up for the class.
Status (per try)	Text that describes the completion status for the job.

Table 18. Job Details - Detailed Status Tab

Column	Meaning
Attempt	Shows the number of attempts if NetBackup tried more than once. This field does not appear if there is only one attempt.
Job PID	See Table 16.
Storage unit	See Table 16.
Media Server	See Table 16.



Table 18. Job Details - Detailed Status Tab (continued)

Column	Meaning
Started	Time when this attempt began.
Elapsed	Elapsed time for this attempt.
Ended	Time when this attempt ended.
Status	Events that have occurred up to this point. For example, this box has 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 completion status.
Current kilobytes written	Number of kilobytes that have been written to the media up to this time.
Current files written	Files that have been written to the media up to this time.
Kbytes Last Written	(present only for active jobs) Number of bytes written for the last backup of the class and schedule
Files Last Written	(present only for active jobs) Number of files written for the last backup of the class and schedule
Current File	(present only for active jobs) The current file being written.

Using the Troubleshooting Wizard

You can use the troubleshooting wizard to find explanations and corrective actions when a job fails and returns a NetBackup status code.

To Use the Wizard:

1. Select the job in the jobs list.
2. On the Help menu, click on Troubleshooter.
The Troubleshooting wizard appears with an explanation of the error.
3. To change the message context from that of the server that originated the message, click the From Originating Server button and choose the context in which you want to see the error (UNIX server or Windows NT/2000 server).

The explanation and corrective actions can be different depending on whether the server is UNIX or Windows NT/2000.

4. Click Next to see a list of corrective actions.

Note The wizard is also available from the Details tab in the Job Details dialog box.



Managing the Jobs Database

The `/usr/opensv/netbackup/db/jobs` directory has the database files for the job monitor. These files include those for jobs that are in progress and jobs that are DONE.

NetBackup uses the `/usr/opensv/netbackup/bin/admincmd/bpdbjobs` command to periodically clean up the jobs directory. By default, `bpdbjobs` deletes all DONE jobs that are more than three days old and moves newer DONE jobs to the `/usr/opensv/netbackup/db/jobs/done` directory until their three day retention period expires.

`bpdbjobs` runs each time you start `bprd`, the NetBackup request daemon. If `bprd` is active, it also automatically starts `bpdbjobs` when it performs its other cleanup tasks (this occurs the first time that `bprd` wakes up after midnight). These automatic startups occur regardless of whether you choose to run `bpdbjobs` at other times by using `cron` or alternate methods.

You can change the retention period for DONE jobs by using any desired combination of the methods described in the following topics.

Creating a job.conf File

The `job.conf` file, if it exists, is the first place that `bpdbjobs` checks for configuration information.

You create the `job.conf` file in the `/usr/opensv/netbackup/db/jobs` directory. Then, you set the job retention period by adding options from those described in Table 19. The following are example entries:

```
keep_days 7
keep_successful_hours 5
```

Assuming these options are not overridden by `BPDBJOBS_OPTIONS` or `bpdbjobs` command line options (see the next two topics), the activity monitor keeps unsuccessful jobs for seven days and successful jobs for five hours.

Table 19. Jobs Retention Period Options

Option ¹	Description
² <code>keep_days</code> <i>days</i>	Specifies how many days <code>bpdbjobs</code> keeps DONE jobs. This value can range between 1 and 30. Values outside this range are ignored. The default is three days.



Table 19. Jobs Retention Period Options (continued)

Option ¹	Description
² <code>keep_hours</code> <i>hours</i>	<p>Specifies how many hours <code>bpdbjobs</code> keeps DONE jobs. This value can range from 3 to 720. Values outside this range are ignored.</p> <p>The default is 72 hours.</p>
³ <code>keep_successful_days</code> <i>days</i>	<p>Specifies how many days <code>bpdbjobs</code> keeps successful DONE jobs. This value can range from 1 to 30 but must be less than <code>keep_days</code>.</p> <p>Values outside the 1 to 30 range are ignored. The default is three days.</p>
³ <code>keep_successful_hours</code> <i>hours</i>	<p>Specifies how many hours <code>bpdbjobs</code> keeps successful DONE jobs. This value can range from 3 to 720 but must be less than <code>keep_hours</code>.</p> <p>Values outside the 3 to 720 range are ignored. The default is 72 hours.</p>
<code>verbose</code>	<p>Causes <code>bpdbjobs</code> to log additional information in the activity log in the <code>/usr/openv/netbackup/logs/bpdbjobs</code> directory if this directory exists.</p>

Notes

1. The `keep_hours`, `keep_successful_hours` and other parameters in Table 19 can be either upper or lower case in any of the following:

- ◆ `job.conf` file
- ◆ `BPDBJOBS_OPTIONS` environment variable
- ◆ command line parameters

2. `keep_days` and `keep_hours` are mutually exclusive. If you specify both values, `bpdbjobs` uses only the last one it finds.

3. `keep_successful_days` and `keep_successful_hours` are mutually exclusive. If you specify both values, `bpdbjobs` uses only the last one it finds. A successful DONE job shows a status of 0 in the status column of the Jobs List. The status message is the requested operation was successfully completed.

4. The retention period values are measured against the time the job ended.



Specifying a BPDBJOBS_OPTIONS Environment Variable

The BPDBJOBS_OPTIONS environment variable is the second place that `bpdbjobs` checks for configuration information. This variable offers same options as described in Table 19, and provides a convenient way to set them up in a script. Options on BPDBJOBS_OPTIONS override corresponding options in the `job.conf` file.

The following is an example script named `cleanjobs`:

```
setenv BPDBJOBS_OPTIONS "-keep_days 5 -keep_successful_hours 3 -clean"
/usr/openv/netbackup/bin/admincmd/bpdbjobs ${*}
```

Executing the `cleanjobs` script overrides corresponding options in `job.conf`, and results in `bpdbjobs` using options in the BPDBJOBS_OPTIONS environmental variable, unless they are overridden by `bpdbjobs` command line options.

Using bpdbjobs Command Line Options

`bpdbjobs` options are the last thing that `bpdbjobs` checks for configuration information and these options override corresponding options in either `job.conf` or BPDBJOBS_OPTIONS. You can use the `bpdbjobs` command to delete or move DONE job files. The syntax for the jobs retention options to `bpdbjobs` usage is as follows:

```
bpdbjobs -version
bpdbjobs -clean [-keep_days days | -keep_hours hours ]
[-keep_successful_days days | -keep_successful_hours hours]
[-verbose]
```

Where:

- ◆ `-version` causes `bpdbjobs` to print the version string and halt. Any other switches are ignored.
- ◆ `-clean` causes `bpdbjobs` to delete DONE jobs that are older than a specified time period. DONE jobs that are not older than the specified time period are moved to the `jobs/done` directory.

The additional options that you can use with `-clean` to specify a retention period are:

```
-keep_days days
-keep_hours hours
-keep_successful_days days
-keep_successful_hours hours
-verbose
```

These `-clean` options have the same definition and defaults as those described in Table 19.



For example, the command

```
bpdbjobs -clean -keep_jobs 30
```

overrides any `keep_jobs` value that exists in either `job.conf` or `BPDBJOBS_OPTIONS`.

bpdbjobs Activity Log

If you need more detailed information, you can enable activity logging for `bpdbjobs`. To enable this activity logging, create the following directory:

```
/usr/opensv/netbackup/logs/bpdbjobs
```

Note Before using this or other activity logs, read the guidelines in the Activity Logs section of the *NetBackup Troubleshooting Guide - UNIX*.

This chapter contains topics related to the administration and management of NetBackup.

- ◆ Powering Down and Rebooting NetBackup Servers
- ◆ Managing Daemons
- ◆ Managing the Restore of Client Files
- ◆ Administering NetBackup Licenses
- ◆ Duplicating, Verifying, and Importing Backup Images
- ◆ Goodies Scripts
- ◆ Server Independent Restores
- ◆ Load Balancing
- ◆ Using NetBackup with Storage Migrator
- ◆ Allowing Nonroot Users to Administer NetBackup
- ◆ Authorizing NetBackup-Java Users
- ◆ Configuration Options for jbpSA
- ◆ Administrator's Quick Reference



Powering Down and Rebooting NetBackup Servers

Powering Down a NetBackup Server

1. Before powering down a NetBackup server:
 - ◆ Use the `/usr/opensv/netbackup/bin/bpps` script to ensure that no backups or restores are running.
 - ◆ Stop the NetBackup request daemon by executing:

```
/usr/opensv/netbackup/bin/bprdreq -terminate
```

Note The request daemon does not run on a media server.

2. Power down the server.

Rebooting a NetBackup Server

To reboot a NetBackup server that is powered down, do the following:

Master Server

1. Restart the system.
2. Start `ltid` if it is not already running.
3. Ensure that `bprd` and `bpdbm` are running by using the `/usr/opensv/netbackup/bin/bpps` script. If necessary, start `bprd` and `bpdbm`.

Media Server

1. Restart the system.
2. Start `ltid` if it is not already running.

Managing Daemons

Starting and Stopping bprd

The NetBackup request daemon, `bprd`, starts the scheduler and the NetBackup database manager, `bpdbm`, in addition to controlling other functions.

To enable `bprd` activity logging create the `/usr/opensv/netbackup/logs/bprd` directory before starting `bprd`.



Starting bprd

The networks and network daemons must be fully operational and the Media Manager daemons (if applicable) must be running before `bprd` is started. Execute `/usr/opensv/netbackup/bin/bpps` to determine the status of `bprd` and other NetBackup or Media Manager processes.

To start `bprd` execute:

```
/usr/opensv/netbackup/bin/bprd
```

You can also start `bprd` by executing the script:

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

By modifying this script you can easily change the `bprd` start options.

To automatically start `bprd` when the master server is booted, modify your master server's system initiation scripts as follows:

- ◆ For a SPARC server, change the `/etc/rc.local` file by adding lines similar to those shown below. In testing, these lines were placed after the test for `/etc/exports` and before the test for `/tftpboot`.

```
if [ -f /usr/opensv/netbackup/bin/initbprd ] then
    /usr/opensv/netbackup/bin/initbprd &
    sleep 2
    echo "NetBackup request daemon started." > /dev/console
else
    echo "NetBackup request daemon not started." > /dev/console
fi
```

- ◆ For a Hewlett-Packard server, add lines, similar to those for the SPARC server to the `localrc` function in the `/etc/rc` file.
- ◆ For a Solaris or SGI server, the

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

directory has scripts that you can put in the server's `/etc/rc2.d` and `/etc/rc0.d` directories, respectively.

Stopping bprd

To stop `bprd`, execute:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```



If the daemon has started any activities, the command lets them complete. With `bprd` stopped, NetBackup cannot perform any backup, archive, or restore operations. Stopping `bprd` does not stop `bpdbm`.

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 execute:

```
/usr/opensv/netbackup/bin/initbpdbm
```

To stop `bpdbm`, use:

```
bpdbm -terminate
```

For more information, see the `bpdbm(1M)` man page.

Displaying Active Processes With `bpps`

NetBackup provides a script that determines which NetBackup processes are active on a UNIX system. This script is named:

```
/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
```

You can prevent this script from showing processes that you do not want to check by adding them to an exclude list. Refer to comments within the script itself for more information.

If you add the `-a` option

```
/usr/opensv/netbackup/bin/bpps -a
```

the `bpps` script shows both NetBackup and Media Manager options.

Managing the Restore of Client Files

The discussions in this section cover the following aspects of managing restores of client files.

- ◆ Allowing Restores to an Alternate Client
- ◆ Setting Client List and Restore Permissions



- ◆ Reduce Restore Times by Indexing the Image Catalog
- ◆ Server-Directed Restores
- ◆ Set Original atime for Files During Restores

Note A related topic is “Rules for Using Host Names in NetBackup” on page 604. Incorrectly specified host names are often a factor in file restore problems.

Allowing Restores to an Alternate Client

The user interfaces on NetBackup clients have options for restoring files that were backed up by other clients. The client that did not back up the files is called an alternate client and the operation is called an alternate-client restore.

A client can restore files belonging to other clients only if the necessary configuration is performed on the NetBackup master server. This configuration consists of creating the `/usr/opensv/netbackup/db/altnames` directory on the master server and adding files to it as explained in this section. To undo the changes, remove the `altnames` directory and its files.

Caution The `/usr/opensv/netbackup/db/altnames` directory presents a potential breach of security, because it permits users to select and restore files from other clients, if those users have the permissions to locally create the files found in the backup.

How NetBackup Enforces Restore Restrictions

By default, NetBackup permits restores only to the client from which the files were backed up. NetBackup enforces this restriction by ensuring that:

The name specified by the NetBackup client name setting on the requesting client matches the peername used in the connection to the NetBackup server.

Where:

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 dialog box. To display this dialog box, start the NetBackup user interface on the client and click Specify machines on the Actions menu.
- ◆ On NetWare target clients, specify the client name in the `bp.ini` file.
- ◆ On Macintosh and UNIX clients, specify the client name in the user interface.



peername 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 *peername* is equivalent to the client's *host name*. When a client connects through a gateway, the gateway can use its own *peername* to make the connection.

Allowing All Clients to Perform Alternate-Client Restores

The administrator can allow all clients to restore backups belonging to other clients by creating the following file on the NetBackup master server:

```
/usr/opensv/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 peername of the requesting client does not have to match the NetBackup client name setting.

For example, assume that a user on a UNIX client named *freddie* wants to restore a file that was backed up by a client named *oscar*:

1. The administrator creates the following file on the master server:

```
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

2. The user on *freddie* changes the NetBackup client name setting in the client user interface on *freddie* to *oscar*.

Allowing a Single Client to Perform Alternate-Client Restores

The administrator can give a single client permission to restore backups belonging to other clients by creating an empty file named

```
/usr/opensv/netbackup/db/altnames/peername
```

for the client that is to have the restore privileges. This file must be on the master server.

In this case, the client named by *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.

For example, assume that a user on a UNIX client named *freddie* wants to restore files that were backed up by a client named *oscar*:

1. The administrator creates the following file on the master server:

```
/usr/opensv/netbackup/db/altnames/freddie
```

2. The user on freddie changes the NetBackup client name setting in the client user interface on freddie to oscar.

Allowing Alternate-Client Restores of Specific Client's Files

The administrator can give a single client permission to restore backups belonging to specific other clients by creating and then adding client names to the single client's `/usr/opensv/netbackup/db/altnames/peername` file.

In this case, the client named by *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.

For example, assume that a user on a UNIX client named freddie wants to restore files backed up by a client named oscar:

1. The administrator creates the following file on the master server:

```
/usr/opensv/netbackup/db/altnames/freddie
```

2. The administrator enters the name oscar on a separate line in the freddie file.
3. The user on freddie changes the NetBackup client name setting in the client user interface on freddie to oscar.

Examples of Restoring Files to an Alternate Client

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/opensv/netbackup/db/images/client_name
```

or be a member of an existing NetBackup class.

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_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_client* that were backed up from *backed_up_client*. The *restore_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_client` so it includes the name of *backed_up_client*.
 - or
 - ◆ Execute the `touch` command on the file:

```
/usr/opensv/netbackup/db/altnames/No.Restrictions
```
2. Log in on *restore_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 `peername` used in the connection to the NetBackup server. When the NetBackup client name is the host name for the client and matches the `peername` (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 on Figure 3.

Figure 3. Example Restore from Token Ring Client



In this example network, restore requests coming from the clients, alf, cheech, and chong, are routed through the TCP gateway. Because the gateway uses its own peername 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 peername 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 peername used on the request. Entries in the report will be similar to:


```
01/29/93 08:25:02 bpserver - request from invalid
server or client bilbo.dvlp.null.com
```

In this example, the peername is `bilbo.dvlp.null.com`.

2. After determining the peername, create the file


```
/usr/opensv/netbackup/db/altnames/peername
```

on the NetBackup master server.

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 Alternate-Client Restores” on page 212.



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 Alternate-Client Restores of Specific Client’s Files” on page 213.

Note that this example requires no changes on the clients.

Example 3

If the files cannot be restored by using the method in Example 2, perform the following steps:

1. On the NetBackup master server, add the `VERBOSE` entry to the `bp.conf` file.
2. Create the activity log directory for `bprd` by executing:

```
mkdir /usr/opensv/netbackup/logs/bprd
```

3. On the NetBackup server, stop the NetBackup request daemon, `bprd`, and restart it in verbose mode by executing:

```
/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 peername connection used by *restore-to-client*.

Examine the failure as logged in the All Log Entries report or examine the `bprd` activity 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.

or

- ◆ Execute the `touch` command on the

```
/usr/opensv/netbackup/db/altnames/peername
```

file. This lets *restore_client* access any *backed-up-client* backups by changing its NetBackup client name setting to specify the *backed-up-client* client.

or

- ◆ 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.

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:

- ◆ Adding Clients to the NetBackup Client Database
- ◆ Setting the List and Restore Permissions
- ◆ Examples

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 286 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 use:

```
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 execute 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 use:

```
bpclient -delete -client client_name
```
- ◆ To list a client entry use:

```
bpclient -L -client client_name
```
- ◆ To list all client entries use:

```
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.

```
/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 named 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_LIST` 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_LIST` changes the default to deny lists.

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 and 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 that you have three clients: shark, eel, and whale.

Example 1

Prevent lists and restores on all three 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 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 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 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.

Reduce Restore Times by Indexing the Image Catalog

If you have large numbers of backups, you can 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.

You can generate indexes for one client or all clients and for up to nine levels of directories by using the following command:

```
/usr/opensv/netbackup/bin/index_clients level client_name
```

Where:

- ◆ *level* is the number of directory levels you want to index (1 to 9). These levels refer to the directories from which files were backed up on the client. For example, if the search is for `/payroll/smith/taxes/97` and *level* is 2, then NetBackup starts the search at `/payroll/smith`. The default is 9.



- ◆ *client_name* is the name of the client whose backups you want to index. The default is all clients.

Execute 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 a client named mars to index level 5 (five levels of directories), execute:

```
/usr/opensv/netbackup/bin/index_clients 5 mars
```

- ◆ To index selected clients, execute 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, execute:

```
/usr/opensv/netbackup/bin/index_clients 3
```

- ◆ To index all NetBackup clients to index level 9, execute:

```
/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

- ◆ 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.

Improve 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
```

Where *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 percent 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.

Server-Directed Restores

An administrator can use the backup, archive, and restore interface on the NetBackup master server to direct restores to any client, providing NetBackup on the client is configured to permit them. For instructions, see the NetBackup user's guide (UNIX).

Set Original atime for Files During Restores

During a restore NetBackup by default sets the `atime` for each file to the current time. If you want NetBackup to set the `atime` for each restored file to the value it had when it was backed up, create the following special file on the client.



```
/usr/opensv/netbackup/RESTORE_ORIGINAL_ETIME
```

Note If you are using VERITAS Storage Migrator, do not create the `RESTORE_ORIGINAL_ETIME` file. If you do, it is possible that restored files will be immediately migrated because of their older `etime`.

Administering NetBackup Licenses

The license key for NetBackup on each computer is initially entered when the software is installed. You can view, add, or delete licenses later by using the NetBackup License Key utility. This can be necessary, for example, when changing to a different level of NetBackup or adding separately-priced options. You can also list the licenses on a remote host.

To start the NetBackup License Key utility:

1. Execute the `/usr/opensv/netbackup/bin/admincmd/get_license_keys` command.

The following 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
```

2. Choose from the following actions:
 - ◆ To list only the licenses that are currently active, choose F. Licenses that are expired do not appear in this listing. You can specify either the local or a remote host.
 - ◆ To list all registered licenses (active or inactive) choose L. You can specify either the local or a remote host.
 - ◆ To add a new license key, choose A and then type the key at the prompt.
 - ◆ To delete a license from the list, choose D and then type the license key at the prompt. Note that you cannot delete evaluation license keys.
 - ◆ For help on the utility, choose H.
 - ◆ To quit the utility, choose q.



Duplicating, Verifying, and Importing Backup Images

Duplicating Backup Images

NetBackup can duplicate backups that are already on NetBackup volumes and are not expired.

It is possible to duplicate backups from:

- ◆ One storage unit to another
- ◆ One media density to another
- ◆ One server to another
- ◆ Multiplex to nonmultiplex format
- ◆ 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 multiplexing session.)

You cannot:

- ◆ Create a copy while the backup is being created
- ◆ Duplicate a backup when its retention period has expired
- ◆ Use the NetBackup scheduler to automatically schedule duplications
- ◆ Duplicate a backup of the NetBackup catalogs
- ◆ Create a multiplexed duplicate of the following:
 - ◆ Auspex FastBackup
 - ◆ Flashbackup
 - ◆ NDMP backup
 - ◆ Backups from disk type storage units
 - ◆ Backups to disk type storage units
 - ◆ Nonmultiplexed backups

Note Do not use the duplicate feature while a NetBackup catalog backup is running. This results in the catalog backup not having information about the duplication.

The duplicate operation requires a minimum of two drives: one to read the original and another on which to create the copy. 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.



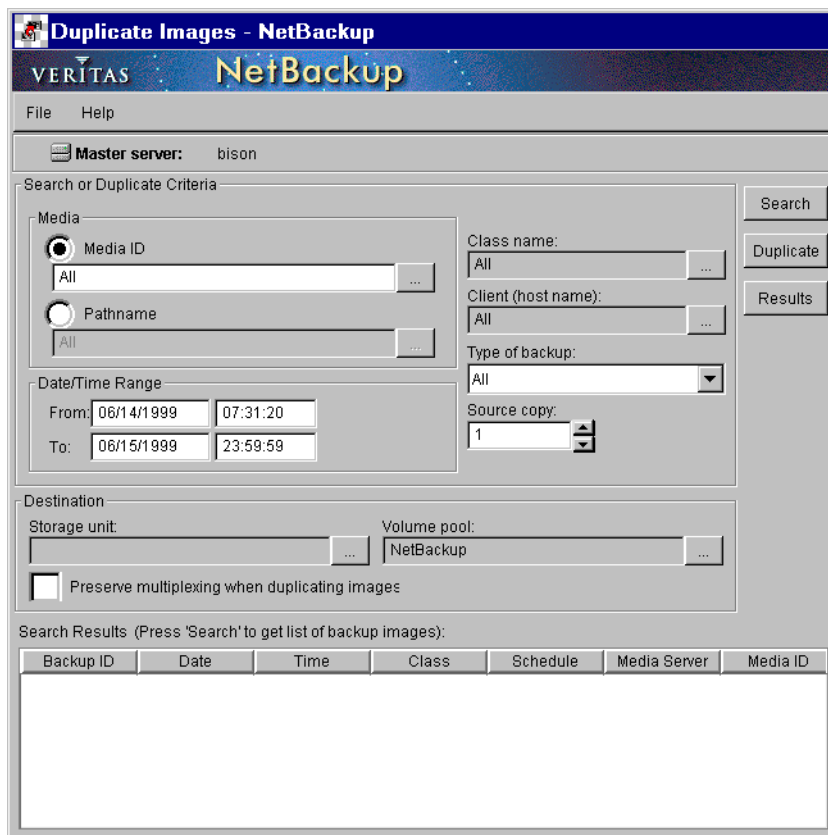
By default, duplication is done serially and attempts to minimize media mounts and positioning time. Only one backup is processed at a time. If multiplex duplication is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated (see list above). Then, multiplexed backups are duplicated.

The duplicate copy shares many attributes of the primary (original) copy, including backup ID. Initially, the expiration date is the same for the duplicate and primary copies. You can use the `bpxpdate` command to change the expiration date of the duplicate. Other attributes, such as elapsed time, apply only to the primary. It is the primary copy that NetBackup uses to satisfy restore requests (see “Restoring from Duplicated Backups” on page 231).

To Duplicate Backup Images

1. In the NetBackup Administration window, click Duplicate on the Images menu.

The NetBackup Duplicate Images dialog box appears.



Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), select the master server in the tree in the left pane to make the dialog box appear. Then, skip to step 3.

2. If you are not on a master of masters and there is more than one NetBackup master server, check the Master Server line below the menu bar and ensure that it names the master server with the backups that you want to duplicate. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.
 - b. In the dialog box, type the name of the desired server and click OK.
 - c. Click OK.

If you cannot change to the specified server, refer to “If You Cannot Change to Another Master Server” on page 23.

3. In the Search or Duplicate area of the dialog box, specify the criteria explained in Table 20.

NetBackup uses this criteria to build a list of backups from which you can make your selections.

Table 20. Search or Duplicate Criteria for Backup Images

Search Criteria	Description
Media ID	Media ID for the volume that has the desired backups. To change the value, type a media ID in the box or click the button and select from the list. If the original is fragmented, NetBackup duplicates only the fragments that exist on the specified volume.
Pathname	For a disk storage unit, select the check box and specify the file path that includes the originals.
Date/time range	Range of dates and times that includes all the backups that you want to duplicate. The default range is determined by the Interval for status reports global attribute. The default for that attribute is the 24 hour period prior to the From date and time.
Class name	Class for which the selected backups were performed. To change the value, click the button and select from the list.
Client (host name)	Host name of the client that produced the originals. To change the value, click the button and select from the list.



Table 20. Search or Duplicate Criteria for Backup Images (continued)

Search Criteria	Description
Type of backup	Type of schedule that created the backups you are duplicating. To change the value, click the button and select from the list.
Source copy number	<p>NetBackup allows only two copies. You can create a duplicate from either copy 1 or 2. However, if a copy 1 and copy 2 already exist and are not expired, NetBackup does not permit further duplication.</p> <p>When two copies exist and one is expired, the duplication assigns the number of the expired copy to the duplicate. For example, if copy 1 is expired enter 2 as the Source Copy. NetBackup then duplicates copy 2 and the duplicate becomes the new copy 1.</p> <p>If only one copy currently exists, select copy 1 as the Source Copy and the duplicate becomes copy 2.</p> <p>The default Source Copy Number is 1.</p>

Note If you are running the administration interface on a Master of Masters, skip to step 8.

4. Specify the storage unit on which to store the copies.

If a storage unit has multiple drives, you can use it for both the source and destination. The destination storage unit is a required parameter (except for searches).

5. Specify the volume pool for the copies (not applicable to disk type storage units).

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.

Choose a volume pool that already exists or add a new volume pool. If you add a new pool, also add a volume to this pool so there is one available for the duplicate.

6. Select the backups and start the duplication as follows:

Note If your selections include multiplexed backups and you want to keep them multiplexed in the duplicate, check the Preserve multiplexing box. 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 multiplexing session.)

To select and duplicate all backups that meet the criteria you specified, click Duplicate and go to step 7 to view the results.



To select and duplicate only specific backups from those that meet the criteria you specified:

a. Click Search.

No duplication is done at this point. NetBackup searches for backups that meet the criteria. The results of the search appear in the Search Results box. For example, if you used the defaults for the selection criteria, the list shows all the media IDs that NetBackup has used for backups.

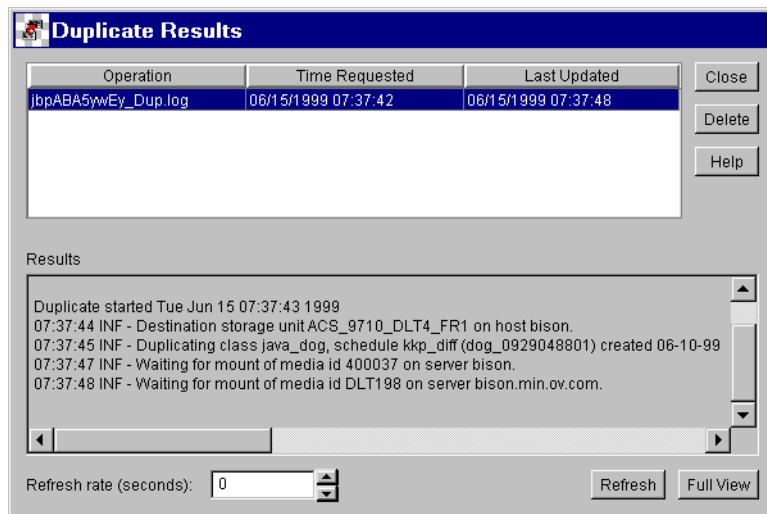
b. In the Search Results box, select the backups that you want to duplicate.

If you do not select any backups, NetBackup duplicates all the backups in the list.

c. Click Duplicate to start duplicating your selections and go to the next step.

7. To view the results, click the Results button.

The Duplication Results dialog box appears.



Note The following steps apply only if you are running the administration interface on a Master of Masters.

8. Select the backups and start the duplication as follows:

Note If your selections include multiplexed backups and you want to keep them multiplexed in the duplicate, check the Preserve multiplexing box. If you do not duplicate all the backups in a multiplexed group, the duplicate group will have a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together in a single multiplexing session.)



To duplicate all backups that meet the criteria you specified, click Duplicate and go to step 9.

To duplicate only specific backups from those that meet the criteria you specified:

- a. Click Search.

NetBackup searches for backups that meet the criteria. The results of the search appear in the Search Results box. For example, if you used the defaults for the selection criteria, the list shows all the media IDs that NetBackup has used for backups.

- b. In the Search results box, select the backups that you want to duplicate.

If you do not select any backups, NetBackup duplicates all the backups in the list.

- c. Click Duplicate.

The Duplicate Images Dialog box appears. The list shows all media IDs that will be duplicated.

9. In the Duplicate Images dialog box:

- a. If there are unwanted backups in the list, select them and press Delete.

- b. Specify the storage unit on which to store the copies.

If a storage unit has multiple drives, you can use it for both the source and destination. The destination storage unit is a required parameter (except for searches).

- c. Specify the volume pool for the copies (not applicable to disk type storage units).

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.

Choose a volume pool that already exists or add a new volume pool. If you add a new pool, also add a volume to this pool so there is one available for the duplicate.

- d. Click OK to start duplicating backups.

10. In the Duplicate Images dialog box, click Results to view the results of the duplication.

The Duplication Results dialog box appears.

Duplication Results Dialog Box

The selection list in the upper section of the dialog box shows all existing log files. To view a log file, select it from this list.



The Results section shows the contents of the selected log file. To refresh the Results section for an operation in progress, click the Refresh button. To set a refresh interval, specify a value in the Refresh Rate (seconds) box at the bottom of the dialog box. To view an entire log file, click Full View. The following figure shows an example verification log.

```

09:58:50 INF - Destination storage unit odl on host eel.
09:58:52 INF - Duplicating class ODL0, schedule full (giskard_0784364892) created 11 /04/94 01:08:12,
media id ODL01A.
09:59:01 INF - Waiting for positioning of media id ODL07B on server eel.
09:59:05 INF - Waiting for positioning of media id ODL01A on server eel.
09:59:15 INF - Waiting for positioning of media id ODL07B on server eel.
09:59:17 INF - Beginning duplication on server eel of client giskard image,
copy 2
09:59:22 INF - Waiting for positioning of media id ODL01A on server eel.
09:59:25 INF - Beginning duplicate on server eel of client giskard.
10:00:05 INF - Duplicate of backupid giskard_0784364892 successful.

10:00:07 INF - Duplicating class ODL0, schedule full (yak_0784364986) created 11 /09/94 01:09:46, media id
ODL01A.
10:00:31 INF - Waiting for positioning of media id ODL07B on server eel.
10:00:33 INF - Beginning duplication on server eel of client yak image, copy 2
10:00:33 INF - Waiting for positioning of media id ODL01A on server eel.
10:01:48 INF - Duplicate of backupid yak_0784364986 successful.

10:01:50 INF - Duplicating class ODL0, schedule full (raistlin_0784365070) created 11/09/94 01:11:10, media
id ODL01A.

```

Restoring from Duplicated Backups

Each backup is assigned a *primary copy*. NetBackup uses the primary copy to satisfy restore requests. If the primary copy is unavailable and you have created a duplicate, change the primary copy by entering the following command (all on one line):

```
/usr/openv/NetBackup/bin/admincmd/bpduplicate -npc pcopy -backupid bid
```

Where:

pcopy is the copy number you want to be the new primary.

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.

The `bpduplicate` command writes all of its output to the NetBackup logs and nothing appears in the command window.



After changing the duplicate to be the primary copy, you can use the NetBackup interface on the client to list and restore files from the backup. See the NetBackup user's guide for the client for instructions.

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 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 is important only for SQL-BackTrack backups.

- ◆ When you duplicate an image, both the primary and duplicate copies have the same expiration date. To change the expiration date of either copy, use the `bpxpdate` command. See the `bpxpdate(1M)` man page for more information.

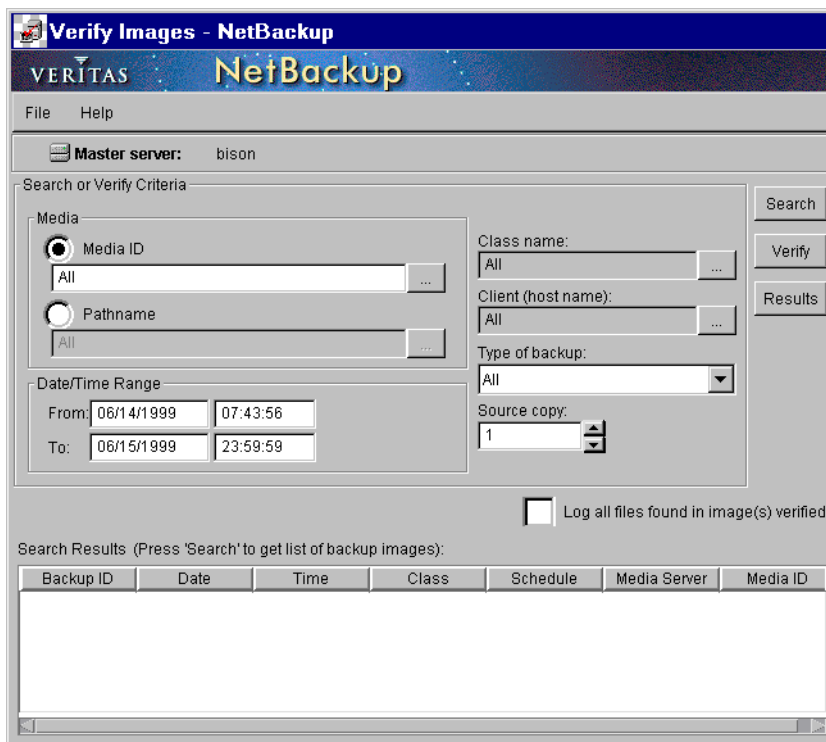
Verifying Backup Images

NetBackup can verify the contents of a backup by reading the volume and comparing its contents to what is recorded in the NetBackup catalog. This operation does not compare the data on the volume with the contents of the client disk. However, it does read each block in the image thus verifying that the volume is readable (though there could be data corruption within a block). NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

The procedure is as follows:



1. In the NetBackup Administration window, click Verify on the Images menu.
The Verify Images dialog box appears.



Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), select the master server in the tree in the left pane to make the dialog box appear. Then, skip to step 3.

2. If you are not on a master of masters and there is more than one NetBackup master server, check the Master Server line below the menu bar and ensure that it names the master server that has the backups you want to verify. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.
 - b. In the dialog box, type the name of the desired server.
 - c. Click OK.

If you cannot change to the specified server, refer to “If You Cannot Change to Another Master Server” on page 23.



3. In the search or verify section of the dialog box, specify the search criteria as explained in Table 21.

NetBackup uses this criteria to search the image catalog and accumulate a list of the backups from which you can make your selections.

Table 21. Search Criteria for Verifying Backup Images

Search Criteria	Description
Media ID	Media ID for the volume that has the desired backups. To change the value, type a media ID in the box or click the button and select from the list. Backups that have fragments on another volume are included, as they exist in part on the specified volume.
Pathname	For a disk storage unit, select the check box and specify the file path that has the desired backups.
Date/time range	Range of dates and times that includes all the backups you want to verify. The default range is determined by the NetBackup global attribute, Interval for status reports. The default for that attribute is the 24 hour period prior to the From date and time.
Class name	To change the value, click the button and select from the list.
Client (host name)	To change the value, click the button and select from the list.
Type of backup	Type of schedule that created the backups you are verifying. To change the value, click the button and select from the list.
Source Copy	You can verify either copy 1 or 2. The default Source Copy is 1.
Log all files found in image(s) verified	Provides information on each file that NetBackup verifies.

4. To have NetBackup verify all the backups that meet the search or verify criteria that you have specified, click **Verify**. Then, skip to step 8 to view the results of the operation.
5. To have NetBackup build a selection list of the backups that meet the search criteria, click **Search**.



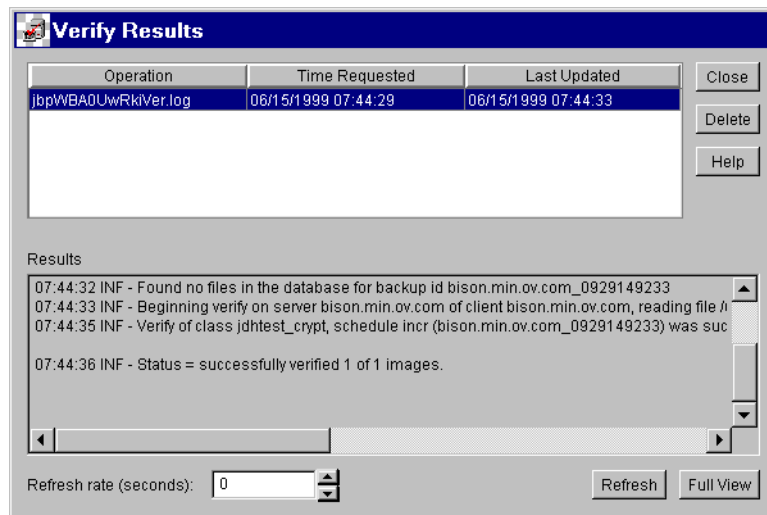
No verification is done at this point. NetBackup searches for backups that meet the criteria you entered earlier. The results of the search appears in the Search Results section of the dialog box. For example, if you used the defaults for the selection criteria, the list shows all the media IDs that NetBackup has used for backups.

6. In the Search Results section of the dialog box, select the backups that you want to verify.

If you do not select any backups, NetBackup verifies all the backups in the list.

7. Click Verify to start verifying your selections.
8. To view the results, click the Results button.

The Verification Results dialog box appears.



Verification Results Dialog Box

The selection list in the upper section of the dialog box shows all existing log files. To view a log file, select it from this list.

The Results section shows the contents of the selected log file. To refresh the Results section for an operation in progress, click the Refresh button. To set a refresh interval, specify a value in the Refresh rate (seconds) box at the bottom of the dialog box. To view an entire log file, click Full View.



Importing Backup Images

NetBackup can import backups that are expired or are from another NetBackup server. During an import operation, NetBackup recreates NetBackup catalog entries for the backups that are on the imported volume. This option is useful for moving volumes from one site to another and also for recreating NetBackup catalog entries for expired backups.

The expiration date for the imported items will be the current date plus the retention period. For example, if a backup is imported on November 14, 1999 and its retention period is one week, its new expiration date is on November 21, 1999.

Notes About Importing Backup Images

- ◆ The media contents and image list reports contain entries for the duplicates.
- ◆ You cannot import images generated for clients in Apollo wbak classes.
- ◆ 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 executing the following command:

```
/usr/opensv/NetBackup/admincmd/bin/bpexpdate -d 0 -ev 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 system administrator's guide for Media Manager has instructions for deleting and adding volumes.


To avoid this problem in the future, use unique prefix characters for media IDs on all servers.

- ◆ You cannot import a backup if an unexpired copy of it already exists on the server where you are trying to import it.

The procedure for importing backups is as follows:



1. Add the media IDs that have the backups to Media Manager on the server where you are going to import the backups (see the Media Manager system administrator's guide for instructions on adding media IDs).
2. In the Netbackup Administration window, click Import on the Images menu and then on the submenu click Step 1: Read Catalog From Media. The Step 1: Read Catalog From Media dialog box appears.



- ◆ In the Media host box, specify the name of the host that has the volume you are going to import (type the name or click the button and select from those listed).
 - ◆ In the Media ID (to import) box, enter the Media ID of the volume that has the backups you are importing.
 - ◆ On a Master of Masters (requires the Global Data Manager option), specify in the Master server box, the name of the master server to which you want to import the backups.
3. Click OK to start the process of reading the catalog information from the source volume.

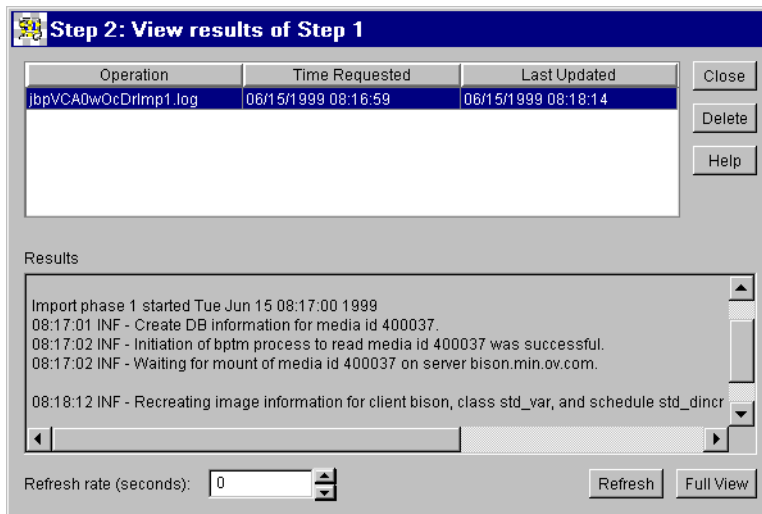
The result of this step is a list from which you can select what to import (no importing occurs at this stage).

Note Note that reading the catalog and building the list can take a long time to complete, since it is necessary to mount and read the tape.

4. In the NetBackup Administration window, click Import on the Images menu and then on the submenu click Step 2: View Results of Step 1.



The Import Results dialog box appears (see Import Results Dialog Box).

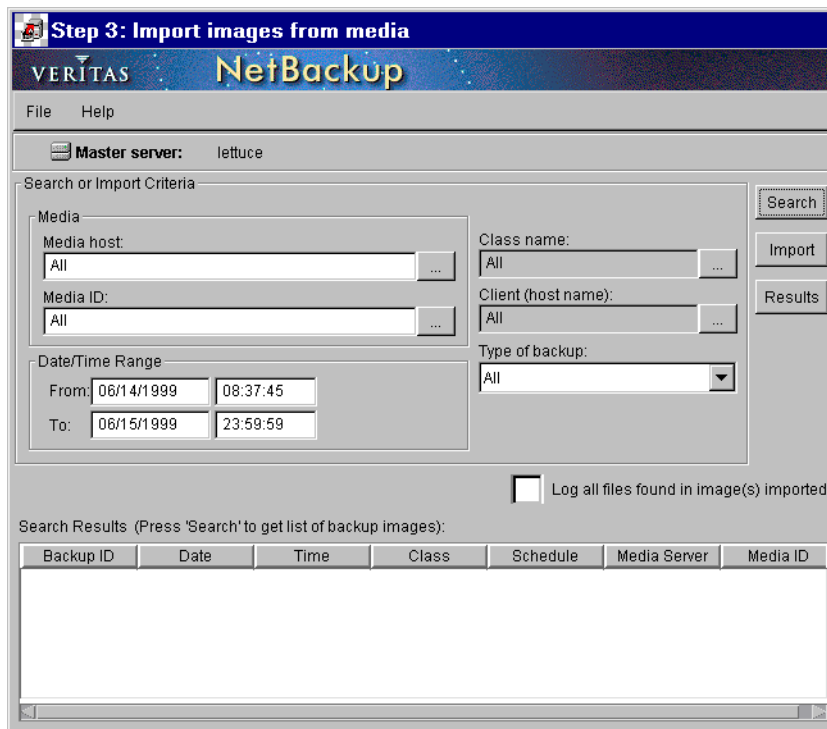


5. When the catalog has been read from the media, close the Import Results dialog box.

Note When importing backups that have fragments on more than one tape, do not start the import until you have read the catalog for *all* the tapes containing fragments. Otherwise, the import will fail with a message similar to: "Import of backupid failed, fragments are not consecutive."

6. In the NetBackup Administration window, click Import on the Images menu and then on the submenu click Step 3: Import Images From Media.

The Import Images dialog box appears. This is the dialog where you make your selections and start the import operation.



Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), select the master server from the tree in the left pane to make the dialog box appear. Then, skip to step 8.

7. If you are not on a master of masters and there is more than one NetBackup master server, check the Master Server line below the menu bar and ensure that it names the master server to which you want to import the backups. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.
 - b. In the dialog box, type the name of the desired server and click OK.
 - c. Click OK.

If you cannot change to the specified server, refer to “If You Cannot Change to Another Master Server” on page 23.

8. In the Import Images dialog box, complete the Search or Import criteria as explained in Table 22.



NetBackup uses this criteria to generate a list of backups from which you can select what you want to import.

Table 22. Search Criteria for Importing Backup Images

Search Criteria	Description
Media Host	NetBackup server where the volume resides, as shown in the Images on Media report. To change the value, type a name in the box or click the button and select from the list.
Media ID	Media ID for the volume that has the desired backups. The default is all media IDs that have been processed by Step 1 of the import operation. To change the value, type a media ID in the box or click the button and select from the list. If a backup begins on a media ID that has not been processed by Step 1, it is not imported. If a backup ends on a media ID that has not been processed by Step 1, the imported backup is incomplete.
Date/time range	Range of dates and times that includes all the backups that you want to import. The default range is determined by the Interval for status reports global attribute. The default for that attribute is the 24 hour period prior to the From date and time.
Class name	Class for which the selected backups were performed. To change the value, click the button and select from the list.
Client (host name)	Host name of the client. To change the value, click the button and select from the list.
Type of Backup	Type of schedule that created the backups that you are importing. To change the value, click the button and select from the list.
Log all files found in image(s) imported	Lists each file imported.

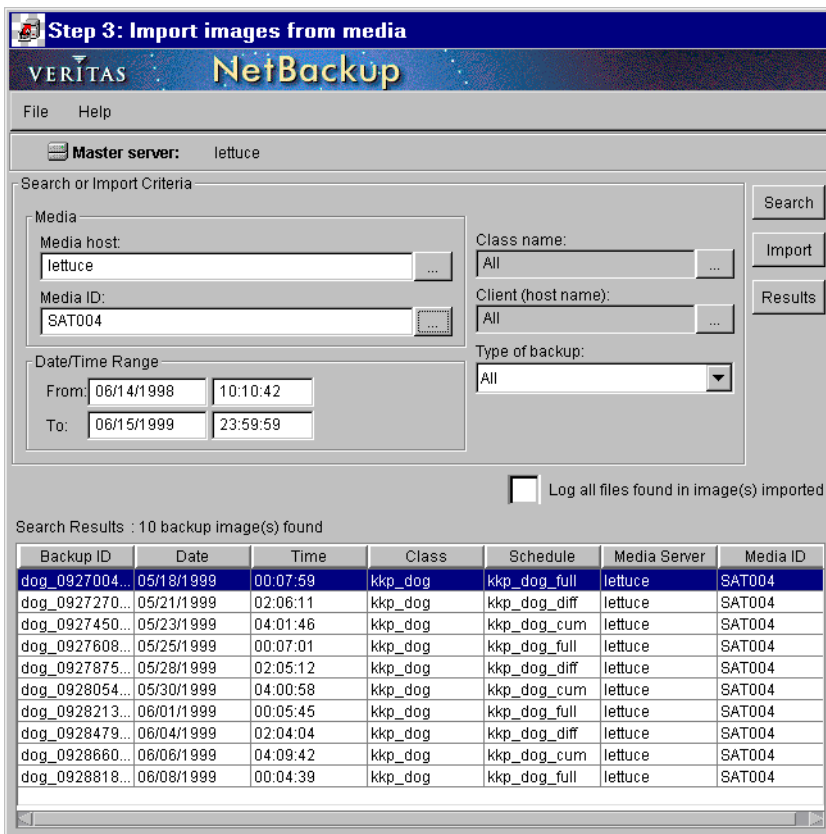
9. To have NetBackup import all the backups that meet the search or verify criteria that you have specified, click Import. Then, skip to step 11 to view the results of the operation.

10. To select specific backups to import, proceed as follows:



a. Click Search.

No importing is done at this point. NetBackup searches for backups that meet the criteria you entered previously. The results of the search appears in the Search Results section of the dialog box. For example, if you used the defaults for the selection criteria, the list shows backups on all the media IDs that NetBackup has used.



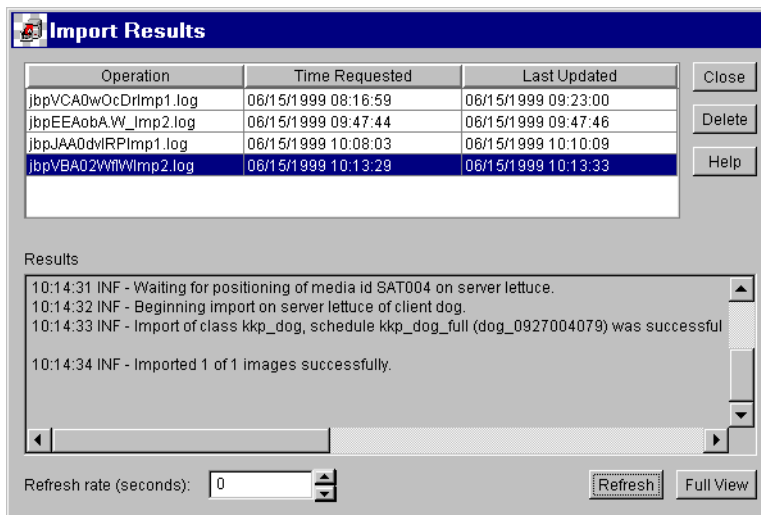
b. In the Search Results box, select the backups you want to import.

c. Click Import to start importing.

11. To view the progress of the import operation, click the Results button (you can also use Step 4: View results of Step 3 on the Images menu).



The Import Results dialog box appears.



Import Results Dialog Box

The selection list in the upper section of the dialog box shows all existing log files. To view a log file, select it from this list.

The Results section shows the contents of the selected log file. To refresh the Results section for an operation in progress, click the Refresh button. To set a refresh interval, specify a value in the Refresh rate (seconds) box at the bottom of the dialog box. To view an entire log file, click Full View. The following figure shows an example verification log file.

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 `crontab` 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 executing `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.

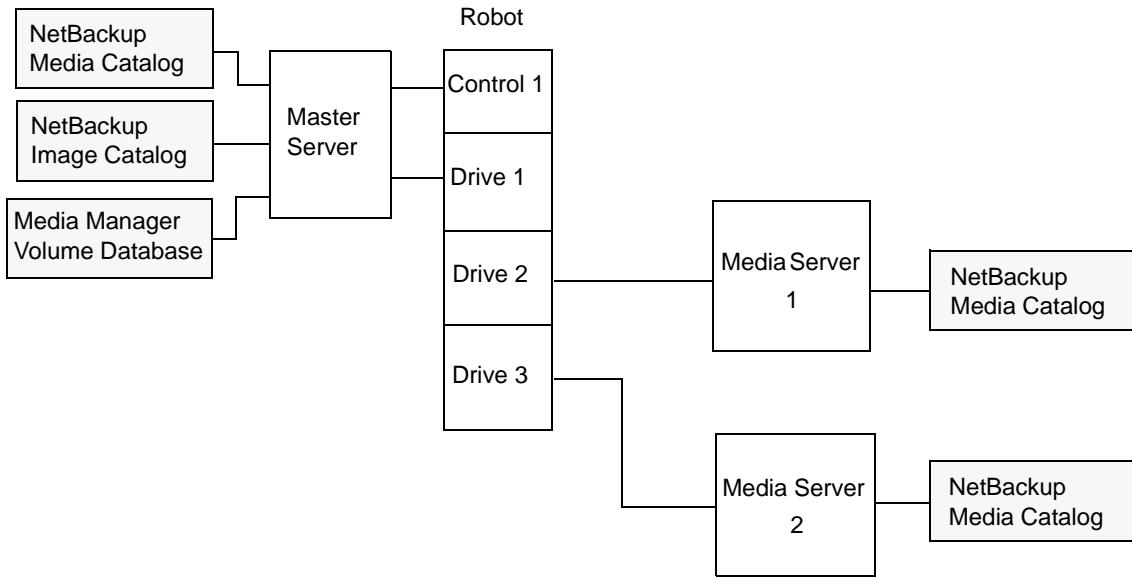
Due to the existence of the media catalog on each server where backups are written, restoring data through a device on another server is more involved than other restores but can be accomplished by using the methods described in this section. These methods do not require you to expire and import backup images; although, that can be useful in some instances (see “Related Topics” later in this section).

Supported Configurations

Figure 4 and Figure 5 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.



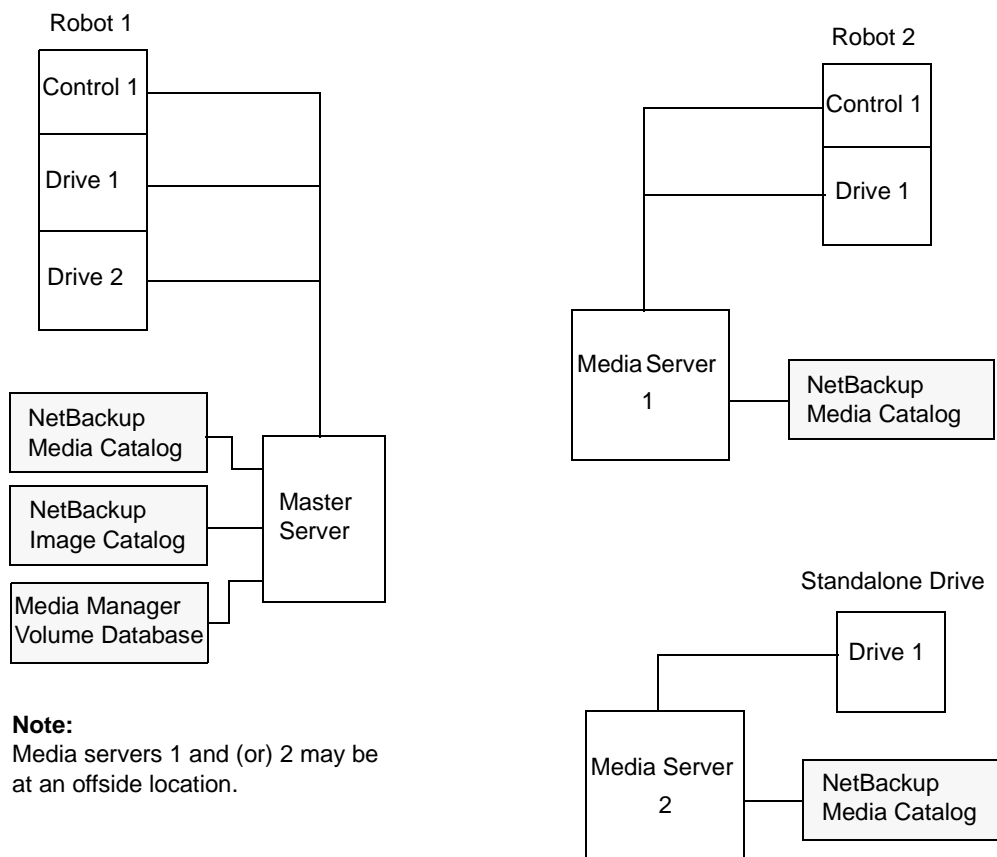
Figure 4. NetBackup Servers Sharing Robotic Peripherals



In Figure 4, 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.

Figure 5. NetBackup Servers with Separate Non-shared Peripherals



In Figure 5, 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
- ◆ Method 2 -- Overriding the Original Server
- ◆ Method 3 -- Automatic Failover to Alternate Server

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 -ev media_id -newserver hostname
      -oldserver hostname
```


(the `admincmd` command above must be on one line)

Windows NT/2000 NetBackup server (as administrator, from the MSDOS prompt):

```
cd install_path\NetBackup\bin\admincmd
bpmedia.exe -movedb -ev 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 NT/2000 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.

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 NT/2000 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.

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 NT/2000 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.

Related Topics

- ◆ 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.

- ◆ Utility for 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 from the NetBackup master server:

From a UNIX NetBackup server (as root user):

```
cd /usr/opensv/netbackup/bin/admincmd  
bpimmedia -spangroups -U -mediaid media_id
```

From a Windows NT/2000 NetBackup server (as administrator, from the MSDOS prompt):

```
cd install_path\NetBackup\bin  
bpimmedia.exe -spangroups -U -mediaid media_id
```

To display all media in all spanning groups, omit `-mediaid media_id` from the command.

Load Balancing

NetBackup provides ways to balance loads between servers, clients, classes, 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 class attribute for one or more of the classes that the server is backing up. For example, decreasing Limit Jobs Per Class reduces the load on a server on a specific network segment. Reconfiguring classes 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 execute 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 class or classes that are backing up the clients.

Give Preference To a Class

Increase the Limit jobs per class attribute for the preferred class relative to other classes. Or, increase the priority for the class.

Adjust Load Between Fast and Slow Networks

Increase the Limit jobs per class and Maximum jobs per client for classes 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, class, 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 class 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

Note This section applies only to Storage Migrator for UNIX and Storage Migrator Remote for UNIX.



If you require a storage migration product, VERITAS recommends that you use Storage Migrator for UNIX or Storage Migrator Remote for UNIX. NetBackup can use a disk type storage unit that is being managed by these products and can also back up the managed disks. NetBackup proceeds as follows when backing up a file that has been migrated by Storage Migrator:

- ◆ For user backups by nonroot users, NetBackup first caches the file and then backs it up.
- ◆ For scheduled backups and user backups by a root user, NetBackup backs up only the migration information for the files. The file data itself is neither backed up nor retrieved.

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 or restores will not be possible.

NetBackup restores migrated files to a migrated state. If the file has only been premigrated and the data does not reside on secondary storage, the file is migrated in and backed up. Premigrated means it will be migrated soon, but hasn't been copied to secondary storage nor had its disk space released.

Some NetBackup notification scripts also test for usage of Storage Migrator (see Appendix D”).

Be certain to observe the following when using Storage Migrator with NetBackup:

Set a Large Enough Media Mount Timeout

When NetBackup restores files to a disk storage unit managed by Storage Migrator, the Media mount timeout global attribute value is in effect during the caching of the (potentially) migrated backups. 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 (see “Defining NetBackup Global Attributes” on page 268).

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 223).

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.

Allowing Nonroot Users to Administer NetBackup

The following sections explain how to:

- ◆ Configure nonroot usage of all NetBackup administrator applications.
- ◆ Create a group specifically for Media Manager tape users.

Configuring Nonroot Usage of All NetBackup Administrator Applications

This section explains how to configure nonroot usage of all NetBackup administrator applications. This includes NetBackup-Java and all other NetBackup administration commands and interfaces (such as `bpadm` or `tpconfig`).

You must always configure nonroot usage on the system where you will run the administrator applications. For NetBackup-Java, this is the system that you specify in the login dialog box when starting the NetBackup-Java interface. In addition, if you are using `jnbSA`, you must allow nonroot usage of that command on the system where you start the interface (this step is not necessary for `jbP`SA).

For NetBackup-Java administration, you must configure nonroot usage on each system you plan to use.

Example 1

Assume you plan to start `jnbSA` on a Solaris system named `shark` and then specify an HP-UX system named `dolphin` in the login dialog box. Here, you must:

- ◆ Allow nonroot usage of the `jnbSA` command on `shark`
- ◆ Configure nonroot usage of NetBackup administrator applications on `Dolphin`.



Example 2

Assume you plan to start `jnbSA` on a Solaris system named `shark` and then specify that same system in the login dialog box. Here, you must allow nonroot usage of both the `jnbSA` command and NetBackup administrator applications on `shark`.

Example 3

Assume the NetBackup-Java Windows Display Console is installed on a Windows NT system named `trout`. You start the interface on `trout` and specify an AIX system named `whale` in the login dialog box. Here, you must configure nonroot usage of NetBackup administrator applications on `whale`.

Procedure

Perform the following steps as root to allow nonroot users to administer NetBackup with NetBackup-Java or any other administrator application or command (such as `bpadm` or `tpconfig`).

1. On the UNIX system where you are going to run the `jnbSA` command to start the NetBackup-Java administration interface, execute the following:

```
chmod 755 /usr/opensv/java/jnbSA
```

Note that this step allows nonroot users to start the interface on this system but does not provide administrative privileges. Also note that this step is not required for `jbpsa` because its normal permissions allow nonroot usage.

2. On the UNIX system that you will specify in the login dialog box when starting the NetBackup-Java interface, create distinct file-system groups as desired for the applications that will have nonroot usage. If you want all nonroot administrators to have privileges for all applications, create only one distinct file-system group.
 - ◆ If the system is Solaris or HP-UX, you can have two separate groups. One for each of the following:
 - ◆ NetBackup-Java administrator applications, including administrator capabilities in the Backup, Archive, and Restore application.
 - ◆ Administrator capabilities for only `jbpsa`
 - ◆ If the system is not Solaris or HP-UX, create a group that includes all the other NetBackup commands, such as `bpadm`.
3. On the UNIX system that you will specify in the login dialog box when starting the NetBackup-Java interface, execute
`/usr/opensv/netbackup/bin/nonroot_admin.`

You are now asked to provide the group names you created.

Rerun this script any time a patch is installed that replaces any file in `/usr/opensv/netbackup/bin/admincmd` or files `bpbackup`, `bplist` or `bprestore` in `/usr/opensv/netbackup/bin`.

4. Change the NetBackup-Java authorization file, `/usr/opensv/java/auth.conf`, to provide the desired capabilities for the affected users (this file does not exist by default on UNIX master servers that are not Solaris or HP-UX so you must create it first on those systems).

For details, refer to “Authorizing Nonroot Users for Specific Applications” on page 258 and “Capabilities Authorization for jbpSA” on page 259.

5. Ask all affected users on the system where you ran the `nonroot_admin` script to restart the NetBackup-Java application.

A nonroot user that is not authorized for some of the applications per the `auth.conf` file, sees the following warning message dialog after logging in:

```
You are not authorized to use some of the applications.
Access to those applications has been disabled.
```

A nonroot user that logs in when step 2 and step 3 are not completed, (that is, the permissions on executables in the respective directories are not changed), sees a warning message dialog containing messages such as the following:

```
Administrator applications are not accessible due to the following error:
/bin/sh: /usr/opensv/netbackup/bin/admincmd/bpclvalid: cannot execute
```

This disables all NetBackup-Java administrator applications.

Creating a Group Specifically for Media Manager Tape Users

To create a group specifically for users that must use the `tpreq` and `tpunmount` commands to mount and unmount media, but do not need other administrative privileges, proceed as explained in the following:

The `tpreq` and `tpunmount` commands are released with permission modes set to 500 (`-r-x-----`).

- ◆ To allow any user in the system to mount tapes, the mode can be changed to 555 (`-r-xr-xr-x`).
- ◆ To allow a specific group of users to mount tapes, perform commands similar to the following:

```
cd /usr/opensv/volmgr/bin
chgrp tapeusers tpreq tpunmount
chmod 550 tpreq tpunmount
```



- ◆ Using the Media and Device Management utility, also change the volume pool configuration to specify ANY user and the group in the Group Name field.

If a group was created with administrative privileges as explained in step 2 in the previous topic, that group must be a member of the same group used for the `tpreq` and `tpunmount` commands.

Authorizing NetBackup-Java Users

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 3.2 or later Solaris and HP-UX 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.

As is explained later in this section, you can validate nonroot users to administer NetBackup and can also validate users for specific capabilities of the NetBackup Java applications.

Note To administer a NetBackup Windows master server with NetBackup-Java, set up user authorization as explained in the “Using NetBackup-Java” appendix in the NetBackup system administrator’s guide for Windows NT/2000.

Authorization File

The released version of the `/usr/opensv/java/auth.conf` file that is installed on Solaris and HP-UX clients and servers 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.

- ◆ 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.”
 - ◆ The `JBP` keyword specifies what the user can do with the Backup, Archive, and Restore 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 259.
 - ◆ 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 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/.nbjava 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/.nbjava directory the first time an application is exited. Only NetBackup-Java applications use the .nbjava 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 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

All NetBackup-Java Applications

To authorize nonroot usage of all NetBackup-Java applications, see “Configuring Nonroot Usage of All NetBackup Administrator Applications” on page 253.

Authorizing Nonroot Users for Specific Applications

It is possible to authorize nonroot users for a subset of the NetBackup-Java administrator applications. This done as part of allowing nonroot usage of the NetBackup-Java administrator applications (see step 4 in “Configuring Nonroot Usage of All NetBackup Administrator Applications” on page 253).



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:

AM - Activity Monitor
BPM - Backup Policy Management
JBP - Backup, Archive, and Restore
DM - Device Monitor
MM - Media Management
REP - Reports
SUM - Storage Unit 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
```

Capabilities Authorization for `jbpSA`

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` - only authorized for restore capabilities; from true image, archive or regular backups plus alternate client restores
- ◆ `BU` - allowed to perform backup tasks
- ◆ `ARC` - allowed to perform archive tasks (`BU` capability required for this)
- ◆ `RAWPART` - allowed to perform raw partition restores
- ◆ `ALL` - allowed for all of the above including restoring to a different client from the one you are logging into (that is, server directed restores). This normally requires execution from the root account or an account set up for nonroot administration. Alternate client restores, also require the NetBackup configuration changes explained in “Allowing Restores to an Alternate Client” on page 211.

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`.

The following example entry allows a user named `bill` to restore but not back up or archive files:

```
bill ADMIN=JBP JBP=ENDUSER
```



Converting Previous auth.conf Files

A conversion script, `/usr/opensv/java/convert_auth_file.sh`, is available to convert pre-NetBackup 3.2 `auth.conf` files to be usable in 3.2 or later configurations. The `JBP=ALL` part of each nonroot user's entry in the pre-NetBackup 3.2 version 3.1.1J2 `auth.conf` files must be converted. You can choose either of two options:

- ◆ Provide all end user features (back up, archive and restore) by converting `JBP=ALL` to `JBP=ENDUSER+BU+ARC`.
- ◆ Provide only restore capabilities to end users by converting `JBP=ALL` to `JBP=ENDUSER`.

This script also removes the password field that existed in pre 3.2 `auth.conf` files.

Configuration Options for jbpSA

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 325.

Default NetBackup Host to Log Into

To change the default NetBackup host name that appears in the login dialog when starting the NetBackup-Java administrator interface, edit the following line in the `/usr/opensv/java/nbj.conf` file.

```
SERVER_HOST=eagle
```

For example, change host `eagle` to `hawk`.

```
SERVER_HOST=hawk
```

Similarly, for the end user Backup, Archive and Restore application (`jbpSA`), change the following option in this file.

CLIENT_HOST

Changing Fonts and Font Sizes on UNIX Platforms

The fonts in the NetBackup-Java applications are commonly available on a large number of UNIX platforms. The Backup, Archive, and Restore application also has good resolution and allows the application to be rendered on an average size screen.

However, if the default fonts are not available on your machine(s) you must edit the `font.properties` file. This file is in `/usr/opensv/java/jre/lib`. In this file, you must change the entries for the “Dialog” font, which are found after the “Dialog font definition” comment.

To find the fonts that are available for your machine, execute `/usr/bin/X11/xlsfonts`. This lists the font styles and sizes. You may have to try several fonts to find one you like.

As an example, the initial `font.properties` file (see the fragment below) uses the `-b&h-lucida-medium` font.

```
.
.
# Dialog font definition
#
dialog.plain.0=-b&h-lucida-medium-r-normal-sans-*-%d-75-*-%p-*-%iso8859-1
dialog.1=-urw-its zapfdingbats-medium-r-normal--*-%d-*-%p-*-%sun-fontspecific
dialog.2=-symbol-medium-r-normal--*-%d-*-%p-*-%sun-fontspecific

dialog.italic.0=-b&h-lucida-medium-i-normal-sans-*-%d-75-*-%p-*-%iso8859-1

dialog.bold.0=-b&h-lucida-bold-r-normal-sans-*-%d-75-*-%p-*-%iso8859-1

dialog.bolditalic.0=-b&h-lucida-bold-i-normal-sans-*-%d-75-*-%p-*-%iso8859-1
.
.
```

If you run `jbp` and don't see any text, execute `xlsfonts`. You will discover that the `-b&h-lucida-medium` fonts is not available on you machine, but you do have the `-adobe-courier-medium` fonts. You can then change the `font.properties` file to use it as follows.

```
.
.
```



```
# Dialog font definition
#
dialog.plain.0=-adobe-courier-medium-r-normal--*-%d-75-*-*m-*-*iso8859-1
dialog.1=-urw-its zapfdingbats-medium-r-normal--*-%d-*-*p-*-*sun-fontspecific
dialog.2=--symbol-medium-r-normal--*-%d-*-*p-*-*sun-fontspecific

dialog.italic.0=-adobe-courier-medium-o-normal--*-%d-75-*-*m-*-*iso8859-1

dialog.bold.0=-adobe-courier-bold-r-normal--*-%d-75-*-*m-*-*iso8859-1

dialog.bolditalic.0=-adobe-courier-bold-o-normal--*-%d-75-*-*m-*-*iso8859-1
.
.
```

For information on fonts, see *X Window System User's Guide* from O'Reilly and Associates, Inc.

You can also change the font size used for the Backup, Archive, and Restore application (jbpSA). To change this setting, add the line

```
FontSize=font_size
```

to the `JBPSimple.properties` and `Launch.properties` files in `/usr/opensv/java`. Where *font_size* specifies the size of the font. Verify that the size you choose is available for the font set in the `font.properties` file.

If you change the font size, it can also be necessary to change the `WIDTH` and `HEIGHT` values in `JBPSimple.properties` so the application is sized to fit the new font.

Example

Assume you decide to use a larger font for jbpSA, for instance size 12 instead of the default size 10 and add

```
FontSize=12
```

to `JBPSimple.properties` (and to `Launch.properties`).

Also, assume the `WIDTH` and `HEIGHT` options in your original `JBPSimple.properties` file are:

```
WIDTH=760
```

```
HEIGHT=580
```



In this example, changing the font size makes it necessary to change the size of the window. So you set `WIDTH` to 820 and `HEIGHT` to 600 in the `JBPSimple.properties` file:

```
WIDTH=820
HEIGHT=600
```

NetBackup-Java Performance Improvement Hints

Performance of the NetBackup-Java applications depends on the environment where they are running. Following are guidelines for improving performance:

- ◆ Run NetBackup-Java on a 256 MB machine that has at least 128 MB of RAM available to the application. In some instances, the application does not even initiate due to insufficient memory. These failures can be indicated by an `Abort - core dumped` message. See “Insufficient Memory” on page 264.
- ◆ Run NetBackup-Java on a machine that has a low level of activity. For example, there can be dramatic differences in response time when a copy of Netscape Navigator is running on the machine. Multiple instances of NetBackup-Java on the same machine have the same effect.
- ◆ Since startup of the Java virtual machine and some of the applications can take a while, leaving NetBackup-Java running (iconified) rather than exiting and restarting is beneficial.
- ◆ Increasing the `-ms` argument to the Java runtime environment (`jre`) command in the `jbpsa` or `jnbSA` scripts can be helpful. This argument specifies how much memory is allocated for the heap when the Java interpreter starts. The value used in these scripts is `4m` - the `m` suffix refers to megabytes.
- ◆ Increasing the `-mx` argument to the Java runtime environment (`jre`) command in the `jbpsa` or `jnbSA` scripts can be helpful. This argument specifies the maximum heap size the runtime environment 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). See “Insufficient Memory” on page 264.
- ◆ Increasing the amount of swap space available to the system where you are running the applications can increase performance, especially if there is a lot of other activity on the machine. This can also alleviate hangs or other problems related to insufficient memory for the applications. See “Insufficient Memory” on page 264.
- ◆ If sufficient machine resources (see previous items) on the administrative server (one where you logged in when starting the interface) cannot be made available, run the interface on another machine.



Insufficient Memory

If the following error appears in the log file, it indicates that insufficient memory is available:

```
java.lang.OutOfMemoryError
```

Upon receiving this error, all applications are probably hung so you have to restart by using the `jnbSA` command.

An attention dialog box with an `Out of memory error exception occurred message` also indicates an insufficient memory problem.

In the above instances, increase the value of the `-mx` argument to the `jre` command in the scripts in `/usr/opensv/netbackup/bin`. The value in these scripts is `32m` (32 megabytes).

Increasing the amount of swap space or running on a machine that has more memory available or less activity can also provide better results.

Note that insufficient memory for NetBackup-Java applications can also result in a `SIGSEGV 11` segmentation violation in the output file (that is, `Abort - core dumped` in the xterm window where you started the application). We have seen this occur on a machine with 64MB of memory and more than one Netscape Navigator processes running. Shutting down Navigator allowed NetBackup-Java to initiate and run.

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 shown in this figure.

Command	Description
Administrator Utilities	
<code>xnb</code>	Opens the NetBackup Administration window on the server. The other administration interfaces can be started from this window.
<code>xbpadm</code>	Starts X windows administrator's interface on the server.
<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>xbp</code>	Starts X windows client-user interface.

Command	Description
<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>bprdregr -terminate</code>	Stops <code>bprd</code> (request daemon)
<code>initbpdbm</code>	Starts <code>bpdbm</code> (database manager).
<code>xbpadm</code>	Has option for starting and stopping <code>bprd</code> .
<code>bpadm</code>	Has option for starting and stopping <code>bprd</code> .
Monitor Processes	
<code>bpps</code>	Lists active NetBackup processes.

File	Description
<code>/usr/opensv/netbackup/bp.conf</code>	Configuration options (server and client).
<code>\$HOME/bp.conf</code>	Configuration options for user (on client).





This chapter explains settings that in many instances are optional either because its default setting is appropriate or a site does not use the feature. The topics included here are:

- ◆ Defining NetBackup Global Attributes
- ◆ Redefining Retention Periods
- ◆ Multiplexing
- ◆ Using Multiple NetBackup Servers
- ◆ Dynamic Host Name and IP Addressing
- ◆ Bandwidth Limiting
- ◆ Busy-File Processing (UNIX Clients Only)
- ◆ Open Transaction Management (Microsoft Windows Clients Only)
- ◆ Enhanced Authentication
- ◆ NetBackup Authorization
- ◆ Configuring Email Notifications
- ◆ Adjust Time Zone
- ◆ Specifying the Locale of the NetBackup installation
- ◆ NetBackup Configuration Options



Defining NetBackup Global Attributes

NetBackup global attributes affect all operations for all classes and clients. The default values are adequate for most installations but you can change them if necessary.

1. In the NetBackup Administration window, click NetBackup System Configuration on the Configure menu.

The NetBackup System Configuration dialog box appears.

System Configuration - NetBackup

VERITAS NetBackup

File Help

Master server: lettuce

Global Attributes Retention Periods

E-mail address for notifications:
gla@min.ov.com

Maximum jobs per client:

Media mount timeout: minutes (0 = no timeout)

Wakeup interval: minutes

Interval for status reports: hours

Schedule backup attempts: tries per hours

Delay to compress database: days (0 = do not compress)

Duration to retain logs: days

How long to keep TIR information: days

2. If there is more than one NetBackup master server, check the Master Server line below the menu bar and ensure that it names the master server where you want to change the configuration. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.
 - b. In the dialog box, specify the master server.

Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), click the browse button to the right of the Select New Master Server box and select from the list. The list shows the master servers that have a `KNOWN_MASTER` entry in the `bp.conf` file on the Master of Masters system.

- c. Click OK.

The Master Server line below the menu bar now shows the master server that you selected. If you cannot change to the specified server, refer to “If You Cannot Change to Another Master Server” on page 23.

3. Complete the entries on the Global Attributes tab.

See “Defining NetBackup Global Attributes” on page 268 for an explanation of the global attributes.

4. On the File menu:
 - a. Click Save to save the changes.
 - b. Click Exit to close the dialog box (Exit does not perform a save).

NetBackup Global Attributes

E-mail address for notifications

Specifies the address 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. Specify the address of the NetBackup administrator. The default is no address.

On a Windows NT/2000 NetBackup server, it might 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 NT and 2000, 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. By default, the `nbmail.cmd` script does not send E-mail.

Maximum jobs per client

Specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. The default is 1.


NetBackup can process concurrent backup jobs from different classes on the same client only if:

- ◆ There is more than one storage unit available.



or

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



Files and directories that are on the same client but in different classes, can be backed up concurrently to different storage devices.

You can specify any number of concurrent jobs within the following constraints (the default is 1):

- ◆ 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 classes, 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.

Wakeup interval

Specifies how often in minutes 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. The minimum setting is 1 minute. The default is 10 minutes.

Schedule backup attempts

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

Specifies the number of times that NetBackup tries to complete a scheduled backup job during the specified time period. This lets you 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 and must occur within the backup window. If the backup window closes before the retry starts, the job fails with a status code 196.

The number of tries must be greater than 0 in order for scheduled backups to occur. Specifying 0 for number of tries is legal but stops all scheduled backups.

The time period must *always* be greater than 0. The default is 2 tries in 12 hours.

Duration to retain logs

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

Keep the logs as long as you need them to evaluate failures. For example, if you check the backups every day you can delete the logs sooner than if you check them once a month. However, the logs can consume a lot of disk space so do not keep them any longer than necessary.

The minimum setting is 0 days. The default is 28 days.

Media mount timeout

Specifies the number of minutes that NetBackup waits for the requested media to be mounted, positioned, and ready on backups and restores. The default setting is 0 (no timeout). If you do not specify 0, the value must be 5 minutes or greater.

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.

Interval for status reports

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.

The minimum setting is 1 hour. The default is 24 hours.

Delay to compress catalog

Specifies the number of days that NetBackup waits after a backup before compressing the image catalog file that has information about the backup. On Windows 2000 and NT, NetBackup uses NTFS file compression and the catalog must be in an NTFS partition for compression to occur.

The default is 0, which turns off compression and keeps all image catalog files in uncompressed format.

How long to keep TIR information

Specifies the number of days to keep true image restore information on disk. This applies to all classes for which NetBackup is collecting true image restore information. For more information, see “True image restore information” on page 70. The default is one day.

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

- ◆ Backed up files
- ◆ True image restore information

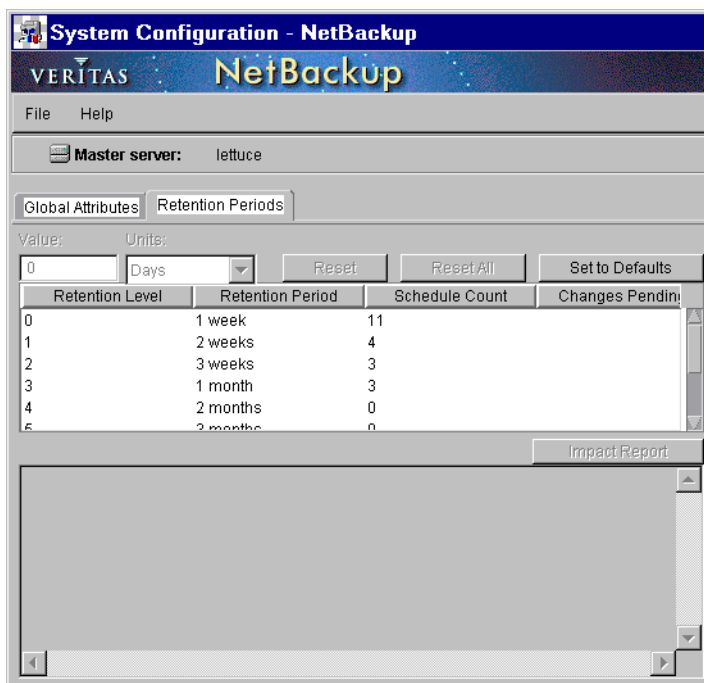
NetBackup also stores the true image restore information on disk in the `/usr/opensv/netbackup/db/images` directory and keeps it for the number of days specified by this global attribute. Keeping the information on disk speeds up restores. If a user requests a true image restore after the information 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.

Redefining Retention Periods

When setting up a schedule, you select a retention period that determines how long NetBackup retains the backups or archives that it creates according to that schedule. There are 10 possible levels from which to select and you can change the period associated with each level as explained below.

1. In the NetBackup Administration window, click NetBackup System Configuration on the Configure menu.

The NetBackup System Configuration dialog box appears.



2. If there is more than one NetBackup master server, check the Master Server line below the menu bar and ensure that it names the master server where you want to change the configuration. To change master servers:
 - a. On the File menu, click Change Master Server.
A dialog box appears.



Note If you are administering NetBackup from a Master of Masters (requires the Global Data Manager option), click the down arrow on the Select New Master Server button in the dialog box and select from the list. The list shows the master servers that have a `KNOWN_MASTER` entry in the `bp.conf` file on the Master of Masters system.

- b. In the dialog box, specify the master server.
- c. Click OK.

The Master Server line below the menu bar now shows the master server that you selected. If you cannot change to the specified server, refer to “If You Cannot Change to Another Master Server” on page 23.

3. Click the Retention Periods tab.

The list in the upper half of the dialog box shows the current definition for the 10 possible levels of retention (0 through 9). The Schedule Count column indicates how many schedules are using each level. If you change the retention period for a level, it affects all schedules that use that level.

4. Select the retention level that you want to change. You cannot change the period for level 9; it will always be infinite.

The list box in lower half of the dialog box now shows the names of all schedules that are using the selected retention level and also the class to which each schedule belongs.

5. Click Units and then select the units for the new retention period.
6. In the Value box, type the new retention period.

Retention periods can range from 0 (no retention) to 30 years.

Note When you change either Units or Value, an asterisk (*) appears in the Changes Pending column to indicate that the period was changed. NetBackup does not change the actual configuration until you click Save.

7. Click Impact Report.

The lower half of the dialog box lists 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.

8. To discard your changes, click one of the following:

- ◆ Reset - If you select a retention period that was changed (has an asterisk in the Modified column), clicking Reset restores the selection to the value that was set with the last Save.
 - ◆ Reset All - Restores all retention periods that were changed (have an asterisk in their Modified column), to the values that were set with the last Save.
 - ◆ Set to Defaults - Sets all levels of retention to their standard defaults.
 - ◆ Exit on the file menu: This discards all changes since the last Save and closes the dialog box.
9. To save your changes, click Save on the File menu. Then, click Exit to close the dialog box.

