

Symantec NetBackup™ Administrator's Guide, Volume II

Windows

Release 7.5

Symantec NetBackup™ Administrator's Guide, Volume II

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Documentation version: 7.5

PN: 21220059

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

www.symantec.com/business/support/

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

Contacting Technical Support

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

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America supportsolutions@symantec.com

Contents

Technical Support	4	
Chapter 1	Traditional licensing	13
	About traditional licensing	13
	About using nbdeployutil for traditional licensing	13
	Gathering data	14
	Analyzing the gathered data	16
	About the traditional licensing report	17
	How to reconcile the traditional licensing report	18
	Verify the Summary tab	18
	Complete the Hosts tab	19
	Resolve the NDMP tab	20
	Update the Virtual Servers tab	20
	Confirm the Drives tab	20
	Final steps	21
Chapter 2	Capacity licensing	23
	About capacity licensing	23
	Requirements	23
	About Front-end Terabytes	24
	About capacity usage calculation tools	24
	About using nbdeployutil for capacity licensing	25
	Gathering capacity data	26
	Reporting on the gathered capacity data	27
	Business unit reporting	28
	About the capacity licensing report	28
	Factors influencing performance	29
	About the capacity licensing report	30
	Examining the results	30
	Verify the completeness of the report inputs	30
	Eliminate redundant data due to client aliases and multiple IP addresses	31
	Examine the Itemization tab for flagged conditions in the Accuracy column	31

Verify correct grouping and summation of multistreamed backup images	33
How to reconcile the capacity licensing report results	33
Verify completeness of the report	34
Locate policy full backup	34
Review compressed image information	34
Eliminate redundant counting	34
Determine affect of multistreamed backups	35
Confirm the accuracy of any database backups	35
Locate full backup for snapshot image	35

Chapter 3	Additional configuration	37
	About multiple NetBackup master servers	37
	About multiple media servers with one master server	38
	About software on each server	40
	About NetBackup catalogs	40
	About direct I/O for backups	40
	Disabling direct I/O	41
	About dynamic host name and IP addressing	41
	About setting up dynamic IP addresses and host names	43
	Configuring the NetBackup master server	44
	bpclient commands that control client entries	46
	Configuring a dynamic Microsoft Windows client	46
	Configuring a dynamic UNIX NetBackup client	47
	About specifying the locale of the NetBackup installation	49
	Exporting PureDisk data to NetBackup	50
	Required software and licenses	50
	Exporting PureDisk data	50
	Restoring PureDisk export data	52
	About Shared Storage Option configuration	53
	About Shared Storage Option components	54
	About reserving or releasing shared devices	57
	How to share robotic libraries without using the Shared Storage Option	58
	Shared Storage Option terms and concepts	58
	About the Shared Storage Option license key	59
	Licensing the Shared Storage Option	59
	About Shared Storage Option prerequisites	60
	About hardware configuration guidelines	61
	About installing and configuring drivers	62
	Verifying the connectivity	62
	About configuring the Shared Storage Option in NetBackup	63

Verifying your Shared Storage Option configuration	65
Device Monitor and Shared Storage Option	70
Viewing SSO summary reports	71
Operating system assistance	72
Common configuration issues with Shared Storage Option	72
Frequently asked questions about Shared Storage Option	74
About the vm.conf configuration file	74
ACS_mediatype entry in vm.conf	74
ACS_SEL_SOCKET entry in vm.conf	75
ACS_SSI_HOSTNAME entry in vm.conf	75
ACS_SSI_SOCKET entry in vm.conf	75
ADJ_LSM entry in vm.conf	75
API_BARCODE_RULES entry in vm.conf	77
AUTHORIZATION_REQUIRED entry in vm.conf	77
AUTO_PATH_CORRECTION entry in vm.conf	78
AUTO_UPDATE_ROBOT entry in vm.conf	78
AVRD_PEND_DELAY entry in vm.conf	79
AVRD_SCAN_DELAY entry in vm.conf	79
CLEAN_REQUEST_TIMEOUT entry in vm.conf	79
CLIENT_PORT_WINDOW entry in vm.conf	80
CLUSTER_NAME entry in vm.conf	80
CONNECT_OPTIONS entry in vm.conf	80
DAS_CLIENT entry in vm.conf	81
DAYS_TO_KEEP_LOGS entry in vm.conf	81
EMM_RETRY_COUNT entry in vm.conf	82
EMM_CONNECT_TIMEOUT entry in vm.conf	82
EMM_REQUEST_TIMEOUT entry in vm.conf	82
ENABLE_ROBOT_AUTH entry in vm.conf	83
INVENTORY_FILTER entry in vm.conf	83
MAP_ID entry in vm.conf	83
MAP_CONTINUE_TIMEOUT entry in vm.conf	84
MEDIA_ID_BARCODE_CHARS entry in vm.conf	84
MEDIA_ID_PREFIX entry in vm.conf	85
MM_SERVER_NAME entry in vm.conf	86
PREFERRED_GROUP entry in vm.conf	86
PREVENT_MEDIA_REMOVAL entry in vm.conf	86
RANDOM_PORTS entry in vm.conf	87
REQUIRED_INTERFACE entry in vm.conf	87
SERVER entry in vm.conf	87
SSO_DA_REREGISTER_INTERVAL entry in vm.conf	88
SSO_DA_RETRY_TIMEOUT entry in vm.conf	88
SSO_HOST_NAME entry in vm.conf	89
TLH_mediatype entry in vm.conf	89