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 associated with 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 -ev media_ID
```

Multiplexing

NetBackup multiplexing sends concurrent backups from one or several clients to a single storage device (see figure below). NetBackup multiplexes the backups sequentially onto the media. Multiplexed and unmultiplexed backups can reside on the same volume. It is not necessary to create separate volume pools or media IDs.



No special action is required to restore a multiplexed backup. NetBackup finds the media and restores the requested backup.



When to Use Multiplexing

Multiplexing is generally used to reduce the amount of time required to complete backups. The following are situations where multiplexing can improve backup performance.

- ◆ Slow clients. This includes instances where NetBackup is using software compression, which normally reduces client performance.
- ◆ Multiple slow networks. The parallel data streams take advantage of whatever network capacity is available
- ◆ Many short backups (for example, incrementals). In addition to providing parallel data streams, multiplexing reduces the time each job spends waiting for a device to become available, and therefore better utilizes the transfer rate of storage devices.

However, when you use multiplexing, expect reduced performance on restores due to the extra time spent reading the images.

Note To reduce the impact of multiplexing on restore times, set maximum fragment size for the storage units to a value other than zero. Also, enable fast-tape positioning (locate block), if it applies to the tape drives you are using.

How to Configure Multiplexing

Multiplexing must be set in two places in the NetBackup configuration:

- ◆ Storage unit
- ◆ Schedule

Note If you change these values, it does not take effect until the next time a schedule runs.

Maximum Multiplexing Per Drive for Storage Unit

The Maximum multiplexing per drive setting for a storage unit specifies how many backups NetBackup can multiplex onto any single drive in the storage unit. You set this value for each storage unit (see Chapter 2). The number can range from 1 through 32, where 1 is the default and specifies no multiplexing.

Choose a value based on the ability of your CPU to handle parallel jobs. Because extra buffers are required, memory is also important. If the server cannot perform other tasks or runs out of memory or processes, reduce the Maximum multiplexing per drive setting for the storage unit. Consider the following when estimating the load that multiplexing can potentially put on your CPU:

- ◆ The maximum number of concurrent-backup jobs that NetBackup is allowed to attempt, equals the sum of the concurrent-backup jobs that can run on the storage units.
- ◆ The maximum number of concurrent-backup jobs that can run on a single storage unit equals:

Maximum Multiplexing per drive x number of drives

Media Multiplexing for a Schedule

In addition to the Maximum multiplexing per drive setting for a storage unit, you specify a Media multiplexing value for each schedule (see “Media multiplexing” on page 129). This setting specifies the maximum number of backups from the schedule that you can multiplex onto any single drive in the configuration.

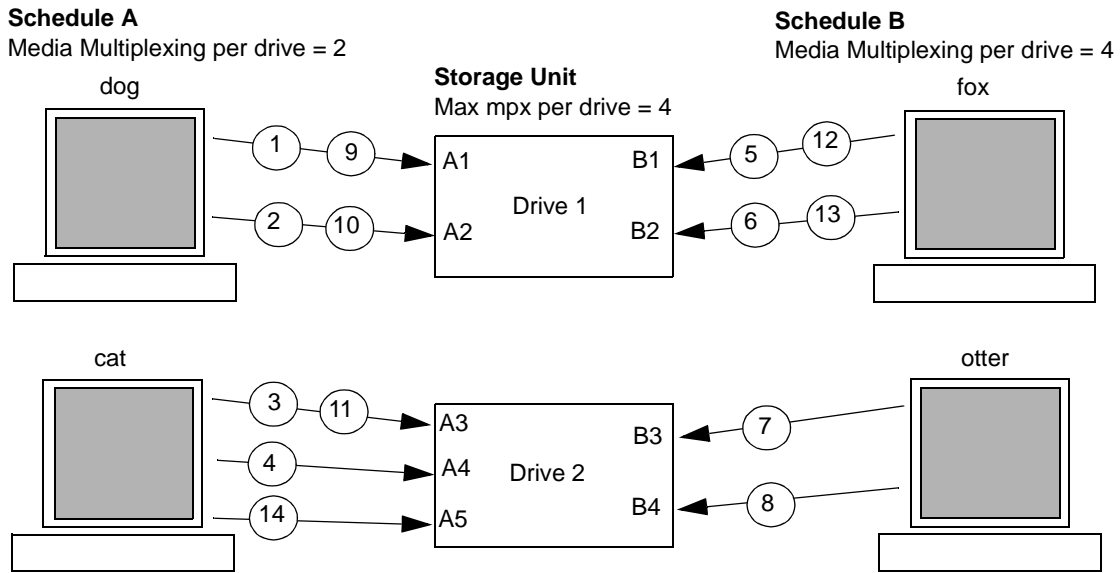
The Media multiplexing setting can range from 1 through 32, where 1 is the default and specifies no multiplexing. Regardless of the setting on a schedule, the maximum jobs that NetBackup starts never exceeds the storage unit’s Maximum multiplexing per drive.

When NetBackup multiplexes jobs, it continues to add jobs to a drive until the number of jobs on the drive matches either of the following:

- ◆ This schedule’s Media multiplexing setting.
If the limit is reached for a drive, NetBackup starts sending jobs to another drive. In the following figure, when the Schedule A limit is reached on Drive 1, NetBackup starts adding Schedule A jobs to Drive 2.
- ◆ The storage unit’s Maximum multiplexing per drive.



NetBackup can add jobs from more than one schedule to a drive. In the figure below, NetBackup adds jobs from both schedules to Drive 1 until the total on that drive reaches four.



Assume schedule A begins first (note that the schedules can be in the same or different classes). Also, assume that Allow Multiple Data Streams is enabled, so a client can have multiple data streams.

- ① ② Jobs A1 and A2 from client dog start on drive 1. Schedule A Media Multiplexing limit of 2 is reached for this drive.
- ③ ④ Jobs A3 and A4 from client cat start on drive 2. Schedule A Media Multiplexing limit of 2 is reached for this drive.
- ⑤ ⑥ Jobs B1 and B2 for client fox start on drive 1. Storage unit max mpx is reached for this drive.
- ⑦ ⑧ Jobs B3 and B4 from client otter start on drive 2. All jobs are now running for schedule B. Storage Unit Max mpx is reached for drive 2.
- ⑨ ⑩ Jobs A1 and A2 from client dog finish on drive 1. However, jobs B1 and B2 for client fox are still running so Schedule A Media Multiplexing limit of 2 still prevents job A5 from starting on drive 2.
- ⑪ ⑫ ⑬ Job A3 from client cat finishes on drive 2 and jobs B1 and B2 for client fox finish on drive 1. No jobs are currently running on drive 1.
- ⑭ Job A5 from client cat starts on drive 2. This is the last job for schedule A. When adding jobs to drives, NetBackup attempts to add multiplex jobs to drives that are already using multiplexing. This leaves other drives available for non-multiplex jobs.

When adding jobs to drives, NetBackup attempts to add multiplex jobs to drives that are already using multiplexing. This leaves other drives available for non-multiplex jobs.

Note If the backup window closes before NetBackup can start all the jobs in a multiplexing set, NetBackup completes only the jobs that have actually started. For example, on the figure above, assume that the NetBackup Activity Monitor shows A1 through A5 as queued and active. If only A1 and A2 actually start before the window closes, NetBackup does not perform the other jobs that are in the set. If the window closes before any jobs have started, then only the first queued and active job starts and completes (A1 in this example).

Other Configuration Settings to Consider

Limit Jobs per Class

Set Limit jobs per class high enough to support the specified level of multiplexing. See “Limit jobs per class” on page 68.

Maximum Jobs per Client

The Maximum jobs per client global attribute limits the number of backup jobs that can run concurrently on any NetBackup client. Usually, its setting does not affect multiplexing. However, to illustrate its effect, consider a case where there are jobs from different schedules on the same client and all are going to the same storage unit. In this case, it is possible for the maximum number of jobs permitted on the client to be reached before the multiplexing limit is reached for the storage unit. If this occurs, it prevents NetBackup from fully utilizing the storage unit’s multiplexing capabilities.

Maximum Jobs this Client

You can also set the maximum number of jobs that are allowed on a specific client without affecting other clients. You can set this with the `bpconfig` command. The instructions for setting this value that is given under “Number of Streams That Can Run Concurrently” on page 84 also applies here (except that concurrent data streams do not necessarily apply).

MPX Restore Delay

The NetBackup configuration option, `MPX_RESTORE_DELAY`, applies to multiplexed restores and specifies how long (in 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.



Demultiplexing

Demultiplexing speeds up future restores and is also useful for creating a copy for offsite storage.

To demultiplex a backup, select Duplicate from the Images menu in the NetBackup Administration window. This command lets you copy one multiplexed backup at a time from the source media to the target media. When duplication is complete, the target contains a single demultiplexed copy of each backup you duplicated (the target can also have other backups). If desired, you can make the duplicate copy the primary copy.

Note If you use the `bpduplicate` command instead of the administration interface, do not include the `-mpx` option on that command.

Example

Assume you multiplexed clients A, B, and C to media ID MPX001. This requires three separate duplicate operations. In each of them, you limit the selection of backups to a specific client and media ID. The general procedure is as follows. See “Duplicating Backup Images” on page 225 for detailed instructions.

1. Click Duplicate on the Images menu and select the storage unit and volume pool.
2. In the Search Criteria section of the Duplicating Images dialog box, select client A and media ID MPX001.

Ensure that the date and time range covers the period of the multiplexed backup.

3. Click Search.

NetBackup lists the backups that were created for client A on MPX001.

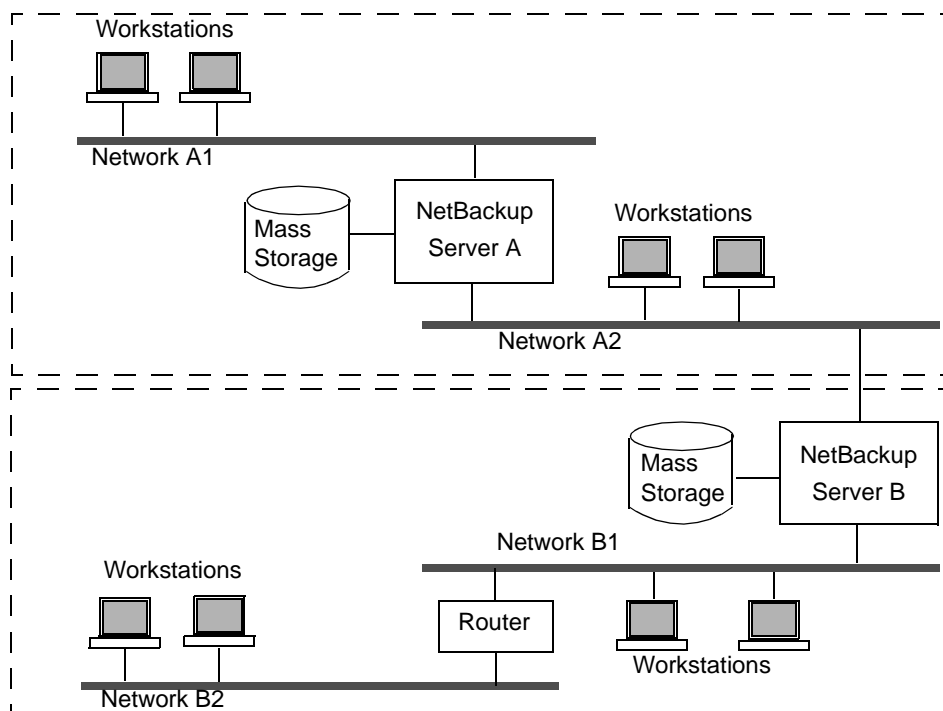
4. Click Select All to select all the backups
5. Clear the MPX box if it is checked. Otherwise, the selected backups are duplicated as multiplexed backups and no demultiplexing occurs.
6. Click Duplicate and check the progress log for results.
7. Repeat these steps for the clients B and C.

Using Multiple NetBackup Servers

A large site that has more than one master server can divide the clients between the servers as necessary to optimize the backup loads. The figure below shows a multiple-server configuration where the two sets of networks (A1/A2 and B1/B2) each



have enough clients to justify separate servers. In this environment, the two NetBackup server configurations are completely independent. You can also create a configuration where one server is the master and the other is a media server (see the next topic).



Configuring a Master and Media Server Cluster

NetBackup lets you set up a cluster of NetBackup servers where one server is the master and the others are used only as media servers and have peripherals to provide additional storage. The master server controls all backup scheduling and the other media servers provide additional storage.

The term cluster refers collectively to the master and its media servers. In a cluster of NetBackup servers, a client can have its backup directed to any device on any server in the cluster.

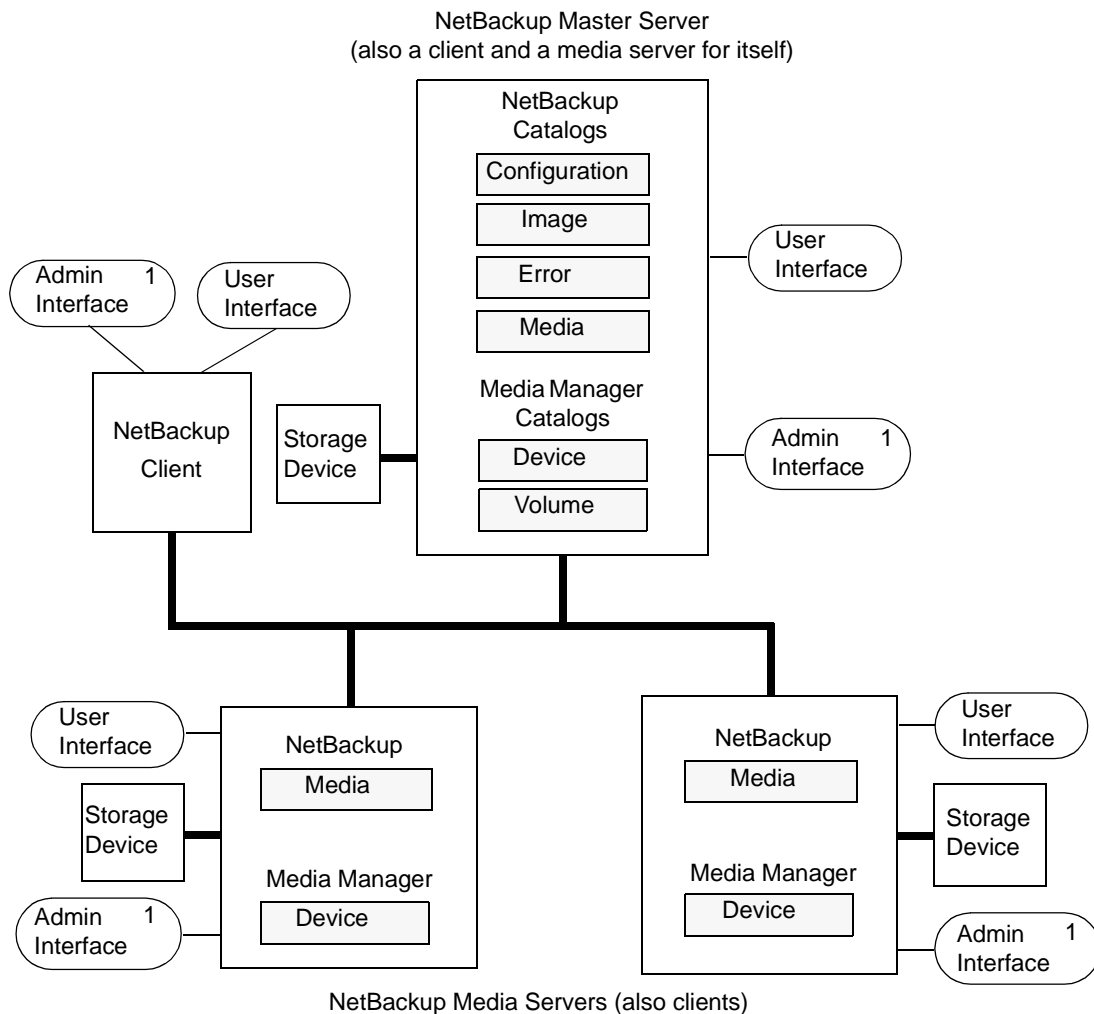
A common strategy is to install extra peripherals on clients that produce large amounts of data and make them media servers. The data from the client is then directed to its own peripherals. This reduces network traffic by allowing the data to be backed up without transferring it over the network. It also distributes the backup load between the master and the media servers.

Two important points to remember about master and media servers:



- ◆ There can be only one master server in a cluster.
- ◆ A NetBackup server is a media server for itself but cannot be a media server for another master.

The following figure shows where the software is installed and also where the NetBackup catalogs are located. The following topics provide more details on master and media servers along with a procedure to configure them.



NOTES

1. You can also use the Windows NT NetBackup administration interface from a Windows NT administration client (a Windows NT client that has the administration interface program installed).

Software on Each Server

You install NetBackup server software on each NetBackup server that has a peripheral that you want to include in a storage unit. The NetBackup install program has choices for master and media server installation.

NetBackup Catalogs

The master server has a complete set of NetBackup and Media Manager catalogs. Additional media servers have only two NetBackup catalogs:

- ◆ Media - Contains information on media used for backups on the drives attached to the media server.
- ◆ Device - Contains information on drives and robots configured on the media server.

Special Note on Configuring Volumes

To simplify administration, it is recommended that you use only one server as a volume database host and add all your volumes on that host. Multiple configurations make administration more complicated and it is not possible to merge the volume information later.

Note Regardless of which server is the volume database host, the one that has the drive always requests the media. If the media isn't available, the mount request shows up on that server.

NetBackup and Media Manager Daemons

The NetBackup database manager (bpdbm) is used only on master servers.

On Windows NT/2000 servers, the equivalents to the NetBackup and Media Manager daemons are as follows:

- ◆ NetBackup Device Manager service = Media Manager device daemon (ltid)
- ◆ NetBackup Client service = NetBackup client daemon (bpcd)
- ◆ NetBackup Request Manager service = NetBackup request daemon (bprd)
- ◆ NetBackup Volume Manager service = Media Manager volume daemon (vmd)
- ◆ NetBackup Database Manager service = NetBackup database manager (bpdbm)



To Add Media Servers

Note For information on configuring a NetBackup Windows NT/2000 server, see the Windows NT/2000 version of the NetBackup system administrator's guide.

1. Install the following software packages on the media server as explained in the vendor's documentation:
 - ◆ Any software required to drive the storage devices. This refers to software that the storage-device vendor provides.
 - ◆ NetBackup server software as explained in the NetBackup Getting Started Guide - *UNIX*.

If you decide to make a UNIX media server a client, install the client software from the master server, not from the distribution media. When the installation script asks you if the host is the master server, reply no and enter the name of the master server when prompted for it.
2. On a Windows NT/2000 media server, configure the drives and robots as explained in the system administrator's guide for NetBackup Media Manager (Windows NT/2000 version).
3. On a UNIX media server, configure the drives and robots per the following steps.
 - a. Log in to the media server.
 - b. Create the necessary device files, if this was not done at the time the devices were installed. If necessary, refer to the documentation for the device and your host system. The NetBackup Media Manager Device Configuration Guide has advice on creating device files.
 - c. Configure the robots and drives within Media Manager as instructed in the system administrator's guide for NetBackup Media Manager (UNIX version).

Remember, the server that you specify as the Volume Database Host is the one that will keep records of volumes used in this device. The Volume Database Host can be any of the following:

 - ◆ Master server
 - ◆ media server that you are currently adding
 - ◆ Another media server

Note Use only one server as a volume database host and add all your volumes on that host. Multiple volume database hosts complicate administration and it is not possible to merge the volume information later.

4. Add the volumes for each robot or nonrobotic drive configured in step 2 or step 3.



Always add the volumes on the server that you specified as the volume database host for the devices in step 2 or step 3. See the *Media Manager System Administrator's Guide - UNIX* for instructions on adding volumes.

Note Defining a separate volume pool for volumes used on the media server can simplify administration.

5. On the master server, make the following changes to the NetBackup configuration:

a. Add the media server's storage units.

Remember, when adding the storage units, always specify the media server as the media server for the storage unit.

b. Add the catalog paths for the media server to the NetBackup catalog backup configuration. For instructions, see Chapter 4.

For a Windows NT/2000 media server, the paths are:

```
media_server_name: install_path\NetBackup\db
```

```
media_server_name: install_path\Volmgr\database
```

Where *install_path* is the directory where the NetBackup software is installed on the media server.

For a UNIX media server, the paths are:

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

```
media_server_name: /usr/opensv/volmgr/database
```

c. Configure the NetBackup classes (policies) and schedules that use the storage units you configured on the media server.

6. Execute the `/usr/opensv/netbackup/bin/add_slave_on_clients` script from the master server to add the name of the new media server to the `/usr/opensv/netbackup/bp.conf` file on each UNIX client.

If the update fails on any client, either:

◆ Fix the problem and rerun `add_slave_on_clients`.

or

◆ Edit the `/usr/opensv/netbackup/bp.conf` files on the clients that fail by adding the following below the entry for the master server (which is always the first `SERVER` entry):

```
SERVER = media_server_name
```

where *media_server_name* is the host name of the media server.



Note You will encounter problems with NetBackup, unless the host names in the `bp.conf` file match those shown in the `/etc/hosts` file (or appropriate NIS, or DNS file).

The host names must also match throughout your network. If you are using NIS, then this applies to the NIS hosts file. See “Rules for Using Host Names in NetBackup” on page 328, for more information on choosing host names for NetBackup hosts and clients.

In addition, the `SERVER` entries **MUST** be the same on all servers in a master and media server cluster. It is recommended (but not mandatory) that all other `bp.conf` entries, except `CLIENT_NAME`, also match on all servers.

7. On each PC client, add `SERVER` entries for the new media server.

- ◆ For most PC clients, you can add these entries through the user interface on the client. For instructions on a specific client, see the NetBackup user’s guide for that client.
- ◆ The first `SERVER` entry in each case must be that of the master server. An entry for the media server must appear lower in the list:

```
SERVER = master_name
```

```
SERVER = media_server_name
```

8. On the master server, stop and restart the NetBackup request daemon (`bprd`) and database daemon (`bpdbm`).

a. To stop `bprd`, execute:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate.
```

b. To stop `bpdbm` execute:

```
/usr/opensv/netbackup/bin/bpdbm -terminate
```

c. To restart both daemons, execute:

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

9. Test your configuration by performing a user backup or a manual backup that uses a schedule specifying a storage unit on the media server.

Dynamic Host Name and IP Addressing

By default, a NetBackup server assumes that a NetBackup client name is the same as the network host name of the client machine. This makes it difficult to back up clients that have network host names that might change; examples of this are portable machines that



plug into a LAN and obtain IP addresses from a DHCP server or remote machines that dial into a PPP server. NetBackup dynamic host name and IP addressing allows you to define NetBackup clients that do not have fixed IP addresses and host names.

Note If you use dynamic addressing, remember that the NetBackup servers still require fixed IP addresses and host names.

Note All clients configured to use dynamic addressing and host names must trust each other in a way similar to that provided by the NetBackup altnames feature.

The following steps are required to support configurations that use dynamic IP addressing for NetBackup. Read all sections of this topic prior to making any changes to your configuration.

1. Configure your network to use a dynamic IP addressing protocol like DHCP.

NetBackup requires that IP addresses of clients have a network host name. Be sure to define network host names for the range of dynamic IP addresses in the `hosts` file, NIS, and (or) DNS on your network.

2. Determine the NetBackup client names for the machines that have dynamic IP addresses and network host names.

You will use these NetBackup client names in step 3 and step 6 of this procedure. Each NetBackup client must have a unique NetBackup client name. The NetBackup client name assigned to a client is permanent--do not change it.

3. Make changes on the master server:

- a. Create NetBackup classes with client lists that include the names from step 2.
- b. Create entries in the NetBackup client database for the client names from step 2.

Create the entries by using the `bpclient` command as explained in "Configuring the NetBackup Master Server" on page 289.

4. Make changes on each dynamic NetBackup Windows 2000, NT, 98, or 95 client:

- a. Start the user interface on the client and click **Configure** on the **Actions** menu. Then, on the **General** tab of the NetBackup Configuration dialog box, change the **Client Name** so it is the correct NetBackup client name for the machine.
- b. In the registry, modify the NetBackup configuration option, `Announce_DHCP_Interval`, so it contains a value other than 0. This option is in the following registry key on the client.

`HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\Config`

5. Make changes on each dynamic NetBackup Macintosh client:



- a. Modify the `bp.conf` file so it includes a `CLIENT_NAME` entry with the correct NetBackup client name for the machine.
 - b. Modify the `mac.conf` file so it includes a `DYNAMICNOTIFY` entry to periodically notify the NetBackup master server of the machine's NetBackup client name and current network host name.
6. Make changes on each dynamic NetBackup UNIX client:
- a. Modify the `bp.conf` file to include a `CLIENT_NAME` entry with the correct NetBackup client name for the machine.
 - b. Configure the system to notify the master server of the machine's NetBackup client name and current network host name during startup.

The `bpdynamicclient` command is used to notify the master server as explained in "Configuring a Dynamic UNIX NetBackup Client" on page 292.
 - c. Configure the system to periodically notify the master server of the machine's NetBackup client name and current network host name.

Setting up Dynamic IP Addresses and Host Names

Configure your network to use a dynamic IP addressing protocol. A protocol like DHCP will have a server and several clients. For example, when a DHCP client starts up, it requests an IP address from the DHCP server. The server then assigns an IP address to the client from a range of predefined addresses.

NetBackup requires that the IP addresses of NetBackup clients have corresponding network host names. Ensure that each IP address that could be assigned to NetBackup clients has a network host name defined in the `host` file, NIS, and (or) DNS on your network.

As an example, suppose that you have 10 dynamic IP addresses and host names available. The dynamic IP addresses and host names might be:

```
123.123.123.70 dynamic00
123.123.123.71 dynamic01
123.123.123.72 dynamic02
123.123.123.72 dynamic03
.
.
.
123.123.123.79 dynamic09
```

Assign a unique NetBackup client name to each NetBackup client that might use one of these dynamic IP addresses. The NetBackup client name assigned to a client is permanent and should not be changed. The client name assigned to NetBackup clients with dynamic IP addressing must not be the same as any network host names on your network. If the NetBackup client names are changed or are not unique, backup and restore results are unpredictable.

For example, suppose you have 20 machines that will share the IP addresses defined above. If you want these machines to be NetBackup clients, you might assign them these NetBackup client names:

```
nbclient01
nbclient02
nbclient03
nbclient04
.
.
.
nbclient20
```

Configuring the NetBackup Master Server

On the master server, create your NetBackup classes (backup policies) as you would otherwise. For client name lists, use the NetBackup client names (for example, nbclient01) rather than the dynamic network host names (for example, dynamic01).

Next, create the client database on the master server. The client database 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 dynamic client entry use:

```
bpclient -add -client client_name -dynamic_address 1
```

where *client_name* is the NetBackup client name. The `-dynamic_address 1` argument means that the client uses dynamic IP addressing. You can create entries with `-dynamic_address 0` for static IP addressing, but that is unnecessary and will adversely affect performance.

- ◆ To delete a client entry use:



```
bpclient -delete -client client_name
```

- ◆ To list a client entry use:

```
bpclient -L -client client_name
```

- ◆ To list all client entries use:

```
bpclient -L -All
```

In our example, you can enter these commands to create the 20 clients:

```
cd /usr/openv/netbackup/bin/admincmd
```

```
bpclient -add -client nbclient01 -dynamic_address 1
```

```
bpclient -add -client nbclient02 -dynamic_address 1
```

```
bpclient -add -client nbclient03 -dynamic_address 1
```

```
bpclient -add -client nbclient04 -dynamic_address 1
```

```
.
```

```
.
```

```
.
```

```
bpclient -add -client nbclient20 -dynamic_address 1
```

To see what's currently in the client database, execute:

```
/usr/openv/netbackup/bin/admincmd/bpclient -L -All
```

The output is similar to the following:

```
Client Name: nbclient01
```

```
Current Host:
```

```
Hostname: *NULL*
```

```
IP Address: 0.0.0.0
```

```
Connect on non-reserved port: no
```

```
Dynamic Address: yes
```

```
Client Name: nbclient02
```

```
Current Host:
```

```
Hostname: *NULL*
```

```
IP Address: 0.0.0.0
```

```
Connect on non-reserved port: no
```

```
Dynamic Address: yes
```

```
.
```



```
.  
. Client Name: nbclient20  
Current Host:  
Hostname: *NULL*  
IP Address: 0.0.0.0  
Connect on non-reserved port: no  
Dynamic Address: yes
```

After the NetBackup client notifies the NetBackup server of its NetBackup client name and network host name, the Current Host, Hostname, and IP Address fields will display the values for that NetBackup client.

Configuring a Dynamic Microsoft Windows Client

If it is not already installed, install NetBackup for Windows 2000, NT, 98, or 95.

Start the NetBackup user interface on the client and click Configure on the Actions menu. Then, on the General tab of the NetBackup Configuration dialog box, change the Client Name to specify the NetBackup client name for the Windows client. For example:

```
Client Name = nbclient06
```

On the client, set `ANNOUNCE_DHCP_INTERVAL` to specify how much time in seconds must elapse before the client will attempt to notify the server that it is using a different IP address. The notification is sent only if the client is using a different IP address than the last time it was checked.

On all NetBackup clients that are running Windows 2000, NT, 98, or 95, you can add this option to the following registry key on the client:

```
HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config
```

The server is not notified if the default value of 0 is used. For a DHCP client, a good value to use is one-half of the lease period.

On the client, stop and restart the NetBackup Client service to have the changes take effect.

Configuring a Dynamic Macintosh NetBackup Client

If not already installed, install NetBackup for Macintosh.



Edit the `bp.conf` document with a text editor like SimpleText. The `bp.conf` document is in the NetBackup folder in the Preferences folder in the System Folder of the startup disk. Use the `CLIENT_NAME` entry to specify the NetBackup client name for the Macintosh. For example:

```
CLIENT_NAME = nbclient02
```

Edit the `mac.conf` document with a text editor like SimpleText. The `mac.conf` document is in the NetBackup folder in the Preferences folder in the System Folder of the startup disk. Use the `DYNAMICNOTIFY` entry to specify how often (in seconds) to notify the NetBackup server of the NetBackup client name and network host name of the Macintosh. For example, to notify the server every hour:

```
dynamicnotify = 3600
```

The server is not notified if the default value of 0 is used. For a DHCP client, a good value to use is one-half of the lease period.

Restart the Macintosh.

Configuring a Dynamic UNIX NetBackup Client

If not already installed, install the NetBackup client software.

Edit the `/usr/opensv/netbackup/bp.conf` file. Use the `CLIENT_NAME` entry to specify the NetBackup client name for the machine. For example:

```
CLIENT_NAME = nbclient00
```

You must execute the `bpdynamicclient` command once when the system first starts up. `bpdynamicclient` notifies the NetBackup server of the machine's NetBackup client name and current network host name. The `bpdynamicclient` command is in the directory:

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

The format of the `bpdynamicclient` command is:

```
bpdynamicclient -last_successful_hostname file_name
```

When `bpdynamicclient` starts up, it checks for the existence of *file_name*. If *file_name* does exist, `bpdynamicclient` determines if the host name written in the file is the same as the current network host name of the machine. If the host names match, `bpdynamicclient` exits and does not connect to the master server. If the host names do not match, `bpdynamicclient` connects to the master server and informs the server of its NetBackup client name and host name. If `bpdynamicclient` successfully informs the server, `bpdynamicclient` writes the current network host name into *file_name*. If `bpdynamicclient` cannot inform the server, `bpdynamicclient` deletes *file_name*.

Most UNIX systems provide a facility to define startup scripts. For example, on a Solaris system, you can create a script in the `/etc/rc2.d` directory:

```
# cat > /etc/rc2.d/S99nbdynamicclient <<EOF
#! /bin/sh

rm /usr/opensv/netbackup/last_successful_hostname
/usr/opensv/netbackup/bin/bpdynamicclient -last_successful_hostname \
/usr/opensv/netbackup/last_successful_hostname
EOF
# chmod 544 /etc/rc2.d/S99nbdynamicclient
```

Ensure that the dynamic client startup script is called after the machine obtains its IP address.

You must also create a root `crontab` entry to periodically call the `bpdynamicclient` command. For example, the following entry (one line) calls `bpdynamicclient` at seven minutes after each hour:

```
7 * * * * /usr/opensv/netbackup/bin/bpdynamicclient -last_successful_hostname
/usr/opensv/netbackup/last_successful_hostname
```

If you are using DHCP, a good interval to use between calls to `bpdynamicclient` is one-half of the lease period.

Bandwidth Limiting

The bandwidth limiting feature lets you restrict the amount of network bandwidth consumed by one or more NetBackup clients on a network. The actual limiting occurs on the client side of the backup connection.

This feature only restricts bandwidth during backups. Restores are unaffected.

Read This First

- ◆ NetBackup does not currently support bandwidth limiting on the following clients:
 - ◆ Apollo wbak clients
 - ◆ Auspex-FastBackup clients
 - ◆ NetBackup for Oracle clients
 - ◆ NetBackup for DataTools SQL-BackTrack clients
 - ◆ NetBackup for Microsoft SQL-Server clients
- ◆ Bandwidth limiting has no effect on a local backup (where the server is also a client and data does not go over the network).



- ◆ Bandwidth limiting restricts maximum network usage and does not imply required bandwidth. For example, if you set the bandwidth limit for a client to 500 kilobytes per second, the client can use up to that limit. It does not mean, however, that the client requires 500 kilobytes per second.
- ◆ You cannot use bandwidth limiting to load-balance active backups by having NetBackup pick the most-available network segment. NetBackup does not pick the next client to run based on any configured bandwidth limits.

How Bandwidth Limiting Works

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

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

As the number of active backups increase or decrease on a subnet, NetBackup dynamically adjusts the bandwidth limiting on that subnet. If additional backups are started, the NetBackup server instructs the other NetBackup clients running on that subnet to decrease their bandwidth setting. Similarly, bandwidth per client is increased if the number of clients decreases. Changes to the bandwidth value occur on a periodic basis rather than as backups stop and start. This can reduce the number of bandwidth value changes that are required.

Configuration

You enable bandwidth limiting by adding one or more `LIMIT_BANDWIDTH` entries to the `/usr/opensv/netbackup/bp.conf` file on the master server. These entries let you designate bandwidth values and the IP addresses of the clients and networks to which they apply. For information on adding these entries, see “`LIMIT_BANDWIDTH`” on page 335.

Rules for IP Address Ranges

The IP address ranges can specify individual clients or entire subnets. The following are some specific rules on addresses:

- ◆ An IP address can have any one of the following forms:
 - ◆ `a.b.c.d`

Where `a`, `b`, `c`, and `d` are integers in the range 0-255.



- ◆ `128.net.host`
Class B address (16 bit host).
- ◆ `net.host`
Class A address (24 bit host).
- ◆ `a`
A 32 bit integer, representing the full IP address in network byte order (that is, big endian, the most significant byte is first on the wire).
- ◆ You can enter IP addresses as decimal, octal or hexadecimal numbers. Numbers beginning with 0 are assumed to be octal, numbers beginning with 0x are hexadecimal and all others are assumed to be decimal.
- ◆ Neither the net nor the host part of an IP address can be zero.
- ◆ Only ordinary IP addresses are accepted (class A, B & C, no multicast or reserved addresses).
- ◆ Do not create multiple entries that specify the same range of IP addresses. If you do, NetBackup uses the last one it finds. In the following example, NetBackup uses the second entry.

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 200
```

This rule also applies to multiple entries that specify an exact client address, such as

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 200
```

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 100
```

- ◆ Do not specify IP address ranges that overlap one another. For example, if you specify


```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```

```
LIMIT_BANDWIDTH = 111.222.333.5 111.222.333.255 500
```

 the ranges overlap and bandwidth limiting results are unpredictable.
- ◆ You can specify a range of addresses in one entry and an address for a specific client in other entries.

If a client is covered by an entry that specifies its exact IP address and by another entry that specifies a range of IP addresses, NetBackup uses the bandwidth value in the entry with the exact IP address.

For example:

- ◆ The following sets the bandwidth for a range of IP addresses:

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```



- ◆ The following sets the bandwidth for a specific address that is within the above range.

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 200
```

In this case, NetBackup uses the specific entry (bandwidth of 200) for the client whose address is 111.222.333.111. You can also use this capability to exclude specific clients from bandwidth limiting (see Example 3 below). The order in which the range and specific address entries appear in the `bp.conf` file is not significant.

Rules for Setting Bandwidth Values

When setting bandwidth values for individual clients, you must set it to either:

- ◆ 0 (no bandwidth limiting)
- or
- ◆ Less than or equal to any value set for the IP address range containing the IP address for the client.

For example, the following is valid:

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 300
```

If you set the bandwidth higher for an individual client than it is for the range, NetBackup ignores that setting and uses the value for the range. In this case, the client gets its share of the bandwidth specified for the network.

If the bandwidth limit for an individual client is equal to or lower than the value for the range, the client uses the lower of:

- ◆ Its share of the network bandwidth value
- ◆ Its individual bandwidth value

The bandwidth value that NetBackup uses for a client will always be at least one kilobyte per second.

Examples

Example 1

To configure a bandwidth limit of 500 kilobytes per second for all machines on the subnet 111.222.333, use:

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```



Example 2

To configure a bandwidth limit of 700 kilobytes per second for a particular client (111.222.333.111), use:

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 700
```

Example 3

To disable bandwidth limiting for a client in a subnet that has a bandwidth limit, specify 0 for the kilobytes per second. For example:

```
LIMIT_BANDWIDTH = 111.222.333.1 111.222.333.255 500
```

```
LIMIT_BANDWIDTH = 111.222.333.111 111.222.333.111 0
```

In this case, no limiting occurs for the client with IP address 111.222.333.111

Busy-File Processing (UNIX Clients Only)

Note Busy-file processing applies only to UNIX clients. If you have Microsoft Windows clients, see “Open Transaction Management (Microsoft Windows Clients Only)” on page 303.

A busy file is a file that was detected as changed during a user or scheduled backup. Typically, this occurs if a process is writing to a file while NetBackup is attempting to back it up. The backup usually completes with a status of 1 indicating that the backup was partially successful. The busy-file processing feature lets the user control the actions of NetBackup when busy files are detected.

To enable busy-file processing, you add the `BUSY_FILE_PROCESSING` option to the client `/usr/opensv/netbackup/bp.conf` file. You then add other busy-file options to control the processing of busy files. These other options can exist in both the client `/usr/opensv/netbackup/bp.conf` file and a user's `bp.conf` (the user's `bp.conf` file takes precedence when the options are in both places).

NetBackup creates several files and directories when processing busy files. Initially, a working directory named `busy_files` is created under `/usr/opensv/netbackup`. NetBackup then creates an actions directory under `busy_files` and places action files in that directory. An action file has the information that NetBackup uses to control the processing of busy files. By default, the contents of the action file is derived from the `BUSY_FILE_ACTION` options in `bp.conf`. A user can also create an action file in order to control a specific backup class and schedule. NetBackup creates a logs directory under `busy_files` for storing busy file status and diagnostic information.



Getting Started

Perform the following steps to enable the busy files feature:

- ◆ Modify the `bp.conf` file options as described in “Modifying `bp.conf`.” later in this section.
- ◆ Copy the script

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

to

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

Be sure to set the execute permissions to allow *group* and *other* to execute `bpend_notify`.

- ◆ Configure a class with a user-backup schedule to be used by busy-file backups.

This class will service the backup requests generated by the `repeat` option in the actions file. The class name is significant, since by default, NetBackup searches alphabetically (upper-case characters first) for the first available class with a user backup schedule and an open backup window. For example, a class name of *AAA_busy_files* is selected ahead of *B_class*.

Modifying `bp.conf`

The user can direct busy-file processing by setting the following in the `bp.conf` file.

BUSY_FILE_PROCESSING

Used in a `/usr/opensv/netbackup/bp.conf` file on a client, this option enables the NetBackup busy-file-processing feature. By default, this option is not in `bp.conf`, thus disabling busy-file processing.

BUSY_FILE_DIRECTORY

Used in a `/usr/opensv/netbackup/bp.conf` or `$HOME/bp.conf` file on a client, this option specifies the path to the busy files working directory. By default, `bp.conf` does not contain this option and NetBackup creates the `busy_files` directory in `/usr/opensv/netbackup`.

BUSY_FILE_ACTION

Used in a `/usr/opensv/netbackup/bp.conf` or `$HOME/bp.conf` file on a client, this option directs the action that NetBackup performs on busy files. There can be multiple entries of the form:



`BUSY_FILE_ACTION = filename_template action_template`

Where

- ◆ *filename_template* is the absolute pathname and file name of the busy file. The shell language metacharacters `*`, `?`, `[]`, `[-]` can be used for pattern matching of filenames or parts of filenames.
- ◆ *action_template* is one of:

`MAIL | mail`

Directs NetBackup to mail a busy file notification message to the user specified by the `BUSY_FILE_NOTIFY_USER` option.

`REPEAT | repeat [repeat_count]`

Directs NetBackup to retry the backup on the specified busy file. A repeat count can be specified to control the number of backup attempts. The default repeat count is 1.

`IGNORE | 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.

`BUSY_FILE_NOTIFY_USER`

Used in a `/usr/opensv/netbackup/bp.conf` or `$HOME/bp.conf` file on a client, this option specifies the recipient of the busy file notification message when `BUSY_FILE_ACTION` is set to `MAIL` or `mail`. By default, `BUSY_FILE_NOTIFY_USER` is not in `bp.conf` and the mail recipient is `root`.

Examples

Example 1

```
BUSY_FILE_PROCESSING
BUSY_FILE_NOTIFY_USER = kwc
BUSY_FILE_ACTION = /usr/* mail
BUSY_FILE_ACTION = /usr/local ignore
```

NetBackup will send a mail notification message to the user `kwc` for all busy files that it finds under `/usr`, except for those in `/usr/local`.

Example 2

```
BUSY_FILE_PROCESSING
```



```
BUSY_FILE_ACTION = /usr/opencv mail  
BUSY_FILE_ACTION = /usr/* repeat 2  
BUSY_FILE_ACTION = /usr/local ignore
```

This set of options causes NetBackup to take the following actions when it encounters busy files:

- ◆ Send a busy-file-notification message to root for busy files in `/usr/opencv`.
- ◆ Repeat the backup up to a maximum of two times for all busy files that it finds under `/usr`, except for those in `/usr/opencv` and `/usr/local`.
- ◆ Exclude the busy files in `/usr/local` from all actions.

Creating Action Files

When a backup operation begins, NetBackup creates a default action file named `actions` in the `busy_files/actions` directory. The contents of the actions file are derived from the `BUSY_FILE_ACTION` options in the `bp.conf` file.

NetBackup refers to the default action file for all future busy-file processing, unless you override the default by creating an action file to control a specific backup class and schedule. The naming convention for the class and schedule action files is:

```
actions.class_name.schedule_name
```

or

```
actions.class_name
```

Where *class_name* and *schedule_name* correspond to a predefined backup class and schedule.

When searching for an action file, NetBackup:

1. Checks for file that names a specific class and schedule, such as:

```
actions.class_name.schedule_name
```

2. If a file for a specific class and schedule is not found, NetBackup searches for a less-specific name such as:

```
actions.class_name
```

3. If a less-specific name does not exist, NetBackup refers to the default action file.

The contents of user-created action files are similar to the default. Optional comment lines can be included and the specification is the same as for the `BUSY_FILE_ACTION` option:

```
# comment_line  
filename_template action_template
```



Example 1

If the `bp.conf` file contains

```
BUSY_FILE_ACTION = /usr/opencv mail
BUSY_FILE_ACTION = /usr/* repeat 2
BUSY_FILE_ACTION = /usr/local ignore
```

the default actions file, named `actions`, will contain:

```
/usr/opencv mail
/usr/* repeat 2
/usr/local ignore
```

Example 2

An action file name for a backup class `production_servers` with a schedule name `full` is:

```
actions.production_servers.full
```

If the `actions` file contains

```
/bin/* repeat
```

NetBackup repeats the backup for busy files in `/bin` directory.

Logs Directory

During busy-file processing NetBackup creates a number of files under the `busy_files/logs` directory. These files contain status and diagnostic information that is recorded by NetBackup. NetBackup derives the names of these files from the class name, schedule name, and process id (PID) of the backup.

- ◆ Busy-file log

NetBackup records the names of any busy files in the busy file log. The name of the busy-file log has the form:

```
class_name.schedule_name.PID
```

- ◆ Diagnostic-log file

NetBackup generates a log file that contains diagnostic information. The name of the log file has the form:

```
log.class_name.schedule_name.PID
```

- ◆ Retry-log file



NetBackup also generates a retry file that contains diagnostic information that is recorded when the `repeat` option is specified. The name of the retry file has the form:

```
class_name.schedule_name.PID.retry.retry_count
```

Where *retry_count* starts at zero and is incremented by one every time a backup is repeated. Processing stops when *retry_count* is one less than the number specified on the `repeat` option

Example

To service busy-file backup requests, the administrator defined a class named *AAA_busy_files* that has a user backup schedule named *user*. A scheduled backup is initiated with the class named *production_servers*, schedule named *full*, and PID of 1442.

If busy files are detected, NetBackup generates the following files in the `/usr/opensv/netbackup/busy_files/logs` directory:

```
production_servers.full.1442
log.production_servers.full.1442
```

If the actions file has repeat count set to 2, NetBackup generates the following files:

```
production_servers.full.1442.retry.0
AAA_busy_files.user.10639
log.AAA_busy_files.user.10639
```

If a second repeat backup is attempted, NetBackup generates the following files:

```
production_servers.full.1442.retry.1
AAA_busy_files.user.15639
log.AAA_busy_files.user.15639
```

Modifying `bpend_notify_busy`

The administrator can modify busy-file processing by changing the `bpend_notify_busy` script. The *only* recommended changes are:

- ◆ Changing the `RETRY_CLASS` and `RETRY_SCHED` variables from `NONE` to the busy-file-backup class name and schedule name.
- ◆ Remove the files in the logs directory after busy-file processing (these logs are not removed automatically):

- a. At the end of the `busy_files()` function, add the command:

```
/bin/rm -f $LOG_FILE
```


- b. After the call to the `busy_files()` function in main, add the commands:

```
/bin/rm -f $BUSYFILELOG
```

```
/bin/rm -f $RETRY_FILE
```

Open Transaction Management (Microsoft Windows Clients Only)

On Microsoft Windows clients, NetBackup uses OTM (Open Transaction Manager) to back up files, databases, and applications that are open or active. Open Transaction Management software is included on the same CD-ROM as the server software and can be optionally installed on Microsoft Windows clients. To use OTM, you must have access to the administration interface on a Windows NT/2000 NetBackup server or Administration Client. You cannot configure OTM from a NetBackup UNIX server.

The instructions for configuring OTM are in the online help for the Windows NT/2000 NetBackup administration interface and in the NetBackup system administrator's guide for Windows NT/2000.

Enhanced Authentication

This section describes additional authentication that you can provide for NetBackup programs that communicate through sockets. For example, this authentication is used when a backup or restore operation is started from a client and also during remote administration.

The authentication method that NetBackup uses is called vopie (VERITAS One-time Passwords In Everything). There are two levels of authentication: host and user. First, the hosts authenticate one another. Then, if the user attempting the connection is a non-root user on UNIX or a non-administrator on Windows NT/2000, the user is authenticated as well. Authentication takes place after a NetBackup connection has been established but before any NetBackup transactions have taken place.

Authentication is performed through a series of challenge and responses that require the exchange of secret password information. Passwords are defined during installation and configuration so users do not have to enter them each time they start a backup, archive, or restore.

The authentication software must be installed and configured on all NetBackup servers and clients that require additional authentication.

The following topics are covered here:

- ◆ **Configuration Files** - describes the purpose and content of the configuration files used by authentication. This information is useful if you have to modify these files or if troubleshooting is necessary.



- ◆ **Commands** - provides a brief description of the commands that are available for configuring authentication. Detailed descriptions are in Appendix A.
- ◆ **To Configure Authentication** - describes the basic steps for configuring authentication.
- ◆ **Examples** - contains examples of common configurations.

Configuration Files

The following configuration files are used by the authentication software. Some of them may have to be modified during configuration.

- ◆ `methods.txt`
- ◆ `methods_allow.txt`
- ◆ `methods_deny.txt`
- ◆ `names_allow.txt`
- ◆ `names_deny.txt`

These files are in the following locations:

Windows 2000, NT, 98, 95: `install_path\NetBackup\var\auth` directory

UNIX: `/usr/opensv/var/auth` directory

Macintosh: `:System Folder:Preferences:NetBackup:var:auth` folder

methods.txt

The `methods.txt` file defines the supported authentication methods and must exist for authentication to occur. One method is currently supported:

`vopie` - one-time password authentication

The method is on a separate line that shows the method number, method name, and the path to a shared library:

Windows 2000, NT, 98, 95:

```
128 vopie install_path\NetBackup\lib\libvopie.dll
```

UNIX (except HP-UX):

```
128 vopie /usr/opensv/lib/libvopie.so
```

UNIX (HP-UX only):

```
128 vopie /usr/opensv/lib/libvopie.sl
```

Macintosh:

```
128 vopie libvopie.dll
```



The syntax rules for this file are as follows:

- ◆ Empty lines are ignored
- ◆ The # character and all following characters on a line are ignored.

methods_allow.txt

The `methods_allow.txt` file defines the authentication methods that NetBackup servers and clients can use. When a client or server attempts a connection, it specifies the authentication method it is using. The other server or client then checks its `methods_allow.txt` file to determine if that method is allowed for the system that is attempting the connection. If an entry in this file matches the host and method, the method is allowed. Otherwise, NetBackup checks the `methods_deny.txt` file (see next topic).

The following is an example file:

```
# All hosts in the ourcompany.com domain and host name
# bob.theircompany.com can use the vopie method.
vopie : .ourcompany.com, bob.theircompany.com
#
# Hosts with IP addresses in the 12.123.56 network and IP address
# 2.123.57.23 can use all methods.
ALL : 12.123.56
ALL : 12.123.57.23
```

The keyword `ALL` is used to specify all valid methods (see example above) or all possible hosts.

The default file contains only the following entry:

```
ALL : ALL
```

This means that all methods are allowed for all hosts. If the default file is empty, the `methods_deny.txt` file is always checked.

Other syntax rules for this file are as follows:

- ◆ Each entry must be on a separate line.
- ◆ Empty lines are ignored.
- ◆ The # character and all following characters on a line are ignored.
- ◆ If a domain name is preceded by a dot (.), all hosts in that domain will match.
- ◆ If a network number is followed by a dot (.), all IP numbers in that network will match.



- ◆ A comma separated list of name patterns and number patterns can be specified on a single line.

methods_deny.txt

The `methods_deny.txt` file defines the authentication methods that NetBackup servers and clients cannot use. NetBackup checks this file only if the `methods_allow.txt` file does not have a matching entry for the host and method (see previous topic). If a matching entry is found in `methods_deny.txt` the method is not allowed and authentication is not used. Otherwise, the method is used and authentication proceeds.

The following is an example file:

```
# All hosts in the ourcompany.com domain cannot use the vopie method.
vopie : .ourcompany.com
#
# Hosts with IP addresses in the 12.123.56 network cannot use all
# methods.
ALL : 12.123.56.
```

The default file contains only the following entry:

```
ALL : ALL
```

This means that all methods are denied for all hosts, unless it is specified otherwise in the `methods_allow.txt` file.

The syntax rules for `methods_deny.txt` are the same as for `methods_allow.txt`.

names_allow.txt

The `names_allow.txt` file defines the network host names that a NetBackup client or server can use when establishing connections. This file is required when NetBackup client or server names do not correlate to their host names and IP addresses. For example, when:

- ◆ NetBackup clients are using DHCP or another dynamic addressing scheme. Here, a client probably uses a different IP address each time it attempts a connection.
- ◆ A NetBackup server or client has more than one network interface. Here, the host name associated with the IP address can be different than the NetBackup client name.
- ◆ A NetBackup client connects through a gateway. Here, the peername for the gateway can be different than the NetBackup client name.

In the above instances, when a client or server attempts a connection, NetBackup checks the `names_allow.txt` file to determine if the network-host name for the connection correlates to a NetBackup name. If a match is found, the connection is allowed. Otherwise, NetBackup checks the `names_deny.txt` file (see next topic).

If NetBackup client and server names correlate to their host names and IP addresses, then neither `names_allow.txt` file or `names_deny.txt` are used.

Each line in `names_allow.txt` contains a logical name (usually, a NetBackup client name) followed by a colon and then a list of host names or IP addresses. The following is an example file:

```
# The next three client entries can match IP numbers in the
# 123.123.56 network.
client1 : 123.123.56.
client2 : 123.123.56.
client3 : 123.123.56.
#
# The entry below permits the name fred to be used for hosts
# dhcp0 and dhcp1 in the ourcompany.com domain.
fred : dhcp0.ourcompany.com, dhcp1.ourcompany.com
```

The default file contains only the following entry:

```
ALL : ALL
```

This means that all names are allowed for all hosts. If the default file is empty, the `names_deny.txt` file is always checked.

The syntax rules for `names_allow.txt` are the same as for `methods_allow.txt`. The only variation is the `ALL` keyword, which in this case specifies all valid names (see example above) or all possible hosts.

names_deny.txt

The `names_deny.txt` file defines the NetBackup client or server names that hosts cannot use. NetBackup checks this file only if the `names_allow.txt` file does not have a matching entry for the host and name (see previous topic). If a matching entry is found in `names_deny.txt` the name is not allowed and authentication fails. Otherwise, the name is used and authentication proceeds.

The following is an example file:

```
# The entry below prevents the name fred to be used for hosts
# in the theircompany.com domain.
fred : .theircompany.com
#
# The entry below prevents any names from being used for hosts
# with IP addresses in the 12.123.53 network.
```



ALL : 123.123.53.

The default file contains only the following entry:

ALL : ALL

This means that all names are denied for all hosts, unless it is specified otherwise in the `names_allow.txt` file.

The syntax rules for `names_deny.txt` are the same as for `names_allow.txt`.

Library Files

The library files that are required for authentication depend on the platform. See “methods.txt” on page 304.

Authentication Commands

The following commands are used to configure and manage authentication. For more information on these commands, see Appendix A.

bpauthsync

Executed on the master server and sets up authentication for one or more clients and media servers. For example, it ensures that the hashed and unhashed files contain the correct information. It is in the following directory:

install_path\NetBackup\bin\admincmd\ (Windows NT/2000)

/usr/opensv/netbackup/bin/admincmd/ (UNIX)

vopie_util

Executed on the client and manages the hashed and unhashed files. This command generates the secret key for the local system and also the information that must be placed into the hashed file on systems that want to access this one. It is in the following directory:

install_path\NetBackup\bin\ (Windows NT/2000)

/usr/opensv/bin/ (UNIX)

vopied Daemon

The vopie daemon, `vopied`, manages the authentication of nonroot users on Windows NT/2000 and UNIX clients and servers. By default, NetBackup configures the system to automatically start `vopied` when the system is started. To start `vopied` directly, execute on the client or server:



install_path\NetBackup\bin\vopied (Windows NT/2000)

/usr/opensv/bin/vopied (UNIX)

vopie Files

The vopie processes use the following files during authentication.

hashed (public key) Files

The hashed files contain the authentication challenges that the local system presents to remote systems. These files are as follows:

Windows 2000, NT, 98, 95:

install_path\NetBackup\var\auth\vopie\hashed*localhost**remotehost*.txt

UNIX:

/usr/opensv/var/auth/vopie/hashed/*localhost*/*remotehost*.txt

Macintosh

:System Folder:Preferences:

NetBackup:var:auth:vopie:hashed:auth:*localhost:remotehost*.txt

Where:

- ◆ *localhost* is the local system.
- ◆ *remotehost* contains the challenges for the remote system named *remotehost*.

There is a *remotehost*.txt file for each remote system that can be authenticated. Only root on the local system can read or write these files.

unhashed (secret key) Files

The unhashed files contains the secret key that NetBackup uses when it responds to challenges from remote systems. These files are as follows:

Windows 2000, NT, 98, 95:

install_path\NetBackup\var\auth\vopie\unhashed*localhost**remotehost*.txt

UNIX:

/usr/opensv/var/auth/vopie/unhashed/*localhost*/*remotehost*.txt

Macintosh

:System Folder:Preferences:

NetBackup:var:auth:vopie:hashed:auth:*localhost:remotehost*.txt

Where:

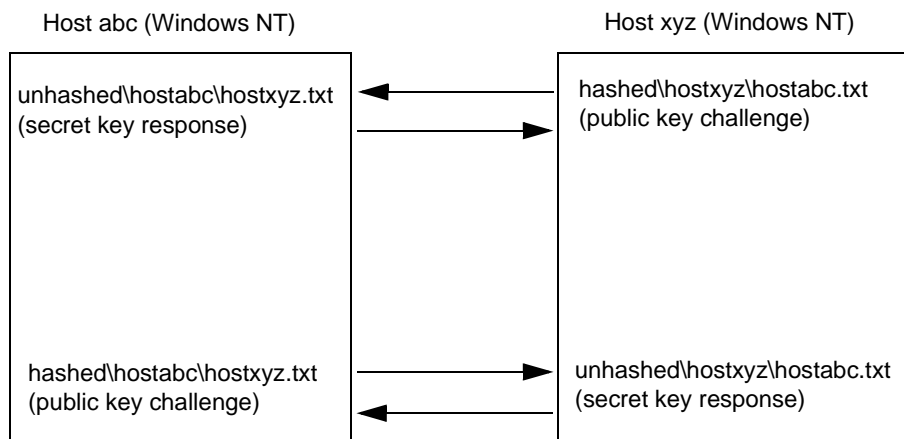


- ◆ *localhost* is the local system.
- ◆ *remotehost.txt* contains the responses for the remote system named *remotehost*.

There is a *remotehost.txt* file for each remote system that can request authentication. These files are created during installation and only root on the local system can read or write these files.

Caution Protect the unhashed files by allowing access only by the root on the local system. Also, do not NFS mount them on UNIX or place them on a network drive on Windows NT/2000.

The `bpauthsync` command synchronizes the information between the hashed files on the local host with the unhashed files on remote systems. This enables the remote host to offer the correct response when it is challenged. The following figure illustrates this exchange between Windows NT systems.



temp File

On a Windows NT/2000 or UNIX system, the `vopie` daemon, `vopied`, creates a temporary file where it stores the challenges and responses required to authenticate nonroot users. This is necessary because nonroot users cannot access the files in the hashed and unhashed directories. The temporary files are valid for only one connection and are automatically deleted. They are in:

Windows NT/2000:

`install_path\NetBackup\var\auth\vopie\temp\username\tempname.txt`

UNIX:

`/usr/opensv/var/auth/vopie/temp/username/tempname.txt`

To Configure Authentication

1. Install NetBackup 3.2 (or later) on each system requiring authentication.

The NetBackup install process installs the necessary files and commands. The administrator then uses commands to set up the files so they contain the proper authentication information.


2. Execute the `/usr/opensv/netbackup/bin/admincmd/bpauthsync` command on the master server.

This command sets up authentication files on the NetBackup servers and clients. For more information, see:

- ◆ Examples following this procedure
- ◆ `bpauthsync` command description in Appendix A

Example 1 - Typical Configuration

Assume that this is an initial installation and you want to configure authentication for all systems in the figure below. NetBackup server and client software has already been installed.



1. Execute the following on the master server (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -vopie -servers
-clients
```

This synchronizes the key files on all the systems.

2. On the master server, copy the `methods_allow.txt` to a temporary file. For example, `/tmp/ma.txt`.
3. To the temporary file, add an entry for each host that requires authentication:

```
vopie : shark
vopie : whale
vopie : dog
vopie : cat
```



4. Synchronize the `methods_allow.txt` files on the servers and the new client by executing the following on the master server (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -methods
-methods_allow /tmp/ma.txt -servers -clients
```

The information in `/tmp/ma.txt` is written in the `methods_allow.txt` files on the servers and clients.

Example 2 - Disable Authentication for a Client

To disable authentication for the client named `cat` in the previous figure:

1. Push an empty `methods_allow.txt` file to the client by executing the following on the master server (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -methods
-methods_allow /dev/null -clients cat
```

This disables authentication on the client.

2. On the master server, remove the entry for `cat` from the `/usr/opensv/var/auth/methods_allow.txt` file.
3. Synchronize the methods files on all servers by executing the following on the master server (all on one line):


```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -methods -servers
```

Authentication is no longer performed when communicating with this client.

Example 3 - Add a Client

Assume that all systems below are configured for authentication, except for the client named `cat`.

To add authentication for `cat`:



1. On the master server, copy the `methods_allow.txt` to a temporary file. For example, `/tmp/ma.txt`.
2. Add an entry for the new client to the temporary file:

```
vopie : cat
```

3. Synchronize the methods files on the servers and the new client by executing the following on the master server (all on one line):


```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -vopie -methods  
-methods_allow /tmp/ma.txt -servers -clients cat
```

The information in `/tmp/ma.txt` is written in the `methods_allow.txt` files on the servers and the client.

Example 4 - Restore Authentication After Client Disk Crash

Assume that `cat` was configured for authentication and the disk failed.

To restore authentication so all files can be recovered, proceed as follows:



1. On the master server, copy the current `methods_allow.txt` file to another file. For example, copy it to `/usr/opensv/var/auth/methods_allow.txt.save`
2. Remove the entry for the failed client from `methods_allow.txt` on the master server.
3. Push the modified `methods_allow.txt` file to the other servers by executing the following (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -methods -servers
```

This disables authentication for the failed client so the servers can communicate with it during recovery.

4. Reinstall the operating system (Windows 2000, NT, or UNIX) and NetBackup 3.2 (or later) on the failed client by following the instructions in the troubleshooting guide. However, do not restore any NetBackup or user files at this time.
5. On the master server, execute the following to synchronize and push the original methods to the servers and the failed client (the command is all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -vopie -methods  
-servers -clients cat -methods_allow  
/usr/opensv/var/auth/methods_allow.txt.save
```




The information in `methods_allow.txt.save` is written in the `methods_allow.txt` files on servers and the client. The original authentication methods are now restored.

Note Do not restore the files in the `/usr/opensv/var/auth` directory on the client or authentication will have to be resynchronized.

6. Complete the client recovery by restoring the original NetBackup and user files as explained in the *NetBackup Troubleshooting Guide - UNIX*.

Example 5 - Restore Authentication on NetBackup Master Server

Assume that authentication was configured on all servers and clients and the disk fails on the master server shark.



If the NetBackup catalog backup was written to a storage unit on the master server shark, proceed as follows:

1. On the master server, recover the disk `s` explained in the troubleshooting guide and reinstall NetBackup.
2. Restore all files to the master server.
3. Synchronize all clients and servers by executing the following on the master server (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -vopie -servers  
-clients
```

If the NetBackup catalog backup was written to a storage unit on whale, shark cannot recover the catalogs because the two servers cannot authenticate one another. In this instance, the following steps are required:

1. Install netbackup 3.2 or later on the master server (do not restore any files at this time).
2. Disable authentication between the master server and the media server where the catalog backup was written, by modifying their `methods_allow.txt` files:

- a. On the master server, remove the entry for the media server from the `methods_allow.txt` file (if an entry is present).
 - b. On the media server, remove the entry for the master server from the `methods_allow.txt` file.
3. On the master server, execute the `bprecover` command to restore the catalog files.
 4. Restore all files to the master server, including those in the `/usr/opensv/var/auth` directory.
 5. On the media server, add back the entry for the master server from the `methods_allow.txt` file.
 6. Synchronize all servers and clients by executing the following on the master server (all on one line):

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync -vopie -servers
-clients
```

The original configuration is now restored.

Troubleshooting Authentication

If you have problems with authentication, perform the following steps:

1. Look for status code 160 (authentication failed). If you see this status code, go to the *NetBackup Troubleshooting Guide - UNIX* for corrective actions.
2. Create activity log directories for the processes involved in communication between NetBackup systems. These include:
 - ◆ On the server create activity log directories for `bprd`, `bpdbm`, `bpcd`.
 - ◆ On the client, create activity log directories for `bpbackup`, `bprestore`, `bpbkar` (Windows NT/2000 only).
3. Retry the operation and check the logs.

NetBackup Authorization

NetBackup user authorization provides a platform-independent mechanism for selected users (or groups of users) to administer a NetBackup server from a remote NetBackup *administrative console*.

An administrative console can be any of the following:

- ◆ A Windows NT/2000 or UNIX NetBackup master or media server.



- ◆ An administration client. This is a Windows NT client that has NetBackup administration interface software installed and can be used as a remote console for administering NetBackup.

If you are using NetBackup-Java through the Windows Display Console on a Windows NT/2000 platform, the administrative console is the UNIX or Windows NT/2000 computer that you log into when starting the NetBackup-Java interface.

Configuring NetBackup Authorization

Step 1: Add NetBackup Servers To Each Others Server Lists

The NetBackup master and all the media servers that you will be administering (including a server used as an administrative console) must be listed in each others server and `vm.conf` files (SERVER entries). However, when you are using authorization, it is not necessary for the server lists on the NetBackup servers to contain an entry for an administration client that you are using for an administrative console.

Step 2: Enable NetBackup Authentication

To use the authorization feature, you must enable NetBackup authentication between administrative consoles and the NetBackup servers to be administered. To perform administrative tasks on clients, such as client configuration, you must also enable NetBackup authentication between the clients and administrative consoles.

For more on authentication, see:

- ◆ “Enhanced Authentication” on page 303.
- ◆ “vmd Security” in the system administrator’s guide for Media Manager.

Step 3: Create an `authorize.txt` file

To enable the authorization feature on a NetBackup master or media server, create an `authorize.txt` file on that server, as explained in the following sections. After this file is created, the server requires authorization from any administrative console that attempts remote administration.

Note To ensure a secure NetBackup server, always restrict access to the `authorize.txt` file.

`authorize.txt` File Location

On a UNIX NetBackup master or media server, create the file in the following location:

```
/usr/opensv/var/authorize.txt
```



On a Windows NT/2000 NetBackup master or media server, create the file in the following location:

```
install_path\NetBackup\var\authorize.txt
```

authorize.txt File Format

Use the following format for authorization entries in the `authorize.txt` file:

```
user_name:host_name:domain_group_name[ :local ]
```

If the administrative console is UNIX:

- ◆ *user_name* is the UNIX user name, or * for all users.
- ◆ *host_name* is the remote NetBackup administrative console name, or * for all hosts.
- ◆ *domain_group_name* is a netgroup name or a local group name. For information about netgroups refer to the `netgroup` man page.
- ◆ `local` (if specified) indicates that the *domain_group_name* is a local group name.

You can use a * symbol in the *user_name* and *host_name* fields to authorize all users and/or hosts. For comments, use a # symbol.

If the administrative console is Windows NT/2000:

- ◆ *user_name* is the Windows NT/2000 Administrator name, or * for all users.
- ◆ *host_name* is the remote NetBackup Administrative console host name, or * for all users.
- ◆ *domain_group_name* is the Windows NT/2000 domain and group name in the form *domain\group*.
- ◆ `local` (if specified) indicates the group is not a domain group, but is local to the host specified by *host_name*.

For comments, use a # symbol.

Example authorize.txt File Entries

```
# Authorize 'root' with a local group name
# of 'admin' on the UNIX server
dogroot:dog:admin:local

#

# Authorize all NT Administrators that are
#members of NETBACKUP\Domain Admins
*:*:NETBACKUP\Domain Admins
```



Step 4: Specify Preferred Group on Administrative Console (optional)

On the *administrative console*, you can specify a *preferred group* for authorizing administrative users. The preferred group entry is intended specifically for use with NetBackup authorization and determines the *domain_group_name* that is sent to the NetBackup server.

Some NetBackup processes also use the preferred group entry for Media Manager authorization. For more information on this, see the “Media Manager Configuration File (vm.conf)” section in the system administrator’s guide for Media Manager.

Note To facilitate a platform-independent implementation, the string in the preferred group entry is case sensitive for both UNIX and Windows NT/2000. An example entry for a Windows NT/2000 computer would be: “NTDOMAINNAME\Backup Operators”

On UNIX Administrative Consoles:

On UNIX administrative consoles, specify the preferred group by adding the `PREFERRED_GROUP` entry to the `bp.conf` file. This entry must be of the form:

```
PREFERRED_GROUP = netgroup name
```

- ◆ If the `bp.conf` configuration file has a `PREFERRED_GROUP` entry, the `innetgr()` function is used to determine if the user is in the netgroup (for further details refer to the `innetgr` 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 that netgroups are not supported for Pyramid and Sequent systems.

On Windows NT/2000 Administrative Consoles:

On Windows NT/2000 administrative consoles, specify the preferred group with the Preferred group setting on the Universal Settings tab of the client and server dialog boxes in the Configure - NetBackup window as follows:

1. Click NetBackup Configuration on the Start menu in the NetBackup Administration window.
The Configure-NetBackup window appears.
2. Select the client or server
3. Click the Properties (read/write) command on the File menu.
4. On the Universal Settings tab, select the Preferred group check box and type the group name in the text box. An example entry for a Windows NT computer is:



NTDOMAINNAME\Backup Operators


The entry is interpreted as follows:

- ◆ If Preferred group is specified, a check is made to determine if the user is a member of *domain\group*. This check is limited to Windows NT global groups. That is, if Preferred group specifies a local group, a match does not occur and the user's primary *domain\group* is used.
- ◆ If Preferred group is not specified or the user is not a member of the *domain\group*, the user's primary *domain\group* is obtained.

When the domain name is an empty string or is the name of the local machine, it is considered to be local.

Example Configuration

The following explains how to set up NetBackup Authorization between the computers in the figure below.



1. Update the server lists and `vm.conf` files as follows:
 - ◆ On shark, add dolphin to the server list and `vm.conf` file.
 - ◆ On dolphin, add shark to the server list and `vm.conf` file.
 - ◆ On whale, add shark and dolphin to the server list.
2. Enable NetBackup authentication:
 - a. On shark, execute:


```
bpauthsync -vopie -servers shark dolphin whale
```
 - b. On shark, edit the `C:\tmp_file` to contain:


```
vopie: shark
vopie: dolphin
vopie: whale
```



- c. On shark, execute (all on one line):

```
bpauthsync -methods_allow c:\tmp_file -servers shark dolphin
whale
```

- 3. Create a global group named:

MYDOMAIN\NetBackup Admins

- 4. Edit the `authorize.txt` files on shark and dolphin so they contain:

```
*:*:MYDOMAIN\NetBackup Admins
```

- 5. On whale, set the preferred group to be:

MYDOMAIN\NetBackup Admins


NetBackup User Authorization Process Description

The following describes the flow for a request from a remote NetBackup administrative console to a NetBackup master server.

Gaining Access to a Server

When an administrator on a remote NetBackup administrative console makes a request to a NetBackup server, and authentication is enabled between the two systems, the `user_name`, `host_name`, `domain_group_name`, and `local` flag are passed from the requesting NetBackup administrative console to the NetBackup master server accepting the request.

After passing authentication, the accepting NetBackup master server checks for the existence of the `authorize.txt` file and for an entry in the file that matches the information passed by requester. If a match exists, the request is authorized (that is, allowed). If request is not authorized, the request can proceed only if the NetBackup administrative console making the request has a `SERVER` entry in the accepting server's NetBackup configuration file. Otherwise, the request fails, indicating a request from invalid server. You also need an entry in the `vm.conf` file in order to use Media Manager applications (see the system administrator's guide for Media Manager).

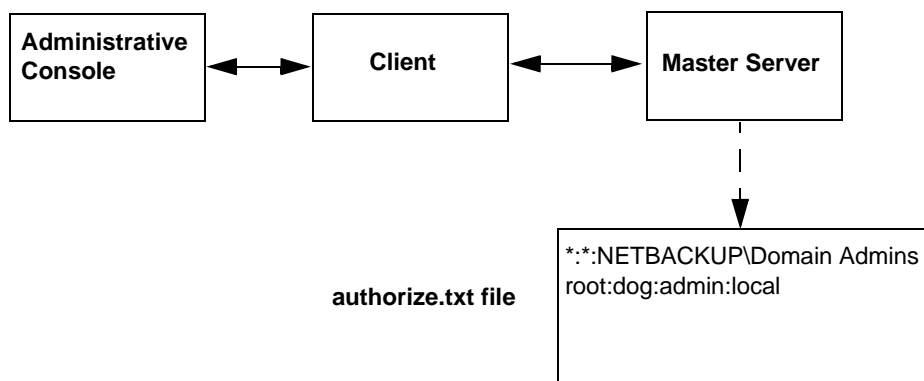


Gaining Access to a Client

Some requests, such as client configuration, are made directly to a client. These types of requests do not require an `authorize.txt` file on the client. The following describes the flow for a request from a remote NetBackup administrative console to a NetBackup client.

When an administrator on a remote NetBackup administrative console makes a request to a NetBackup client, and authentication is enabled between the two systems, the `user_name`, `host_name`, `domain_group_name`, and `local` flag are passed from the requesting NetBackup administrative console to the NetBackup client accepting the request.

If the requesting NetBackup administrative console is not in the client's server list, the client requests authorization from its master server (this is the first server listed in the server list). The NetBackup administrative console authorization information is passed to the master server. The master server checks for the existence of the `authorize.txt` file and for an entry in the file that matches the information passed. If a match exists, authorization is granted, otherwise authorization is denied.



Configuring Email Notifications

You can configure NetBackup to send E-mail notifications to users and administrators on the results of backup, archive, and restore operations. The types of notifications you can configure are:

- ◆ Notification to server administrator when a scheduled backup, administrator-directed manual backup, or a backup of the NetBackup databases occurs.

Configure NetBackup to mail these notifications by specifying the server administrator's address with the NetBackup global attribute, E-mail address for notifications.



If you customize the `dbbackup_notify` script, to include an email message and recipient, this script also sends a message after each NetBackup database backup.

- ◆ Notification to users on UNIX clients as to the success or failure of their user operations.

To configure these notifications, specify the user's E-mail address with the `USEMAIL` option in the user's personal `bp.conf` file. This file is located in the user's home directory (create one if necessary).

- ◆ Notification to system administrators on UNIX clients about the success or failure of scheduled or manual backups.

To configure these notifications, specify the client administrator's address with the `USEMAIL` option in the `/usr/opensv/netbackup/bp.conf` file on the client.

You can also set up mail notifications with the scripts provided with NetBackup UNIX server software (see "Goodies Scripts" on page 242).

Adjust Time Zone

When working with sites that have geographically dispersed NetBackup server and client machines, it can be necessary to adjust the time zone for `jnbSA` when restoring files. Set the time zone relative to the server's time zone.

1. On the Configure menu in the NetBackup Administration window, click Adjust Application Time Zone.

The Adjust Timezone dialog box appears.

2. Adjust the time to reflect how many hours/minutes the server's time zone is behind or ahead of Greenwich Mean Time.
3. To use daylight savings time, select Use Daylight Savings Time.
4. Indicate when Daylight Savings Time should begin.
 - a. Select the method you wish to use for indicating when Daylight Savings Time should begin.
 - ◆ To have DST begin on a specific date, select Absolute date and indicate the desired month and day.
 - ◆ To have DST begin on the first occurrence of a day in a month, select First day of week in month and indicate the desired day of the week and the month.
 - ◆ To have DST begin on the first occurrence of a day in a month and after a specific date, select First day of week in month after date and indicate the desired day of the week and the month and day.



- ◆ To have DST begin on the last occurrence of a day in a month, select Last day of week in month and indicate the desired day of the week and the month.
 - ◆ To have DST begin on the last occurrence of a day in a month and before a specific date, select Last day of week in month after date and indicate the desired day of the week and the month and day.
- b. Select the appropriate Day of week, Month, Day, and Time.

To have DST begin/end:

Enter this:

April 5

First Monday in April

First Monday after April 5

Last Thursday in April

Last Thursday before April 30

5. As you did in step 4, indicate when Daylight Savings Time should end.
6. If you wish to have the time zone settings applied to the current session and all future sessions, select Save as default time zone.
7. Click OK.

Specifying the Locale of the NetBackup installation

NetBackup applications can display a wide range of international date and time formats as determined by the locale of the installation. To help ensure consistency among the applications, NetBackup uses a single configurable source to define the locale conventions.



For NT/98/2000 platforms:

To access the regional settings, double-click Regional Settings in the Windows Control Panel. This provides access to the predefined Number and Date/Time formats.

See the Microsoft Help pages for further assistance.

For Macintosh platforms:

Use the Date & Time Control Panel to change the values for the current date and current time, as well as to customize the date and time formats.

See the Mac OS System Software manual pages for further assistance.

For UNIX platforms:

The `/usr/opensv/msg/.conf` file contains information on the supported locales. This file defines the date and time formats for each supported locale.

The `.conf` file contains very specific instructions on how to add or modify the list of supported locales and formats. However, the format of the file is summarized here.

The `.conf` file is divided into two parts, the TL lines and the TM lines.

TL lines

The third field of the TL lines defines the case-sensitive locales that the NetBackup applications support. The fourth and fifth fields define the date and time fields and associated separators for that supported locale.

You can modify the existing formats to change the default output. For example, the TL line for the C locale is:

```
TL 1 C : hh: mn: ss/ mm/ dd/ yyyy
```

An alternate specification the order of months, days, and years could be:

```
TL 1 C : hh: mn: ss - yyyy- mm- dd
```

or:

```
TL 1 C : hh: mn: ss/ dd/ mm/ yy
```

To add more TL lines, see the comments in the `.conf` file for more information.

If the `.conf` file is not accessible, the default locales (TL lines) are:

```
TL 1 C : hh:mn:ss / mm/ dd/ yyyy
```

```
TL 2 ov : hh: mn: ss/ mm/ dd/ yyyy
```

Note that C and ov are synonymous.



TM lines

The TM lines define a mapping from unrecognized locales to those supported by NetBackup, as defined by the TL lines.

The third field of the TM lines define the unrecognized locale and the fifth field defines the supported equivalent identified in the TL lines.

For example, to map the unrecognized locale *french* to the supported locale *fr*, the TM line is:

```
TM 6 french 2 fr
```

or, to map french to C

```
TM 6 french 1 C
```

To add more TM lines, see the specific instructions in the `.conf` file.

If the `.conf` file is not accessible, there are no default TM lines as the default locale will be C (ov)

NetBackup Configuration Options

The NetBackup configuration options allow the administrator to customize NetBackup to meet specific site preferences and requirements. In most instances, the internal software defaults provide satisfactory results. However, if settings must be changed from their defaults, do so according to the following instructions.

Method for Specifying the Configuration Options

The method to use for specifying the configuration options depends on the type of server or client you are configuring.

- ◆ On NetBackup UNIX servers and clients, specify the configuration options in the `bp.conf` file as explained in this chapter.
- ◆ On NetBackup Windows NT/2000 servers, these options are referred to as NetBackup properties and are explained in the NetBackup system administrator's guide for Windows NT/2000 Server.
- ◆ On PC clients, specify configuration options as explained in the NetBackup user's guide for the client.

Note After making a change to the `/usr/opensv/netbackup/bp.conf` file on the master server, stop and restart all NetBackup daemons and utilities. This ensures that the new `bp.conf` values will be used by all the NetBackup processes that



require them (a process reads `bp.conf` only when it begins). This action is not required for changes to `bp.conf` files on a client or to a `$HOME/bp.conf` file on the master server.

Syntax Rules for `bp.conf` Options

When creating entries in a `bp.conf` file, you can use:

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

`bp.conf` Options for UNIX Servers

The `bp.conf` options for NetBackup UNIX servers are located in the following file:

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

If a single UNIX system is running as both a client and a server, the `/usr/opensv/netbackup/bp.conf` file will contain both server and client options.

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

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

See the `bp.conf` discussion for UNIX clients later in this chapter for an explanation of client options which of these can be in a personal `bp.conf` file.

The following are the options that you can specify in the `/usr/opensv/netbackup/bp.conf` file on a NetBackup UNIX server.

Note The `SERVER` option *must* be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX clients and servers. It is also the *only required* entry in these `bp.conf` files. As installed, NetBackup uses internal software defaults for all options in the `bp.conf` file, except `SERVER`. During installation, NetBackup sets the `SERVER` option to the name of the master server where the software is installed. The `SERVER` entries *MUST* be the same on all servers in a master and media server cluster. It is recommended (but not mandatory) that all other entries, except `CLIENT_NAME`, also match on all servers.

ALLOW_MEDIA_OVERWRITE

Overrides NetBackup's overwrite protection for the following media formats on removable media used for backups:

- ◆ ANSI - ANSI labeled media
- ◆ AOS/VS - Data General AOS/VS backup format
- ◆ CPIO - `cpio` format
- ◆ DBR - format from a former backup product
- ◆ MTF1 - Backup Exec
- ◆ TAR - `tar` format

You can add this option to the

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

file on NetBackup servers.

For example, to permit overwriting the `cpio` format, add the following on the master server (and media servers if applicable):

```
ALLOW_MEDIA_OVERWRITE = CPIO
```

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 block on a variable length media be less than or equal to 32 kilobytes.

If media contains one of the above formats and you do not permit media overwriting, NetBackup takes the following actions:

- ◆ If the volume has not been previously assigned for a backup, NetBackup:
 - ◆ 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, NetBackup:
 - ◆ 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 catalogs, 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.

When using `bplabel`, you receive a prompt for user action if NetBackup encounters:

- ◆ Any of the protected formats listed above (unless you use `bplabel -o` to unconditionally overwrite these formats)
- ◆ A NetBackup label
- ◆ The NetBackup catalogs

See the `bplabel(1M)` man page for more information on using the `bplabel` command.

ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA

Allows NetBackup to mix retention levels on media. It applies to media in both robotic and nonrobotic drives. By default, this option is not present and each volume can contain backups of only a single retention level.

You can add this option to the

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

file on NetBackup servers.

ALLOW_NON_RESERVED_PORTS

Specifies that the NetBackup client daemon (`bpcd`) can accept remote connections from nonprivileged ports (port numbers 1024 or greater). If this entry is not present, then `bpcd` requires remote connections to come from privileged ports (port numbers less than 1024). This option can be useful when NetBackup clients and servers are on opposite sides of a firewall.

You can add this option to the

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

file on a UNIX server or client (for use on a client, see “ALLOW_NON_RESERVED_PORTS” on page 344).

APOLLO_RESTORE_TIMEOUT

Note This option applies only to Apollo clients. It also has a reasonable default and will have to be changed only under very specific circumstances.

Specifies the number of seconds to use for client read timeouts for restores on Apollo clients.

You can add this option to the `/usr/opensv/netbackup/bp.conf` file on UNIX NetBackup servers.

By default, `APOLLO_RESTORE_TIMEOUT` is not specified and the option has a value of 0 (no timeout).

BPEND_TIMEOUT

Note If you change this option, verify that the `CLIENT_READ_TIMEOUT` option is set to the same or higher value.

Specifies the number of seconds to wait for the `bpend_notify` script on a client to complete. The default timeout is 300 seconds.

You can add this option to the

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

file on NetBackup servers.

BPSTART_TIMEOUT

Note If you change this option, verify that the `CLIENT_READ_TIMEOUT` option is also set to the same or higher value.

Specifies the number of seconds to wait for the `bpstart_notify` script on a client to complete. The default timeout is 300 seconds.

You can add this option to the

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

file on NetBackup servers.

BPTM_QUERY_TIMEOUT

Determines the time that the scheduler waits for a drive-count query to `bptm` to complete. If you have problems with timeouts you can modify this setting to extend the time that the scheduler waits. The default is 480 seconds (8 minutes). You can add this option to the

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

file on NetBackup master servers. The following is an example entry:

```
BPTM_QUERY_TIMEOUT=80
```

Also see “Automatic Drive Availability Checking” on page 36.



CLIENT_CONNECT_TIMEOUT

Specifies the number of seconds that the server waits before timing out when connecting to a client. The default timeout is 300 seconds.

You can add this option to the

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

file on NetBackup servers.

CLIENT_PORT_WINDOW

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 a client configured to accept nonreserved ports (for information on client configuration, see “ALLOW_NON_RESERVED_PORTS” on page 344).

You can add this option to the

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

files on NetBackup servers or clients.

For example:

```
CLIENT_PORT_WINDOW = 4800 5000
```

permits ports from 4800 through 5000.

If you specify 0 for the first number, the operating system determines the nonreserved port to use. The default is 0 0, which means the operating system determines the nonreserved port.

CLIENT_READ_TIMEOUT

Note Use this option only on a server or a database-extension client (such as NetBackup for Oracle). This option has a reasonable default and has to be changed only if problems are encountered.

Specifies the number of seconds to use for the client-read timeout.

You can add this option to the

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

file on NetBackup servers.

You can also add this option on database-extension clients (such as NetBackup for Oracle).



The `CLIENT_READ_TIMEOUT` on a database-extension client is a special case because these types of clients can initially require more time to get ready than other clients. This is the case because database backup utilities frequently start several backup jobs at the same time, which slows the CPU.

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 value to use initially. If the option is not set, the standard default of five minutes is used.
- ◆ When the database-extension API receives the server's value, it uses it as the `CLIENT_READ_TIMEOUT`.

By default, `CLIENT_READ_TIMEOUT` is not specified on either a server or database-extension client and the timeout is 300 seconds.

Note We suggest that you set `CLIENT_READ_TIMEOUT` on the database-extension client to a value greater than 5 minutes. A setting of 15 minutes has been found to be adequate for many installations.

CLIENT_RESERVED_PORT_WINDOW

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 a client configured to accept only reserved ports (for information on client configuration, see “ALLOW_NON_RESERVED_PORTS” on page 344).

You can add this option to the

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

files on NetBackup servers or clients.

For example:

```
CLIENT_RESERVED_PORT_WINDOW = 900 1023
```

permits ports from 900 through 1023.

The default range is 512 through 1023. Note that if you specify 0 for the first number, a nonreserved port is used instead and is chosen by the operating system.

DISABLE_JOB_LOGGING

Disables the logging of job information required by the NetBackup job monitor. By default, job logging occurs.

You can add this option to the

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



file on NetBackup servers.

DISABLE_STANDALONE_DRIVE_EXTENSIONS

Disables the nonrobotic drive operations described in “How NetBackup Uses Media in Standalone Drives” on page 640. This means that during a backup, NetBackup does not automatically attempt to use whatever labeled or unlabeled media it finds in a nonrobotic drive. By default, standalone drive extensions are enabled.

You can add this option to the

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

file on NetBackup servers.

DISALLOW_BACKUPS_SPANNING_MEDIA

Prevents backups from spanning media. If the end of media is encountered and this option is present, the media is set to FULL and the operation terminates abnormally (applies to both robotic and nonrobotic drives). By default, backups can span media.

You can add this option to the

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

file on NetBackup servers.

DISALLOW_CLIENT_LIST_RESTORE

Note To override the `DISALLOW_CLIENT_LIST_RESTORE` option for individual clients by changing their `list_restore` setting. See “Setting Client List and Restore Permissions” on page 217.

Denies list and restore requests for all clients. When this option is present, clients cannot list or restore files that they have backed up through this master server. By default, this option is not present and clients can list and restore their files.

You can add this option to the

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

file on NetBackup master servers.

DISALLOW_CLIENT_RESTORE

Note You can override the `DISALLOW_CLIENT_RESTORE` option for individual clients by changing their `list_restore` setting. See “Setting Client List and Restore Permissions” on page 217.

Denies restore requests for all clients. When this option is present, clients cannot restore files that they have backed up through this master server. By default, this option is not present and clients can restore their files.

You can add this option to the

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

file on NetBackup master servers.

FAILOVER_RESTORE_MEDIA_SERVERS

Specifies automatic failover to another NetBackup server if a server is temporarily inaccessible for a restore. This failover does not require administrator intervention. By default, NetBackup does not perform automatic failover. The format for the entry is:

```
FAILOVER_RESTORE_MEDIA_SERVERS = failed_host host1 host2 ... hostN
```

Where:

failed_host is the server that is not operational.

host1 ... *hostN* are the servers that provide failover capabilities.

When automatic failover is necessary for a given server, NetBackup searches from left to right through the associated `FAILOVER_RESTORE_MEDIA_SERVERS` list to find until it finds one that is 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.

You can add this option to the

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

file on the NetBackup master server.

After adding the `FAILOVER_RESTORE_MEDIA_SERVERS` entry, stop and restart the NetBackup Request daemon on the master server where you are changing the configuration. Also, see “Server Independent Restores” on page 243.

FORCE_RESTORE_MEDIA_SERVER

Forces restores to go to a specific server, regardless of where the files were backed up. The format for the entry is:

```
FORCE_RESTORE_MEDIA_SERVER = fromhost tohost
```

where *fromhost* is the server that performed the original backup and *tohost* is the server to use for the restore.



You can add this option to the

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

file on the NetBackup master server.

After adding the `FORCE_RESTORE_MEDIA_SERVER` entry, stop and restart the NetBackup Request daemon on the master server. Before attempting a restore, physically move the media to *tohost* and update the Media Manager volume database to reflect the move.

This setting applies to all storage units on the original server. Restores for any storage unit on *fromhost* will go to *tohost*. To revert to the original configuration for future restores, delete the entry. Also, see “Server Independent Restores” on page 243.

INITIAL_BROWSE_SEARCH_LIMIT

Specifies the number of days back that NetBackup searches for files to restore. The value is in days. For example, to limit the browse range to the seven days prior to the current date specify:

```
INITIAL_BROWSE_SEARCH_LIMIT = 7
```

This option can be specified on the master server and applies to all NetBackup clients. It can also be specified on a UNIX client. When specified on a UNIX client, it 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 class, then the browse starts with the earliest of the set of last-full backups.

You can add this option to the

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

file on master servers (and on clients).

KNOWN_MASTER

Specifies the master servers that can be administered by a master of masters. You add this option to the

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

file on a NetBackup server that has the Global Data Manager option installed and is acting as a master of masters.

For example, assume that Global Data Manager software is installed on a UNIX server named `omega`. Also assume that `omega` will be the master of masters for two existing UNIX master servers named `alpha` and `beta`. In this instance, you add the following entry to the `bp.conf` file on `omega`:

```
KNOWN_MASTER = alpha beta
```

In addition, in the `bp.conf` file on `alpha` and `beta`, you add a `SERVER` entry and a `MASTER_OF_MASTERS` entry for `omega`. For example the `bp.conf` file on `alpha` must have at least the following entries

```
SERVER = alpha
```

```
SERVER = omega (the entry for omega must follow the alpha entry)
```

```
MASTER_OF_MASTERS = omega
```

Beta has equivalent `bp.conf` entries.

The `SERVER` entry allows `omega` to communicate with `alpha` and `beta`. The `MASTER_OF_MASTERS` entry grants `omega` permission to administer `alpha` and `beta`.

LIMIT_BANDWIDTH

Note Read “Bandwidth Limiting” on page 293 before setting this option.

Specifies a limit for the network bandwidth used by one or more NetBackup clients on a network. The actual limiting occurs on the client side of the backup connection. This feature limits only backups. Restores are unaffected.

Each `LIMIT_BANDWIDTH` entry specifies the bandwidth value and the IP address of the clients and networks to which it applies. The syntax is:

```
LIMIT_BANDWIDTH = xxx.xxx.xxx.xxx yyy.yyy.yyy.yyy zzz
```

```
LIMIT_BANDWIDTH = xxx.xxx.xxx.xxx yyy.yyy.yyy.yyy zzz
```

For example: `10.0.0.2 10.0.0.49 200`

Where:

- ◆ `xxx.xxx.xxx.xxx` is the beginning of the IP address range.
- ◆ `yyy.yyy.yyy.yyy` is the end of the IP address range.
- ◆ `zzz` is the bandwidth limitation in kilobytes per second. A value of 0 disables throttling for the individual client or the range of IP addresses covered by this entry.

You can add `LIMIT_BANDWIDTH` entries to the

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

file on NetBackup master servers.



By default, the bandwidth is not limited.

MASTER_OF_MASTERS

Specifies the master of masters servers that can administer this server. You add add this option to the

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

file on on a NetBackup server that will be administered by a master of masters.

For example, assume that a master server named alpha is to be administered by a master of masters named omega. In this instance, you add the following entry to the `bp.conf` file on alpha:

```
MASTER_OF_MASTERS = omega
```

In addition, you add a `SERVER` entry for omega so alpha's `bp.conf` file has at least the following entries:

```
SERVER = alpha
```

```
SERVER = omega (the entry for omega must follow the alpha entry)
```

The `SERVER` entry allows omega to communicate with alpha. The `MASTER_OF_MASTERS` entry grants omega permission to administer alpha.

To complete the configuration, you add a `KNOWN_MASTER` entry for alpha to the `bp.conf` file on omega.

MAX_APOLLO_RESTORE_ARG_CHARS

Note This option applies only to Apollo clients. It has a reasonable default and will have to be changed only under very specific circumstances.

Specifies the maximum number of characters to allow on an `rbak` command on Apollo clients.

You can add this option to the

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

file on NetBackup servers.

By default, the maximum number of characters is 9000.

MEDIA_ID_PREFIX

Applies to media in nonrobotic drives and specifies the media ID prefix that is used to create media IDs when unlabeled media is found in a nonrobotic drive. The prefix must be one to three alpha-numeric characters. NetBackup appends remaining numeric characters.

The following is an example entry:

```
MEDIA_ID_PREFIX = FEB
```

NetBackup appends remaining numeric characters so the assigned media IDs become FEB000, FEB001, and so on.

The default media ID prefix is:

```
MEDIA_ID_PREFIX = A
```

For the default, NetBackup assigns A00000, then A00001, and so on.

You can add this option to the

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

file on NetBackup servers.

MEDIA_UNMOUNT_DELAY

Applies only to user operations (including backups and restores of database-extension clients, such as those running NetBackup for Oracle). When you specify `MEDIA_UNMOUNT_DELAY`, the media unload is delayed for the specified number of seconds after the requested operation has completed. This 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 (the default is 180 seconds). If you specify 0, the media unmount occurs immediately upon completion of the requested operation. The maximum delay is 1800 seconds. Values greater than 1800 are set to 1800.

You can add this option to the

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

file on NetBackup servers.

MEDIA_REQUEST_DELAY

Applies only to nonrobotic drives and specifies the number of seconds that NetBackup waits for a drive to become ready. 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.



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 the delay is 150 seconds:

```
MEDIA_REQUEST_DELAY = 150
```

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 before looking for another drive. 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.

By default, the delay is 0 seconds.

You can add this option to the

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

file on NetBackup servers.

MPX_RESTORE_DELAY

Applies to multiplexed restores and specifies how long (in 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). By default, the delay is 30 seconds.

You can add this option to the

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

file on NetBackup servers.

The following is an example entry:

```
MPX_RESTORE_DELAY = 60
```

MUST_USE_LOCAL_DRIVE

If the client is a server and this entry is present, backups for this client must occur on a local drive. If a client is not a server, this entry has no effect.

You can add this option to the

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

file on NetBackup servers.

QUEUE_ON_ERROR

Causes jobs to enter the queued state on startup, if the required storage unit is not available. The jobs will then run when the storage unit becomes available. If this entry is not present, the job fails with a 219 status. By default, this option is not present and jobs fail with a status code 219 if the storage unit is not available. You can add this option to the

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

file on NetBackup master servers. The following is an example entry:

```
QUEUE_ON_ERROR
```

This entry requires that the `WAIT_IN_QUEUE` entry also exist or the job will fail immediately anyway with a 219 status if the storage unit is not available. Also see “Automatic Drive Availability Checking” on page 36.

RANDOM_PORTS

Specifies whether NetBackup chooses port numbers randomly or sequentially when it requires one for communication with NetBackup on other computers.

- ◆ If `RANDOM_PORTS = YES` (default), NetBackup chooses port numbers randomly 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.
- ◆ If `RANDOM_PORTS = NO`, NetBackup chooses numbers sequentially, starting with highest number that is available in the allowed range. For example, if the range is from 1023 through 5000, NetBackup chooses 5000 (assuming it is free). If 5000 is being used, port 4999 is chosen.

You can add this option to the

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

file on NetBackup servers and clients.

By default, this option is not present and NetBackup uses the random method for selecting port numbers.

RE_READ_INTERVAL

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. The default value is 300 seconds (5 minutes). You can add this option to the

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

file on NetBackup master servers. The following is an example entry:



```
RE_READ_INTERVAL = 350
```

Also see “Automatic Drive Availability Checking” on page 36.

REQUIRED_INTERFACE

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 and, by default, the operating system determines the one to use. To force NetBackup connections to be through a specific network interface, use this entry to specify the network host name of that interface. For example:

```
REQUIRED_INTERFACE = host1
```

In this example, `host1` is the network host name of the interface.

You can add this option to the

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

file on a NetBackup client or server. By default, the entry does not exist and the operating system determines the interface to use.

Example 1 - Client with multiple network interfaces

Assume you have a NetBackup client with two network interfaces. One is for the regular network and one is for the backup network:

- ◆ The host name for the regular interface is `fred`
- ◆ 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 users on `fred` start a backup, restore, or list operation, the request ideally always goes out on the `fred_nb` interface and over the backup network. This assumes that `fred` and the network are set up for this. However, 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 alternate client restores for `barney`. This allows the server to accept the request, but leaves NetBackup traffic on the regular network.

A better solution is to add the following entry to the `bp.conf` file on `fred`:

```
REQUIRED_INTERFACE = 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 you have a NetBackup server with two network interfaces. One is for the regular network and one is for the backup network:

- ◆ The host name for the regular interface is barney
- ◆ The host name for the backup interface is barney_nb

The `bp.conf` file on all NetBackup servers and clients have a `SERVER = barney_nb` entry.

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 for this. However, 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 a `SERVER = barney` entry to the `bp.conf` file on the client. The client now accepts requests from barney, but NetBackup traffic is still on the regular network.

A better solution is to add the following entry to the `bp.conf` file on barney:

```
REQUIRED_INTERFACE = barney_nb
```

Now, when barney connects to a client, the connection is always through the `barney_nb` interface and everything works as intended.

SERVER

During installation, NetBackup sets `SERVER` to the name of the system where you are installing the server software. NetBackup uses the `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.

If you configure media servers, the `bp.conf` file must have a `SERVER` entry for each of them. The first entry in the list designates the master server to which requests are directed. The remaining entries are for media servers.

```
SERVER = master_server
```

```
SERVER = media_server
```

```
.  
.
.
```

The `SERVER` entries must be the same on all servers in a master and media server cluster.



If you modify or add a `SERVER` entry in the `bp.conf` file on the master server, stop and restart both the NetBackup request daemon (`bprd`) and NetBackup database manager (`bpdbm`) so NetBackup will recognize the change.

The `SERVER` option *must* be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX servers and clients. It is also the only required entry in these `bp.conf` files. This option is not used in `$HOME/bp.conf` files on a client.

SERVER_PORT_WINDOW

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 (for information on client configuration, see “`ALLOW_NON_RESERVED_PORTS`” on page 344).

You can add this option to the

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

files on NetBackup servers. It can also be useful on clients that are running the NetBackup-Java application server.

Example:

```
SERVER_PORT_WINDOW = 4900 5000
```

permits ports from 4900 through 5000.

The default range is 1025 through 5000. Note that if you specify 0 for the first number, the operating system determines the nonreserved port to use.

SERVER_RESERVED_PORT_WINDOW

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 (for information on client configuration, see “`ALLOW_NON_RESERVED_PORTS`” on page 344).

You can add this option to the

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

files on NetBackup servers. This entry is usually not useful on clients.

For example:

```
SERVER_RESERVED_PORT_WINDOW = 900 1023
```

permits ports from 900 through 1023.

The default range is 512 through 1023. Note that if you specify 0 for the first number, a nonreserved port is used instead and is chosen by the operating system.

SLAVE_CONNECT_TIMEOUT

Specifies the number of seconds that the master server waits before timing out when connecting to a media server. By default, the timeout period is 30 seconds.

You can add this option to the

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

file on NetBackup master servers.

The following is an example entry:

```
SLAVE_CONNECT_TIMEOUT = 60
```

TIMEOUT_IN_QUEUE

Determines how long a job can be requeued while NetBackup waits for a required storage unit if it is currently unavailable. The default is 36000 seconds (10 hours). You can add this option to the

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

file on NetBackup master servers. The following is an example entry:

```
TIMEOUT_IN_QUEUE = 30000
```

Also see “Automatic Drive Availability Checking” on page 36.

VERBOSE

Causes NetBackup to include more information in its logs. By default, verbose logging is disabled. You can add this option to the `/usr/opensv/netbackup/bp.conf` file on NetBackup servers.

WAIT_IN_QUEUE

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 the storage unit becomes available. A job fails if the `TIMEOUT_IN_QUEUE` time expires or its backup window closes before the storage unit becomes available. By default, this option is not present and the job is not requeued. You can add this option to the

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

file on NetBackup master servers. The following is an example entry:

```
WAIT_IN_QUEUE
```

Also see “Automatic Drive Availability Checking” on page 36.



bp.conf File on UNIX Clients

On NetBackup UNIX clients, the main `bp.conf` file is located in

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

As installed, NetBackup uses internal software defaults for all options in the `bp.conf` file, except `SERVER`. During installation, NetBackup sets the `SERVER` option to the name of the master server where the software is installed.

Note The `SERVER` option must be in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX clients. It is also the only required entry in this file.

If a single UNIX system is running as both a client and a server, both the server and client options are in the `/usr/opensv/netbackup/bp.conf` file.

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

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

The options in personal `bp.conf` files apply only to user operations. During a user operation, NetBackup checks the `$HOME/bp.conf` file before `/usr/opensv/netbackup/bp.conf`. Root users do not have personal `bp.conf` files. NetBackup uses the `/usr/opensv/netbackup/bp.conf` file for root users.

The following topics describe the options that you can specify in the `/usr/opensv/netbackup/bp.conf` and `$HOME/bp.conf` files on a NetBackup UNIX client.

Note PC clients provide similar options that you can change either through the client-user interface or in a configuration file, depending on the client. For instructions, see the NetBackup user's guide for the client.

ALLOW_NON_RESERVED_PORTS

Specifies that the NetBackup client daemon (`bpcd`) can accept remote connections from non-privileged ports (port numbers 1024 or greater). If this entry is not present, then `bpcd` requires remote connections to come from privileged ports (port numbers less than 1024). This option can be useful when NetBackup clients and servers are on opposite sides of a firewall.

You can add this option to the

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

file on a UNIX client.

In addition to adding `ALLOW_NON_RESERVED_PORTS` to the client, execute the following commands as root on the master server.

```
cd /usr/opensv/netbackup/bin/admincmd
./bpclient -client client_name -add -connect_nr_port 1
```

Where *client_name* is the name of the client where you added the `ALLOW_NON_RESERVED_PORTS` option. These commands instruct the master server to use nonprivileged ports. On the master server, specify the range of open ports with the `SERVER_PORT_WINDOW` `bp.conf` entry.

BARCHIVE_CLASS

Specifies the name of the class to use for user archives. For example

```
BARCHIVE_CLASS = arch_1
```

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and $HOME/bp.conf
```

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BARCHIVE_CLASS` is not in any `bp.conf` file and NetBackup uses the first class that it finds that has the client and a user archive schedule.

BARCHIVE_SCHED

Specifies the name of the schedule for user archives. For example

```
BARCHIVE_SCHED = user_arch1
```

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and $HOME/bp.conf
```

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BARCHIVE_SCHED`, is not in any `bp.conf` file and NetBackup uses the first archive schedule in the first class that it finds that has this client.

BPBACKUP_CLASS

Specifies the name of the class name to use for user backups. For example,

```
BPBACKUP_CLASS = userback_1
```

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and (or) $HOME/bp.conf
```



files on a UNIX client.

The value in user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BPBACKUP_CLASS`, is not in any `bp.conf` file and NetBackup uses the first class it finds that has both the client and a user backup schedule.

BPBACKUP_SCHED

Specifies the name of the schedule to use for user backups.

```
BPBACKUP_SCHED = user_back1
```

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and $HOME/bp.conf
```

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BPBACKUP_SCHED`, is not in any `bp.conf` file and NetBackup uses the first class it finds that has both the client and a user backup schedule.

BUSY_FILE_ACTION

Note Does not apply to Apollo clients.

Directs the action that NetBackup performs on busy files when busy-file processing is enabled (see `BUSY_FILE_PROCESSING`).

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and $HOME/bp.conf
```

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists.

There can be multiple entries of the form:

```
BUSY_FILE_ACTION = filename_template action_template
```

Where

- ◆ *filename_template* is the absolute pathname and file name of the busy file. The shell language metacharacters `*`, `?`, `[]`, `[-]` can be used for pattern matching of filenames or parts of filenames.

- ◆ *action_template* is one of:

```
MAIL | mail
```

Directs NetBackup to mail a busy file notification message to the user specified by the `BUSY_FILE_NOTIFY_USER` option.



REPEAT | repeat [repeat_count]

Directs NetBackup to retry the backup on the specified busy file. A repeat count can be specified to control the number of backup attempts. The default repeat count is 1.

IGNORE | ignore

Directs NetBackup to exclude the busy file from busy file processing.

BUSY_FILE_DIRECTORY

Note Does not apply to Apollo clients.

The `BUSY_FILE_DIRECTORY` option specifies the path to the busy-files working directory when busy-file processing is enabled (see `BUSY_FILE_PROCESSING`).

You can add this option to the

`/usr/opensv/netbackup/bp.conf` and `$HOME/bp.conf`

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `BUSY_FILE_DIRECTORY` is not in any `bp.conf` file and NetBackup creates the `busy_files` directory in `/usr/opensv/netbackup`.

BUSY_FILE_NOTIFY_USER

Note Does not apply to Apollo clients.

The `BUSY_FILE_NOTIFY_USER` option specifies the recipient of the busy file notification message when `BUSY_FILE_ACTION` is set to `MAIL` or `mail`.

You can add this option to the

`/usr/opensv/netbackup/bp.conf` and `$HOME/bp.conf`

files 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`.

BUSY_FILE_PROCESSING

Note Does not apply to Apollo clients.



The `BUSY_FILE_PROCESSING` option lets the user control the actions that NetBackup performs when it determines that a file is changing while it is being backed up. See “Busy-File Processing (UNIX Clients Only)” on page 297 for instructions on setting this option.

You can add this option to the

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

file on a UNIX client.

By default, `BUSY_FILE_PROCESSING` option is not in `bp.conf` and busy-file processing does not occur.

CLIENT_NAME

Specifies the name of the client as it is known to NetBackup. There can be one `CLIENT_NAME` entry and it must match the name used in the class that is backing up the client. The only exception is for an alternate client restore, where the name must match that of the client whose files are being restored (see “Allowing Restores to an Alternate Client” on page 211). The client installation procedures automatically set `CLIENT_NAME` to the value specified on the `ftp_to_client` or `install_client` command in the installation scripts.

You can add this option to the

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

file on a UNIX client.

It can also be added to a `$HOME/bp.conf` file on a UNIX client but this is normally done only for alternate-client restores.

If the value is not in any `bp.conf` file, NetBackup uses the value returned by the `gethostname()` library function.

CLIENT_PORT_WINDOW

Specifies the range of nonreserved ports on this computer that are used for connecting to NetBackup on other computers. See “`CLIENT_PORT_WINDOW`” on page 330 for a description.

CLIENT_READ_TIMEOUT

Specifies the number of seconds for the client-read timeout on a server or a database-extension client. See “`CLIENT_READ_TIMEOUT`” on page 330 for a description.



CLIENT_RESERVED_PORT_WINDOW

Specifies the range of reserved ports on this computer that are used for connecting to NetBackup on other computers. See “CLIENT_RESERVED_PORT_WINDOW” on page 331 for a description.

COMPRESS_SUFFIX

Note This option has a reasonable default and has to be changed only if problems are encountered. This option does not apply to Apollo clients.

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]

You can add this option to the

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

file on a UNIX client.

By default, COMPRESS_SUFFIX is not in the `bp.conf` file. See “Compression” on page 67 for more information on compressing files.

CRYPT_OPTION

Note CRYPT_OPTION applies only to clients that have the NetBackup Encryption option installed. See the *NetBackup Encryption System Administrator’s Guide*.

CRYPT_OPTION specifies the encryption options on NetBackup clients. NetBackup creates this entry automatically in the

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

file on a UNIX client when you run the `bpinst_crypt` command on the NetBackup master server.

Do not alter the entry or create it manually unless it has been accidentally deleted. The allowable values are:

DENIED|denied



Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, it is considered an error. This option is the default for a client that has not been configured for encryption.

ALLOWED | allowed

Specifies that the client allows either encrypted or unencrypted backups.

REQUIRED | required

Specifies that the client requires encrypted backups. If this value is specified and the server requests an unencrypted backup, it is considered an error.

CRYPT_STRENGTH

Note CRYPT_STRENGTH applies only to clients that have the NetBackup Encryption option installed. See the *NetBackup Encryption System Administrator's Guide*.

Specifies the encryption strength on NetBackup clients. NetBackup creates this entry automatically in the

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

file on a UNIX client when you run the `bpinst_crypt` command on the NetBackup master server.

Do not alter the entry or create it manually unless it has been accidentally deleted. The possible values are:

DES_40 | des_40

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

DES_56 | des_56

Specifies 56-bit DES encryption.

CRYPT_LIBPATH

Note CRYPT_LIBPATH applies only to clients that have the NetBackup Encryption option installed. See the *NetBackup Encryption System Administrator's Guide*.

Specifies the directory that contains the encryption libraries for NetBackup clients. NetBackup creates this entry automatically in the

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

file on a UNIX client when you run the `bpinst_crypt` command on the NetBackup master server.



Do not alter the entry or create it manually unless it has been accidentally deleted.

- ◆ The default value on UNIX systems is:

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

- ◆ The default value on Windows 2000, NT, 98, and 95 systems is:

```
install_path\bin\
```

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

- ◆ The default value on Macintosh systems is:

```
:System Folder:Extensions:
```

CRYPT_KEYFILE

Note CRYPT_KEYFILE applies only to clients that have the NetBackup Encryption option installed. See the *NetBackup Encryption System Administrator's Guide*.

Specifies the file that contains the encryption keys on NetBackup clients. NetBackup creates this entry automatically in the

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

file on a UNIX client when you run the `bpinst_crypt` command on the NetBackup master server.

Do not alter the entry or create it manually unless it has been accidentally deleted. The default values are:

- ◆ On UNIX systems:

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

- ◆ On Windows 2000, NT, 98, and 95 systems:

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

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

- ◆ On Macintosh systems:

```
:System Folder:Preferences:NetBackup:keyfile
```

DISALLOW_SERVER_FILE_WRITES

Prevents the NetBackup server from creating files on the NetBackup client. For example, this prevents server-directed restores or server-directed updates of the `bp.conf` file on the client.



You can add this option to the

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

file on a UNIX client.

By default, server writes are allowed.

DO_NOT_RESET_FILE_ACCESS_TIME

Note `DO_NOT_RESET_FILE_ACCESS_TIME` does not apply to Apollo clients.

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.

Specifies that if a file is backed up its access time (atime) will show the time of the backup. By default, NetBackup preserves the access time by resetting it to the value it had before the backup.

You can add this option to the

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

file on a UNIX client.

INFORMIX_HOME

Specifies the path to the Informix home directory and is required when the client is using NetBackup for Informix.

You must add this option to the

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

file on UNIX clients that are running NetBackup for Informix.

INITIAL_BROWSE_SEARCH_LIMIT

Reduces the default number of days back that NetBackup searches for files to restore. You can add this option to the

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

file on a UNIX client. See “INITIAL_BROWSE_SEARCH_LIMIT” on page 334 for an explanation of this option.

KEEP_DATABASE_COMM_FILE

Causes NetBackup to keep database-extension client logs for seven days.

You can add this option to the

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

file on a UNIX database-extension client (for example, a client that is running NetBackup for Informix).

By default, NetBackup keeps database-extension client logs for only one day.

KEEP_LOGS_DAYS

Specifies the number of days to keep job and progress logs generated by the NetBackup Java program, Backup, Archive, and Restore. NetBackup writes these files in 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 Backup, Archive, and Restore program.

You can add this option to the

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

file on a UNIX client.

The default is three days.

LIST_FILES_TIMEOUT

Specifies the number of minutes to wait for a response from the NetBackup server when listing files by using the client-user interface or `bplist`. If this time is exceeded, the user receives a `socket read failed` error even if the server is still processing the user's request.

You can add this option to the

```
/usr/opensv/netbackup/bp.conf and $HOME/bp.conf
```

files on a UNIX client.

The value in the user's `$HOME/bp.conf` file takes precedence if it exists. By default, `LIST_FILES_TIMEOUT` is not in any `bp.conf` file and NetBackup uses a value of 30 minutes.

LOCKED_FILE_ACTION

Note Does not apply to Apollo clients.



Specifies the behavior of NetBackup when it tries to back up a file that has mandatory file locking enabled in its file mode (see `chmod(1)`). If `LOCKED_FILE_ACTION` is specified and has a value of `SKIP` (the only legal value), NetBackup skips files that currently have mandatory locking set by another process and logs a message to this effect.

You can add this option to the

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

files on a UNIX client.

By default, NetBackup waits for files to become unlocked.

MEGABYTES_OF_MEMORY

Note This option does not apply to Apollo clients. This option has a reasonable default and has to be changed only if problems are encountered.

Specifies how much memory is 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 the compress code, the greater the compression. The percentage of machine resources used is also greater. 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.

You can add this option to the

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

file on a UNIX client.

By default, NetBackup assumes a value of one megabyte.

NFS_ACCESS_TIMEOUT

Specifies the number of seconds that the backup process waits when processing an NFS mount table before considering an NFS file system unavailable.

You can add this option to the

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

file on a UNIX client.

By default, the timeout period is five seconds.

RANDOM_PORTS

Specifies whether NetBackup chooses port numbers randomly or sequentially when it requires one for communication with NetBackup on other computers. For a description, see “RANDOM_PORTS” on page 129.

RESTORE_RETRIES

Note This option has a reasonable default and will have to be changed only if problems are encountered.

Specifies the number of times to retry a restore after a failure.

You can add this option to the

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

file on a UNIX client. By default, there are no retries.

REQUIRED_INTERFACE

Specifies the network interface that NetBackup uses when connecting to another NetBackup client or server. For more information, see “REQUIRED_INTERFACE” on page 130.

SERVER_PORT_WINDOW

Specifies the range of nonreserved ports on which this computer accepts connections from NetBackup on other computers.

SERVER

Defines the list of servers that can access the client and also the server to which the client must connect in order to list and restore files. During installation, NetBackup sets `SERVER` to the name of the server where you are installing the software. The `SERVER` option must be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX clients. It is also the only required entry in this `bp.conf` file. This option is not used in a `$HOME/bp.conf` file.

If you configure media servers, you must have a `SERVER` entry for each of them. The first entry must be for the master server, which is the server to which requests are directed. The remaining entries are for media servers. The `SERVER` entries must be the same on all servers in a master and media server cluster.

Note On NetBackup UNIX servers, the `SERVER` entry applies to the both the client and the server.



SYBASE_HOME

Specifies the path to the Sybase home directory and is required when using NetBackup for Sybase to back up Sybase databases.

You must add this option to the

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

file on a NetBackup for Sybase client. By default, SYBASE_HOME is not in the bp.conf file.

USE_CTIME_FOR_INCREMENTALS

Note If you specify USE_CTIME_FOR_INCREMENTALS, you must also specify DO_NOT_RESET_FILE_ACCESS_TIME.

DO NOT use these options 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.

Causes NetBackup client software to use both modification time (mtime) and inode change time (ctime) during incremental backups to determine if a file has changed.

You can add this option to the

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

file on a UNIX client.

By default, NetBackup uses only mtime.

USEMAIL

Specifies the E-mail address where NetBackup sends status on the outcome of operations for a UNIX client.

Note You can use multiple addresses or a mail alias as long as there are no blanks or white space between them.

You can add this option to the /usr/opensv/netbackup/bp.conf and \$HOME/bp.conf files on a UNIX client.

- ◆ If the /usr/opensv/netbackup/bp.conf file specifies an address, NetBackup sends automatic backup and manual backup status to that address.
- ◆ If the \$HOME/bp.conf file specifies an address, NetBackup also sends status on the success or failure of user operations to that address.

By default, `USEMAIL` is not present in any `bp.conf` file and no E-mail is sent.

VERBOSE

Causes NetBackup to include more information in its logs.

You can add this option to the `/usr/opensv/netbackup/bp.conf` file on a UNIX client.

By default, verbose logging is disabled.

Examples - UNIX Client

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

```
SERVER = hare
CLIENT_NAME = freddie
USEMAIL = abc@bdev.com
COMPRESS_SUFFIX = .Addrs
COMPRESS_SUFFIX = .Counts
VERBOSE
RESTORE_RETRIES = 1
BPBACKUP_CLASS = U1userdir
BPBACKUP_SCHED = userbackups
BPARCHIVE_CLASS = U1userdir
BPARCHIVE_SCHED = userarchives
LOCKED_FILE_ACTION = SKIP
```

Example `$HOME/bp.conf` File

Nonroot users on UNIX clients can have a personal `bp.conf` file in their home directory. A personal `bp.conf` file can have any of the following options

Note A root user, cannot have a personal `bp.conf` file. For root users, NetBackup uses the `/usr/opensv/netbackup/bp.conf` file.

```
USEMAIL = mars@bdev.com
BPBACKUP_CLASS = user1
BPBACKUP_SCHED = userback
```



```
BPARCHIVE_CLASS = user1
BPARCHIVE_SCHED = userarch
LIST_FILES_TIMEOUT = 10
CLIENT_NAME
```

(specify `CLIENT_NAME` only when doing restores to an alternate client as described in “Examples of Restoring Files to an Alternate Client” on page 213)



This section describes man pages specific to the NetBackup product.

The following are special conventions used in the command descriptions.

- ◆ Brackets [] mean that the enclosed command line component is optional.
- ◆ A vertical bar (or pipe) symbol | separates optional arguments from which the user can choose. For example, assume that a command has the following format:

```
command [arg1 | arg2]
```

Here, the user can choose either arg1 or arg2 (but not both).

- ◆ Italics indicate that the information is user supplied. For example, the user supplies *class*, *schedule*, and *filename* in the following command:

```
bpbackup -c class -s schedule filename
```

- ◆ An ellipses (...) means that you can repeat the previous parameter. For example, consider the following command:

```
bpbackup [-S master_server [,master_server,...]] filename
```

Here, the -S option requires the first master server name. Additional names can be added, separated by commas and followed by a file name as in:

```
bpbackup -S mars,coyote,shark,minnow memofile.doc
```



bp(1)

NAME

bp - start the NetBackup menu interface for users

SYNOPSIS

```
/usr/opensv/netbackup/bin/bp [-a | -ra | -b | -r | -rr | -o | -ro  
    | -s | -rs | -i | -ri | -k | -rk | -rti | -p | -rp | -2  
    | -r2 | -n | -rn] [-verbose]  
  
/usr/opensv/netbackup/bin/bp [ -b | -a | -r | -ra] [-verbose]
```

DESCRIPTION

The `bp` command starts a menu interface that lets users archive, back up, and restore files, directories, or raw partitions from their client workstations. This interface can be run from any character-based terminal (or terminal emulation window) for which the user has a `termcap` or `terminfo` definition.

The first form of the command above applies to all except Apollo clients. The second form applies to Apollo clients (note Apollo clients are supported only by NetBackup DataCenter master servers).

The *NetBackup User's Guide - UNIX* and the `bp` online help provide detailed operating instructions.

OPTIONS

The menu that appears at startup depends on the option used with the `bp` command. Executing the `bp` command without specifying an option starts the utility at the main menu. To start the utility at a secondary menu, specify one of the following options:

- a Starts `bp` in the Archive of Files and Directories menu.
- ra Starts `bp` in the Restore Archives menu.
- b Starts `bp` in the Backup of Files and Directories menu.
- r Starts `bp` in the Restore Backups menu.
- rr Starts `bp` in the Restore Raw Partitions Backups menu.

- o Starts `bp` in the Backup Oracle DB menu.
- ro Starts `bp` in the Restore Oracle DB menu.

- s Starts `bp` in the Backup Sybase DB menu.



-
- rs Starts bp in the Restore Sybase DB menu.
 - i Starts bp in the Backup Informix DB menu.
 - ri Starts bp in the Restore Informix DB menu.
 - rti Starts bp in the Restore True Image Backups menu.

Note The following options for SAP, DB2, and SQL-BackTrack apply only to NetBackup DataCenter.

- p Starts bp in the Backup SAP DB menu.
- rp Starts bp in the Restore SAP DB menu.
- 2 Starts bp in the Backup DB2 DB menu.
- r2 Starts bp in the Restore DB2 DB menu.
- k Starts bp in the Backup SQL-BackTrack DB menu.
- rk Starts bp in the Restore SQL-BackTrack DB menu.
- verbose Provides a verbose response.

FILES

/usr/opensv/netbackup/help/bp/*
 /usr/opensv/netbackup/logs/bp/*
 /usr/opensv/netbackup/bp.conf

SEE ALSO

bparchive(1), bpbackup(1), bplist(1), bprestore(1)



bpadm(1M)

NAME

bpadm - start the NetBackup menu interface for administrators

SYNOPSIS

```
/usr/opensv/netbackup/bin/bpadm
```

DESCRIPTION

The `bpadm` utility has a menu interface that an administrator can use to configure NetBackup and monitor its operations. `bpadm` requires root privileges. This interface can be used from any character-based terminal (or terminal emulation window) for which the administrator has a `termcap` or `terminfo` definition.

See your NetBackup system administrator's guide and the `bpadm` online help for detailed operating instructions.

FILES

```
/usr/opensv/netbackup/help/bpadm/*
```

```
/usr/opensv/netbackup/logs/admin/*
```

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

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

SEE ALSO

`bprd(1M)`



bparchive(1)

NAME

bparchive - archive files to the NetBackup server

SYNOPSIS

```
/usr/opensv/netbackup/bin/bparchive [-c class] [-s schedule] [-S
  master_server [ , master_server, ... ] [-t class_type] [-L
  progress_log] [-w [ hh:mm:ss ] [-help] [-k
  "keyword_phrase" ] -f listfile | filenames
```

DESCRIPTION

bparchive processes files that you list on the command line or in the file you specify with the `-f listfile` option. Any file path entered can be a file or directory name. If the list of files includes a directory, bparchive archives all files and subdirectories of that directory starting at the directory itself.

By default, you are returned to the system prompt after bparchive is successfully submitted. The command works in the background and does not return completion status directly to you. The `-w` option lets you change this behavior so bparchive works in the foreground and returns completion status after a specified time period.

bparchive writes its informative and error messages to a progress-log file. You must create this file prior to the execution of the bparchive command and then specify it with the `-L progress_log` option. If bparchive cannot archive any of the requested files or directories, you can use the progress log to determine the reason for the failure.

If you create a `/usr/opensv/netbackup/logs/bparchive/` directory with public-write access, bparchive creates an activity log file in this directory that you can use for troubleshooting.

In addition, if a nonroot user specifies `USEMAIL = mail_address` in their `$HOME/bp.conf` file or a root user specifies it in the `/usr/opensv/netbackup/bp.conf` file, NetBackup sends mail on the archive completion status to *mail_address*. This message is sent when the archive process is complete.

The following restrictions apply to this command:

- ◆ To archive files with the bparchive command, you must be root to delete the file and the file must not be read-only. Otherwise, NetBackup saves the files but cannot reset their access time (utime) and does not delete them from the disk.
- ◆ If you specify a UNIX file that is a link, bparchive archives only the link itself, not the file to which it links.



- ◆ bparchive does not archive the "." or ".." directory entries, and also does not archive raw partitions.

OPTIONS

- c *class* Names the class to use for the user archive. If it is not specified, the NetBackup server uses the first class it finds that includes the client and a user archive schedule.
- s *schedule*
Names the schedule to use for the user archive. If it is not specified, the NetBackup server uses the first user archive schedule it finds in the class it is using (see the -c option).
- S *master_server* [, *master_server*, . . .]
Specifies the name of the NetBackup master server. The default is the first SERVER entry in the /usr/opensv/netbackup/bp.conf file.
- t *class_type*
Specifies one of the following numbers corresponding to the class type (the default is 0 on all clients except Apollo, where the default is 3):
 - 0 = Standard
 - 4 = Oracle
 - 6 = Informix-On-BAR
 - 7 = Sybase
 - 10 = NetWare
 - 13 = MS-Windows-NT
 - 14 = OS/2
 - 15 = MS-SQL-Server
 - 16 = MS-Exchange-Server
 - 19 = NDMP

Note The following class types (Apollo-wbak, DataTools-SQL-BackTrack, Auspex-FastBackup, SAP, DB2, FlashBackup, Split-Mirror, and AFS) apply only to NetBackup DataCenter.

- 3 = Apollo-wbak
- 11 = DataTools-SQL-BackTrack
- 12 = Auspex-FastBackup
- 17 = SAP
- 18 = DB2
- 20 = FlashBackup
- 21 = Split-Mirror

22 = AFS

- L *progress_log*
Specifies the name of an existing file in which to write progress information. The file name must begin with /.
For example: /home/tlc/proglog.
The default is to not use a progress log.
- w [*hh:mm:ss*]
Causes NetBackup to wait for a completion status from the server before returning you to the system prompt.

Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section later in this command description.

You can optionally specify a wait time in hours, minutes, and seconds. The maximum wait time you can specify is 23:59:59. If the wait time expires before the archive is complete, the command exits with a timeout status. The archive, however, still completes on the server.

If you use `-w` without specifying the wait time or if you specify a value of 0, NetBackup waits indefinitely for the completion status.

- help Displays a synopsis of command usage when it is the only option on the command line.
- k *keyword_phrase*
Specifies a keyword phrase that NetBackup associates with the image created by this archive operation. You can then restore the image by specifying the keyword phrase with the `-k` option on the `bprestore` command.

The keyword phrase is a textual description of the archive that is a maximum of 128 characters in length. All printable characters are permitted including space (" ") and period ("."). Enclose the phrase in double quotes ("...") or single quotes ('...') to avoid conflict with the UNIX shell.

The default keyword phrase is the null (empty) string.
- f *listfile* Specifies a file (*listfile*) containing a list of files to be archived and can be used instead of the *filenames* option. In *listfile*, place each file path on a separate line.

The format required for the file list depends on whether the files have spaces or newlines in the names.

To archive files that do not have spaces or newlines in the names, use this format:

filepath



Where *filepath* is the path to the file you are archiving. For example:

```
/home  
/etc  
/var
```

To archive files that have spaces or newlines in the names, use this format:

```
filepathlen filepath
```

Where *filepath* is the path to the file you are archiving and *filepathlen* is the number of characters in the file path.

For example:

```
5 /home  
4 /etc  
4 /var  
19 /home/abc/test file
```

filenames

Names one or more files to be archived and can be used instead of the `-f` option.

Any files that you specify must be listed at the end, after all other options.

For Apollo clients, specify absolute file paths (Apollo clients are supported only by NetBackup DataCenter master servers).

NOTES

The format to use for date and time values in NetBackup commands depends the locale setting. The examples in this command description are for a locale setting of C.

If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the usage. For example:

```
/usr/opensv/netbackup/bin/bparchive -help  
USAGE: bparchive [-c class] [-s schedule] [-k "keyword phrase"]  
      [-L progress_log] [-S master_server [,master_server,...]]  
      [-t class_type] [-w [hh:mm:ss]] -f listfile | filenames
```

Notice the hours:minutes:seconds requirements for the `-w` option. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.



EXAMPLES

EXAMPLE 1

To archive a single file, enter:

```
bparchive /usr/user1/file1
```

EXAMPLE 2

To archive files listed in a file named `archive_list`, enter:

```
bparchive -f archive_list
```

EXAMPLE 3

To associate the keyword phrase "Archive My Home Directory 01/01/97" to the archive of the directory `/home/kwc` and use a progress log named `/home/kwc/arch.log` enter the following (all on one line):

```
bparchive -k "Archive My Home Directory 01/01/97" -L  
/home/kwc/arch.log /home/kwc
```

FILES

`$HOME/bp.conf`

`/usr/opensv/netbackup/logs/bparchive/log.mmdyy`

SEE ALSO

`bp(1)`, `bpbackup(1)`, `bplist(1)`, `bprestore(1)`



bpauthsync(1M)

NAME

bpauthsync - synchronize authentication files on NetBackup servers and clients

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpauthsync [-verbose]
        [-methods] [-names] [-vopie] [-methods_allow path_name]
        [-methods_deny path_name] [-names_allow path_name ]
        [-names_deny path_name] [-clients [client1 client2 ...
        clientN ] ] [-servers [server1 server2 ... serverN ] ]
```

DESCRIPTION

This command is available only on NetBackup master servers and sets up authentication files on NetBackup servers and clients according to the options that are specified on the command.

OPTIONS

- verbose Issue additional messages.
- methods Push the `methods_allow.txt` and `methods_deny.txt` files to the specified clients and servers.
- names Push the `names_allow.txt` and `names_deny.txt` files to the specified clients and servers.
- vopie Synchronize the VOPIE key files between the specified servers and the specified clients.

Note If none of `-methods`, `-names`, and `-vopie` is specified, all three are assumed.

- methods_allow *path_name*
Specifies the local copy of the `methods_allow.txt` file to push to the servers and clients. If this option is not included, NetBackup uses the `/usr/opensv/var/auth/methods_allow.txt` file.
- methods_deny *path_name*
Specifies the local copy of the `methods_deny.txt` file to push to the servers and clients. If this option is not included, NetBackup uses the `/usr/opensv/var/auth/methods_deny.txt` file.



- `-names_allow` *path_name*
 Specifies the local copy of the `names_allow.txt` file to push to the servers and clients. If this option is not included, NetBackup uses the `/usr/opensv/var/auth/names_allow.txt` file.
- `-names_deny` *path_name*
 Specifies the local copy of the `names_deny.txt` file to push to the servers and clients. If this option not included, NetBackup uses the `/usr/opensv/var/auth/names_deny.txt` file.
- `-clients` [*client1 client2 ... clientN*]
 Names the clients to update. If `-clients` is specified without listing any client names, all unique client names in the NetBackup catalog are updated. A client name can also be specified in this format:
name:host
 Where *name* is the client name and *host* is the network host name of the client. This is useful for specifying NetBackup clients that use dynamic network addressing like DHCP.
- `-servers` [*server1 server2 ... serverN*]
 Names the servers to update.
 If `-servers` is specified but no server names are listed, all server names in the NetBackup configuration are updated.

Note The following cases also apply to using the `-clients` and `-servers` options:

If neither `-clients` nor `-servers` is used, all clients and all servers are updated.

If `-servers` is used but `-clients` is not, no clients are updated.

If `-servers` is not used but `-clients` is used along with `vopie` (either specifically or by default), the local server is updated.

If `-servers` is not used but `-clients` is used along with `-names` or `-methods`, no servers are updated.

FILES

`/usr/opensv/netbackup/logs/admin/log.*`

`/usr/opensv/var/auth/methods.txt`

`/usr/opensv/var/auth/methods_allow.txt`

`/usr/opensv/var/auth/methods_deny.txt`

`/usr/opensv/var/auth/names_allow.txt`

`/usr/opensv/var/auth/names_deny.txt`



`/usr/opensv/var/auth/vopie/*`

SEE ALSO

`vopied(1M)`, `vopie_util(1M)`

bpbackup(1)

NAME

bpbackup - back up files to the NetBackup server

SYNOPSIS

```
/usr/opensv/netbackup/bin/bpbackup [-c class] [-s schedule] [-S
master_server [ , master_server, ... ] [-t class_type] [-L
progress_log] [-w [ hh:mm:ss ] [-help] [-k
"keyword_phrase" ] -f listfile | filenames
```

```
/usr/opensv/netbackup/bin/bpbackup -c class -i [-h hostname] [-s
schedule] [-S master_server [ , master_server, ... ] [-t
class_type] [-w [ hh:mm:ss ] [-k "keyword_phrase" ]
```

DESCRIPTION

bpbackup starts either of the following:

- ◆ A user backup that is the equivalent to what is performed by using the interface on the client (see the first form of the command shown above). This type of backup can be started from any NetBackup client in order to back up files from that client.

The bpbackup command processes the files that you list on the command line or in the file that you specify with the -f *listfile* option. A file path can be a file or directory name. If the named files include a directory, bpbackup backs up all files and subdirectories of that directory starting at the directory itself.

- ◆ An immediate-manual backup of a client (see the second form of the command shown above). This variation requires the -i option on the bpbackup command and is available only to the administrator on the master server. It is the equivalent starting a manual backup from the NetBackup administrator's interface. Use the -h option to specify the host.

The following restrictions apply to this command:

- ◆ You must be the owner of the file or an administrator to back up a file with bpbackup.
- ◆ You can back up files and directories owned by other users if you have the necessary UNIX file permissions.
- ◆ If you specify a UNIX file that is a link, bpbackup backs up only the link itself, not the file to which it links.
- ◆ bpbackup does not back up the "." or ".." directory entries.



By default, you are returned to the system prompt after `bpbackup` is successfully submitted. The command works in the background and does not return completion status directly to you. The `-w` option lets you change this behavior so the command works in the foreground and returns completion status after a specified time period.

`bpbackup` writes informative and error messages to a progress-log file if you create the file prior to the execution of the `bpbackup` command and then specify the file with the `-L progress_log` option. If `bpbackup` cannot back up the requested files or directories, use the progress log to determine the reason for the failure.

If you create a directory named `/usr/opensv/netbackup/logs/bpbackup/` with public-write access, `bpbackup` creates an activity log file in this directory that can be used for troubleshooting.

In addition, if a nonroot user specifies `USEMAIL = mail_address` in their `$HOME/bp.conf` file or a root user specifies it in the `/usr/opensv/netbackup/bp.conf` file, NetBackup sends mail on the backup completion status to `mail_address`. This message is sent when the backup process is complete.

OPTIONS

- `-c class` Names the class to use for the backup.
If this option is not specified for a user backup, NetBackup uses the first class it finds that includes the client and a user backup schedule.
The `-c` option is required for an immediate-manual backup (`-i` option).
- `-i` Starts an immediate-manual backup. This is the equivalent of starting a manual backup from the NetBackup administrator interface. You must be the administrator on the master server to use the `-i` option.
- `-h hostname`
Use this option only in conjunction with the `-i` option. It names the client host on which to run the backup. If it is not specified, NetBackup runs the backup on all clients in the class.
- `-s schedule`
Names the schedule to use for the backup. If it is not specified, the NetBackup server uses the first user backup schedule it finds for the client in the class it is using (see the `-c` option).
- `-S master_server [, master_server, . . .]`
Specifies the name(s) of the NetBackup master server(s). The default is the first `SERVER` entry found in the `/usr/opensv/netbackup/bp.conf` file.
- `-t class_type`
Specifies one of the following numbers corresponding to the class type (the default is 0 on all clients except Apollo, where the default is 3):



0 = Standard
 4 = Oracle
 6 = Informix-On-BAR
 7 = Sybase
 10 = NetWare
 13 = MS-Windows-NT
 14 = OS/2
 15 = MS-SQL-Server
 16 = MS-Exchange-Server
 19 = NDMP

Note The following class types (Apollo-wbak, DataTools-SQL-BackTrack, Auspex-FastBackup, SAP, DB2, FlashBackup, Split-Mirror, and AFS) apply only to NetBackup DataCenter.

3 = Apollo-wbak
 11 = DataTools-SQL-BackTrack
 12 = Auspex-FastBackup
 17 = SAP
 18 = DB2
 20 = FlashBackup
 21 = Split-Mirror
 22 = AFS

-L *progress_log*

Specifies the name of an existing file in which to write progress information.

For example: /home/tlc/proglog

The default is to not use a progress log.

-w [*hh:mm:ss*]

Causes NetBackup to wait for a completion status from the server before returning you to the system prompt.

Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section later in this command description.

You can optionally specify a wait time in hours, minutes, and seconds. The maximum wait time you can specify is 23:59:59. If the wait time expires before the backup is complete, the command exits with a timeout status. The backup, however, still completes on the server.



If you use `-w` without specifying a wait time or you specify a value of 0, NetBackup waits indefinitely for the completion status.

If you include `-i` with `-w`, NetBackup waits until all initiated jobs have completed before returning status. However, if more than one job starts, the status is unpredictable. If the multiple jobs are due to there being more than one client and the class does not have Allow Multiple Data Streams selected, you can include the `-h` option to restrict the operation to one client and obtain predictable status. However, if the class has Allow Multiple Data Streams selected and there is more than one job from the selected client, the status is still unpredictable.

`-help` Displays a synopsis of command usage when it is the only option on the command line.

`-k keyword_phrase`

Specifies a keyword phrase that NetBackup associates with the image being created by this backup operation. You can then restore the image by specifying the keyword phrase with the `-k` option on the `bprestore` command.

If you use the `-i` option with `-k`, NetBackup establishes an association between the keyword phrase and the backup class and image.

The keyword phrase is a textual description of the backup that is a maximum of 128 characters in length. All printable characters are permitted including space (" ") and period ("."). Enclose the phrase in double quotes ("...") or single quotes ('...') to avoid conflict with the UNIX shell.

The default keyword phrase is the null (empty) string.

`-f listfile`

Specifies a file (*listfile*) containing a list of files to be backed up. This option can be used instead of the *filenames* option, but cannot be used with the `-i` option. List each file on a separate line.

The format required for the file list depends on whether the files have spaces or newlines in the names.

To back up files that do not have spaces or newlines in the names, use this format:

filepath

Where *filepath* is the path to the file you are backing up. For example:

/home

/etc

/var

To back up files that have spaces or newlines in the names, use this format:

filepathlen filepath

Where *filepath* is the path to the file you are backing up and *filepathlen* is the number of characters in the file path.

For example:

```
5 /home
4 /etc
4 /var
19 /home/abc/test file
```

filenames

Names one or more files to be backed up. This option can be used instead of the `-f` option, but cannot be used with the `-i` option. Any files that you specify must be listed at the end, following all other options.

For Apollo clients, specify absolute file paths (Apollo clients are supported only by NetBackup DataCenter master servers).

NOTES

The format that you must use for date and time values in NetBackup commands depends on the locale setting. The examples in this command description are for a locale setting of C.

If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the USAGE. For example:

```
/usr/opensv/netbackup/bin/bpbackup -help
USAGE: bpbackup [-c class] [-s schedule] [-k "keyword phrase"]
      [-L progress_log] [-i] [-h hostname]
      [-S master_server[,... ,master_server]]
      [-t class_type] [-w [hh:mm:ss]] -f listfile | filenames
```

Notice the hours:minutes:seconds requirements for the `-w` option. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES**EXAMPLE 1**

To perform a user backup of a single file, enter:

```
bpbackup /usr/user1/file1
```



EXAMPLE 2

The following command starts a user backup of the files that are listed in a file named `backup_list`.

```
bpbackup -f backup_list
```

EXAMPLE 3

The following command (all on one line) starts an immediate-manual backup of the client host named `diablo`, in the class named `cis_co`. The class type is Standard class and is in the configuration on the master server named `hoss`. The progress log is `/home/hrp/prog`.

```
bpbackup -c cis_co -i -h diablo -S hoss -t 0 -L /home/hrp/prog
```

EXAMPLE 4

The following command (all on one line) associates the keyword phrase "Backup My Home Directory 01/01/97" to the user backup of the directory `/home/kwc`. The progress log is `/home/kwc/bkup.log`.

```
bpbackup -k "Backup My Home Directory 01/01/97" -L  
/home/kwc/bkup.log /home/kwc
```

EXAMPLE 5

The following command (all on one line) associates the keyword phrase "Class Win NT 01/01/97" to the immediate-manual backup of the client host named `slater` in the class named `win_nt_class`. The progress log is `/tmp/bkup.log`.

```
bpbackup -k "Class Win NT 01/01/97" -i -h slater -c win_nt_class -t  
13 -L /tmp/bkup.log
```

FILES

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

```
/usr/opensv/netbackup/logs/bpbackup/log.mmddyy
```

SEE ALSO

```
bp(1), bparchive(1), bplist(1), bprestore(1)
```



bpbackupdb(1M)

NAME

bpbackupdb - back up NetBackup image catalogs

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpbackupdb [-dpath disk_path]
        [-nodbpaths] [-v] [path...]
```

```
/usr/opensv/netbackup/bin/admincmd/bpbackupdb [-tpath
        tape_device_path [-rv recorded_vsn]] [-nodbpaths] [-v]
        [path..]
```

```
/usr/opensv/netbackup/bin/admincmd/bpbackupdb [-opath
        optical_device_path [-rv recorded_vsn]] [-nodbpaths] [-v]
        [path...]
```

DESCRIPTION

bpbackupdb initiates a backup of one or more NetBackup image catalogs. bpbackupdb backs up the set of catalogs specified on the bpbackupdb command line. bpbackupdb also backs up the default set of NetBackup catalogs, unless the command line contains -nodbpaths. If the command line specifies a destination, the backup is stored there.

Otherwise, the backup is stored at the default location for backups of the NetBackup internal databases, which are called catalogs.

You can specify the default set of catalogs and the backup destination:

- ◆ The default paths to the NetBackup image catalogs are part of the NetBackup configuration. bpbackupdb uses the set of configured NetBackup catalog paths as the default value for the path option.
- ◆ The NetBackup configuration includes two destinations (media IDs or disk pathnames) for NetBackup catalog backups. bpbackupdb uses the less-recently used of the two destinations as its default value for the backup destination.

Your NetBackup system administrator's guide explains how to configure and display these values.

This command requires root privileges.

Only one copy of bpbackupdb can run at a time. The bpbackupdb command fails if it determines that a NetBackup catalog backup is already running. If bpbackupdb fails because other backups are in progress, retry when no other NetBackup activity is occurring.



If `bpbackupdb` fails with the message “cannot find Internet service `bpcd/tcp`,” then the service/protocol pair `bpcd, tcp` is not among the set of services defined on the local system. On UNIX, `netstat -a` displays the defined set of services. On Windows NT/2000, look for a `bpcd/tcp` entry in the `install_path\system32\drivers\etc\services` file.

Your NetBackup system administrator’s guide provides additional information on backing up NetBackup catalogs. The NetBackup utility `bprecover` recovers catalogs that `bpbackupdb` has backed up. The NetBackup troubleshooting guide (UNIX version) provides information on restoring the NetBackup catalogs if a disaster recovery is required.

OPTIONS

You can either specify a list of NetBackup image catalogs with the following options or default to the catalogs specified in the NetBackup configuration:

`-dpath` *disk_path*

`-tpath` *tape_device_path*

`-opath` *optical_device_path*

`-tpath` specifies a tape raw device path as the destination for the backup.

`-opath` specifies an optical raw device path as the destination for the backup.

`-dpath` Specifies a raw disk path as the destination for the backup.

If the media for the catalog backup is non-robotic, a mount request occurs and the catalog backup waits until the mount request is either granted or denied. The `MEDIA_MOUNT_TIMEOUT` attribute does not apply to this request.

The Media Manager device and volume daemons,

`/usr/opensv/volmgr/bin/ltid` and

`/usr/opensv/volmgr/bin/vmd`, need not be active when you use one of the destination options.

On UNIX, NetBackup assumes it is using a Berkeley-style close device for the `-tpath` option. This is the device path with `b` in the device name. For example, on Solaris the device name could be `/dev/rmt/0cbn`.

`bpbackupdb` will fail with an I/O error if it does not use a Berkeley-style close device on a platform that requires it. See the *Media Manager Device Configuration Guide* for more information.

If `-tpath` or `-opath` is used, the device name can be an NDMP device name. The syntax for an NDMP device name is `client:drivename`. An NDMP device name can contain `/` but it cannot contain `/ndmp`.

- rv** *recorded_vsn*
 This is the recorded volume serial number (RVSN). This option is meaningful if either **-tpath** or **-opath** is used. Media Manager uses the RVSN for removable media to verify that the correct platter is mounted. The RVSN is the same value as the media ID. The RVSN's string length is between one and six characters and the string can be either uppercase or lowercase.
- help** Prints a synopsis of command usage when it is the only option on the command line.
- nodbpaths**
 Do not back up the configured NetBackup catalogs. If this option is present, you must specify at least one catalog path on the command line. If this option is absent, bpbackupdb backs up the catalogs configured by NetBackup for catalog backups, as well as any catalog listed by the *path* option.
- v** Selects verbose mode. This option causes bpbackupdb to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (/usr/opensv/netbackup/logs/admin directory defined).
- path...** Back up these NetBackup catalogs. This is a list of absolute pathnames. The catalog backup paths must not contain any soft links. When NetBackup backs up its catalogs, it does not follow soft links. If you have moved any of the catalog files or directories and created soft links to their new locations, you must delete any path that has a link in it and add the actual path. Otherwise, the catalog backup aborts.
 To back up a NetBackup catalog on the master server, specify the catalog backup path as an absolute pathname, for instance,
 /usr/opensv/volmgr/database.
 To back up a NetBackup catalog on a media server other than the master server (this configuration is supported only by NetBackup DataCenter), specify the catalog backup path as *hostname:pathname*. For instance,
hostname:/usr/opensv/volmgr/database.
 There must be at least one path specified if **-nodbpaths** is present.

EXAMPLES

These examples assume that NetBackup has been configured, so that bpbackupdb can use the default values for catalogs and destination.

Example 1

This example backs up the NetBackup catalogs



```
example% bpbackupdb
```

- ◆ If the backup succeeds, the NetBackup mail administrator receives email containing the details of the backup.
- ◆ If the backup fails, the NetBackup mail administrator receives email containing the reason for the failure.

Example 2

This example backs up the NetBackup catalogs to the tape device `/dev/rmt/0mbn`.

```
example% bpbackupdb -tpath /dev/rmt/0mbn
```

EXIT STATUS

`bpbackupdb` returns the following exit values:

0 The backup succeeded.

Non-zero The backup failed.

DIAGNOSTICS

If `bpbackupdb` succeeds, it logs one of the following messages:

NB database backup to path *destination* SUCCEEDED

NB database backup to media id *destination* SUCCEEDED

NB database backup SUCCEEDED

If `bpbackupdb` fails, it logs one of the following messages:

NB database backup to path *destination* FAILED

NB database backup to media id *destination* FAILED

NB database backup FAILED

`bpbackupdb` also sends mail to the NetBackup administrator reporting the results of the backup.

FILES

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

```
/usr/opensv/netbackup/logs/admin/log.mmddyy
```

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

SEE ALSO

`bpadm(1M)`, `bprecover(1M)`, `netstat(1M)`, `services(4)`



bpclassnew(1M)

NAME

bpclassnew - create, copy, or rename a NetBackup class

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bpclassnew class_name
    [-verbose] [-M master_server,...]

/usr/opensv/netbackup/bin/admincmd/bpclassnew class_name -sameas
    existing_class_name [-verbose] [-M master_server,...]

/usr/opensv/netbackup/bin/admincmd/bpclassnew existing_class_name
    -renameto class_name [-verbose] [-M master_server,...]

```

DESCRIPTION

bpclassnew performs one of the following operations on a NetBackup class:

- ◆ Create a new NetBackup class with default attribute values
- ◆ Create a new NetBackup class with the same attributes as an existing class *r*
- ◆ Rename an existing NetBackup class

When bpclassnew runs without `-sameas` or `-renameto`, it creates a new NetBackup class with default attribute values. If `-M` is present, the defaults used for the class definition on each master server are the defaults for that master server.

bpclassnew copies a class by adding a new class to the NetBackup database. The clients, files, schedules, and attributes for the new class are the same as those for the existing class. bpclassnew does not create a class copy with the same name as an existing class.

If bpclassnew renames a class, the existing association of images with the class is lost. This means that an image listing for the renamed class does not include the images that were created before the class was renamed. bpclassnew does not rename a class to have the same name as an existing class.

The NetBackup command `bpclinfo` replaces the class-attribute defaults with new values. `bpclclients`, `bpclinclude`, and `bpclsched` define the clients, backup files, and schedules for the class. A class needs to have at least one client, one file specification, and one automatic schedule before it can run automatic backups.

bpclassnew sends its error messages to `stderr`. bpclassnew sends a log of its execution to the NetBackup admin log file for the current day.

This command requires root privileges.

See your NetBackup system administrator's guide for additional information on classes.



OPTIONS

class_name

The name of a NetBackup class which `bpclassnew` creates or the name to which `bpclassnew` changes an existing class. There is no default value. This class name must differ from any existing class name. It is composed of numeric, alphabetic, plus, minus, underscore, and period characters. Do not use a minus as the first character or leave any spaces between characters.

existing_class_name

The name of a NetBackup class which already exists when `bpclassnew` runs. There is no default value.

`-renameto`

Change the existing class's name to the new class name.

`-sameas`

Create a new class, copying its characteristics from the existing class.

`-help`

Print a command-line usage message. When `-help` is present, it is the only option on the command line.

`-M master_server,...`

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.

`-verbose`

Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).

EXAMPLES

Example 1

Create a class with default attribute values on the master server plum:

```
bpclassnew ishkabibble -M plum
```

```
bpclist ishkabibble -U -M plum
```

```
-----  
Class Name:          ishkabibble  
Class Type:          Standard  
Active:              yes
```




```

Client Compress:      no
Follow NFS Mounts:   no
Cross Mount Points:  no
Collect TIR info:    no
Block Incremental:   no
Mult. Data Streams:  no
Client Encrypt:      no
Class Priority:       0
Max Jobs/Class:      99
Disaster Recovery:   0
Residence:           (specific storage unit not required)
Volume Pool:         NetBackup
Keyword:             (none specified)

Clients:             (none defined)

Include:             (none defined)

Schedule:           (none defined)

```

Example 2

Create a new class, `myclass_copy` from the existing class `myclass`. `bpcllist` shows that `myclass_copy` has the same attributes as `myclass`. For brevity, most of the schedule information is omitted here:

```
bpclassnew myclass_copy -sameas myclass
```

```
bpcllist myclass -U
```

```

-----
Class Name:          myclass
Class Type:          Standard
Active:              yes
Client Compress:     no
Follow NFS Mounts:   no
Cross Mount Points:  no

```



Collect TIR info: no
Block Incremental: no
Mult. Data Streams: no
Client Encrypt: no
Class Priority: 0
Max Jobs/Class: 99
Disaster Recovery: 0
Residence: myunit
Volume Pool: NetBackup
Keyword: (none specified)

HW/OS/Client:	CRAY_J90	UNICOS.10.0	ixnay
	Linux	RedHat	zippity
	SGI	IRIX6	mango

Include: /tmp/my

Schedule: full
Type: Full Backup
Frequency: every 7 days
Maximum MPX: 1
Retention Level: 0 (1 week)
Residence: (specific storage unit not required)
Volume Pool: (same as class volume pool)
Daily Windows:
 Sunday 00:00:00 --> Sunday 08:00:00
 Monday 00:00:00 --> Monday 08:00:00
 Tuesday 00:00:00 --> Tuesday 08:00:00
 Wednesday 00:00:00 --> Wednesday 08:00:00
 Thursday 00:00:00 --> Thursday 08:00:00
 Friday 00:00:00 --> Friday 08:00:00
 Saturday 00:00:00 --> Saturday 08:00:00



```
Schedule:      incr
Type:          Differential Incremental Backup
```

```
bpclist myclass_copy -U
```

```
-----
Class Name:    myclass_copy
Class Type:    Standard
Active:        yes
Client Compress: no
Follow NFS Mounts: no
Cross Mount Points: no
Collect TIR info: no
Block Incremental: no
Mult. Data Streams: no
Client Encrypt: no
Class Priority: 0
Max Jobs/Class: 99
Disaster Recovery: 0
Residence:     myunit
Volume Pool:   NetBackup
Keyword:       (none specified)

HW/OS/Client:  CRAY_J90      UNICOS.10.0  ixnay
                Linux       RedHat       zippity
                SGI         IRIX6       mango
```

```
Include: /tmp/my
```

```
Schedule:      full
Type:          Full Backup
Frequency:     every 7 days
Maximum MPX:   1
Retention Level: 0 (1 week)
```



```

Residence:      (specific storage unit not required)
Volume Pool:    (same as class volume pool)
Daily Windows:
    Sunday      00:00:00  -->  Sunday      08:00:00
    Monday      00:00:00  -->  Monday      08:00:00
    Tuesday     00:00:00  -->  Tuesday     08:00:00
    Wednesday   00:00:00  -->  Wednesday   08:00:00
    Thursday    00:00:00  -->  Thursday    08:00:00
    Friday      00:00:00  -->  Friday      08:00:00
    Saturday    00:00:00  -->  Saturday    08:00:00

Schedule:       incr
Type:           Differential Incremental Backup
    
```

Example 3

Rename a class from `class_nfs` to `new_nfs`. Before and after the renaming, `bpclist` shows the classes in the NetBackup configuration database:

```

bpclist
myclass
class_nfs
test

bpclassnew class_nfs -renameto new_nfs

bpclist
myclass
new_nfs
test
    
```

EXIT STATUS

- = 0 The command executed successfully.
- <> 0 An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/openv/netbackup/logs/admin` in the form:

```
bpclassnew: EXIT status = exit status
```



If an error occurred, a diagnostic precedes this message.

FILES

`/usr/opensv/netbackup/logs/admin/*`

`/usr/opensv/netbackup/db/class/class_name`

SEE ALSO

`bpclclients(1m)`, `bpclinfo(1m)`, `bpclsched(1m)`, `bpcldelete(1m)`,
`bpcllist(1m)`

NetBackup system administrator's guide for more information on NetBackup classes.



bpclclients(1M)

NAME

bpclclients - administer the clients within NetBackup classes.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpclclients  
  
/usr/opensv/netbackup/bin/admincmd/bpclclients [class_name |  
-allunique [-ct class_type]] [-L | -l | -U | -noheader]  
[-M master_server,...] [-v]  
  
/usr/opensv/netbackup/bin/admincmd/bpclclients class_name [-M  
master_server,...] [-v] -add host_name hardware os [priority]  
  
/usr/opensv/netbackup/bin/admincmd/bpclclients class_name [-M  
master_server,...] [-v] -delete host_name ...  
  
/usr/opensv/netbackup/bin/admincmd/bpclclients class_name [-M  
master_server,...] [-v] -modify host_name [-hardware  
hardware] [-os os] [-priority priority]  
  
/usr/opensv/netbackup/bin/admincmd/bpclclients class_name -rename  
old_client_name new_client_name [-os os] [-hardware hardware]
```

DESCRIPTION

bpclclients will do one of the following:

- ◆ Produce a listing of clients.
- ◆ Add a new client to a class.
- ◆ Delete a list of clients from a class.
- ◆ Modify an existing client in a class.

For the `-add`, `-delete`, and `-modify` options, bpclclients returns to the system prompt immediately after it submits the client change request to NetBackup. To determine whether the change was successful, run bpclclients again to list the updated client information.

When the listing option is used, the list is ordered alphabetically by client name. Each client entry is on a single line, and there is a single entry for each client.

This command requires root privileges.

OPTIONS

The options used with bpclclients depend on the form of bpclclients being used.



The first form of `bpclclients` has no options and produces a listing of information about the clients for all classes.

The second form of `bpclclients` produces a listing of information about the clients for a single class or for all classes. The following options apply to this form:

`class_name` | `-allunique` [`-ct class_type`]
`class_name` specifies the name of a class and lists client information only for the class with this name.

`-allunique` without [`-ct class_type`] lists client information for all classes defined for NetBackup on the master server.

If you use `-allunique -ct class_type`, where `class_type` is a specific class type (such as Sybase), the command lists the client information only for the clients that belong to that type of class.

If the command line contains neither the `class_name` nor `-allunique` option, the listing contains client information for all classes.

These options, if used, must be the first option on the command line.

`-L` List in long format. There is no two-line header at the top of the listing; the header is embedded in the line for each client. The line for each client includes the following fields:

Client/HW/OS/Pri: (the header)

Client name

Hardware type

Operating system

Priority

There are also four additional fields which can be ignored. These fields are either unused or used for internal processing.

`-l` List in short format; this produces a terse listing and is also called *raw output mode*. There is no two-line header at the top of the listing; the header is embedded in the line for each client. The line for each client includes the following fields:

CLIENT (the header)

Client name

Hardware type

Operating system

Priority

There are also four additional fields which can be ignored. These fields are either unused or used for internal processing.

This option is useful for scripts or programs that rework the listing contents into a customized report format.



- U List in user format. The listing consists of one line for each client, containing the hardware type, operating system, and client name. A two-line header begins the listing. This is the default format for the listing.
- noheader List without any header. The listing consists of one line for each client, containing the hardware type, operating system, and client name.
- M *master_server, . . .*

A list of alternative master servers. This is a comma-delimited list of host names. If this option is present, each master server in the list executes the `bpclclients` command. Each master server in the list must allow access by the system issuing the `bpclclients` command. If an error occurs for any master server, processing stops at that point.

If `bpclclients` is producing a listing, the listing is the composite of the information returned by all the master servers in this list.

If `bpclclients` is adding, deleting, or modifying a client (explained later), the change is made on all the master servers in this list.
- v Selects verbose mode. This option causes `bpclclients` to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/openv/netbackup/logs/admin` directory defined).

The next three forms of `bpclclients` affect one or more clients in a single class. The client will be added, deleted, or have its attributes modified within the class. This form of `bpclclients` uses the following options:

- class_name*

Change the client information for this class. This option must be the first option on the command line.
- M *master_server, . . .*

Explained earlier. This option must precede the `-add`, `-delete`, or `-modify` option on the command line.
- v Explained earlier. This option must precede the `-add`, `-delete`, or `-modify` option on the command line.

Note The next three options, `-add`, `-delete`, and `-modify`, determine the change that `bpclclients` makes to the clients for the class. Any of these options, with its accompanying client information, must be the final option on the command line. Only one of these options can be used at a time.



- `-add host_name hardware os [priority]`
 Add a client to the class. If the local system already has the maximum number of clients defined, an error is returned. The installation default for the maximum number of clients is unlimited for DataCenter and 4 for BusinessServer. Specify the host name, hardware type, and operating system (see the definitions below). (*priority* is not implemented at this time)
- `-delete host_name ...`
 Delete one or more clients from the class. Up to twenty clients can be deleted at a time. The clients are provided as a space-delimited list of host names.
- `-modify host_name ...`
 Modify the attributes for a client within a class. The client has been added to the class previously. The attribute values that follow the client name replace the previous equivalent attribute values for this client. At least one of the client's attributes must be modified. `-priority` is not implemented at this time.
- `-hardware hardware`
 The hardware type of this client. Use one of the hardware types as displayed in the dialog box used for adding clients to a class with the Backup Policy Management utility.
- `-os os`
 The operating system of this client. Use one of the operating-system types as displayed in the dialog box. Use one of the hardware types as displayed in the dialog box used for adding clients to a class with the Backup Policy Management utility.
 The values chosen for the hardware and os options must form a valid combination.
- `-priority priority`
 Not implemented.

The following form of `bpclclients` changes the name of the client in a class and can also change the operating system and hardware type that is specified for the client. This form of `bpclclients` uses the following options:

- `class_name`
 The class that has the client. This option must be the first option on the command line.
- `-rename old_client_name new_client_name`
old_client_name specifies the current name of the client and *new_client_name* specifies the new name.



`-hardware hardware`

Specifies a different operating system for the client. Use one of the hardware types as displayed in the dialog box used for adding clients to a class with the Backup Policy Management utility.

`-os OS`

Specifies a different operating system for the client. Use one of the hardware types as displayed in the dialog box used for adding clients to a class with the Backup Policy Management utility.

The values chosen for the hardware and os options must form a valid combination.

EXAMPLES

Example 1

While running on the master server, list the clients known to the master server.

```
bpclclients
```

The output returned will look like the following:

Hardware	OS	Client
-----	-----	-----
C910_920	IRIX5	boris
C910_920	IRIX6	hat
Novell	NetWare	marge
PC	WindowsNT	marmot
HP9000-800	HP-UX10.20	squash
PC	WindowsNT	tiger

This command could also be entered on a client of hat, with the same results.

Example 2

List the clients defined for the class oneclass:

```
bpclclients oneclass
```

Hardware	OS	Client
-----	-----	-----
Sun4	Solaris2.6	buffalo
Sun4	Solaris2.5	jeckle
RS6000	AIX	streaky
HP9000-800	HP-UX	chilly

SGI	IRIX5	yak
ALPHA	OSF1	alpha
Sun4	Solaris2.5	heckle
HP9000-700	HP-UX	shark
NCR	UNIX	cougar
RS6000	AIX	whale
Sun4	SunOS4	oahu

Example 3

Add the client marmot to the class twoclass on the master servers serv1 and serv2. marmot's hardware type is C910_920, and marmot's operating system is IRIX6. The default priority is used. (the command is all on one line)

```
hat% bpclclients twoclass -M serv1,serv2 -add marmot C910_920
IRIX6
```

Example 4

Delete the clients marmot and vole from the class twoclass on the master servers serv1 and serv2. (the command is all on one line)

```
bpclclients twoclass -M serv1,serv2 -delete marmot vole
```

Example 5

While running on the master server hat, list client information for class BackTrack on the master server beaver:

```
bpclclients BackTrack -M beaver
```

Hardware	OS	Client
-----	-----	-----
Sun4	Solaris2.5	saturn

EXIT STATUS

0	The command executed successfully.
>0	An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form

```
bpclclients: EXIT status = exit status
```

If an error occurred, a diagnostic precedes this message.

Example 6



Assume you have a class called `my_class` that has 1 client defined. The client name is `pear`, the operating system is `Solaris2.6`, and the hardware type is `Solaris`.

```
bpclclients my_class -rename pear apple -os MacOS -hardware  
MACINTOSH
```

This command changes the client name `pear` in `my_class` to `apple`. It also changes the `os` from `Solaris2.6` to `MacOS` and hardware from `Solaris` to `MACINTOSH`.

FILES

```
/usr/opensv/NetBackup/logs/admin/*
```

```
/usr/opensv/NetBackup/db/class/class_name/clients
```

SEE ALSO

`bpadm(1M)`, `bpclinfo(1M)`



bpclinclude(1M)

NAME

bpclinclude - maintain the list of files automatically backed up by a NetBackup class

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bpclinclude class_name [-v]
    [-M master_server,...] -add path_name

/usr/opensv/netbackup/bin/admincmd/bpclinclude class_name [-v]
    [-M master_server,...] -delete path_name

/usr/opensv/netbackup/bin/admincmd/bpclinclude class_name [-v]
    [-M master_server,...] -modify {old_path_name
    new_path_name}

/usr/opensv/netbackup/bin/admincmd/bpclinclude class_name [-v]
    [-M master_server,...] -L | -l

```

DESCRIPTION

bpclinclude maintains the class file list for a NetBackup class. This is the list of files backed up when NetBackup runs an automatic backup for the class. The class file list does not apply to user backups or archives since users select the files when they start those operations.

bpclinclude performs one of the following operations:

- ◆ Adds pathnames to the class file list
- ◆ Deletes pathnames from the class file list
- ◆ Modifies pathnames in the class file list
- ◆ Displays the class file list for a class

The `-add`, `-delete`, and `-modify` options include a list of pathnames. The list of pathnames must be the final part of the bpclinclude command line. The pathname must be the entire path from the root of the filesystem to the desired location. For the absolute pathname syntax for your client type, refer to the File-Path Rules topics in your NetBackup system administrator's guide. The last part of the path can be a filename, a directory name, or a wildcard specification. You can enclose pathnames in quotes. Use enclosing quotes if the pathname contains special characters or a wildcard specification.

File-Path Rules for does not verify the existence of the input directories or files.

NetBackup backs up only the files it finds and does not require that all entries in the list be present on every client.



See your NetBackup system administrator's guide for additional information on class file lists.

For database extensions, the input entries are scripts. NetBackup executes these during the backup. See the NetBackup guide that comes with the extension product for additional information.

For certain class attributes (such as Allow Multiple Data Streams) and extension products (such as NetBackup for NDMP), the entries added to the class file list may be directives, rather than pathnames. Refer to your NetBackup system administrator's guide or the NetBackup guide for the extension product.

The options `-l` and `-L` produce nearly identical displays of the class file list.

`bpclinclude` sends its error messages to `stderr`. `bpclinclude` sends a log of its execution to the NetBackup admin log file for the current day.

This command requires root privileges.

OPTIONS

`-add path_name`

Add these `path_names` to the class file list. A pathname must be enclosed in quotes (" ") if it contains special characters, such as blank(" "), or a wildcard specification. Use a blank to separate two pathnames, not a comma. `bpclinclude` interprets a comma as part of the pathname. This means that `bpclinclude` concatenates two or more comma-delimited pathnames into a single pathname with embedded commas. `bpclinclude` does not verify the syntax or the existence of the pathnames. This option must be the final entry on the command line.

`-delete path_name`

Delete these `path_names` from the class file list. Refer to `-add` for the pathname-list syntax. Deleting a pathname from the class file list does not prevent you from recovering any backups or archives for that pathname. This option must be the final entry on the command line.

`-help`

Print a command-line usage message. When `-help` is present, it is the only option on the command line.

`-L`

Display the contents of the class file list in long format.

`-l`

Display the contents of the class file list in compact format. (Note: the `-l` and `-L` displays are similar.)

`-modify {old_path_name new_path_name}`

Modify an entry in the class file list. The values are a list of pathname pairs `{old_path_name new_path_name}`. For each pathname pair, `new_name_path` replaces `old_name_path` in the class file list. If no list



entry matches `old_path_name`, then `new_path_name` is not entered into the class file list. Refer to the `-add` option for the pathname syntax. Delimit the list entries with spaces, both within a pathname pair and between pathname pairs. This option must be the final entry on the command line.

`-M master_server,...`

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.

`-v`

Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).

OPERANDS

`class_name`

Specifies the class for which the class file list is to be set.

EXAMPLES

Example 1:

While running on another master server `kiwi`, display the class file list for class `oprdoc_class` on the master server `plum`:

```
bpclininclude oprdoc_class -L -M plum
Include:                c:\oprdoc
```

Example 2:

Illustrate `bpclininclude`'s interpretation of wildcards by adding and deleting pathnames that include one wildcard entry:

```
bpclininclude mkbclass -add /yap /y*
bpclininclude mkbclass -L
Include: /yap
Include: /y*
bpclininclude mkbclass -delete /y*
bpclininclude mkbclass -L
Include: /yap
```



Note that the wildcard entry `/y*` for `-delete` is not interpreted by `bpclinclude` as meaning that both `/yap` and `/y*` should be deleted. Only `/y*` is deleted from the include list for `mkbclass`. The interpretation of the wildcard occurs when `NetBackup` is selecting files to be backed up, during the actual backup.

Example 3:

Add two entries to the class file list for a class, and then modify them:

```
bpclinclude mkbclass -add "/ima file" "/ura file"
bpclinclude mkbclass -L
    Include: /ima file
    Include: /ura file
bpclinclude mkbclass -modify "/ima file" "/ima file 2" "/ura file"
"/ura file 2"
bpclinclude mkbclass -L
    Include: /ima file 2
    Include: /ura file 2
```

Example 4:

Add a raw partition to the class file list for the class `rc` (UNIX clients). The full path name for the device is used (the command is all on one line):

```
bpclinclude rc -add
/devices/sbus@2,0/dma@2,81000/esp@2,80000/sd@6,0:h,raw
```

(see the [Adding Unix Raw Partitions to the File List](#) section of your `NetBackup` system administrator's guide).

Example 5:

Display the class file list for the class `mkb_class`:

```
bpclinclude mkb_class -l
    INCLUDE /etc/services
    INCLUDE /etc/aliases
    INCLUDE /usr/bin
```

EXIT STATUS

= 0 The command executed successfully.



<> 0 An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form:

 bpclinclude: EXIT status = exit status

If an error occurred, a diagnostic precedes this message.

FILES

`/usr/opensv/netbackup/logs/admin/*`

`/usr/opensv/netbackup/db/class/class_name/includes`

SEE ALSO

`bpclclients(1m)`, `bpclinfo(1m)`, `bpclsched(1m)`, `bpcldelete(1m)`,
`bpcllist(1m)`

NetBackup system administrator's guide for more information on NetBackup classes.



bpclinfo(1M)

NAME

bpclinfo - manage or display class attributes for NetBackup.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpclinfo class_name -L | -l |
-U [-v] [-M master_server,...]
```

```
/usr/opensv/netbackup/bin/admincmd/bpclinfo class_name -set |
-modify [-residence label] [-pool label] [-priority
priority] [-rfile flag] [-blkincr flag] [-classjobs
max_jobs] [-multiple_streams flag] [-keyword "keyword
phrase" ] [-ct class_type] [-encrypt flag] [-active |
-inactive] [-collect_tir_info value] [-compress flag]
[-crossmp flag] [-follownfs flag] [-v] [-M
master_server,...]
```

DESCRIPTION

bpclinfo initializes, modifies, or displays the attribute values for a NetBackup class. Only root can execute this command.

OPTIONS

The options used with bpclinfo depend on the form of bpclinfo being used.

The first form of bpclinfo displays a class. The following options apply to this form:

```
class_name -L | -l | -U
```

List information for this class. This is a required option.

-L specifies a long list type and produces a listing with one class attribute per line, in the format *class_attribute: value*. The value may be expressed both in numeric and name form. Fields in the list include:

Class Type

Active

Follow NFS Mounts

Cross Mount Points

Client Compress

Collect TIR Info

Class Priority

Ext Security Info

File Restore Raw



Client Encrypt

Max Jobs/Class

Multiple Streams

Disaster Recovery

Max Frag Size

Residence

Volume Pool

-l specifies a short list type and produces a terse listing. This option is useful for scripts or programs that rework the listing contents into a customized report format. A short listing contains the following information for the specified class:

Line 1: "INFO", client_type, follow_nfs_mounts, client_compress, priority, proxy_client, client_encrypt, disaster recovery, max_jobs_per_class, cross_mount_points, max_frag_size, active, collect_tir_info, block_incr, ext_sec_info, i_f_r_f_r, streaming, ci_unused1

Line 2: "KEY",keyword

Line 3: "BCMD",backup_command

Line 4: "RCMD",restore_command

Line 5: "RES",residence

Line 6: "POOL", pool

-U specifies a user list type and produces a listing with one class attribute per line, in the format *class_attribute: value*. This listing is similar to the -L listing, but contains fewer fields:

-M *master_server,...*

A list of alternative master servers. This is a comma-delimited list of hostnames. If this option is present, each master server in the list executes the `bpclinfo` command. Each master server in the list must allow access by the system issuing the `bpclinfo` command. If an error occurs for any master server, processing terminates at that point.

For the display form of `bpclinfo`, the report is the composite of the information returned by all the master servers in this list. `bpclinfo` queries each of these master servers. The master server returns information from its class catalog.

For the class-definition form of `bpclinfo`, the class is created or modified on each master server in the list.

The default is the master server for the system running `bpclinfo`.



- `-v` Selects verbose mode. This option causes `bpclinfo` to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/openv/netbackup/logs/admin` directory defined).

The second form of `bpclinfo` initializes attribute values for a class or modifies the attribute values for a class. The following options apply to this form:

Note Not all options apply to every class type. For instance, if the class type is *MS-Windows-NT*, `bpclinfo` accepts the options `-compress` and `-crossmp`. When `bpclinfo` completes, it returns a zero status. However, NetBackup's subsequent handling of the class with the *MS-Windows-NT* class type is as though the options had not been set.

`class_name` `-set` | `-modify`
Initialize or modify attributes for this class. This is a required option.
`-set` initializes (or reinitializes) attributes for the class to their default values, except for those attributes set by options on the current command line.
`-modify` modifies attributes for the class. Attributes that are not explicitly set by options on the current command line do not change their values.

`-ct class_type`
Specify the class type by entering one of the following character strings (the default is Standard):
Informix-On-BAR
MS-Exchange-Server
MS-SQL-Server
MS-Windows-NT
NDMP
NetWare
Oracle
OS/2
Standard
Sybase

Note The following class types (AFS, Apollo-wbak, Auspex-FastBackup, DataTools-SQL-BackTrack, DB2, FlashBackup, SAP, and Split-Mirror) apply only to NetBackup DataCenter.

AFS



Apollo-wbak
 Auspex-FastBackup
 DataTools-SQL-BackTrack
 DB2
 FlashBackup
 SAP
 Split-Mirror

-pool *volume_pool_label*

Specifies the volume pool for the class. The default is NetBackup. The volume pool should be one of the volume pools for the class storage unit. This attribute is not relevant if a disk storage unit is the residence for the class. If the class storage unit is Any_available (Residence: - appears on the bpclinfo display), then the volume pool for any storage unit can be selected. If "*NULL*" is specified, the volume pool is set to NetBackup. To display the configured volume pools, run
 /usr/opensv/volmgr/bin/vmpool -listall.

-residence *storage_unit_label*

Specifies the label of the storage unit for storing the backups created according to this schedule. The default is Any_available. This allows the class to use any storage unit which has the attribute On Demand Only? set to No. If the class needs to use a specific storage unit or the storage unit desired has the attribute On Demand Only? set to Yes, then specify the storage unit. If "*NULL*" is specified, the residence for the schedule is set (or reset) to Any_available. The class residence determines the residence for the class schedules, unless the he Override Class Storage Unit setting on an individual schedule specifies a residence. Run bpstulist to display the set of defined storage units.

-priority *flag*

The priority of this class in relation to other classes. Priority is greater than or equal to 0. This value determines the order in which classes are executed. The higher the value, the earlier the class is executed. The default is 0, which is the lowest priority.

-rfile *flag*

0 (disabled) or 1 (enabled).

If 1, allow Individual File Restore From Raw.

If 0, do not allow Individual File Restore From Raw.

For a FlashBackup class, this option is ignored, since the attribute is always enabled. (Note that FlashBackup is available only if you are running NetBackup DataCenter and have the separately-priced FlashBackup option.)



`-blkincr` *flag*

Note This option applies only if you are running NetBackup DataCenter and also have VERITAS Oracle Edition, which supports block-level incrementals.

0 (disabled) or 1 (enabled). Perform block-level-incremental backups for clients in this class.

If 1, do perform block-level-incremental backups.

If 0, do not perform block-level-incremental backups.

`-classjobs` *max_jobs*

The maximum number of concurrent jobs that NetBackup allows for this class (corresponds to the Limit Jobs per Class setting in the administration interface). *max_jobs* is always greater than or equal to 0.

For the default or when `-classjobs` is 0, `bpclinfo` sets *max_jobs* to a value that corresponds to unlimited. The effective maximum number of jobs in this instance is 8 for NetBackup BusinessServer and 99 for NetBackup DataCenter.

`-multiple_streams` *flag*

0 (disabled) or 1 (enabled). Allow Multiple Data Streams.

If 1, allow multiple data streams.

If 0, do not allow multiple data streams.

`-keyword` "*keyword phrase*"

The value will be associated with all backups created using this class. The keyword phrase can be used to link related classes. It can also be used during restores to search only for backups that have the keyword phrase associated with them.

`-encrypt` *flag*

0 (disabled) or 1 (enabled). Specifies whether files should be encrypted or not.

If 1, encryption is enabled.

If 0, encryption is disabled.

`-active` | `-inactive`

Set the class to active or inactive. If the class is active, NetBackup executes all its automatic schedules and permits user-directed backups and archives to be used. A class must be active for an automatic backup to occur. This is the default.

If the class is inactive, NetBackup does not execute any automatic schedules or permit user-directed schedules to be used. This option is useful for temporarily inactivating a class to prevent schedules from being used.

`-collect_tir_info` *value*

Collect true-image-recovery (TIR) information. True-image recovery allows NetBackup to restore a directory to exactly what it was at the time of any scheduled full or incremental backup. Files deleted before the time of the selected backup are not restored. After enabling this attribute, NetBackup starts collecting additional information beginning with the next full or incremental backup for the class.

If 0, NetBackup does not keep track of true-image-recovery information.

If 1, NetBackup collects TIR information.

If 2, NetBackup collects TIR information and tracks client files.

`-compress` *flag*

0 (disabled) or 1 (enabled). Specifies whether to compress files or not. If 1, the files selected are compressed by the client software onto the media. Compression may increase total backup time. If 0, the files are not compressed onto the media. This is the default.

This option has no effect on the hardware compression that may be available on the storage unit.

Image compression is not available on Apollo clients (Note that Apollo clients are supported only by NetBackup DataCenter servers.)

`-crossmp` *flag*

0 (disabled) or 1 (enabled). Specifies whether to cross mount points during backups or not.

If 1, NetBackup backs up or archives all files and directories in the selected path regardless of the file system on which they reside.

If 0, NetBackup backs up or archives only those files and directories that are on the same file system as the selected file path. This is the default.

This attribute can affect the Follow NFS class attribute (applies only to NetBackup DataCenter). Refer to NetBackup DataCenter system administrator's guide for more details.

This attribute does not affect Apollo clients. Those clients always behave as if the attribute is enabled. (Note that Apollo clients are supported only by NetBackup DataCenter servers.)

`-follownfs` *flag*

Note The `follownfs` options applies only to NetBackup DataCenter

0 (disabled) or 1 (enabled). Specifies whether to follow NFS mount points or not. For class types MS-Windows-NT and OS/2, setting this flag affects the class attribute Backup Network Drives instead of the Follow NFS attribute.



If 1, NetBackup backs up or archives any NFS-mounted files encountered.

If 0, NetBackup does not back up or archive any NFS-mounted files encountered. This is the default.

The behavior of this attribute varies somewhat depending on the setting of the Cross Mount Points attribute. Refer to your NetBackup system administrator's guide for more details.

This attribute does not affect Apollo clients. Apollo clients always behave as though the attribute is enabled. Therefore, avoid putting NFS-mounted files in the file list for classes containing Apollo clients unless you want them backed up.

`-M master_server, . . .`

Same as explained earlier.

`-v`

Same as explained earlier.

The third form of `bpclinfo` (not shown in the synopsis) shows usage information and has only one option as follows:

`-help` Prints a synopsis of command usage when it is the only option on the command line.

EXAMPLES

EXAMPLE 1

To set the storage unit of the class `tstclass` to `tstunit` and view the results, perform the following:

```
bpclinfo tstclass -modify -residence tstunit
```

```
bpclinfo tstclass -L
```

```
Class Type:           Standard (0)
Active:               no
Follow NFS Mounts:   no
Cross Mount Points:  no
Client Compress:     no
Collect TIR Info:    no
Class Priority:       0
Ext Security Info:   no
File Restore Raw:    no
Client Encrypt:      no
Max Jobs/Class:      8
```




```

Multiple Streams:  1
Disaster Recovery: 0
Max Frag Size:    0 MB (unlimited)
Residence:        tstunit - - -
Volume Pool:      NetBackup - - -

```

EXAMPLE 2

To set the attributes of class `tstclass` back to their default values, perform the following:

```

bpclinfo tstclass -set
bpclinfo tstclass -L
Class Type:          Standard (0)
Active:              yes
Follow NFS Mounts:  no
Cross Mount Points: no
Client Compress:    no
Collect TIR Info:   no
Class Priority:      0
Ext Security Info:  no
File Restore Raw:   no
Client Encrypt:      no
Multiple Streams:   0
Disaster Recovery:  0
Max Jobs/Class:     8
Max Frag Size:      0 MB (unlimited)
Residence:          - - - - -
Volume Pool:        NetBackup - - - -

```

EXAMPLE 3

The following is an example of a short listing for the class named `mkbclass`:

```

bpclinfo mkbclass -l
INFO 0 0 0 0 *NULL* 0 0 99 0 0 0 0 0 0 0 0 *NULL*
KEY my temp directory
BCMD *NULL*
RCMD *NULL*

```



```
RES mkbunit *NULL* *NULL* *NULL* *NULL* *NULL* *NULL* *NULL* *NULL* *NULL*
*NULL*
```

```
POOL NetBackup *NULL* *NULL* *NULL* *NULL* *NULL* *NULL* *NULL* *NULL*
*NULL*
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

```
/usr/opensv/netbackup/db/class/class_name/info
```

bpclsched(1M)

NAME

bpclsched - add, delete, or list NetBackup schedules

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpclsched class_name [-v] [-M
master_server,...] -add sched_label [-st sched_type] [-rl
retention_level] [-freq frequency] [-mpxmax mpx_factor]
[-residence storage_unit_label] [-pool volume_pool_label]
[-window start duration]
```

```
/usr/opensv/netbackup/bin/admincmd/bpclsched class_name [-v] [ M
master_server,... -delete sched_label
```

```
/usr/opensv/netbackup/bin/admincmd/bpclsched class_name [-v] [ M
master_server,... -deleteall
```

```
/usr/opensv/netbackup/bin/admincmd/bpclsched class_name [-v] [-M
master_server...] [-L | -l | -U] [-label sched_label]
```

DESCRIPTION

bpclsched will do one of the following:

- ◆ Add a new schedule to a class.
- ◆ Delete one or more schedules from a class.
- ◆ Delete all the schedules from a class.
- ◆ List one or all schedules in a class.

For the `-add` and `-delete` options, bpclsched returns to the system prompt immediately after it submits the schedule change request to NetBackup. To determine whether the change was successful, run bpclsched again to list the updated schedule information.

When the listing option is used there is a single entry for each schedule, even if the `-M` option is used. The `-l` form lists the information for each schedule on several lines. `-l` does not identify the attributes by name; these are as follows (where the names are not described, they are reserved for internal NetBackup use):

Line 1: SCHED, schedule name, type, max_mpx, frequency, retention level, u_wind/o/d, 2 internal attributes, maximum fragment size, unused5, unused4, unused3, unused2, unused1. Note that u_wind/o/d is a field reserved for future use. This is also true for the u_wind entry in the `-L` display.



Line 2: SCHEDWIN, seven pairs of the form *start,duration*, expressing the start and duration of the window for each day of the week, starting with Sunday.

Line 3: SCHEDRES, residence

Line 4: SCHEDPOOL, pool

If the `-M` option is used, `bpclsched` performs the operation on each of the master servers listed. For instance, if `bpclsched` is adding a schedule, `bpclsched` adds the schedule to the class on each of the master servers listed for `-M`. If the `-M` option is used on a listing request, the listing is the composite of the information returned by all of the master servers in the `-M` list. If the command fails for any of the master servers, execution stops at that point.

To modify an existing NetBackup schedule, use the NetBackup command `bpclschedrep`.

This command requires root privileges.

OPTIONS

These options are common to all forms of `bpclsched`:

class_name

The name of the class that contains the schedules. The class must exist before the execution of this command. This option is required, and must be the first one on the command line.

`-help` Prints a synopsis of command usage when it is the only option on the command line.

`-M master_server, . . .`

A list of alternative master servers. This is a comma-separated list of host names. If this option is present, each master server in the list executes the `bpclsched` command. Each master server in the list must allow access by the system issuing the `bpclsched` command.

If this option is present, the command is executed on each master server in the list. If an error occurs for any master server, processing terminates at that point.

If `bpclsched` is producing a listing, the listing is the composite of the information returned by all the master servers in this list.

If `bpclsched` adds or deletes a schedule, all master servers in this list receive the change.

`-v` Selects verbose mode. This option causes `bpclsched` to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/openv/netbackup/logs/admin` directory defined).



The remaining options depend on the form of `bpclsched`. The first form of `bpclsched` adds a schedule to the named class. The following options apply to this form of `bpclsched`:

- `-add sched_label [suboptions]`
Add a single schedule to the named class.
The suboptions for the `-add` option explained below. These are attributes of the schedule being added. Refer to your NetBackup system administrator's guide for details on schedules and their attributes.
- `-st sched_type`
This is the type of the schedule. The default schedule type is FULL. Here are the possible values, with their meanings, for this attribute:
FULL - full
INCR - differential incremental
CINC - cumulative incremental
UBAK - user backup
UARC - user archive
- `-rl retention_level`
The retention level determines how long to retain backups and archives. The `retention_level` is an integer between 0 and 9. The default retention level is 1. Valid retention levels and their corresponding default retention times are listed below.

Caution Because the retention period associated with each level can be changed by using the NetBackup administration interface, your configuration may have different values for each level than those shown here. Use the NetBackup administration interface to determine the actual retention periods before making any changes with this command. Otherwise, backups could expire sooner than you expect, resulting in loss of data.

- 0 1 week
- 1 2 weeks
- 2 3 weeks
- 3 1 month
- 4 2 months
- 5 3 months
- 6 6 months
- 7 9 months
- 8 1 year
- 9 Infinity



`-freq` *frequency*

The frequency determines how often backups run for this schedule. This value represents the number of seconds between backups for this schedule. This frequency does not apply to user backups and archives. The default is one week.

`-mpxmax` *mpx_factor*

This is the maximum number of jobs for this schedule that NetBackup will multiplex on any one drive. *mpx_factor* is an integer that can range from 1 through 8 for NetBackup BusinessServer and 1 through 32 for NetBackup DataCenter. A value of 1 means that backups for this schedule are not multiplexed. The default is no multiplexing.

`-residence` *storage_unit_label*

This is the name of the storage unit, which specifies the location of the backup images. The value `"*NULL*"` causes NetBackup to use the storage unit specified at the class level. The default is for NetBackup to use the storage unit specified at the class level. If you do not specify a storage unit at either the schedule level or the class level, NetBackup uses the next storage unit available.

`-pool` *volume_pool_label*

This is the name of the volume pool. This choice overrides the class-level volume pool. Entering `"*NULL*"` causes NetBackup to use the volume pool specified at the class level. The default is to use the volume pool specified at the class level. The volume pool label cannot be None. If you do not specify a volume pool at either the schedule level or the class level, NetBackup uses a default value of NetBackup.

`-window` *start duration*

Specifies when NetBackup can run the backups for this schedule. Every day of the week has the same window.

start is the time at which the backup window opens for this schedule.

This is the number of seconds since midnight. This is an integer between 0 and 86399 (there are 86400 seconds in a day).

duration is the length of time that the window remains open. The time unit is seconds. This is a non-negative integer.

The second form of `bpclsched` deletes one or more schedules from the named class. The following option applies to this form of `bpclsched`:

`-delete` *sched_label*

Delete the listed schedules from the named class. The elements of the *sched_label* list must be separated by spaces. There can be up to 25 labels in the list.

The third form of `bpclsched` deletes all schedule from the named class. The following option applies to this form of `bpclsched`:



`-deleteall`
Delete all schedules from the named class.

The fourth form of `bpclsched` produces a listing of information about the schedules for the named class. The following options apply to this form of `bpclsched`:

- `-L` The list type is long. This listing includes all attributes for the schedule. Some attribute values are descriptive terms, rather than numbers.
- `-l` The list type is short. This is the default list type. This produces a terse listing that includes all attributes for the schedule. Each schedule occupies one line of the listing. Most attribute values are expressed numerically. This option is useful for scripts or programs that rework the listing contents into a customized report format.
- `-U` The list type is user. This listing is similar to the long-type listing, but it has fewer entries. Most attribute values are descriptive terms, rather than numbers.
- `-label sched_label`
List the attributes for this schedule in the named class. The default is to list information for all schedules for the named class.

EXAMPLES

Example 1:

In this example, `bpclsched` lists the information for schedule `user` within class `tstclass` in two different ways. The first display is in long mode. The second is in User mode, which shows fewer entries than the Long mode display.

```
bpclsched tstclass -L -label user
Schedule:                user
Type:                    UBAK (2)
Frequency:               1 day(s) (86400 seconds)
Retention Level:        0 (1 week)
u-wind/o/d:              0 0
Incr Type:               DELTA (0)
Incr Depends:           (none defined)
Max Frag Size:          0 MB (unlimited)
Maximum MPX:            1
Residence:               (specific storage unit not required)
Volume Pool:            (same as class volume pool)
Daily Windows:
```



Day	Open	Close	W-Open	W-Close
Sunday	000:00:00	024:00:00	000:00:00	024:00:00
Monday	000:00:00	024:00:00	024:00:00	048:00:00
Tuesday	000:00:00	024:00:00	048:00:00	072:00:00
Wednesday	000:00:00	024:00:00	072:00:00	096:00:00
Thursday	000:00:00	024:00:00	096:00:00	120:00:00
Friday	000:00:00	024:00:00	120:00:00	144:00:00
Saturday	000:00:00	024:00:00	144:00:00	168:00:00

```
bpclsched tstclass -U -label user
```

```
Schedule:          user
Type:             User Backup
Retention Level:  0 (1 week)
Maximum MPX:      1
Residence:        (specific storage unit not required)
Volume Pool:      (same as class volume pool)
Daily Windows:
    Sunday 00:00:00 --> Sunday 24:00:00
    Monday 00:00:00 --> Monday 24:00:00
    Tuesday 00:00:00 --> Tuesday 24:00:00
    Wednesday 00:00:00 --> Wednesday 24:00:00
    Thursday 00:00:00 --> Thursday 24:00:00
    Friday 00:00:00 --> Friday 24:00:00
    Saturday 00:00:00 --> Saturday 24:00:00
```

Example 2:

While running on the system `hat`, list information for the schedule named `full` in class `tstclass`, as defined on the master server `beaver`:

```
bpclsched tstclass -M beaver -L -label full
Schedule:          full
Type:             FULL (0)
Frequency:         0+ day(s) (14400 seconds)
Retention Level:  0 (1 week)
u-wind/o/d:       0 0
```




```

Incr Type:      DELTA (0)
Incr Depends:   (none defined)
Max Frag Size:  0 MB (unlimited)
Maximum MPX:    1
Residence:      (specific storage unit not required)
Volume Pool:    (same as class volume pool)
Daily Windows:
  Day           Open           Close           W-Open          W-Close
  Sunday        000:00:00      024:00:00      000:00:00      024:00:00
  Monday        000:00:00      024:00:00      024:00:00      048:00:00
  Tuesday       000:00:00      024:00:00      048:00:00      072:00:00
  Wednesday     000:00:00      024:00:00      072:00:00      096:00:00
  Thursday      000:00:00      024:00:00      096:00:00      120:00:00
  Friday        000:00:00      024:00:00      120:00:00      144:00:00
  Saturday      000:00:00      024:00:00      144:00:00      168:00:00

```

Example 3:

This example adds a new schedule, `full_2`, to the class `tstclass` on `beaver`, and then lists the new schedule in Long mode. These commands run on the system hat:

```

bpclsched tstclass -M beaver -add full_2
bpclsched tstclass -M beaver -label full_2 -L
Schedule:          full_2
  Type:            FULL (0)
  Frequency:       7 day(s) (604800 seconds)
  Retention Level: 1 (2 weeks)
  u-wind/o/d:      0 0
  Incr Type:       DELTA (0)
  Incr Depends:    (none defined)
  Max Frag Size:  0 MB (unlimited)
  Maximum MPX:    1
  Residence:       (specific storage unit not required)
  Volume Pool:    (same as class volume pool)
  Daily Windows:
    Day           Open           Close           W-Open          W-Close

```



```

Sunday      000:00:00  000:00:00
Monday      000:00:00  000:00:00
Tuesday     000:00:00  000:00:00
Wednesday   000:00:00  000:00:00
Thursday    000:00:00  000:00:00
Friday      000:00:00  000:00:00
Saturday    000:00:00  000:00:00

```

Example 4:

In this example, bpclsched deletes the schedules, full_3, user, user_2, and user_3 from class tstclass:

```
bpclsched tstclass -delete full_3 user user_2 user_3
```

Example 5:

In this example, bpclsched lists the schedule information for class tstclass:

```

bpclsched tstclass -L
Schedule:          full
  Type:            FULL (0)
  Frequency:       1 day(s) (86400 seconds)
  Retention Level: 0 (1 week)
  u-wind/o/d:      0 0
  Incr Type:       DELTA (0)
  Incr Depends:    (none defined)
  Max Frag Size:   0 MB (unlimited)
  Maximum MPX:     1
  Residence:       (specific storage unit not required)
  Volume Pool:     (same as class volume pool)
Daily Windows:
  Day      Open      Close      W-Open      W-Close
  Sunday   000:00:00  024:00:00  000:00:00  024:00:00
  Monday   000:00:00  024:00:00  024:00:00  048:00:00
  Tuesday  000:00:00  024:00:00  048:00:00  072:00:00
  Wednesday 000:00:00  024:00:00  072:00:00  096:00:00
  Thursday 000:00:00  024:00:00  096:00:00  120:00:00
  Friday   000:00:00  024:00:00  120:00:00  144:00:00

```



```

Saturday    000:00:00  024:00:00  144:00:00  168:00:00

```

```

Schedule:      user
Type:          UBAK (2)
Frequency:     1 day(s) (86400 seconds)
Retention Level: 0 (1 week)
u-wind/o/d:   0 0
Incr Type:     DELTA (0)
Incr Depends: (none defined)
Max Frag Size: 0 MB (unlimited)
Maximum MPX:   1
Residence:    (specific storage unit not required)
Volume Pool:  (same as class volume pool)

```

Daily Windows:

Day	Open	Close	W-Open	W-Close
Sunday	000:00:00	024:00:00	000:00:00	024:00:00
Monday	000:00:00	024:00:00	024:00:00	048:00:00
Tuesday	000:00:00	024:00:00	048:00:00	072:00:00
Wednesday	000:00:00	024:00:00	072:00:00	096:00:00
Thursday	000:00:00	024:00:00	096:00:00	120:00:00
Friday	000:00:00	024:00:00	120:00:00	144:00:00
Saturday	000:00:00	024:00:00	144:00:00	168:00:00

Example 6:

In this example, `bpclsched` adds a new schedule, full, with a window from 11 pm to midnight. The second `bpclsched` lists the information for schedule full:

```

bpclsched elevenpm -add full -window 82800 3600
bpclsched elevenpm -U -label full
Schedule:      full
Type:          Full Backup
Frequency:     every 7 days
Retention Level: 1 (2 weeks)

```



Maximum MPX: 1
Residence: (specific storage unit not required)
Volume Pool: (same as class volume pool)
Daily Windows:
 Sunday 23:00:00 --> Sunday 24:00:00
 Monday 23:00:00 --> Monday 24:00:00
 Tuesday 23:00:00 --> Tuesday 24:00:00
 Wednesday 23:00:00 --> Wednesday 24:00:00
 Thursday 23:00:00 --> Thursday 24:00:00
 Friday 23:00:00 --> Friday 24:00:00
 Saturday 23:00:00 --> Saturday 24:00:00

FILES

/usr/opensv/netbackup/logs/admin/*

/usr/opensv/netbackup/db/class/*class_name*/schedule

SEE ALSO

bpclschedrep(1M)

bpclschedrep(1M)

NAME

bpclschedrep - modify the attributes of a NetBackup schedule.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpclschedrep class_name
sched_label [ -M master_server,... ] [-v] [-st sched_type]
[-rl retention_level] [-freq frequency] [-mpxmax mpx_factor]
[-residence storage_unit_label] [-pool volume_pool_label]
[-(0..6) start duration]
```

DESCRIPTION

bpclschedrep changes the attributes of a NetBackup schedule. The schedule and class named by bpclschedrep should already exist when this command is executed. If the -M option is used, bpclschedrep changes the schedule on each of the master servers listed.

This command requires root privileges.

OPTIONS

- class_name*
The name of the class that contains the schedule. This class has been previously created.
- sched_label*
The name of the schedule to be changed. This schedule has been previously created.
- freq *backup_frequency*
The backup frequency controls how much time can elapse between successful automatic backups for clients on this schedule. Frequency does not apply to user schedules because the user can perform a backup or archive any time the backup window is open. This value is a positive integer, representing the number of seconds between successful automatic backups for this schedule.
- help
Prints a synopsis of command usage when it is the only option on the command line.
- mpxmax *mpx_factor*
This is the maximum multiplexing factor for this schedule. Multiplexing sends concurrent, multiple backups from one or several clients to a single drive.



The multiplexing factor can range from 1 through 8 for NetBackup BusinessServer and 1 through 32 for NetBackup DataCenter. A value of 1 specifies no multiplexing and a value greater than 1 means that NetBackup should create multiplexed images on the destination media. The multiplexing factor should be less than or equal to the multiplexing factor for the storage unit.

For more information on multiplexing refer to the multiplexing topic in your NetBackup system administrator's guide.

`-pool` *volume_pool_label*

Specifies the volume pool for the schedule. Do not use this option if a disk storage unit is the residence for the schedule. If `"*NULL*"` is specified, the volume pool for the schedule is the volume pool of the class which contains this schedule.

To display the configured volume pools, run
`/usr/opensv/volmgr/bin/vmpool -listall.`

`-residence` *storage_unit_label*

Specifies the label of the storage unit to be used for storing the backups created according to this schedule. If `"*NULL*"` is specified, the residence for the schedule defaults to the residence of the class which contains this schedule. If the residence value is a storage unit label, the residence for the schedule becomes that storage unit, overriding the residence for the class.

Run `bpstulist` to display the set of defined storage units.

`-rl` *retention_level*

Specifies how long NetBackup retains the backups that it creates using this schedule. Valid retention levels and their corresponding default retention times are listed below.

Caution Because the retention period associated with each level can be changed by using the NetBackup administration interface, your configuration may have different values for each level than those shown here. Use the NetBackup administration interface to determine the actual retention periods before making any changes with this command. Otherwise, backups could expire sooner than you expect, resulting in loss of data.

- 0 1 week
- 1 2 weeks
- 2 3 weeks
- 3 1 month
- 4 2 months
- 5 3 months



6 6 months

7 9 months

8 1 year

9 Infinity

NetBackup keeps the information about the backups for the specified time. Then it deletes information about them. Once deleted, the files in the backups are unavailable for restores. When all the backups on a volume have expired, the volume can be reassigned.

-st *sched_type*

Specifies the type of backup this schedule performs. Schedule types fall into two main categories: automatic and user. Automatic schedules define the windows during which the NetBackup scheduler can initiate a backup for this class.

User schedules define the windows during which a user can initiate a backup or archive.

The values for schedule type are

FULL (full backup)

INCR (differential incremental backup)

CINC (cumulative incremental backup)

UBAK (user backup)

UARC (user archive)

-(0..6) *start duration*

Specifies the window during which NetBackup can run the backups for this schedule. This window applies to a specific day of the week. 0 corresponds to Sunday, 1 to Monday, and so on.

start is the time at which the backup window opens for this schedule.

This is the number of seconds since midnight. It is an integer between 0 and 86400 (the number of seconds in a day).

duration is the length of time that the window remains open. The time unit is seconds. This is a non-negative integer.

-M *master_server, . . .*

A list of alternative master servers. This is a comma-separated list of hostnames. If this option is present, each master server in the list executes the `bpclschedrep` command. Each master server in the list must allow access by the system issuing the `bpclschedrep` command. If an error occurs for any master server, processing terminates at that point.

The schedule attributes will be modified on all the master servers in this list.



`-v` Selects verbose mode. This option causes `bpclschedrep` to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/opensv/netbackup/logs/admin` directory defined).

EXAMPLES

Example 1

Set the frequency for a schedule.

```
bpclschedrep mkbclass incr -freq 604800
```

This sets to 1 week the frequency with which automatic backups will be performed for the schedule `incr` in class `mkbclass`.

Example 2

For Saturday and Sunday of each week, have the window for schedule `full` in class `newclass` open at 10 pm instead of 11 pm. Also, have the window duration be 2 hours instead of 1 hour. `bpclschedrep` resets the windows, and `bpclsched` lists the new schedule values.

```
bpclschedrep newclass full -0 79200 7200 -6 79200 7200
```

```
bpclsched newclass -U -label full
```

```
Schedule:          full
Type:              Full Backup
Frequency:         every 7 days
Retention Level:  1 (2 weeks)
Maximum MPX:      1
Residence:         (specific storage unit not required)
Volume Pool:      (same as class volume pool)

Daily Windows:
    Sunday  22:00:00  -->  Sunday  24:00:00
    Monday  23:00:00  -->  Monday  24:00:00
    Tuesday 23:00:00  -->  Tuesday 24:00:00
    Wednesday 23:00:00 -->  Wednesday 24:00:00
    Thursday 23:00:00 -->  Thursday 24:00:00
    Friday  23:00:00  -->  Friday  24:00:00
    Saturday 22:00:00 -->  Saturday 24:00:00
```


FILES

/usr/opensv/netbackup/logs/admin/*

/usr/opensv/netbackup/db/class/*class_name*/schedule

SEE ALSO

bpclsched(1M)



bpconfig(1M)

NAME

bpconfig - modify or display the global configuration attributes for NetBackup

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpconfig [-cd seconds] [-ha  
  hours] [-kl days] [-kt days] [-ma [address]] [-mdtm drives]  
  [-mhto seconds] [-mj number] [-period hours] [-prep  
  hours] [-to seconds] [-tries times] [-wi minutes]  
  [-verbose] [-M master_server, ...]  
  
/usr/opensv/netbackup/bin/admincmd/bpconfig [-L | -l | -U]  
  [-verbose] [-M master_server, ...]
```

DESCRIPTION

bpconfig modifies or displays the NetBackup global configuration attributes. These attributes affect operations for all classes and clients. With the exception of the NetBackup administrator's email address, the default values for these attributes should be adequate for most installations. The section on NetBackup Global Attributes, in your NetBackup system administrator's guide describes the implications of setting the attribute values.

- ◆ The first form of bpconfig modifies one or more of the NetBackup global configuration attributes. At least one option that changes a NetBackup global configuration attribute must be on the command line.
- ◆ The second form of bpconfig displays the current settings of the NetBackup global configuration attributes. See the section DISPLAY FORMATS for more detail on the displays.

Errors are sent to stderr. A log of the command's execution is sent to the NetBackup admin log file for the current day.

This command requires root privileges.

OPTIONS

-cd *seconds*

The number of seconds that is the Compress-image-Database time interval. When seconds is a positive integer, an image will be compressed after this number of seconds has elapsed since the creation of the image. On Windows NT, NetBackup uses NTFS file compression only if the database is in an NTFS partition. Otherwise, it is not compressed.



The effect of compression is that less disk space is needed for the image database. However, when browsing the image database for restoring, the images need to be decompressed before they can be searched. While browsing for a restore, the compressed images will not be found. To decompress the images, you must use `bpimage(1m)`.

The default is 0, which means no compression is done.

`-ha` *hours*

The number of *hours* ago that is the beginning of the time range for selecting NetBackup report entries. The end of the time range is the current time. For instance, if *hours* ago is 24 and if you request a Backup Status report at 10:00 a.m., the report includes all backups run from 10:00 a.m. yesterday until 10:00 a.m. today. This value is used to calculate the time range for general reports and media reports. General reports include Backup Status, Client Backups, Problems, and All Log Entries. Media reports include Media List, Media Summary, Media Contents, Images on Media, and Media Log Entries. Hours Ago is a positive integer. The default value is 24 hours.

`-kl` *days*

The number of days to keep logs. This determines how long the NetBackup master server keeps its Error database and activity logs. NetBackup derives its Backup Status, Problems, All Log Entries, and Media Log Entries reports from the Error database, so this value limits the period that these reports can cover. Note that this attribute has no effect on remote media servers or clients (remote media servers apply only to NetBackup DataCenter). The default is 28 days.

`-kt` *days*

The number of days to Keep True-image-recovery (TIR) data. This determines how long to keep TIR information for those classes that have specified that TIR information is to be collected. The default is 1 day.

`-L`

The list type is long. See the section DISPLAY FORMATS for more detail.

`-l`

The list type is short. This is the default if the command line has no list-type option (for instance, if you enter "bpconfig" and a carriage return). See the section DISPLAY FORMATS for more detail.

`-M` *master_server,...*

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.



`-ma [address]`

The mail address for the NetBackup administrator. This is the email address to which NetBackup sends notification of failed automatic backups, administrator-directed manual backup operations, and automatic database backups. The default is NULL (no email address).

If no address is provided, the current setting of the Admin Mail Address is cleared. This means that notification no longer will be sent by email to the NetBackup administrator.

`-mdtm drives`

The maximum drives for this master, the maximum number of drives for this master and remote media server cluster that the master server should consider available when scheduling backups. An appropriate value for this attribute is the physical number of drives, counting shared drives only once, in the master and media server cluster. *drives* must be less than or equal to the number permitted by the version of NetBackup that is installed on the server (that is, 2 for NetBackup BusinessServer and unlimited for NetBackup DataCenter). *drives* is a non-negative integer. The default is 0 (unlimited).

`-mhto seconds`

The multihosted-media-mount timeout, the length of time, in seconds, that NetBackup waits for a shared medium to be mounted, positioned, and become ready on backups and restores. Use this timeout to eliminate excessive waits when a shared medium is being used by another server. The default is 0, which means no timeout (unlimited wait time).

MultiHosted Drives is a separately-priced feature of NetBackup. For more information, see Multihosted Drives (Drive Sharing) in the Media Manager Reference Topics section of the Media Manager system administrator's guide.

`-mj number`

Specifies the maximum jobs per client. This is the maximum number of jobs that a client may perform concurrently. *number* must be a positive integer. The default is 1.

`-period hours`

The time interval associated with the configured number of tries for a backup (see `-tries`). This is the period, in hours, during which NetBackup will attempt a backup job for a client/class/schedule combination for as many tries as configured. *hours* must be a positive integer. The default is 12 hours. Note that this attribute does not apply to user-directed backups and archives.

-prep *hours*

The preprocessing interval. This is the minimum time in hours between client queries to discover new paths if NetBackup is using auto-discover-streaming mode. For additional information, see the “Setting the Preprocess Interval for Auto Discovery” section in the topic on File-List Directives for Multiple Data Streams in the NetBackup system administrator’s guide.

The default Preprocessing Interval value is zero (0). This effectively sets the preprocessing interval to 4 hours.

The preprocessing interval can be set for immediate preprocessing by entering Setting the preprocess interval for auto discovery on the `bpconfig` command line. In this case, `bpconfig` displays the preprocessing interval as 0 hours (that is, the displayed value appears to be the same as the default value, although NetBackup treats it differently).

The maximum Preprocessing Interval is 48 hours.

-to *seconds*

This is the media-mount timeout, the length of time, in seconds, that NetBackup waits for the requested media to be mounted, positioned, and become 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).

The default is 0, which means no timeout (unlimited wait time). If seconds is not 0, its value must be 300 (5 minutes) or greater.

-tries *times*

The number of retries for a backup, during the configured time period (see `-period`). NetBackup tries to run a backup job for a given client/class/schedule combination this many times in the configured period. This allows you to limit the number of backup attempts should repeated failures occur. Note that this attribute does not apply to user-directed backups and archives.

Usually the number of tries should be greater than 0. Specifying 0 for the number of tries is legal but stops all scheduled backups. The default is 2 tries. If defaults are used for both `-tries` and `-period`, NetBackup will attempt the backup 2 times in 12 hours.

-U The list type is user. See the section DISPLAY FORMATS for more detail.

-verbose

Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).



`-wi minutes`

This is the wakeup Interval, the length in time in minutes that the scheduler waits before checking if any automatic backups are scheduled to begin. A long wakeup interval can cause the scheduler to miss too much of the backup window to complete its backups. The default is 10 minutes.

DISPLAY FORMATS

`bpconfig` uses three different formats to display the current values of the NetBackup global configuration attributes.

◆ User Display Format (`-U`)

If the command line contains `-U`, the display format is user. The user display format is the format used by `bpadm` and the NetBackup graphical-user interfaces. This option produces a listing with one global attribute per line. Each line has the form *global attribute descriptor: value*. This listing is similar to the `-L` format, except that the global attribute descriptors are more explicit:

Admin Mail Address
Wakeup Interval
Max Simultaneous Jobs/Client
Backup Tries (x tries in y hours)
Keep Error/Debug Logs
Max drives this master
Keep TrueImageRecovery Info
Compress Image DB Files
Media Mount Timeout
Multihost Media Mount Timeout
Display Reports
Preprocess Interval

◆ Long Format (`-L`)

If the command line contains `-L`, the display format is long. This option produces a listing with one global attribute per line, in the format *global attribute descriptor: value*. The fields in this display are as follows:

Mail Admin
Wakeup Interval
Max Jobs/Client

Backup Tries (x in y hours)

Keep Logs

Max drives/master

Compress DB Files

Media Mnt Timeout

Multihost Timeout

Postprocess Image

Display Reports

Keep TIR Info

Prep Interval

◆ Short Format (-l)

If the `bpconfig` command line contains `-l` or contains no `list-format` option, the display format is short. This produces a terse listing. This option can be useful for scripts or programs that rework the listing into a customized report format. The listing layout is a single line containing the values for all global attributes. The attributes appear in the following order, separated by blanks. For those attributes that are expressed in units of time, the time units follow the attributes in parentheses:

NetBackup administrator email address

Wakeup interval (minutes)

Time period (hours)

Maximum simultaneous jobs per client

Tries per period

Keep logs (days)

Maximum drives this master

Compress image database interval (seconds; 0 denotes no compression)

Media mount timeout (seconds; 0 denotes unlimited)

Multihosted-media-mount timeout (seconds; 0 denotes unlimited)

Postprocess images flag (0 denotes deferred, otherwise immediate)

Display reports from <x> hours ago (hours)

Keep TIR information (days)

Preprocessing interval (hours)

◆ Example of How the Formats Differ



Here is an example of how the display formats differ. `bpconfig` runs with each of the three display formats on a NetBackup installation. The NetBackup global attributes are the same for the three displays.

The first display format, `-U`, looks like this:

```
bpconfig -U
Admin Mail Address:
Wakeup Interval:          1 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:             2 time(s) in 12 hour(s)
Keep Error/Debug Logs:   28 days
Max drives this master:   0
Keep TrueImageRecovery Info: 1 days
Compress Image DB Files:  (not enabled)
Media Mount Timeout:      0 minutes (unlimited)
Multihost Media Mount Timeout: 0 minutes (unlimited)
Display Reports:          24 hours ago
Preprocess Interval:      0 hours
```

The second display format, `-L`, looks like this:

```
example% bpconfig -L
Mail Admin:              *NULL*
Wakeup Interval:         1 minutes
Max Jobs/Client:         1
Backup Tries:            2 in 12 hours
Keep Logs:               28 days
Max drives/master:       0
Compress DB Files:       (not enabled)
Media Mnt Timeout:       0 minutes (unlimited)
Multihost Timeout:       0 minutes(unlimited)
Postprocess Image:       immediately
Display Reports:         24 hours ago
Keep TIR Info:           1 days
Prep Interval:           0 hours
```

The third display format, `-l`, looks like this:




```
bpconfig -l
*NULL* 1 12 1 2 28 0 0 0 1 24 1 0
```

The display fields for the `-l` display are interpreted as follows:

NetBackup administrator email address has not been set
 Wakeup interval is 1 minute
 Time period is 12 hours
 Maximum simultaneous jobs per client is 1
 Tries per period is 2
 Keep logs for 28 days
 Maximum drives this master is 0
 Compress image database interval is 0 seconds; 0 denotes no compression
 Media mount timeout is 0seconds; 0 denotes unlimited
 Multihosted-media-mount timeout is 0 seconds; 0 denotes unlimited
 Postprocess images flag is 1 (immediate)
 Display reports from 24 hours ago
 Keep TIR information for 1 day
 Preprocessing interval is 0 hours

EXAMPLES

◆ Example 1

While running on the master server `kiwi`, display the global attribute settings on the master server `plum`:

```
bpconfig -U -M plum
```

```
Admin Mail Address:          ichabod@null.null.com
Wakeup Interval:            10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 8 hour(s)
Keep Error/Debug Logs:      6 days
Max drives this master:      0
Keep TrueImageRecovery Info: 1 days
Compress Image DB Files:     (not enabled)
```



Media Mount Timeout: 30 minutes
Multihost Media Mount Timeout: 0 minutes (unlimited)
Display Reports: 24 hours ago
Preprocess Interval: 0 hours

◆ **Example 2**

Set the Compress-image-database interval to 604800 seconds, so that NetBackup compresses images more than 7 days old:

```
bpconfig -cd 604800  
bpconfig -U
```

Admin Mail Address:
Wakeup Interval: 10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries: 2 time(s) in 12 hour(s)
Keep Error/Debug Logs: 28 days
Max drives this master: 0
Keep TrueImageRecovery Info: 2 days
Compress Image DB Files: older than 7 day(s)
Media Mount Timeout: 0 minutes (unlimited)
Multihost Media Mount Timeout: 0 minutes (unlimited)
Display Reports: 24 hours ago
Preprocess Interval: 0 hours

◆ **Example 3**

Set the Media Mount Timeout to 1800 seconds.

```
bpconfig -to 1800  
bpconfig -U
```

Admin Mail Address: sasquatch@wapati.edu
Wakeup Interval: 10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries: 1 time(s) in 12 hour(s)
Keep Error/Debug Logs: 3 days
Max drives this master: 0



Keep TrueImageRecovery Info: 24 days
Compress Image DB Files: (not enabled)
Media Mount Timeout: 30 minutes
Multihost Media Mount Timeout: 0 minutes (unlimited)
Display Reports: 24 hours ago
Preprocess Interval: 0 hours

EXIT STATUS

◆ = 0

The command executed successfully.

◆ <> 0

An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form:

`bpconfig: EXIT status = exit status`

If an error occurred, a diagnostic precedes this message.

FILES

`/usr/opensv/netbackup/logs/admin/*`

`/usr/opensv/netbackup/db/config/behavior`

SEE ALSO

`bpimage(1m)`

NetBackup system administrator's guide for information on NetBackup global configuration.

NetBackup Media Manager system administrator's Guide for information on MultiHosted Drives.



bpdbm(1M)

NAME

bpdbm - NetBackup database manager daemon

SYNOPSIS

```
/usr/opensv/netbackup/bin/bpdbm [-verbose] [-terminate]
```

DESCRIPTION

bpdbm responds to queries related to the NetBackup internal databases, which are called catalogs. bpdbm must be running in order for NetBackup commands and utilities to work properly. This daemon executes only on the master server and can be started only by the administrator.

The Netbackup request daemon, bprd, starts bpdbm. You can also start it with the `/usr/opensv/netbackup/bin/initbpdbm` script.

The following events occur when bpdbm starts:

1. bpdbm logs a message indicating that it has started, and then verifies that no other instance of bpdbm is executing. If another bpdbm process is found, the program terminates.
2. bpdbm finds its port number by checking the `services` file for an entry that has a service name of bpdbm and a protocol name of tcp. For example:

```
bpdbm 13721/tcp
```

3. After binding to its port, bpdbm starts responding to queries from bprd and the NetBackup administrative utilities. A child process is created to respond to each query.

OPTIONS

`-verbose` Specifies that bpdbm will write additional information in its daily-activity log for debugging purposes.

`-terminate`

Terminates bpdbm. Any currently executing child process continues to execute until its task is complete.

FILES

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

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

```
/usr/opensv/netbackup/logs/bpdbm/*
```



`/usr/opensv/netbackup/bin/initbpdm`

SEE ALSO

`bpadm(1M)`, `bprd(1M)`



bpduplicate(1M)

NAME

bpduplicate - create a copy of backups created by Netbackup.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpduplicate [-dstunit
  destination_storage_unit_label] [-p | -pb | -PM | -PD]
  [-Bidfile file_name] [-v] [-client name] [-st sched_type]
  [-sl sched_label] [-L output_file] [-dp destination_pool_name]
  [-shost source_host] [-class name] [-s mm/dd/yy
  hh:mm:ss] [-e mm/dd/yy hh:mm:ss] [-ct class_type]
  [-hoursago hours] [-cn copy_number] [-M master_server]
  [-backupid backup_id] [-id media_id] [-mpx]

/usr/opensv/netbackup/bin/admincmd/bpduplicate -npc
  new_primary_copy -backupid backup_id [-client name]
```

DESCRIPTION

The `bpduplicate` command allows a copy of a backup to be created. The `bpduplicate` command can also change the primary copy in order to enable restoring from a duplicated backup. The primary copy is used to satisfy restore requests and is initially the original copy.

Multiplexed duplications can be created by using the `-mpx` option. Refer to the discussion of the `-mpx` option for more information.

The duplicated backup has a separate expiration date from the original. Initially, the expiration date of the copy is set to the expiration date of the original. You can change the expiration date of the copy or the original by using the `bpexpdate(1M)` command.

Only two copies of a backup are allowed: the original copy and the duplicated copy.

For more information on duplicating backups, see your NetBackup system administrator's guide.

OPTIONS

```
-dstunit destination_storage_unit_label
  Specifies the destination storage unit. This parameter is required to
  duplicate backups. Do not specify this option to preview backups to be
  duplicated (-p, -pb, -PM, or -PD options) or to change the primary copy
  (-npc option). This option does not have a default.
```



-
- `-Bidfile` *file_name*
file_name specifies a file that contains a list of backup IDs to be duplicated. List one backup ID per line in the file. If this parameter is specified, other selection criteria is ignored.
- `-p` Previews backups to be duplicated according to the option settings, but does not perform the duplication. Displays the media IDs, server name, backups that are not candidates for duplication (and why), and information about the backups to be duplicated.
- `-pb` Previews the duplication but does not perform the duplication. Similar to the `-p` option, but does not display information about the backups.
- `-PM` Displays information on the backups to be duplicated according to the option settings, but does not perform the duplication. This format first displays the backup IDs that cannot be duplicated and why (for example, because the backup already has two copies). It then displays the following information about the backup: date and time of the backup, class, schedule, backup ID, host, and media ID or path.
- `-PD` Same as the `-PM` option, except the backups are sorted and displayed by date and time (newest to oldest).
- `-v` Selects verbose mode. When specified, the debug and progress logs include more information.
- `-client` *name*
Specifies the name of the client that produced the originals and is used as search criteria for backups to duplicate. The default is all clients.
When specified with the `-npc` option in order to change the primary copy, this indicates that NetBackup will first search for the backup ID belonging to the specified client. This is useful if the client name has changed.
- `-st` *sched_type*
Search for backups to duplicate that were created by the specified schedule type. The default is any schedule type.
Valid values are:
FULL (full backup)
INCR (differential-incremental backup)
CINC (cumulative-incremental backup)
UBAK (user backup)
UARC (user archive)
NOT_ARCHIVE (all backups except user archive)



- `-sl` *sched_label*
Search for backups to duplicate that were created by the specified schedule. The default is all schedules.
- `-L` *output_file*
Specifies the name of a file in which to write progress information. The default is to not use a progress file.
- `-dp` *destination_pool_name*
Specifies the volume pool for the duplicates. Netbackup does not verify that the media ID selected for the duplicate copy is not the same media ID where the original resides. Therefore, to avoid the possibility of a deadlock, specify a different volume pool than where the original media ID resides. The default pool name is NB_duplicates.
- `-shost` *source_host*
Specifies that only the backups created on the specified backup server are considered for duplication. The default is to consider all backups regardless of the backup server.
- `-class` *name*
Search for backups to duplicate in the specified class. The default is all classes.
- `-s` *mm/dd/yy* [*hh[:mm[:ss]]*]
Specifies the start of the range of dates and times that include all backups to duplicate. The `-e` option specifies the end of the range. The date and time format depend on the user's locale. For more information, see `locale(1)`. Default is 24 hours prior to the current date and time.
- `-e` *mm/dd/yy* [*hh[:mm[:ss]]*]
Specifies the end of the range of dates and times that include all backups to duplicate. The `-s` option specifies the start of the range. The date and time format depend on the user's locale. For more information, see `locale(1)`. Default is the current date and time.
- `-ct` *class_type*
Search for backups created by the specified class type. The default is any class type.
Valid values are:
Informix-On-BAR
Oracle
Macintosh
MS-Exchange-Server
MS-Windows
MS-Windows-NT

MS-SQL-Server
 NDMP
 Netware
 OS/2
 Standard
 Sybase

Note The following class types (AFS, Apollo-wbak, Auspex-FastBackup, DataTools-SQL-BackTrack, DB2, FlashBackup, SAP, and Split-Mirror) apply only to NetBackup DataCenter.

AFS
 Apollo-wbak
 Auspex-FastBackup
 DataTools-SQL-BackTrack
 DB2
 FlashBackup
 SAP
 Split-Mirror

- hoursago** *hours*
 Specifies number of hours prior to the current time to search for backups. Do not use with the **-s** option.
- cn** *copy_number*
 Determines the copy number to duplicate. Valid values are 1 or 2. The default is 1.
- M** *master_server*
 Specifies the master server that manages the media catalog that has the media ID. If this option is not specified, the default is one of the following:
 For NetBackup BusinessServer:
 NetBackup BusinessServer supports only one server (the master) with no remote media servers. Therefore, the default in this case is always the NetBackup BusinessServer master where you execute the command.
 For NetBackup DataCenter:
 If the command is executed on a master server, then that server is the default.
 If the command is executed on a media server that is not the master, then the master for that media server is the default.



- `-backupid backup_id`
Specifies the backup ID of a single backup to duplicate or for which to change the primary copy.
- `-id media_id`
Search the image catalog for backups to duplicate that are on this media ID. If the original is fragmented between different media IDs, NetBackup duplicates only the fragments that exist on the specified media ID.
- `-mpx`
Specifies that when duplicating multiplexed backups, NetBackup will create multiplexed backups on the destination media. This reduces the time to duplicate multiplexed backups.
Multiplexed duplication is not supported for:
- Non-multiplexed backups
 - Backups from disk type storage units
 - Backups to disk type storage units
 - Auspex-FastBackup, FlashBackup, or NDMP backups
- If backups in the above categories are encountered during duplication, NetBackup duplicates them first and uses non-multiplexed duplication. Then, the multiplexed backups are duplicated by using multiplexed duplication.
- If all the backups in a multiplexed group are not duplicated, the duplicated multiplexed group will have a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single multiplexing session.)
- If this option is not specified, all backups are duplicated using non-multiplexed duplication.
- For more information on multiplexing, see your NetBackup system administrator's guide.
- `-npc new_primary_copy`
Allows the primary copy to be changed. The value can be 1 or 2. The `-backupid` option must be specified with this option.

EXAMPLES

Example 1

The following command (all on one line) lists backups with a copy number of 1, that were backed up by the class named `stdclass`, and created between July 1, 1998, and August 1, 1998.

```
bpduplicate -PM -cn 1 -class stdclass -s 07/01/98 -e 08/01/98
```

Example 2



The following command (all on one line) duplicates copy 1 of the backups listed in file /tmp/bidfile. The destination storage unit is unit1 and the destination pool is dup_pool. Progress information is written to /tmp/bpdup.ls.

```
bpduplicate -dstunit unit1 -Bidfile /tmp/bidfile -L /tmp/bpdup.ls  
-dp dup_pool -cn 1
```

Example 3

The following command (all on one line) is the same as the prior example, except multiplexed backups are duplicated using multiplexed duplication.

```
bpduplicate -dstunit unit1 -Bidfile /tmp/bidfile -mpx -L  
/tmp/bpdup.ls
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

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



bpcerror(1M)

NAME

bpcerror - display NetBackup status and troubleshooting information or entries from the NetBackup error catalog.

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bpcerror {-S | -statuscode
      status_code} [-r|-recommendation] [-p | -platform Unx |
      NTx] [-v]

/usr/opensv/netbackup/bin/admincmd/bpcerror [-all | -problems
      |-media | {-backstat [-by_statcode]}] [-L | -l | -U]
      [-columns ncols] [-d mm/dd/yyyy hh:mm:ss | -hoursago
      hours] [-e mm/dd/yyyy hh:mm:ss] [-client client_name]
      [-server server_name] [-jobid job_id] [-M
      master_server, ...] [-v]

/usr/opensv/netbackup/bin/admincmd/bpcerror [-s
      {severity[+]}|severity ...] [-t type ...] [-L | -l|-U]
      [-columns ncols] [-d mm/dd/yyyy hh:mm:ss | -hoursago
      hours] [-e mm/dd/yyyy hh:mm:ss] [-client client_name]
      [-server server_name] [-jobid job_id] [-M
      master_server, ...] [-v]

```

DESCRIPTION

bpcerror displays information from either the same source as the online troubleshooter (in the Activity Monitor or Reports applications) or from the NetBackup error catalog. bpcerror provides the following types of displays:

- ◆ A display of the message that corresponds to a status code and, optionally, a recommendation on how to troubleshoot the problem. In this case, the display results come from the same source as the online troubleshooter for the local system.
- ◆ A display of the error catalog entries that satisfy the command-line options. For instance, bpcerror can display all the problem entries for the previous day.
- ◆ A display of the error catalog entries that correspond to a particular message severity and/or message type.

For information on details of the displays, see DISPLAY FORMATS later in this command description.

bpcerror writes its activity log information to the `/usr/opensv/netbackup/logs/admin` directory. You can use the information in this directory for troubleshooting.



The output of bpcerror goes to standard output.

Only root can execute this command.

OPTIONS

`-all`

`-backstat [-by_statcode]`

`-media`

`-problems`

These options specify the type and severity of log messages to display. The default type is ALL. The default severity is ALL.

For `-all`: The type is ALL, and severity is ALL. Executing bpcerror with this option and `-U` produces an All Log Entries report.

For `-backstat`: The type is BACKSTAT, and severity is ALL. If `-by_statcode` is present, the display contains one entry for each unique status code. Line 1 of the entry contains the status code and the corresponding message text. Line 2 of the entry contains the list of clients for which this status code occurred. `-by_statcode` is only valid when the command line contains both `-backstat` and `-U`. Executing bpcerror with this option and `-U` produces a Backup Status report.

For `-media`: The type is MEDIADEV, and severity is ALL. Executing bpcerror with this option and `-U` produces a Media Logs report.

For `-problems`: The type is ALL, and severity is the union of WARNING, ERROR, and CRITICAL. Executing bpcerror with this option and `-U` produces a Problems report.

`-client client_name`

Specifies the name of a NetBackup client. This name must be as it appears in the NetBackup catalog. By default, bpcerror searches for all clients.

`-columns ncols`

For the `-L` and `-U` reports, `-columns` provides an approximate upper bound on the maximum line length. bpcerror does not attempt to produce lines exactly *ncols* characters in length.

`-columns` does not apply to the `-l` report.

ncols must be at least 40. The default is 80.

`-d mm/dd/yy [hh:mm:ss]`

`-e mm/dd/yy [hh:mm:ss]`

Specifies the start and end date range for the listing.



Note The locale setting for your system affects the way you must specify dates and times. See the NOTES section later in this command description.

`-d` specifies a start date and time for the listing. The resulting list shows only images in back ups or archives that occurred at or after the specified date and time. Use the following format:

`mm/dd/yy [hh[:mm[:ss]]]`

The valid range of dates is from 01/01/70 00:00:00 to 01/19/2038 03:14:07. The default is 24 hours before the current date and time.

`-e` specifies an end date and time for the listing. The resulting list shows only files from backups or archives that occurred at or before the specified date and time. Use the same format as for the start date. The default is the current date and Time. The end date must be greater than or equal to the start date.

`-help` Print a command-line usage message. When `-help` is present, it is the only option on the command line.

`-hoursago hours`

Specifies a start time of this many hours ago. This is equivalent to specifying a start time (`-d`) of the current time minus hours. Hours is an integer. The default is 24, meaning a start time of 24 hours before the current time.

`-jobid job_id`

Specifies a NetBackup job ID. By default, `bpcerror` searches for all job IDs.

`-L` Report in long format.

`-l` Report in short format. This produces a terse listing. This option is useful for scripts or programs that rework the listing contents into a customized report format. This is the default list type.

`-M master_server`

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.

- `-p Unx | NTx`
- `-platform Unx | NTx`
 Display the message that applies to the platform (UNIX or Windows NT) for the specified status code. The default is to display the message for the platform on which `bpcerror` is running. The `-S` or `-statuscode` option must be specified when using this option.
- `-r | -recommendation`
 Display the recommended action for the specified status code from the troubleshooting guide. The default is not to display the recommended action. The `-S` or `-statuscode` option must be specified when using this option.
- `-S status_code`
- `-statuscode status_code`
 Display the message that corresponds to the status code. There is no default for this option.
- `-s severity`
- `-s severity+`
 Specifies the severity of log messages to display. The defined values are ALL, DEBUG, INFO, WARNING, ERROR, and CRITICAL.
 There are two ways to specify severity. The first way is a list of one or more severity values. For instance, "`-s INFO ERROR`" displays the messages with either severity INFO or severity ERROR. The delimiter between the elements in the list must be a blank (" "). The second way is a single severity value with "+" appended, meaning this severity or greater. For instance "`-s WARNING+`" displays the messages with severity values WARNING, ERROR, and CRITICAL.
 The default is ALL. The severity value can be in either upper or lower case.
- `-server server_name`
 Specifies the name of a NetBackup server. This name must be as it appears in the NetBackup catalog. The display is limited to messages logged for this server, which also satisfy the other criteria specified by `bpcerror` options. For instance, if `-server plum` and `-hoursago 2` are `bpcerror` options, the display contains messages logged for the media server plum in the past two hours.
 The server name must match the server name recorded in the log messages. For instance, if the logs record the server name as plum.null.null.com, specifying `-server plum` will not display the logs, but `-server plum` will.



The query goes to the error catalog residing on the master server (either the local master server or the master server specified by `-M`). The master server must allow access by the system running `bpcerror`.

The default is to display log messages for all media servers known to the master server(s).

`-t type`

Specifies the type of log messages to display. The defined values are `ALL`, `BACKSTAT`, `MEDIADEV`, `GENERAL`, `BACKUP`, `ARCHIVE`, `RETRIEVE`, and `SECURITY`. The default is `ALL`. The type value can be in either upper or lower case. The type value is entered as a list of one or more values. For instance, "`-t BACKSTAT MEDIADEV`" displays the messages with either type `BACKSTAT` or type `MEDIADEV`. The delimiter between the elements in the list must be a blank (" ").

`-U`

Report in user format. This is the report format used by NetBackup report-generating tools such as the NetBackup-Java Reports application.

`-v`

Selects verbose mode. This option causes `bpcerror` to log additional information for debugging purposes. The information goes into the NetBackup-administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/openv/netbackup/logs/admin` directory defined). The default is not to be verbose.

DISPLAY FORMATS

STATUS CODE DISPLAY (for example, `bpcerror -S status_code`):

`bpcerror` queries the NetBackup online troubleshooter on the local system for the message that corresponds to the status code. `bpcerror` displays the message text on one line and an explanation on a second line.

If `-r` or `-recommendation` is an option, `bpcerror` also queries for the troubleshooting recommendation that corresponds to the status code. `bpcerror` displays the recommendation following the status message, on one or more lines.

ERROR CATALOG DISPLAY (for example, `bpcerror -all`; `bpcerror -s severity`):

`bpcerror` queries the NetBackup error catalog on either the local master server or the master servers in the `-M` option list. The display consists of the results returned from querying the error catalog on the master server(s). The results are limited to catalog entries that satisfy all the `bpcerror` options. For instance, if the `bpcerror` command line contains options for client, start time, and end time, then `bpcerror` reports only the jobs run for that client between the start and end times. For the display variant that shows individual message entries from the error catalog, the display can appear in long (`-L`),

user (-U), or short (-l) format. For the display variant that categorizes by status code, the display can appear in user (-U) format only. The display content for each of these formats is as follows:

- ◆ Error catalog display, individual message entries, long format (for example, `bpcerror -media -L`). This report produces several lines per log entry, with the following contents:

Line 1: Date and time

V:NetBackup version

S:Server

C:Client

J:Job ID

(U:Job group ID and an unused field) If multi-streaming is enabled for a class, the job group ID is the job ID of the first job that spawned a collection of multi-streaming backups; if multi-streaming is disabled for the class, the job group ID is always zero.

Line 2: Severity (severity name and severity code in base 16)

Type (type name and type code in base 16)

Who (name of the entity that added the log entry)

Line 3: Text (beginning of the log message text, continued on succeeding lines if necessary)

- ◆ Error catalog display, individual message entries, user format (for example., `bpcerror -media -U`). The user format produces a header line showing column names, and then one or more lines per log entry, with the following contents:

Line 1: Date and time

Server

Client

Text (beginning of the log message text, continued on succeeding lines if necessary)

- ◆ Error catalog display, individual message entries, short format (for example., `bpcerror -media -l`). The short format produces a single line per log entry, with the following contents:

Line 1: Time (internal system representation)

NetBackup version

Type code (base 10)



Severity code (base 10)

Server

Job ID

Job Group ID

An unused field

Client

Who

Text (the entire log message text, with no truncation of the line length)

- ◆ Error catalog display categorized by status code. This display reports only each unique status code, instead of listing every log entry for that status code (for example, `bpcerror -backstat -by_statcode -U`). This produces two or more lines per status code, with the following contents:

Line 1: Status code

Text (the beginning of the log message text, continued on succeeding lines if necessary)

Line 2: The list of clients for which this status occurred.

NOTES

The format that you must use for date and time values in NetBackup commands vary according to your locale setting. The examples in this man page are for a locale setting of C. If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the USAGE. For example, here is the beginning of the `bpcerror` usage output:

```
USAGE: bpcerror {-S|-statuscode status_code}
        [-r|-recommendation] [-p|-platform Unx|NTx] [-v]
bpcerror [-all|-problems|-media|{-backstat [-by_statcode]}]
        [-L|-l|-U] [-columns ncols]
        [-d mm/dd/yyyy hh:mm:ss|-hoursago hours]
        [-e mm/dd/yyyy hh:mm:ss] [-client client_name]
        [-server server_name][-jobid job_id]
        [-M master_server,...] [-v]
bpcerror [-s {severity[+]}|severity ...] [-t type ...]
        [-L|-l|-U] [-columns ncols]
        [-d mm/dd/yyyy hh:mm:ss|-hoursago hours]
```

```
[-e mm/dd/yyyy hh:mm:ss][-client client_name]
[-server server_name] [-jobid job_id]
[-M master_server,...] [-v]
```

Notice the month/day/year and hours:minutes:seconds requirements for the `-d` and `-e` options. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

◆ Example 1

Here `bpcerror` displays the error for a job that failed because the NetBackup encryption package was not installed. Status code 9 is the NetBackup status code for this failure. The second execution of `bpcerror` displays the action recommended for NetBackup status code 9.

```
bpcerror -d 12/23/99 16:00:00 -e 12/23/99 17:00:00 -t backstat -U
STATUS CLIENT      CLASS      SCHED      SERVER      TIME COMPLETED
   9    plum      jdencrypt  user      plum      12/23/99 16:38:09
                                (an extension package is needed, but was not installed)
```

```
bpcerror -S 9 -r
```

```
an extension package is needed but was not installed
```

```
A NetBackup extension product is required in order to perform the
requested operation.
```

```
Install the required extension product.
```

◆ Example 2

Here `bpcerror` reports, in User format, the problems that have occurred in the previous 24 hours.

```
bpcerror -U -problems
```

```
      TIME          SERVER/CLIENT          TEXT
11/23/99 16:07:39 raisin - no storage units configured
11/23/99 16:07:39 raisin - scheduler exiting - failed reading storage
unit database information (217)
11/23/99 16:17:38 raisin - no storage units configured
11/23/99 16:17:38 raisin - scheduler exiting - failed reading storage
unit database information (217)
```



```
11/23/99 16:26:17 raisin - WARNING: NetBackup database backup is
currently disabled
11/23/99 18:11:03 raisin nut bpcd on nut exited with status 59:
access to the client was not allowed
11/23/99 18:11:20 raisin - WARNING: NetBackup database backup is
currently disabled
```

◆ Example 3

This example displays status for type `backstat` for jobs run in the previous 24 hours. The option `-by_statcode` produces a display organized by status code.

The display shows that one or more jobs for each of the clients `chives`, `guava`, `plum`, and `raisin` completed successfully (the status code is 0). In addition, one or more jobs for client `nut` failed because `nut` did not allow access by the master or media server (the status code is 59).

```
bpcerror -U -backstat -by_statcode
 0 the requested operation was successfully completed
    chives guava plum raisin
59 access to the client was not allowed nut
```

◆ Example 4

This example identifies and retrieves the results for a particular user job. It first lists the log entries with job Ids other than zero. It then runs a `User-format` report on the job of interest.

```
bpcerror -hoursago 2000 -L | grep 'S:' | egrep 'J\:[1-9]'
```

```
12/21/99 17:24:14 V1 S:plum C:plum J:1 (U:0,0)
12/23/99 16:31:04 V1 S:plum C:plum J:1 (U:0,0)
12/23/99 16:31:06 V1 S:plum C:plum J:1 (U:0,0)
12/23/99 16:38:04 V1 S:plum C:plum J:3 (U:0,0)
12/23/99 16:38:07 V1 S:plum C:plum J:3 (U:0,0)
12/23/99 16:38:08 V1 S:plum C:plum J:3 (U:0,0)
12/23/99 16:38:09 V1 S:plum C:plum J:3 (U:0,0)
01/07/00 13:12:31 V1 S:plum C:plum J:34 (U:0,0)
01/07/00 13:12:36 V1 S:plum C:plum J:34 (U:0,0)
01/07/00 13:12:40 V1 S:plum C:plum J:34 (U:0,0)
01/07/00 13:12:41 V1 S:plum C:plum J:34 (U:0,0)
```

```
bpcerror -d 1/7/00 -jobid 34 -U
```

TIME	SERVER/CLIENT	TEXT
01/07/00 13:12:31	plum plum	started backup job for client plum, class jdencrypt, schedule user on storage unit jdencrypt
01/07/00 13:12:36	plum plum	begin writing backup id plum_0947272350, copy 1, fragment 1
01/07/00 13:12:40	plum plum	successfully wrote backup id plum_0947272350, copy 1, fragment 1, 32 Kbytes at 11.057 Kbytes/sec
01/07/00 13:12:41	plum plum	CLIENT plum CLASS jdencrypt SCHED user EXIT STATUS 0 (the requested operation was successfully completed)

◆ Example 5

This example shows the media entries in the error catalog for the past 2000 hours.

```
bpcerror -hoursago 2000 -media -U
```

TIME	SERVER/CLIENT	TEXT
12/23/99 16:31:04	plum plum	media manager terminated during mount of media id A00000, possible media mount timeout
12/24/99 04:31:20	plum -	media id A00000 removed from media manager database (manual deassign)

◆ Example 6:

This example tallies and reports the total number of bytes backed up in the past 24 hours.

```
bpcerror -all -hoursago 24 | grep "successfully wrote backup id" | awk
'bytes= bytes + $20} END {print "backed up",bytes," Kbytes of data"}'
backed up 64 Kbytes of data
```

◆ Example 7

This example reports the performance, in Kbytes per second, for each of today's backups:

```
bpcerror -all | grep Kbytes
0912013673 1 4 4 hat 0 0 0 hat bptm successfully wrote backup id
hat_0912013584, copy 1, fragment 1, 32256 Kbytes at 891.222 Kbytes/sec
```



```
0912014210 1 4 4 hat 0 0 0 hat bptm successfully wrote backup id
hat_0912014132, copy 1, fragment 1, 32256 Kbytes at 1576.848
Kbytes/sec
```

```
0912016068 1 4 4 hat 0 0 0 hat bptm successfully wrote backup id
hat_0912015780, copy 1, fragment 1, 603136 Kbytes at 2645.960
Kbytes/sec
```

◆ Example 8

Here `bperror` displays the status message and the recommended action for status code 0:

```
bperror -S 0 -r
```

```
the requested operation was successfully completed
```

```
There were no problems detected with the requested operation.
```

None, unless this was a database backup performed through a database extension product (for example, NetBackup for Oracle or NetBackup for SQL Server). In those instances, code 0 means the backup script that started the backup ran without error. However, you must check other status as explained in the related NetBackup manual to see if the database was successfully backed up.

FILES

```
/usr/opensv/netbackup/logs/admin/log.mmddyy
```

```
/usr/opensv/netbackup/db/error/log files
```

```
/usr/opensv/msg/locale/netbackup/TrbMsgs
```

```
/usr/opensv/msg/C/netbackup/TrbMsgs
```

```
/usr/opensv/msg/.conf
```



bpexpdate(1M)

NAME

bpexpdate - change the expiration date of backups in the image catalog and media in the media catalog.

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bpexpdate -ev media_id -d
    date_time [-host name] [-force] [-M
    master_server, . . . , master_server]

/usr/opensv/netbackup/bin/admincmd/bpexpdate -deassignempty [-ev
    media_id] [-host name] [-force] [-M
    master_server, . . . , master_server]

/usr/opensv/netbackup/bin/admincmd/bpexpdate -backupid backup_id
    -d date_time [-client name] [-copy number] [-force] [-M
    master_server, . . . , master_server]

/usr/opensv/netbackup/bin/admincmd/bpexpdate -recalculate -d
    date_time [-ret retention_level] [-backupid backup_id]
    [-client name] [-class name] [-sched type] [-M
    master_server, . . . , master_server]

```

DESCRIPTION

NetBackup maintains internal databases with backup image and media information. These internal databases are called catalogs. Both an image record in the image catalog and a media ID in the media catalog contain an expiration date. The expiration date is the date and time when NetBackup removes the record for a backup or media ID from the corresponding catalog.

The `bpexpdate` command allows the expiration date and time of backups to be changed in the NetBackup image catalog. It is also used to change the expiration of removable media in the NetBackup media catalog. If the date is set to zero, `bpexpdate` immediately expires backups from the image catalog or media from the media catalog. When a media ID is removed from the NetBackup media catalog, it is also deassigned in the Media Manager volume database, regardless of the media's prior state (FROZEN, SUSPENDED, and so on).

Changing the expiration can be done on a media ID basis or on an individual backup ID basis. Changing the expiration date of a media ID also causes the expiration date of all backups on the media to be changed. `bpexpdate` also provides options to deassign media from the media catalog if they no longer contain valid backups and to recalculate the expiration date based on the configured or a supplied retention level.

The different formats of the command are described below.



- ◆ **ev**

Changes the expiration date or removes the media ID from the media catalog and associated backups from the Netbackup catalog. A separate expiration date is maintained in the image catalog for each copy of a backup. When this format is used, only the expiration of the copy on the media is affected. If the media ID is removed from the media catalog by specifying a zero date, the media ID is also deassigned in the Media Manager volume database.
- ◆ **deassignempty**

Searches the catalog for removable media that no longer contain valid backups, removes it from the media catalog, and deassigns the media IDs in the Media Manager catalog. The media is then available to be used again. You can use the NetBackup Images on Media report to determine if there are assigned media that no longer contain valid backups.
- ◆ **backupid**

Changes the expiration of a single backup. If the date is zero, the backup is removed from the image catalog. If the backup is on removable media and the expiration date given by the `-d` option is greater than the current expiration of the media ID, the expiration date of the media ID in the media catalog is also changed. The change affects all copies of a backup, unless the `-copy` option is used. The `-copy` option causes only the specified copy to be affected.
- ◆ **recalculate**

Allows the expiration date of backups to be changed based on the specified retention level or you can specify a new expiration date. When the expiration is changed according to retention level, the new date is calculated based on the creation date of the backup plus the value of the retention level. The expiration can be changed for a single backup, or for all backups for a particular client, class, or schedule type.

If the backup is on removable media, the expiration date of the media ID in the media catalog is changed, providing the expiration date on this command is greater than the current expiration of the media ID. For more information on retention levels, see your NetBackup system administrator's guide.

OPTIONS

`-ev media_id`

Specifies the media ID that is affected by the expiration date change. The expiration dates of the backups on the media ID are also changed. The `-d` option must be included with this option.

This option can also be used when the `-deassignempty` option is specified to check if valid backups exist on this particular media ID. In this case, do not include the `-d` option.

The media ID must be six or less characters and must be in the NetBackup media catalog.

- deassignempty
Expires removable media from the media catalog when that media no longer contains valid backups and also deassigns the media ID in the Media Manager catalog.
- backupid *backup_id*
Specifies the backup ID of the backup that is affected by the expiration date change. The `-d` option must be included with this option.
This option can also be used to specify a backup ID for the `-recalculate` option. In this case, the `-d` option is not required.
- recalculate
Recalculates the expiration of backups based on the retention level or you can specify a new expiration date. Other options can be included in order to change the expiration for a single backup, or for all backups for a specific client name, class name, or schedule type. Either the `-d` or `-ret` option must be specified with this option.
- d *date_time*
Specifies the expiration date and time. *date_time* can be any one of the following:
mm/dd/yy hh:mm:ss
or
0
or
infinity
If 0 is specified, the backup or media is expired immediately. If *infinity* is specified the backup is never expired. The date and time format is also dependent on the user's locale. For more information, see `-backupid`.
- host *name*

Note For NetBackup BusinessServer this option is not required because there is only one server (the master), so if you do use the option specify the host name of that server.

Specifies the host name of the server where the media catalog resides. This option is required only if the master has remote media servers and the volume was not written on the server where you execute the `bpexpdate` command. In this case, the media ID is in the NetBackup media catalog on the server where the media was written and you must specify the name of that server on the `bpexpdate` command.



For example, assume you have a master server named whale and a media server named eel. You execute the following `bpexpdate` command on whale in order to manually remove media ID BU0001 from the media catalog, and all corresponding backups from the image catalog:

```
bpexpdate -ev BU0001 -d 0 -host eel
```

You can use the NetBackup Media List report to determine which server's media catalog has the volume.

- `-force` Prior to executing the specified operation, `bpexpdate` queries before starting the operation. This option forces `bpexpdate` to carry out the operation without querying the user.
- `-client name`
Specifies the client name for the `-backupid` and `-recalculate` operations.
For the `backupid` operation, this option causes NetBackup to first search for the backup ID for the specified client, which is useful if the client name has changed.
For the `recalculate` operation, this option causes NetBackup to recalculate the expiration date based on the retention level for all the specified client backups.
- `-copy number`
Expires or changes the expiration date of the specified copy number and is valid only with the `-backupid` option. Valid values are 1 or 2.
If the primary copy is expired, the other copy becomes the primary copy. If this option is not specified, the expiration affects both copies of the backup.
- `-class name`
Specifies the class name and is valid with the `-recalculate` option. When specified, the expiration is recalculated based on the retention level for all backups created in this class.
- `-ret retention_level`
Specifies the retention level to use when recalculating expiration dates and is valid with the `-recalculate` option. Levels range from 0 to 9. The new expiration date is determined by adding the configured retention level value to the backup's creation date. Either the `-backupid` or `-class` option must be specified with this option.
- `-sched type`
Specifies the schedule type and is valid with the `-recalculate` option. When specified, the expiration is recalculated based on the retention level for all backups created with this schedule type. Enter a numeric value for type as follows:

- 0 = Full
- 1 = Differential Incremental
- 2 = User Backup
- 3 = User Archive
- 4 = Cumulative Incremental

The `-class` option must be specified with `-sched`.

`-M master_server, . . . , master_server`

Specifies the master server that manages the media catalog that has the media ID. If this option is not specified, the default is one of the following:

For NetBackup BusinessServer:

NetBackup BusinessServer supports only one server (the master) with no remote media servers. Therefore, the default in this case is always the master server where you execute the command.

For NetBackup DataCenter:

If the command is executed on a master server, then that server is the default.

If the command is executed on a media server that is not the master, then the master for that media server is the default.

NOTES

Some options in large environments can take a significant amount of time to complete. Changes that cause backups or media to expire are irrevocable; importing backups and (or) recovering previous versions of the catalogs can be required if mistakes are made using this command.

The `bpexpdate` command itself does not necessarily make modifications to the catalogs. Therefore, aborting the command will not produce the desired or expected result.

EXAMPLES

Example 1

The following command, executed on the master server, removes media ID BU0002 from the media catalog, and deassigns the media ID in the Media Manager catalog. It also expires associated image records in the image catalog.

```
bpexpdate -ev BU0002 -d 0
```

Example 2

The following command (all on one line) changes the expiration of copy 2 of backupid classA_0904219764_FULL. The expiration of copy 1 of the backup is not affected.



```
bpexpdate -backupid classA_0904219764_FULL -d 12/20/99 08:00:00  
-copy 2
```

Example 3

The following command removes the backup from the image catalog. Since the `-copy` option is not specified, all copies are removed.

```
bpexpdate -backupid classA_0904219764_FULL -d 0
```

Example 4

The following command checks for all media in host cat's media catalog that are still assigned but no longer contain valid backups. If any such media are found, the command removes them from the media catalog and deassigns them in the Media Manager catalog.

```
bpexpdate -deassignempty -host cat
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

```
/usr/opensv/netbackup/db/media/*
```

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

bpimagerlist(1M)

NAME

bpimagerlist - lists backed up and archived files on the NetBackup server

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpimagerlist -l | -L | -U |
  -idonly [-d mm/dd/yy hh:mm:ss | -hoursago hours ] [-e
  mm/dd/yy hh:mm:ss] [-keyword "keyword phrase"] [-client
  client_name] [-backupid backup_id] [-option option_name]
  [-class class_name] [-ct class_type] [-rl retention_level]
  [-sl sched_label] [-st sched_type] [-M master_server,...]
  [-v]
```

```
/usr/opensv/netbackup/bin/admincmd/bpimagerlist -media [-l | -L |
  -U | -idonly] [-d mm/dd/yy hh:mm:ss | -hoursago hours]
  [-e mm/dd/yy hh:mm:ss] [-server server_name] [-keyword
  "keyword phrase"] [-client client_name] [-option
  option_name] [-class class_name] [-ct class_type] [-rl
  retention_level] [-sl sched_label] [-st sched_type] [-M
  master_server,...] [-v]
```

DESCRIPTION

bpimagerlist queries the NetBackup catalog and produces a report on the status of the NetBackup images. It will produce one of two types of reports:

- ◆ Report images satisfying a set of criteria (if `-media` is absent)
- ◆ Report on removable media satisfying a set of criteria (if `-media` is present).

bpimagerlist shows a list of previously archived or backed up files according to the options that you specify. You can choose the file or directory and the time period that you want the listing to cover. Directories can be recursively displayed to a specified depth.

The list shows only the files that you have read access to. You also must have read access to all directories in the file paths or you must own the directories. You can list files that were backed up or archived by another client if you are validated to do so by the NetBackup administrator.

bpimagerlist writes its activity log information to the `/usr/opensv/netbackup/logs/admin` directory. You can use the information in this directory for troubleshooting.

The output of bpimagerlist goes to standard output.

This command requires root privileges.



OPTIONS

Report-type options

- `-media` Specifies that the listing reports on removable media satisfying a set of criteria. If `-media` is not present, the report is on images, not media, satisfying a set of criteria.

Report-format options:

- `-U` Report in User mode. The report is formatted, it includes a banner listing the column titles, and the status is a descriptive term instead of a number.
- `-L` Report in Long mode. For instance, for the Media List report, the report lists the information for each media ID as a series of *attribute = value* pairs, and the density value is provided as both a descriptive term and a number.
- `-l` Report in Short mode. This produces a terse listing. This option is useful for scripts or programs that rework the listing contents into a customized report format.
- `-idonly` Produce an abbreviated listing. For an image listing, the listing contains the creation time, backup ID, and schedule type of each image. For instance, if the listing criterion is a window of time image listing contains, for each image created in this window, only the creation time, backup ID, and schedule type of the image.
- For a media listing, the listing contains only the applicable media IDs. For instance, if the listing criterion is a window of time, the listing contains only the media IDs written in this window.

The following options represent the criteria that determine which images or media are selected for the report. Where images are discussed in these options, media can be substituted if this is a media report.

- `-hoursago` *hours*
Include images written up to this many hours ago. This is equivalent to specifying a start time (`-d`) of the current time minus *hours*. *hours* must be 1 or greater.
- `-option` *option_name*
Specifies a criterion for finding images to list. *option_name* is one of the following character strings, in either upper-or lower-case:
- INCLUDE_PRE_IMPORT - Include images that have completed phase one of an import. Refer to the `bpimport(1M)` command description or your NetBackup system administrator's guide for more information.
- ONLY_PRE_IMPORT - Include only images that have completed phase one of an import.



INCLUDE_TIR - Include images that were created by true-image-recovery backups. Refer to the `bpclinfo(1M)` command description or your NetBackup system administrator's guide for more information on this topic.

ONLY_TIR - Include only images that were created by true-image-recovery backups.

The default is that there are no restrictions on the images selected.

- backupid *backup_id*
Specifies a backup ID to use for finding applicable images (applies only to image listing).
- client *client_name*
Specifies a client name to use for finding backups or archives to list. This name must be as it appears in the NetBackup catalog. By default, `bpimagelist` searches for all clients.
- server *server_name*
Specifies the name of a NetBackup server or ALL. This option applies to the media report (`-media`). If `-server` specifies a server name, the media in the report are only the media which reside on that server and which also satisfy the other criteria specified by `bpimagelist`. For instance, if `-hoursago 2` is specified, the media must contain an image created in the past two hours.

The query goes to the image catalog residing on the local master server. The master server must allow access by the system running `bpimagelist`.

The default is to report all media in the image catalog on the local master server. This is equivalent to specifying `-server ALL`.
- M *master_server, ...*
A list of alternative master servers. This is a comma-delimited list of hostnames. If this option is present, each master server in the list executes the `bpimagelist` command. If an error occurs for any master server, processing stops at that point.

The report is the composite of the information returned by all the master servers in this list. `bpimagelist` queries each of these master servers. The master server returns image or media information from the image catalogs. Each master server must allow access by the system issuing the `bpimagelist` command.

The default is the master server for the system running `bpimagelist`.
- ct *class_type*
Specifies a *class_type*. By default, `bpimagelist` searches for all class types. *class_type* is one of the following character strings:



Informix-On-BAR
MS-Exchange-Server
MS-SQL-Server
MS-Windows-NT
NetWare
Oracle
OS/2
Standard
Sybase
NDMP

Note The following class types (AFS, Apollo-wbak, Auspex-FastBackup, DataTools-SQL-BackTrack, DB2, FlashBackup, SAP, and Split-Mirror) apply only to NetBackup DataCenter.

AFS
Apollo-wbak
Auspex-FastBackup
DataTools-SQL-BackTrack
DB2
FlashBackup
SAP
Split-Mirror

- `-r1 retention_level`
Specifies the *retention_level*. The *retention_level* is an integer between 0 and 9. By default, `bpimagelist` searches for all retention levels.
- `-d mm/dd/yy [hh:mm:ss]`
- `-e mm/dd/yy [hh:mm:ss]`
Specifies the start and end date range for the listing.

Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section later in this man page.

`-d` specifies a start date and time for the listing. The resulting list shows only images in backups or archives that occurred at or after the specified date and time. Use the following format:

`mm/dd/yy [hh[:mm[:ss]]]`

The valid range of dates is from 01/01/70 00:00:00 to 01/19/2038 03:14:07. The default is midnight beginning the current day.



`-e` specifies an end date and time for the listing.

The resulting list shows only files from backups or archives that occurred at or before the specified date and time. Use the same format as for the start date. The default is the current date and time.

`-keyword "keyword_phrase"`

Specifies a keyword phrase for NetBackup to use when searching. The phrase must match the one that has been previously associated with the image. For instance, the `-k` option of the `bpbackup(1)` or `bparchive(1)` command associates a keyword with the image when the image is created.

`-sl sched_label`

Specifies a schedule label for the image selection. The default is all schedules.

`-st sched_type`

Specifies a schedule type for the image selection. The default is any schedule type. Valid values are:

FULL (full backup)

INCR (differential-incremental backup)

CINC (cumulative-incremental backup)

UBAK (user backup)

UARC (user archive)

NOT_ARCHIVE (all backups except user archive)

`-class name`

Searches for backups to import in the specified class. The default is all classes.

Other options:

`-help` Prints a synopsis of command usage when it is the only option on the command line.

`-v` Selects verbose mode. This option causes `bpimagelist` to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (`/usr/opensv/netbackup/logs/admin` directory defined).

NOTES

The format that you must use for date and time values in NetBackup commands will vary according to the locale setting. The examples in this man page are for a locale setting of C.



If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the USAGE. For example:

```
/usr/openv/netbackup/bin/admincmd/bpimagelist -help
USAGE: bpimagelist [-media] -l|-L|-U|-idonly
[-d mm/dd/yy hh:mm:ss] [-e mm/dd/yy hh:mm:ss]
[-hoursago hours] [-keyword "keyword phrase"]
[-client client_name] [-server server_name]
[-backupid backup_id] [-option option_name]
[-class class_name] [-ct class_type] [-rl retention_level]
[-sl sched_label] [-st sched_type] [-M master_server,...] [-v]
```

Notice the month/day/year and hours:minutes:seconds requirements for the `-d` and `-e` options. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

Example 1

The first example shows the last time each media ID available to a server had a backup image written today:

```
bpimagelist -media -U
Media ID      Last Written      Server
-----
IBM000       01/06/99 01:06   hat
AEK800       01/06/99 03:01   hat
C0015        01/06/99 02:01   hat
```

Example 2

This example shows the last time the media IDs available to the server had a backup image written during the specified time:

```
bpimagelist -media -d 01/05/99 18:00:46 -e 01/06/99 23:59:59 -U
```

```
Media ID      Last Written      Server
-----
IBM000       01/06/99 01:06   hat
AEK800       01/06/99 03:01   hat
```



```
C0015      01/06/99 02:01   hat
143191     01/05/99 23:00   hat
```

This example lists all images written today:

```
bpimagelist -U
```

Backed Up	Expires	Files	KB	C	Sched	Type	Class
01/27/99 01:08	02/03/99	1122	202624	N	Full	Backup	3590Grau
01/27/99 01:01	02/03/99	1122	202624	N	Full	Backup	IBM3590class
01/27/99 03:01	02/03/99	531	1055104	N	Full	Backup	DELLclass
01/27/99 02:01	02/03/99	961	31776	N	Full	Backup	QUALclass
01/27/99 01:08	02/03/99	2063	603328	N	Full	Backup	IBM3590class
01/27/99 01:01	02/03/99	2063	603328	N	Full	Backup	3590Grau

Example 3

This example lists media written information for 01/05/99:

```
bpimagelist -media -d 01/05/99 -e 01/05/99 -U
```

Media ID	Last Written	Server
IBM000	01/05/99 01:13	hat
143191	01/05/99 23:00	hat
AEK800	01/05/99 03:07	hat
C0015	01/05/99 02:06	hat

FILES

```
/usr/opensv/netbackup/logs/admin/log.mmddyy
```

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

SEE ALSO

```
bp(1), bparchive(1), bpbackup(1), bprestore(1)
```



bpimmedia(1M)

NAME

bpimmedia - display information about the NetBackup images on media

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpimmedia [-l | -L] [-class  
class_name] [-client client_name] [-d mm/dd/yyyy  
hh:mm:ss] [-e mm/dd/yyyy hh:mm:ss] [-mediaid vsn |  
path_name] [-mtype image_type] [-option option_name] [-r1  
retlevel] [-s1 sched_label] [-t sched_type] [-verbose] [-M  
master_server,...]
```

```
/usr/opensv/netbackup/bin/admincmd/bpimmedia -spangroups  
[-mediaid vsn] [-U] [-cn copy_number]
```

DESCRIPTION

bpimmedia queries the NetBackup image catalog and reports on the NetBackup images. bpimmedia produces two reports:

- ◆ An Images-on-Media report
- ◆ A Spangroups report

The first form of bpimmedia in the SYNOPSIS displays a set of NetBackup images in the Images-on-Media report. The Images-on-Media report lists the contents of media as recorded in the NetBackup image catalog. You can generate this report for any medium (including disk), filtering the report contents according to client, media ID or path, and so on. Refer to the section on NetBackup Reports in your NetBackup system administrator's guide for more information, including details about the fields in the Images on Media report. The Images on Media report does not show information for media used in backups of the NetBackup catalogs.

The second form of bpimmedia in the SYNOPSIS uses the `-spangroups` option to list media id groups that are *related* because images span from one volume to another. The output lists, for each media server in the cluster, the media ids that have spanning images. The `-spangroups` form of bpimmedia must be run on the NetBackup master server that administers the volumes. (See the Spanning Media topic in your NetBackup system administrator's guide.) Only removable media types are processed.

bpimmedia sends its error messages to stderr. bpimmedia sends a log of its execution to the NetBackup admin log file for the current day.

This command requires root privileges.

OPTIONS

- class *class_name*
Class name. By default, bpimmedia searches for images for all classes.
- client *client_name*
Client name. This name must be as it appears in the NetBackup catalog. By default, bpimmedia searches for all clients.
- cn Copy number (1 or 2) of a backup ID. The default is copy 1. This option is used only in combination with `-spangroups`.
- d *mm/dd/yyyy [hh:mm:ss]*
The start and end date. These specify the time range during which an image must have been created to be included in the report. Note: The locale setting for your system affects the way you specify dates and times. See the NOTES section later in this command description.
-d specifies a start date and time. The resulting list shows only images from backups or archives that occurred at or after the specified date and time. Use the following format:
mm/dd/yyyy [hh[:mm[:ss]]]
The valid range of dates is from 01/01/70 00:00:00 to 01/19/2038 03:14:07. The default is midnight beginning the current day.
- e *mm/dd/yyyy [hh:mm:ss]*
-e specifies an end date and time. The resulting list shows only images from backups or archives that occurred at or before the specified date and time. Use the same format as for the start date. The default is the current date and time.
- help Print a command-line usage message. When `-help` is present, it is the only option on the command line.
- L The list type is long. See the section DISPLAY FORMATS for more detail.
- l The list type is short. This is the default if the command line has no list-type option (for instance, if you enter bpimmedia and a carriage return). See the section DISPLAY FORMATS for more detail.
- M *master_server,...*
A list of alternative master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.



`-mediaid` *vsn* | *pathname*

This is either a VSN or an absolute pathname. If the media ID is a VSN, it is a one- to six-character string. If the media ID is a pathname, it is the absolute pathname of the filesystem for a disk storage unit.

When `-mediaid` is specified, the Images-on-Media report displays only images stored on this VSN or pathname. By default, the report displays images stored on all media IDs and pathnames.

For the Spangroups report (`-spangroups`), `-mediaid` can only be followed by a VSN. If `-mediaid` is omitted when `-spangroups` is present, `bpimmedia` displays all media in all spanning groups.

`-mtype`

Image type. The defined values, and their interpretations, are

0 = Regular backup (scheduled or user-directed backup)

1 = Pre-imported backup (phase 1 completed)

2 = Imported backup

`-option` *option_name*

Specifies a criterion for finding images to list. *option_name* is one of the following character strings, in either upper-or lower-case:

INCLUDE_PRE_IMPORT - Include images that have completed phase one of an import. Refer to the `bpimport(1M)` command description or the NetBackup system administrator's guide for more information.

ONLY_PRE_IMPORT - Include only images that have completed phase one of an import.

The default is INCLUDE_PRE_IMPORT.

`-rl` *retention_level*

The *retention_level*. The *retention_level* is an integer between 0 and 9. By default, `bpimmedia` searches for all retention levels.

`-sl` *sched_label*

The schedule label. By default, `bpimmedia` searches for images for all schedule labels.

`-spangroups`

Specifies that `bpimmedia` should create a Spangroups report. The default is to create an Images-on-Media report.

`-t` *sched_type*

Specifies a schedule type for the image selection. The default is any schedule type. Valid values, in either upper- or lower-case, are:

FULL (full backup)

INCR (differential-incremental backup)

CINC (cumulative-incremental backup)



- UBAK (user backup)
 UARC (user archive)
- U The list type is user. This option is used only in combination with `-spangroups`. See the section DISPLAY FORMATS for more detail.
- verbose Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).

DISPLAY FORMATS

IMAGES-ON-MEDIA REPORT

For the Images-on-Media report, there are two formats, short (`-l` or default) and long (`-L`).

◆ Long Display Format (`-L`)

If the command line contains `-L`, the display format is long. The `-L` display format contains a multi-line entry for each backup image. The number of lines for an entry is $n+1$, where n is the number of fragments for the image. The fields for an entry are listed in the table below. The first line of the entry contains the fields Backup_ID...Expires. Then, for each fragment in the image, there is a line containing the fields Copy_Media ID. The report has a two-line header. The first header line lists the field names for line 1 of each entry. The second header line lists the field names for the lines that contain fragment information.

See `bpduplicate(1m)` for more information on the terms *copy number* and *primary copy*.

Fields and meanings for the `-L` format are as follows:

Line 1

Backup-ID - Unique identifier for the backup that produced this image

Class - Class name (may be truncated if long)

Type - Schedule type (FULL, etc.)

RL - Retention level (0..9)

Files - Number of files in the backup

C - Compression (Y or N)

E - Encryption (Y or N)

T - Image type

R is Regular (scheduled or user-directed backup)



P is Pre-imported backup (phase 1 completed)

I is Imported backup

PC - Primary copy, 1 or 2. Designates which copy of the backup NetBackup chooses when restoring.

Expires - Expiration date of the first copy to expire, which is indicated by the Expires field of the fragment which is described below

Line 2_n+1

Copy - Copy number of this fragment

Frag - Fragment number, or IDX for a true-image-restore (TIR) fragment

KB - Size of the fragment, in kilobytes. This value does not include the size of tape headers between backups. A fragment size of 0 is possible for a multiplexed backup.

Type - Media type (Rmed for removable media; Disk otherwise)

Density - Density of the device that produced the backup (applies only to removable media)

Fnum - File number; this is the n-th backup on this medium (applies only to removable media)

Off - The byte offset on the medium where the backup begins (applies only to optical disk; ignore this value for tapes and magnetic disk)

Host - Server whose catalog contains this image

DWO - Device Written On; device where the backup was written. The DWO matches the drive index as configured in Media Manager (applies only to removable media).

MPX - Flag indicating whether this copy is multiplexed, Y or N (applies only when fragment number is 1)

Expires - Expiration date of this copy (applies only when fragment number is 1)

MediaID - Media ID or absolute path where the image is stored

Example of Long display format:

```
bpimmedia -L -class regr1_guava -t FULL
```

Backup-ID	Class	Type	RL	Files	C	E	T	PC	Expires		
Copy	Frag	KB	Type	Density	FNum	Off	Host	DWO	MPX	Expires	MediaID
guava_0949949902	regr1_guav	FULL	3	25			N	N	R	1	12:58 03/09/2000
1	1	256	RMed	dlt	13	0	plum	0	Y	12:58 03/09/2000	A00002



◆ Short Display Format (-1)

If the `bpconfig` command line contains `-1` or contains no `list-format` option, the display format is short. This produces a terse listing. This option can be useful for scripts or programs that rework the listing into a customized report format. The `-1` display format contains a multi-line entry for each backup image. The number of lines for an entry is $n+1$, where n is the number of fragments for the image. The layout of an entry is a first line, containing information about the image, followed by a line containing information about each fragment of the image. The attributes appear in the following order, separated by blanks.

Fields and Meanings for the `-1` format are as follows:

Line 1

IMAGE - Identifies the start of an image entry

Client - Client for the backup that produced this image

Version - Image-version level

Backup-ID - Unique identifier for the backup that produced this image

Class - Class name

Class type - 0 denotes Standard, etc. Run `bpimmedia -L` or refer to `bpbackup(1m)` to interpret the class-type value as a class-type name.

Schedule - Schedule name

Type - Schedule type (full, etc.)

RL - Retention level (0..9)

Files - Number of files

Expires - Expiration date of the first copy to expire, which is indicated by the Expires field of the fragment which is described below (system time); 0 denotes an image "in progress" or failed.

C - Compression; 1 (yes) or 0(no)

E - Encryption; 1 (yes) or 0(no)

Line 2_n+1

FRAG - Identifies a fragment line in an image entry

Copy - Copy number of this fragment

Frag - Fragment number, or -1 for a TIR fragment

KB - Size of the fragment, in kilobytes

(Internal) Internal value, not documented



Type - Media type (2 for removable media; 0 for disk)

Density - Density value (applies only to removable media) Run `bpimmedia -L` or `bpmedialist -mlist -L -ev mediaid` to interpret the density value as a density label.

Fnum - File number; this is the n-th backup on this medium (applies only to removable media)

MediaID - Media ID or absolute path where the image is stored

Host - Server whose catalog contains this image

Block size - Number of kilobytes per block for this medium

Off - Offset

Media dateTime this medium was allocated (system time)

DWO - Device Written On (applies only to removable media)

(Internal) - Internal value, not documented

(Internal) - Internal value, not documented

Expires - Expiration date of this copy in system time (applies only when fragment number is 1)

MPX - Flag indicating whether this copy is multiplexed, 1(yes) or 0(no) (applies only when fragment number is 1)

Example of the short display format:

```
bpimmedia -l -class regr1_guava -t FULL
IMAGE guava 3 guava_0949949902 regr1_guava 0 full 0 3 25 952628302 0 0
FRAG 1 1 10256 512 2 13 13 A00002 plum 65536 0 949616279 0 0 *NULL* 952628302 1
```

SPANGROUPS REPORT

For the Spangroups report, there are two formats: user (`-U` option) and short (the default). Both formats list, for each server, the server name, and the group data for that server. For each group of media that share spanned backup images, the media Ids are listed. When `-mediaid` appears on the command line, only the server and media group related to that media ID are displayed.

The user (`-U`) display format looks like this:

```
bpimmedia -spangroups -U
```

```
Related media groups containing spanned backup images, server plum:
```

```
Group:
```

```
A00002 A00003
```

```
Group:
```

```
400032
```

The short display format looks like this

```
bpimmedia -spangroups
```

```
SERVER plum
```

```
GROUP A00002 A00003
```

```
GROUP 400032
```

NOTES

The format that you must use for date and time values in NetBackup commands will vary according to your locale setting. The examples in this command description are for a locale setting of C.

If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the USAGE.

For example:

```
bpimmedia -help
```

```
USAGE: bpimmedia: [-l | -L] [-client <name>] [-t <sched_type>]
        [-class class_name] [-sl <sched_label>] [-d mm/dd/yyyy hh:mm:ss]
        [-e mm/dd/yyyy hh:mm:ss] [-verbose] [-mtype <media type>]
        [-M <master_server,...,master_server>]
        [-mediaid <vsnp/path>] [-option <name>]
USAGE: bpimmedia: -spangroups [-mediaid <vsnp>] [-U] [-cn
        <copy number>]
```

Notice the month/day/year and hours:minutes:seconds requirements for the `-d` and `-e` options. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

◆ Example 1

List the images for class `c_NDMP`. This request runs on a NetBackup media server. The report is based on the image catalog on the media server's master server, almond.

```
bpimmedia -L -class c_NDMP
```

```
Backup-ID          Class Type RL Files C E T PC Expires
```



bpimmedia(1M)

Copy Frag	KB	Type	Density	FNum	Off	Host	DWO	MPX	Expires	MediaID

t_0929653085		c_NDMP	FULL	3	5909	N N R 1	15:58		07/18/99	
1	844	RMed	dlt	2	0	almond		3		CB7514
1	1	9136	RMed	dlt	1	0	almond	3	N 15:58 07/18/99	CB7514

◆ Example 2

This example displays the tapes required to restore a particular file. If the `bpimmedia` command line provides the criteria to identify an individual backup, the output shows which media were used for the backup.

In this case, the command line provides the client, the date of the backup and the schedule type. The output shows that tape A00002 on the server plum contains the backup.

```
bpimmedia -L -client guava -d 2/7/00 -t UBAK
```

Backup-ID	Class	Type	RL	Files	C	E	T	PC	Expires	MediaID
Copy Frag	KB	Type	Density	FNum	Off	Host	DWO	MPX	Expires	MediaID

guava_0949949686		regrl_guav	UBAK	3	25	N N R 1	12:54		03/09/2000	
1	1	10256	RMed	dlt	11	0	plum	0	Y 12:54 03/09/2000	A00002

◆ Example 3

List, in long format, all the backups in the image catalog on the master server guava.

```
bpimmedia -L -M guava
```

Backup-ID	Class	Type	RL	Files	C	E	T	PC	Expires	MediaID
Copy Frag	KB	Type	Density	FNum	Off	Host	DWO	MPX	Expires	MediaID

guava_0949599942		test-class	FULL	1	15	N N R 1	11:45		02/17/2000	
1	1	224	Disk	-	-	guava	-	N	11:45 02/17/20	
/var/qatest/storage_unit//guava_0949599942_C1_F1										

◆ Example 4

List, in long format, the backups on media ID CB7514.



```
bpimmedia -L -mediaid CB7514
```

Backup-ID	Class	Type	RL	Files	C	E	T	PC	Expires				
Copy Frag	KB	Type	Density	FNum	Off	Host	DWO	MPX	Expires	MediaID			

toaster1_0929679294	tort_class	FULL	3	5898		N	N	R	1	23:14	07/18/99		
1	IDX	839	RMed	dlt	4	0				almond	6		CB7514
1	1	27154	RMed	dlt	3	0				almond	6	N	23:14 07/18/99 CB7514
toaster1_0929653085	NDMP_class	FULL	3	5909		N	N	R	1	15:58	07/18/99		
1	IDX	844	RMed	dlt	2	0				almond	3		CB7514
1	1	9136	RMed	dlt	1	0				almond	3	N	15:58 07/18/99 CB7514

◆ Example 5

This example uses the `-spangroups` option to list groups of media where images span between the media. `bpimmedia` executes on the NetBackup media server called `hat`. `hat`'s master server, `almond`, provides the results from its image catalog.

```
hat# bpimmedia -spangroups
SERVER almond
GROUP ISV043
GROUP ISV040
GROUP 00125B
GROUP 00126A
GROUP 00127A
SERVER hat
GROUP ISV044
GROUP 143191
```

EXIT STATUS

= 0 The command executed successfully.

<> 0 An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form:

```
bpimmedia: EXIT status = exit status
```



If an error occurred, a diagnostic precedes this message.

FILES

`/usr/opensv/netbackup/logs/admin/*`

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

SEE ALSO

`bpbackup(1m)`, `bpduplicate(1m)`, `bpimport(1m)`

NetBackup system administrator's guide

bpimport(1M)

NAME

bpimport - import backups that are expired or are from another Netbackup server.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpimport -create_db_info -id
  media_id [-server name] [-L output_file]
```

```
/usr/opensv/netbackup/bin/admincmd/bpimport [-l] [-p] [-pb]
  [-PD] [-PM] [-v] [-client name] [-Bidfile file_name] [-M
  master_server] [-st sched_type] [-sl sched_label] [-L
  output_file] [-class name] [-s startdate] [-e enddate] [-ct
  class_type] [-hoursago hours] [-cn copy_number] [-backupid
  backup_id] [-id media_id]
```

DESCRIPTION

The `bpimport` command allows backups to be imported. This command is useful for importing backups that have expired or are from another Netbackup server.

The import operation consists of two steps:

- ◆ Step 1 is performed with the first form of the command shown above (`-create_db_info` option) and recreates catalog entries for the backups that are on the specified media.
- ◆ Step 2 is performed with the second form of the command shown above and imports the backups from the media.

The expiration date for imported backups is the current date plus the retention period. For example, if a backup is imported on 11/14/98 and its retention level is one week, its new expiration date is 11/21/98.

You can import a backup only if all copies of it are expired. For more information on importing backups, see your NetBackup system administrator's guide.

OPTIONS

`-create_db_info`

This option recreates catalog entries for the backups that are on the specified media. It skips backups that are already in the catalog. This option only creates information about backups that are candidates for import, and does not perform the import operation. The `bpimport` command must be executed with this option prior to importing any backups.

The `-id` parameter is required with this option.



`-id media_id`

For step 1 (`-create_db_info`), this option specifies the media ID that has the backups you are going to import. This option is required with `-create_db_info`.

For step 2, this option designates a specific media ID from which to import backups. The default is all media IDs that were processed by step 1 of the import operation.

A backup ID that begins on a media ID that was not processed by step 1 is not imported. A backup that ends on a media ID that was not processed by step 1 will be incomplete.

`-server name`

Note For NetBackup BusinessServer there is only one server (the master) so specify the name of that server.

Specifies the name of the media server. The volume database for this server must have a record of the media ID that contains the backups to be imported. The default is the media server where the command is executed.

`-L output_file`

Specifies the name of a file in which to write progress information. The default is to not use a progress file.

`-Bidfile file_name`

file_name specifies a file that contains a list of backup IDs to be imported. List one backup ID per line in the file. If this option is included, other selection criteria is ignored.

`-l` Produces output in the progress log that lists each file imported.

`-p` Previews backups to be imported according to the option settings, but does not perform the import. Displays the media IDs, server name, and information about the backups to be imported.

`-pb` Previews the backups to import but does not perform the import. Similar to the `-p` option, but does not display the backups.

`-PM` Displays information on the backups to be imported according to the option settings, but does not perform the import. It displays the following information about the backup: date and time of the backup, class, schedule, backup ID, host, and media ID.

`-PD` Same as the `-PM` option, except the backups are sorted by date and time (newest to oldest).

`-v` Selects verbose mode. When specified, the debug and progress logs display more information.

`-client name`
 The host name of the client for which the backups were performed. The default is all clients.

`-M master_server`

Note For NetBackup BusinessServer, this option is not required because there is only one server, the master. If you do use this option in this case, specify the NetBackup BusinessServer master where you execute the command.

Specifies the master server that manages the media catalog that has the media ID. If this option is not specified, the default is one of the following:

If the command is executed on a master server, then that server is the default.

If the command is executed on a media server that is not the master, then the master for that media server is the default.

`-st sched_type`
 Search for backups to import which were created by the specified schedule type. The default is any schedule type.

Valid values are:

FULL (full backup)

INCR (differential-incremental backup)

CINC (cumulative-incremental backup)

UBAK (user backup)

UARC (user archive)

NOT_ARCHIVE (all backups except user archive)

`-sl sched_label`
 Search for backups to import which were created by the specified schedule. The default is all schedules.

`-class name`
 Search for backups to import in the specified class. The default is all classes.

`-s startdate`
 Specifies the start of the range of dates and times that include all backups to import. The syntax for *startdate* is:

mm/dd/yy [hh[:mm[:ss]]]

The `-e` option specifies the end of the range. The date and time format are dependent on the user's locale. For more information, see `locale(1)`.

Default is 24 hours prior to the current date and time.



`-e enddate`

Specifies the end of the range of dates and times that include all backups to import. The syntax for *enddate* is:

mm/dd/yy [hh[:mm[:ss]]]

The `-s` option specifies the start of the range. The date and time format are dependent on the user's locale. For more information, see `locale(1)`. Default is the current date and time.

`-ct class_type`

Search for backups created by the specified class type. The default is any class type.

Valid values are:

Informix-On-BAR

MS-Exchange-Server

MS-SQL-Server

MS-Windows-NT

NDMP

NetWare

Oracle

OS/2

Standard

Sybase

Note The following class types (AFS, Auspex-FastBackup, DataTools-SQL-BackTrack, DB2, FlashBackup, SAP, and Split-Mirror) apply only to NetBackup DataCenter.

AFS

Auspex-FastBackup

DataTools-SQL-BackTrack

DB2

FlashBackup

SAP

Split-Mirror

`-hoursago hours`

Specifies number of hours to search prior to the current time for backups. Do not use with the `-s` option.

`-cn copy_number`

Specifies the source copy number of the backups to import. You can import copy 1 or 2 or both. The default is both.



`-backupid backup_id`
Specifies the backup ID of a single backup to import.

EXAMPLES

Example 1

The following command (all on one line) creates catalog information for backups on media ID A00000. The media host hostname is cat. The progress file is `/tmp/bpimport.ls`.

```
bpimport -create_db_info -id A00000 -server cat -L  
/tmp/bpimport.ls
```

Example 2

The following command (all on one line) displays information about the backups that are candidates for import. The backups displayed would have been created between 11/01/98 and 11/10/98. The `bpimport` command with the `-create_db_info` option must be executed prior to this command.

```
bpimport -PM -s 11/01/98 -e 11/10/98
```

Example 3

The following command imports the backups specified in the `/tmp/import/images` file. The progress is entered in the `/tmp/bpimport.ls` file.

```
bpimport -Bidfile /tmp/import/image -L /tmp/bpimport.ls
```

FILES

`/usr/opensv/netbackup/logs/admin/*`

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



bplabel(1M)

NAME

bplabel - write a NetBackup label on tape media

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bplabel -ev evsn -d density
[-o] [-p volume_pool_name] [-n drive_name | -u
device_number]
```

DESCRIPTION

bplabel writes a NetBackup label on the specified media. Labeling is required only for media that were last used for NetBackup catalog backups or by a non-NetBackup application. You can also use it to assign specific media IDs. The NetBackup Device Manager daemon (ltid) must be active for bplabel to succeed. You must also manually assign the drive by using the NetBackup Device Monitor unless you include the -u option on the bplabel command.

Caution Ensure that the media does not contain required backups. After the media is relabeled, any backups that were on it cannot be restored.

The following are some facts about using this command:

- ◆ The -ev and -d options are required.
- ◆ The -p option is required if the *evsn* (media ID) is not in the NetBackup volume pool.
- ◆ If the data already on the media is in a recognized format and the -o option is not specified, bplabel prompts you to confirm the overwrite. Data format recognition works only if the first block on a variable length media is less than or equal to 32 kilobytes.
- ◆ Use the bplabel command only for tapes. For optical disk media, use the tpfomat command on a UNIX server.

OPTIONS

-ev *evsn* A required option that specifies the external volume serial number that is written to the tape label as a media ID . You can enter the evsn in either uppercase or lowercase. Internally, it is always converted to uppercase. The evsn must be six or fewer alphanumeric characters.

-d *density*

A required option that specifies the density of the tape drive on which the media is mounted. The tape mount request must be performed on a drive type that satisfies the -d option.



The valid densities are as follows:

4mm (4-mm cartridge)
 8mm (8-mm cartridge)
 dlt (dlt cartridge)
 dlt2 (dlt cartridge alternate)
 qscsi (1/4 in cartridge)

Note The following densities are supported only by NetBackup DataCenter servers.

dtf (dtf cartridge)
 hcart (1/2 Inch cartridge)
 hcart2 (1/2 Inch cartridge alternate)
 odiskwm (Optical disk-write many)
 odiskwo (Optical disk-write once)

- o Unconditionally overwrites the selected media ID. If this option is not specified, `bplabel` prompts for permission to overwrite media that meets any of the following conditions:
 - Contains a NetBackup media header.
 - Is NetBackup catalog backup media.
 - Is in TAR, CPIO, DBR, AOS/VS, or ANSI format.
- p *volume_pool_name*
 This option is required if the media ID is defined in the Media Manager volume database but is not in the NetBackup volume pool. *volume_pool_name* must specify the correct pool.
- u *device_number*
 Unconditionally assigns the standalone drive specified by *device_number*. The drive must contain media and be ready. By using this option, manual operator assignment is not required. The number for the drive can be obtained from the Media Manager configuration.
- n *drive_name*
 Unconditionally assigns the standalone drive specified by *drive_name*. The drive must contain media and be ready. By using this option, manual operator assignment is not required. The name for the drive can be obtained from the Media Manager configuration.

SEE ALSO

ltid(1M), vmadm(1M)



bplist(1)

NAME

bplist - lists backed up and archived files on the NetBackup server

SYNOPSIS

```
/usr/opensv/netbackup/bin/bplist [-A | -B] [-C client] [-S  
    master_server] [-t class_type] [-k class] [-F] [-R [n]] [-b  
    | -c | -u] [-l] [-r] [-Listseconds] [-T] [-unix_files]  
    [-nt_files] [-s mm/dd/yy [hh:mm:ss]] [-e mm/dd/yy  
    [hh:mm:ss]] [I] [PI] [-help] [-keyword "keyword_phrase"]  
    [filename]
```

DESCRIPTION

bplist shows a list of previously archived or backed up files according to the options that you specify. You can choose the file or directory and the time period that you want the listing to cover. Directories can be recursively displayed to a specified depth.

The list shows only the files that you have read access to. You also must own or have read access to all directories in the file paths. You can list files that were backed up or archived by another client only if you are validated to do so by the NetBackup administrator.

If you create directory `/usr/opensv/netbackup/logs/bplist/` with public-write access, bplist creates an activity log file in this directory that you can use for troubleshooting.

The output of bplist goes to standard output.

OPTIONS

- A | -B Specifies whether to produce the listing from archives (-A) or backups (-B). The default is -B.
- C *client* Specifies a client name to use for finding backups or archives to list. This name must be as it appears in the NetBackup configuration. The default is the current client name.
- S *master_server*
Specifies the name of the NetBackup server. The default is the first SERVER entry found in the `/usr/opensv/netbackup/bp.conf` file.
- t *class_type*
Specifies one of the following numbers corresponding to the class type (the default is 0 on all clients except Apollos, where it is 3):
 - 0 = Standard
 - 4 = Oracle

6 = Informix-On-BAR
 7 = Sybase
 10 = NetWare
 13 = MS-Windows-NT
 14 = OS/2
 15 = MS-SQL-Server
 16 = MS-Exchange-Server
 19 = NDMP

Note The following class types (Apollo-wbak, DataTools-SQL-BackTrack, Auspex-FastBackup, SAP, DB2, FlashBackup, Split-Mirror, and AFS) apply only to NetBackup DataCenter.

3 = Apollo-wbak
 11 = DataTools-SQL-BackTrack
 12 = Auspex-FastBackup
 17 = SAP
 18 = DB2
 20 = FlashBackup
 21 = Split-Mirror
 22 = AFS

- k *class* Names the class to search to produce the list. If not specified, all classes are searched.
- F Specifies that in the list output, symbolic links (applies only to UNIX clients) will end with a trailing @ and executable files with a trailing *.
- R [*n*] Recursively lists subdirectories encountered to a depth of *n*. The default for *n* is 999.
- b | -c | -u Specifies an alternate date and time to be used for printing with the -l option:
- b displays the backup date and time of each file.
 - c displays the last inode modification date and time for each file.
 - u displays the last access date and time of each file.
- The default is to display the time of last modification of each file.



- l Lists in long format, giving mode, owner, group, size in bytes, and time of last modification for each file (see the EXAMPLES section of this man page). The list shows the mode of each file as 10 characters that represent the standard UNIX file permissions. The first character is one of the following:
- d (specifies a directory)
 - l (specifies a link)
 - (specifies a file)
- The next nine characters show the three sets of permissions. The first set shows the owner's permissions, the next set shows the user-group permissions, and the last set shows permissions for all other users. Each set of three specifies the read, write, and execute permissions as follows:
- r means the file is readable
 - w means the file is writable
 - x means the file is executable
 - means the indicated permission is not granted
- Listseconds Specifies that seconds granularity be used for the time stamp when the -l option is used.
- r Lists raw partitions that were backed up. The default is to list file systems.
- T Lists directories in true-image backups. The default is to list non-true-image backups.
- unix_files Lists the files and directories in UNIX format. For example:
/C/users/test.
- nt_files Lists the files and directories in Windows NT/2000 format. For example:
C:\users\test.
- s *mm/dd/yy* [*hh:mm:ss*]
- e *mm/dd/yy* [*hh:mm:ss*]
Specifies the start and end date range for the listing.

Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section later in this command description.

-s specifies a start date and time for the listing. The resulting list shows only files in backups or archives that occurred at or after the specified date and time. Use the following format for the start date and time:



mm/dd/yy [*hh[:mm[:ss]]*]

The valid range of dates are from 01/01/70 00:00:00 to 01/19/2038 03:14:07. The default is the current date minus six months.

-e specifies an end date and time for the listing. The resulting list shows only files from backups or archives that occurred at or before the specified date and time. Use the same format as explained above for start date and time. The default is the current date and time.

- I Specifies a case-insensitive search. This means that capitalization is not considered when comparing names (for example, Cat matches cat).
- PI Specifies a path-independent search, which means that NetBackup searches for a specified file or directory without regard to the path. For example, if a file named `test` exists in the three directories shown below, a search for `test` finds all three instances of the file:

`/tmp/junk/test`

`/abc/123/xxx/test`

`/abc/123/xxx/yyy/zzz/test`

- help Displays a synopsis of command usage when it is the only option on the command line.
- keyword "*keyword_phrase*"

Specifies a keyword phrase for NetBackup to use when searching for backups or archives from which to restore files. The phrase must match the one that was previously associated with the backup or archive by the -k option of the `bpbbackup` or `bpararchive` command.

You can use this option in place of or in combination with the other restore options in order to make it easier to restore your backups and archives. The following meta characters can be used to simplify the task of matching keywords or parts of keywords in the phrase:

* matches any string of characters.

? matches any single character.

[] matches one of the sequence of characters specified within the brackets.

[-] matches one of the range of characters separated by the "-".

The keyword phrase can be up to 128 characters in length. All printable characters are permitted including space (" ") and period ("."). The phrase must be enclosed in double quotes ("...") or single quotes ('...') to avoid conflict with the UNIX shell.

The default keyword phrase is the null (empty) string.

- filename* Names the file or directory to list. If you do not specify a path, the default is the current working directory.



Any files that you specify must be listed at the end, following all other options.

NOTES

The format that you must use for date and time values in NetBackup commands depends on your locale setting. The examples in this command description are for a locale setting of C.

If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the usage. For example:

```
/usr/openv/netbackup/bin/bplist -help
USAGE: bplist [-A | -B] [-C client] [-S master_server]
        [-t class_type] [-k class] [-keyword "keyword phrase"]
        [-F] [-R [n]] [-b | -c | -u] [-l] [-r] [-T] [-I] [-PI]
        [-unix_files | -nt_files]
        [-s mm/dd/yy [hh:mm:ss]] [-e mm/dd/yy [hh:mm:ss]] [filename]
```

Notice the month/day/year and hours:minutes:seconds requirements for the `-s` and `-e` options. These are for a locale setting of C and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

EXAMPLE 1

To list recursively, in long format, the files that were backed up in `/home/user1`.

```
bplist -l -R /home/user1
lrwxrwxrwx  user1    eng      0    Apr  5 12:25 /home/user1/dirlink
drwxr-xr-x  user1    eng      0    Apr  4 07:48 /home/user1/testdir
drwxr-x---  user1    eng      0    Apr  4 07:49 /home/user1/dir
-rwxr----- user1    eng    1002  Apr  2 09:59 /home/user1/dir/file
lrwxrwxrwx  user1    eng      0    Apr  4 07:49 /home/user1/dir/link
```

EXAMPLE 2

To list, with details, the files that were backed up and associated with all or part of the keyword phrase

```
"My Home Directory"
```

in directory `/home/kwc`, enter the following:

```
bplist -keyword "*My Home Directory*" -l /home/kwc
```



EXAMPLE 3

To list, with details, the files that were archived and associated with all or part of the keyword phrase

"My Home Directory"

in directory /home/kwc enter the following:

```
bplist -A -keyword "*My Home Directory*" -l /home/kwc
```

EXAMPLE 4

To list, recursively and with details, the files that were backed up on drive D of Windows NT client slater and associated with all or part of the keyword phrase

"Win NT"

enter the following:

```
bplist -keyword "*Win NT*" -C slater -t 13 -R -l /D
```

FILES

/usr/opensv/netbackup/logs/bplist/log.*mmddy*

SEE ALSO

bp(1), bparchive(1), bpbackup(1), bprestore(1)



bpmedia(1M)

NAME

bpmedia - freeze, unfreeze, suspend, or unsuspend NetBackup media

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpmedia -freeze | -unfreeze |  
-suspend | -unsuspend -ev media_id [-h host] [-v]  
  
/usr/opensv/netbackup/bin/admincmd/bpmedia -movedb -ev media_id  
-newserver hostname [-oldserver hostname] [-v]
```

DESCRIPTION

bpmedia allows an individual NetBackup media ID to be controlled in terms of allowing or disallowing future backups or archives to be directed to the media. Note that this command applies only to media managed by Media Manager.

Note Under certain media or hardware error conditions, NetBackup automatically suspends or freezes media. If this happens, the reason is logged in the NetBackup Problems report. If necessary, you can use the bpmedia -unfreeze or -unsuspend options to reverse this action.

OPTIONS

- freeze Freezes the specified media ID. When an active NetBackup media ID is frozen, NetBackup stops directing backups and archives to the media. All unexpired images on the media continue to be available for restores. NetBackup never deletes a frozen media ID from the NetBackup media catalog, nor is it unassigned in the NetBackup volume pool when it expires.
- unfreeze Unfreeze the specified media ID. This reverses the action of freeze and allows the media to be used for backups or archives again if it has not expired. If a media is expired when it is unfrozen, it is immediately unassigned in the NetBackup volume pool.
- suspend Suspend the specified media ID. The action is the same as freeze except that when the media ID expires, it is immediately unassigned in the NetBackup volume pool.
- unsuspend Unsuspend the specified media ID. This reverses the action of suspend and allows the media to be used for backups or archives again.



```
-movedb -newserver hostname [-oldserver hostname]
```

Note You cannot use the `-movedb` option with NetBackup BusinessServer.

Moves a media catalog entry from one server to another in a master and media server cluster. This command moves the media catalog entry for the specified media ID from *oldserver* to *newserver* and updates the NetBackup image catalog to reflect that the media ID was moved. It is assumed that after the move, *newserver* has access to the media.

`-newserver hostname` specifies the name of the host to which the entry is moved.

`-oldserver hostname` specifies the name of the host where the catalog entry to be moved currently resides. If you do not specify `-oldserver`, the system where the command is being executed is considered to be the old server.

The `-movedb` option is most meaningful in configurations where a master and its media servers are sharing a robotic library and have access to all the media in the robot. If this is not the case, at a minimum, all NetBackup servers must use the same Media Manager volume database, so the media can be moved from one robotic library to another without losing their attributes and assignment status.

`-ev media_id`

Specifies the media ID that requires action. The media ID must be six or fewer characters and must be in the NetBackup media catalog.

`-h host`

Specifies the host name of the server where the media catalog resides. This option is required only if the volume was not written on the server where you execute the `bpmedia` command. In this case, the media ID is in the NetBackup media catalog on the other server and you must specify the name of that server on the `bpmedia` command.

For example, assume you have a master server named `whale` and a media server named `eel`. You execute the following `bpmedia` command on `whale` in order to suspend media ID `BU0001` that is in the media catalog on `eel`:

```
bpmedia -suspend -ev BU0001 -h eel
```

Use the NetBackup Media List report to determine the host that has the volume in its media catalog.

`-v`

Select verbose mode. This is only meaningful when running with activity logging turned on (that is, when the `/usr/opensv/netbackup/logs/admin` directory exists).



EXAMPLE (MOVEDB)

Note You cannot use the `-movedb` option with NetBackup BusinessServer.

Assume that the master server is HOSTM, with HOSTS1 and HOSTS2 being media servers. The following command (all on one line), executed on HOSTM, moves the media catalog entry for media ID DLT001 from HOSTS1 to HOSTS2 and updates the NetBackup image catalog:

```
bpmedia -movedb -ev DLT001 -newserver HOSTS2 -oldserver HOSTS1
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

```
/usr/opensv/netbackup/db/media/*
```

bpmedialist(1M)

NAME

bpmedialist - display NetBackup media status

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bpmedialist [-m list] [-U |
  -l | -L] [-ev media_id] [-rl ret_level] [-d density] [-p
  pool_name] [-h host_name | -M master_server,...] [-v]

/usr/opensv/netbackup/bin/admincmd/bpmedialist -summary [-U |
  -L] [-brief [-p pool_name] [-h host_name | -M
  master_server,...] [-v]

/usr/opensv/netbackup/bin/admincmd/bpmedialist -m contents -ev
  media_id [-U | -l | -L] [-d density] [-h host_name | -M
  master_server,...] [-v]

/usr/opensv/netbackup/bin/admincmd/bpmedialist -count -rt
  robot_type -rn robot_number [-d density] [-U | -l] [-h
  host_name | -M master_server] [-v]

```

DESCRIPTION

bpmedialist queries one or more NetBackup media catalogs and produces a report on the status of the NetBackup media. This command requires root privileges.

bpmedialist produces one of four reports:

MEDIA LIST REPORT

Media List (-m list) report, provides information on either a single volume or all volumes in the NetBackup media catalog. This report does not apply to disk storage units. The report lists, for each volume in the report, the volume's media Id, media server, and other attributes. This is the default report type.

If -U is an option, the status field appears as English text. Otherwise, the status appears as a hexadecimal integer. This is a three-digit value. The interpretation of the two upper-order digits is given here. Any or all of these flags can be set. Settings other than those listed here correspond to unreported states.

>= 0x200 Multiplexing is TRUE.

>= 0x080 Imported is TRUE.

>= 0x040 Multiple retention levels is TRUE.

The interpretation for the low-order status digit is determined by comparing the digit to the following values in order.



- >= 0x008 The status is Full.
- >= 0x004 This is an unreported state.
- >= 0x002 The status is Suspended.
- == 0x001 The status is Frozen.
- == 0x000 The status is Active.

The reported status is the status for the low-order digit combined with the status for the upper-order digits. For instance, for a status value of 0x040, the media ID is active, and multiple retention levels are in effect.

The `-1` option produces a report in Short mode. Each media ID occupies one line of the report. The fields on this line are listed below. The section on the Media List Report in your NetBackup system administrator's guide describes the fields in detail. Any fields listed below that are not documented in that section are reserved for NetBackup internal use.

- ◆ media id
- ◆ partner id
- ◆ version
- ◆ density
- ◆ time allocated
- ◆ time last written
- ◆ time of expiration
- ◆ time last read
- ◆ Kbytes
- ◆ nimages
- ◆ vimages (unexpired images)
- ◆ retention level
- ◆ volume pool
- ◆ number of restores
- ◆ status (described above)
- ◆ hsize
- ◆ ssize
- ◆ l_offset
- ◆ reserved

- ◆ psize
- ◆ reserved
- ◆ 4 reserved fields

MEDIA SUMMARY (-SUMMARY) REPORT

The Media Summary report lists, by server, summary statistics for active and inactive media grouped according to expiration date. The report shows the expiration date for the media and the number of media at each retention level, and the status of each media ID.

MEDIA CONTENTS REPORT

The Media Contents report lists the contents of media as read directly from the media. It lists the backup IDs that are on a single media ID. It does not list each individual file. This report does not apply to disk storage units. Note that if you attempt to abort the command by entering `ctl-c` and the media requested are still being mounted or positioned, the storage unit may stay in use for some time after the break. Each entry in the report appears as that area of the storage unit is read.

The `-l` format for the Media Contents report produces one line for each backup ID, containing the fields below. The section on the Media Contents Report in your NetBackup system administrator's guide contains more details. Any fields not described in that section are reserved for NetBackup internal use.

- ◆ version (1 denotes a DB backup image, 2 denotes a regular backup image)
- ◆ backup id
- ◆ creation time
- ◆ expiration time
- ◆ retention level
- ◆ fragment number
- ◆ file number
- ◆ block size (in bytes)
- ◆ status
- ◆ media_id
- ◆ size
- ◆ reserved
- ◆ data_start
- ◆ reserved
- ◆ client_type *



- ◆ copy_num *
- ◆ sched_type *
- ◆ flags *
- ◆ opt_extra
- ◆ mpx_headers
- ◆ res1
- ◆ class name *
- ◆ schedule label *

* These fields are significant only if version is 2.

MEDIA COUNT (-COUNT) REPORT

The Media Count report shows a count of the number of UP devices matching all the criteria specified. The robot type and the robot number are mandatory criteria for this report. The `-U` format provides a title, Number of UP devices for $rt(rn) = value$. The `-l` format provides only the value.

OPTIONS

Report-type Options

`bpmédialist` produces one of four types of reports. An option on the command line determines the type of report produced. The report-type options are as follows:

- `-mlist`
Produce a Media List report. This is the default report type.
- `-summary`
Produce a Media Summary report.
- `-mcontents`
Produce a Media Contents report.
- `-count`
Produce a Media Count report. This report also displays the media attribute `ALLOW_MULT_RET_PER_MEDIA` and its value, 0 (do not allow) or 1 (allow).

Report-format Options

The `bpmédialist` report can appear in one of several formats. The report-format options are as follows:



- brief Produce a brief report. This option is available for the Media Summary report only. The default is a full report, which includes a breakdown of active and non-active media, reporting on each media ID's status within these categories.
- U Report in user mode. This is the default report mode. The report includes a banner listing the column titles, and the report style is descriptive, rather than terse.
- L Report in long mode. This format produces the report with the most complete information. For instance, for the Media List report, the report lists the attributes of each media ID as a series of *keyword = value* pairs, one attribute per line. A value may be expressed as both a numeric value and a descriptive value.
- l Report in short mode. This format produces a terse report. This option is useful for scripts or programs that rework the listing contents into a customized report format.

Other Options

The following are the remaining options used by `bpmédialist`:

- density *density_type*
Report on media of this density type. If the robot type is specified on the command line, the value for density should be consistent with the robot type. Available density types are:
 - dlt - DLT Cartridge
 - 8mm - 8mm Cartridge
 - 4mm - 4mm Cartridge
 - qscsi - 1/4 Inch Cartridge
 - dlt2 - DLT Cartridge alternate
-
- Note** The following densities are supported only on NetBackup DataCenter servers.
- hcart - 1/2 Inch Cartridge
 - dtf - DTF Cartridge
 - odiskwm - Optical Disk Write-Many
 - odiskwo - Optical Disk Write-Once
 - hcart2 - 1/2 Inch Cartridge alternate
- ev *media_id*
Report on this media ID only. This is a required option for the Media Contents report.



For the Media List report, this option is optional, and, by default, all media IDs are included in that report. The media ID can be provided in either upper- or lower-case. The media ID must be six or fewer characters and must be in the NetBackup media catalog (that is, assigned from the NetBackup volume pool).

-h host_name

Note For NetBackup BusinessServer, there is only one server (the master) so use the name of that server for *host_name*.

host_name is either the name of a host, or the character string ALL. If *host_name* is the name of a host, the query goes to the media catalog residing on the system *host_name*. For the *-mcontents* and *-count* options, this option can appear once. For the *-mlist* and *-summary* options, this option can appear more than once.

The system *host_name* must allow access by the system running *bpmedialist*. *host_name* can be a media server for a master server other than the local master server. The default is the master server of the local cluster.

For a media server for a master server other than the local master, if a *bpmedialist* query is made using *-h the_media_server*, and an equivalent *bpmedialist* query uses *-M the_media_servers_master*, the *bpmedialist* using *-h* may complete faster. This difference in response time can be significant if the master server addressed by *-M* is located remotely, and the media server addressed by *-h* is local.

If *host_name* is ALL, the query goes to the local master server and its media servers.

-help Prints a synopsis of command usage when it is the only option on the command line.

-M master_server, . . .

A list of alternative master servers. This is a comma-delimited list of host names. If this option is present, each master server in the list executes the *bpmedialist* command. If an error occurs for any master server, processing stops at that point.

The report is the composite of the information returned by all the master servers in this list. *bpmedialist* queries each of these master servers. Each master server in the list must allow access by the system issuing the *bpmedialist* command.

For *-mcontents* (Media Contents report) only, the master server returns media information from the media catalogs. This media information is for both the master and its media servers (except for NetBackup BusinessServer which does not support remote media servers). For

example, if a media ID exists on a media server of one of the master servers in the `-M` list, the master retrieves the media information from the media server and returns it to the system running `bpmedialist`. In this case, both the master server and the media server must allow access by the system issuing the `bpmedialist` command.

The default is the master server for the server running `bpmedialist` (Note that NetBackup BusinessServer supports only one server, the master; so the default, in this case, is always the NetBackup BusinessServer master where you run `bpmedialist`).

`-p` *pool_name*

Report on the media IDs that belong to this volume pool. The default is all pools.

`-rl` *retention_level*

Report on media that are using this retention level. The retention level determines how long to retain backups and archives. The *retention_level* is an integer between 0 and 9. The default retention level is 1.

Here are the retention levels with the installation values for the corresponding retention periods. Note that your site may have reconfigured the retention periods corresponding to the retention levels.

0	1 week
1	2 weeks
2	3 weeks
3	1 month
4	2 months
5	3 months
6	6 months
7	9 months
8	1 year
9	Infinity

`-rn`

Report on the robot using this robot number. This is a required option when the `-count` option is used. The robot number can be obtained from the Media Manager device configuration. For rules concerning the use of this number, see your Media Manager system administrator's guide.

`-rt` *robot_type*

Report on a robot of this type. This is a required option when the `-count` option is used. For non-robotic (standalone) devices select NONE. Valid robot types include the following

TL4 - Tape Library 4MM



TL8 - Tape Library 8MM
 TLD - Tape Library DLT
 TS8 - Tape Stacker 8MM
 TSD - Tape Stacker DLT
 NONE - Not robotic

Note The following robot types apply only to NetBackup DataCenter.

ACS - Automated Cartridge System
 LMF - Library Management Facility
 ODL - Optical Disk Library
 TC4 - Tape Carousel 4MM
 TC8 - Tape Carousel 8MM
 TLH - Tape Library Half-Inch
 TLM - Tape Library Multimedia
 TSH - Tape Stacker Half-Inch

-v Select verbose mode. This option causes bpmedialist to log additional information for debugging purposes. The information goes into the NetBackup administration daily activity log. This option is meaningful only when NetBackup has activity logging enabled (/usr/opensv/netbackup/logs/admin directory defined).

EXAMPLES

Example 1

This example produces a media report for all media IDs defined for the master server of the local system and any media servers. (Note that for NetBackup BusinessServer the report includes only media IDs for the master server because remote media servers are not supported.)

```
hat 36# ./bpmedialist
Server Host = hat
```

id	rl	images	allocated	last updated	density	kbytes	restores
		vimages	expiration	last read		<----- STATUS ----->	
143191	0	28	12/03/98 23:02	12/22/98 23:00	dlt	736288	1
		7	12/29/98 23:00	12/09/98 10:59			



144280	0	9	11/25/98 11:06	12/01/98 23:03	dlt	290304	0
		0	12/08/98 23:03	N/A	EXPIRED	FROZEN	
AEK800	0	22	12/06/98 03:05	12/23/98 03:01	dlt	23213184	0
		7	12/30/98 03:01	12/09/98 10:48			
C0015	0	28	11/26/98 02:09	12/23/98 02:01	dlt	896448	0
		7	12/30/98 02:01	N/A			
IBM001	0	16	12/16/98 01:01	12/23/98 01:07	dlt	6447360	0
		14	12/30/98 01:07	N/A			
L00103	0	20	12/07/98 08:33	12/23/98 01:07	dlt	7657728	0
		9	12/30/98 01:07	N/A			
L00104	0	9	12/11/98 01:09	12/21/98 01:04	dlt	5429504	0
		5	12/28/98 01:04	N/A			

Example 2

This example produces a media count report for robot type TLD and robot number 0:

```
hat 40# ./bpmedialist -count -rt TLD -rn 0
ALLOW_MULT_RET_PER_MEDIA 0
Number of UP devices for TLD(0) = 2
```

Example 3

This example produces a media contents report for media ID AEK802. The report is partially listed below.

```
% ./bpmedialist -mcontents -ev AEK802
media id = AEK802, allocated 01/08/99 03:10, retention level = 0
```

File number 1

```
Backup id = hat_0915786605
Creation date = 01/08/99 03:10
Expiration date = 01/15/99 03:10
```



```
Retention level = 0
Copy number = 1
Fragment number = 2
Block size (in bytes) = 65536
```

File number 2

```
Backup id = hat_0915809009
Creation date = 01/08/99 09:23
Expiration date = 01/15/99 09:23
Retention level = 0
Copy number = 1
Fragment number = 1
Block size (in bytes) = 65536
```

Example 4

In this example, bpmedialist runs on the master server buffalo. bpmedialist produces a Media List report for master servers hat and duo.

```
buffalo# ./bpmedialist -M hat,duo
```

Server Host = hat

id	rl	images	allocated		last updated		density	kbytes	restores
		vimages	expiration	last read	<----- STATUS ----->				
143191	0	51	12/03/98	23:02	01/11/99	23:04	dlt	1436686	2
		9	01/18/99	23:04	01/08/99	10:26			
144280	0	9	11/25/98	11:06	12/01/98	23:03	dlt	290304	0
		0	12/08/98	23:03	01/12/99	16:10	EXPIRED FROZEN		
AEK800	0	38	12/06/98	03:05	01/08/99	03:10	dlt	39229824	0
		3	01/15/99	03:10	12/09/98	10:48	FULL		
AEK802	0	6	01/08/99	03:10	01/12/99	03:05	dlt	6140544	0
		6	01/19/99	03:05	01/12/99	16:12			




```
C0015  0    48  11/26/98 02:09  01/12/99 02:11    dlt  1531968    0
        7    01/19/99 02:11    N/A

IBM000  0    19  01/01/99 01:09  01/12/99 02:05    dlt  8284224    0
        13  01/19/99 02:05  01/09/99 05:41
```

Server Host = duo

id	rl	images	allocated	last updated	density	kbytes	restores
		vimages	expiration	last read	<-----	STATUS	----->
A00004	0	0	11/16/97 05:31	N/A	4mm	0	0
		0	N/A	N/A	FROZEN		
DLT210	1	5	12/09/98 06:10	01/08/99 06:04	dlt	2560	0
		2	01/22/99 06:04	N/A			
DLT215	0	124	12/08/98 14:57	01/12/99 08:07	dlt	9788072	4
		28	01/19/99 08:07	12/31/98 15:42			

Example 5

In this example, `bpmedialist` reports which of two hosts has a given media ID configured. Since the host hat does not have A00004 configured in its media catalog, it reports, the requested media ID was not found in the NetBackup media catalog or Media Manager volume database

The host duo does have A00004 configured, so it produces a Media List report for A00004 (the command is all on one line).

```
buffalo# ./bpmedialist -mlist -h hat -h duo -ev A00004
```

```
requested media id was not found in NB media database and/or
MM volume database
```



Server Host = duo

id	rl	images	allocated	last updated	density	kbytes	restores
		vimages	expiration	last read	<-----	STATUS	----->
A00004	0	0	11/16/97 05:31	N/A	4mm	0	0
		0	N/A	N/A	FROZEN		

FILES

/usr/opensv/netbackup/logs/admin/*

/usr/opensv/netbackup/db/media/mediaDB



bprd(1M)

NAME

bprd - initiates the NetBackup request daemon

SYNOPSIS

```
/usr/opensv/netbackup/bin/bprd [-verbose]
```

DESCRIPTION

bprd is responsible for starting automatic client backups and responding to client requests for file restores and user backups and archives. bprd executes only on the master server and can be started only by the administrator.

The following steps occur when bprd starts:

1. After disassociating itself from the terminal, the daemon
 - ◆ Logs a message indicating that it has started.
 - ◆ Starts bpdbm (NetBackup Database Manager).
 - ◆ Verifies that no other instance of bprd is executing. If another instance of bprd is found, the program terminates.
2. The program reads the NetBackup configuration attributes and recycles older error and activity log files. Activity and error logs are also recycled on a daily basis.
3. bprd determines its port number by checking the `services` file for an entry with a service name of `bprd` and a protocol name of `tcp`. For example:


```
bprd 13720/tcp
```
4. After binding to its port, the program starts scheduling automatic client backups, accepting requests from client machines for file restores or user backups or archives, and accepting administrative requests from the server.

You can use `bprdreq -terminate` to terminate bprd. Terminating bprd does not terminate bpdbm.

OPTIONS

`-verbose` Specifies that bprd will write additional information in its daily-activity log for debugging purposes.

FILES

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

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



`/usr/opensv/netbackup/logs/bprd/*`

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

`/usr/opensv/netbackup/bin/initbpdbm`

SEE ALSO

`bpadm(1M)`, `bpdbm(1M)`

bprecover(1M)

NAME

bprecover - recover selected NetBackup related catalogs

SYNOPSIS

```

/usr/opensv/netbackup/bin/admincmd/bprecover [-v]
-l -ev media_id -d density [-v]
-l -dpath disk_path [-v]
-l -tpath tape_device_path [-v]
-l -opath optical_device_path [-v]
-r [all | ALL | image_number] -ev media_id -d density [-stdout]
  [-dhost destination_host] [-v]
-r [all | ALL | image_number] -dpath disk_device_path [-stdout]
  [-dhost destination_host] [-v]
-r [all | ALL | image_number] -tpath raw_tape_device_path [-stdout]
  [-dhost destination_host] [-v]
-r [all | ALL | image_number] -opath optical_device_path [-stdout]
  [-dhost destination_host] [-v]

```

DESCRIPTION

bprecover initiates the NetBackup utility for restoring the NetBackup internal databases called catalogs and recovers catalogs that were backed up by using the procedures described in the NetBackup system administrator's guide. Use bprecover only if catalogs were destroyed on disk.

The command has two main modes: list and recover. List shows the contents of a backup media or disk path. Recover recovers the catalog files.

Only root can execute this command.

OPTIONS

-l Lists the header information from the specified media or disk path.

-ev *media_id* -d *density*
 Specifies the media ID and the density of the media from which to recover files.
media_id must be six or less characters and must be defined in the Media Manager volume database.



density must be one of the following:

4mm (4-mm cartridge)
8mm (8-mm cartridge)
dlt (dlt cartridge)
dlt2 (dlt cartridge alternate)
qscsi (1/4-in cartridge)

Note The following densities apply only to NetBackup DataCenter servers.

hcart (1/2 Inch cartridge)
hcart2 (1/2 Inch cartridge alternate)
dtf (DTF cartridge)
odiskwm (Optical disk-write many)
odiskwo (Optical disk-write once)

-dpath *disk_path*

-tpath *tape_path*

-opath *optical_path*

Specifies a raw device path. If *-ev* and *-d* are not specified, you must use *-dpath*, *-opath*, or *-tpath* to specify a raw device path. The Media Manager device and volume daemons (*ltid* and *vmd*) do need not to be active when you use one of these options.

Note Some platforms require a Berkeley-style close device for the *tpath* option. This is the path with *b* in the device name (for example on a Solaris system, it could be */dev/rmt/0c**b**n*). You will get an I/O error if you do not specify a Berkeley style close device on platforms that require it.

-r [*all* | ALL | *image_number*]

Recovers images from the specified media or disk path. There are three modes of recovery available with *-r*:

If *-r all* (or ALL) is specified, recover all the images contained in the specified media or disk path.

If *-r image_number* is specified, recover only the selected image number from the specified media or disk path.

If *-r* is specified by itself, *bprecover* interactively prompts and asks if you want to recover the images contained in the specified media or disk path.

`-stdout` Specifies that the selected backup image be written to stdout instead of automatically being restored. This option is useful, for example, if only one individual file was lost and you want to restore it without restoring the rest of the catalog files in the image.

Note You cannot specify `-r ALL` with `-stdout` because the `-stdout` option permits only one file image to be read at a time.

`-dhost destination_host`
Specifies the host to which the selected catalog is restored. Normally, catalogs are restored to the host where the data originated (as displayed with the `-l` option). The `-d` option makes it possible to restore the catalog to another host.

Caution Use the `dhost` option with EXTREME caution, since it can overwrite existing catalogs on the destination host. To permit recovery in case you unintentionally overwrite the wrong catalogs, you can move existing catalogs to a temporary directory on the destination host.

The following NetBackup client software must be installed on the destination host:

```
/usr/opensv/netbackup/bin/bpcd
and
/usr/opensv/netbackup/bin/tar
```

Note Do not specify `-r all` (or `ALL`) with `-dhost`. Either explicitly specify an image (for example, `-r 2`) or use the interactive mode (`-r`).

`-v` Selects verbose mode. This is meaningful only when running with activity logging turned on (that is, when the `/usr/opensv/netbackup/logs/admin` directory exists).

EXAMPLES

Example 1

List the backup header information for catalog backup that was done to disk path `/disk/bpbackup`.

```
#bprecover -l -dpath /disk1/bpbackup
Database Backup Information from /disk1/bpbackup
```

```
Created:      03/30/93 11:31:34
Server:      bphost
```



```

                Path
                ----
IMAGE1         /usr/opensv/netbackup/db
IMAGE2         /usr/opensv/volmgr/database

```

Example 2

List the backup header information from media ID JBL29, which is density 8mm.

```

#bprecover -l -ev JBL29 -d 8mm
Database Backup Information from JBL29

Created:       04/02/93 05:50:51
Server:       bphost
Block size:   32768

```

```

                Path
                ----
IMAGE1         /usr/opensv/netbackup/db
IMAGE2         /usr/opensv/volmgr/database

```

Example 3

Recover the `/usr/opensv/netbackup/db` files from disk path `/disk1/bpbackup`.

```

#bprecover -r 1 -dpath /disk1/bpbackup
Recovering bphost:/usr/opensv/netbackup/db

```

Example 4

Recover all the backed up catalogs from media ID JBL29.

```

#bprecover -r ALL -ev JBL29 -d 8mm
Recovering bphost:/usr/opensv/netbackup/db
Recovering bphost:/usr/opensv/volmgr/database

```

Example 5

Interactively restore selected images. Use raw tape path `/dev/rmt/1cbn`. Assume the media that is loaded into the drive is the same one as in example 4.

```

#bprecover -r -tpath /dev/rmt/1cbn
Recover bphost:/usr/opensv/netbackup/db y/n (n)? n

```




```
Recover bphost:/usr/opensv/volmgr/database y/n (n)? y
Recovering bphost:/usr/opensv/volmgr/database
```

Example 6

Recover a single file from image 1 on JBL29.

```
#bprecover -r 1 -ev JBL29 -d 8mm -stdout | /bin/tar -xvf
- /usr/opensv/netbackup/file_to_recover
Writing bphost:/usr/opensv/netbackup/db to stdout
```

Example 7

Restore an image to another host by using the `-dhost destination_host` option.

```
#bprecover -r -ev ODL08B -d odiskwm -dhost giskard
Recover bphost:/usr/opensv/netbackup/db to host giskard y/n (n)? n
Recover bphost:/usr/opensv/volmgr/database to host giskard y/n (n)? y
Recovering bphost:/usr/opensv/volmgr/database to host giskard
```

ERRORS

If any errors occur during the recover operation, error messages are written to stderr.

FILES

```
/usr/opensv/netbackup/logs/admin/*
/usr/opensv/netbackup/db/*
/usr/opensv/volmgr/database/*
```

SEE ALSO

`tpreq(1)` (Media Manager command)

NetBackup system administrator's guide (UNIX version) for information on backing up the NetBackup catalogs.

NetBackup troubleshooting guide (UNIX version) for information on disaster recovery.



bprestore(1)

NAME

bprestore - restores files from the NetBackup server

SYNOPSIS

```
/usr/opensv/netbackup/bin/bprestore [-A | -B] [-K] [-l | -H |  
-y] [-r] [-T] [-L progress_log] [-R rename_file] [-C client]  
[-D client] [-S master_server] [-t class_type] [-c class] [-s  
mm/dd/yy [hh:mm:ss]] [-e mm/dd/yy [hh:mm:ss]] [-w  
[hh:mm:ss]] [-k "keyword_phrase"] -f listfile | filenames
```

DESCRIPTION

bprestore lets users restore a backed up or archived file or list of files. You can also name directories to restore. If you include a directory name, bprestore restores all files and subdirectories of that directory.

By default, you are returned to the system prompt after bprestore is successfully submitted. The command works in the background and does not return completion status directly to you. The -w option lets you change this behavior so bprestore works in the foreground and returns completion status after a specified time period.

The bprestore command restores the file from the most recent backups within the time period you specify, except for a true-image restore (see the -T option description)."

. bprestore overwrites any file of the same name that already exists on the local client disk, unless you include the -K option. It is also possible to restore files that were backed up or archived on another client (-C option). You must be validated by the NetBackup administrator to restore from other clients.

Use the bplist command to display information on the files and directories that were backed up or archived.

bprestore writes informative and error messages to a progress-log file if you create the file prior to the execution of the bprestore command and then specify the file with the -L *progress_log* option. If bprestore cannot restore the requested files or directories, you can use the progress log to find the reason for the failure.

For detailed troubleshooting information, create a directory named /usr/opensv/netbackup/logs/bprestore with public-write access. bprestore then creates an activity log file in this directory.

In addition, if a nonroot user specifies USEMAIL = *mail_address* in their \$HOME/bp.conf file, NetBackup sends mail on the restore completion status to *mail_address*. This message is sent when the restore process is complete.

The following restrictions apply to bprestore:



- ◆ You can restore files and directories that you own and those owned by other users if you have read access. You need write access to another user's directories and files to restore that user's files to their original location.
- ◆ The operating system restricts the number of files and directories that you can specify on a single `bprestore` command line. If this is a problem, use the `-f` option to restore the files.

OPTIONS

- `-A` | `-B` Specifies whether to restore from archives (`-A`) or backups (`-B`). The default is `-B`.
- `-K` Specifying this option causes `bprestore` to keep existing files rather than writing over them when restoring files with the same name. The default is to overwrite existing files.

Note The `-l` | `-H` | `-y` options apply only when restoring UNIX files to a UNIX system.

- `-l` | `-H` | `-y`
 - Specifying `-l` renames the targets of UNIX links by using the `-R rename_file` option in the same way as when renaming files.
 - Specifying `-H` renames UNIX hard links by using the `-R rename_file` option in the same way as when renaming files. Soft links are unchanged.
 - Specifying `-y` renames UNIX soft links by using the `-R rename_file` option in the same way as when renaming files. Hard links are unchanged.
 - See Example 5 in the EXAMPLES section.
- `-r` Specifying this option restores raw partitions instead of file systems.
- `-L progress_log`
 - Specifies the name of an existing file in which to write progress information.
 - For example: `/home/tlc/proglog`
 - The default is to not use a progress log.
- `-R rename_file`
 - Specifies the name of a file with name changes for alternate-path restores.
 - Use the following form for entries in the rename file:
`change backup_filepath to restore_filepath`
 - Where:
 - The file paths must start with `/` (slash)
 - The first *backup_filepath* that is matched is replaced with the *restore_filepath* string. The default is to restore using the original path.



For example, the following entry renames `/usr/fred` to `/usr/fred2`:
change `/usr/fred` to `/usr/fred2`

- C *client* Specifies a client name to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog. The default is the current client name.
- D *client* Specifies a destination client. This can be done by a root user on the master server in order to direct the restored files to a machine other than the client specified with the -C option.
- S *master_server*
Specifies the name of the NetBackup server. The default is the first server found in the `/usr/opensv/netbackup/bp.conf` file.
- t *class_type*
Specifies one of the following numbers corresponding to the class type (the default is 0 on all clients except Apollos, where it is 3):
 - 0 = Standard
 - 4 = Oracle
 - 6 = Informix-On-BAR
 - 7 = Sybase
 - 10 = NetWare
 - 13 = MS-Windows-NT
 - 14 = OS/2
 - 15 = MS-SQL-Server
 - 16 = MS-Exchange-Server
 - 19 = NDMP

Note The following class types (Apollo-wbak, DataTools-SQL-BackTrack, Auspex-FastBackup, SAP, DB2, FlashBackup, Split-Mirror, and AFS) apply only to NetBackup DataCenter.

- 3 = Apollo-wbak
- 11 = DataTools-SQL-BackTrack
- 12 = Auspex-FastBackup
- 17 = SAP
- 18 = DB2
- 20 = FlashBackup
- 21 = Split-Mirror
- 22 = AFS

- C *class* Specifies the class for which the backups or archives were performed.



`-s mm/dd/yy [hh:mm:ss]`

`-e mm/dd/yy [hh:mm:ss]`

Specifies the start and end date range for the listing. The `bprestore` command restores only files from backups or archives that occurred within the specified start and end date range.

Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section of this command description.

`-s` specifies a start date and time for the restore window. `bprestore` restores files only from backups or archives that occurred at or after the specified date and time. Use the following format:

`mm/dd/yy [hh[:mm[:ss]]]`

The valid range of dates are from 01/01/70 00:00:00 to 01/19/2038 03:14:07. The default start date is 01/01/70 00:00:00.

`-e` specifies an end date and time for the restore window. `bprestore` restores only files in backups or archives that occurred at or before the specified date and time. Use the same format as for the start date and time.

The end backup date and time do not need to be exact, except for a true-image restore (see the `-T` option description). The `bprestore` command restores the file that has the specified backup date and time or the file that is the most recent backup preceding the end date and time. The default is the current date and time."

If you do not specify either `-s` or `-e`, `bprestore` restores the most recently backed up version of the file.

`-T` Specifies a true-image restore, where only files and directories that existed in the last true-image backup are restored. This option is useful only if true-image backups were performed. If this option is not specified, all files and directories meeting the specified criteria are restored, even if they were deleted.

When the `-T` option is specified, the image requested must be uniquely identified. Unique identification is accomplished by using the `-e` option with seconds granularity. The `-s` option, if any, is ignored. The seconds granularity of an image can be retrieved by using the `bplist` command with the `-l` and `-Listseconds` options.

`-w [hh:mm:ss]`

Causes NetBackup to wait for a completion status from the server before returning you to the system prompt.



Note The locale setting for the system affects the way you must specify dates and times. See the NOTES section of this command description.

You can optionally specify a wait time in hours, minutes, and seconds. The maximum wait time you can specify is 23:59:59. If the wait time expires before the restore is complete, the command exits with a timeout status. The restore, however, still completes on the server.

Specifying 0 or not specifying a time, means wait indefinitely for the completion status.

-k *"keyword_phrase"*

Specifies a keyword phrase for NetBackup to use when searching for backups or archives from which to restore files. The phrase must match the one that was previously associated with backup or archive by the **-k** option of the `bpbbackup` or `bparcarchive` command.

You can use this option in place of or in combination with the other restore options in order to make it easier to restore your backups and archives. The following meta characters can simplify the task of matching keywords or parts of keywords in the phrase:

* matches any string of characters.

? matches any single character.

[] matches one of the sequence of characters specified within the brackets.

[-] matches one of the range of characters separated by the "-".

The keyword phrase can be up to 128 characters in length. All printable characters are permitted including space (" ") and period ("."). The phrase must be enclosed in double quotes ("...") or single quotes ('...') to avoid conflict with the UNIX shell.

The default keyword phrase is the null (empty) string.

-f *listfile*

Specifies a file (*listfile*) containing a list of files to be restored and can be used instead of the *filenames* option. In *listfile*, list each file path on a separate line.

The format required for the file list depends on whether the files have spaces or newlines in the names.

To restore files that do not have spaces or newlines in the names, use this format:

filepath

Where *filepath* is the path to the file that you are restoring. For example:

/home

/etc

/var

To restore files that have spaces or newlines in the names, use one of the following formats:

filepathlen filepath

filepathlen filepath start_date_time end_date_time

filepathlen filepath -s datetime -e datetime

Where:

filepath is the path to the file you are restoring.

filepathlen is the total number of characters in the file path.

start_date_time and *end_date_time* are the decimal number of seconds since 01/01/70 00:00:00.

datetime is the same as the command line (*mm/dd/yy [hh[:mm[:ss]]]*).

The start and end date and time specified on the command line is used unless a line in *listfile* overrides it. The dates may change from line to line.

The following is an example that uses *filepathlen filepath*:

```
5 /home
4 /etc
4 /var
19 /home/abc/test file
```

filenames Names one or more files to be restored and can be used instead of the *-f* option.

Any files that you specify must be listed at the end, following all other options. You must also specify absolute file paths.

NOTES

The format that you must use for date and time values in NetBackup commands depends on the locale setting. The examples in this command description are for a locale setting of C.

If you are uncertain of the NetBackup command requirements for your locale, enter the command with the *-help* option and check the usage. For example:

```
/usr/openv/netbackup/bin/bprestore -help
USAGE: bprestore [-A | -B] [-K] [-l | -H | -y] [-r] [-T]
        [-L progress_log] [-R rename_file] [-C client]
        [-D client] [-S master_server] [-t class_type]
        [-c class] [-k "keyword phrase"]
        [-s mm/dd/yy [hh:mm:ss]] [-e mm/dd/yy [hh:mm:ss]]
        [-w [hh:mm:ss]] -f listfile | filenames
```



Notice the month/day/year and hours:minutes:seconds requirements for the `-s` and `-e` options. These are for a locale setting of `C` and can be different for other locales.

For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

EXAMPLE 1

To restore a file from backups of `/usr/user1/file1` that were performed between `04/01/93 06:00:00` and `04/10/93 18:00:00`, enter the following (all on one line):

```
bprestore -s 04/01/93 06:00:00 -e 04/10/93 18:00:00
/usr/user1/file1
```

EXAMPLE 2

To restore files listed in a file named `restore_list` by using the most recent backups, enter the following:

```
bprestore -f restore_list
```

EXAMPLE 3

To restore the directory `/home/kwc` from the backups that are associated with a keyword phrase that contains "My Home Directory" and use a progress log named `/home/kwc/bkup.log`, enter the following (all on one line):

```
bprestore -k "My Home Directory*" -L /home/kwc/bkup.log
/home/kwc
```

EXAMPLE 4

To restore the `D` drive on the Windows NT client `slater` from the backups that are associated with a keyword phrase that contains "My Home Dir" and use a progress log named `/home/kwc/bkup.log`, enter the following (all on one line):

```
bprestore -k "My Home Dir*" -C slater -D slater -t 13 -L
/home/kwc/bkup.log /D
```

EXAMPLE 5

Assume you have a rename file named `/home/kwc/rename` on a UNIX client and it contains the following:

```
change /home/kwc/linkback to /home/kwc/linkback_alt
```

To restore the hard link named `/home/kwc/linkback` to alternate path `/home/kwc/linkback_alt` on that client, execute:

```
bprestore -H -R /home/kwc/rename /home/kwc/linkback
```



FILES

`$HOME/bp.conf`

`/usr/opensv/netbackup/logs/bprestore/log.mmddy`

SEE ALSO

`bp(1)`, `bparchive(1)`, `bpbackup(1)`, `bplist(1)`



bpstuadd(1M)

NAME

bpstuadd - create a NetBackup storage unit

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpstuadd -label
    storage_unit_label -path path_name | {-density density [-rt
    robot_type -rn robot_number]} [-host host_name] [-cj
    max_jobs] [-odo on_demand_only] [-mfs max_fragment_size]
    [-maxmpx mpx_factor] [-nh NDMP_attach_host] [-verbose]
    [-M master_server, ...]
```

DESCRIPTION

bpstuadd creates a NetBackup storage unit. The command line must include a label for the new storage unit and either the `-density` or the `-path` option. bpstuadd will not create the storage unit if the master server has already created the maximum number of storage units allowed by its NetBackup configuration. bpstuadd will not create a storage unit that specifies the same destination medium as an existing storage unit.

There are several types of storage units. The storage-unit type affects how NetBackup stores the data. The options on the bpstuadd command line determine the storage-unit type, which is one of the following:

- ◆ **Disk.** The storage destination is a disk filesystem directory.
- ◆ **Media Manager.** The storage destination is a medium (a tape or optical device) managed by the Media Manager.
- ◆ **NDMP.** An NDMP storage unit is controlled by Media Manager. The NetBackup for NDMP option must be installed. Where the Media Manager storage-unit type is discussed in this command description, the discussion also applies to the NDMP storage-unit type, unless it is specifically excepted. The media for an NDMP storage unit always attach directly to an NDMP host and cannot be used to store data for other NetBackup clients. When defining an NDMP storage unit, bpstuadd must be executed on the master server. Refer to the NetBackup for NDMP System Administrator's Guide for more information on adding NDMP storage units.

Errors go to stderr. A log of the command's execution goes to the NetBackup admin log file for the current day.

This command requires root privileges.

See the NetBackup system administrator's guide for additional information on storage units.



OPTIONS

`-cj max_jobs`

The maximum number of concurrent jobs permitted for this storage unit. *max_jobs* is a non-negative integer. The appropriate value depends on your server's ability to comfortably execute multiple backup processes and the available space on the storage media. Also, refer to Maximum Jobs per Class in the NetBackup system administrator's guide.

0 means that this storage unit will never be selected when a job is being scheduled. The default is 1.

`-density density`

If this option is present, the storage unit type is Media Manager. There is no default for this option. Either `-density` or `-path` must be on the command line. Do not use `-path` when `-density` is being used. If the robot type is specified on the command line, the value for *density* should be consistent with the robot type.

Valid *density* types are:

`d1t` - DLT Cartridge

`d1t2` - DLT Cartridge alternate

`8mm` - 8mm Cartridge

`4mm` - 4mm Cartridge

`qscsi` - 1/4 Inch Cartridge

Note The following densities are supported only on NetBackup DataCenter servers.

`hcart` - 1/2 Inch Cartridge

`hcart2` - 1/2 Inch Cartridge alternate

`dtf` - DTF Cartridge

`odiskwm` - Optical Disk Write-Many

`odiskwo` - Optical Disk Write-Once

`-help`

Print a command-line usage message. When `-help` is present, it is the only option on the command line.

`-host host_name`

Note NetBackup BusinessServer does not support remote media servers.

The NetBackup host that is associated with the destination media. The default is the hostname of the local system.



The host you select must be either your NetBackup master server or a remote media server (if you are configuring remote media servers). The host name must be the network name for the server as known by all NetBackup servers and clients.

If *host_name* is a valid network name, but it has not been configured in NetBackup previously, *host_name* will be added to NetBackup's configuration as a media server. On UNIX, this shows up as a `SERVER` entry in the `bp.conf` file; on Windows NT, this shows up on the Servers tab in the server properties dialog box in the NetBackup configuration window. If *host_name* is not a valid network name, you must configure it manually.

`-label` *storage_unit_label*

The name of the storage unit. This is a required option. The maximum length of a storage-unit label is 128 characters.

`-mfs` *max_fragment_size*

The maximum fragment size specifies, in megabytes, how large a fragment for a NetBackup image can be.

For a Media Manager storage unit, this value is either zero (the fragment size is unlimited, meaning there is no fragmentation) or any integer greater than or equal to 50 megabytes (MB). The default value is 0.

For a Disk storage unit, this value ranges from 20 megabytes to 2000 megabytes (2 gigabytes). The default value is 2000 (2 gigabytes).

`-maxmpx` *mpx_factor*

The maximum multiplexing factor. Multiplexing sends concurrent, multiple backups from one or several clients to a single drive. Refer to the topic "Multiplexing (MPX)" in the NetBackup system administrator's guide.

The multiplexing factor can range from 1 to 32. 1 means no multiplexing. A value greater than 1 means that NetBackup can create multiplexed images on the destination medium. Licensing determines the effective subset of the 1..32 range for the local NetBackup installation.

The default is 1.

`-M` *master_server*

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point. The default is the master server for the system where the command is entered.

- `-nh NDMP_attach_host`
Specifies the hostname of the NDMP server. If this option is present, the storage unit type is set to NDMP. The default is no NDMP server.
- `-odo on_demand_only`
The On-Demand-Only flag controls the condition under which NetBackup uses the storage unit:
To make the storage unit available only to classes or schedules that request it, set the flag to 1 (enabled).
To make the storage unit available to any class or schedule, set the flag to 0 (disabled).
If the storage unit's type is Disk, the default is 1; NetBackup uses the storage unit only when explicitly requested. Otherwise, the default is 0.
- `-path path_name`
The path to a disk filesystem, expressed as an absolute pathname. This is the data storage area for this storage unit. When this option is present, the storage unit type is Disk. There is no default for this option. Either `-path` or `-density` must be on the command line. Do not use `-density` when `-path` is being used.
In general when this option is used, it is recommended that the On-Demand-Only flag be enabled (see `-odo`). Otherwise, any NetBackup class that does not require a specific storage unit has the opportunity to fill the disk filesystem `path_name`. This can cause serious system problems. For instance, if the system swap area happens to be on the same filesystem, new processes may fail.
- `-rn robot_number`
The robot number for this storage unit. The robot number must be greater than or equal to 0. The robot number can be obtained from the Media Manager device configuration. The system administrator's guide for Media Manager discusses the rules concerning the use of this number. This option is ignored unless the `-rt` option is present. There is no default for this option.
- `-rt robot_type`
The robot type for this storage unit. For non-robotic (standalone) devices select `NONE` or omit this option. The default value is `NONE` (Not Robotic). The value for density should be consistent with the robot type.
If this option is set to any value other than `NONE`, the `-rn` option is required. Available robot type codes are:
`NONE` - Not Robotic
`TLD` - Tape Library DLT
`TSD` - Tape Stacker DLT



ACS - Automated Cartridge System
TS8 - Tape Stacker 8MM
TL8 - Tape Library 8MM
TL4 - Tape Library 4MM
ODL - Optical Disk Library
TSH - Tape Stacker Half-inch
TLH - Tape Library Half-inch
TLM - Tape Library Multimedia
LMF - Library Management Facility
RSM - Removable Storage Manager

-verbose

Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).

EXAMPLES

Example 1:

Create a new storage unit, named `hatunit`. Its storage unit type is `Disk`. The path for the storage unit is `/tmp/hatdisk`:

```
bpstuadd -label hatunit -path C:\tmp\hatdisk/tmp/hatdisk -verbose
<2>bpstuadd: INITIATING: NetBackup 3.2Beta created: 98121513
<2>bpstuadd: EXIT status = 0.
```

Example 2:

Note This example refers to remote media servers and applies only to NetBackup DataCenter. NetBackup BusinessServer supports only a master server, not remote media servers.

Create a storage unit using a UNIX server, which has not been configured previously in NetBackup:

```
mo% bpstuadd -label parrot_stu -host parrot -density dlt -rt TLD -rn 2
```

The remote media server `parrot` was added to the `bp.conf` file.

You must also install NetBackup and Media Manager on `parrot` and run the `add_slave_on_clients` shell script on `mango`.

```
mango% grep parrot /usr/opensv/netbackup/bp.conf
SERVER = parrot
SERVER = parrot
```

EXIT STATUS

= 0 The command executed successfully.
> 0 An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form:

`bpstuadd: EXIT status = exit status`

If an error occurred, a diagnostic precedes this message.

FILES

`/usr/opensv/netbackup/logs/admin/*`

`/usr/opensv/netbackup/db/config/storage_units`

SEE ALSO

`bpstudel(1m)`, `bpstulist(1m)`, `bpsturep(1m)`

The NetBackup system administrator's guide has more information on NetBackup storage units.



bpstudel(1M)

NAME

bpstudel - delete a NetBackup storage unit

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpstudel -label  
storage_unit_label [-verbose] [-M master_server, ...]
```

DESCRIPTION

bpstudel deletes a NetBackup storage unit. The command must include a label for the storage unit.

If bpstudel cannot delete the storage unit, for instance, if the storage unit label is mistyped on the command line, bpstudel does not return an error message. You can run bpstulist to verify that the storage unit was deleted.

Errors are sent to stderr. A log of the command's execution is sent to the NetBackup admin log file for the current day.

This command requires root privileges.

See your NetBackup system administrator's guide for additional information on storage units.

OPTIONS

-label *storage_unit_label*

The name of the storage unit. This is a required option. The maximum length for a storage-unit label is 128 characters.

-M *master_server_*

A list of master servers. This is a comma-separated list of host names. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point. The default is the master server for the system where the command is entered.

-verbose Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the /usr/opensv/netbackup/logs/admin directory is defined).

EXAMPLES

Delete the storage unit named tst.dsk, listing the existing storage units before and after the deletion:




```
bpstulist
stuunit 0 mango 0 -1 -1 1 0 /tmp/stuunit 1 1 2000 *NULL*
tst.dsk 0 mango 0 -1 -1 3 0 /hsm3/dsk 1 1 2000 *NULL*
```

```
bpstudel -label tst.dsk
```

```
bpstulist
stuunit 0 mango 0 -1 -1 1 0 /tmp/stuunit 1 1 2000 *NULL*
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
/usr/opensv/netbackup/db/config/storage_units
```

SEE ALSO

bpstuadd(1m), bpstulist(1m), bpsturep(1m)

The NetBackup system administrator's guide has more information on NetBackup storage units.



bpstulist(1M)

NAME

bpstulist - display one or all of the NetBackup storage units

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpstulist [-label  
storage_unit_label] [-L|-l|-U] [-verbose] [-M  
master_server, . . .]
```

DESCRIPTION

bpstulist displays the attributes for a NetBackup storage unit. If no storage unit label is specified, bpstulist displays the attributes for all NetBackup storage units.

Errors are sent to stderr. A log of the command's execution is sent to the NetBackup admin log file for the current day.

This command requires root privileges.

See your NetBackup system administrator's guide for additional information on storage units.

OPTIONS

List-type options:

- L The list type is long. This option produces a listing with one storage-unit attribute per line, in the format *storage-unit attribute: value*. Some attribute values are expressed in both interpreted and raw form. For instance, a robot-type entry might be TL4 (7) (7 is NetBackup's internal value for a TL4 robot).
For a disk storage unit, a long listing has these attributes for each storage unit:
 - Label
 - Media Type (this is the storage-unit type)
 - Host Connection
 - Concurrent Jobs
 - On Demand Only
 - Path
 - Robot Type (not robotic)
 - Max Fragment Size
 - Max MPX

For a Media Manager storage unit, a long listing has these attributes for each storage unit:

- Label
- Media Type (this is the storage-unit type)
- Host Connection
- Number of Drives
- On Demand Only
- Density
- Robot Type/Number
- Max Fragment Size
- Max MPX/drive

-l The list type is short. This produces a terse listing. This option is useful for scripts or programs that rework the listing contents into a customized report format. This is the default list type.

A single line contains the information for a storage unit, with all attribute values expressed in raw form. The fields on this line are:

- label
- storage unit type
- host
- robot_type
- robot_number
- density
- concurrent_jobs
- initial_mpx
- path
- on_demand_only
- max_mpx
- maxfrag_size
- ndmp_attach_host

-U The list type is user. This option produces a listing with one storage-unit attribute per line, in the format *storage-unit attribute: value*. Attribute values are expressed in interpreted form. For instance, a robot-type value might be TL4, instead of 7.

For a disk storage unit, a user-type listing has these attributes for each storage unit:

- Label



- Media Type (this is the storage-unit type)
- Host Connection
- Concurrent Jobs
- On Demand Only
- Max MPX
- Path
- Max Fragment Size

For a Media Manager storage unit, a user-type listing has these attributes for each storage unit:

- Label
- Media Type (this is the storage-unit type)
- Host Connection
- Number of Drives
- On Demand Only
- Max MPX/drive
- Density
- Robot Type/Number
- Max Fragment Size

Here are the remaining options for `bpstulist`:

- label *storage_unit_label*
The name of the storage unit. If this option is not present, the listing is for all storage units. The maximum length for a storage-unit label is 128 characters.
- M *master_server, . . .*
A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.
- verbose Select verbose mode for logging. This is only meaningful when running with activity logging turned on (the `\admin/usr/opensv/netbackup/logs/admin` directory is defined).

EXAMPLES

List the storage units defined on the master server `apricot`, using the `-U` display option:

```
bpstulist -U -M apricot
```



```
Label:                redtest
Storage Unit Type:    Disk
Host Connection:      apricot
Concurrent Jobs:      1
On Demand Only:       yes
Max MPX:              4
Path:                 /usr/redtest
Max Fragment Size:    2000 MB
```

```
Label:                bluetest
Storage Unit Type:    Media Manager
Host Connection:      apricot
Number of Drives:     6
On Demand Only:       yes
Max MPX/drive:        1
Density:              4mm - 4mm Cartridge
Robot Type/Number:    TL4 / 0
Max Fragment Size:    (unlimited)
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
/usr/opensv/netbackup/db/config/storage_units
```

SEE ALSO

bpstuadd(1m), bpstudel(1m), bpsturep(1m)

The NetBackup system administrator's guide has more detailed information on NetBackup storage units.



bpsturep(1M)

NAME

bpsturep - replace selected NetBackup storage unit attributes

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpsturep -label
    storage_unit_label [-host host_name] [-cj max_jobs] [-odo
on_demand_only] [-mfs max_fragment_size] [-maxmpx
mpx_factor] [{-path path_name} | {-density density_type
[-rt robot_type -rn robot_number]}] [-nh NDMP_attach_host]
[-verbose] [-M master_server, ...]
```

DESCRIPTION

bpsturep modifies an existing NetBackup storage unit by replacing selected storage-unit attributes in the NetBackup catalog. The command line must include a label for the storage unit. The label is the only storage-unit attribute that bpsturep cannot modify.

Use bpsturep with care. The changes to the storage unit must be compatible with the existing attributes of the storage unit. Make sure resulting attribute combinations are valid, especially for the following attributes:

robot_type

robot_number

density

max_fragment_size

path

NDMP_attach_host

The safest way to modify these attributes is to run bpsturep once for each attribute to be replaced.

bpsturep makes the changes by deleting the old storage unit and adding a new storage unit with the specified attribute changes. Therefore, if bpsturep specifies invalid options or an invalid combination of options, the storage unit may be deleted without being re-added. It is best to run bpstulist after bpsturep to determine whether the intended changes were actually applied.

Errors go to stderr. A log of the command's execution goes to the NetBackup admin log file for the current day.

This command requires root privileges.



See your NetBackup system administrator's guide for additional information on storage units.

OPTIONS

`-cj` *max_jobs*

The maximum number of concurrent jobs permitted for this storage unit. *max_jobs* is a non-negative integer. The appropriate value depends on your server's ability to comfortably execute multiple backup processes and the available space on the storage media. Also, refer to the Maximum Jobs per Class topic in your NetBackup system administrator's guide.

0 means that this storage unit will never be selected when a job is being scheduled. The default is 1.

`-density` *density_type*

If this option is present, the storage unit type is Media Manager. There is no default for this option. One of `-density` or `-path` must be on the command line, but not both. If the command line includes a robot type, the value for density should be consistent with the robot type.

Valid density types are:

`d1t` - DLT Cartridge

`d1t2` - DLT Cartridge alternate

`8mm` - 8mm Cartridge

`4mm` - 4mm Cartridge

`qscsi` - 1/4 Inch Cartridge

Note The following densities apply only to NetBackup DataCenter servers.

`hcart` - 1/2 Inch Cartridge

`hcart2` - 1/2 Inch Cartridge alternate

`dtf` - DTF Cartridge

`odiskwm` - Optical Disk Write-Many

`odiskwo` - Optical Disk Write-Once

`-host` *host_name*

Note NetBackup BusinesServer does not support remote media servers.

The NetBackup host to which the destination media is attached. The default is the hostname of the local system.



The host you select must be either your NetBackup master server or a media server (if you are configuring media servers). The host name must be the network name for the server as known by all NetBackup servers and clients.

If *host_name* is a valid network name and is not yet configured in NetBackup, the value *host_name* will be added to NetBackup's configuration as a media server. On UNIX, this shows up in `bp.conf`; on Windows NT, this shows up in the Configuration window for Servers. If *host_name* is not a valid network name, you must configure it manually.

`-label storage_unit_label`

The name of a storage unit. This is the storage unit whose attributes `bpsturep` replaces. This is a required option. The maximum length of a storage-unit label is 128 characters.

`-mfs max_fragment_size`

The maximum fragment size specifies, in megabytes, how large a fragment for a NetBackup image can be. For a Media Manager storage unit, this value is either zero (the fragment size is unlimited, meaning there is no fragmentation) or any integer greater than or equal to 50 megabytes (MB). The default value is 0.

For a Disk storage unit, this value ranges from 20 megabytes to 2000 megabytes (2 gigabytes). The default value is 2000 (2 gigabytes).

`-maxmpx mpx_factor`

The maximum multiplexing factor. Multiplexing sends concurrent, multiple backups from one or several clients to a single drive. Refer to the topic "Multiplexing (MPX)" in your NetBackup system administrator's guide.

The multiplexing factor can range from 1 to 32, where 1 means no multiplexing. A value greater than 1 means that NetBackup can create multiplexed images on the destination medium. Depending on the licensing of the local NetBackup installation, it may not be possible to assign multiplexing factors in the entire range 1..32.

The default is 1.

`-M master_server_`

A list of master servers. This is a comma-separated list of hostnames. If this option is present, the command is executed on each of the master servers in this list. The master servers must allow access by the system issuing the command. If an error occurs for any master server, processing stops at that point in the list. The default is the master server for the system where the command is entered.

- `-nh NDMP_attach_host`
Specifies the hostname of the NDMP server. If this option is present, the storage unit type is set to NDMP. The default is no NDMP server.
- `-odo on_demand_only`
The *on-demand-only* flag controls whether the storage unit is used only for backups that explicitly request (demand) the storage unit:
To make the storage unit available only to classes or schedules that request it, set the flag to 1 (enabled).
To make the storage unit available to any class or schedule, set the flag to 0 (disabled).
If the storage unit's type is Disk, the default is 1; NetBackup uses the storage unit only when explicitly requested. Otherwise, the default is 0.
- `-path path_name`
The path to a disk filesystem, expressed as an absolute pathname. This is the data storage area for this storage unit. When this option is present, the storage unit type is Disk. There is no default for this option. One of `-density` or `-path` must be on the command line, but not both.
In general when this option is used, it is recommended that the *on-demand-only* flag be enabled (see `-odo`). Otherwise, any NetBackup class that does not require a specific storage unit has the opportunity to fill the disk filesystem *path_name*. This can cause serious system problems. For instance, if the system swap area happens to be on the same filesystem, new processes may fail.
- `-rn robot_number`
The robot number for this storage unit. The robot number must be greater than or equal to 0. The robot number can be obtained from the Media Manager device configuration. The system administrator's guide for Media Manager discusses the rules concerning the use of this number. This option is ignored unless the `-rt` option is present. There is no default for this option.
- `-rt robot_type`
The robot type for this storage unit. For non-robotic (standalone) devices select NONE or omit this option. The default value is NONE (Not Robotic). The value for density should be consistent with the robot type.
If this option is set to any value other than NONE, the `-rn` option is required.
Available robot type codes are:
NONE - Not Robotic
TLD - Tape Library DLT
TSD - Tape Stacker DLT



ACS - Automated Cartridge System

TS8 - Tape Stacker 8MM

TL8 - Tape Library 8MM

ODL - Optical Disk Library

TSH - Tape Stacker Half-inch

TLH - Tape Library Half-inch

TLM - Tape Library Multimedia

LMF - Library Management Facility

RSM - Removable Storage Manager

`-verbose` Select `verbose` mode for logging. This is only meaningful when running with activity logging turned on (the `/usr/opensv/netbackup/logs/admin` directory is defined).

EXAMPLES

Change the path for a disk storage unit, `mkbunit`. The path is changed from `/tmp/mkbunit` to `/tmp/mkbunit2`:

```
bpstulist
```

```
mkbunit 0 beaver 0 -1 -1 1 0 /tmp/mkbunit 1 1 2000 *NULL*
```

```
bpsturep -label mkbunit -path /tmp/mkbunit2
```

```
bpstulist
```

```
mkbunit 0 beaver 0 -1 -1 1 0 /tmp/mkbunit2 1 1 2000 *NULL*
```

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

```
/usr/opensv/netbackup/db/config/storage_units
```

SEE ALSO

```
bpstuadd(1m), bpstudel(1m), bpstulist(1m)
```

The NetBackup system administrator's guide has more detailed information on NetBackup storage units.



bpverify(1M)

NAME

bpverify - verify the backups created by NetBackup.

SYNOPSIS

```
/usr/opensv/netbackup/bin/admincmd/bpverify [-l] [-p] [-pb] [-v]
[-client name] [-st sched_type] [-sl sched_label] [-L
output_file] [-class name] [-s mm/dd/yyyy hh:mm:ss] [-e
mm/dd/yyyy hh:mm:ss] [-M master_server] [-Bidfile
file_name] [-ct class_type] [-hoursago hours] [-cn copy
number] [-backupid backup_id] [-id media_id | path]
```

DESCRIPTION

bpverify verifies the contents of one or more backups by reading the backup volume and comparing its contents to the NetBackup catalog. This operation does not compare the data on the volume with the contents of the client disk. However, it does read each block in the image, thus verifying that the volume is readable. NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

If either `-Bidfile` or `-backupid` is specified, bpverify uses this option as the sole criterion for selecting the set of backups it will verify. If the command line does not contain `-Bidfile` or `-backupid`, then bpverify selects the backups that satisfy all the selection options. For instance, if the command line looks like

```
bpverify -ct Standard -hoursago 10
```

then bpverify verifies the set of backups with class type `Standard` that have been run in the past 10 hours.

If `-p` or `-pb` is specified, bpverify previews the set of backups that meet the selection criteria. In this case, bpverify displays the backup IDs, but does not perform the verification.

bpverify sends its error messages to `stderr`. bpverify sends a log of its execution to the NetBackup admin log file for the current day.

This command requires root privileges.

For more information on verifying backups, see your NetBackup system administrator's guide.



OPTIONS

- `-Bidfile` *file_name*
file_name specifies a file that contains a list of backup IDs to be verified. The file contains one backup ID per line. If this option is specified, other selection criteria are ignored. The default is no file of backup IDs, meaning any backup can be verified.
- `-backupid` *backup_id*
Specifies the backup ID of a single backup to verify. This option takes precedence over any other selection criteria except `-Bidfile`. The default is any backup.
- `-class` *name*
Search for backups to verify in the specified class. The default is any class.
- `-client` *name*
Specifies the name of the client that produced the original backup. The default is any client.
- `-cn` *copy_number*
Copy number (1 or 2) of the backup ID to verify. The default copy number is 1.
- `-ct` *class_type*
Specifies the class type for selecting backups to verify. The default is any class type.
The valid class types are the following:
AFS
Apollo-wbak
Auspex-FastBackup
DataTools-SQL-BackTrack
DB2
Extensible-Client
FlashBackup
Informix-On-BAR
Lotus-Notes
MS-Exchange-Server
MS-SQL-Server
MS-Windows-NT
NCR-Teradata
NDMP
NetWare



-
- Oracle
 - OS/2
 - SAP
 - Split-Mirror
 - Standard
 - Sybase
- e** *mm/dd/yyyy [hh[:mm[:ss]]]*
 Specifies the end of the time range for selecting backups to verify. The **-s** option or the **-hoursago** option specifies the start of the range. The date and time format depend on the user's locale. For more information, see `locale(1)`.
 The default ending time is the current date and time.
- help**
 Prints a synopsis of command usage when **-help** is the only option on the command line.
- hoursago** *hours*
 Specifies the number of hours before the current time to search for backups. This is equivalent to specifying a start time (**-s**) of the current time minus hours. Do not use both this option and the **-s** option.
 Hours is a non-negative integer. The default starting time is 24 hours ago.
- id** *media_id | path*
 Search the image catalog for backups to verify that are on this media ID or pathname. If a backup has some fragments on this media ID and some fragments on another media ID, NetBackup skips verifying that backup. For images stored on disk rather than removable media, specify an absolute pathname instead of *media_id*. The default is any media ID or pathname.
- L** *output_file*
 Specifies the name of a file in which to write progress information. The default is to not use a progress file, in which case the progress information is written to `stderr`. For additional information, see `DISPLAY FORMATS` later in this command description.
- l**
 Specifies that the list type is long, causing `bpverify` to write additional information to the progress log. The default list type is short. For additional information, see `DISPLAY FORMATS` later in this command description.



- `-M master_server`
Specifies the master server that provides the bpverify image data. The master server must allow access by the system issuing the bpverify command. The default is the master server for the system where bpverify is entered:
For NetBackup BusinessServer:
The default is always the master server where the command is entered.
For NetBackup DataCenter:
If the command is entered on a master server, then that server is the default.
If the command is entered on a remote media server, then the master for that media server is the default.
- `-p`
Previews the verification, but does not perform the verification. For additional information, see DISPLAY FORMATS later in this command description.
- `-pb`
Previews the verification but does not perform the verification. This is similar to the `-p` option, but `-pb` does not display information about the individual backups. For additional information, see DISPLAY FORMATS later in this command description.
- `-s mm/dd/yyyy [hh[:mm[:ss]]]`
Specifies the start of the range of dates and times that include all backups to verify. The `-e` option specifies the end of the range. The date and time format depend on the user's locale. For more information, see locale(1). The default is 24 hours ago.
- `-sl sched_label`
Search for backups to verify that were created by the specified schedule. The default is all schedules.
- `-st sched_type`
Search for backups to verify that were created by the specified schedule type. The default is any schedule type.
Valid values are:
FULL (full backup)
INCR (differential-incremental backup)
CINC (cumulative-incremental backup)
UBAK (user backup)
UARC (user archive)
NOT_ARCHIVE (all backups except user archive)
- `-v`
Selects verbose mode. When `-v` is specified, the debug and progress logs include more information. The default is not verbose.

DISPLAY FORMATS

PREVIEW DISPLAYS:

`bpverify` runs a preview by searching for backups and displaying them. `bpverify` does not actually verify the backups.

- ◆ The `-p` display lists backup IDs that meet the criteria set by the `bpverify` command-line options. The `-p` display is ordered by volume. For each volume containing a selected backup, the media ID and server are displayed, followed by the selected backup IDs that reside on that volume
- ◆ The `-pb` display is a brief version of the `-p` display. It lists the media ID and server for each volume that contains backups that meet the selection criteria.

VERIFICATION DISPLAYS:

`bpverify` creates these displays as it verifies images. If the `bpverify` command line contains no option to set the list format, the display format is short. If the command line contains `-l`, the display format is long. If the command line contains both `-l` and `-L`, `bpverify` creates a file containing the progress log.

The verification display is ordered by volume.

- ◆ In long format, `bpverify` displays the following information for each selected backup ID:
 - Class, schedule, backup ID, media ID or path, and creation time
 - Files backed up
 - Any problems that `bpverify` detects while verifying the image
 - Whether the image verification is successful or not
- ◆ In short format, `bpverify` omits listing the files backed up.

NOTES

The format that you must use for date and time values in NetBackup commands vary according to your locale setting. The examples in this command description are for a locale setting of C. If you are uncertain of the NetBackup command requirements for your locale, enter the command with the `-help` option and check the USAGE.

For example:

```
/usr/openv/netbackup/bin/admincmd/bpverify -help
bpverify: [-l] [-p] [-pb] [-v] [-client <name>]
          [-st <sched_type>] [-sl <sched_label>] [-L <output_file>]
          [-class <name>] [-s mm/dd/yyyy hh:mm:ss] [-e mm/dd/yyyy hh:mm:ss]
          [-M master_server]
```



```
[-Bidfile <file_name>]
[-ct <class_type>] [-hoursago <hours>] [-cn <copy number>]
[-backupid <backup_id>] [-id <media_id or path>]
```

Notice the month/day/year and hours:minutes:seconds requirements for the `-s` and `-e` options. These are for a locale setting of C and can be different for other locales. For more information on locale, see the `locale(1)` man page for your system.

EXAMPLES

Example 1:

This example verifies the backups run in the past 36 hours:

```
bpverify -hoursago 36
Verify started Thu Feb  3 11:30:29 2000
INF - Verifying class mkb_class, schedule Full (plum_0949536546), path
/tmp/mkbunit, created 02/02/00 18:09:06.
INF - Verify of class mkb_class, schedule Full (plum_0949536546) was
successful.
INF - Status = successfully verified 1 of 1 images.
```

Example 2:

This example compares the two preview displays, `-p` and `-pb`:

```
bpverify -p -hoursago 200
Media id = A00002  Server = plum
Bid = plum_0949616279  Kbytes = 32800  Filenum = 1  Fragment = 1
Bid = guava_0949681647  Kbytes = 12191  Filenum = 2  Fragment = 1
Bid = guava_0949683298  Kbytes = 161  Filenum = 3  Fragment = 1
Bid = guava_0949683671  Kbytes = 11417  Filenum = 4  Fragment = 1
Bid = guava_0949684009  Kbytes = 11611  Filenum = 5  Fragment = 1
Bid = guava_0949684276  Kbytes = 806  Filenum = 6  Fragment = 1
Bid = guava_0949688704  Kbytes = 9869  Filenum = 7  Fragment = 1
Bid = guava_0949688813  Kbytes = 9869  Filenum = 8  Fragment = 1
Bid = guava_0949949336  Kbytes = 10256  Filenum = 9  Fragment = 1
Bid = plum_0949949337  Kbytes = 6080  Filenum = 9  Fragment = 1
Bid = plum_0949949337  Kbytes = 4176  Filenum = 10  Fragment = 2
Bid = guava_0949949686  Kbytes = 10256  Filenum = 11  Fragment = 1
Bid = plum_0949949687  Kbytes = 5440  Filenum = 11  Fragment = 1
```



```
Bid = plum_0949949687  Kbytes = 4816  Filenum = 12  Fragment = 2
Bid = guava_0949949902  Kbytes = 10256  Filenum = 13  Fragment = 1
Bid = plum_0949949901  Kbytes = 8832  Filenum = 13  Fragment = 1
Bid = plum_0949949901  Kbytes = 1424  Filenum = 14  Fragment = 2
Bid = plum_0950053561  Kbytes = 10256  Filenum = 15  Fragment = 1
```

```
Media id = 400032  Server = plum
```

```
Bid = toaster2_0950199621  Kbytes = 298180  Filenum = 1  Fragment = 1
```

```
Bid = toaster2_0950199901  Kbytes = 298180  Filenum = 3  Fragment = 1
```

```
bpverify -pb -hoursago 200
```

```
Media id = A00002  Server = plum
```

```
Media id = 400032  Server = plum
```

EXIT STATUS

◆ = 0

The command executed successfully.

◆ <> 0

An error occurred.

If administrative logging is enabled, the exit status is logged in the administrative daily log under the directory `/usr/opensv/netbackup/logs/admin` in the form:

`bpverify: EXIT status = exit status`

If an error occurred, a diagnostic precedes this message.

FILES

```
/usr/opensv/netbackup/logs/admin/*
```

```
/usr/opensv/netbackup/db/error/*
```

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

SEE ALSO

NetBackup system administrator's guide



vopied(1M)

NAME

vopied - daemon to provide VERITAS One-time Password user authentication

SYNOPSIS

```
/usr/opensv/bin/vopied [-standalone] [-debug] [-portnum number]
                        [-max_time seconds] [-log_dir path] [-severity mask]
```

DESCRIPTION

This program is available on Windows NT/2000 and UNIX NetBackup clients. It accepts connections from remote NetBackup servers and clients that are attempting to verify the identity of requests from the local NetBackup system. The authentication method is VERITAS One-time Password (vopie). Normally, vopied is started by the NetBackup Client service on Windows NT/2000 and inetd on UNIX.

When you install NetBackup on a Windows NT/2000 client or UNIX client, the installation process adds entries for vopied to

C:\WINNT\system32\drivers\etc\services on Windows NT/2000 and /etc/services and /etc/inetd.conf on UNIX.

The services entry looks like this:

```
vopied 13783/tcp      vopied
```

The inetd.conf entry on UNIX looks like this:

```
vopied stream tcp      nowait  root    /usr/opensv/bin/vopied vopied
```

OPTIONS

-standalone

Available only on UNIX clients and specifies that vopied will run continuously rather than being started by inetd.

-debug

Available only on UNIX clients and implies -standalone (that is, vopied runs continuously). This option prevents vopied from forking and does not disconnect it from standard input, output, and error.

-portnum *number*

Available only on UNIX clients and implies -standalone (that is, vopied runs continuously). Specifies the port number where vopied listens for requests. The default is the vopied entry in:

```
/etc/services
```



-
- `-max_time` *seconds*
Specifies a time out value for network connections. The default is 60 seconds.
- `-log_dir` *path*
Specifies the directory where the `vopied` log directory resides. The default is:
`install_path\NetBackup\logs` (Windows NT/2000)
`/usr/opensv/logs` (UNIX)
To enable logging, create a `vopied` directory in the *path* directory before starting `vopied`. For example:
`/usr/opensv/logs/vopied`
- `-severity` *mask*
Specifies the type of messages to be logged. *mask* is the sum of zero or more of these values:
1 Unknown
2 Debug
4 Information
8 Warning
16 Error
32 Critical
The default is 48 decimal (0x30 hexadecimal or 060 octal), which specifies critical and error.

SEE ALSO

`bpauthsync(1M)`, `vopie_util(1M)`



vopie_util(1M)

NAME

vopie_util - manage local vopie authentication files

SYNOPSIS

```
/usr/opensv/bin/vopie_util [-log_dir path] [-severity mask]  
[-debug] [-local_name name] [-always_write] [-hashed |  
-unhashed] remote_name [sequence seed hash]
```

DESCRIPTION

The vopie_util program is available on Windows NT/2000 and UNIX NetBackup servers and clients. It updates the hashed (public) and unhashed (secret) key files for the vopie authentication method on the local system. Typically, vopie_util is used to synchronize the vopie key files between two systems.

OPTIONS

-log_dir *path*

Specifies the directory where the vopie_util log directory resides. The default is:

install_path\NetBackup\logs (Windows NT/2000)

/usr/opensv/logs (UNIX)

To enable logging, create a vopie_util directory in the *path* directory before starting vopie_util. For example:

/usr/opensv/logs/vopie_util

-severity *mask*

Specifies the type of messages to be logged. *mask* is the sum of zero or more of these values:

1 Unknown

2 Debug

4 Information

8 Warning

16 Error

32 Critical

The default is 48 decimal (0x30 hexadecimal or 060 octal), which specifies critical and error.

-debug

Specifies that additional information is logged to standard error.



- `-local_name name`
 Specifies the name of the local system. The default is the network host name of the system. We recommend that this name match the NetBackup client name for the system.
- `-always_write`
 Always update the file even if it already exists. The default is to not overwrite existing files.
- `-hashed`
 Updates the hashed (public) key file. This file contains the challenges that this system presents to other systems during authentication. If the *sequence*, *seed*, and *hash* options described below are not specified, the hashed-key file data matches any secret key.
- `-unhashed`
 Updates the unhashed (secret) key file. A secret key is randomly generated and written to the unhashed key file. The unhashed file contains the responses that the system returns when challenged by another system.
 The corresponding hashed-key file data is displayed after executing the command with this option.
- `remote_name`
 Specifies the name of the remote system with which this one is being synchronized.
- `sequence seed hash`
 Can be used with the `-hashed` option and specifies data that is written in the hashed (public) key file:
sequence is a decimal number between 10 and 499.
seed is a 6 to 20 character string.
hash is a 16 digit hexadecimal number.

EXAMPLES

Example 1

In this example, the vopie key files are set up so the first connection between systems red and blue is not fully authenticated. After the connection, the key files are updated so full authentication is required. This is the easiest way to synchronize the key files but it leaves a small window of insecurity.

1. On system red:
 - a. Create a secret key file on red by executing:


```
vopie_util -local_name red -unhashed blue
```



The public key (hashed) file information for red is displayed:

```
red 0167 jp0167 0aa47eae2d86231d
```

This information can be ignored in this example.

- b.** Create a public key file on red that will match any secret key on blue:

```
vopie_util -local_name red -hashed blue
```

- 2.** On system blue:

- a.** Create a secret-key file on blue by executing:

```
vopie_util -local_name blue -unhashed red
```

The public key (hashed) file information for blue is displayed:

```
blue 0431 gw3251 0aa47eae2d86231d
```

This information can be ignored in this example.

- b.** Create a public key file on blue that will match any secret key on red by executing:

```
vopie_util -local_name blue -hashed red
```

Example 2

In this example, the vopie key files on systems green and yellow are synchronized. Full authentication is required immediately. This is a more secure method than in example 1.

- 1.** On system green, create a secret key file on green by executing:

```
vopie_util -local_name green -unhashed yellow
```

The public key (hashed) file information for green is displayed:

```
green 0209 fz9365 f852019bde05e92f
```

yellow uses this key when it issues challenges.

- 2.** On system yellow:

- a.** Create a public key file on yellow that matches the secret key file on green by executing the following (all on one line):

```
vopie_util -local_name yellow -hashed green 0209 fz9365  
f852019bde05e92f
```

- b.** Create a secret key file on yellow by executing:

```
vopie_util -local_name yellow -unhashed green
```

The public key (hashed) file information for yellow is displayed:

```
yellow 0468 yq0860 82723984b43bf474
```

green uses this key when it issues challenges.

3. On system green, create a public key file on green that matches the secret key file on yellow by executing the following (all on one line):

```
vopie_util -local_name green -hashed yellow 0468 yq0860  
82723984b43bf474
```

SEE ALSO

bpauthsync(1M), vopied(1M)



xbp(1)

NAME

xbp - start the X Windows based interface for NetBackup users

SYNOPSIS

```
/usr/opensv/netbackup/bin/xbp [-r] [-ra] [-rr]
[-nl][-browselimit files] [X options]
```

DESCRIPTION

The `xbp` command starts a graphical user interface that lets users archive, back up, and restore files, directories, or raw partitions from their client workstations. You can use `xbp` only from an X terminal or X server that is compatible with MIT release X11.R4 (or later) of the X Window system.

The `xbp` interface follows OSF/Motif conventions. If you are unfamiliar with these conventions, see the *OSF/Motif User's Guide*, authored by the Open Software Foundation and published by Prentice-Hall, Inc., ISBN 0-130640509-6.

The *NetBackup User's Guide - UNIX* and the `xbp` online help provide detailed operating instructions.

OPTIONS

`xbp` has separate modes for backups, archives, and restores. The backup and archive modes display the file system. By default, `xbp` starts in filesystem mode. The following options allow you to directly control the startup mode:

- `-r` Start with display of backups for possible restore.
- `-ra` Start with display of archives for possible restore.
- `-rr` Start with display of raw-partition backups for possible restore.
- `-nl` Specifies that `xbp` does not resolve links during the search. The default is to resolve links.
- `-browselimit files`
Specifies the limit for implicit searching.
When switching to restore mode, if the number of files and directories that were backed up during the specified date range is large (10000 by default), `xbp` pops up a warning dialog saying that searching is delayed until the user explicitly selects Update Display from the Edit menu.
By using the `-browselimit` parameter when invoking `xbp`, a user can increase this limit beyond 10000 files.



Also, `xbp` supports the standard command-line options for X programs. One of these is the `-d` option, which forces the name of the X terminal or server. Most users already have their `DISPLAY` environment variable defined and can routinely ignore the `-d` option.

Other useful X options are:

- `-bg color` Specifies the color to use for the background of the window. The default is `white`.
- `-fg color` Specifies the color to use for displaying text. The default is `black`.
- `-font` Allows you to enlarge text for visibility. It is best to use fixed-pitch fonts because `xbp` formats some text into columns. These columns can appear uneven with proportional fonts.
- `-geometry` Allows you to control the initial size and position of the `xbp` window.
- `-title` Controls the window manager title bar and is useful if you run several instances of `xbp` at once.

FILES

`/usr/opensv/netbackup/help/xbp/*`

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

SEE ALSO

`bp(1)`, `bparchive(1)`, `bpbackup(1)`, `bplist(1)`, `bprestore(1)`



xbpadmin(1M)

NAME

xbpadmin - start the X Windows based NetBackup interface for administrators

SYNOPSIS

```
/usr/opensv/netbackup/bin/xbpadmin [-notmwm] [X options]
```

DESCRIPTION

The `xbpadmin` utility provides a graphical user interface that the administrator can use to configure NetBackup and monitor its operations. This utility requires root-user privileges and can be run from any X terminal or X server that is compatible with MIT release 11.R5 (or later) of the X Window system.

The `xbpadmin` interface follows OSF/Motif conventions. If you are unfamiliar with these conventions, see the *OSF/Motif User's Guide*, authored by the Open Software Foundation and published by Prentice-Hall, Inc., ISBN 0-130640509-6.

The *NetBackup System Administrator's Guide - UNIX* and the `xbpadmin` online help provide detailed operating instructions.

OPTIONS

`-notmwm` Specifies that you are using a window manager other than `mwm`. This option eliminates display problems that can appear when using with a window manager such as `twm`.

You can also include this option in an environmental variable so you do not have use it each time you start `xbpadmin`. The following is an example entry for a `.login` file on a SunOS system:

```
setenv XBPADMIN_OPTIONS "-notmwm"
```

`-load_max_bytes`

Specifies the size in bytes at which the NetBackup truncates report listings. If a log is too large for viewing, a message similar to the following appears at the top of the report:

```
"The text has been truncated to 100000 bytes. The actual size is 207513 bytes. /tmp/xbpadmin.9488.r2.out."
```

This message means that `xbpadmin` has saved a copy in a temporary file. In this example, the file name is `/tmp/xbpadmin.9488.r2.out`.

To change the size at which the NetBackup truncates reports use the `-load_max_bytes` option when starting `xbpadm`. You can also include this option an environmental variable so you do not have use it each time you start `xbpadm`. The following is an example entry for a `.login` file on a SunOS system:

```
setenv XBPADM_OPTIONS "-load_max_bytes 2000000"
```

X options The most commonly used X option is `-d`. This option defines the name of the X terminal or server on which the interface is displayed. If you define your `DISPLAY` environment variable prior to starting `xbpadm`, you do not need the `-d` option.

FILES

`/usr/opensv/netbackup/help/xbpadm/*`

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

SEE ALSO

`bpadm(1M)`, `bprd(1M)`



xbpmon(1M)

NAME

xbpmon - start the NetBackup job monitor

SYNOPSIS

```
/usr/opensv/netbackup/bin/xbpmon [X options]
```

DESCRIPTION

The `xbpmon` utility allows the system administrator to check the progress of backups and also to exercise limited control over those operations. This utility requires root-user privileges and can be run from any X terminal or X server that is compatible with MIT release 11.R5 (or later) of the X Window system. You can also start `xbpmon` from `xbpadm`.

The `xbpmon` display follows OSF/Motif conventions. If you are unfamiliar with these conventions, see the *OSF/Motif User's Guide*, authored by the Open Software Foundation and published by Prentice-Hall, Inc., ISBN 0-130640509-6.

The *NetBackup System Administrator's Guide - UNIX* and the `xbpmon` online help provide detailed operating instructions.

OPTIONS

`X options` The most commonly used X option is `-d`. This option defines the name of the X terminal or server on which the interface is displayed. If you define the `DISPLAY` environment variable prior to starting `xbpmon`, you do not need the `-d` option.

FILES

```
/usr/opensv/netbackup/help/xbpmon/*
```

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

SEE ALSO

`xbpadm(1M)`



The NetBackup `bpadm` administrator utility has a character-based, menu-driven interface that you can use at any terminal (or terminal emulation window) for which you have a `termcap` or `terminfo` definition.

This appendix describes procedures for configuring and managing NetBackup using `bpadm`. The areas covered are as follows:

- ◆ Starting `bpadm`
- ◆ Defining and Managing Storage Units
- ◆ Defining and Managing Classes
- ◆ Defining NetBackup Global Attributes
- ◆ Installing NetBackup Software on All Trusting Client Hosts
- ◆ Displaying Reports
- ◆ Managing `bprd` (NetBackup Request Daemon)
- ◆ Redefining Retention Levels
- ◆ Performing Manual Backups
- ◆ Backing Up the NetBackup Databases (catalogs)

For further information on the subjects described in this chapter, see the corresponding topics in Chapter 3.



Starting bpadm

Note Use `bpadm` only on the master server and ensure that no other instances of `bpadm` are active when you are modifying the configuration. If you attempt to modify the configuration by using more than one instance or a combination of these utilities the results will be unpredictable.

Start the `bpadm` program by entering the following command (you must be a root user):

```
/usr/opensv/netbackup/bin/bpadm
```

When `bpadm` starts, the main menu appears on your screen.

```
NetBackup Server: bunny
```

```
NetBackup Administration
```

```
-----  
s) Storage Unit Management...  
c) Class Management...  
g) Global Configuration...  
r) Reports...  
m) Manual Backups...  
x) Special Actions...  
u) User Backup/Restore...  
v) Volume Administration...  
h) Help  
q) Quit
```

```
ENTER CHOICE:
```

You can probably perform any operation you want by responding to the prompts that `bpadm` provides. All menus have online help available. If you need more information, the topics in this chapter provide detailed instructions on common operations. You can abort many operations by pressing the Escape key.



Defining and Managing Storage Units

The system administrator's guide for NetBackup Media Manager (UNIX version) explains how to define storage devices and media to Media Manager. The procedures here explain how to define and manage them within NetBackup. The Storage Unit Management menu (see below), has options for defining and managing storage units. You display this menu by selecting s (Storage Unit Management) from the bpadm main menu.

Storage Unit Label: <none>

Storage Unit Host: <none>

Storage Unit Type: <none>

Output Destination: SCREEN

Storage Unit Management

a) Add Storage Unit...

m) Modify Storage Unit...

d) Delete Storage Unit

b) Browse Storage Units

l) List/Display Storage Units

o) Output Destination (SCREEN or FILE)

h) Help

q) Quit Menu

ENTER CHOICE:



Adding a Removable or Robotic Storage Unit

Before adding a Removable or Robotic type storage unit, you must configure the related devices and media within Media Manager. When that configuration is complete, you can add a storage unit so NetBackup can direct data to those devices and media.

The example below shows the dialog that occurs when adding an 8-mm tape stacker (user responses are in bold. <CR> means carriage return). To start the series of prompts in this example, enter a (Add Storage Unit) from the Storage Management menu.

Adding Storage Unit (<ESC> to abort)

Enter Storage Unit Label: TSD_1 **<CR>**

Enter Host Name: (bunny) bunny **<CR>**

Storage unit type:

- 1) Disk
- 2) Removable or Robotic Media

Enter Choice [1-2]: **2 <CR>**

Robot Type Selections:

- 0) NONE - Not Robotic
- 1) TS8 - Tape Stacker 8MM
- 2) TSD - Tape Stacker DLT

.
. .
.

Enter Robot Type [0-7]:**2 <CR>**

Enter this device's robot number: **2 <CR>**

Density Selections:

- 1) dlt - DLT Cartridge
- 2) dlt2 - Cartridge alternate

Enter Density [1-3]:**1 <CR>**

Determine the number of drives you wish to use for backups and archives. The number you use must be less than or equal to the number of drives installed.

Enter number of drives: **1 <CR>**

Use this storage unit only if required by a class or schedule?



(y/n) (n): <CR>

Maximum fragment size for backup images is configurable.

What maximum multiplexing factor should be used per drive?

(A value of 1 indicates to not do multiplexing)

Enter value [1-32]: (1)<CR>

Allowable values are in the range of 50 MB to unlimited.

Enter maximum fragment size (in MB) or 0 for unlimited:(0)<CR>

Add Storage Unit? (y/n): y

Adding Storage Unit...

Storage unit add completed

The procedure for adding a removable or robotic storage unit is as follow:

1. Enter a unique label for the storage unit (no spaces allowed in the label). This is the label you use to associate the unit with a class or schedule. Choose a label that is descriptive of the type of storage you are defining.
2. Enter the name of the host that is controlling the storage unit. This must correspond to the host to which the drives attach. The default host appears in parentheses. Either press Return to accept the default or enter a new name.
3. Enter the storage unit type. Choose 2 for Removable or Robotic Media. This brings up a list of choices for robot types.
4. Choose the storage unit's robot type. The resulting prompts vary depending on the type you choose. The following defines the actions for each choice:
 - ◆ Choosing 0 (NONE - Not Robotic) brings up the list of choices for the density of the unit.

Specify the density according to the value configured in Media Manager, then enter the number of drives of this density that you want to use. All nonrobotic drives of a given density must belong to the same storage unit. Specifying more than one drive can make it possible for the storage unit to handle more than one job at a time.

- ◆ Choosing a robot brings up a prompt for the device's robot number. This number must match the number you configured in Media Manager.

If you are prompted for density, set it according to the configuration in Media Manager. Then, specify how many of the robot's drives that you want to use for NetBackup operations. This number must be less than or equal to the number of drives that are installed in the robot.



5. Decide whether you want to use the storage unit only when a class or schedule specifies it, or to make it available for any schedule.
 - ◆ y reserves the unit for use only by classes or schedules that specify it.
 - ◆ n makes the storage unit available for any class or schedule.

The example accepts the default, which is n.

6. Specify the maximum image multiplexing (MPX) factor to use.

Image multiplexing sends concurrent, multiple backups from one or several clients to a single disk storage unit and multiplexes the images onto the media.

Enter a value from 1 to 32. A value of 1 (default) disables multiplexing by allowing only one backup job at a time to go to any given drive.

See “Multiplexing” on page 275 for more information on multiplexing.

7. Enter a value, in megabytes, for the maximum fragment size.

This is the largest size fragment that you want NetBackup to create when fragmenting images. A value of 0 specifies unlimited fragment size (no fragmentation). This parameter is most useful for disk type storage units as explained in the next topic, “Adding a Disk Type Storage Unit.”

Also see “Fragmented Backups” on page 643.

8. Enter y to confirm the addition or n to cancel.
9. Review the addition by using the l (List/Display Storage Units) option from the Storage Management menu. To change attributes, use the m (Modify Storage Unit) option or else delete the storage unit and then add it again.

If you are configuring NetBackup for the first time and are satisfied with your storage unit configuration, go to “Adding a Class” on page 565.

Adding a Disk Type Storage Unit

To add a disk type storage unit, enter a (Add Storage Unit) from the Storage Management menu and follow the prompts. Example prompts and responses are shown below (user responses are in bold. <CR> means carriage return).

Adding Storage Unit (<ESC> to abort)

Enter Storage Unit Label: **unixdisk_1** <CR>

Enter Host Name: (bunny) **bunny** <CR>

Storage unit type:

1) Disk

2) Removable or Robotic Media

Enter Choice [1-2]: **1** <CR>

Enter full path to image directory: **/bpimages**<CR>

Enter Number of Concurrent Jobs: (1) **2**<CR>

Use this storage unit only if required
by a class or schedule? (y/n) (n): <CR>

What maximum multiplexing factor should be used?
(A value of 1 indicates to not do multiplexing)

Enter value [1-32]: (1) **<CR>**

Maximum fragment size for backup images is configurable.

Allowable values are in the range of 20 MB to 2000 MB (2GB).

Enter maximum fragment size (in MB): (2000) **<CR>**

Add Storage Unit? (y/n): **y**

The procedure for adding a disk type storage unit is as follows:



1. Enter a unique label for the storage unit (no spaces allowed in the label). This is the label you use to associate the unit with a class or schedule. Choose a label that is descriptive of the type of storage you are defining. In the example above, we intend to configure a storage unit on magnetic disk that is managed by the UNIX operating system and use `unixdisk_1` for the label.
2. Enter the name of the server that is controlling the disk. This is the network name of the server as returned by the UNIX `hostname` command.
3. Enter the storage unit type. Choose 1 for Disk. This brings up a prompt to enter the path name.

Note For Disk type storage units, see “Adding Disk Type Storage Units” on page 32 for information on disk space requirements.

4. Specify the directory path for the backup and archive images. This can be anywhere on your disk that you have room.
5. Specify the number of concurrent jobs that you are going to allow. This number depends on your server’s ability to comfortably execute multiple backup processes (see “Limit jobs per class” on page 68).
6. Decide whether you want to use the storage unit only when a class or schedule specifies it, or to make it available for any class or schedule.
 - ◆ `y` reserves the unit for use only by classes or schedules that specify it.
 - ◆ `n` makes the storage unit available to any class or schedule.

The example accepts the default, which is `n`.

7. Specify the maximum image multiplexing (MPX) factor to use.

Image multiplexing sends concurrent, multiple backups from one or several clients to a single drive and multiplexes the images onto the media.

Enter a value from 1 to 32. A value of 1 (default) disables multiplexing by allowing only one backup job at a time to go to any given drive.

See “Multiplexing” on page 275 for more information on multiplexing.

8. Enter a value, in megabytes, for the maximum fragment size.

This is the largest size fragment that you want NetBackup to create when fragmenting images. The value can range from 20 to 2000 (default).

The Maximum Fragment Size setting is normally used to ensure that the backup images do not exceed the maximum size allowed by the filesystem. For example, on a filesystem managed by Storage Migrator, this breaks up the image so that Storage Migrator does not have to wait for the entire image to be on disk before starting its migration process. See “Fragmented Backups” on page 643 for more information.

9. Enter **y** to confirm the addition or **n** to cancel. This returns you to the Storage Unit Management menu.
10. Review the addition by using the **l** (List/Display Storage Units) option from the Storage Management menu. To change attributes, use the **m** (Modify Storage Unit) option or else delete the storage unit and add it again.

If you are configuring NetBackup for the first time and are satisfied with your storage unit configuration, go to “Adding a Class” on page 565.

Displaying and Changing Storage Unit Configurations

The Storage Unit Management menu has options for viewing the attributes for currently configured storage units or directing the list of attributes to a file. This menu also has options for modifying the configuration by either deleting storage units or changing their attributes.

The basic procedure for all these options is as follows:

1. Select the desired storage unit by pressing **b** (Browse Storage Units) on the Storage Unit management menu until the name of that storage unit appears on the Label line at the top of the screen. The next two lines show the host to which the storage unit connects and the type of storage unit.
2. Choose the desired option:
 - ◆ To modify the number of concurrent jobs or whether the storage unit is used only on demand, choose **m** (Modify Storage Unit) from the Storage Unit Management menu and follow the prompts (existing values are in parentheses). To modify other attributes, you must delete and then re-add the storage unit.
 - ◆ To delete a storage unit, choose **d** (Delete Storage Unit) from the Storage Unit Management menu. At the prompt, check to ensure that you are deleting the correct storage unit and enter **y** if you want to delete it. Deleting a storage unit from the NetBackup configuration does not prevent you from restoring files that are stored on that unit. A restore requires only that the same type of storage unit is available (in Media Manager for a removable or robotic type storage unit).
 - ◆ To view the attributes for the storage unit, choose **l** (List/Display Storage Units) from the Storage Unit Management menu. Use the controls at the bottom of the screen to page forward and backward or move up and down a line at a time.
 - ◆ To direct the list of attributes to a file, choose **o** (Output Destination) from the Storage Unit Management menu and enter the desired file path at the prompt. Then, choose **l** to write the list to the file.



Defining and Managing Classes

The procedures in this section explain how to define and manage NetBackup classes. To display the Class Management menu, select c (Class Management) from the bpadm main menu.

```
Class: <ALL>
Clients: <ALL>
Schedules: <ALL>
Output Destination: SCREEN
```

```
Class Management
```

```
-----
a) Add Class...
d) Delete Class
m) Modify Class Attributes...
s) Schedule Management...
c) Client List Management...
f) File List Management...
```

```
b) Browse Classes Forward
r) Browse Classes Reverse
e) Enter Class
l) List/Display Classes
o) Output Destination (SCREEN or FILE)
h) Help
q) Quit Menu
```

```
ENTER CHOICE:
```



Adding a Class

To add a class to the configuration, select the Add Class option from the Class Management menu.

1. Press a from the Class Management menu to start a series of prompts for adding a class (see example below). Some choices, such as Cross Mount Points, have default values in parentheses. Press Return to accept default values. (Example user responses are in bold. <CR> means carriage return.)

```
Adding Class (<ESC> to abort)
```

```
-----
```

```
Enter Unique Class Name: W2 <CR>
```

```
Use an existing class as a template; if yes, all
attributes and schedules will be duplicated: (y/n)?n
```

```
Class Type
```

```
-----
```

```
1) Standard
```

```
.
```

```
. (the actual menu will show more than is listed here)
```

```
.
```

```
Enter Choice: (1) <CR>
```

```
Active? (y/n) (y): <CR>
```

```
Collect True Image Recovery information
```

```
0 = No
```

```
1 = Yes
```

```
2 = Yes with move detection
```

```
Enter Choice [0-2]: (0) <CR>
```

```
Cross mount points? (y/n) (n): <CR>
```

```
Follow NFS mounts? (y/n) (n): <CR>
```

```
Client Compression? (y/n) (n): <CR>
```

```
Enter priority as compared to other classes: (0) <CR>
```



```
Block level incremental? (y/n) (n): <CR>
Enter maximum number of jobs per class (99) <CR>
Require images to be written to a specific storage unit? (y/n) (n): y
Enter Storage Unit label: ts8_1
Enter the volume pool images should be directed to:(NetBackup) <CR>
Associate a keyword with this class? (y/n) (n): <CR>
Allow multiple data streams? (y/n): n
Add class now? (y/n): y
```

2. Enter the name for the class. This name must be unique to the configuration (no spaces allowed in name).
3. Choose whether you want to use an existing class as a template. This is convenient if another class has many of the same attributes. You can subsequently make the necessary changes for the class you are adding. If you use another class for a template, NetBackup duplicates the following:
 - ◆ Class attributes
 - ◆ Files list
 - ◆ Client list
 - ◆ All schedules
4. Select the class type from the list.
5. Specify whether to activate the class. A class must be active for NetBackup to execute any of its schedules (automatic or user-directed). The example uses y, which sets the class to active.
6. Choose whether to collect True Image Recovery Information (see “True image restore information” on page 70). The example uses the default which is n (no).
7. Choose whether to follow NFS mounts.
 - ◆ y enables NetBackup to back up and archive NFS-mounted files and directories.
 - ◆ n prevents NetBackup from backing up or archiving NFS-mounted files and directories.The example uses the default, which is n.
8. Choose whether to cross mount points when doing backups and archives. The example uses the default, which is n.
9. Choose whether to compress the files that you archive or back up from that client. The example uses the default, which is n (no).

10. Enter the priority for this class relative to other classes. You can enter any positive integer. The class with the highest value has the highest priority. The default is 0.
11. Choose whether to use block level incremental backups. The example uses the default, which is n (no).
12. Enter the maximum number of jobs that this class can perform concurrently. The default is unlimited (entering 0 also sets it to unlimited). See “Limit jobs per class” on page 68 for more information.
13. Choose whether to specify a default storage unit for the class. The example specifies TS8_1. Therefore, NetBackup directs backups and archives for this class to TS8_1, except for schedules that specify a storage unit.

Choose whether to specify a default volume pool for the class. If you do not specify a volume pool for either the class or the schedule, the *NetBackup* volume pool is used.
14. Choose whether to use a keyword phrase (see “Keyword phrase (optional)” on page 69). The example uses the default, which is n (no).
15. Enter y to add the class or n to cancel.
16. Review the addition by using the l (List/Display Classes) option from the Class Management menu. To change attributes, use the m (Modify Class Attributes) option.

If you are configuring NetBackup for the first time and are satisfied with your class configuration, go to “Adding Clients to a Class” on page 569.

Displaying and Changing Class Configurations

The Class Management menu (see “Defining and Managing Classes” on page 564) has options for viewing the attributes for currently configured classes or directing the list of attributes to a file. This menu also has options for modifying the configuration by either deleting classes or changing their attributes.

The basic procedure for all these options is as follows:

1. Select the desired class by browsing with the b and r options on the Class Management menu until the name of that class appears on the Class line at the top of the screen. You can also use the e option to enter the class name.
2. Choose the desired option:
 - ◆ To modify the attributes, choose m (Modify Class Attributes). At the prompt, check the top line on the screen to ensure you are modifying the correct class. Enter new values at the prompts or simply press Return to accept the existing values (shown in parentheses).



- ◆ To delete a class, choose d (Delete Class). At the prompt, check to ensure that you are deleting the correct class and enter y if you want to delete it. Deleting a class from the NetBackup configuration does not prevent you from restoring files that were backed up or archived by clients in that class.
- ◆ To list the attributes for the class, choose l (List/Display Classes). Use the controls at the bottom of the screen to page forward and backward or move up and down a line at a time.
- ◆ To direct the list of attributes to a file, choose o (Output Destination) and enter the desired file path at the prompt. Then, choose l to write the list to the file.



Defining and Managing the Client List for a Class

The procedures in this section explain how to define and manage the list of clients for each class.

Adding Clients to a Class

The following explains how to add clients to a class.

1. Select the class for which you are going to define the client list. You accomplish this by pressing **b** (Browse Classes) on the Class Management menu (see “Defining and Managing Classes” on page 564) until the name of the desired class appears on the Class line at the top of the screen.
2. Press **c** to bring up the Client List Management menu. This menu has options for managing your client list. The class you selected in the previous step appears on the Class line at the top of the screen. The example below shows class W2.

```
Class: W2
Clients: <none>
Schedules: <none>
Output Destination: SCREEN

Client List Management
-----
a) Add Clients
d) Delete Clients

l) List/Display Class
o) Output Destination (SCREEN or FILE)
h) Help
q) Quit Menu

Enter Choice:
```



3. Choose a from the Client List Management menu. This brings up the list of client types currently installed at your site. In the example below, user responses are in bold. <CR> means carriage return.

```

Class: W2
  Adding Clients (<ESC> to abort)
  -----
  1) C910_920, IRIX
  2) HP9000-700, HP-UX
  3) HP9000-800, HP-UX
  4) RS6000, AIX
  5) Solaris Solaris2.5
  6) Solaris Solaris2.6

Enter Selection (or 'q' to end): 4 <CR>
Enter clients of RS6000, AIX type: (empty line to end)
  Enter Client Name: mars <CR>
  Enter Client Name: <CR>
  Adding clients to class W2
  mars
  Install client software (y/n) n
  [Menu of choices reappears]

Enter Selection (or 'q' to quit): 1 <CR>
Enter clients of C910_920, IRIX type: (empty line to end)
  Enter Client Name: jupiter <CR>
  Enter Client Name: <CR>
  Adding clients to class W2
  jupiter
  Install client software (y/n) n
  [Menu of choices reappears]

Enter Selection (or 'q' to quit): q<CR>
  
```

4. Enter the number corresponding to the type of client you are adding.
5. Specify the names of the clients of this type (one per line). When choosing client host names, always observe the following rules:



- ◆ Use the same name if you put the client in multiple classes.
- ◆ Use a name by which the server knows the client. This name should be 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`.

When you finish naming the clients, leave a blank line and press Return. This brings up a message informing you that the client is being added. You are then prompted as to whether you want to install client software.

- ◆ If you added *trusting* clients and want to install software now, choose `y` to have `bpadm` immediately push client software from the server to the client. A *trusting client* is one that does have an `/.rhosts` file with an entry for the NetBackup server. This software installation occurs after the clients are added to the class. If the software installation fails on any of the clients, NetBackup notifies you, but still keeps the client in the class. Note that client software installation can take a minute or more per client.
- ◆ If you added *secure* clients, you should enter `n` and then install them later as explained under “Installing Software on Secure UNIX Clients” on page 90. A *secure* client is one that does not have an entry for the NetBackup server in its `/.rhosts` file.
- ◆ If you added *trusting* clients but want to install software later, enter `n` at the installing software prompt. You can install the software later by using the Install All Clients option from the Special Actions menu (see “Installing NetBackup Software on All Trusting Client Hosts” on page 583).

If you enter `n` at the prompt or if software installation is complete, `bpadm` returns you to the list of choices so you can add another type of client.

Note The prompt appears only if client software was loaded on the master server during NetBackup installation and is therefore available for installation on clients.

6. Repeat step 4 and step 5 until your list is complete, then press `q` to return to the Client List Management menu.
7. Review the addition by using the `l` (List/Display Class) option from the Client List Management menu.

If you are configuring NetBackup for the first time and are satisfied with your client list for this class, go to “Adding to a File List” on page 572.



Displaying Client Lists and Deleting Clients from a Class

The Client List Management menu (see “Adding Clients to a Class” on page 569) has options for viewing a client list for a currently configured class or directing the list to a file. This menu also has an option for deleting clients from a class.

The basic procedure for these options is as follows:

1. Select the desired class by pressing **b** (Browse Classes) on the Class Management menu (see “Defining and Managing Classes” on page 564) until the name of that class appears on the Class line at the top of the screen.
2. Press **c** to bring up the Client List Management menu (see “Adding Clients to a Class” on page 569). The class you selected in the previous step appears at the top of the screen.
3. Choose the desired option:
 - ◆ To delete clients, choose **d** (Delete Clients) from the Client List Management menu. Check to ensure that you are deleting clients from the correct class and follow the prompts. Deleting a client does not delete any backups or archives that belong to the client.
 - ◆ To list the attributes for the class (including the clients), choose **l** (List/Display Class) from the Client List Management menu (or the Class Management menu). Use the controls at the bottom of the screen to page forward and backward or move up and down a line at a time.
 - ◆ To direct the list of class attributes (including the clients) to a file, choose **o** (Output Destination) from the Client List Management menu. Enter the desired file path at the prompt, then choose **l** (List/Display Class).

Defining and Managing the File List for a Class

The file list for a class applies to all full and incremental backups for the clients in that class. The procedures in this section explain how to define and manage the list of files.

Adding to a File List

1. Select the class for which you are going to define the file list. You accomplish this by pressing **b** (Browse Classes) on the Class Management menu (see “Defining and Managing Classes” on page 564) until the name of the desired class appears on the Class line at the top of the screen (W2 for our example).
2. Press **f** to bring up the File List Management menu. This menu has options for managing your client list. The class you selected in the previous step appears on the Class line at the top of the screen. The example below is for class W2.



```

Class: W2
Clients: mars jupiter ...
Schedules: <none>
Output Destination: SCREEN

```

```
File List Management
```

```
-----
```

```

a) Add Files
d) Delete Files

l) List/Display
o) Output Destination (SCREEN or FILE)
h) Help
q) Quit Menu

```

```
ENTER CHOICE:
```

- 3.** Choose a (Add Files) from the File List Management menu. The following is an example of the screen that appears (example user responses are in bold. <CR> means carriage return).

```

Class: W2
Client(s): mars jupiter ...
Schedule(s): <none>
File Paths: <none>

```

```

Adding File Paths (<ESC> to Abort, Blank line to end)
(NOTE: Spaces, ` ` , are significant in path names)

```

```
-----
```

```

Enter File Path: /usr <CR>
Enter File Path: /home <CR>
Enter File Path: /var <CR>
Enter File Path: <CR>

```

```
Adding file paths . . .
```



```
getting class list . . .
```

4. Enter the file paths at the prompts. You can enter one path per line and they must be full (absolute) file paths. When you finish, leave a blank line and press Return. This returns you to the File List Management menu (pressing Escape aborts the operation without altering the configuration).

You can use meta-characters or wildcard characters when specifying file lists. See “To Add Files For Automatic Backups” on page 90 for details.

To back up a raw partition, enter the path to the block or character device file. For example:

```
/dev/rdisk/isc0d2s6
```

The character device is preferred as it generally is faster than the block device. See also, “UNIX Raw Partitions” on page 100.

For some database extension class types, such as Oracle, you specify the scripts that control the back up (see the NetBackup guide for the respective product for more information).

5. Review the additions by using the l (List/Display) option from the File List Management menu. To make changes, use the a (Add Files) or d (Delete Files) options.

If you are configuring NetBackup for the first time and are satisfied with your file list, go to “Adding a Schedule” on page 576.



Displaying and Changing a File List

The File List Management menu has options for viewing the file list for currently configured classes or directing the list to a file. This menu also has an option for deleting files from a class.

The basic procedure for these options is as follows:

1. Select the desired class by pressing **b** (Browse Classes) on the Class Management menu (see “Defining and Managing Classes” on page 564) until the name of that class appears on the Class line at the top of the screen.
2. Press **f** to bring up the File List Management menu (see “Adding to a File List” on page 572). The class you selected in the previous step appears at the top of the screen.
3. Choose the desired option:
 - ◆ To delete files, choose **d** (Delete Files) from the File List Management menu. Check to ensure that you are deleting files from the correct class and follow the prompts. Deleting a file from the file list does not prevent you from recovering any backups or archives for that file.
 - ◆ To list the attributes for the class (including the files), choose **l** (List/Display) from the File List Management menu. Use the controls at the bottom of the screen to page forward and backward or move up and down a line at a time.
 - ◆ To direct the list of class attributes (including the file list) to a file, choose **o** (Output Destination) from the File List Management menu. Enter the desired file path at the prompt, then choose **l** (List/Display) to write the attributes to the file.

Defining and Managing Schedules for a Class

Each class must have a set of schedules to control its backup and archive operations. The procedures in this section explain how to define and manage those schedules with `bpadm`.



Adding a Schedule

The following describes how to add either an automatic or user-directed schedule. The process is the same for both, except where noted.

1. Select the class where you are adding the schedule. You accomplish this by pressing **b** (Browse Classes) option on the Class Management menu (see “Defining and Managing Classes” on page 564) until the name of the desired class appears on the Class line at the top of the screen.
2. Press **s** to bring up the Schedule Management menu. This menu has options for managing schedules. The class you selected in the previous step appears on the Class line at the top of the screen. The example below shows class W2.

```
Class: W2
Schedule: mars jupiter ...
Clients: <none>
Output Destination: SCREEN
```

Schedule Management

```
-----
a) Add Schedule...
d) Delete Schedule
m) Modify Schedule...

b) Browse Schedules
l) List/Display Schedule
o) Output Destination (SCREEN or FILE)
h) Help
q) Quit Menu
```



3. Choose a (Add Schedule) from the Schedule Management menu. The example below shows the addition of an automatic schedule for class W2 (example user responses are in bold. <CR> means carriage return). All choices, except Schedule Label, have default values in parentheses. Press Return to accept default values.

```

Class:                W2
Add Schedule (<ESC> to abort)
-----
Enter Schedule Label: W2_daily_differential <CR>
Schedule Type
  0) Full
  1) Differential Incremental Backup
  2) Cumulative Incremental Backup
  3) User Backup
  4) User Archive
Enter Choice [0-3]:(0)1 <CR>
Backup Frequency -- in hours, days, or weeks? (h/d/w):(d)<CR>
Enter Frequency (in days):(7) 1 <CR>
Retention levels:
  0) one week
  1) two weeks
  2) three weeks
  .
  .
  .
  8) one year
  9) indefinite
Enter Retention Level [0-9]:(3)0 <CR>
Require images to be written to a specific storage unit?(y/n) (n):<CR>
Do you want to override the class volume pool? (y/n)(n):<CR>
Use multiplexing if able? (y/n) (n):y
What maximum multiplexing factor should be used?(A value of 1
indicates to not do multiplexing)
Enter value [1-32]: (1)2

```



Should the backup window be the same every day of the week?(y/n)(y):**n**
<CR>

Enter Daily windows:

start time, then duration in hours

```
Sunday      (20:00:00 0): 22 0
Monday      (20:00:00 6): 22 8
Tuesday     (20:00:00 6): 22 8
Wednesday   (20:00:00 6): 22 8
Thursday    (20:00:00 6): 22 8
Friday      (20:00:00 6): 22 8
Saturday    (20:00:00 0): 22 0
```

Add schedule W2_daily_incr to class W2 now? (y/n)(y): **y**

4. Enter a unique label for the schedule (no spaces allowed in the label). This name appears on screens and messages from NetBackup, so choose a name with a meaning you can remember.
5. Choose the schedule type. Choices 0, 1, and 2 select automatically scheduled backups. Choices 3 and 4 are user directed. The example specifies 1 for Differential Incremental backup.

If the class type is for database backups, such as an Oracle-Obackup class, you see a set of choices similar to the following:

```
Schedule Type
1. Scheduled Obackup script
2. Obackup initiated script
```

Where choice 1 is for an automatically scheduled database backup that is started by the NetBackup scheduler. Choice 2 is started by the obackup process on the client. See the installation guide for the respective product for more information on these schedule types.

6. Specify the units for the backup frequency you will specify in step 7 (does not apply to user-directed backups and archives). In the example, we press Return to select the default, which is days.
7. Specify the backup frequency (does not apply to user-directed backups and archives). This is the time that should occur between successful backups and is expressed in terms of the units selected in step 6. The example selects 1 day.
8. Specify the retention level for the backups or archives that this schedule creates (also see "Retention" on page 126). The example retains its daily incrementals for 1 week.



If you are configuring a schedule for an Oracle backup class, retention level 9 (indefinite) is automatically selected. Oracle controls image expiration for these backups.

9. Choose whether you want to direct the backup images for this schedule to a specific storage unit.
 - ◆ Entering *y* brings up a prompt for the name of the storage unit.
 - ◆ Entering *n* accepts the storage unit as specified at the class level.

If you did not specify one at the class level, NetBackup uses the next storage unit available.

10. Choose whether to specify a volume pool for the schedule.
 - ◆ If you enter a volume pool name, this choice overrides the class level volume pool.
 - ◆ If you do not enter a volume pool name, NetBackup uses the volume pool specified at the class level. If you do not specify one at either the schedule or class level, NetBackup uses a default of *NetBackup*.
11. Specify whether you want to use multiplexing.

Multiplexing sends concurrent, multiple backups from one or several clients to a single drive and multiplexes the images onto the media.

If you answer *y* to this prompt, then you are asked to specify the multiplexing factor. The multiplexing factor is the maximum number of jobs from this schedule that you want to multiplex on any one drive. The number can range from 1 to 32, where 1 specifies no multiplexing.

See “Multiplexing” on page 275 for more information on multiplexing.

12. Specify the start times and durations for the backup window:
 - ◆ Entering *y* (or accepting the default, which is *y*) specifies that the backup window opens on each day of the week. NetBackup can attempt backups each day and during the same time frame. The prompts ask you to define when the window opens and how long it remains open each day.
 - ◆ Entering *n* brings up prompts for specifying a different window for each day of the week. Specify time in terms of the 24-hour clock. For example, 00:00:00 is 12 am, 12:00:00 is 12 pm, and 23:30:00 is 11:30 pm. The duration is in hours.

You can enter the time in *hours*, *hours:minutes*, or *hours:minutes:seconds*. For example, if you enter just the hours or hours and minutes, `bpadm` completes the entry. Entering 22 results in a time of 22:00:00 and entering 22:30 results in a time of 22:30:00.



When completing the daily windows, remember to leave a blank space between the hours and the duration. Entering 22 8 results in a time of 22:00:00 and duration of 8 hours. Entering 2 8 results in a time of 02:00:00 and a duration of 8 hours. Entering 0 for the duration results in no backup window. Entering 0 for the time results in a start time of 00:00:00.

13. Enter **y** to add the schedule to this class or **n** to cancel.
14. Review the addition by using the **l** (List/Display Schedule) option from the Schedule Management menu. To change any attributes use the **m** (Modify Schedule) option.

If you are configuring NetBackup for the first time and are satisfied with the schedules for this class, return to “Adding a Class” on page 565 and repeat the procedures in this chapter as necessary for the next class.

Displaying and Modifying a Schedule

The Schedule Management menu (see “Adding a Schedule” on page 576) has options for modifying the list of schedules for currently configured classes or directing the list to a file. This menu also has options for modifying schedules or deleting them from a class.

The basic procedure for these options is as follows:

1. Select the desired class by pressing **b** (Browse Classes) on the Class Management menu until the name of that class appears on the Class line at the top of the screen.
2. Press **s** to bring up the Schedule Management menu (see “Adding a Schedule” on page 576). The class you selected in the previous step appears at the top of the screen.
3. Select the desired schedule by pressing **b** (Browse Schedules) on the Schedule Management menu until the name of the desired schedule appears on the Schedule line at the top of the screen.
4. Choose the desired option:
 - ◆ To modify a schedule, choose **m** (Modify Schedule) from the Schedule Management menu. Check the top line on the screen to ensure that you are modifying the correct schedule. Enter new values at the prompts or simply press Return to accept the existing values (shown in parentheses).
 - ◆ To delete a schedule, choose **d** (Delete Schedule) from the Schedule Management menu. At the prompt, check to ensure that you are deleting the desired schedule and enter **y** if you want to delete it.
 - ◆ To list the attributes for the schedule selected in step 3, choose **l** (List/Display Schedule) from the Schedule Management menu. Use the controls at the bottom of the screen to page forward and backward or move up and down a line at a time.

- ◆ To direct the list of class attributes for the schedule selected in step 3 to a file, choose o (Output Destination) from the Schedule Management menu. Enter the desired file path at the prompt, then choose l (List/Display Schedule) to write the attributes to the file.

Defining NetBackup Global Attributes

The global attributes define aspects of NetBackup operation not defined elsewhere in the configuration. To list or modify them, select g (Global Attributes) from the NetBackup Administration menu. The screen that appears will be similar to the one below (in the example, user responses are in bold. <CR> means carriage return).

```

Keep Logs: 4days
Admin Mail Address: lxx@freddie.bdev.null.com,txz@mars...
Wakeup Interval:60 minutes
Backup Tries:2 times in 12 hours
Output Destination:SCREEN
Global Configuration
-----
m) Modify Configuration Parameters...
l) List/Display All Configuration Parameters
o) Output Destination (SCREEN or FILE)
h) Help
q) Quit Menu
ENTER CHOICE: m
Modify Configuration
-----
m) Mail Address: lfk@freddie.bdev.null.com,rjh>window...
w) Wakeup Interval: 60 minutes
j) Max Jobs/Client: 2
b) Backup Tries: 2 times in 12 hours
k) Keep Logs: 4 days
i) Keep TIR Info: 1 days
t) Media Mount Timeout: 0 minutes
h) Display Reports: 12 hour(s) ago
c) Compress Image Database Files: older than 7 days

```



d) Notify Request Daemon of Changes

q) Quit Menu

ENTER CHOICE:**k**

Enter Days to Keep Logs:(4)3 <CR>

Changing global attribute....

- ◆ To list the current values, select l (List/Display) from the Global Attributes menu.
- ◆ To modify the values, select m from the Global Attributes menu. The example above changes the value of Keep Logs from 4 days to 3.
- ◆ To have the request daemon (bprd) reread the configuration files, choose d Notify Request Daemon of Changes from the Modify Configuration menu.

Table 23 defines each of the NetBackup global attributes.

Table 23. NetBackup Global Attributes

Mail Address	Address to which NetBackup sends notifications on results of failed automatic backups, administrator- directed manual backup operations, and automatic database backups. This address should be that of the administrator. Default is none.
Wakeup Interval	Length of time, in minutes, that the scheduler waits before checking if any backups are scheduled to begin. Long wake up intervals can cause the scheduler to miss too much of a backup window to complete its backups. Default is 10 minutes.
Maximum Jobs/Client	Maximum number of jobs that NetBackup clients can perform concurrently. The default is 1.
Backup Tries	Number of times that NetBackup tries a backup job for a client/class/schedule combination during the specified time period. The time period must be greater than 0. The number of tries should also be greater than 0. Specifying 0 for number of tries is legal but stops all scheduled backups. The default is 2 tries in 12 hours. Note that this attribute does not apply to user-directed backups and archives.
Keep Logs	Length of time, in days, that the NetBackup server keeps its Error Database, Job database, and Activity logs. NetBackup derives its Backup Status, Problems, All Log Entries, and Media Log Entries reports from the error database so Keep Logs limits the time period that these reports can cover. The default is 28 days.
Keep TIR Info	Length of time to keep true image recovery information for those classes that use it.



Table 23. NetBackup Global Attributes

Media Mount Timeout	Length of time, in minutes, that NetBackup waits for the requested media to be mounted. This timeout will eliminate excessive waits for operations with nonrobotic devices (operator must mount media) or for media that is outside the robot or off site. The default is 0 (unlimited).
Display Reports	Default time period that NetBackup uses as it searches for information to put into a report. For example, a setting of 8 hours provides a report covering the previous 8 hour period. The minimum setting is 1 hour. The default is 24 hours.
Compress Image Database Files	Number of days that must elapse (since the image was created) before NetBackup compresses its image database files (also called image catalog files). The image database has information about client backups and archives. A value of 0 means that no compression should be done.
Preprocess Interval	Minimum time that can elapse between client queries to discover new paths if NetBackup is using auto-discover streaming mode (see “File-List Directives for Multiple Data Streams” on page 112). The default is 4 hours (a value of 0 sets it to the default).
Max Drives this Master	Maximum number of drives (on this master and its media servers) that the master server should consider available when scheduling backups.

Installing NetBackup Software on All Trusting Client Hosts

You can install software on *trusting* clients when you add them to a class or on all of them at once, by using the *c* (Install All Clients) option from the Special Actions menu. A *trusting* client is one that has an *.rhosts* file with an entry for the NetBackup server. The *c* option pushes client software from the server to the client.

1. Select *x* (Special Actions) from the main menu to bring up the Special Actions menu.

Special Actions

- c) Install All Clients...
- d) Backup Databases...
- r) View and Change Retention Levels
- i) Initiate Request Daemon
- t) Terminate Request Daemon
- h) Help



q) Quit Menu

ENTER CHOICE:

2. Choose c (Install All Clients) from the Special Actions menu. This starts the installation of software on all clients. Note that client software installation takes a minute or more per client.

Displaying Reports

The Reports menu has options for viewing problem or status reports from one or more NetBackup servers or clients. To reach this menu, choose r from the NetBackup Administration menu.

```
Server Name:bunny
```

```
Client Name:ALL
```

```
Start Date:03/18/93 08:38:59
```

```
End Date:03/19/93 23:59:59
```

```
Output Destination:SCREEN
```

```
Reports
```

```
-----
```

```
b) Backup Status
```

```
l) List Client Backups
```

```
p) Problems
```

```
a) All Log Entries
```

```
m) Media...
```

```
d) Change Dates
```

```
c) Change Client
```

```
s) Change Server
```

```
o) Output Destination (SCREEN or FILE)
```

```
h) Help
```

```
q) Quit Menu
```

The following procedure explains how to use the reports menu.



1. Select the server that has the reports you want to view by using the **s** (Change Server) option.

The name you enter with this option appears on the Server Name line at the top of the menu. Entering ALL (default) provides a report for all servers (except when using Media reports option **m**).

2. Select the client by entering **c** (Change Client).

The name you enter with this option appears on the Client Name line at the top of the menu. Entering ALL gives you reports for all clients and the selected server.

3. Specify the time period that you want the reports to cover by pressing **d** (Change Dates) and following the prompts.

The dates you enter appear on the Start Date and End Date lines at the top of the screen. The resulting report shows information ranging from the start date to the end date.

NetBackup derives the Backup Status, Problems, All Log Entries, and Media Log Entries reports from the error database; therefore, the Keep Logs attribute sets the maximum time period that these reports can cover (see “Defining NetBackup Global Attributes” on page 581). The maximum time limit for other Media reports and the List Client Backups report depends on the retention period for the associated backup images.

4. Choose the desired report option from the menu.

The following topics provide brief descriptions of each report. See “Report Descriptions” on page 177 for detailed information about each report.

Backup Status Report

The **b** (Backup Status) option from the Reports menu provides status and error information on backups completed successfully or failed within the specified time period.

Client Backups Report

The **l** (List Client Backups) option, from the Reports menu, shows detailed information on successful backups completed within the specified time period.

Problems Report

The **p** (Problems) option from the Reports menu lists the problems that the server has logged during the specified time period. The information is a subset of the information you get from the All Log Entries option.



All Log Entries Report

The a (All Log Entries) option from the Reports menu lists all the log entries for the specified time period.

Media Reports

Choosing m (Media) from the Reports menu brings up the Media Reports menu (see below). Prior to executing a media report option, you can select the servers and also clients (if necessary) for which you want the report. For Media Log entries, you also can select the range of dates that you want the report to cover.

```

Server:  ALL
Client:  ALL
Media ID/Path:  ALL
Start Date:  12/11/95  07:23:35
End Date:    12/12/95  23:59:59
Output Destination:  SCREEN
    
```

```

Media Reports          Change Parameters
-----
l)  Media List        s)  Change Server
u)  Media Summary    c)  Change Client
m)  Media Contents   p)  Change Media ID/Path
i)  Images on Media  d)  Change Dates
e)  Media Log Entries o)  Output Destination(SCREEN or FILE)
w)  Media Written

h)  Help
q)  Quit Menu
ENTER CHOICE:
    
```

The following is a basic procedure for choosing each type of report:



1. Select the server that you want to show reports for by using the `s` (Change Server Name) option. The name you enter with this option appears on the Server line at the top of the menu. Except for choice `m`, entering `ALL` provides a report for all servers. For `m`, entering `ALL` selects the server on which you are running `bpadm`.

When you change the server, the server initiating the request (one on which you are running `bpadm`) must be able to access the server you select. Otherwise, you get an access is not allowed message. Access to a server is controlled by the `SERVER` entry in its `bp.conf` file (see “NetBackup Configuration Options” on page 325).

2. For Images on Media reports, select the client by using the `c` (Change Client Name) option. The name you enter with this option appears on the Client line at the top of the menu. Entering `ALL` gives you reports for all clients and the selected server.
3. For Media Log entries, specify the time period that you want the reports to cover by pressing `d` (Change Dates) and following the prompts. The dates you enter appear on the Start Date and End Date lines at the top of the screen. The resulting report shows information ranging from the start date to the end date.
4. Choose the desired report option from the menu. The following topics explain each report.

Media List from NetBackup Media Catalog

The `l` (Media List) option shows either a single or all media IDs in the NetBackup media catalog. This option does not apply to Disk type storage units, nor does it show media assigned for the purpose of backing up NetBackup catalogs. You can get information for images on those storage units by using the `i` (Images on Media) option.

Media Summary

The `u` (Media Summary) option lists all media in the specified server’s catalog, according to whether it is active. The report also shows the expiration date for the media and shows the number of media that are at each retention level.

Media Contents

The `m` (Media Contents) option lists the contents of a single media ID. You must select only one media ID to use this option. The resulting report shows the contents of the media header and backup headers that are recorded on the media. You cannot use this option for disk type storage units.

The media contents report is useful for determining the backup IDs that are on a specific media ID by reading them from the media itself rather than the catalog. Because it requires a media mount, this option involves a greater delay for tape than for optical disk.



Images on Media

The **i** (Images on Media) option can list the contents of media as recorded in the NetBackup catalog. You can use this option to list the contents of any type of media (including disk). You can select by client, media ID, or path.

Media Log Entries

The **e** (Media Log Entries) option lists the media errors or informational messages relating to media that are recorded in the NetBackup error database. You can use the **d** (Change Dates) option to select errors by date.

Media Written

The **w** (Media Written) option lists media in the specified server's catalog that has been used for backups within the specified time period. This report does not show media used for image duplication if the original image was created prior to the specified time period.



Managing bprd (NetBackup Request Daemon)

The x (Special) option from the main menu displays the Special Actions menu, which provides choices for managing bprd (NetBackup request daemon). This daemon starts the scheduler and the NetBackup database daemon (bpdbm), in addition to controlling other functions within NetBackup.

Special Actions

```

-----
c)  Install Client Software...
b)  Backup Databases...
r)  View and Display Retention Levels

i)  Initiate Request Daemon
t)  Terminate Request Daemon

h)  Help
q)  Quit Menu

```

ENTER CHOICE:

The following explains the choices related to managing the request daemon.

Initiate Request Daemon

Starts bprd, if it was not running when you started bpadm. Normally, this should not be necessary if bprd is started at boot time. You do need this option, however, when you stop the daemon to alter the configuration (see Terminate Request Daemon, below). Starting bprd also starts bpdbm if bpdbm is not already executing.

Terminate Request Daemon

Stops bprd. If the daemon has started any activities, they are allowed to complete. With bprd stopped, NetBackup is unable to perform any backup, archive, or restore operations.

You should always stop the NetBackup request daemon (bprd) before making any changes to classes or schedules. This eliminates the possibility of a previously scheduled backup or archive operation invoking the scheduler and reading the configuration while you are making changes.

Use the `/usr/opensv/netbackup/bin/bpps` script to verify that bprd has terminated. Note that terminating bprd does not terminate bpdbm. You use bpdbm `-terminate` to stop bpdbm (see bpdbm(1M)).



Redefining Retention Levels

The **r** option on the Special Actions menu allows you to change the retention period associated with any retention level. Choose **x** (Special) from the main menu to display the Special Actions menu.

```
Special Actions
-----
c)  Install Client Software...
b)  Backup Databases...
r)  View and Display Retention Levels

i)  Initiate Request Daemon
t)  Terminate Request Daemon

h)  Help
q)  Quit Menu
```

ENTER CHOICE: **r**

Then, choose the **r** (View and Display Retention Levels) option to display the following menu:

```
Current retention definitions
```

```
      level   Period
      ----   -
0      0      1 week
1      1      2 weeks
2      2      3 weeks
3      3      1 month
4      4      2 months
5      5      4 months
6      6      6 months
7      7      9 months
8      8      1 year
9      9      infinity
```

```
ENTER 'r' to restore defaults
```



'*' indicates the retention is used in a current schedule
 Select the retention level you wish to change (0-8,r,q=quit)>

Note If an asterisk appears in front of a retention level, it indicates that the retention level is referenced in a currently defined schedule and that changing it could have adverse effects on the schedules using it.

1. Select the retention level. A prompt appears for you to specify the units.

The retention level can be any number between 0 and 8. You cannot change Level 9, it must remain as infinite (infinite for this application is defined to be 30 years).

2. Specify the units to be used (for example, days).

3. After selecting the units, you are prompted for the period. Specify a period and press Return.

The period may be either infinite (which for this application is defined to be 30 years) or a value from 0 (no retention) up to 30 years.

When you press Return, the screen is updated with the new definition and the following prompt appears (the new definition is not saved yet however).

Select another retention level to change. (0-8, r ,q=quit)

- ◆ To edit another retention level enter a number from 0 through 8.

- ◆ To restore all the levels to their default values, enter r.

4. When you are through changing retention levels, press q.

You will see a message that says: *Building Schedule Report*. After a short wait, a report appears that summarizes the retention period changes and any possible problems that these changes could cause.

Press f to move forward through the report and then press q again to get the following prompt:

Do you want to save this definition? (y/n/r=resume editing)>

- ◆ y to save the changes and exit the menu.

- ◆ n to discard the changes and return to the Special Actions menu.

- ◆ r to make further changes to retention levels.



Performing Manual Backups

To perform a manual backup of the files associated with any class, client, and schedule choose m (Manual Backups) on the main menu. The Manual Backups menu appears.

```
Class: W2
Client:<ALL>
Schedule:w2_daily_incr (Incremental)
```

Manual Backups

- i) Initiate Backup
- b) Browse Classes Forward
- r) Browse Classes Reverse
- s) Browse Schedules
- c) Browse Client Workstations
- e) Enter Class/Client/Schedule...
- h) Help
- q) Quit Menu

ENTER CHOICE:

Manual backups are useful when you do not want to wait for the next scheduled backup to occur. You can use either of two methods to select the class, client, and schedule for a manual backup:

- ◆ Use the e (Enter Class/Client/Schedule) option to identify a specific class, client, and schedule.
- or
- ◆ Use the browse options as described below.
1. Select the class for which you want to perform the backup by browsing forward or backward (b or r options) until the name of the desired class appears on the Class line at the top of the screen.
 2. Select either a single client or all clients by pressing c (Browse Client Workstations) until the name of the desired client (or ALL for all clients) appears on the Client line at the top of the screen.



3. Select the schedule or schedules by pressing `s` (Browse Schedules) until the name of the schedule appears on the Schedule line at the top of the screen (you cannot do manual backups of user-directed schedules).
4. Start the backup by pressing `i` (Initiate Backup).

Backing Up the NetBackup Databases (catalogs)

The `b` (Backup Databases) option from the Special Actions menu provides options for backing up the NetBackup internal databases (also called catalogs).

```

Backup When: never
Output Destination: SCREEN

Backup Databases
-----
m)  Modify DB Backup Settings...
d)  Delete DB Backup Media ID...
b)  Backup DB Now...

a)  Add DB Backup File Path...
r)  Remove DB Backup File Path...

l)  List/Display DB Backup Settings
o)  Output Destination (SCREEN or FILE)
h)  Help
q)  Quit Menu

ENTER CHOICE:
```



There are two information lines above the menu:

Backup When

Shows how often the current database backup settings cause the NetBackup databases to be automatically backed up. The three possible values are:

never - must be manually initiated

after each backup schedule

after any successful backup/archive

The “Modify DB Backup Settings” procedure, later in this topic, explains each of these settings.

Output Destination

Determines where `bpadm` sends the output of a List/Display DB Backup Settings. If the word `SCREEN` appears on this line, the output appears on your terminal screen. If a file path appears (for example, `/tmp/bp_db_backup`), the output goes to that file. You can change the output setting with the `o` option.

The following procedures explain how to use the options from the Backup Databases menu. See Chapter 4 for more information on the NetBackup databases.

Cautions to Observe for Backing Up NetBackup Databases

Read the precautions listed under “Important Precautions to Observe” on page 144.



Listing Database Backup Settings

To list the current settings for backing up the NetBackup internal databases (also called catalogs), choose **l** from the Backup Databases menu. This brings up a screen similar to the one below. Table 24 defines the information on this screen.

```
Frequency of DB Backup:  after each successful backup session
```

```
Server: bunny
```

```
Sequence # 1      Last Media Used: AA0018
```

```
Written          Allocated          Type   Density  Media
-----          -
1  03/25/93 13:05:15  02/25/93 09:33:45  RMedia odiskwm  AA0016
2  03/26/93 13:06:33  02/25/93 09:33:45  RMedia odiskwm  AA0018
```

```
Paths Included:
```

```
/usr/openv/netbackup/db
```

```
/usr/openv/volmgr/database
```

```
(B)ack  (F)orward  (U)p  (D)own  (Q)uit
```

Table 24. List DB Backup Settings

Setting	Description
Frequency	How often the current database backup settings cause the NetBackup scheduler to automatically backup the databases. The paths for these databases are listed under Paths Included at the bottom of the screen. The three possibilities are as follows: never - must be manually initiated after each successful backup schedule after any successful backup/archive You select one of these options with the m (Modify DB Backup Settings) option as explained in the next procedure.
Server	NetBackup server with the database.
Sequence #	This value currently cannot be changed and is always 1.



Table 24. List DB Backup Settings

Setting	Description
Last Media Used	Path (for disk) or media id (for removable or robotic media) that was used to store the last database backup. This path or media ID is one of the two listed, unless you changed media since the last backup. For example, assume AA0018 has been used many times and you want to start using a different tape. Here, you can use Modify DB Backup Settings (see next procedure) to set the media ID to another value, such as AA0019. This erases AA0018 from line 2 and replaces it with AA0019. Last Media Used still shows AA0018 until after the next database backup.
1 and 2	The two media IDs that you assign for use in database backups. If you assign both IDs, NetBackup alternates between them, always using the one that was not used for the previous backup (based on the time in the Written column). If 1 or 2 are removable or robotic type media (see Type below), they must be in the <i>NetBackup</i> media pool in Media Manager's volume database. Their media IDs, however, cannot be among those that NetBackup uses for backup or archive images.
Written	Date and time the media was last used and is <i>never</i> if it has not been written.
Allocated	If the Type column indicates that the media is removable or robotic (RMedia), the Allocated column shows the date and time the media was assigned as a NetBackup database backup tape. If the Type column indicates that the media is disk, the Allocated column shows <i>n/a</i> because an assignment is not done for disk.
Type	Type of media that this media ID represents and is either RMedia (removable or robotic) or Disk.
Density	Empty if the media type is disk. Otherwise, it shows the density of the media for this ID.
Media	Media ID (if removable or robotic media) or path (if disk) of the assigned media.
Paths Included	Paths for the databases you are backing up.

Modify Database Backup Settings

To modify current settings or initially configure the media and other settings for backing up the NetBackup internal databases (also called catalogs), choose **m** from the Backup Databases menu and follow the prompts. The prompts you see will be similar those in the screen on the next page (example user responses are in bold. <CR> means carriage return).

The following describes the steps to follow when responding to the prompts.

Caution If you modify any information regarding a media ID previously used for backups, the Written date and time for this media ID is overwritten in the database. The contents of the media itself is not destroyed unless used again. For example, assume you change to a different media ID in order to make a extra copy of the databases. When you change to the new media ID, NetBackup replaces the old ID with the new ID and no longer tracks the old ID in its database. This results in the media associated with old ID being made available for reassignment by Media Manager.

For example, assume that you change to a different media ID in order to make a extra copy of the databases. When you change to the new media ID, NetBackup replaces the old ID with the new ID and no longer tracks the old ID in its database. This results in the media associated with old ID being made available for reassignment by Media Manager.

1. Select when you want database backups to occur. The choices are:

never - must be manually initiated: NetBackup NEVER automatically backs up its databases. You must do it yourself with the **b** (Backup DB Now) option (see “Performing Manual Database Backups” on page 600).

after each successful backup schedule: Back up the databases after any regularly scheduled backup sessions that result in the creation of at least one successful backup image. Database backup DOES NOT occur after a manual or user-directed backup or archive. This is the recommended method.

after any successful backup/archive: Back up the databases after any backup session that results in the creation of at least one backup or archive image. This includes scheduled, manual, and user-directed, backups and archives.

In the following example, we choose 2 (the current state is 1, as shown in parentheses).



```
(<ESC> to exit)
Backup DB When Selections:
1) never - must be manually initiated
2) after each successful backup schedule
3) after any successful backup/archive
Enter Selection [1-3]: (1) 2 <CR>
Enter Server Name:(bunny)elk
    Remember to check paths backed
    up to assure they are correct.
Modify ID 1? (y/n): y
Storage Unit Type Selections:
    1) Disk
    2) Removable or Robotic Media
Enter Type [1-2]: (1) <CR>
Enter ID (path): /nb/dbbackup <CR>
Modify ID 2? (y/n): y
Storage Unit Type Selections:
    1) Disk
    2) Removable or Robotic Media
Enter Type [1-2]: (1) 2
Density Selections
1) dlt - DLT Cartridge
2) 8mm - 8mm Cartridge
.
.
.
8) dlt2 - DLT Cartridge alternate
Enter Choice [1-9]: 1 <CR>
    Enter ID (vsname): RR1005
    Make change now? (y/n): y
```



2. Specify the server to which these backups will be sent.

The default is the current value shown in parentheses after the Enter Server Name prompt. During initial configuration the default is always the master server.

If you are changing the destination to a media server, ensure that the server has been previously configured (that is, named in the `bp.conf` file on the master server when you started `bprd` and `bpdbm`).

Also, if you are backing up to a media server, do not forget to modify the database-backup paths for the master server as explained in step 7.

3. Choose whether you want to modify the first of the two available media IDs (ID1).

- ◆ Answer `n` to leave the media ID unchanged, then go to step 5.
- ◆ Answer `y` to change the ID, then go to step 4.

Caution A database backup DOES NOT span a tape volume. All the backup data must fit on one tape. Therefore, it is EXTREMELY important for the administrator to choose a media type that can hold all the data to be backed up. The size requirement is dependent on the size of the databases. NetBackup notifies you if the backup fails.

4. Choose the storage unit type (number in parentheses shows the current type).

- ◆ Choose 1 for Disk type and enter the path to which you want to write the database backup. This should be to a subdirectory. NetBackup creates the path if it does not exist and produces an error if the path exists and is a file rather than a directory (the error occurs when the backup is done, *not* when you specify the path).
- ◆ Choose 2 for a Removable or Robotic type storage unit and select the density (5 in the example). Then, enter the media ID (volume serial number) of the media you want to use.

5. Choose whether you want to modify the second media ID (ID 2). If you answer `y`, you are prompted as shown for media ID 1 in step 3.

6. Choose whether you want to make the changes:

- ◆ Choose `y` to change the configuration.
- ◆ Choose `n` to abort the operation and leave the configuration unchanged.

Either choice returns you to the Backup Databases menu.



7. If you are backing up the databases to a media server (see step 2), then modify the database-backup paths for the master server as follows:
 - a. Remove each database-backup path for the master server by using the Removing Database Backup File Paths option on the Backup Databases menu (see “Removing Database Backup File Paths” on page 602).
 - b. Add back each database-backup path for the master server by using the Add DB Backup File Path option on the Backup Databases menu (see “Adding Database Backup File Paths” on page 601).

When you add the paths back, be certain to specify them as follows:

```
master_name:database_backup_path
```

For example, if the platform is named bunny, the paths are:

```
bunny: /usr/opensv/netbackup/db
```

```
bunny: /usr/opensv/volmgr/database
```

Delete Database Backup Media ID

To delete a media ID from those used for backing up the NetBackup internal databases (also called catalogs), choose d from the Backup Databases menu and follow the prompts. The example below deletes ID 2. Use the Modify DB Backup Settings option from the Backup Databases menu to add or modify a media ID.

```
Delete ID 1 (AA0016)? (y/n): n
```

```
Delete ID 2 (AA0018)? (y/n): y
```

```
Are you sure you want to delete ID2? (y/n): y
```

Performing Manual Database Backups

To manually start an immediate backup of the NetBackup internal databases (also called catalogs), choose b (Backup DB Now) from the Database Management menu.

If you choose this selection, the following prompt appears:

```
WARNING: Backing up the database may take a while.
```

```
Are you sure you want to continue? (y/n):
```

Note If the media ID used for the database backup is not in a robot, you get a mount request for that media ID. If the mount request is not honored, a manual database backup must wait for the mount before proceeding. A scheduler driven database

backup must also wait for the mount and, because the scheduler is waiting, all other backups and archives must also wait until the database backup is complete (see “Modify Database Backup Settings” on page 597 for more on scheduler driven backups).

- ◆ Choose **y** to start the database backup. NetBackup uses the least recently used of the two media IDs you have assigned for backups. You must wait for completion of the backup to regain control of your terminal session.
- ◆ Choose **n** to abort the operation.

Adding Database Backup File Paths

To add database-backup paths, choose a from the Backup Databases menu. This option lets you add NetBackup internal database files (also called catalogs) to the list of files that you back up. In some cases, it will be a new addition and, in other cases, it will be to change existing paths. For example, if you back up your databases to a media server, you use this option to add the new path specifications for the master server. The following is an example with user responses in bold (<CR> means carriage return).

Adding new DB backup file paths (<ESC> to abort, Blank Line to End)

```
-----
Enter File Path:elk:/usr/opensv/netbackup/db/media
Enter File Path:elk:/usr/opensv/volmgr/database/ltidevs
Enter File Path:elk:/usr/opensv/volmgr/database/robotic_def
Enter File Path: <CR>
```

Proceed with the change? (y/n): **y**

The procedure is as follows:

1. Enter the file paths at the Enter File Path prompt, using one of the following formats:
 - ◆ For master-server databases that you are backing up to the master server:

file_path

For example:

```
/usr/opensv/netbackup/db
/usr/opensv/volmgr/database
```
 - ◆ For master-server databases that you are backing up to a media server:



master_name:file_path

For example, if the master server is named bunny:

bunny: /usr/opensv/netbackup/db

bunny: /usr/opensv/volmgr/database

- ◆ For media-server databases:

server_name:file_path

For example, if the media server is named elk without a volume database:

elk: /usr/opensv/netbackup/db/media

elk: /usr/opensv/volmgr/database/ltidevs

elk: /usr/opensv/volmgr/database/robotic_def

2. End your list of absolute or full file path entries by entering a Return at the Enter File Path prompt to leave a blank line. This brings up the following prompt:

Proceed with the change? (y/n):

3. To confirm your entries, type y at the Proceed with the change? (y/n): prompt. To abort the operation and leave the configuration unchanged, enter n.

Removing Database Backup File Paths

To remove database-backup file paths (also called catalog-backup file paths), choose r from the Backup Databases menu and follow the prompts as in the example below (example user responses are in bold. <CR> means carriage return):

Do you want to remove /usr/opensv/netbackup/db? (y/n): **n**

Do you want to remove /usr/opensv/volmgr/database? (y/n): **n**

Do you want to remove elk:/usr/opensv/netbackup/db/media: **y**

Deleting elk:/usr/opensv/netbackup/db/media.....

Proceed with the change? (y/n): **y**

This option lets you delete server database files from the list of files that you back up. In some cases, the removal will be permanent and, in other cases, it will be part of a change. For example, if you back up your databases to a media server, you use this option to delete the old path specifications for the master server and then add the new path specification with the Add DB Backup File Path option (see “Modify Database Backup Settings” on page 597).



The topics in this appendix provide additional information about various aspects of NetBackup configuration and management.

- ◆ Rules for Using Host Names in NetBackup
- ◆ Terminal Configuration on UNIX
- ◆ Media Management Concepts
- ◆ Overview of Auspex FastBackup Classes
- ◆ Factors Affecting Backup Time
- ◆ Determining NetBackup Transfer Rate
- ◆ Guidelines for Setting Retention Periods
- ◆ Guidelines for Setting Backup Frequency
- ◆ Determining Backup Media Requirements
- ◆ How NetBackup Builds Its Automatic-Backup Worklist
- ◆ Incremental Backups Overview
- ◆ Storage Management Overview
- ◆ Media Management Concepts



Rules for Using Host Names in NetBackup

NetBackup uses host names to identify, communicate with, and initiate processes on NetBackup client and server computers. The correct use of host names during configuration is essential to the proper operation of NetBackup. Also, see “Dynamic Host Name and IP Addressing” on page 286.

Qualifying Host Names

A major consideration when configuring host names is the extent to which you qualify them. In many cases, using a computer’s short host name is adequate. If the network environment is or will eventually be multi-domain, qualify host names to the extent that servers and clients can identify each other in a multi-domain environment.

For example, use a name such as

mercury.bdev.null.com

or

mercury.bdev

rather than just mercury.

The following two discussions provide more information by explaining:

- ◆ How NetBackup uses host names
- ◆ How to update NetBackup for client host name changes

How NetBackup Uses Host Names

The following discussions explain where NetBackup stores host names and how it uses them. These discussions also mention factors to consider when choosing host names.

Server and Client Name on UNIX Servers and Clients

On both UNIX servers and clients, the `SERVER` entries in the `bp.conf` file define the NetBackup servers that are allowed access. The first `SERVER` entry identifies the master server and it is to this server that client requests are made. For this reason, the `SERVER` name must be one by which all clients can connect to the server.

If more than one `SERVER` entry exists, the additional entries identify other NetBackup servers that can initiate scheduled backups on the client. The `bp.conf` file must have multiple `SERVER` entries if any remote media servers are configured. The NetBackup Request daemon and NetBackup Database Manager do not run on any server other than the master.

When a client makes a list or restore request to the server, the NetBackup client name as specified on the client is used to determine whether to allow the operation. The client name used is usually the `CLIENT_NAME` from the client's `bp.conf` file. However, in the case of alternate client restores, it can also be a name specified through the user interface or with a parameter on the `bprestore` command.

For a list or restore request to be successful, the NetBackup client name must match the name that is specified for the client in the NetBackup configuration on the server. The only exception to this rule is if the server is configured to allow alternate client restores.

Host Names on Windows NT/2000 Servers and PC Clients

Windows NT/2000 NetBackup servers and PC clients, also have `SERVER` and `CLIENT_NAME` settings. On these systems, you specify them either in a configuration file or through the user interface.

Class Configuration

The host name that you specify for a client when you add it to a class is called the client's *configured name*, and is the client's host name as it appears in the NetBackup configuration. NetBackup also adds a `CLIENT_NAME` entry to a UNIX client's `bp.conf` file when software is first installed on the client and sets the entry to match the configured name.

The server uses the client's configured name to connect to the client and start the processes that satisfy client requests. When adding clients to a class always use host names that are qualified to the extent that all NetBackup servers can connect to the clients.

When a client makes a user backup, archive, or restore request to the NetBackup server, the server uses the peername of the client (identified from its TCP connection) to determine the client's configured name.

If you add a client to more than one class, always use the same configured name in all cases. Otherwise, the client cannot view all files backed up on its behalf and file restores are complicated because both user and administrator action is required to restore from some of the backups.

Image Catalog

A subdirectory in the image catalog is created for a client when a backup is first created for that client. The subdirectory's name is the client's configured name.

Every backup for a client has a record in this subdirectory. Each of these backup records contains the host name of the server on which the backup was written.



Error Catalog

NetBackup uses entries in the error catalog for generating reports. These entries contain the host name of the server generating the entry and the client's configured name, if applicable. The server host name is normally the server's short host name (for example, shark instead of shark.null.com).

Scheduler

The NetBackup scheduler uses the server host name associated with the storage units to start a process on the server. When you specify this host name, always qualify it to the extent necessary for the master server to make a connection to the server that has the storage units. Normally, a short host name is adequate (for example, shark instead of shark.null.com).

Catalog Backup Information

If you configure media servers and include catalog files from the media server in your NetBackup catalog backups, qualify the host name portion of the media server's catalog file path to the extent necessary to allow the master server to make a connection to the media server.

How to Update NetBackup After Host Name Changes

Note Do not change the host name of a NetBackup server. This practice is not recommended because it can be necessary to import all previously used media to the server before you can use it under the new host name.

Follow these steps to update the NetBackup configuration if a client's host name is changed.

1. On the master server:

- ◆ Delete the client's old name from all classes in which it exists and add the client's new name to those classes. You do not have to reinstall NetBackup software on the client. The client also still has access to all previous backups.
- ◆ Create a symbolic link from the client's old image directory to its new image directory. For example,

```
cd /usr/opensv/netbackup/db/images
ln -s old_client_name new_client_name
```

2. On the client:

- ◆ On PC clients, you can change the client name setting either through the user interface or in a configuration file.



- ◆ On UNIX clients, change the `CLIENT_NAME` value in the `bp.conf` file to the new name.

Note If users on UNIX clients have a `bp.conf` file in their `$HOME` directory, they must change `CLIENT_NAME` in that file to the new name.

3. On the client, change the client name setting either through the user interface or in a configuration file (see the users guide for the client).

Special Considerations For Domain Name Service (DNS)

In some requests to the master server, client software sends the name that it obtains through its `gethostname(2)` library function. If this (possibly unqualified) name is unknown to the Domain Name Service (DNS) on the master server, it is possible that the master server cannot reply to client requests.

Whether this situation exists, depends on how the client and the server are configured. If `gethostname(2)` on the client returns host names that are not qualified to the extent that DNS on the master server can resolve them, you will encounter problems.

A possible solution is to reconfigure the client or the master server DNS hosts file. However, because this is not always desirable, NetBackup allows you to create a special file named

```
/usr/opensv/netbackup/db/altnames/host.xlate
```

in order to force the desired translation of NetBackup client host names.

Each line in the `host.xlate` file has three elements, a numeric key and two hostnames. Each line is left-justified, and each element of the line is separated by a space character.

```
key hostname_from_client client_as_known_by_server
```

Where

- ◆ *key* is a numeric value used by NetBackup to specify the cases where translation is to be done. Currently this value must always be 0, indicating a configured name translation.
- ◆ *hostname_from_client* is the value to translate. This must correspond to the name obtained by the client's `gethostname(2)` and be sent to the server in the request.
- ◆ *client_as_known_by_server* is the name to substitute for *hostname_from_client* when responding to requests. This name must be the name configured in the NetBackup configuration on the master server and must also be known to the master server's network services.

For example, the line

```
0 danr danr.eng.aaa.com
```



specifies that when the master server receives a request for a configured client name (numeric key 0), the name `danr` is always replaced by the name `danr.eng.aaa.com`. This resolves the problem mentioned above, assuming that:

- ◆ The client's `gethostname(2)` returned `danr`.
- ◆ The master server's network services `gethostbyname(2)` library function did not recognize the name `danr`.
- ◆ The client was configured and named in the NetBackup configuration as `danr.eng.aaa.com` and this name is also known to network services on the master server.

Terminal Configuration on UNIX

The following provides information that will be useful in setting up and resolving problems with your UNIX displays.

Modifying the XKeysymDB File

The X interface programs depend on the `/usr/lib/X11/XKeysymDB` file to correctly use editing keystrokes such as backspace and delete. The NetBackup installation procedures are designed to install a nominal version of this file if it does not already exist. It does not, however, overwrite any existing version of the file, thereby ensuring that user-customized versions of XKeysymDB are not destroyed.

Sun systems typically already have a version of XKeysymDB that does not define some editing keys. Analyze the current content of your XKeysymDB file and compare it to the NetBackup standard version which is in the `/usr/opensv/netbackup/bin/XKeysymDB` file on NetBackup servers. Then append the NetBackup version onto the existing one, or replace the existing one with the NetBackup version, or hand-edit the existing version to achieve your preferred keyboard support.

Modifying the terminfo File

The following discussion provides information on terminfo files that will be useful in resolving terminal problems with the character-based interfaces.

To modify a terminfo source file, start with an existing source file. If one is not readily available, obtain one by using `infocmp(1)` and `infocmp(8)` to print out `terminfo(4)` descriptions. For example:

```
infocmp term-type > /tmp/terminfo.file
```

The `terminfo(4)` and `terminfo(5)` man pages show the symbols that are valid for a terminfo source file, along with explanations of their use.

To find the actual character sequence sent by special keys that can subsequently be used in a terminfo source file, enter the following at the command line prompt:

```
stty -echo; cat -v; stty echo
```

and then type the special keys, following each with a carriage return. Type CTRL-D when you are finished. Another possible way to capture the character sequence is to type CTRL-V while in the insert mode of `vi`. This causes `vi` to echo the character sequence generated by the next keypress.

Once you have a suitable terminfo source file, use the following steps to install the file:

1. Move the terminfo source file to the desired machine
2. If this terminfo change is not to be globalized for all machine users, make a directory to contain the compiled terminfo files and set an environment variable to affect the terminfo search path:

```
mkdir ~/terminfo
setenv TERMINFO ~/terminfo
```

Note that if the terminfo is to be used by all users on this machine you must execute these commands as root in order to have the correct permissions to install the compiled terminfo entries.

3. Use `tic(1)` or `tic(8)` to compile the terminfo file:

```
tic /tmp/terminfo.file
```

To make use of the new terminfo file use one of the following commands:

```
setenv TERM new_terminfo
set term=new_terminfo
```

Changing X Resources for `xbp`

The X-Windows-based user-interface program has enough space in its text fields to accommodate user-selected fonts up to the 10x20 size. The windows can also fit on a 1024x768 screen with some room for window manager borders.

The `xbp` program uses default resources for your X server. You might want to change this default if it is not satisfactory to you. For example, the default font might be too small for comfortable viewing. A font that has proportional spacing can also be unsatisfactory because proportional fonts cause column misalignment on certain fields that have columnar text.

For an explanation on how to specify X resources to force a particular font, look near the end of the `xbp` Tutorial help file. If your X server defaults to proportional fonts and you do not know what font names are available for your system, try the generic font name *fixed*.



Because the default window sizes are small enough to fit on a 1024x768 screen, if you use larger fonts, `xbp`'s *verbose* display mode can force you to use a horizontal scrollbar to see all information. Widening the window at runtime with the window manager may eliminate the scrollbar. You can also change X resources as described in the Tutorial help file to enlarge the default geometry and use a smaller font.

Reading Backup Images with tar

NetBackup uses a modified GNU `tar` for reading backup images. This `tar` can understand compressed files, sparse files, long pathnames, and has features similar to those in `cpio`. If you want to read NetBackup tapes manually, it is best to use `/usr/opensv/netbackup/bin/tar`.

You can also use most other versions of `tar` to read NetBackup created tapes after using the `mt` command to position to the proper tape location. Refer to "Media Format" on page 642 for information on the location of the tapemarks.

If you use a version of `tar` other than the one provided by NetBackup, it will not support all the features provided by NetBackup and as a result you will encounter the following problems:

- ◆ You cannot recover a backup that was compressed.
- ◆ If the backup has pathnames longer than 100 characters, `/usr/opensv/netbackup/bin/tar` generates files with names of the form:

`@@MaNgLeD.nnnn`

that contain the real file

and a file named

`@@MaNgLeD.nnnn_Rename`

and for a long symbolic link, there will be a file named

`@@MaNgLeD.nnnn_Symlink`

The `@@MaNgLeD.nnnn_Rename` files explain how to rename the `@@MaNgLeD.nnnn` file in order to get it back to the proper location. View the `@@MaNgLeD.nnnn_Rename` and files and perform the file renaming.

The `@@MaNgLeD.nnnn_Symlink` files contain descriptions of the symbolic links that need to be made to get a link back to the proper file. View the `@@MaNgLeD.nnnn_Symlink` files and create the proper symbolic links.

- ◆ Multiplexed backups cannot be read by any version of `tar`.
- ◆ You cannot recover a backup that contains raw partitions (this includes FlashBackup and FastBackup images).



- ◆ NDMP client backup images cannot be restored using any version of tar, but NDMP vendors may have tools or a utility which could perform a restore directly from the media.
- ◆ If the backup contains sparse files, use the NetBackup version of tar. The `/bin/tar` on most systems have trouble with the sparse files and skip them.
- ◆ HP, AIX, and Sequent ACLs are restored in a separate file of the form:

.SeCuRiTy.nnnn

and the file has to be read and the ACLs regenerated by hand.

- ◆ VxFS extent attributes are restored in a separate file of the form:

.ExTeNt.nnnn

and the file has to be read and the extent attributes regenerated by hand.

- ◆ HP CDFs are restored, but the directory will no longer be hidden and the name of the directory has a + appended to it.
- ◆ If the backup spans more than one media, you must read the fragments from the media and concatenate the fragments to give to tar. The system's dd command might be useful in accomplishing this.

Another possibility is to use tar on the fragments. This probably will allow you to recover any file in the backup other than the one that spanned the media.

Some versions of the HP9000-800 `/bin/tar` command are known to give a *directory checksum error* for the second fragment of a backup that crossed media.

- ◆ Some versions of Solaris tar will combine the atime, mtime, and ctime strings with the file name and create file paths that are not desirable. Use the NetBackup tar instead.
- ◆ Backups from Apollo/DomainOS clients are created using their native wbak utility; the rbak utility must be used to restore from these backups.

The following process explains the commands necessary if you do decide to use another tar to read a backup from a NetBackup tape. This sequence of commands assumes that the media is known to Media Manager and that the tape drive is under Media Manager's control (see note 5 at the end of this procedure).

Before starting, you must obtain the required information:

- ◆ Media id of the tape containing the required backup
- ◆ Tape file number of the backup on the tape (see the NetBackup Images on Media report for this tape)
- ◆ Tape type/density
- ◆ Tape pool



Then, execute the following commands:

1. `tpreq -ev media_id -a r -d density -p poolname -f /tmp/tape`

Where:

- ◆ *media_id* is the media id of tape containing the backup.
- ◆ *density* is the density of the tape.
- ◆ *poolname* is the volume pool to which the tape belongs

2. `mt -f /tmp/tape rew`

3. `mt -f /tmp/tape fsf file_#`

Where:

file_# is the tape file number of the backup on tape Determine the tape file number by checking the NetBackup Images on Media report for the tape.

4. `mt -f /tmp/tape fsr`

5. `/bin/tar -tvfb /tmp/tape blocksize`

Where:

- ◆ *blocksize* is 64 (assuming that the tape is written with 32K blocks)

6. `tpunmount /tmp/tape`

Notes:

1. This procedure will NOT work for optical platters
2. This procedure will NOT work if you have compressed backups using NetBackup client software compression.
3. This procedure will NOT work if the backups were encrypted by NetBackup Encryption. In this case, the backups will be recovered but they will be encrypted and you will not be able to decrypt them.

To determine if a backup is encrypted, execute `tar -t` prior to the recovery. The output for an encrypted backup will be similar to the following:

```
erw-r--r-- root/other Nov 14 15:59 1997 .EnCryYpTiOn.388
-rw-r--r-- root/other Oct 30 11:14 1997 /etc/group.10-30
```

Where the e at the beginning of line one indicates that the backup is encrypted. There will also be other messages if you attempt the recovery.

4. This procedure will NOT work on multiplexed backup tapes.
5. This procedure will NOT work as-is if the backup you desire spans tapes.



6. This procedure will NOT work on Solaris. You cannot use the system `tar` (`/usr/sbin/tar`) on Solaris to read NetBackups because that `tar` command uses the `ctime` and `atime` fields differently than other `tar` commands.

When trying to restore using `/usr/sbin/tar`, you will see directories with large numbers being created at the top level. These directories are from the `ctime` and `atime` fields being read as path names.

You can, however, use `/usr/openv/netbackup/bin/tar` or GNU `tar` to read the backups on Solaris platforms.

7. Steps 1 and 6 are optional in a standalone environment. If step 1 is skipped, DOWN the drive and then substitute the `/dev` path of the drive in place of `/tmp/tape` in the other steps. Remember to UP the drive when you are done.

Example

The following example was successful on an HP9000-800 using a DOWNed 4-mm standalone drive and the NetBackup `tar`.

```
mt -t /dev/rmt/0hncb rew
mt -t /dev/rmt/0hncb fsf 1
mt -t /dev/rmt/0hncb fsr 1
/usr/openv/netbackup/bin/tar tvfb /dev/rmt/0hncb 64
```

Some platforms require other options on the `tar` command. The following was required on a Solaris 2.4:

```
/usr/openv/netbackup/bin/tar -t -v -f /dev/rmt/0hncb -b 64
```

Overview of Auspex FastBackup Classes

An Auspex FastBackup class is a special NetBackup class that supports very-fast raw backups of disks attaching to Storage Processors on Auspex NetServer clients. These backups are essentially raw-partition backups, but much faster because data transfer is from device to device without directly involving either NetBackup or the UNIX kernel in I/O operations.

The following topics provides further explanation of how backups work in an Auspex FastBackup class and how to configure NetBackup for them.

Note These topics discuss Auspex NetServers only to the extent necessary to explain their role as a NetBackup client. For more information on NetServer operation or administration, see the documentation provided by Auspex.



Functional Overview

The following figure shows an Auspex NetServer and Table 25 provides a brief description of the major components in that figure.

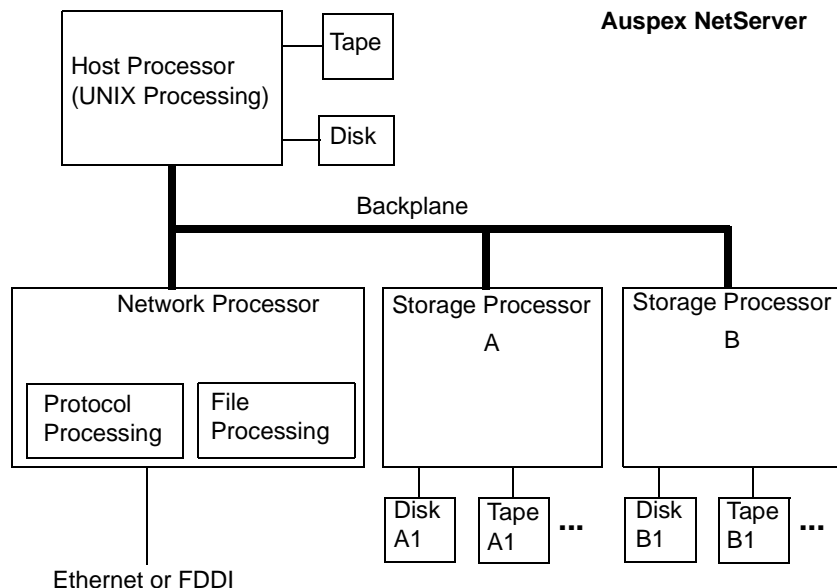


Table 25. Auspex FastBackup Client

Component	Description
Host Processor	Contains the main UNIX processing unit and is the processor that runs the NetBackup client and server software.
Network Processor	Provides interfaces to the network(s) and controls network protocol and file processing. There are separate processors for network protocol and file processing.
Storage Processor	Controls input and output for storage devices attaching to these processors. A single client can have multiple storage processors, each supporting multiple storage devices.

The following explains backup operations for Auspex clients by first considering the case where the client is in a Standard class and then where it is in an Auspex FastBack class.

Auspex Client in a Standard Class

If you configure an Auspex client in a Standard class, you can back up file systems or raw partitions from disks on a Storage Processor to storage units that are on another Storage Processor or the Host Processor. To a NetBackup user or administrator, the operation is the same as for other UNIX clients. There are no special considerations outside of Auspex specific device configuration.

In a Standard class, the data passes through both the NetBackup application and the Host Processor. A file system backup also interacts with the File System Processor.

For a raw partition backup in a Standard class, you must unmount raw partitions before backing them up in order to get a consistent view of the file system.

Auspex Client in an Auspex FastBackup Class

Auspex FastBackup classes support raw partition backups from a disk on a given Storage Processor to any other storage device on that same Storage Processor. For example, in the previous figure, a backup from Disk A1 can go to a tape on Storage Processor A, but not to a tape on Storage Processor B. In addition, a client in an Auspex FastBackup class, cannot utilize devices on the Auspex Host Processor as storage units.

Note Because the client disk must always be on the same Storage Processor as the storage unit, an Auspex FastBackup client must always be on a NetBackup server.

An Auspex FastBackup operation is essentially a device-to-device data transfer that does not directly involve either the NetBackup application or the UNIX kernel (Host Processor). The result is much higher performance than is possible with a file system or standard raw partition backup.

Users can continue making changes while the backup is occurring. It is not necessary to unmount the disk or mount it read only as with a standard raw partition backup. This “live backup” capability is made possible by a *freeze* mechanism that preserves a view of the file system for the backup.

The *freeze* mechanism requires a cache partition to store user changes made during the backup. This cache partition must be on the same Storage Processor as the storage device and the partition being backed up. The file list for the class must specify the location of the cache, as explained under “File List and Cache” on page 617.

To restore the backups, a user or administrator uses `jbpsA` or `bp`.

Master and Media Server Configuration

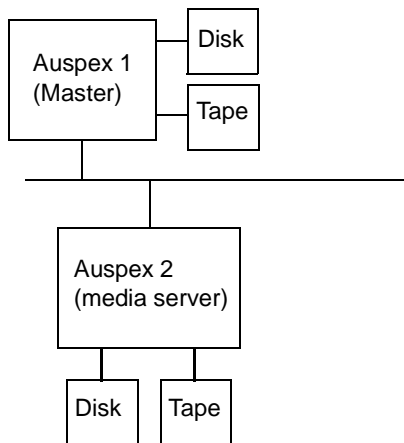
If you have more than one Auspex FastBackup client, it can be advantageous to set up a master and media server configuration, so the scheduler is on a central host; although, the same restrictions apply to the client disks and storage units. For example, in the figure below:



- ◆ The scheduler is on Auspex1
- ◆ Auspex1 backups must go between a disk and storage unit on the same Storage Processor on Auspex1.
- ◆ Auspex2 backups must go between a disk and storage unit on the same Storage Processor on Auspex2.

Auspex FastBackup Classes

Auspex1
Auspex2



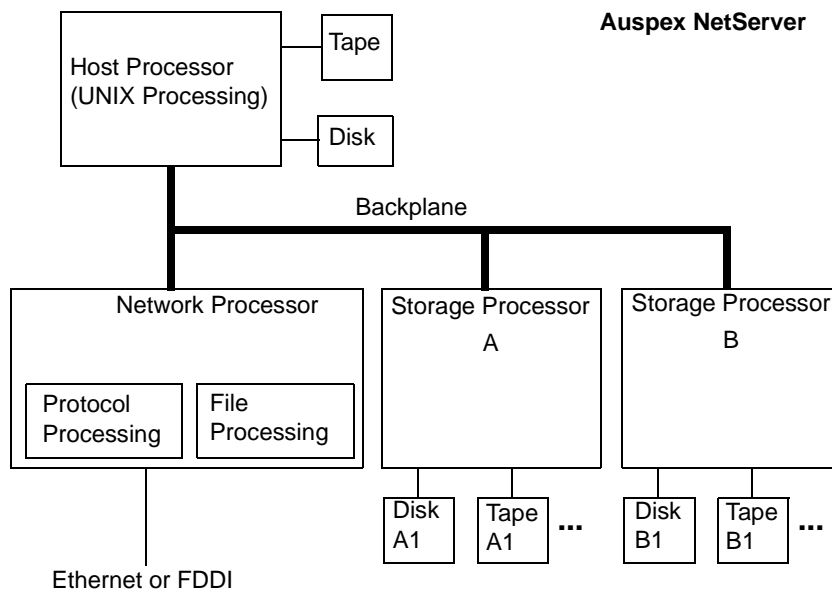
Configuring Auspex FastBackup Classes

The following explains only the configuration differences between standard and Auspex FastBackup classes.

Storage Units

A storage unit for an Auspex FastBackup class or schedule, must be on the same Storage Processor as the disk you are backing up with that class. For example, in the figure below, if you are backing up partitions from disk A1, you must use tape A1 on Storage Processor A. You cannot use a tape on Storage Processor B on the Host Processor.

There are no restrictions on using a specific storage unit for both Auspex FastBackup and standard backups. In the figure below, you can use Tape A1 on processor A to store Auspex FastBackup raw partitions and also to store files backed up by a Standard class type. The files for the Standard class can even be from a client on another host.



Class Attributes

You select the class type when you specify the general attributes for the class. For an Auspex FastBackup class, you always set the Class Type attribute to Auspex FastBackup.

NetBackup forces the following settings to No. You cannot change them.

- ◆ Cross Mount Points
- ◆ Follow NFS
- ◆ Compression
- ◆ Collect True Image Recovery Information

An Auspex FastBackup class also has an attribute called "Individual File Restore From Raw" that will enable you to restore individual files from raw partition backups.

File List and Cache

The entries in the file list for an Auspex FastBackup class are the same as for a standard raw partition backup, except that there must be a special `CACHE` entry at the top of the list.



For example:

```
CACHE=/dev/rad8h  
/dev/rad15a  
/dev/rad15d
```

The `CACHE` entry specifies the cache partition where the operating system will store any changes that users make while the backup is in process. This cache is part of the mechanism that preserves a consistent view of the partition for NetBackup.

You need to specify only the location of the cache when configuring NetBackup. The Auspex operating system controls the writing and reading of the cache and configuring it is an Auspex system administration task.

The basic requirements for the cache partition are:

- ◆ It must be on the same Storage Processor as the disk you are backing up.
- ◆ It cannot be the partition being backed up or one that is used as part of another file system or virtual partition.
- ◆ It must be unmounted.
- ◆ It must have enough space to hold all the information that users will write during the backup. For example, backups during nonworking hours normally require a smaller cache than a backup during peak activity times.

Schedules

When you configure a schedule, NetBackup forces the Maximum MPX to 1, thus disabling multiplexing. You cannot change this setting. (You can still select multiplexing for the storage unit, however, because non-Auspex FastBackup classes may be using it for this purpose.)

Restoring Files to an Auspex FastBackup Client

Users restore these raw partition backups through `jbpsA` or `bp` in the same manner as for other class types. If the “Individual File Restore From Raw” attribute is set for the Auspex FastBackup class, then it is also possible to restore individual files from these backups if the Individual File Restore From Raw attribute is set for the class.

Factors Affecting Backup Time

The time NetBackup requires to complete a backup is an important factor in scheduling. This is particularly true for sites that deal with large amounts of data. For example, the total backup time can exceed the time allotted to complete backups and interfere with



normal network operations. Longer backup times also increase the possibility of a problem disrupting the backup. The time to back up files can also give you an indication of how long it takes to recover them.

The following formula shows the major factors that affect backup time:

$$\text{Backup Time} = \frac{\text{Total data}}{\text{Transfer rate}} \times \text{Compression Factor} + \text{Device Delays (optional)}$$

Total data

The amount of data you must back up depends on the size of the files for each client in the class you are backing up. It also depends on whether it is a full or incremental backup.

- ◆ Full backups involve all the data. Therefore, a full backup usually takes longer than an incremental.
- ◆ Differential-incremental backups include only the data that has changed since the last full or intervening incremental.
- ◆ Cumulative-incremental backups include all the data that has changed since the last full backup.

With both differential- and cumulative-incremental backups, the amount of data in the backups, depends on the frequency with which files change. If a large number of files change frequently, incrementals are larger.

Transfer rate

Transfer rate depends on factors such as:

- ◆ Speed of the backup device. For example, sending backups to a tape having a maximum transfer rate of 400 kilobytes per second normally takes less time than to a tape that transfers at only 200 kilobytes per second (assuming other factors allow taking advantage of the faster transfer rate).
- ◆ Available network bandwidth. The theoretical network bandwidth is about 10 megabits per second for Ethernet and 100 megabits per second for FDDI and 100 base T. The available bandwidth, however, is less than this and depends on how much other network traffic is present. For example, multiple backups occurring on the same network compete for bandwidth.
- ◆ Speed with which the client can process the data. This varies with the hardware platform and depends on the other applications running on the platform. File size is also an important factor. Clients can process larger files faster than smaller ones. You can back up 20 files that are 1 megabyte in size faster than 20,000 files that are 1 kilobyte in size.



- ◆ Speed with which the server can process the data. Like client speed, server speed also varies with the hardware platform and depends on the other applications running on the platform. The number of concurrent backups being performed also affects server speed.

See the next topic “Determining NetBackup Transfer Rate” for ways to compute the transfer rate for your clients.

Compression

If you use software compression, it often multiplies the backup time by a factor of two or three for a given set of data.

Device delays

Device delays are due to factors such as the device being busy, loading the media, and finding the place on the media at which to start writing the backup. These delays depend on the devices and computing environments and can vary widely.

Determining NetBackup Transfer Rate

You can calculate three different variations of the backup transfer rate by using the data provided in NetBackup reports. The three rates and the methods for calculating them are as follows.

Network-Transfer Rate

The network-transfer rate considers only the time required to transfer data over the network from client to server. This rate ignores:

- ◆ Time to load and position media before a backup.
- ◆ Time to gracefully close the tape file and write an additional NetBackup information record to the tape.

The network-transfer rate is the one provided in the All Log Entries report.

Network-Transfer Plus End-of-Backup-Processing Rate

This rate ignores the time it takes to load and position media before a backup, but includes the end-of-backup processing that is ignored in the network transfer rate. To determine this rate, use the All Log Entries report and calculate the time from the message:

```
begin writing backup id xxx
```

to the message

```
successfully wrote backup id xxx
```

Then, divide this time (in seconds) into the total bytes transferred (as recorded in the All Log Entries report) to calculate the transfer rate.

Total-Transfer Rate

This transfer rate includes the time for loading and positioning the media as well as the end-of-backup processing. Using the List Client Backups report, calculate the transfer rate by dividing Kilobytes by Elapsed Time (converted to seconds).

Examples

Assume that the reports provide the following data.

All Log Entries Report

```
TIME                SERVER/CLIENT      TEXT
04/28/94 23:10:37 windows giskard begin writing backup
                    id giskard_0767592458, fragment 1 to
                    media id TL8033 on device 1 . . .
04/29/94 00:35:07 windows giskard successfully wrote
                    backup id giskard_0767592458,
                    fragment 1, 1161824 Kbytes at
                    230.325 Kbytes/sec
```

List Client Backups Report

```
Client:                giskard
Backup ID:              giskard_0767592458
Class:                  production_servers
Client Type:            Standard
Sched Label:            testing_add_files
Schedule Type:          Full
Backup Retention Level: one week (0)
Backup Time:            04/28/94 23:07:38
Elapsed Time:           001:27:32
Expiration Time:        05/05/94 23:07:38
Compressed:             no
Kilobytes:              1161824
```



Number of Files: 78210

Using the backup data from the example reports above, you get the following three rates:

Network Transfer Rate

1161824 Kbytes at 230.325 Kbytes per second

Network Transfer Plus End of Backup Processing Rate

23:10:30 - 00:35:07 = 01:24:30 = 5070 seconds

1161824 Kbytes/5070 = 229.157 Kbytes per second

Total Transfer Rate

Elapsed time = 01:27:32 = 5252 seconds

1161824 Kbytes/5252 = 221.216 Kbytes per second



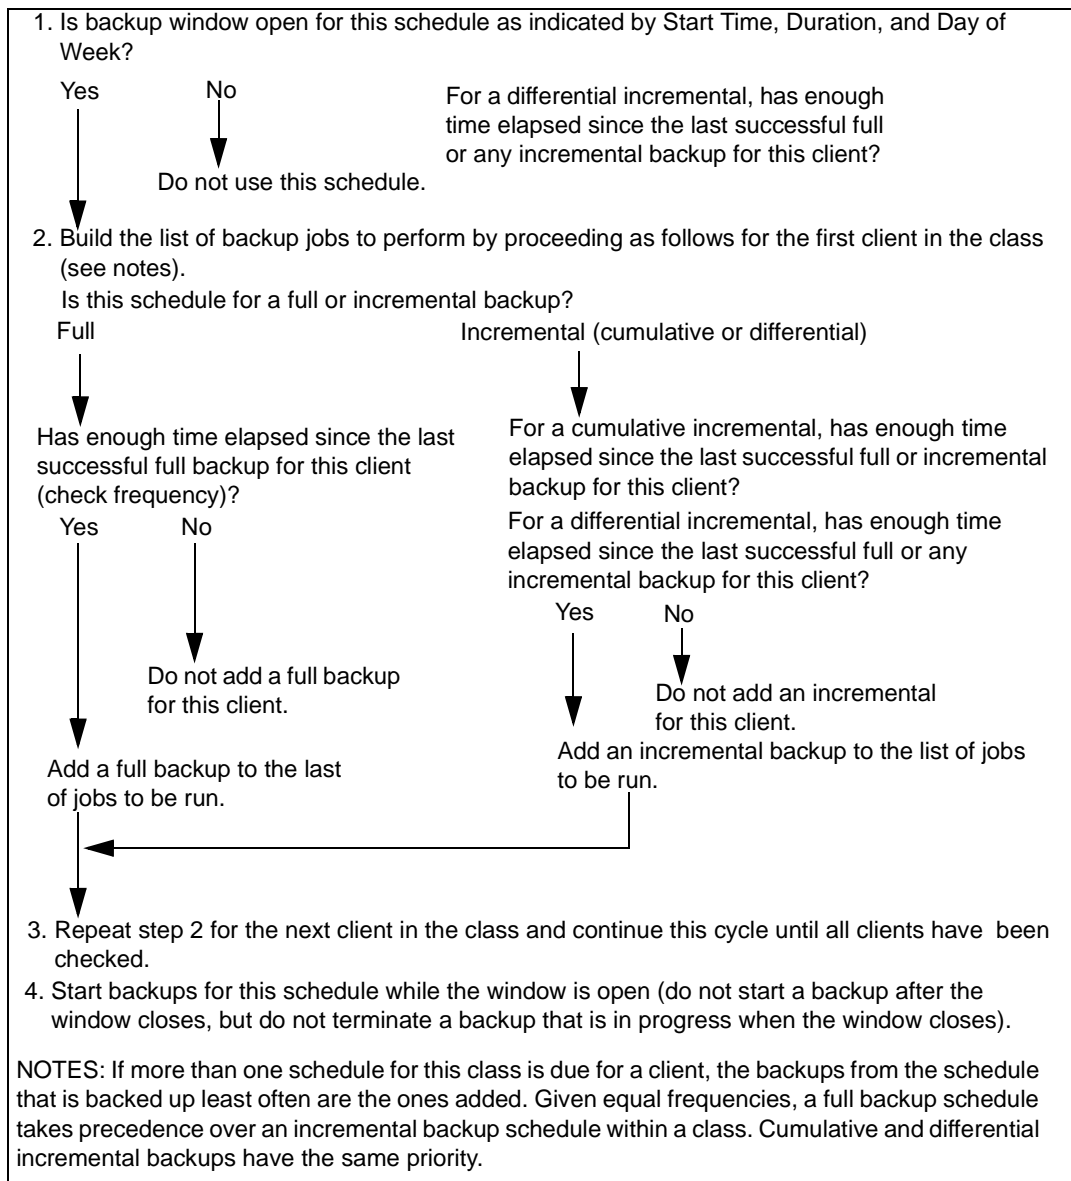
How NetBackup Builds Its Automatic-Backup Worklist

The following topics explain how NetBackup determines the order in which automatic backups occur for each client. This information is for reference only but will be useful in evaluating problems with your schedules.



Building the Worklist (Queue)

When the backup window opens for an automatic-backup schedule, NetBackup proceeds as shown in the following figure to determine whether to add the client backups for that schedule to the worklist (queue).



Prioritizing the Worklist

The worklist typically contains jobs from different classes and schedules. NetBackup checks for the following when determining the order in which to run the backups that are in the worklist:

1. Highest priority backup as determined by the class Priority attribute.

Backup jobs from the class with the highest priority runs first.

For example, assume that clients named Ant and Beetle are in different classes and that Ant is in the class with the highest priority. Here, the jobs for Ant always run before the Beetle jobs.

2. Backup with a retention level that is the same as a tape that is currently mounted.

If class priorities are equal, NetBackup tries to start a backup job that has the same retention period as a tape that is currently mounted. This reduces delays in waiting for tape mounts.

For example, assume that Ant and Beetle are in the same class but their schedules have different retention periods. Also, assume that the Ant job is the most overdue. However, a tape is mounted that has the same retention level as Beetle.

Here, the Beetle job runs first because it can be stored on a tape that is already mounted, thus making the most efficient use of resources. If there is another drive of the correct type available, a tape will be mounted on that drive for the Ant job.

3. Most overdue backup job.

If the priorities and retention level are equal, NetBackup prioritizes backups according to how long they are overdue. The clients that are the most overdue have the highest priority.

NetBackup determines how long a backup is overdue by subtracting the backup frequency (on the schedule) from the length of time since the last successful backup for that client.

For example, assume that Ant and Beetle have backup jobs that are in the same class and have the same retention level. Also assume that the schedules for these backup jobs both have a frequency of 1 day. If the last backup for Ant ran 25 hours ago and the last backup for Beetle ran 26 hours ago, then both clients are overdue for a backup. However, the Beetle job is the most overdue and will run first.

This approach ensures that a backup that was not successful during its previous backup window has priority over backups that were successful. This is important on a busy system where the backup window can sometimes close before all backups can begin.



Guidelines for Setting Retention Periods

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

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

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

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

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

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

Guidelines for Setting Backup Frequency

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



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

Daily backups are usually incrementals that record the changes since the last incremental or full backup. This conserves resources because incrementals use less storage and take less time to perform than full backups.

Full backups usually occur less frequently than incrementals but should occur often enough to avoid accumulating too many consecutive incrementals. Too many incrementals between full backups increases restoration time because of the effort required to merge those incrementals when restoring files and directories. When setting the frequency for full backups:

- ◆ Choose longer times between full backups for files that seldom change. This uses fewer system resources. It also does not significantly increase recovery time because there should be smaller incremental backups.
- ◆ Choose shorter times between full backups for files that change frequently. This decreases restore time. It can also use less resources because it reduces the cumulative effect of the longer incrementals that are necessary to keep up with frequent changes in the files.

To achieve the most efficient use of resources, ensure that most of the files in a given class change at about the same rate. For example, assume that about half the files in a class file list change frequently enough to require a full backup every week, but the rest rarely change and require only monthly full backups. Here, if all the files are in the same class, you must perform full backups weekly on all the files. This wastes system resources and media because half the files need full backups only once a month. A better approach is to divide them into two classes, each with the appropriate backup schedule.

Determining Backup Media Requirements

To assist you in determining how much media is available, NetBackup provides:

- ◆ The NetBackup Media Summary report, which lists the active and nonactive media that is available to a server.
- ◆ The `available_media` script in the `/usr/opensv/netbackup/bin/goodies` directory, which lists all the media IDs that are available on the server where you run the script.

To efficiently manage your backup environment, however, you must also know the amount of media that is required for both daily and long-term use. The daily requirement must be known to ensure that enough tape volumes and disk space are available for each backup session. The long-term requirements are necessary to assess costs for acquisition of new media, storage devices, and off-site storage (if required).



For daily requirements, you must first determine the approximate amount of data in the files that you will back up to each type of media each day. Then, you can check the Media Summary report and the results from running the `available_media` script to verify that enough media IDs and disk space are available.

For long term planning, you must also consider the following:

- ◆ How long you retain the data. A related consideration is that all backups on a given tape or optical have the same retention level. This means that if you have many different retention levels, you need more tapes or optical disks, unless you add `ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA` to the `bp.conf` file.
- ◆ Duplicates for off-site storage or extra security.
- ◆ New software releases and other special backups.
- ◆ Replacing worn out media.
- ◆ Changes in disk usage patterns over the time period under consideration. If your disk usage and capacity increase, your backup needs will also probably increase.
- ◆ Number of backups that are on a tape. Because tape marks are created between backups, a tape with many small backups (as with incrementals) contains less real data than if it contains fewer large backups. The size of the tape marks vary depending on the media type. A large number of small files will also have a higher percentage of overhead in the backup because each file requires an extra 512 bytes for catalog information on the tape or disk.
- ◆ If you have many different volume pools, ensure that enough media is defined in each one to accommodate the data.


Incremental Backups Overview

NetBackup supports two types of incremental backups:

- ◆ Differential
- ◆ Cumulative


A differential incremental backs up only the data that has changed since the last full or incremental. The following example shows the data that is included in a series of backups between January 1 and January 4. The January 1 backup is a full and includes all files and directories in the class file list. The subsequent backups are differential incrementals 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 incrementals are required for the recovery.





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

A cumulative incremental backs up all the data that has changed since the last full backup. The following example shows the data that is included in a series of backups between January 1 and January 4. The January 1 full backup includes all files and directories in the class file 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 and the last cumulative-incremental are required for the recovery.



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

The next two topics compare the relative retention requirements and the backup and restore times of these two types of incremental backups. The third topic in this section explains how NetBackup determines the files to include in an incremental backup.

Retention Requirements

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.



Type	Retention Requirement	Comments
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 incrementals since the last full backup are required 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 incrementals 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 incrementals 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 following series of backups.

day 1	day 2	day 3	day 4	day5	day 6	day 7	day 8
Full	Diff	Cum	Diff	Cum	Diff	Full	Diff

- ◆ Every other day a differential-incremental backup occurs, which will usually have a minimum backup time.
- ◆ On alternate days, a cumulative-incremental backup occurs, which will require more time than the differential but not as much as a full. The differential 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 incrementals were used. The fulls can be done less often if the amount of data being backed up by the incrementals is small.

Determining Files Due for Backup - Windows NT/2000 Clients

On Windows NT/2000 clients, NetBackup performs incremental backups of files based on the Perform Incrementals Based on Archive Bit setting in the NetBackup Configuration dialog box. To open this dialog box, start the client-user interface on the client, click Configure on the Actions menu, and go to the General tab.

If the Perform incrementals based on archive bit box is checked, incrementals for this client are based on the state of each file's archive bit. 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, depends 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 box is clear, NetBackup includes a file in an incremental backup only if the file's datetime stamp 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 file's datetime stamp against the last full or incremental backup.
- ◆ For a cumulative-incremental backup, NetBackup compares the file's datetime stamp 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 - UNIX Clients

The following explains how NetBackup determines that a file on a UNIX client is due for an incremental backup.

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
- ◆ `atime` -- the file access time
- ◆ `ctime` -- the inode change time

UNIX man pages contain a definition of these attributes.

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 file's `mtime`, and then reset it after the modification using the `utime(2)` system call.

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 file's `atime`, and then reset it after the file access using the `utime(2)` system call.

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 can not 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.

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 file's access time. For this reason, NetBackup saves the file's `atime` and `mtime` prior to reading the file, and (by default) resets the `atime` and `mtime` using the `utime(2)` system call. By "covering it's tracks,"

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 file's ctime.

As an option to a NetBackup configuration, customers can choose to have NetBackup not reset the file's access time after it reads a file. Additionally, customers can choose to have NetBackup use the file's ctime, 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 file's mtime to determine what files and directories to back up.

When a file is moved from one location to another, the file's ctime 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 directory's ctime 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.

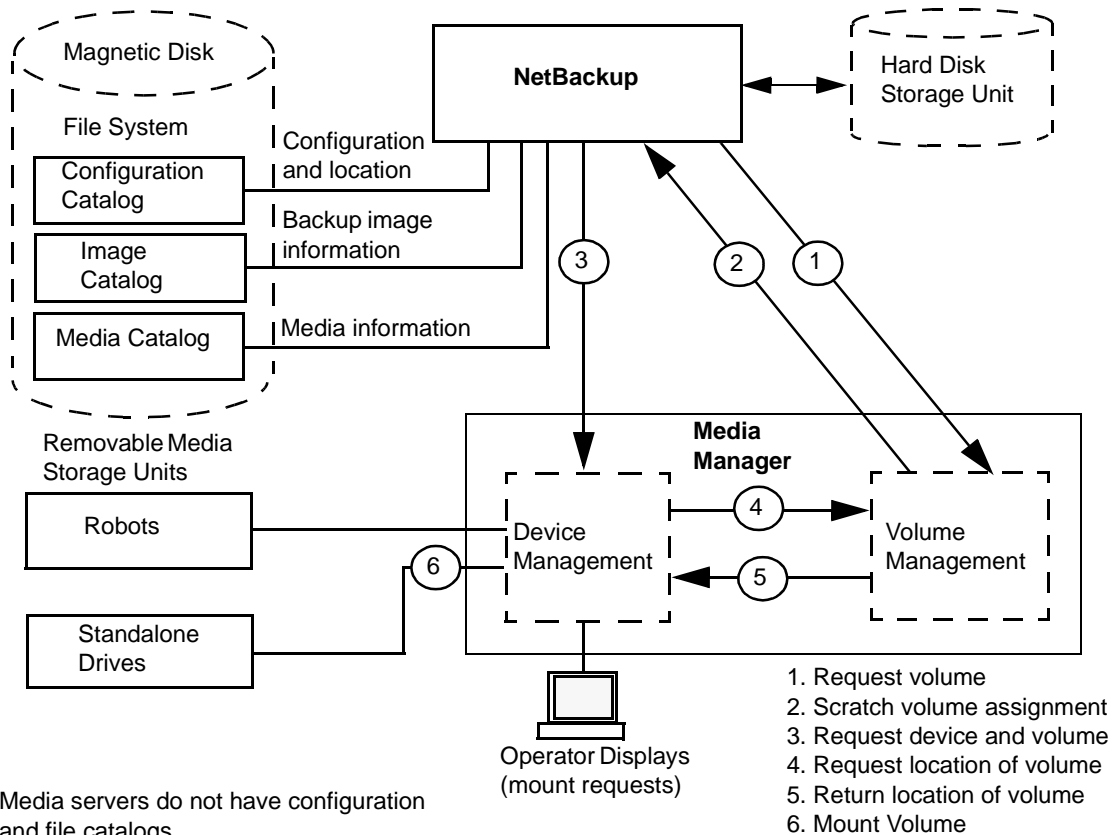
In either case, these moved files and directories will obviously be included in subsequent full backups.

Also, see "True image restore information" on page 70.



Storage Management Overview

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



Storage Units

The peripheral that stores the backup data for NetBackup is referred to as a *storage unit*. In this context, the term *storage unit* means a group of one or more storage devices of a specific type and density that attach to a NetBackup server. The storage devices can be for removable media such as tape or a directory on a hard disk. Removable media devices can be robotic or standalone.

The administrator defines the storage units that are available and which storage unit to use for each class. For example, it is possible to specify a robot as the storage unit for one class and a standalone tape drive for another class.

Media Manager

NetBackup keeps records about the files in the backups and also about the media where they 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 volume's media ID and device density. 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 Media Manager is managed separately and can also be used by other applications, such as Storage Migrator.

Retention

The administrator specifies the retention period for the files associated with each schedule. It is possible to keep all backups with the same retention period on the same volume or to mix different retention periods on a volume.

Volume Pools

For Media Manager storage units, NetBackup supports a concept called volume pools. A *volume pool* is a distinct set of media that can be used *only* by the users and hosts that you designate when configuring the pool. You create volume pools and assign media to them when you configure Media Manager. 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 named *NetBackup* is always created by default and, unless you specify otherwise in the class or schedule, all backups go to media in the *NetBackup* pool. You can, however, create other pools for NetBackup to use. For example, if you create *Auto* and *User* volume pools, you can specify that automatic backups will use media from the *Auto* pool and user backups will go to media in the *User* 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 system administrator's guide for Media Manager.

Media Management Concepts

This section discusses concepts involved in managing NetBackup media.

NetBackup and Media Manager Catalogs

NetBackup and Media Manager use internal databases to keep information about media and device configuration. With the exception of the volume database, these internal databases are usually referred to as catalogs.

Caution Do not remove or manually edit the NetBackup or Media Manager catalogs. These files are for internal program use only and altering them in any way can result in permanent loss of data.

Volume Database

The volume database has information about volumes that have been configured for use by Media Manager. When you add volumes they are recorded in the volume database. The volume database resides in the `/usr/opensv/volmgr/database` directory.

When you add new volumes, you do it on the NetBackup server that has the volume database. A part of this process is assigning media IDs.

Media IDs must be unique and can consist of six or less alphanumeric characters. Optical disks each have two media IDs, one for side A and one for side B. The terms media ID and External Volume Serial Number (EVSN) are equivalent. Storage areas on disk are identified by their pathname.

Media Catalog

NetBackup keeps a media catalog with information that correlates backups to the volumes where they are stored. Each NetBackup server maintains a media catalog for the storage units attaching to that server.



During installation, the media catalog is created in the `/usr/opensv/netbackup/db/media` directory. NetBackup refers to the media catalog when it needs a volume for a backup or restore. If the media catalog does not contain a suitable volume, NetBackup has Media Manager assign one. In this manner, the catalog is populated as NetBackup uses new volumes for backups.

When the retention period has ended for all backups on a volume, NetBackup deletes it from the media catalog. Media Manager then deassigns the volume so it is available for reassignment at a future date.

Note Volumes for backups of the NetBackup catalogs are a special case and do not appear in the media catalog. You must track the media IDs for these volumes separately so you can find them in case the media catalog is damaged. However, they do appear in the Media Manager volume catalog and are listed as assigned to NetBackup (they are deassigned only if you delete them from your catalog backup settings).

Device Catalogs

The device catalogs have information about the drives and robots that are in NetBackup storage units. When you configure drives and robots, Media Manager stores this information in its device catalogs. These catalogs are located under `/usr/opensv/volmgr/database`.

Media States

The following media states apply to volumes that are in the NetBackup media catalog but are not active. That is, the volumes cannot be used for both backups and restores.

Table 26. Media States

State	Description
FULL	NetBackup sets FULL status if it encounters an end of media (EOM) during a backup. A full volume is unavailable for future backups until the retention period expires for all backups that are on it. At that time, the volume is deleted from the NetBackup media catalog and deassigned from NetBackup.



Table 26. Media States (continued)

State	Description
SUSPENDED	<p>NetBackup can automatically set volumes to the <i>SUSPENDED</i> state as explained in “ALLOW_MEDIA_OVERWRITE” on page 327 or “Automatic Media Suspend Or Device Down” on page 645.</p> <p>You can also use the <code>bpmmedia</code> command to manually suspend or unsuspend volumes (see the <code>bpmmedia(1M)</code> man page).</p> <p>When an active NetBackup volume is suspended:</p> <ul style="list-style-type: none">◆ Backups are still available for restores (if they are expired you must import them first).◆ The volume is unavailable for future backups until the retention period ends for all backups that are on it. At that time, the suspended volume is deleted from the NetBackup media catalog and deassigned from NetBackup.
FROZEN	<p>NetBackup can automatically set a volume to the <i>FROZEN</i> state (see “ALLOW_MEDIA_OVERWRITE” on page 327. You can also use the <code>bpmmedia</code> command to manually freeze or unfreeze volumes (see the <code>bpmmedia(1M)</code> man page).</p> <p>When an active NetBackup volume is <i>FROZEN</i>:</p> <ul style="list-style-type: none">◆ Backups are still available for restores (if they are expired you must import them first).◆ The volume is unavailable for future backups.◆ The volume never expires, even when the retention period ends for all backups that are on the media. This means that the media ID is never deleted from the NetBackup media catalog and, therefore, remains assigned to NetBackup.
IMPORTED	<p>NetBackup automatically sets volumes to the <i>IMPORTED</i> state if they are imported to this server.</p> <p>When an active NetBackup volume is <i>IMPORTED</i>:</p> <ul style="list-style-type: none">◆ Backups are still available for restores (if they are expired you must import them first).◆ The volume is unavailable for future backups until the retention period ends for all backups that are on it. At that time, the imported volume is deleted from the NetBackup media catalog and deassigned from NetBackup.

How NetBackup Selects Media in a Robot

When NetBackup automatically selects a volume in a robot, it proceeds as follows:



1. Searches the NetBackup media catalog for a volume that is already mounted in a drive and meets the following criteria:
 - ◆ Configured to contain backups at the retention level required by the schedule (unless `ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA` is present in the `bp.conf` file).
 - ◆ In the volume pool required by the backup being performed.
 - ◆ Not in a FULL, FROZEN, IMPORTED, or SUSPENDED state.
 - ◆ Of the same density required by the requested backup and, in the case of a robotic storage unit, in the robot requested by the backup.
 - ◆ Not currently in use by another backup or a restore.
 - ◆ Not written in a protected format (see “ALLOW_MEDIA_OVERWRITE” on page 327). This is detected after the volume is mounted. If the volume is in a protected format, it is unmounted and NetBackup resumes the search.
2. If NetBackup cannot find a mounted volume that satisfies the above conditions, it checks its media catalog for any volume that is suitable.
3. If the media catalog does not have a suitable volume, NetBackup requests Media Manager to assign one. Media Manager then assigns a volume to NetBackup that meets all of the following criteria:
 - ◆ Is the correct media type
 - ◆ Is for the correct robot type (if applicable)
 - ◆ Is located in the requested robotic peripheral (if applicable)
 - ◆ Resides on the requested host
 - ◆ Is in the correct volume pool
 - ◆ Is not currently assigned (not already allocated to NetBackup)
 - ◆ Is not expired (if an expiration date is defined in Media Manager)
 - ◆ Has not exceeded the maximum number of mounts allowed
4. If more than one volume qualifies, Media Manager chooses the one with the least mounts. NetBackup then adds it to the media catalog and assigns it the specified retention level.
5. If there are no unassigned volumes of the requested type, the backup terminates with an error indicating no available media.

Spanning Media

After an end of media condition, automatic media selection is a special case and depends on whether NetBackup is configured to allow backups to span media.



- ◆ NetBackup spans media if `DISALLOW_BACKUPS_SPANNING_MEDIA` is not included in the `bp.conf` file. Here, NetBackup uses another volume to start the next fragment and the resulting backup is composed of fragments on different volumes.
- ◆ NetBackup does not span media if `DISALLOW_BACKUPS_SPANNING_MEDIA` is specified. Here, the backup terminates abnormally and the operation is retried according to the Schedule backup attempts global attribute.

How NetBackup Uses Media in Standalone Drives

The section explains media selection and other aspects of standalone drive operations.

Media Selection Using Standalone Drive Extensions

When the standalone-drive extensions feature is enabled, NetBackup tries to use whatever labeled or unlabeled media happens to be in a standalone drive. This capability is enabled by default during installation. The selection process is as follows:

1. If a backup is requested and an appropriate standalone drive does not contain a volume, NetBackup selects one in the same way as explained in “How NetBackup Selects Media in a Robot” on page 638.

The Device Monitor shows the mount request and an operator must manually insert the volume and assign it to a drive.

2. If an appropriate drive contains a volume, NetBackup tries to select and use the volume that is in the drive.

- ◆ A volume that has been previously used for backups must:
 - ◆ Not be FULL, FROZEN, or SUSPENDED
 - ◆ Be at the same retention level and in the same volume pool as the backup being performed, unless you include `ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA` in the `bp.conf` File
- ◆ Previously unused media are used by NetBackup.

If the unused media is unlabeled, you can label it ahead of time by using the `bplabel` command. When using this command, you can specify the `-u` parameter in order to force assignment of a specific drive index. This eliminates the need to manually assign the drive. Refer to the `bplabel(1M)` man page for more information.

If the media is unlabeled:

- ◆ NetBackup labels the media.

- ◆ Media Manager adds a media ID to the volume configuration, if necessary. If a media ID is added, the NetBackup `bp.conf` entry `MEDIA_ID_PREFIX` is used as the first characters of the media ID. If a media ID is added, the NetBackup property Media ID Prefix is used as the first characters of the media ID. If `MEDIA_ID_PREFIX` is not specified, the default prefix is A. For example, A00000.
- ◆ Media Manager adds the requested volume pool to the volume configuration (if the backup class specifies a volume pool).

Disabling Standalone Drive Extensions

You can disable the standalone drive extensions by including `DISABLE_STANDALONE_DRIVE_EXTENSIONS` in the `bp.conf` file. This causes NetBackup to use the same method to select media for standalone drives as it uses for robotic drives (see “How NetBackup Selects Media in a Robot” on page 638).

Spanning Media

Media selection following an end of media condition is a special case and depends on whether NetBackup is configured to allow backups to span media.

- ◆ NetBackup spans media if `DISALLOW_BACKUPS_SPANNING_MEDIA` is not present in the `bp.conf` file. Here, NetBackup selects another volume to begin the next fragment and the resulting backup has fragments on more than one volume.
- ◆ NetBackup does not span media if `DISALLOW_BACKUPS_SPANNING_MEDIA` is specified. Here, the backup terminates abnormally when the end of media is reached and the operation is rescheduled according to the Schedule backup attempts global attribute.

When an end of media is encountered on a standalone drive that uses a gravity feed stacker (a stacker not controlled by software), you can have NetBackup continue on the next volume loaded by the stacker, rather than looking for another drive. To do this, add the `MEDIA_REQUEST_DELAY` option in the `bp.conf` file. This setting specifies a number of seconds for NetBackup to pause before looking for another drive.

Keeping Standalone Drives in the Ready State

To leave standalone drives in a ready condition after a backup or restore completes, use the `-nsu` (no standalone unmount) option when executing the `ltid` command. This option prevents `ltid` from ejecting the tape when Media Manager issues a `tpunmount` after an operation completes. The tape does eject if end of media (EOM) is reached. See the `ltid(1M)` man page for more information on the `ltid` command.



It is possible for more than one standalone drive to be ready and contain suitable media. If that occurs, drive selection occurs in logical drive index number order. For example, if drives 2 and 3 are the same type and both contain suitable media, NetBackup selects drive 2.

Media Format

NetBackup writes to media in a format that allows position to be verified before appending new backups. The format for tape and optical media differ slightly due to characteristics of the media itself.

To determine the contents of tape or optical media, use the Media Contents report. For optical media, the offsets and sizes are shown, along with the backup ID. For tape media, the file number is shown.

Although, manually reading the media is normally required, it is possible to use the information from the Media Contents report (file position, offsets, block size, etc) in programs that manually position and read NetBackup created media. You can do this with a customized program or a standard utility such as `mt` or `dd`. The output from a customized program or standard utility are normally piped to `/usr/opensv/netbackup/bin/tar`, since the backups are `tar` compatible (or `rbak` if it is an Apollo client backup).

Non-QIC Tape Format

For all tape media except QIC, the format for backups that are not multiplexed is:

MH * BH Image * BH Image * BH Image * EH *

Where:

MH = Media Header (1024 bytes)

* = Tape Mark

BH = Backup Header (1024 bytes)

Image = Data from the backup

EH = Empty Backup Header, used for position validation.

When adding a new backup to the above example, the tape is positioned to EH and position is verified. The EH is overwritten by a BH and the backup proceeds. When complete, a new EH is written for future positioning validation. When NetBackup encounters the end of media during write, it terminates the tape with two tape marks and does not write an EH.

For the multiplexed format, see “Multiplexing Format” on page 644.

QIC Tape Format

For QIC tape media, NetBackup does not write empty backup headers (EH) so the format nonmultiplexed backups is:

MH * BH Image * BH Image * BH Image . . .

For optical media, the QIC format is:

MH BH Image EH BH Image EH BH Image EH

To append backup images to QIC media, NetBackup positions to the end of data (EOD) and then starts the next backup.

For the multiplexed format, see “Multiplexing Format” on page 644.

Note Optical disk media has no tape marks to delimit backups. The data on the optical disk are recorded in successive sectors and the offsets maintained by Media Manager. Since optical disk can seek to a random position, finding and verifying position is a very fast operation.

Fragmented Backups

For fragmented backups the media format is the same as described for QIC and non-QIC tapes, except that NetBackup breaks the backup into fragments of the size that you specify when you configure the storage unit.

For example:

MH * BH1 Image (frag 1)* BH1 Image (frag 2)* BH1 Image (frag n) * EH *

Fragmentation is intended primarily for storing large backup images on a disk type storage unit. In these instances, fragmenting images allows you to avoid exceeding the two gigabyte size limit that applies to most UNIX file systems.

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.

Fragmenting tape backups can also speed up restores because NetBackup can skip to the specific fragment before starting its search for a file. Otherwise, it has to start at the very beginning of the backup and read tar headers until it finds the desired file.

For the multiplexed format, see “Multiplexing Format” on page 644.

Note 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.



Spanning Tapes

By default, NetBackup spans a backup image to another tape if it encounters the end of media during a backup. The format is the same as explained above for fragmented backups and the first fragment on the next tape begins with the buffer of data where the end of media occurred:

First tape: MH * ... *BHn Image (frag 1) * *

Next tape: MH * BHn Image (frag2)* ... * EH *

On the first tape, NetBackup does not write an EH, and terminates the tape with two tape marks.

Multiplexing Format


The tape format for multiplexed backups is:

MH * BH1 ... BHn Data...

Where:

- ◆ MH is the Media Header (1024 bytes).
- ◆ * is a Tape Mark.
- ◆ BH1 ... BHn are backup Headers (1024 bytes). One for each job that is part of the set of jobs being multiplexed.
- ◆ Data is the data from the backup. By default, the data is in 64 kilobyte blocks (63 kilobytes on Sun). Each block also contains 512 bytes that are reserved for multiplexing control information (less than 1% of the total data) and to identify the backup that the block corresponds to.

Each time a job ends or a new job is added to the multiplexing set, NetBackup writes a tape mark and starts multiplexing the revised set of jobs. The figure below shows an example.



Labeling Media

You normally do not have to label media.

- ◆ For a robot, you select the media IDs when you configure the robot in Media Manager and tape labeling is done automatically when NetBackup uses the media. For optical media, you have the option of formatting and labeling when you add the media to the robot. Or, you can do it manually with the Media Manager `tpformat` command.
- ◆ For standalone drives, the standalone drive extension feature makes it unnecessary to label media in a standalone drive, unless you desire to do so. You can, however, prelabel tapes by using the `bplabel` command. See the `bplabel(1M)` man page for further information. Also, see “How NetBackup Uses Media in Standalone Drives” on page 640.

Note Automatic labeling does not occur if the media was last used for NetBackup catalog backups. It also does not occur if the media contains data from a recognized non-NetBackup application and you are not using the NetBackup `bp.conf` option, `ALLOW_MEDIA_OVERWRITE`. In either of these instances, you must label the media by using the `bplabel` command.

Mounting and Unmounting Media

For robots, Media Manager automatically mounts and unmounts the volume. Operator intervention is usually required only if the required volume is not in the drive.

For example, if a restore requires a volume that has been removed from a robot, the Device Monitor shows a mount request. The operator can then insert the proper volume and assign it to the request.

Automatic Media Suspend Or Device Down

NetBackup can automatically *suspend* the use of volumes or *down* a device if it suspects failures are due to the volume or the device. The reason for the action is logged in the NetBackup error catalog. Repeated write failures are usually the cause for setting a volume to the SUSPENDED state or a device to DOWN. A volume is also set to SUSPENDED if the write failure occurs in a is such that future positioning attempts will be unreliable. Write failures are frequently caused by a tape device with dirty heads or deteriorating media.

If investigation shows that the suspend or down action was incorrect, you can reverse it by:

1. Using the `bpmedia` command to unsuspend the volume.
2. Using the NetBackup Device Monitor to set the device to Up.





NetBackup Notify Scripts

D

Note Before using the notify scripts, ensure that they are executable by *other*. Do this by executing `chmod 755 script_name`. Where *script_name* is the name of the script.

NetBackup provides the following scripts (batch files on Windows 2000 and NT) for collecting information and providing notification of events.

Scripts that run on a server:

```
backup_notify
backup_exit_notify
dbbackup_notify
diskfull_notify
restore_notify
session_notify
session_start_notify
userreq_notify
```

Scripts that run on clients:

```
bpstart_notify (UNIX clients only)
bpend_notify (UNIX clients only)
bpstart_notify.bat (Microsoft Windows clients only)
bpend_notify.bat (Microsoft Windows clients only)
```

The scripts that run on a server are installed during NetBackup server installation and are in:

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

On a UNIX client, you can run only the `bpstart_notify` and `bpend_notify` scripts. Before using these scripts, you must copy them from

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



on the server to

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

on the client.

On a Windows 2000 or NT client, you can run only the `bpstart_notify.bat` and `bpend_notify.bat` scripts. These scripts are not supplied with the software. You must create them on the client per the criteria in the “`bpstart_notify.bat` (Microsoft Windows clients only).” and “`bpend_notify.bat` (Microsoft Windows clients only).” discussions.

For further information, refer to the comments in the scripts.

Caution If you use either the `bpstart_notify` or `bpend_notify` scripts, do not include commands that write to stdout. If you do, NetBackup sends this output to the server as part of the backup and the resulting backup can abort with an error message pertaining to block sizes. Also, ensure that all commands in the scripts are appropriate to the client platform. For example, the `-s` parameter is invalid for the `UNIX mail` command on some UNIX platforms and its use can cause data to be written to stdout or stderr, resulting in the same problem noted above.

backup_notify

The `backup_notify` script runs on the NetBackup server where the storage unit is located and is called each time a backup is successfully written to media. The parameters that NetBackup passes to this script are:

- ◆ The name of the program doing the backup
- ◆ The backup-image name or path

For example:

```
backup_notify bptm bilbo_0695316589
```

Note If NetBackup backed up files to a UNIX disk storage unit that is being managed by Storage Migrator, the `backup_notify` script notifies Storage Migrator to perform migration as quickly as possible. The released script does not, however, have commands to force a backup of the managed file system after NetBackup has stored its backups. To back up the managed file system, modify the script as necessary to meet site requirements for backup.

backup_exit_notify

The `backup_exit_notify` script runs on the master server. The NetBackup scheduler on the master server calls this script to do site specific processing when an individual backup has completed from the perspective of the client, Media Manager, and the image catalog.

NetBackup passes the following parameters to the script:

Parameter	Description
<code>clientname</code>	Name of the client from the NetBackup catalog.
<code>classname</code>	Class name from the NetBackup catalog.
<code>schedname</code>	Schedule name from the NetBackup catalog.
<code>schedtype</code>	One of the following: FULL INCR (differential incremental) CINC (cumulative incremental) UBAK UARC
<code>exitstatus</code>	Exit code for the entire backup job.

For example:

```
backup_exit_notify freddie production fulls FULL 0
backup_exit_notify danr production incrementals INCR 73
```

bpstart_notify (UNIX clients only)

On UNIX clients, NetBackup calls the `bpstart_notify` script each time the client starts a backup or archive operation. To use this script, copy

```
/usr/opensv/netbackup/bin/goodies/bpstart_notify
```

from the server to

```
/usr/opensv/netbackup/bin/bpstart_notify
```

on the UNIX client. Then, modify the script as desired and ensure that you have execute permission.



The `bpstart_notify` script executes each time a backup or archive starts and initialization is completed (but before the tape positioning). This script must exit with a status of 0 for the calling program to continue and for the backup or archive to proceed. A nonzero status causes the client backup or archive to exit with a status of `bpstart_notify failed`.

If the `/usr/opensv/netbackup/bin/bpstart_notify` script exists, it executes in the foreground and the `bpbkcar` process on the client waits for it to complete before continuing. Any commands in the script that do not end with an `&` character execute serially.

The server expects the client to respond with a `continue` message within the period of time specified by the NetBackup `BPSTART_TIMEOUT` option on the server.

The default for `BPSTART_TIMEOUT` is 300. If the script needs more time than 300 seconds, increase the value to allow more time.

NetBackup passes the following parameters to the script:

Parameter	Description
<code>clientname</code>	Name of the client from the NetBackup catalog.
<code>classname</code>	Class name from the NetBackup catalog.
<code>schedname</code>	Schedule name from the NetBackup catalog.
<code>schedtype</code>	One of the following: FULL INCR (differential incremental) CINC (cumulative incremental) UBAK UARC

For example:

```
bpstart_notify freddie cd4000s fulls FULL
bpstart_notify danr cd4000s incrementals INCR
bpstart_notify hare cd4000s fulls FULL
bpstart_notify freddie cd4000s user_backups UBAK
bpstart_notify danr cd4000s user_archive UARC
```

To create a `bpstart_notify` script for a specific class or class and schedule combination, create script files with a `.classname` or `.classname.schedulename` suffix. The following are two examples of script names for a class named *production* that has a schedule named *fulls*:

```
/usr/opensv/netbackup/bin/bpstart_notify.production
```

```
/usr/opensv/netbackup/bin/bpstart_notify.production.fulls
```

The first script affects all scheduled backups in the class named production. The second script affects scheduled backups in the class named production only when the schedule is named fulls.

Note For a given backup, NetBackup uses only one `bpstart_notify` script and that is the one with the most specific name. For example, if there are both `bpstart_notify.production` and `bpstart_notify.production.fulls` scripts, NetBackup uses only `bpstart_notify.production.fulls`.

The `bpstart_notify` script can use the following environment variables:

```
BACKUPID
```

```
UNIXBACKUPTIME
```

```
BACKUPTIME
```

The NetBackup `bpbkar` process creates these variables. The following are examples of strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526
```

```
UNIXBACKUPTIME=0857340526
```

```
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

In addition to the above, the following environment variables can be used for the support of multiple data streams:

`STREAM_NUMBER` indicates the stream number. The first stream started from a class, client, and schedule will be 1. A value of 0, indicates that multiple data streams is not enabled.

`STREAM_COUNT` specifies the total number of streams to be generated from this class, client, and schedule.

`STREAM_PID` is the pid (process ID) number of `bpbkar`.

bpstart_notify.bat (Microsoft Windows clients only)

For Windows 2000, NT, 98, and 95 clients, you can create batch scripts that provide notification whenever the client starts a backup or archive. These scripts must reside on the client and in the same directory as the NetBackup client binaries:

```
install_path\NetBackup\bin
```

Where *install_path* is the directory where NetBackup is installed.

You can create `bpstart_notify` scripts that provide notification for all backups or just for backups of a specific class or schedule.



To create a script that applies to all backups, name the script:

```
install_path\netbackup\bin\bpstart_notify.bat
```

Note On Windows 98 and 95 systems, use a `.pif` suffix on the batch scripts. For example, `bpstart_notify.pif`. The `.bat` suffix, as shown in the examples, applies only to Windows 2000 and NT systems.

To create a `bpstart_notify` script that applies only to a specific class or class and schedule combination, add a `.classname` or `.classname.schedulename` suffix to the script name.

- ◆ The following script applies only to a class named *days*:

```
install_path\netbackup\bin\bpstart_notify.days.bat
```

- ◆ The following script applies only to a schedule named *fulls* that is in a class named *days*:

```
install_path\netbackup\bin\bpstart_notify.days.fulls.bat
```

The first script affects all scheduled backups in the class named *days*. The second script affects scheduled backups in the class named *days* only when the schedule is named *fulls*.

For a given backup, NetBackup calls only one `bpstart_notify` script and checks for them in the following order:

```
bpstart_notify.class.schedule.bat
```

```
bpstart_notify.class.bat
```

```
bpstart_notify.bat
```

For example, if there are both `bpstart_notify.class.bat` and `bpstart_notify.class.schedule.bat` scripts, NetBackup uses only the `bpstart_notify.class.schedule.bat` script.

Note If you are also using `bpend_notify` scripts, they can provide a different level of notification than the `bpstart_notify` scripts. For example, if you had one of each, they could be `bpstart_notify.class.bat` and `bpend_notify.class.schedule.bat`.

When the backup starts, NetBackup passes the following parameters to the script:

Parameter	Description
%1	Name of the client from the NetBackup catalog.
%2	Class name from the NetBackup catalog.
%3	Schedule name from the NetBackup catalog.

Parameter	Description
%4	One of the following: FULL INCR CINC UBAK UARC
%5	Status of the operation is always 0 for <code>bpstart_notify</code> .
%6	Results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script. If the script applies to a specific class and schedule, the results file must be named <i>install_path\netbackup\bin\BPSTART_RES.class.schedule</i> If the script applies to a specific class, the results file must be named <i>install_path\netbackup\bin\BPSTART_RES.class</i> If the script applies to all backups, the results file must be named <i>install_path\netbackup\bin\BPSTART_RES</i> An <code>echo 0 > %6</code> statement is one way for the script to create the file. NetBackup deletes the existing results file before calling the script. After the script executes, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.

The server expects the client to respond with a `continue` message within the period of time specified by the NetBackup `BPSTART_TIMEOUT` option on the server. The default for `BPSTART_TIMEOUT` is 300. If the script needs more than 300 seconds, increase the value to allow more time.

For Windows 2000 and NT clients, the `bpstart_notify` script can use the following environment variables for the support of multiple data streams:

`STREAM_NUMBER` indicates the stream number. The first stream started from a class, client, and schedule will be 1. A value of 0, indicates that multiple data streams is not enabled.

`STREAM_COUNT` specifies the total number of streams to be generated from this class, client, and schedule.

`STREAM_PID` is the pid (process ID) number of `bpbkar`.



bpend_notify (UNIX clients only)

Caution The `bpend_notify` script is executed when the client is finished sending data, but the server has not yet completed writing to media.

For a UNIX client, if you need notification whenever the client completes a backup or archive operation, copy

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

from the server to

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

on the UNIX client. Then, modify the script as desired, and ensure that you have execute permission.

The `bpend_notify` script executes each time a backup or archive completes. For archives, it executes after the backup but before the files are removed.

If `bpend_notify` exists, it executes in the foreground and `bpbkar` on the client waits until it completes. Any commands that do not end with an `&` character execute serially.

The server expects the client to respond within the period of time specified by the `BPEND_TIMEOUT` NetBackup configuration option on the server. The default for `BPEND_TIMEOUT` is 300.

If the script needs more than 300 seconds, set `BPEND_TIMEOUT` to a larger value. Avoid too large a value or you will delay the server from servicing other clients.

NetBackup passes the following parameters to the `bpend_notify` script:

Parameter	Description
<code>clientname</code>	Name of the client from the NetBackup catalog.
<code>classname</code>	Class name from the NetBackup catalog.
<code>schedname</code>	Schedule name from the NetBackup catalog.
<code>schedtype</code>	One of the following: FULL INCR (differential incremental) CINC (cumulative incremental) UBAK UARC

Parameter	Description
exitstatus	Exit code from bpbkar. This is only client status and does not mean that the backup is complete and successful (see “backup_exit_notify” on page 649). For example, the client can show a status 0 when, due to a failure on the server, the All Log Entries report shows a status 84.

For example:

```
bpend_notify freddie cd4000s fulls FULL 0
bpend_notify danr cd4000s incrementals INCR 73
```

To create a `bpend_notify` script for a specific class or class and schedule combination, create script files with a `.classname` or `.classname.schedulename` suffix. The following are two examples of script names for a class named *production* that has a schedule named *fulls*:

```
/usr/opensv/netbackup/bin/bpend_notify.production
/usr/opensv/netbackup/bin/bpend_notify.production.fulls
```

The first script affects all scheduled backups in the class named *production*. The second script affects scheduled backups in the class named *production* only when the schedule is named *fulls*.

Note For a given backup, NetBackup uses only one `bpend_notify` script and that is the one with the most specific name. For example, if there are both `bpend_notify.production` and `bpend_notify.production.fulls` scripts, NetBackup uses only `bpend_notify.production.fulls`.

If the UNIX client is running NetBackup 3.0 or later software, the `bpend_notify` script can use the following environment variables:

```
BACKUPID
UNIXBACKUPTIME
BACKUPTIME
```

The NetBackup `bpbkar` process creates these variables. The following are examples of strings that are available to the script for use in recording information about a backup:

```
BACKUPID=freddie_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 1997
```

In addition to the above, the following environment variables can be used for the support of multiple data streams:



`STREAM_NUMBER` indicates the stream number. The first stream started from a class, client, and schedule will be 1. A value of 0, indicates that multiple data streams is not enabled.

`STREAM_COUNT` specifies the total number of streams to be generated from this class, client, and schedule.

`STREAM_PID` is the pid (process ID) number of bpbkar.

bpend_notify.bat (Microsoft Windows clients only)

For Windows 2000, NT, 98, and 95 clients, you can create batch scripts that provide notification whenever the client completes a backup or archive. These scripts must reside on the client and in the same directory as the NetBackup client binaries:

```
install_path\NetBackup\bin
```

Where *install_path* is the directory where NetBackup is installed.

You can create `bpend_notify` scripts that provide notification for all backups or just for backups of a specific class or schedule.

To create a `bpend_notify` script that applies to all backups, name the script:

```
install_path\netbackup\bin\bpend_notify.bat
```

Note On Windows 98 and 95 systems, use a `pif` suffix on the batch scripts. For example, `bpend_notify.pif`. The `bat` suffix, as shown in the examples, applies only to Windows 2000 and NT systems.

To create a script that applies only to a specific class or class and schedule combination, add a `.classname` or `.classname.schedulename` suffix to the script name.

- ◆ The following script applies only to a class named *days*:

```
install_path\netbackup\bin\bpend_notify.days.bat
```

- ◆ The following script applies only to a schedule named *fulls* that is in a class named *days*:

```
install_path\netbackup\bin\bpend_notify.days.fulls.bat
```

The first script affects all scheduled backups in the class named *days*. The second script affects scheduled backups in the class named *days* only when the schedule is named *fulls*.

For a given backup, NetBackup calls only one `bpend_notify` script and checks for them in the following order:

```
bpend_notify.class.schedule.bat
```

```
bpend_notify.class.bat
```

`bpend_notify.bat`

For example, if there are both `bpend_notify.class.bat` and `bpend_notify.class.schedule.bat` scripts, NetBackup uses only `bpend_notify.class.schedule.bat`.

Note If you are also using `bpstart_notify` scripts, they can provide a different level of notification than the `bpend_notify` scripts. For example, if you had one of each, they could be `bpstart_notify.class.bat` and `bpend_notify.class.schedule.bat`.

When the backup completes, NetBackup passes the following parameters to the script:

Parameter	Description
%1	Name of the client from the NetBackup catalog.
%2	Class name from the NetBackup catalog.
%3	Schedule name from the NetBackup catalog.
%4	One of the following: FULL INCR CINC UBAK UARC
%5	Status of the operation and is same as sent to the NetBackup server. This is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.



Parameter	Description
%6	<p>Results file that NetBackup checks for a return code from the script. NetBackup uses %6 to pass the file name and then expects the script to create the file in the same directory as the script.</p> <p>If the script applies to a specific class and schedule, the results file must be named <i>install_path\netbackup\bin\BPEND_RES.class.schedule</i></p> <p>If the script applies to a specific class, the results file must be named <i>install_path\netbackup\bin\BPEND_RES.class</i></p> <p>If the script applies to all backups, the results file must be named <i>install_path\netbackup\bin\BPEND_RES</i></p> <p>An <code>echo 0> %6</code> statement is one way for the script to create the file.</p> <p>NetBackup deletes the existing results file before calling the script. After the script executes, NetBackup checks the new results file for the status. The status must be 0 for the script to be considered successful. If the results file does not exist, NetBackup assumes that the script was successful.</p>

The server expects the client to respond with a *continue* message within the period of time specified by the NetBackup `BPEND_TIMEOUT` option on the server. The default for `BPEND_TIMEOUT` is 300. If the script needs more than 300 seconds, increase the value to allow more time.

For Windows 2000 and NT clients, the `bpend_notify` script can use the following environment variables for the support of multiple data streams:

`STREAM_NUMBER` indicates the stream number. The first stream started from a class, client, and schedule will be 1. A value of 0, indicates that multiple data streams is not enabled.

`STREAM_COUNT` specifies the total number of streams to be generated from this class, client, and schedule.

`STREAM_PID` is the pid (process ID) number of `bpbkar`.



dbbackup_notify

The `dbbackup_notify` script runs on the master server and is called each time NetBackup completes an attempt to back up its catalogs. NetBackup passes the following parameters to this script:

Parameter	Description
<code>device</code>	Device type the backup was written to.
<code>vsn_or_path</code>	Volume serial number (for tape) or path (for disk) used for the backup.
<code>status</code>	Specifies whether the backup was successful and must have a value of either <code>SUCCESS</code> or <code>FAIL</code> .

For example:

```
dbbackup_notify DISK /disk1/bpsync1 SUCCESS
dbbackup_notify OPTICAL AA0001 FAIL
dbbackup_notify TAPE XYZ047 SUCCESS
```

You must be able to identify the most recent catalog backup. Therefore, consider modifying this script to produce a printed copy of the media ID to which the catalog backup was done.

Note If the NetBackup catalog files are backed up to a UNIX disk storage unit that is being managed by Storage Migrator, the `dbbackup_notify` script notifies Storage Migrator to perform migration as quickly as possible. The script does not, however, have commands to force Storage Migrator to back up its own catalog after a backup of the NetBackup catalog. You must modify the script to meet site requirements for backup of the Storage Migrator catalog.



diskfull_notify

The `diskfull_notify` script runs on the NetBackup server that has the storage unit. The disk-media manager (`bpdm`) calls this script if it encounters a disk full condition when writing a backup to a disk type storage unit. The default action is to sleep five minutes and retry the write (file being written is kept open by the active `bpdm`).

You can modify the script to send a notification to someone or to perform actions such as removing other files in the affected directory or file system. NetBackup passes the following parameters to this script:

Parameter	Description
<code>programname</code>	Name of the program (always <code>bpdm</code>).
<code>pathname</code>	Path to the file being written.

For example:

```
diskfull_notify bpdm /disk1/images/host_08193531_c1_F1
```

restore_notify

Note If the NetBackup catalog files are backed up to a UNIX disk storage unit that is being managed by Storage Migrator, the `restore_notify` script notifies Storage Migrator to perform migration as quickly as possible after the restore is complete.

The `restore_notify` script runs on the server that has the storage unit. The NetBackup tape or disk manager (`bptm` or `bpdm`) calls the script when it is finished sending data to the client during a restore (regardless of whether data is actually sent). NetBackup passes the following parameters to this script:

Parameter	Description
<code>programname</code>	Name of the program doing the restore or other read operation.
<code>pathname</code>	Path to the backup.
<code>operation</code>	One of the following: <code>restore</code> <code>verify</code> <code>duplication</code> <code>import</code>

For example:

```
restore_notify bptm bilbo_0695316589 duplication
```

session_notify

The `session_notify` script runs on the master server and is called at the end of a backup session if at least one scheduled backup has succeeded. NetBackup passes no parameters to this script. The scheduler is suspended until this script completes, thus no other backups can start until that time.

session_start_notify

The `session_start_notify` script runs on the master server. When a set of backups is due to run, NetBackup calls this script to do any site specific processing prior to starting the first backup. NetBackup passes no parameters to this script.

userreq_notify

The `userreq_notify` script runs on the master server and is called by NetBackup each time a request is made to:

- ◆ List files that are in backups or archives
- ◆ Start a backup, archive, or restore

You can alter this script to gather information about user requests to NetBackup. NetBackup passes the following parameters to this script.

Parameter	Description
<code>action</code>	Defines the action and can have values of: backup archive manual_backup restore list
<code>clientname</code>	Defines the client name.
<code>userid</code>	Defines the user ID.



For example:

```
userreq_notif backup mercury jdoe
userreq_notify archive mercury jdoe
userreq_notify manual_backup mercury jdoe
userreq_notify restore mercury jdoe
userreq_notify list mercury jdoe
```



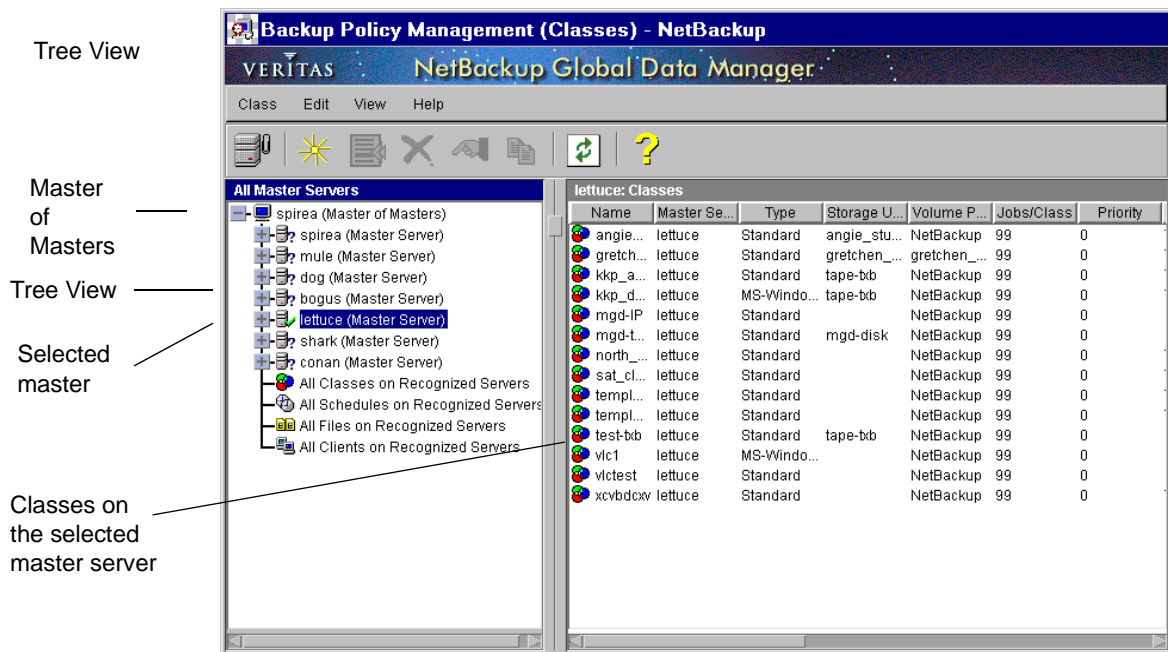
VERITAS Global Data Manager (GDM) provides the ability to monitor and administer remote NetBackup and Backup Exec servers from a single console. Backup Exec administration is limited to monitoring devices and backup jobs.

The server where the GDM option is installed is called a *Master of Masters* and can be any NetBackup master server. The Master of Masters sends administration requests to the master servers that it administers. The master servers execute the requests and return results to the Master of Masters. All backup, scheduling, and configuration information is stored on the individual master servers, not the Master of Masters.

You can run the GDM administration interface on either the Master of Masters (if it is Solaris, HP-UX, or Windows NT/2000) or a remote computer (see “Requirements” on page 665). The computer where you are running the interface is called the GDM console. If you run the interface on a remote computer, the first operation you must perform after starting the interface is to log into the Master of Masters.

In the GDM interface, the left pane in all the main administration windows has a tree view that shows the master servers that you can manage. Selecting a master server in the tree establishes a connection to that server, thus allowing standard administration tasks to be performed on it. For example, selecting a master server in the tree in the Backup Policy Management window makes it possible to add classes to the configuration on that server.





You can change the configuration on only one master server at a time. You cannot select multiple master servers and then perform operations such as adding a class to all of them.

The right pane is a detail view that shows information about what is selected in the tree view. If you select the Master of Masters (top of the tree), the detail view shows all the master servers that you can currently administer. With this exception, the details are generally the same as in the interface on other master servers. For example, if you select a master server in the Backup Policy Management window, the detail view lists the classes on that server.

The tree views in most cases have special branches that provide a view of information across master servers. For example in Backup Policy Management, if you select the All Classes on Recognized Servers branch in the tree, the detail view shows all classes on all the masters recognized by the Master of Masters. The other administration utilities have similar selections.

The tree and detail views also have shortcut menus that appear when you press the right mouse button. These menus contain frequently used commands that apply to what is currently selected.

At the master server level and below, the GDM interface behaves in exactly the same way as when you are administering the NetBackup-Java applications on a single master server. For information about using Global Data Manager in a specific NetBackup administration utility, see the online help.

GDM Terminology

- ◆ Master of Masters

The NetBackup server where the Global Data Manager option is installed. This can be any UNIX or NT/2000 NetBackup master server. The Master of Masters is sometimes called the GDM host.

- ◆ GDM console

The computer where the GDM interface is run. The Master of Masters and GDM console are the same computer if you run the GDM interface on the Master of Masters.

Installing GDM

Requirements

You must have the following in order to use GDM.

1. A NetBackup 3.4 master server to be the Master of Masters.
 - ◆ This server can be on any of the platforms listed in the Supported Platforms section of the NetBackup release notes.
 - ◆ You must choose a master server, not a media server, for the Master of Masters.
2. A GDM license key from VERITAS. You activate this key on the Master of Masters. No additional software is required because all the required components were installed with the NetBackup 3.4 server software.
3. One of the following computer types on which to run the GDM interface:
 - ◆ A Solaris or HP-UX computer that has NetBackup 3.4 server software installed. The interface software does not run on any of the other UNIX platforms.
 - ◆ Any Windows NT/2000/98/or 95 computer that has the NetBackup 3.4 Windows Display Console installed. This software is included with NetBackup 3.4 server software. For installation instructions, see the installation guide.

Note that the computer where you run the interface must have the same version of NetBackup as the Master of Masters.

Procedure

Choose the Master of Masters and activate the license on that computer. No other steps are required. All the required components were installed with the NetBackup 3.4 server software.



Configuring GDM

You must configure:

- ◆ The Master of Masters so it knows which master servers it will administer.
- ◆ All master servers and their media servers to allow the Master of Masters to access their data.

▼ **The procedure is as follows:**

1. Configure the Master of Masters:

- ◆ If the Master of Masters is a UNIX system, use the `gdm_config` utility in the `/usr/opensv/netbackup/bin/admincmd` directory. This utility takes you through all the steps that you have to perform.
- ◆ If the Master of Masters is a Windows NT/2000 system, add a `KNOWN_MASTER` entry to the registry for every master server to be administered. Add these entries to the

`HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config`

registry key on the Master of Masters. The entry has a type value of `string`.

Note If the Master of Masters has an empty list of master servers, then by default GDM shows the Master of Masters as the only master server.

2. Configure UNIX NetBackup master and media servers.

- c. On each UNIX NetBackup master server, add a `MASTER_OF_MASTERS` entry for the Master of Masters to the `bp.conf` file. The following is an example entry for a Master of Masters named `owl`.

```
MASTER_OF_MASTERS= owl
```

- d. On each UNIX NetBackup master or media server, add a `SERVER` entry for the Master of Masters to the `bp.conf` file. The following is an example entry for a Master of Masters named `owl`.

```
SERVER = owl
```

- e. On a master server, stop and restart the NetBackup `bpdbm` daemon after adding the `bp.conf` entries.

3. Configure Windows NT/2000 NetBackup Master and Media Servers.

On each master server, specify the Master of Masters:

- a. Start the NetBackup administration interface.
- b. In the NetBackup Administration window, click NetBackup Configuration on the Start menu.



- c. In the Configure - NetBackup window, select the NetBackup master server, then, on the File menu click Properties (Read/Write).

The Master Server Properties dialog box appears.

- d. On the GDM tab, each master server you selected appears in a separate column. Find the column for the master server you want to administer and add the name of the Master of Masters to the Global Data Managers section.

On each master or media server, add the Master of Masters to the server list:

- a. Start the NetBackup administration interface.
- b. In the NetBackup Administration window, click NetBackup Configuration on the Start menu.
- c. In the Configure - NetBackup window, select the NetBackup master or media server, then, on the File menu click Properties (Read/Write).

The Master Server or Media Server Properties dialog box appears, depending on the server you selected.

- d. On the Servers tab, find the column for the server you want to administer and add the name of the Master of Masters to the Additional Servers section of that column.

4. Configure Device Host Security, if required.

If any of the NetBackup master or media servers are using device host security, the `vm.conf` file on that server must contain a `SERVER` entry for the Master of Masters.

- ◆ On UNIX, the `vm.conf` file is located in the `/usr/opensv/volmgr` directory.
- ◆ On Windows NT/2000, the `vm.conf` file is located in the `install_path\volmgr` directory (where the default `install_path` is `C:\VERITAS\Program Files`).

The example below shows a `SERVER` entry for a Master of Masters named owl:

```
SERVER = owl
```

For more information on device host security and the `vm.conf` file, see the system administrator's guide for Media Manager.

5. Configure NetBackup authentication and authorization, if required.

If any of the NetBackup master or media servers use NetBackup authentication or authorization, you must set this up between those servers and the Master of Masters. For instructions, see the authentication and authorization topics in your NetBackup system administrators guide.

6. Configure Backup Exec Master Servers.



See the Backup Exec documentation for information on configuring a Backup Exec Windows NT, NetWare, or Network Storage Edition server for use with Global Data Manager.

Starting GDM

1. Start the GDM interface program as follows:

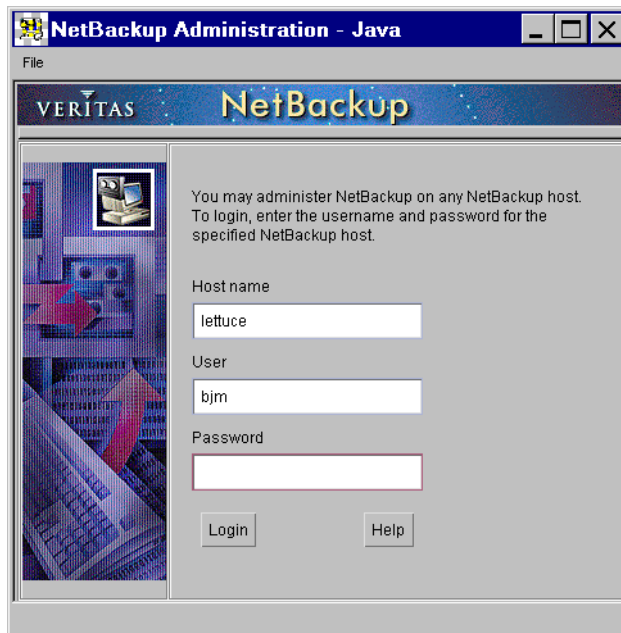
- ◆ On a Solaris or HP-UX computer that has NetBackup 3.4 server software installed, execute the following command:

```
/usr/obj/RS60000/netbackup/bin/jnbSA
```

- ◆ On a Windows NT/2000/98/95 computer that has the NetBackup 3.4 Windows Display Console installed, click the Start button and point to Programs. Then point to VERITAS NetBackup and click NetBackup - Java on *host* on the submenu. *host* is the name that appears by default in the Host name box of the login window (see step 2).

If there are no NetBackup - Java on *host* entries in the VERITAS NetBackup menu, then the Windows Display Console (WDC), is not installed on that system. For installation instructions, see the Installation Guide. The software is included with NetBackup 3.4 server software.

When the interface program starts, the login window appears.



2. Type the name of the Master of Masters in the Host name box.

For example, assume that the Master of Masters is shark and you start the interface on a Windows NT/2000 system named bear. In this example, bear is the GDM console and you specify shark (the Master of Masters) in the Host name box.

3. Type your user name and password and click Login.

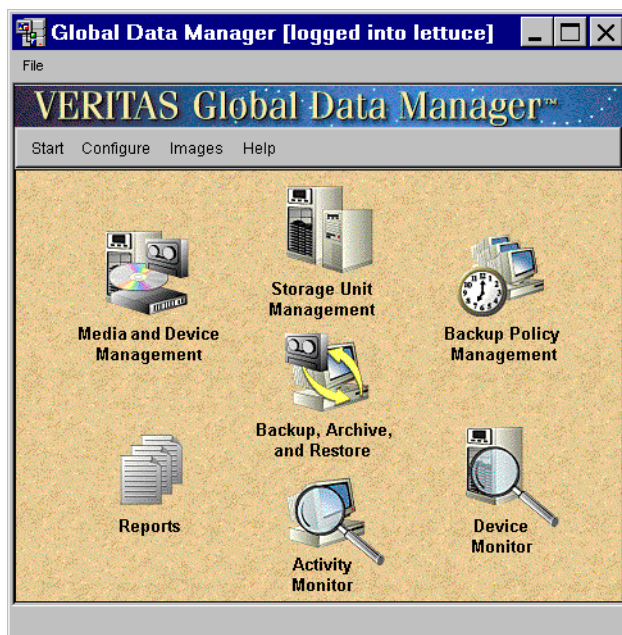
This logs you into GDM on the specified server and the Global Data Manager window appears. The interface program continues to communicate through the server you specified for the remainder of this session.

When logging into a Windows NT/2000 server, you must enter both the server's domain and the user name in the form:

domain_name\user_name

domain_name specifies the domain of the NetBackup host and is required only if the NetBackup host is a member of a domain.

4. Start the desired utility by clicking its icon or selecting its command from a menu in the Global Data Manager window.



For a description of all available commands and instructions on performing specific operations (such as adding classes) see the online help or the other chapters in this manual. The next topic in this chapter describes commands that apply only to GDM.



Interface Commands for GDM

The following describes commands that apply only to the GDM interface.

Append Master Server

Note In the Media and Device Management window, the command is Append Media Manager and works the same as Append Master Server(s) in that they append the specified host to the tree.

If the master server that you want to configure does not appear in the tree, you can temporarily add it as follows:

1. Position the pointer over the left pane.
2. Press the right mouse button and on the shortcut menu, click Append Master Server. Then, follow the instructions in the dialog box to select the server to be appended.

The new server now appears in the tree and you can perform the same operations as for other master servers. However, the configuration is not modified and the server is not in the tree the next time you start the NetBackup Administration interface.

For the other server to be appended and recognized, it must be operational. In addition, the Master of Masters must be in the NetBackup server list on the other server. If necessary, add the server list entry.

For example, assume you start the GDM interface on a system named bear. The Login window appears and you enter shark in the Host name box (where shark is the Master of Masters to which you want to connect).

Here, you are running the GDM interface on bear, but the operations are executing on shark (the Master of Masters). In instances like this, when you do not run the GDM interface directly on the Master of Masters (where the GDM license is installed), the Master of Masters (shark) is sometimes called the GDM host and the system where you run the GDM interface (bear) is called the GDM display console.

Now assume you append a master server named tiger and select or expand the tiger node in the tree. Also, suppose the GDM interface reports that it cannot load data from tiger. In most instances, the interface has enough information about the cause of the problem to suggest a solution in the message that it displays. Most solutions will be to verify that tiger is operational and its services or daemons are active. If the interface reports that tiger did not permit access from shark, you must add shark to tiger's server list:

- ◆ If tiger is a Windows NT NetBackup system, start the NetBackup administrator's interface on tiger and make the addition on the Servers tab in the Master Server Properties or Media Server Properties dialog box (see the online help on that

server for further instructions). Also, add a `SERVER=shark` entry (or verify that it exists) to the `install_path\Volmgr\vm.conf` file. Then, stop and restart the NetBackup Database Manager and NetBackup Request Manager services on tiger.

- ◆ If tiger is a UNIX NetBackup system, add a `SERVER=host` entry below the existing entries in the `/usr/opensv/netbackup/bp.conf` file on tiger. In this example, *host* is shark. Also, add a `SERVER=shark` entry (or verify that it exists) to the `/usr/opensv/volmgr/vm.conf` file. Then, stop and restart the NetBackup database manager (`bpdbm`) and NetBackup request daemon (`bprd`) on tiger.
- ◆ If tiger is a Backup Exec Windows NT system or NetWare system, see the Backup Exec manuals for instructions on configuring shark as the Global Data Manager host (Master of Masters).

Ignore Master Server(s)

Note In the Media and Device Management window, the command is `Ignore Host(s)` and in the Device Monitor window it is `Ignore Server(s)`. These commands work the same as `Ignore Master Server(s)` in that they set the host to the ignored state.

It is sometimes useful to set a server to the ignored state so the Master of Masters will not attempt connections to it. For example, ignoring a server that you know is down, eliminates waiting for the timeout during refreshes. A server is automatically set to the ignored state if the Master of Masters cannot connect to it.

To set a server to the ignored state:

1. Click the Master of Masters in the tree.
The detail view shows the list of master servers.
2. Position the pointer over the master server name in the detail view.
3. Press the right mouse button and on the shortcut menu left-click `Ignore Master Server(s)`.

The branch of the tree for the ignored server and the associated information in the detail view appears dimmed. When it is in the ignored state, you cannot query the server or refresh its detail information. The only operation you can perform on an ignored server is to set it to the recognized state by using the `Recognize Master Server(s)` command.



Recognize Master Server(s)

Note In the Media and Device Management window, the command is Recognize Host(s) and in the Device Monitor window it is Recognize Device Server(s). These commands work the same as Recognize Master Server(s) in that they cause the Master of Masters to recognize an ignored host.

To recognize an ignored server:

1. Position the pointer over the master server name in either the tree or detail view.
2. Press the right mouse button and on the shortcut menu left-click Recognize Master Server(s).

The branch of the tree for the recognized server and the associated information in the detail view is no longer dimmed. Refreshes now apply and you can also select the server and change its configuration (for example, add a class).

Using NetBackup With AFS

F

This appendix explains how to install, configure, and use NetBackup to back up AFS file servers (AFS is an acronym for Andrew File System).

Installation

System Requirements

- ◆ AFS file servers that will be NetBackup AFS clients:
 - ◆ Solaris 2.6, HP-UX 10.20 platform
 - ◆ NetBackup 3.2 or later
 - ◆ AFS level 3.4 installed
- ◆ NetBackup servers that will be backing up the clients must have NetBackup 3.2 or later installed.

Installation on the Server

The AFS software on the server is installed automatically with the rest of the NetBackup server software. There are no special instructions.

Installation on the Clients

Client installation consists of distributing the appropriate NetBackup for AFS library to the NetBackup clients.

For example, to distribute software to a Solaris 2.6 NetBackup client machine named `dog`, execute the following (all on one line):

```
rcp /usr/opensv/lib/client/Solaris/Solaris2.6/libvafs.so  
dog:/usr/opensv/lib
```



Configuration

To configure backups for NetBackup AFS clients, add an AFS class to the NetBackup configuration on the master server. Except for the differences mentioned here, the requirements are the same as for other NetBackup classes. To back up files and directories that are not in AFS volumes, create a separate class.

General Class Attributes

When you select the general attributes for the class, specify AFS for the class type.

Client List

In the client list, specify the names of the AFS file servers to be backed up. These systems must have NetBackup client and NetBackup AFS software installed.

File List

In the file list for the AFS class, specify the AFS volumes and (or) vice partitions to be backed up by the schedules in this class. The following example shows both volumes and vice partitions:

```
user.abc
/vicepb
/vicepc/user.*
```

In this instance, NetBackup backs up:

- ◆ The volume `user.abc`
- ◆ All volumes in vice partition `vicepb`
- ◆ All volumes in `vicepc` that begin with `user`.

When the list includes a vice partition, all the volumes in the partition are backed up one at a time.

File List Directives

The following directives can be in the file list in an AFS class:

- ◆ `CREATE_BACKUP_VOLUMES`

This directive causes NetBackup to create `.backup` volumes prior to performing the backup. If a `.backup` volume already exists, NetBackup overwrites it, thus creating a more recent copy.



Because NetBackup backs up only the .backup copy of AFS volumes, this directive is useful if an automated mechanism is not in place to create .backup copies. Creating .backup copies also ensures that the backups include the latest changes.

Caution If an automated mechanism is not in place to create .backup copies, you must include the `CREATE_BACKUP_VOLUMES` directive in the file list or AFS volumes are not backed up.

◆ `SKIP_SMALL_VOLUMES`

This directive allows you to skip small or empty volumes during backups. For example:

```
SKIP_SMALL_VOLUMES=5
```

(do not include spaces on either side of the = sign)

In this example, NetBackup skips volumes that have five or fewer kilobytes. You can specify any number for the size. Although, the practical range is from 2 to 200 000 (2 kilobytes is an empty volume and 2 gigabytes is the maximum allowable size for an AFS volume).

If you do not specify a number, the size defaults to two kilobytes. For example:

```
SKIP_SMALL_VOLUMES
```

The following rules also apply to the directives:

- ◆ They must be all upper case.
- ◆ They can be anywhere in the file list but it is best to place them at the top. For example:

```
CREATE_BACKUP_VOLUMES
SKIP_SMALL_VOLUMES
/user.abc
/vicepb
```

Regular Expressions

NetBackup supports regular expressions in file list entries. These are useful if you want to:

- ◆ Add or move volumes without having to change the file list.
- ◆ Add vice partitions without having to change the file list.
- ◆ Split volumes and (or) vice partitions on AFS file servers into groups that can be backed up by separate classes. This allows concurrent backups or multiplexing.

The following examples use regular expressions:



```
user.[a-m]*  
/vicep[a-c]
```

Exclude and Include Lists

Exclude lists can be created on the client in order to exclude certain specific volumes from automatic backups. An exclude list cannot contain vice partitions but it can contain individual volumes within a vice partition.

An include list adds back volumes specified in the exclude list. For example, if a range of volumes is excluded, the include list can add back specific volumes within the range.

For more information on exclude and include lists, see “Excluding Files From Automatic Backups” on page 119.

Backups and Restores

Backups

Note You cannot perform user backups or archives of AFS volumes.

Automatic Backup

The most convenient way to back up NetBackup for AFS clients is to configure an AFS class and set up schedules for automatic, unattended backups.

Manual Backup

The administrator on the master server can use the NetBackup administration interface to manually execute a backup for an AFS class. For information about manual backups, see “Performing Manual Backups” on page 141.

Restores

All restores must be performed by the administrator either on the NetBackup AFS client or the master server. Restores are performed on the basis of volumes. To restore a vice partition, the administrator must select all the volumes in that partition.

Caution If you select the Overwrite existing files option, the volumes are overwritten and all changes or files created since the last backup are lost.



Restore From the NetBackup for AFS Client

An administrator on a NetBackup AFS client (AFS file server) can use the NetBackup client interface to restore volumes to that client. It is also possible to perform an alternate-path restore. An alternate-path restore will restore a volume to another volume or vice partition.

Restore From the NetBackup Master Server

The administrator can use the NetBackup client interface on the master server to restore volumes to the same or an alternate NetBackup AFS client (AFS file server). This is called a server-directed restore. For instructions, see the *NetBackup User's Guide for - UNIX*.

Notes About Restores

- ◆ On UNIX, the NetBackup Java user interface, `jbpSA`, provides a convenient mechanism for specifying an alternate name for a volume and (or) vice partition. Specifying an alternate volume name prevents an existing volume from being overwritten by the restore.
- ◆ If the administrator does not specify **Overwrite existing files** or an alternate name for the volume, then NetBackup adds an *R* to the name of the restored volume as follows:

- ◆ If the volume name is less than 22 characters long, NetBackup adds a leading *R* to the name of the restored volume. For example:

If the volume name is

```
/AFS/shark/vicepa/user.abc
```

The restored name is

```
/AFS/shark/vicepa/Ruser.abc
```

- ◆ If the volume name is 22 characters long (maximum allowable length for a volume name), the first character of the original volume name is replaced with an *R*. For example:

If the volume name is

```
/AFS/shark/vicepa/engineering.documents1
```

The restored name is

```
/AFS/shark/vicepa/Rngineering.documents1
```

- ◆ If you are restoring to an alternate path and specify an existing volume, you must select the **Overwrite existing files** option for the restore to succeed. In this case, the entire volume is overwritten. If you do not select the **Overwrite existing files** option, the restore fails.



- ◆ When restoring a volume to an alternate vice partition, the vice partition must exist or the restore fails.

Troubleshooting

The following sections provide tips and information for troubleshooting problems with NetBackup for AFS. See the *NetBackup Troubleshooting Guide - UNIX* for overall troubleshooting information.

Troubleshooting Backups

To increase the level of detail in the logs:

- ◆ Add the `VERBOSE` option to the `/usr/opensv/netbackup/bp.conf` file on the NetBackup for AFS client.
- ◆ Create the following activity log directory on the NetBackup for AFS client:

```
/usr/opensv/netbackup/logs/bpbkar
```

If the AFS backup terminates with a status code of 9 (an extension package is needed, but was not installed), it means that NetBackup AFS client software was not properly installed on the client. Refer to “Installation” on page 673 for installation instructions.

If the AFS backup terminates with a status code of 78 (`afs/dfs` command failed), it indicates an AFS `vos` command failure. The NetBackup Problems Report provides additional information on why the command failed. The `bpbkar` activity log shows the command that was executed. You can execute the `vos` command manually to attempt to duplicate the problem.

Also, examine the `/usr/opensv/netbackup/listvol` file on the NetBackup client for irregularities. The `vos listvol` command can be very demanding on system resources so NetBackup caches the output of the `vos listvol` command in this file. If the cached `listvol` file was created less than four hours prior to the backup, NetBackup uses it to obtain the list of volumes instead of executing another `vos listvol` command.

Troubleshooting Restores

If the restore of an AFS volume fails, check the restore process log for additional information. If a `vos restore` command failure is indicated, create a `/usr/opensv/netbackup/logs/tar` activity log directory, retry the operation, and check the resulting log to see the `vos restore` command that was executed.

NetBackup Intelligent Disaster Recovery (IDR) for Windows NT/2000 is a fully-automated disaster recovery solution that allows you to quickly and efficiently recover your Windows NT/2000 computers after a disaster. The IDR wizards guide you in preparing for disaster recovery and in recovering your computer to its pre-disaster state.

All information in this appendix applies to both Windows NT and Windows 2000 unless otherwise specified.

This appendix contains the following sections:

- ◆ “Requirements for IDR” explains the prerequisites for using IDR.
- ◆ “Overview of IDR Use” explains the main steps involved in using the disaster recovery software.
- ◆ “About the DR Files” introduces the DR (Disaster Recovery) files and explains their importance in disaster recovery.
- ◆ “Configuring NetBackup Classes for IDR” explains how to configure classes that contain clients that are using IDR.
- ◆ “Preparing the IDR Bootable Media” explains how to use this wizard to prepare the bootable media that is used to recover your data.
- ◆ “Updating IDR Media” explains how and when to update the IDR media so it is always ready when you need it.
- ◆ “Recovering Your Computer” explains how to perform disaster recovery.
- ◆ “Notes on Recovering Specific Platforms” provide information on recovering data on specific types of platforms.
- ◆ “IDR’s Frequently Asked Questions” answers questions that are frequently asked about IDR.

Supported Windows NT/2000 Editions

IDR allows you to recover Windows NT 4.0 Enterprise Server, Small Business Server, Terminal Server, and Workstation editions with Service Pack 3 or later. You can also recover Windows 2000 Server, Advanced Server, and Professional.



Requirements for IDR

- ◆ NetBackup 3.4 for Windows NT/2000 or later must be installed on the NetBackup server and clients that you are going to protect.
- ◆ The Intelligent Disaster Recovery software must be installed on the client and also on the master server if it is Windows NT/2000. The IDR software is provided with your NetBackup server software.
- ◆ Intel platform running Windows NT 4.0 (with Service Pack 3 or later) or Windows 2000.
- ◆ At least 40 MB of hard drive space to hold the minimal recovery system.
- ◆ Sufficient space for the data that is being restored.
- ◆ Sufficient swap space to support your system's RAM.

For example, if you have 128 MB of RAM, the minimum swap used is 128 MB. For a 2 GB partition storing 1.8 GB of data, the required hard drive space for that partition is 1.8 GB plus 128 MB plus 40 MB, for a total of 1.97 GB.
- ◆ The partition on the first physical drive must be the boot partition and must also be labeled `c:\`.
- ◆ A protected computer must use a network card that does not require a Windows NT/2000 service pack to be installed in order to work. For a list of cards that have passed Microsoft compatibility tests without service packs, see the "Network LAN Adapters" section of the "Hardware Compatibility List" that comes with the Microsoft Windows NT/2000 software.
- ◆ The driver required by the CD-ROM on a protected computer must be supported by Windows NT/2000 in order to use Intelligent Disaster Recovery. A possible workaround is to choose Use SCSI Drivers Currently installed on this system when prompted by the IDR preparation wizard about the SCSI Drivers (assuming that the CD-ROM driver in question is a SCSI miniport driver).

Overview of IDR Use

Using IDR involves the following steps:

- ◆ Installation:

The IDR software is provided with your NetBackup server software.

 - ◆ The IDR software must be installed on each client that will be protected.
 - ◆ The IDR software must be installed on master servers running Windows NT/2000 (but not on UNIX master servers).

- ◆ Configuration. On the NetBackup server, the class configuration for protected clients must be set up to collect disaster recovery information.
- ◆ Preparing the bootable media. The IDR preparation wizard guides you through the preparation of bootable media used to recover protected systems.
- ◆ Backup. Back up your data files frequently.
- ◆ Recovery. A wizard guides you through the steps for restoring data to a protected system.

The installation, configuration, preparation, and backup steps are prerequisites for successfully recovering a Windows NT/2000 system through a network connection to a NetBackup server

About the DR Files

The DR files are mentioned frequently in this appendix and in the screens that you see in the wizards. A DR (Disaster Recovery) file contains specific information about the computer you are protecting, including:

- ◆ Hard disk partition information
- ◆ Network interface card information
- ◆ NetBackup configuration information required to restore data files

To fully automate the recovery of an IDR protected computer, you need a copy of the DR file for that computer. If IDR software is installed on the server and client, NetBackup creates a DR file and stores a copy on the client and the master server after every:

- ◆ Full backup
- ◆ Incremental (differential or cumulative) backup
- ◆ User backup
- ◆ User archive

NetBackup stores the DR file for each client in the *install_path*\NetBackup\ldr\data directory on the client. The DR files generated after a backup are named in the format *netbackup_client_name.dr*. For example, if the client name is bison, the DR file is *bison.dr*.

Note IDR requires that the DR file name match the computer name of the client. That is, if the computer name is recognized by the network as bison, then the DR file must be named *bison.dr*. If the NetBackup client name is different for some reason, you must manually rename a DR file created after each backup to *computer_name.dr* before you can use it in a recovery.



If a full backup has been performed for a client, you can also run `install_path\NetBackup\bin\drfile.exe` on the client itself to create or update the client's DR file. In this instance the name of the resultant DR file always matches the computer name of the client (which is the name required by IDR), even if this name happens to be different than the one used in the NetBackup class configuration. For more information on this method, see "Using drfile.exe To Create or Update a DR File" on page 692.

On the master server, the DR files for all clients are stored in the NetBackup catalogs on the server. However, you can run the IDR preparation wizard on the master server, a media server, or a NetBackup administration client, and choose the option that places all client DR files in the server's `install_path\NetBackup\Idr\data` directory. This allows you to easily obtain the latest copy of a DR file if a client fails and you did not get the latest DR file prior to the failure.

Other sections of this appendix and the wizards provide more information on using the DR files.

Configuring NetBackup Classes for IDR

Set up the configuration on the NetBackup master server as follows:

- ◆ Ensure that each protected client is in an MS-Windows-NT type class.
- ◆ Select the Collect Disaster Recovery Information attribute for at least one of the MS-Windows-NT classes that is backing up protected clients.
 - ◆ If the master server is running Windows NT/2000, ensure that IDR is installed on that server. Otherwise, you cannot select the Collect Disaster Recovery Information attribute.
 - ◆ Ensure that all the clients in this class have IDR installed. If a client in a class that is collecting disaster recovery information does not have IDR installed, backups performed for that client by this class can never end with a status of 0. A successful backup in this instance shows a status of 1 (partially successful). This is a result of NetBackup not finding a DR file to store in its catalog after each backup.
 - ◆ Ensure that the client names used in the NetBackup class configuration match the client's computer name. If these names do not match, you must manually rename the DR file that is created after each backup to `computer_name.dr` before you can use it in a recovery.

See the Chapter 3 for instructions on modifying classes.

Preparing the IDR Bootable Media

The IDR preparation wizard guides you in creating the bootable media required for recovering a Windows NT/2000 computer. To use this wizard, you must have:

- ◆ The Windows NT/2000 installation CD for the version and language installed on the protected system.
- ◆ Administrative privileges for the protected system.
- ◆ One of the following for media:

For Windows NT, your choice of the following:

- ◆ Diskettes
- ◆ CD-R (CD Recordable CD-ROM)
- ◆ CD-RW (CD Rewritable CD-ROM)

For Windows 2000:

- ◆ Diskettes (CD-ROM bootable media is not supported)

More information on media is provided later.

You must prepare your bootable media before a disaster. For CD-R or CD-RW (Windows NT only), also try booting from the media before a disaster occurs to ensure that your hardware can boot from it (see “Step 1: Boot Your Computer” on page 693).

If an IDR protected NetBackup client is available, you can prepare IDR bootable diskettes on an emergency basis. However, if the DR files are not available, it may be necessary to manually repartition your hard drives, manually install networking, and manually submit the restore request.

Choosing the Bootable Media

For Windows 2000, the IDR preparation wizard can create bootable diskettes. For Windows NT, the IDR preparation wizard, can create both bootable diskettes and bootable CD-Recordable (CR-R) or CD-Rewritable (CR-RW) media.

When choosing between diskettes and CD-ROM media for Windows NT, consider the type of Windows NT system or systems you are protecting, the available hardware, and your system BIOS.

- ◆ Diskettes work on most systems but require more time for preparation and recovery. The Windows NT installation CD is also required during recovery.
- ◆ Diskettes will hold SCSI driver information for only one computer (due to space limitations).



This means you must pick a computer that represents the set of computers to be protected and then create the bootable media for that computer. This works fine but if you have a variety of driver configurations it means creating a set of diskettes for each variation.

CD media has enough space to add SCSI driver information for multiple systems, so you can use a single CD for multiple systems during disaster recovery.

With both diskettes and CDs, you must prepare separate media for each operating system level and language being used.

- ◆ CD media is restricted to recovering Windows NT 4.0 systems.
- ◆ CDs require less time for preparation and recovery than diskettes. However, they also need:
 - ◆ BIOS that supports booting from a CD
 - ◆ Third party CD writing hardware and software for writing ISO 9660 CD images

Creating Bootable Diskettes

If you select diskettes for the bootable media, you need four (for Windows NT) or five (for Windows 2000) blank, formatted 1.44 MB diskettes for each set of disaster recovery diskettes. You do not need a separate set of disaster recovery diskettes for each computer. However, a separate set is required for each version and language of Windows that is used.

In each set:

- ◆ One is the diskette on which the IDR preparation wizard copies the Disaster Recovery wizard that guides you through an actual disaster recovery.
- ◆ The rest are Windows NT/2000 Setup diskettes. These are initially created by a utility that is on the Windows NT/2000 installation CD. IDR then modifies these setup diskettes for use specifically with NetBackup for Windows NT/2000.

Note The Windows NT/2000 installation CD is needed both to prepare disaster recovery diskettes and for disaster recovery using those diskettes.

▼ To Create Bootable Diskettes

1. Format the diskettes that you are going to use (four for Windows NT or five for Windows 2000).
2. On the computer where you are going to prepare the diskettes, go to the *install_path*\NetBackup\bin folder and double-click *drprepwizard.exe* (*install_path* is C:\Program Files\VERITAS by default).

The Welcome screen for the IDR preparation wizard appears.

3. Click Next to continue.
The Create or Update IDR Boot Media screen appears.
4. Select Diskettes to boot the Windows NT or Windows 2000 setup CD and click Next.
The Intelligent Disaster Recovery Preparation - Diskettes screen appears.
5. Read this screen and click Next to continue.
The Create or Update IDR Boot Media screen appears.
6. Select Yes and click Next to continue.
The Select Machine For Diskette Preparation screen appears.
7. Specify the name of the computer for which the disaster recovery diskettes are being created and click Next.
The Enter Windows NT CD Path screen displays.
8. Place the Windows NT/2000 installation CD into the CD-ROM drive.
9. Enter the path of the install directory that is on the Windows NT/2000 installation CD (for example, D:\i386) and then click Next.
The default path shown is the path from which you installed Windows NT/2000 on the computer where you are preparing the diskettes. To use a different path, click Browse and select a directory.
If you receive a message that an invalid path was specified, make sure the Windows NT/2000 installation CD is in the CD-ROM drive, and try again.

Note If you are creating the diskettes on a Windows NT computer perform step 10. If you are creating the diskettes on a Windows 2000 computer, proceed to step 11.

10. If you are on a Windows NT computer and entered the correct path in the previous step, the Windows NT Upgrade/Installation screen appears.
 - a. Click Continue and follow the Windows NT instructions for creating and labeling the setup diskettes.
 - b. During the SCSI device detection phase of this utility, if a SCSI driver version is detected on the selected computer that is different from the version on the Windows NT installation CD, you are prompted to use either the default SCSI drivers that are on the Windows NT installation CD or the SCSI drivers that are installed on the selected computer.
 - ◆ If you are creating disaster recovery diskettes for multiple computers or for a computer other than the one you selected, select Use Default SCSI drivers that are available on the inserted CD.



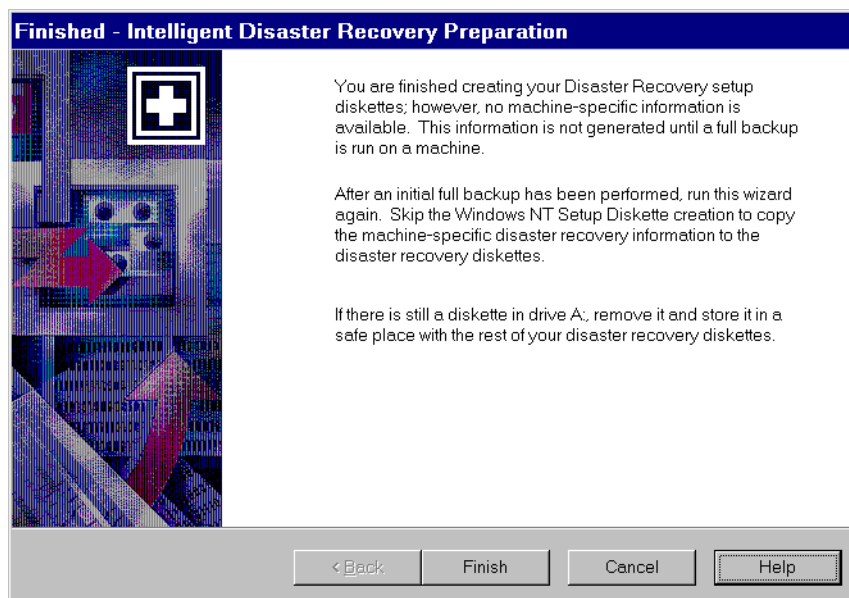
- ◆ If you are creating disaster recovery diskettes for the selected computer, select the default, Use SCSI Drivers currently installed on this system. This option adds the drivers currently installed on the selected computer to the Setup diskettes.

Caution It is strongly recommended that you use the SCSI drivers currently installed on the computer being protected because the drivers contained on the Windows NT CD-ROM may not be up to date. If you have an IDE hard disk greater than eight GB you must select Use SCSI Drivers currently installed on this system.

After making your selection, click Next to continue.

- c. You are now prompted to insert the Windows NT setup diskettes into the drive so information can be updated for the Disaster Recovery wizard. Insert them as prompted and click Next. After the last screen proceed to step 12.
11. If you are on a Windows 2000 computer and entered the correct path in the previous step, the Create Diskettes screen appears. Follow the prompts on this and subsequent Create Diskettes screens to create the Windows 2000 setup diskettes.
12. The last prompt tells you to label a blank formatted diskette as Intelligent Disaster Recovery Diskette, insert it into the drive, and click Next. What happens next depends on whether a DR file already exists on this computer.
 - ◆ If a DR file already exists on this computer skip the remainder of this procedure because it does not apply. Instead, the wizard continues and updates the diskettes as explained in “Updating Disaster Recovery Diskettes” on page 691. Go to that procedure for more information.

- ◆ If a DR file does not exist on this computer, when the IDR Preparation wizard has copied the necessary drivers and the Disaster Recovery wizard to this diskette, the Finished - Intelligent Disaster Recovery Preparation screen appears.



13. If there is still a diskette in drive A, remove it and store it with the rest of the disaster recovery diskettes.
14. Click Finish.

The Disaster Recovery setup diskettes are now complete, except for adding the DR file.

15. Create a DR file for the target computer by running an initial full backup of the entire hard drive (not just the individual directories).

If a full backup has already occurred, you can execute the `drfile.exe` command to create a DR file without waiting for the next backup (see “Using `drfile.exe` To Create or Update a DR File” on page 692). For more information on the DR file, also see “About the DR Files” on page 681.

16. After the initial backup has been performed, run the IDR preparation wizard again to update the Disaster Recovery preparation diskettes with the DR file. For details, see “Updating Disaster Recovery Diskettes” on page 691.

Creating a Bootable CD Image (applies only to Windows NT)

The following are the requirements for using a bootable CD as your disaster recovery media:



- ◆ Computer(s) to be recovered must be able to boot from a CD.
- ◆ Writable (or re-writable) CD device.
- ◆ Third party software that will burn an ISO 9660 image.
- ◆ Windows NT installation CD. The Windows NT operating system version and language on this CD must match those on the computers being protected. If there is more than one operating system level or language, you must create a CD for each variation.

The IDR preparation wizard guides you through the creation of the CD image. You must then use your CD writing system to burn a CD.

To recover your computer with the bootable CD, you need:

- ◆ The bootable CD
- ◆ The latest copy of the DR file for the computer (required for automated recovery)
- ◆ The latest backup images

▼ To Create a Bootable CD Image

1. On the system where you are preparing the media, go to the *install_path*\NetBackup\bin folder and double-click *drprepwizard.exe* (*install_path* is C:\Program Files\VERITAS by default).

The Welcome screen for the IDR preparation wizard appears.

2. Click Next to continue.

The Intelligent Disaster Recovery Preparation Options screen appears.

3. Select CD Image for use with CD Writers and click Next.

The Intelligent Disaster Recovery Preparation - Writable CD screen appears.

4. Read this screen and click Next.

The Select Machine(s) to Protect screen appears.

5. Select the computers for which the bootable CD image is being created. The wizard will gather the SCSI driver information from these computers.

Note All computers selected must be running the same version of Windows NT.

6. Click Next.

If there are different driver versions on the selected computers, the Driver Versions Do Not Match screen appears.

- a. Select which driver version is to be put into the bootable image. Options include:
 - ◆ First version found. Hardware installation diskettes may be required for the highlighted computers.
 - ◆ Latest date time stamp. The newest drivers may not work on older hardware.
 - ◆ Only drivers from Windows NT setup media. Hardware installation diskettes can be required for some computers.
- b. Click Next.

The Select Location for CD Image screen appears.

7. Select the directory where the ISO 9660 CD image file is to be placed.

Note Most CD writer software requires that the image be placed on the same computer as the CD writer software to prevent data underrun problems during burn.

8. Click Next.

The Windows NT Installation Media screen appears.

9. Insert the Windows NT installation CD and specify the drive.

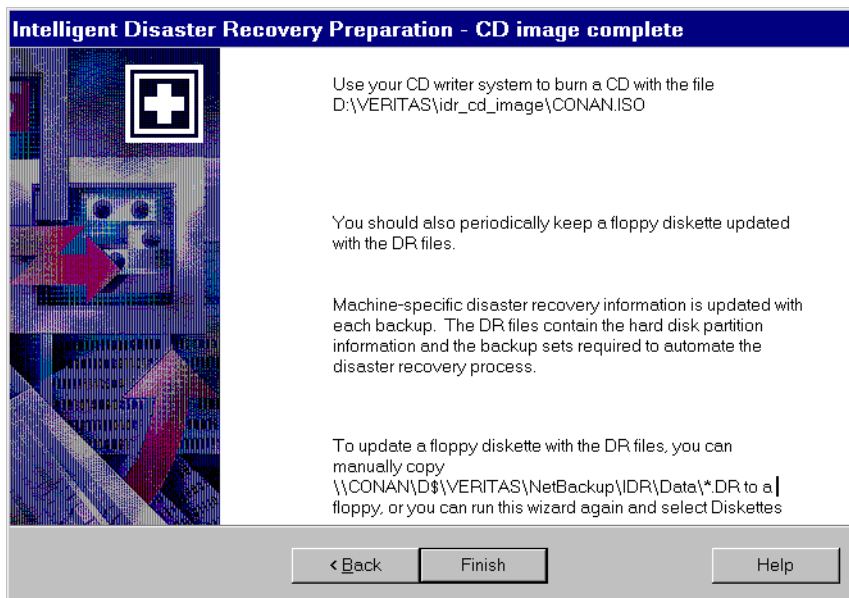
10. Click Next.

The Creating Disaster Recovery Image screen appears. When the creation of the bootable image is complete, a Done status and a Next button appears.

11. Click Next.

The Intelligent Disaster Recovery Preparation Complete - CD Image Complete screen appears.





12. Click Finish.

You must now use your CD writing system to burn a CD. Many popular CD-RW systems ship with both Adaptec Direct CD and Easy CD Creator. Easy CD Creator handles ISO 9660 image files.

Caution Test your bootable CD to ensure that your system can boot from it (see “Step 1: Boot Your Computer” on page 693).

Updating IDR Media

Update your IDR media in the following instances:

- ◆ If your hardware configuration changes.
- ◆ If you are using bootable diskettes, update them after the first full backup that is performed after you created them. This is necessary to add the DR file. It is also recommended that you update them after each subsequent backup in order to maintain the latest DR file.

Updating Disaster Recovery Diskettes

You can update disaster recovery diskettes with the latest DR file by using the IDR preparation wizard. Another method is to run `drfile.exe` from a command prompt and then copy the DR file to the diskette (see “Using `drfile.exe` To Create or Update a DR File” on page 692). The following explains how to use the wizard.

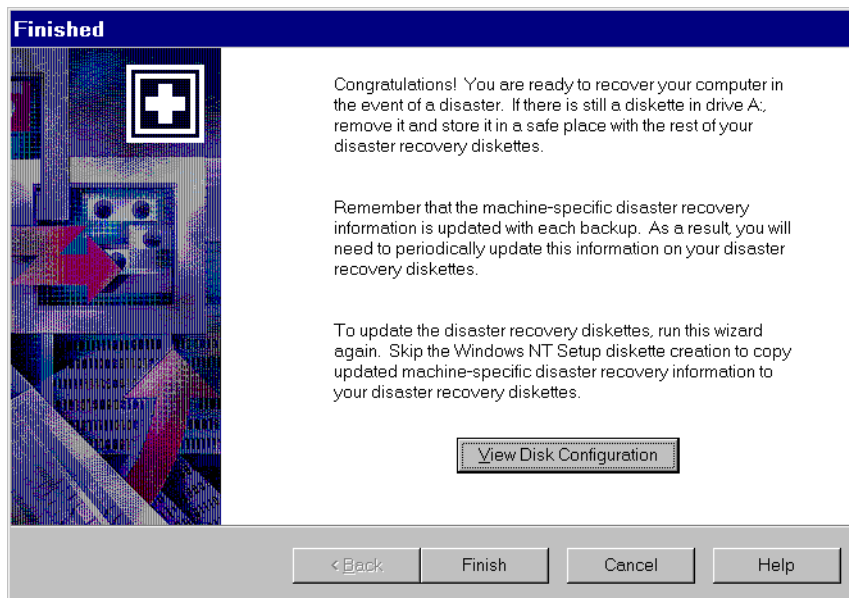
▼ To Update Diskettes By Using the IDR Preparation Wizard

1. Start NetBackup.
2. Run a full backup of the target computer (or wait until NetBackup has completed a scheduled automatic backup).
3. Go to the `install_path\NetBackup\bin` folder and double-click `drprepwizard.exe` (`install_path` is `C:\Program Files\VERITAS` by default).
4. Click Next in the Welcome screen.
The Create or Update IDR Boot Media screen appears.
5. Select Diskettes to boot the Windows NT or Windows 2000 setup CD and click Next.
The Intelligent Disaster Recovery Preparation - Diskettes screen appears.
6. Click Next.
The Create or Update IDR Boot Media screen appears.
7. Click No, I want to update my existing diskettes, and then click Next.
The Update Disks screen appears.
This screen displays the names of the computers that have had a DR file generated by NetBackup during a backup of the computer's hard disk.
8. Select the computer(s) whose DR file you want to store on this set of disaster recovery diskettes and click Next.
The Create Diskettes screen displays.
9. Insert the Windows NT/2000 Setup boot disk and click Next.
This diskette is updated and the Update Disk screen appears.
10. Insert the Intelligent Disaster Recovery diskette into drive A and click Next.
The DR file(s) are written to the diskette.



Note You can store the DR files for all the protected computers on one diskette. You can also store the DR files on a diskette other than the Intelligent Disaster Recovery diskette. In this instance, insert the other diskette when prompted for the DR file during disaster recovery. When you are done, label the diskette and keep it with your set of disaster recovery diskettes.

When the writing is done, the Finished screen appears.



11. Click Finish.

Disaster recovery preparation is now complete.

For subsequent disaster recovery, you will need the disaster recovery diskettes, the Windows NT/2000 installation CD, and the latest backups on the NetBackup server.

Updating a Disaster Recovery CD

If you install a new SCSI card that is not supported by the Windows NT/2000 installation CD, create a new bootable CD as explained in the previous procedure, “Creating a Bootable CD Image (applies only to Windows NT)” on page 687.

Using drfile.exe To Create or Update a DR File

If at least one full backup has been performed, you can run `drfile.exe` to create or update the DR file on your computer.

1. Go to the *install_path*\NetBackup\bin folder and double-click *drfile.exe* (*install_path* is C:\Program Files\VERITAS by default).

This creates (or updates) the DR file that is located in the *install_path*\NetBackup\ldr\Data directory on your computer. Note that the Data directory must exist on the client or the *drfile.exe* command will fail.

The DR file name is of the form *computer_name.dr* as in *bison.dr*. The name of the resultant DR file always matches the computer name of the client (which is the name required by IDR), even if this name happens to be different than the one used in the NetBackup class configuration.

2. Insert the Intelligent Disaster Recovery diskette into your drive and copy the DR file to it or use the wizard as explained earlier.

Note You can also copy the DR file to a diskette other than the Intelligent Disaster Recovery diskette. In this instance, insert the other diskette when prompted for the DR file during disaster recovery.

Recovering Your Computer

Restoring the NetBackup client to its pre-disaster status with IDR includes the following steps:

- ◆ Booting the computer by using previously prepared IDR bootable media.
- ◆ On Windows NT, using the Windows NT Setup program and its Express Setup option to install a minimal version of Windows NT on the system. This step does not apply to Windows 2000.
- ◆ Using the NetBackup IDR Disaster Recovery wizard to restore your system to its pre-disaster state and restore your data files.

Step 1: Boot Your Computer

You can recover a Windows NT/2000 system by using the bootable diskettes or CD created during disaster preparation. The computer being recovered must have a storage device capable of booting from the bootable media.

Note If you have not previously prepared bootable media for a failed client, you can create bootable diskettes on an emergency basis as described in “Preparing Emergency IDR Bootable Diskettes” on page 698.



▼ **To Boot a Computer Using a Bootable Diskette**

1. Insert the bootable diskette.
2. Follow screen instructions.

You will need the Windows NT/2000 installation CD that was used during the preparation of the disaster recovery diskettes.

Note Windows 2000 will require you to log in prior to the time that the disaster recovery wizard appears. For this login, type administrator for the user name. A password is not required.

3. See “Step 2: Run Windows NT Setup (applies only to Windows NT)” on page 694 for additional steps.

▼ **To Boot From a Bootable CD (applies only to Windows NT)**

1. Insert the bootable CD.
2. Follow screen instructions.

Note If you are testing your bootable media, do not continue when the magenta colored NetBackup Intelligent Disaster Recovery Bootstrap screen appears. Instead, remove the CD and press Escape.

3. Press Enter to begin the Windows NT Setup (see “Step 2: Run Windows NT Setup (applies only to Windows NT)” on page 694 for additional steps).

Step 2: Run Windows NT Setup (applies only to Windows NT)

Note This step does not apply to Windows 2000.

Windows NT Setup installs a minimal operating system and also can reformat or repartition your hard drive to prepare space for the recovery system. The Windows NT Setup process is similar for both bootable diskettes and CD.

▼ **To Use Windows NT Setup**

1. If using diskettes for the recovery:
 - a. Replace the preparation diskettes as prompted.
 - b. Place the Windows NT installation CD in the computer’s CD-ROM drive when prompted, and then press Enter.
2. When the Windows NT setup instructions appear, press Enter to choose Express Setup.

Note Express Setup is usually the best choice. Use Custom Setup if the following conditions exist:

- SCSI drivers are not present on the boot media.
 - You have RAID hardware that needs to be reconfigured.
-

3. If a new hard drive is detected on your system, select a file system (FAT or NTFS) to format it, and then press Enter.
-

Note When asked to create a partition on a replacement disk, be sure to select FAT format for the C drive. IDR cannot repartition to the old layout if you build the partition as NTFS.

4. Ensure no diskettes or CDs are in the drives and press Enter to reboot the system. After the reboot, the Disaster Recovery Wizard automatically starts.

Step 3: Run the Disaster Recovery Wizard

To fully automate the recovery with the Disaster Recovery wizard, you need:

- ◆ A NetBackup server that can restore the latest backups to the computer being recovered.
- ◆ The latest DR file for the machine you are recovering.

If you booted from a bootable CD, the DR file on that media was created when you ran the IDR Preparation Wizard and can contain out-of-date hard disk partition, network-interface-card driver, or backup set information.

The DR file can also be out of date on bootable diskettes, unless you recently updated it.

The latest DR files (required to fully automate the recovery) are stored on both the client and on the NetBackup server that did the last backup (see “About the DR Files” on page 681). If necessary, you can copy either of these DR files to a diskette and use it for automated recovery. The other option is to not use a DR file and perform a manual recovery as explained in the procedure.

▼ To Use the Disaster Recovery Wizard

1. When the wizard appears, insert the diskette that contains the latest DR file into drive A and click Next. If you do not intend to use a DR file, just click Next.
2. If you have a DR file, select the DR file for the computer you are recovering and click Next.

The name of a DR file matches the computer for which it was created. For example, if the computer is named carrot look for a file named `carrot.dr`.



Note If you do not have a DR file, click Next anyway to proceed. A message stating that the recovery file was not selected appears. Click Yes to continue in manual mode.

3. For Windows NT (not Windows 2000), if your hard disk partition layout changed, the wizard prompts you to either replace the current hard drive partition with the partition information contained in the DR file or to keep the current hard drive partitions.
 - a. The next wizard screen allows you to run the Windows NT Disk Administrator program. Disk Administrator allows you to make additional changes to your partition information, if necessary (for more information on Disk Administrator and fault tolerant configurations, refer to the Windows NT Server 4.0 Resource Kit).
 - b. To make partition changes, click Run Disk Administrator (also see “Notes on Altering Hard Drive Partition Sizes” on page 698). Otherwise, click Next to continue the recovery process.
4. Select either Automatic Restore or Manual Restore.
 - ◆ If you selected Automatic Restore, click Finish to complete the network installation and go to step 12 to continue the recovery.
 - ◆ On Windows NT, if you selected Manual Restore, select Wired to the Network, click Next, and proceed to step 5.
 - ◆ On Windows 2000, if you selected Manual Restore, a message box appears instructing you to manually configure the network connection. To do this, click the Start button on the Windows 2000 task bar and then the Network and Dial-Up Connections command on the menu that appears. Then, select the Local Area Connection icon for your network card and configure the properties for that connection according to the requirements for your network. When the network connection is configured, click the OK button in the IDR message box and proceed to step 12.
5. To select your network adapter, do one of the following:
 - ◆ If your network adapter requires a manufacturer supplied setup diskette, click Select from list and then click Have Disk.
 - ◆ If your network adapter does not require a manufacturer supplied setup diskette, either click Select from list or Start search.

A list of network adapters appears.

Note If your network adapter is not listed on the screen that appears, click Select from list, and then click Have Disk add an adapter to the Network Adapter List. For automatic network installation to succeed, the Windows NT setup program must be able to recognize the network interface card being used.



6. The next screen lists the default network protocols. Select the networking protocols used on your network and click Next.
7. Windows NT is ready to install the networking components. Insert your Windows NT installation CD or the IDR bootable CD into the CD-ROM drive and click Next to continue.

Note If additional setup screens that specifically address your network interface card appear, follow the prompts.

8. If TCP/IP is selected as the network protocol, you are prompted to use DHCP. If you do not want to use DHCP, enter a TCP/IP number.

The Windows NT Networking Installation dialog box displays.

9. Click Next to start the network and complete the installation of the networking components.
10. Enter the name of the workgroup or domain for your computer and click Next.

Note VERITAS recommends that you enter the name of a temporary workgroup, rather than the name of a domain. When the recovery is complete, the system will be restored to its original workgroup or domain.

11. Click Finish to complete the network installation and continue with recovery.
12. Select either Automatic or Manual:
 - ◆ If you selected Automatic, click Next and proceed to step 13.
 - ◆ If you select Manual, click Next and proceed to step 14.
13. Select the server from which you want to restore files and then:
 - a. Click Start Restore to submit the restore request to the selected server. You can view the progress of the restore after the server has acknowledged the request.
 - b. After the restore is complete, the Next button becomes available. Click Next and proceed to step 15.
14. Select Start NetBackup Interface to start the NetBackup client interface.

Using this interface, you can make changes to the NetBackup configuration and you also have more control over the restore. The following provides basic instructions (see the NetBackup user's guide for more information on using the interface).

- a. Display a restore window and search the NetBackup server of your choice for files and folders to restore.
- b. Select the files and folders to restore.
- c. Submit the restore request.



Before submitting the request, you can specify that NetBackup restore the data to a different path than the one from which it was backed up (sometimes referred to as an alternate path restore).

- d. View the progress of the restore.

When the restore is complete, close the progress viewer and the NetBackup client interface. If the restore is complete, the Next button is now available. Click Next and proceed to step 15.

- 15. Remove any diskettes from drive A and click Finish to reboot the computer.

Preparing Emergency IDR Bootable Diskettes

If you do not have IDR bootable diskettes for a protected client that has failed, you can prepare a set of diskettes on an emergency basis in either of the following ways:

- ◆ Run the IDR Preparation wizard on the server that has been backing up the client.
- ◆ Copy the failed client's DR file from the server to another similarly configured client. Then, run the IDR preparation wizard on the other client.

Note Either of the above methods creates a set of diskettes that will usually work on the client you are trying to recover. However, to increase your chances of successfully recovering a client, always prepare IDR bootable media in advance as explained in "Preparing the IDR Bootable Media" on page 683.

Notes on Altering Hard Drive Partition Sizes

Note This section applies only to Windows NT. Reformatting and repartitioning is not supported on Windows 2000.

IDR defaults to restoring the hard drive partition to the same sizes they were before the disaster. There may be unused and unallocated space. If the hard drive in the target computer is larger than the hard drive that was in place before the disaster occurred, run Windows NT's Disk Administrator program (within the IDR Recovery Wizard) to alter the partition sizes to reflect the larger hard drive size.

An example of why you might want to resize your hard drive partitions: If the pre-disaster computer hardware contained a 4 GB hard drive with two 2 GB partitions, and you have replaced it with a 9 GB model, IDR (using the DR file) will rebuild the hard disk partition table using the partition information found on the original 4 GB hard drive. As a result, only 4 GB of space will be allocated on the new 9 GB hard drive, with a partition map consisting of two 2 GB partitions.



To include additional space, use the Disk Administrator program to repartition the hard drive. For information regarding and fault tolerant configurations, please refer to the Windows NT Server 4.0 Resource Kit.

Notes on Recovering Specific Platforms

Recovering the Dell PowerEdge 6100/200 with RAID

Note Although this section specifically deals with restoring a Dell system, the steps outlined can be used with any system requiring the use of third party drivers.

Recovering a Dell PowerEdge 6100/200 with RAID configuration is different than recovering a regular system with one hard drive.

In order to load Windows NT/2000 on this type of machine, you must manually load the PowerRaid II driver, which is not bundled with the Windows NT/2000 operating system.

After loading the PowerRaid II driver, you must then manually load the Adaptec controller driver. Failure to follow these steps results in Windows NT/2000 not recognizing any hard drive partitions on the system.

▼ Use the Following Steps With Your IDR Recovery Diskette Set

1. When the Windows NT/2000 blue Setup screen appears after booting with the IDR boot diskette, press and hold down the F6 key.
Windows NT/2000 then prompts for IDR diskette 2.
2. Insert IDR diskette 2 and press and hold the F6 key again.
After loading additional drivers, a Setup screen appears, allowing you to specify additional devices.
3. Release the F6 key and then press the S key.
4. Follow the on-screen instructions to load the PowerEdge RAID II controller software.
5. After loading the PowerEdge RAID software, press S again to specify loading another device.
6. Follow the on-screen instructions to load the Adaptec controller software next.
7. After loading both pieces of third party software, press Enter and proceed as normal to recover your system.



Recovering IBM Computers

If you are using an IBM computer and the drive containing the system's configuration information fails, you must reconfigure the system using the IBM Reference Diskette prior to running recovery.

Recovering Compaq Computers

If you are using a Compaq computer and the drive containing the System Configuration Partition fails, Intelligent Disaster Recovery will recreate the partition on the new hard disk; however, you must use the Compaq SmartStart utilities to update the system partition.

IDR's Frequently Asked Questions

1. Can I restore boot managers such as System Commander or OS/2 Boot Manager with Intelligent Disaster Recovery for Windows NT/2000?

No, because boot managers usually are installed at a very low level that NetBackup cannot protect.

For example, the OS/2 boot manager resides in its own hard drive partition that NetBackup cannot access. In fact, because of the many different boot managers on the market, an Intelligent Disaster Recovery restore may render your system unbootable, even though your operating system has been restored. In this case, re-installing the boot manager should fix the problem.

2. I ran a full backup of my system but when I run the IDR Preparation Wizard again, I don't see a disaster recovery file. What happened?

For some reason, the DR file was not generated automatically. Generate it manually as explained in "Using drfile.exe To Create or Update a DR File" on page 692.

3. During recovery, the Windows NT/2000 install fails when attempting to load SCSI drivers. When creating the recovery diskettes, I picked Use SCSI Drivers Currently installed on this system when the IDR Preparation Wizard prompted me to choose drivers.

It could be that Windows NT/2000 does not support your drivers. A possible solution is to run the IDR Preparation Wizard on another system in order to create a new set of recovery disks and this time when prompted to choose drivers, pick Use Default SCSI drivers that are available on the inserted CD.

4. Why does the recovery wizard keep complaining that one or more of my hard drives are smaller than the originals?



If this isn't actually the case, the reason may be because the minimal version of Windows NT/2000 that runs the recovery wizard has detected the hard drives in a different order than what was originally configured under the original version of Windows NT/2000 did.

Be sure that your hard drive and controller configuration matches the original configuration before a disaster occurs.

If the original configuration does not match, then to a certain extent, you can control the hard drive numbering scheme that Windows NT /2000 devises. The following chart lists the normal order that Windows NT/2000 uses to assign disk drive numbers. Keep in mind that this chart can change if third party drivers are used.

Windows NT/2000 Hard Drive Numbering Scheme

Primary IDE	Master
	Slave
Secondary IDE	Master
	Slave
SCSI Adapter 0 (In order of the lowest I/O port address)	SCSI ID 0
	SCSI ID 1
	...
	SCSI ID 7 (or 15 is wide SCSI)
SCSI Adapter 1	SCSI ID 0
	SCSI ID 1
	...
	SCSI ID 7 (or 15 is Wide SCSI)
SCSI Adapter <i>n</i>	SCSI ID 0
	SCSI ID 1
	...
	SCSI ID 7 (or 15 is Wide SCSI)

Other types of mass storage controllers are usually seen as SCSI controllers by Windows NT/2000.

Note On Windows NT (not Windows 2000), if you cannot get the IDR Recovery Wizard to properly detect the hard drive order, you can still manually set up hard drive partitions by using the Windows NT Disk Administrator option within the Disaster



Recovery Wizard. After this is done, you can continue with automated restore of your backup media.

If you have drives greater than eight gigabytes and the recovery wizard reports them as being only eight gigabytes, you must create bootable diskettes with the option Use SCSI drivers currently installed on this system.



Glossary

access control list (ACL)

Security information associated with files on some file systems.

ACS

Automated Cartridge System. This robot type is supported only by NetBackup DataCenter servers.

active job

A job for which NetBackup is currently processing backup or restore data.

activity logs

Logs that can be optionally enabled for specific NetBackup programs and processes and then used to investigate problems.

activity monitor

A NetBackup administration utility that displays information about NetBackup jobs and provides limited control over them.

administrator

A user that is granted special privileges to install, configure, and manage the operation of a system, network, or application

administration client

A Windows NT/2000 NetBackup client that has the administration interface software installed and can be used to administer NetBackup servers.

AIT

Sony Advanced Intelligent Tape, a type of tape drive or media type.



alternate-client restore

Restoring files to your client when they were originally backed up from a different client. The administrator using the interface on the master server can direct a restore to any client (this variation is called a server directed restore).

alternate-target restore

On a Novell NetWare server platform running the NetBackup target version of client software, this operation restores files to a different target than the one from which they were backed up.

alternate path restore

Restores files to a different directory than the one from which they were backed up.

archive

A special kind of backup where NetBackup backs up the selected files, and if the backup is successful, deletes the files from the local disk. In this manual, references to backups also apply to the backup portion of archive operations except where otherwise noted.

archive bit

A file-status bit that the Microsoft based operating system sets when it writes a file, thereby indicating that the file has changed.

attributes for a class

Configuration parameters that control the behavior of NetBackup during operations involving this class.

automatic backup

A scheduled backup by the master server.

back up

The act of copying and saving files and folders to storage media.

backup

Refers to the process of copying and saving files and directories to storage media. For example, *the backup is complete*. This term can also refer to the collection of data that NetBackup saves for a client during a backup or archive. For example, *duplicate the backup*.

Backup is two words when used as a verb. For example, *back up the file*.



backup, archive, and restore interface

The name of the NetBackup Microsoft Windows and Java based user interfaces for clients. On servers, these interfaces can be started through the NetBackup Administration interface.

backup window

The period of time during which backups can begin.

block size

The number of bytes in each block of data written on the media during a backup.

bp

A backup, archive, and restore utility for users on NetBackup UNIX clients. It has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

bpadm

An administrator utility that runs on NetBackup UNIX servers. It has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

bp.conf file

A NetBackup configuration file on UNIX servers and also on UNIX, Macintosh, and OS/2 clients.

bp.ini file

NetBackup initialization file for Novell NetWare target clients.

bpccd

NetBackup Client service on Windows NT/2000 and the NetBackup Client daemon on UNIX.

bprd

NetBackup Request Manager service on Windows NT/2000 and NetBackup Request daemon on UNIX.

catalogs

Internal NetBackup and Media Manager databases. These catalogs contain information about configuration, media, devices, status, errors, and the files and directories in the stored backup images.



CDF

Context-dependent file, which is a type of directory structure on a Hewlett-Packard system.

class

Defines the backup policy for a group of one or more clients that have similar backup requirements.

client

The system with the files to back up, archive, or restore.

client-user interface

The program used to perform user backups, archives, and restores.

cluster

See master and media server cluster.

command lines

Commands that users can execute either from the system prompt or in scripts.

compression

The process of compacting data to enable more efficient transmission and storage.

configuration

The parameters that govern the behavior of an application. This term can also refer to the manner in which a network or system is laid out or connected (for example, a network configuration).

cpio

A UNIX command for formatting data on a tape.

ctime

The time that a UNIX inode was changed.

cumulative-incremental backup

A backup that is scheduled by the administrator on the master server and backs up files that have changed since the last successful full backup. All files are backed up if no prior backup has been done. Also see “differential-incremental backup.”



daemon

A program on a UNIX system that runs in the background and performs some task (for example, starting other programs when they are needed). Daemons are generally referred to as services or processes on Windows NT/2000 systems.

database-agent clients

Clients with additional NetBackup software that is designed to back up relational databases.

database-extension clients

See “database-agent clients.”

debug logs

See “activity logs.”

device delays

Delays caused by the device that are beyond the control of the storage application. An example is the time required to position tape under the read and write heads.

device host

A Media Manager host where a drive or robotic control is attached or is defined.

device monitor

A Media Manager administration utility that provides monitoring and manual control of Media Manager storage devices. For example, an administrator or computer room operator can use this utility to manually reset devices or set them to the UP or DOWN state.

DHCP

Dynamic host configuration protocol. This TCP/IP protocol automatically assigns temporary IP addresses to hosts when they connect to the network.

differential-incremental backup

Scheduled by the administrator on the master server and backs up files that have changed since the last successful incremental or full backup. All files are backed up if no prior backup has been done. Also see “cumulative-incremental backup.”

directory depth

The number of levels below the current directory level that the NetBackup interfaces show in their directory and file list displays.



directory tree

The hierarchical structure in which files are organized on a disk. Each directory lists the files and directories that are directly below it in the tree. On UNIX, the topmost directory is called the root directory.

disaster recovery

Recovering data from backups after a disk crash or other catastrophe.

disk

Magnetic or optical disk storage media.

disk-image backup

A bit-by-bit rather than a file system backup of a disk drive on Windows NT/2000.

DLT

Digital-linear tape or tape drive type.

Domain Name Service (DNS)

A program that handles name translation for network communications.

drive cleaning

The use of a special cleaning tape to clean the heads on a drive.

duplicate image

A copy of a backup image.

encryption

Provides additional security by encrypting backup data on the client. This capability is available only with the NetBackup Encryption option.

entry and exit ports

A slot or other opening in a robot where you can insert or remove a tape without having to access the interior of the robot. After inserting a tape, you move it to a slot by using an inject command. Prior to removing a tape, you move it to the port by using an eject command. The inject and eject commands are supported through the add and move screens in the Media Manager administration interface. Entry and exit ports are sometimes called mailslots, or inports and outports.



exclude list

A list that designates files or directories to exclude from automatic backups.

expiration (image)

The date and time when NetBackup stops tracking a backup image.

expiration (volume)

The date and time when the physical media (tape) is considered to be no longer usable.

EVSN

External volume serial number. This is an identifier written on a media cartridge or canister so the operator can identify the volume before inserting it into a drive or robot. For labeled media, the EVSN must be the same as the RVSN (identifier recorded on the media). For all media, the EVSN is the same as the media ID.

FastBackup

A special type of raw-partition backup that can be performed only on an Auspex client (this option is available only for NetBackup DataCenter).

FlashBackup

A special type of raw-partition backup that requires the NetBackup FlashBackup separately-priced option (this option is available only for NetBackup DataCenter).

flush level

Controls how often Netbackup clears its log files on a Novell NetWare or Microsoft Windows client platform.

fragment

A part of a backup or archive image. NetBackup can be configured to divide images into fragments when they exceed a certain size or span tapes.

frequency (backup)

How often NetBackup performs scheduled backups. For example, if the frequency is seven days then backups occur once a week.

FROZEN media state

If a volume is FROZEN, NetBackup keeps it indefinitely and can restore from it but not use it for further backups or archives.



full backup

A backup that copies, to a storage unit, all files and directories that are beneath a specified directory.

FULL media state

If this appears in a report or listing, it indicates the volume is FULL and cannot hold more data or be used for further backups.

global attributes

NetBackup configuration attributes that affect all classes.

Global Data Manager

A separately-priced option (for UNIX servers) that provides an interface with a tree view where the administrator can view and administer multiple master servers. The server where the option is installed is called a Master of Masters.

GNU tar

A public domain version of the UNIX tar program.

goodies directory

A directory containing programs, scripts, and other files that are not formally supported.

gravity stacker

A robot that relies on gravity to advance to the next required tape.

GUI

Graphical user interface.

hard link

On UNIX, a hard link is a pointer to the inode for the data. On Windows NT/2000 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, and a single file can appear in many directories (or even in the same directory with different names).

heap level

A parameter for memory-heap debugging on a Novell NetWare or Windows NetBackup client.



hierarchical storage management

The process of automatically migrating selected files from a managed file system to specified migration levels on secondary storage, while maintaining transparent access to those files.

host

A computer that executes application programs.

host name

Name by which a host computer is identified by programs and other computers in the network.

HSM

See storage migrator.

image

The collection of data that NetBackup saves for an individual client during each backup or archive. The image contains all the files, directories, and catalog information associated with the backup or archive.

import

The process of recreating NetBackup records of images so the images can be restored.

include list

A list that designates files or directories to add back in from the exclude list.

incremental backup

See “cumulative-incremental backup” and “differential-incremental backup.”

inport

See “entry and exit ports.”

inode

A UNIX data structure that defines the existence of a single file.

install_path

Directory where NetBackup and Media Manager software is installed. The default on Windows NT/2000 is `C:\Program Files\VERITAS` and on UNIX it is `/usr/opensv`.



jbpSA

The Java-based NetBackup interface for performing user backups, archives, and restores.

jnbSA

The Java-based NetBackup interface for administrators.

job

A parcel of work submitted to a computer. NetBackup jobs are backups, archives, or restores.

kernel

The nucleus of an operating system.

keyword phrase

A textual description of a backup.

kill a job

Terminating a job and removing it from the job queue.

label

Identifier of a tape or optical disk volume. A recorded label includes a media ID.

A barcode label allows a barcode scanner to be used for media tracking.

library

Refers to a robot and its accompanying software. A library includes a collection of tapes or optical platters used for data storage and retrieval. For example, a Tape Library DLT (TLD) refers to a robot that has TLD robotic control.

link

See “hard link” or “symbolic link.”

LMF - Library Management Facility

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

This robot type is supported only by NetBackup DataCenter servers.



load

(noun) Amount of work that is being performed by a system or the level of traffic on a network. For example, network load affects performance.

(verb) Copy data to internal memory. For example, load the installation program.

logs

Files where a computer or application records information about its activities.

mailslot

See “entry and exit ports.”

man pages

Online documentation provided with UNIX computer systems and applications.

Master and media server cluster

A NetBackup master server and the remote media servers that it is using for additional storage. It is possible to configure clusters only with NetBackup DataCenter servers. NetBackup BusinessServer supports only a single server, the master.

Master of Masters

A NetBackup host where Global Data Manager software is installed. When logging into this host, the interface has a tree view where the administrator can view and administer multiple master servers.

master server

The NetBackup server that provides administration and control for backups and restores for all clients and servers in a master and media server cluster. NetBackup BusinessServer supports only a single server and it is the master.

media

Physical magnetic tapes, optical disks, or magnetic disks where data are stored.

media host

NetBackup server to which the job (client) is sending the data.

media ID

An identifier that is written on a volume as part of the recorded label.



Media Manager

Software that is part of NetBackup and manages the storage devices and removable media.

Media Manager Host

A host where Media Manager software is installed.

media server

A NetBackup server that provides storage within a master and media server cluster. The master can also be a media server. A media server that is not the master is called a remote media server (or slave server). NetBackup BusinessServer does not support remote media servers.

menu interface

A character-based interface for use on terminals that do not have graphical capabilities.

MHD

See “multihosted drives.”

mount

Make a volume available for reading or writing.

mount point

The point where a file system on a disk logically connects to a system’s directory structure so the file system is available to users and applications.

MPX

See “multiplexing.”

mtime

The point in time when a UNIX or NTFS file is modified.

multihosted drives

A separately priced VERITAS option (Shared Storage Option or SSO) that allows tape drives (standalone or in a robotic library) to be dynamically shared among multiple NetBackup and Storage Migrator servers.

This option is supported only on NetBackup DataCenter servers.



multiplexing

The process of sending concurrent-multiple backups from one or more clients to a single storage device and interleaving those images onto the media.

multiplexed group

A set of backups that were multiplexed together in a single multiplexing session.

NDMP

Network data management protocol. NetBackup requires the NetBackup for NDMP separately-priced option to support NDMP.

NetBackup Client service

NetBackup Windows NT/2000 service that runs on clients and servers and listens for connections from NetBackup servers and clients in the network. When a connection is made, this service starts the necessary programs.

NetBackup configuration options

On UNIX servers and on UNIX and Macintosh, clients, these settings are made in the `bp.conf` file. On NetWare target and OS/2 clients, they are in the `bp.ini` file. On Windows NT/2000 servers and Microsoft Windows clients, these settings are called properties and are made through the Backup, Archive, and Restore interface or the Configure - NetBackup window in the administration interface.

NetBackup databases

See catalogs.

NetBackup Database Manager service

NetBackup Windows NT/2000 service that runs on the master server and manages the NetBackup internal databases (called catalogs). This service must be running on the master server during all NetBackup administrative operations.

NetBackup Device Manager service

The NetBackup Windows NT/2000 service that runs on a NetBackup server and starts the robotic control processes and controls the reservation and assignment of volumes. This service runs only if the server has devices under Media Manager control. The process is `ltid`.

NetBackup properties

Same as NetBackup configuration options but are called NetBackup properties on Microsoft Windows platforms.



NetBackup Request Manager service

The NetBackup Windows NT/2000 service that runs on the master server and starts the scheduler and receives requests from clients.

NetBackup Volume Manager service

A NetBackup Windows NT/2000 service that runs on a NetBackup server, allows remote administration of Media Manager, and manages volume information. The process is `vmd`.

NIS

Network information service.

NLM

NetWare loadable module.

NFS

Network file system.

nonrobotic

See “standalone.”

ODL

Optical disk library. This robot type is supported only by NetBackup DataCenter servers.

OSF and Motif

A set of specifications for user-interface design.

outport

See “entry and exit ports.”

partitions

The logical partitions into which a magnetic disk is divided.

patch

A program that corrects a problem or adds a feature to an existing release of software.

path length

Number of characters in a pathname.



pathname

The list of directories in the path to a destination directory or file.

PC clients

NetBackup clients that have Microsoft Windows (NT/2000, 98, 95), Macintosh, or IBM OS/2 operating systems.

peername

The name by which a computer identifies itself when establishing connections to other systems.

port

A location used for transferring data in or out of a computer.

primary copy

The copy of an image that NetBackup uses to satisfy restores. When NetBackup duplicates an image, the original is designated as the primary copy.

privileges

The tasks or functions that a user, system, or application is authorized to perform.

progress report

Log where NetBackup records events that occur during user operations.

proxy restore

A proxy restore allows the user to restore files, that he has write access to, on a machine other than his desktop. The files must be in a backup of the machine to which they are being restored.

QIC

Quarter-inch-cartridge tape.

queued job

A job that has been added to the list of jobs to be performed.

raw-partition backup

Bit-by-bit backup of a partition of a disk drive on UNIX. On Windows NT/2000, this is called a disk-image backup.



rbak

The program that Apollo clients use to read data from tape during a restore.

registry

A Microsoft Windows 2000, NT, 98, and 95 database that has configuration information about hardware and user accounts.

remote media server

A media server that is not the master. Note that only NetBackup DataCenter supports remote media servers. NetBackup BusinessServer supports only a single server, the master.

residence

In Media Manager, information about the location of each volume is stored in a volume database. This residence entry contains information, such as robot number, robot host, robot type, and media type.

resource

A Novell NetWare term that refers to a data set on the target. For example, in DOS, resources are drives, directories, and files. Also see “target service.”

restore

(verb) The act of restoring selected files and directories from a previous backup or archive and returning them to their original directory locations (or to an alternate directory).

(noun) The process of restoring selected files and directories from a previous backup and returning them to their original directory locations (or to an alternate directory).

retention level

An index number that corresponds to a user-defined retention period. There are 10 levels from which to choose (0 through 9) and the retention period associated with each is configurable. Also see “retention period.”

retention period

The length of time that NetBackup keeps backup and archive images. The retention period is specified on the schedule.

root

The highest level directory in a hierarchical directory structure. In MS-DOS, the root directory on a drive is designated by a backslash (for example, the root on drive C is C:\). On UNIX, the root directory is designated by a slash (/).



Also, a UNIX user name having administration capability.

RS-232

An industry-standard interface for serial communications and sometimes used for communicating with storage peripherals.

RSM Interface

Application in Windows 2000 used to manage Removable Storage Manager (RSM) devices.

RSM - Removable Storage Manager

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

Also, a component of the Windows 2000 operating system that manages storage devices.

RVSN

Recorded volume serial number. This is an identifier recorded as part of the label on a volume and used by Media Manager to ensure that the correct volume is mounted. The RVSN is the same as the media ID.

schedules

Controls when backups can occur in addition to other aspects of the backup, such as: the type of backup (full, incremental) and how long NetBackup retains the image.

SCSI

Small computer system interface. This is a type of parallel interface that is frequently used for communicating with storage peripherals.

slave server

See Remote media server.

server directed restore

Using the client interface on the master server to restore files to any client. Only the administrator can perform this operation.

server independent restore

Restoring files by using a NetBackup server other than the one that was used to write the backup. This feature is available only with NetBackup DataCenter.



server list

The list of servers that a NetBackup client or server refers to when establishing or verifying connections to NetBackup servers. On a Windows NT/2000 server and Microsoft Windows clients, you update the list through a dialog box in the interface. On a UNIX server and UNIX and Macintosh clients, the list is in the `bp.conf` file. On NetWare target and OS/2 clients, the list is in the `bp.ini` file.

service

A program on a Windows NT/2000 system that runs in the background and performs some task (for example, starting other programs when they are needed). Services are generally referred to as daemons on UNIX systems.

session

An instance of NetBackup checking its schedules for backups that are due, adding them to its worklist, and attempting to complete all jobs in the worklist. For user backups and archives, a session usually consists of a single backup or archive.

Shared Storage Option (SSO)

See “multihosted drives.”

SMDR

Storage management data requestor, a Novell NetWare program that provides its services transparently to all SMS modules and lets remote and local modules communicate with one another.

SMS

Novell NetWare storage management services.

standalone

A qualifier used with drives and media to indicate they are not associated with a robot. For example, a standalone tape drive is one where you must manually find and insert tapes before using them. A standalone volume is one that is located in a standalone drive or is stored outside of a drive and designated as standalone in the volume configuration.

status code

A numerical code, usually accompanied by a message, that indicates the outcome of an operation.



storage migrator

Refers to the VERITAS Storage Migrator line of hierarchical storage management products for UNIX and Windows NT/2000. These products make extra room on a disk by transparently moving data to other storage and then transparently retrieving the data when it is needed by a user or application.

Storage Migrator is available only for NetBackup DataCenter servers.

storage unit

Refers to a storage device where NetBackup or Storage Migrator stores files. It can be a set of drives in a robot or consist of one or more single tape drives that connect to the same host.

SUSPENDED media state

If a volume is SUSPENDED, NetBackup can restore from it but cannot use it for backups. NetBackup retains a record of the Media ID until the last backup image on the volume expires.

symbolic link

On a UNIX system, this is a pointer to the name of the file that has the source data.

tape format

The format that an application uses to write data on a tape.

tape marks

A mark that is recorded between backup images on a tape.

tape overhead

The space required for data that is not part of the backup images. For example, tape marks and catalogs of what are on the tape are considered overhead.

tape spanning

Using more than one tape to store a single backup image.

tar

Tape ARchive program that NetBackup uses to extract backup images during a restore.

target

See “target service.”



target service

A Novell NetWare service that needs storage management. The SMS views all services (for example, print services, communication services, workstations) as targets.

Target Service Agent

A Target-service agent is a Novell NetWare agent that prepares the target's data for SMS during a backup and for the target during a restore.

TLD - Tape Library DLT

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

TLH - Tape Library Half-inch

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

This robot type is supported only by NetBackup DataCenter servers.

TLM - Tape Library Multimedia

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

This robot type is supported only by NetBackup DataCenter servers.

TL4 - Tape Library 4MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

TL8 - Tape Library 8MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.



timeout period

The period of time that an application has allotted for an event to occur.

TIR

See “true image restore.”

tpconfig

A Media Manager administration utility for configuring devices and is started from the command line. On UNIX, it has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

transfer rate

The rate at which computer information is transferred between a source and a destination.

true image restore

Restores the contents of a directory to what it was at the time of any scheduled full or incremental backup. Previously deleted files are ignored.

TS8 - Tape Stacker 8MM

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

TSA

“Target Service Agent.”

TSH - Tape Stacker Half-inch

A Media Manager designation for a category of robot. For the specific vendor types and models in this category, see the VERITAS support web page at www.support.veritas.com. Choose NetBackup BusinessServer or NetBackup DataCenter in the VERITAS Support Product List and look under support options.

This robot type is supported only by NetBackup DataCenter servers.

user operation

A backup, archive, or restore that is started by a person on a client system.

verbose flag

Configuration file entry that causes a higher level of detail to be written in the logs.



verify

An operation that compares the list of files that are actually on a volume with what NetBackup has recorded as being on it. The data that is on the media is not verified.

vmadm

A Media Manager administrator utility for managing volumes. It runs on UNIX and has a character-based, menu interface that can be run from terminals that do not have X Windows capabilities.

vm.conf

A Media Manager configuration file with entries that include the servers that can manage local devices and default media ID prefixes for media that do not contain barcodes.

volume

Media Manager volumes are logical units of data storage or cleaning capability on media that have been assigned media IDs and other attributes, which are recorded in the Media Manager volume database.

volume configuration

Refers to configuration information that is stored in the Media Manager volume database.

volume database

An internal database where Media Manager keeps information about volumes. All Media Manager hosts have a volume database. However, the database is empty unless the host is designated as a volume database host.

volume database host

The Media Manager host that contains information about the volumes that Media Manager uses in a device. Because NetBackup BusinessServer supports only a single server, the volume database host is always the Media Manager host.

volume group

A set of volumes that are configured within Media Manager to reside at the same physical location (for example, in a specific robot).

volume pool

A set of volumes that are configured within Media Manager to be used by a single application and are protected from access by other applications and users.



wakeup interval

The time interval at which NetBackup checks for backups that are due.

wbak

The program that Apollo clients use to write data on tape.

wildcard characters

A character that can be used to represent other characters in searches.

WORM media

Write-once, read-many media for optical disks. NetBackup BusinessServer does not support WORM media.

Windows Display Console

A NetBackup-Java interface program that runs on Windows 2000, NT, 98, and 95 computers. Users and administrators can start this interface on their local system, connect to a UNIX system that has the NetBackup-Java software installed, and then perform any user and administrator operations that their permissions allow.

xbp

The X Windows-based backup, archive, and restore program for users on NetBackup UNIX clients.

xbpadm

The X Windows-based NetBackup administration interface on UNIX. This interface is available only with NetBackup DataCenter.

xbpmon

The X Windows-based NetBackup utility for monitoring jobs on UNIX. This utility is available only with NetBackup DataCenter.

xdevadm

The X Windows-based Media Manager utility for managing devices on UNIX. This interface is available only with NetBackup DataCenter.

xvmadm

The X Windows-based Media Manager utility for managing media on UNIX. This interface is available only with NetBackup DataCenter.





Index

A

- Absolute Pathname to Directory box 34
- Access control lists (ACLs) 97
- Access control, to a server or client 341, 355
- ACL (access control list) 97
- Activate class 69
- Active job 197
- Activity monitor
 - BPDBJOBS_OPTIONS 205
 - delete completed jobs 198
 - deleting done jobs 203
 - detailed job status 199
 - disable job logging 331
 - job filter 189
 - job list definitions 196
 - jobs database 203
 - jobs.conf file 203
 - kill uncompleted jobs 198
 - monitoring jobs 196
 - refresh display 194
 - saving job information 199
 - select server to monitor 195
 - starting 186
 - using troubleshooting wizard 202
 - window description 186
- Adding
 - catalog backup file paths 153
 - classes 58
 - clients to class 86
 - files to class 90
 - licenses 224
 - schedules 123
 - storage unit
 - disk type 32
 - media manager type 27
 - NDMP type 18
- Adjust time zone 322
- Administrator
 - definition xxxi
 - nonroot 253, 258
- Administrator-defined streaming mode 113
- AFS
 - class configuration 674
 - class type 65
- All log entries report 179
- Allow multiple data streams
 - file-list directives 112
 - set class attribute 82
 - tuning 84
 - when to use 82
- ALLOW_MEDIA_OVERWRITE 327
- ALLOW_MULTIPLE_RETENTIONS_PER_MEDIA 328
- ALLOW_NON_RESERVED_PORTS
 - client 344
 - server 328
- Alternate client restores
 - allowing 211
 - definition 704
 - host.xlate file 607
- Alternate path restore
 - definition 704
- Alternate target restore
 - definition 704
- altnames file 211
- ANSI format, allow overwrite 327
- AOS/VS format, allow overwrite 327
- Apollo wbak class type 65
- APOLLO_RESTORE_TIMEOUT 328
- Append command 670
- Archive bit
 - definition 704
- Archives
 - definition 704
 - overview 5
- Assistant, NetBackup 14



- atime 223, 252, 352
 - Attributes
 - general for a class 63
 - NetBackup global (see Global Attributes)
 - Attributes, class
 - Auspex FastBackup 617
 - Auspex-FastBackup
 - class type 65
 - overview 613
 - auth.conf file
 - capabilities identifiers 259
 - converting old versions 260
 - description 256
 - entries for specific applications 258
 - overview 256
 - Authentication
 - commands 308
 - configuration files 304
 - overview 303
 - procedure 311
 - Authorization (see NetBackup authorization)
 - Auto-discover streaming mode 114
 - Automounted directories 76
 - available_media script 627
- B**
- Backup frequency
 - definition 709
 - effect on priority 128
 - guidelines for setting 626
 - setting 128
 - Backup images, definition 711
 - Backup network drives 73
 - Backup policy management window 53
 - Backup speed (see Transfer rate)
 - Backup status report 177
 - Backup windows
 - definition 705
 - duration
 - definition 131
 - examples 132
 - end time 131
 - specifying 130
 - start time 130
 - backup_exit_notify script 649
 - backup_notify script 648
 - Backups
 - across mount points 77
 - activating class 69
 - AFS clients 676
 - automatic, introduction to 1
 - backup status report 177
 - backup_exit_notify script 649
 - backup_notify script 648
 - balancing load 250
 - best times for user directed 139
 - bpend_notify script
 - UNIX client 654
 - windows client 656
 - bpstart_notify script
 - UNIX client 649
 - windows client 651
 - catalogs, NetBackup 143
 - clients using storage migrator 252
 - definition 704
 - diskfull_notify script 660
 - duplicating 225
 - estimating time required 618
 - frequency (see Backup frequency)
 - full 125
 - import 236
 - incremental
 - overview 628
 - set on schedule 124
 - manual (see Manual backups)
 - media requirements 627
 - multiplexing (see Multiplexing)
 - network drives 73
 - NFS mounted files 75
 - offsite storage 626
 - registry on Windows clients 104
 - session_notify script 661
 - session_start_notify script 661
 - types of 124
 - user directed
 - overview 1
 - schedules 138
 - verify 232
 - Bandwidth limiting (see LIMIT_BANDWIDTH)
 - Block level incremental 82
 - Block size 181
 - Boot managers and IDR 700
 - Booting a computer
 - with IDR bootable media 693
 - bp command 360



-
- bp.conf file 326
 - definition 705
 - personal
 - for UNIX nonroot user 326, 344, 357
 - for UNIX root user 344
 - UNIX client options 344
 - UNIX server options 326
 - bpadm command 362
 - bpadm, using 555
 - backup frequency, specifying 578
 - backup tries global attribute 582
 - backup windows, specifying 579
 - bpdbm, starting with bprd 589
 - bprd, managing 589
 - classes
 - adding 565
 - adding clients 569
 - deleting 567, 568
 - file list 572
 - modify attributes 567
 - schedules 575
 - clients
 - adding clients 569
 - deleting from class 572
 - install software 571, 583
 - compress backup files 566
 - compress image database files 583
 - cross mount points 566
 - display reports global attribute 583
 - file list
 - adding 572
 - changing 575
 - deleting files 575
 - raw partition backups 574
 - wildcard characters 572
 - follow NFS mounts 566
 - global attributes, specifying 581
 - install client software 571, 583
 - keep logs global attribute 582
 - keep TIR Information, set time 582
 - keyword phrase, specifying 567
 - limit jobs per class 567
 - mail notifications global attribute 582
 - manual backups
 - of clients 592
 - of schedules 592
 - maximum jobs per client 582
 - media mount timeout 583
 - menu overview 556
 - mpx
 - specify for schedule 579
 - specify for storage unit 560
 - NetBackup-database backup
 - adding file paths 601
 - automatic 593
 - changing backup attributes 597
 - delete DB Backup ID 600
 - manual 600
 - removing file paths 602
 - notify request daemon of changes 582
 - printing class properties 567
 - priority for class 567
 - reports, displaying 584
 - retention period, specifying 578
 - retention periods, changing 590
 - schedules
 - adding 576
 - display and modify 580
 - starting bpadm 556
 - storage units
 - adding disk type 561
 - adding Media Manager type 558
 - changing attributes 563
 - deleting 563
 - displaying configuration 563
 - for class 567
 - for schedule 579
 - true image recovery
 - setting 566
 - time to keep TIR information 582
 - volume pool
 - for class 567
 - for schedule 579
 - wakeup interval global attribute 582
 - bparchive command 363
 - BPARCHIVE_CLASS 345
 - BPARCHIVE_SCHED 345
 - bpauthsync command 368
 - bpbackup command 371
 - BPBACKUP_CLASS 345
 - BPBACKUP_SCHED 346
 - bpbackupdb command 377
 - bpclassnew command 381
 - bpclients command 388
 - bpclient
 - add clients to catalog 218
 - delete clients from catalog 218
 - list clients in catalog 218



- preventing lists and restores 218
- bpclinclude command 395
- bpclinfo command 400
- bpclsched command 409
- bpclschedrep command 419
- bpconfig command 115, 424
- bpdbjobs activity log 206
- bpdbjobs command 205
- bpdbm
 - running without bprd 210
 - starting automatically 210
 - stopping bpdbm 210
- bpdbm command 434
- bpduplicate command 436
- bpdynamicclient 292
- bpend_notify script
 - UNIX client 654
 - windows client 656
- BPEND_TIMEOUT 329
- bperror command 442
- bpexpdate command 453
- bpimagelist command 459
- bpimmedia command 466
- bpimport command 477
- bplabel command 482
- bplist command 484
- bpmedia command 490
- bpmedialist command 493
- bppls script 210
- bprd
 - managing 208
 - starting automatically 209
 - terminating 209
- bprd command 505
- bprecover command 507
- bprestore command 512
- bpstart_notify script
 - UNIX client 649
 - Windows client 651
- BPSTART_TIMEOUT 329
- bpstuadd command 520
- bpstudel command 526
- bpstulist command 528
- bpsturep command 532
- BPTM_QUERY_TIMEOUT 329
- bpverify command 537
- Busy file processing
 - Windows clients 303
- BUSY_FILE_ACTION 346

- BUSY_FILE_DIRECTORY 347
- BUSY_FILE_NOTIFY_USER 347
- BUSY_FILE_PROCESSING 348
- Busy-file processing
 - configuration overview 297
 - creating action files 300
 - logs 301
 - logs directory
 - busy log 301
 - logs file 301
 - retry file 301
 - modifying bp.conf 298
 - modifying bpend_notify_busy 302

C

- Cache entry, Auspex file list 617
- Catalog backup
 - automatic 144
 - caution for compressing 165
 - compressing image catalog 165
 - configure with the wizard 145
 - delay to compress 272
 - disk path 150
 - file paths
 - adding 153
 - media server 155
 - NT master 154
 - last media used 147
 - manual backup 156
 - media ID 149
 - media server 147
 - media type 148
 - notification script 659
 - overview 143
 - precautions 144
 - recovery 144
 - setting media density 150
 - setting schedules 151
 - space required 162
 - uncompressing 166
- Catalog indexing 221
- Catalogs
 - device 637
 - media 636
 - NetBackup
 - managing 161
 - volume 636
- Cautions
 - AFS backup volumes 675



- alternate client restores 211
- alternate path restore 213
- catalog backup 144
- database compression 165
- notification scripts 648
- overwriting files on AFS 676
- retention time 127, 139
- wildcards in UNIX raw backups 101
- worm retention 127
- CDE (Common Desktop Environment)
 - set up for NetBackup-Java 7
- Changing
 - catalog backup attributes 146
 - catalog backup paths 154
 - class properties 61
 - licenses 224
 - locale 323
 - server
 - for activity monitoring 195
 - for backup policies (classes) 57
 - for catalog backup settings 146
 - for configuring storage units 22
 - for duplicate images 227
 - for import images 239
 - for verify images 233
 - storage unit 35
- Classes
 - overview 4, 39
 - activating 69
 - adding (see Adding)
 - changing (see Changing)
 - client list (see Clients, NetBackup)
 - configuration wizard 53
 - deleting (see Deleting)
 - example 41
 - file list (see File List)
 - parts of definition 40
 - planning 43
 - select server to manage 57
 - setting IFRFR 85
 - setting priority 69
 - type
 - AFS 65
 - Apollo wbak 65
 - Auspex-FastBackup 65
 - DB2 64
 - Extensible-Client 64
 - Flashbackup 65
 - Informix 65
 - Lotus-Notes 64
 - MS-Exchange 64
 - MS-SQL-Server 64
 - MS-Windows-NT 64
 - NCR-Teradata 64
 - NDMP 64
 - NetWare 64
 - Oracle 65
 - OS2 64
 - SAP 65
 - Split-Mirror 65
 - SQL-BackTrack 65
 - Standard 64
 - Sybase 65
 - user 140
 - user schedules 138
- Client backups report 177
- Client name 211
- Client user interface 706
- Client user, definition xxxi
- CLIENT_CONNECT_TIMEOUT 330
- CLIENT_NAME 348
- CLIENT_PORT_WINDOW
 - on client 348
 - on server 330
- CLIENT_READ_TIMEOUT
 - on client 348
 - on server 330
- CLIENT_RESERVED_PORT_WINDOW
 - on client 349
 - on server 331
- Clients, NetBackup
 - definition 1
 - adding to class 86
 - bp.conf options
 - non-UNIX clients 325
 - UNIX clients 325
 - choosing class type 63
 - definition 706
 - deleting from class 62
 - exclude files list 119
 - exclude list 119
 - host names
 - changing 606
 - include files list 122
 - install software
 - (see Install client software)
 - maximum jobs 269
 - secure clients 90



- setting host names 87
 - software 1
 - trusting clients 88
 - Codes, status 177
 - Collect disaster recovery information 86
 - Commands, NetBackup
 - man pages 359
 - Compaq Computers
 - recovering with IDR 700
 - COMPRESS_SUFFIX 349
 - Compression
 - advantages 79
 - disadvantages 80
 - specifications 80
 - Concurrent jobs
 - on client 269
 - per class 68
 - Configuration 326
 - AFS class 674
 - catalog backups 146
 - classes 39
 - global attributes, NetBackup 268
 - host names 604
 - intelligent disaster recovery 682
 - mail notifications 321
 - options for jbp 260
 - storage units 17
 - terminals 608
 - with NetBackup Assistant 14
 - Copy, primary 226
 - cpio format, allow overwrite 327
 - CREATE_BACKUP_VOLUMES 674
 - Cross mount points
 - effect with UNIX raw partitions 77
 - examples 78
 - separate classes for 77
 - setting 77
 - CRYPT_KEYFILE 351
 - CRYPT_LIBPATH 350
 - CRYPT_STRENGTH 350
 - ctime 356
 - Cumulative incremental
 - overview 629
 - select for schedule 125
 - Custom Setup, when to use in IDR 695
- D**
- Daemons
 - bpdbm
 - starting automatically 210
 - starting with bprd 208
 - terminating 210
 - bprd
 - managing 208
 - starting automatically 209
 - terminating 209
 - checking processes 210
 - Data streams
 - (see Allow multiple data streams) 112
 - Database-extension clients
 - add file paths for 110
 - Databases, NetBackup (see Catalog backup)
 - Datatools (see SQL-BackTrack)
 - Dates, setting for locale 323
 - Datetime stamp 631
 - DB2 class type 64
 - dbbackup_notify script 659
 - DBR format, allow overwrite 327
 - dd 642
 - Delay to compress database 272
 - Deleting
 - classes 62
 - clients from a class 62
 - clients from client catalog 218
 - completed jobs 198
 - done jobs 203
 - files from a class 62
 - licenses 224
 - schedules 62
 - storage unit 36
 - Dell PowerEdge 6100/200 with RAID
 - recovering with IDR 699
 - Detailed job status 199
 - Device delays 620
 - Devices (see Storage devices)
 - DHCP 286
 - Differential incremental
 - overview 628
 - select for schedule 125
 - Directives for AFS 674
 - Directives for file list 111
 - Directives templates 91
 - DISABLE_JOB_LOGGING 331
 - DISABLE_STANDALONE_DRIVE_EXTENSIONS 332
 - DISALLOW_BACKUPS_SPANNING_MEDIA 332
 - DISALLOW_CLIENT_LIST_RESTORE 332



DISALLOW_CLIENT_RESTORE 333
DISALLOW_SERVER_FILE_WRITES 351
Disaster recovery
 catalogs 144
 collect information for 86
Disaster recovery diskettes
 updating 691
Disaster recovery procedure 693
Disk Administrator 698
Disk overhead, for catalogs 628
diskfull_notify script 660
Disk-image backup
 definition 708
 specifying 103
DLT, definition 708
DNS (see Domain Name Service)
DO_NOT_RESET_FILE_ACCESS_TIME 25
3, 352
Domain Name Service
 hostnames 607
DomainOS client exceptions
 APOLLO_RESTORE_TIMEOUT 328
 backup images 642
 MAX_APOLLO_RESTORE_ARG
 _CHARS 336
Done job 197
drfile.exe command 692
Drive, standalone (see Standalone drive)
Duplicate backups
 creating 225
 restoring from 231
Duration of backup window
 examples 132
 setting 131
Duration to retain Logs 271
DWO 183

E

Email address (see Mail notifications)
Emergency IDR bootable diskettes 698
Encryption class attribute 81
End time, schedule 131
Enhanced authentication 303
Escape character, on UNIX 96, 120
Exclude files list 119
 example 121
 for specific classes and schedules 122
Expiration, backup (see Retention period)
Extensible-Client class type 64

External VSN (EVSN) 636

F

FastBackup, Auspex 613
File list
 AFS directives 674
 AFS example entry 674
 Apollo clients 95
 Auspex-FastBack client 102
 disk image on Windows NT 103
 extension clients 110
 links on UNIX 98
 NetWare clients
 nontarget 107
 target 108
 OS/2 clients 107
 raw partitions on UNIX 100
 regular expressions for AFS 675
 standard clients 95
 unix files not backed up 96
 Windows clients 102
File System Processor, Auspex 615
File-list directives 111
Files
 /.rhosts 88
 /usr/lib/X11/XKeysymbDB 608
 catalog space requirements 162, 628
 files in /usr/opencv/netbackup/
 bin/goodies 242
 bp.conf 326
 db/jobs 203
 for catalog backup 153
 host.xlate 607
 linked, UNIX 97
 NFS mounted 75
 No.restrictions 212
 peername 212
 restoring to alternate client 213
 restrictions on restores 211
 terminfo 608
 version xxxi
Filesystem analyzer 13
Filtering activity monitor output 189
Flashbackup class type 65
FNum 183
Follow NFS mounts
 advantages of 76
 disadvantages of 76
 notes on use



- with cross mount points 75
 - with raw partitions 76
 - setting 75
 - with cross mount points 78
 - Fonts
 - changing for xbp 609
 - for Java interfaces 261
 - Format description
 - optical 643
 - tape 642
 - Fragment
 - disk storage unit 35
 - media manager storage unit 31
 - tape format 643
 - freeze media 638
 - Frequency (see Backup frequency)
 - Frozen, media state
 - definition 638
 - unfreezing 638
 - Full backups (see Backups)
 - Full, media state 637
- G**
- GDM Console 665
 - GDM host 665
 - Global attributes, NetBackup 268
 - Global data manager 663
 - GNU tar 610
 - goodies directory 242
 - Gravity stacker 641
- H**
- Hard link on NTFS volumes 105
 - Hard links, UNIX directories 98
 - Hashed file 309
 - Host names
 - changing client name 606
 - changing server name 606
 - client peername 605
 - correct use 604
 - short 606
 - Host Processor, Auspex 615
 - host.xlate file 607
 - How Long to Keep TIR Information 272
 - HSM 711
- I**
- IBM computers, recovering with IDR 700
 - IDR (see Intelligent Disaster Recovery)
 - IDR preparation wizard
 - preparing bootable media 683
 - updating disaster recovery diskettes 691
 - IDX 182
 - IFRFR (see Individual File Restore From Raw)
 - Ignore command 671
 - Images
 - changing primary copy 231
 - duplicating 225
 - fragmentation 643
 - import 236
 - restoring from duplicate 231
 - verify 232
 - Images on media report 182
 - Import backup images 236
 - Include files list 119, 122
 - Incremental backups (see Backups)
 - index_client command 221
 - Indexing, image catalog 221
 - Individual File Restore From Raw 85
 - Informix class type 65
 - INFORMIX_HOME 352
 - initbpdbm 210
 - initbprd 209
 - INITIAL_BROWSE_SEARCH_LIMIT
 - set on server 334
 - set on UNIX client 352
 - Inode change time 356
 - Install client software
 - on PC clients 90
 - on secure clients 90
 - on trusting clients 88
 - install_path 711
 - Installation
 - AFS 673
 - Intelligent disaster recovery
 - bootable media
 - choosing type 683
 - creating CD image 687
 - creating diskettes 684
 - preparing 683
 - collect information for 86
 - configuration 682
 - Custom Setup, when to use 695
 - DR files
 - obtaining from server 682
 - overview 681
 - update with drfile.exe 692



- frequently asked questions 700
- hard disk partition changes 696
- hard drive partition, altering sizes 698
- overview 680
- preparation wizard 683
- recovery wizard 695
- requirements for using 680
- supported Windows NT editions 679
- updating IDR media
 - disaster recovery CD 692
 - recovery diskettes 691
 - using drfile.exe 692
 - when to update 690
- using boot managers 700
- Windows NT Disk Administrator 696
- Windows NT Editions Supported 679
- Windows NT Setup 694
- wizard
 - disaster recovery 695
 - IDR preparation 683

- Intelligent Disaster Recovery diskettes
 - emergency 698
 - preparing 683
 - updating 691
- Interval for status reports 272

J

- Java
 - auth.conf file 256
 - authorizing users 256
 - change default login host 260
 - change fonts 261
 - insufficient memory problems 264
 - jbp.conf file 260
 - jbpSA configuration options 260
 - nonroot usage 253, 254
 - performance improvement hints 263
- Java interface 6
- jbpSA (see Java)
- jnbSA 6
- Jobs
 - (also see Activity monitor)
 - concurrent per disk storage unit 34
 - maximum per class 68
 - maximum per client 269
 - priority for class 69
- jobs.conf file 203

K

- KEEP_DATABASE_COMM_FILE 353

- keep_days 203
- keep_hours 204
- KEEP_LOGS_DAYS 353
- keep_successful_days 204
- keep_successful_hours 204
- Keyboard support
 - terminfo 608
 - XKeysymbDB 608
- Kill
 - uncompleted jobs 198
- KNOWN_MASTER 334

L

- Labeling media 640, 644
- Last media used, catalog backups 147
- Library, definition 712
- License keys 224
- Limit jobs per class, setting 68
- LIMIT_BANDWIDTH
 - configuration option 335
 - configuration overview 293
- Links
 - UNIX hard-linked directories 98
 - UNIX symbolic 97
- LIST_FILES_TIMEOUT 353
- Load balancing 250
- Locale 323
- LOCKED_FILE_ACTION 354
- Logs
 - (also see reports)
 - db extension logs, retention 353
 - progress, for user operations 5
 - setting retention period 271
- Lotus-Notes class type 64

M

- Mail notifications
 - configuration overview 321
 - Email address for admin 269
 - USEMAIL on UNIX clients 356
 - Windows NT nbmail.cmd script 269
- Man pages 359
- Manual backups
 - class 141
 - classes with bpadm 592
 - NetBackup catalogs 156
- Master of masters 663
- Master server (see Server, NetBackup)
- MASTER_OF_MASTERS 336
- Max Drives this Master 583



-
- MAX_APOLLO_RESTORE_ARG
 - _CHARS 336
 - Maximum concurrent jobs
 - disk storage unit 34
 - Maximum fragment size
 - disk storage unit 35
 - media manager storage unit 31
 - Maximum jobs per class
 - (see Limit jobs per class)
 - Maximum jobs per client
 - specifying 269
 - Media
 - active 184
 - automatic suspend 645
 - backup fragmentation 643
 - determining requirements 627
 - format 642
 - freeze 638
 - ID 636
 - labeling 644
 - last used for catalog backup 147
 - media and device information 636
 - mount and unmount 645
 - nonactive 184
 - overwrite protection 327
 - reports (see Reports)
 - selection algorithm 638
 - spanning 640, 641
 - states 637
 - type for catalog backup 148
 - unfreeze 638
 - unsuspend 638, 645
 - using tar to read images 610
 - Media 1 and Media 2, catalog backup 148
 - Media contents report 181
 - Media host, definition 713
 - Media ID
 - for catalog backup 149
 - Media list report 179
 - Media log entries report 183
 - Media Manager
 - overview 635
 - Media Manager host
 - definition 714
 - Media mount timeout 271
 - for storage migrator 252
 - Media server
 - configuring 281
 - definition 714
 - Media summary report 183
 - Media written report 184
 - MEDIA_ID_PREFIX 337
 - MEDIA_REQUEST_DELAY 337
 - MEDIA_UNMOUNT_DELAY 337
 - MEGABYTES_OF_MEMORY 354
 - methods.txt file 304
 - methods_allow.txt file 305
 - methods_deny.txt file 306
 - Mount media 645
 - Mount points and file systems 77
 - Move detection 70
 - MPX (see Multiplexing)
 - MPX_RESTORE_DELAY 338
 - MS-Exchange class type 64
 - MS-SQL-Server class type 64
 - MS-Windows-NT class type 64
 - mt 642
 - MTF1 format, allow overwrite 327
 - mtime 356
 - Multiple data streams
 - (see Allow multiple data streams)
 - Multiple servers 280
 - Multiplexing (MPX)
 - overview 275
 - block sizes 181, 644
 - demultiplexing 280
 - max jobs per client 279
 - overhead 644
 - schedule media multiplexing 277
 - set for schedule 130
 - storage unit max per drive 277
 - tape format 644
 - when to use 276
 - MUST_USE_LOCAL_DRIVE 338
 - Mwm*keyboardFocusPolicy X resource 8
- N**
- Names (see Host names)
 - names_allow.txt file 306
 - names_deny.txt file 307
 - njb.conf file 260
 - njava directory 258
 - nbmail.cmd script 269
 - NCR-Teradata class type 64
 - NDMP
 - class type 64
 - storage units 18
 - NetBackup administration - Java 6



NetBackup Assistant 14
NetBackup authorization
 configuration 315
NetBackup catalogs 165
 (see Catalog backup)
NetBackup Client service
 definition 715
NetBackup configuration options
 configuring 325
 definition 715
NetBackup Database Manager service
 definition 715
NetBackup databases (see Catalogs)
NetBackup Device Manager service
 definition 715
NetBackup Request Manager service
 definition 716
NetBackup Volume Manager service
 definition 716
NetBackup-Java (see Java)
NetServer, Auspex 614
NetWare class type 64
NetWare Loadable Module 716
Network drives, backing up 73
Network file system
 (see Follow NFS mounts)
Network loading 270
NEW_STREAM, file list directive 113
NFS mounted files (see Follow NFS Mounts)
NFS_ACCESS_TIMEOUT 354
NLM (see NetWare Loadable Module)
Nonactive media 184
Nonrobotic, definition 716
Nonroot administration
 all applications 253
 specific applications 258
nonroot_admin script 254
Notification scripts 647
Number of drives setting, storage unit 31

O

Obsolescence date 252
On demand only
 disk storage unit 33
 media manager storage unit 29
Open files (see Busy-file processing)
Open transaction management (OTM) 303
Optical format 643
Oracle class type 65

OS/2 Boot Manager and IDR 700
OS/2 class type 64
OSF/Motif, User's Guide 6
Overhead, for catalogs 628
Override class storage unit 129
Override class volume pool 129
Overwrite protection 327

P

Pathname
 catalog backup to disk 150
 rules for class file list 95
PC NetLink files 97
Peername
 client 605
 file 212
Performance
 measuring (see Transfer rate)
 reducing search time 221
Performance improvement, Java
 applications 263
Planning
 classes 43
 storage units 24
 user schedules 138
 worksheet 48
Pools (see Volume pools)
Power down NetBackup servers 208
Prelabeling media 640
Preprocess interval 115, 583
Primary copy
 changing 231
 definition 226
 in images on Media report 182
Priority
 for a class 69
 for jobs in worklist 625
 for schedule 128
Problems report 178
Processes, show active 210
Progress logs, client 5

Q

QIC 643
QUEUE_ON_ERROR 339
Queued_job 197

R

RANDOM_PORTS
 set use on client 355



- set use on server 339
 - Raw partition backups
 - definition 717
 - on UNIX 100
 - relative speed on UNIX 102
 - when to use on UNIX 101
 - rbak 642
 - RE_READ_INTERVAL 339
 - Rebooting a NetBackup server 208
 - Recognize command 672
 - Redisplay Data
 - activity monitor 192
 - reports window 174
 - Regular expressions, AFS file list 675
 - Remote media server
 - definition 718
 - Reports
 - all log entries 179
 - backup status 177
 - client backups 177
 - images on media 182
 - interval for information 272
 - introduction to 169
 - media contents 181
 - media list 179
 - media log entries 183
 - media summary 183
 - media written 184
 - problems 178
 - running 175
 - settings 175
 - window description 170
 - Re-Queued job 197
 - REQUIRED_INTERFACE
 - set on client 355
 - set on server 340
 - Resource (also see Target service) 718
 - restore_notify script 660
 - RESTORE_RETRIES 355
 - Restores
 - overview 2, 5
 - adjust time zone for 322
 - AFS clients 676
 - catalog backups 144
 - caution for alternate client 211
 - caution for alternate path 213
 - definition 718
 - directed from the server 223
 - from duplicated backups 231
 - notes on AFS 677
 - reducing search time 221
 - restore_notify script 660
 - server independent 243
 - setting client permissions 217
 - symbolic links on UNIX 97
 - to alternate clients 211
 - Retention level (see Retention period)
 - Retention levels
 - default 127
 - Retention period
 - caution for setting 139
 - definition 718
 - guidelines for setting 626
 - mixing on media 128
 - precautions for setting 127
 - redefining 273
 - setting 126
 - user schedule 139
 - Retry
 - backups 271
 - restores 355
 - rl 180
 - Rmed 182
 - root 718
- S**
- SAP class type 65
 - Save
 - job information 199
 - Schedule
 - overview 4, 41
 - adding to class 123
 - automatic, how processed 623
 - catalog backup 151
 - clear button 131
 - default for user backups 346
 - Duplicate button 131
 - examples, of automatic 132
 - frequency 128
 - naming 124
 - priority 128
 - retention level 127
 - retention period
 - guidelines 626
 - setting 126
 - setting backup times 130
 - specify duration 131
 - specify multiplexing 130



- specify start time 130
- storage unit 129
- type of backup 124
- user backup or archive 138
- volume pool 129
- Schedule backup attempts 271
- Scripts
 - add_slave_on_clients 285
 - available_media 627
 - backup_exit_notify 647
 - backup_notify 647
 - bpdbjobs example 205
 - bpend_notify
 - UNIX 647
 - Windows 647
 - bpps 210
 - bprd 209
 - bpstart_notify
 - UNIX 647
 - Windows 647
 - cautions for bpend and bpstart 648
 - dbbackup_notify 647
 - diskfull_notify 647
 - goodies 242
 - initbpdbm 210
 - nonroot_admin 254
 - notification 647
 - rc.local 209
 - rc0.d 209
 - rc2.d 209
 - restore_notify 647
 - session_notify 647
 - session_start_notify 647
 - userreq_notify 647
- SERVER
 - bp.conf option on client 355
 - bp.conf option on server 341
- Server and client speed 270
- Server directed restore
 - definition 719
- Server independent restores 243
- Server list
 - on unix server 341
- Server, NetBackup
 - change for duplicate images 227
 - change for import images 239
 - change for verify images 233
 - configuring bp.conf file 326
 - controlling access 341, 355
 - for catalog backup settings 146
 - for class configuration 57
 - for manage storage units 22
 - host names, changing 606
 - master 281
 - master of masters 663
 - media 281
 - multiple 280
 - power down 208
 - rebooting 208
 - select for monitoring 195
 - software 1
 - using Storage Migrator 252
 - SERVER_PORT_WINDOW 342, 355
 - SERVER_RESERVED_PORT_WINDOW 342
 - Session, NetBackup 720
 - session_notify script 661
 - session_start_notify script 661
 - SKIP_SMALL_VOLUMES 675
 - Slave server (see Media server)
 - SLAVE_CONNECT_TIMEOUT 343
 - SMS (see Storage Management Services)
 - Source Copy Number 228
 - Spanning media
 - enabling 640, 641
 - tape format 644
 - Split-Mirror class type 65
 - SQL-BackTrack class type 65
 - Stacker, gravity 641
 - Standalone
 - definition 720
 - Standalone drive
 - keeping in ready state 641
 - media selection 640
 - standalone extensions
 - disabling 641
 - enabling 640
 - using gravity stacker 641
 - Standard class type 64
 - Status (see Logs)
 - Status code 177
 - Storage Management Services 720
 - Storage Migrator 252, 721
 - Storage Processor, Auspex 615
 - Storage unit management window 18
 - Storage units
 - overview 3, 634
 - adding disk type 32



- adding media manager type 27
- adding NDMP type 18
- Auspex FastBackup Class 616
- changing attributes 35
- changing server to manage 22
- concurrent jobs (see jobs)
- definition 721
- deleting 36
- device configuration wizard 18
- disk type, definition 17
- example media manager type 24
- for class 65
- for schedule 129
- media manager type, definition 17
- NDMP 18
- next available 66
- rules for media manager type 23
- starting the management program 18
- Streaming (see Allow multiple data streams)
- Sun PC NetLink 97
- Suspend media 645
- Suspended, media state 638
- Sybase class type 65
- SYBASE_HOME 356
- Symbolic links, UNIX 97
- System Commander and IDR 700

T

- Tape format
 - overview 642
 - fragmentation 643
 - multiplexed 644
 - non-QIC 642
 - QIC 643
- Tape marks 628
- Tape overhead, for catalogs 628
- Tape spanning 640, 641
- Tapes (see Media)
- tar
 - GNU 610
 - to read backup images 610, 642
- tar format, allow overwrite 327
- Target Service 722
- Templates for directives 91
- Terminals, configuring 608
- terminfo file for bp users 608
- Throttling (see LIMIT_BANDWIDTH)
- Time zone adjustment for restores 322
- Timeout

- bpend 329
- bpstart 329
- bptm query 329
- client connect 330
- client read 330, 348
- DomainOS client read 328
- slave_connect 343
- TIMEOUT_IN_QUEUE 343
- TIMEOUT_IN_QUEUE 343
- TIR (see True image recovery)
- tpconfig, definition 723
- Transfer rate 179, 619, 620
- Tries, backup 271
- Troubleshooting AFS backups 678
- Troubleshooting wizard
 - using in activity monitor 202
- True image restore
 - configuration 70
 - definition 723
 - move detection 70
 - time to keep information 272
- Type of backup 124

U

- Uncompressing NetBackup catalogs 166
- Unfreeze media 638
- Unhashed file 309
- Unmount media 645
- UNSET, file list directive 118
- UNSET_ALL, file list directive 118
- Unsuspend media 638, 645
- Updating IDR bootable media 690
- USE_CTIME_FOR_INCREMENTALS 253, 356
- USEMAIL on UNIX clients 356
- User backups, archives, restores 5
- User schedules
 - planning 138
- User, definition xxxi
- userreq_notify script 661

V

- VERBOSE 343, 357
- Verify, backup images 232
- Version file xxxi
- vimages 180
- vm.conf file
 - definition 724
- vmadm, definition 724
- Volume database 636



Volume database host
 definition 724
Volume database, definition 724
Volume group
 definition 724
Volume pool
 class 67
 definition 724
 for schedule 129
 overview 635
Volumes
 allocation 635
 assignments 635
 scratch 635
vopie_util command 546
vopied 308
vopied command 544

W

WAIT_IN_QUEUE 343
Wakeup Interval
 specifying 271
Wildcard characters
 escaping on UNIX 96, 120
 in AFS file list 675
 in exclude lists 120
 mac clients 110

UNIX file paths 96
windows clients 102
Wizard
 backup policy 53
 catalog backup 145
 device configuration 18
 disaster recovery 695
 IDR preparation 683
 starting from NetBackup Assistant 14
 troubleshooter in activity monitor 202
Worklist, prioritizing 625
Worksheet, planning 48
WORM media
 definition 725
 retention period caution 127

X

X resources, xbp 609
xbp
 changing X resources 609
xbp command 550
xbpadmin
 OSF/Motif, User's Guide 6
xbpadmin command 552
xbpmon command 554
XKeysymDB file for xbp users 608