TLM_mediatype entry in vm.conf	89
VERBOSE entry in vm.conf	89
Example vm.conf file	90
Host name precedence in the vm.conf file	90

Chapter 4

Reference topics	91
Host name rules	91
How NetBackup uses host names	92
Updating NetBackup after changing the host name	93
Special considerations for Domain Name Service (DNS)	94
About reading backup images with tar	96
Consequences of using a non-NetBackup tar	96
About the files that tar generates	97
Factors that affect backup time	97
Total amount of data to back up	98
Transfer rate	98
Methods for determining the NetBackup transfer rate	99
Examples of reports that provide backup data to calculate transfer rates	100
Using the System Monitor with NetBackup	102
NetBackup notify scripts	103
backup_notify.cmd on Windows	105
backup_exit_notify.cmd on Windows	105
bpstart_notify (UNIX clients only)	106
bpstart_notify.bat (Microsoft Windows clients only)	109
bpend_notify (UNIX clients only)	111
bpend_notify.bat (Microsoft Windows clients only)	113
diskfull_notify.cmd on Windows	116
mail_dr_info.cmd	116
media_deassign_notify	117
nbmail.cmd	117
parent_end_notify.cmd on Windows	118
parent_end_notify	118
parent_start_notify.cmd on Windows	119
pending_request_notify	120
restore_notify.cmd on Windows	120
session_notify.cmd on Windows	120
session_start_notify.cmd on Windows	120
shared_drive_notify.cmd on Windows	121
userreq_notify.cmd on Windows	121
Media and device management best practices	122
Media management best practices	123

Device management best practices	123
Media and device performance and troubleshooting	124
About TapeAlert	124
About TapeAlert cleaning (reactive cleaning)	125
About TapeAlert and frequency-based cleaning	125
About TapeAlert requirements	125
TapeAlert logs and codes	126
About tape drive cleaning	129
About library-based cleaning	129
About frequency-based cleaning	130
About operator-initiated cleaning	130
About using a cleaning tape	131
How NetBackup selects drives	131
How NetBackup reserves drives	132
About SCSI persistent reserve	133
About the SPC-2 SCSI reserve process	135
About SCSI reserve requirements	138
About SCSI reserve limitations	139
About SCSI reservation logging	139
About server operating system limitations	139
About checking for data loss	140
About checking for tape and driver configuration errors	140
About configuring SCSI reserve	141
How NetBackup selects media	141
About selecting media in robots	142
About selecting media in standalone drives	145
Volume pool and volume group examples	147
Media formats	150
Media Manager commands	153
Chapter 5	
UNIX reference topics	157
About exclude and include lists on UNIX clients	157
Syntax rules for exclude lists	158
Example of an exclude list	159
Exclude lists for specific policies or schedules	160
About creating an include list on a UNIX client	160
Schedules for user backups or archives	161
Index	163

Traditional licensing

This chapter includes the following topics:

- [About traditional licensing](#)
- [About using nbdeployutil for traditional licensing](#)
- [About the traditional licensing report](#)
- [How to reconcile the traditional licensing report](#)

About traditional licensing

Traditional licensing is based on the total number of clients. Client information is gathered and a report is generated. The information in the report is then reconciled with actual clients in the NetBackup environment. This information then forms the basis for license fees.

About using nbdeployutil for traditional licensing

The utility performs two steps. Data is gathered in the first step and analyzed in the second step. The following table describes the tasks to prepare a traditional license model report.

Table 1-1

Task Number	Description
Task 1	<p>Gather data from one or more master servers.</p> <p>The <code>nbdeployutil</code> utility gathers data remotely for multiple master servers from a central location, provided the master servers granted the initiating server access. The utility supports remotely collecting data from back-level master servers (NetBackup 6.5.6 and later). You must load the engineering binary that is associated with this utility onto all master servers for which you want to gather information.</p> <p>See “Gathering data” on page 14.</p>
Task 2	<p>Run analysis on the gathered data.</p> <p>After the gather process finishes, run the <code>--report</code> option to generate the traditional license report.</p> <p>See “Analyzing the gathered data” on page 16.</p>
Task 3	<p>Examine the results and make the necessary adjustments.</p> <p>See “How to reconcile the traditional licensing report” on page 18.</p>

Depending on your environment, the `nbdeployutil` utility takes from a several seconds to several minutes to complete. This behavior is true for both the `--gather` and the `--report` parameters. In general, the `nbdeployutil` utility runs faster on Linux and Windows servers as compared to other platforms.

Gathering data

The `nbdeployutil` utility contains the following options for collecting traditional data:

```
nbdeployutil --gather [--bpimagelist=options] [--capacity|
--traditional] [--client hostname1, [hostname2, hostname#] |
--clientlist=filename] [--hoursago=number] [--log=filename]
[--master=hostname] [--nolog] [--output=directory] [--runtimestats]
[--start date [--end date]]
```

Refer to the *NetBackup Commands Reference Guide* for a complete description of the parameters.

You can gather capacity data for:

- A single master server
- A remote master server

■ A specific set of clients

Example 1 - Gather information for the local master server

```
root@server_01> admincmd/nbdeployutil --gather
NetBackup Deployment Utility, version
7.1.0.1_EEB1_PET2326556_SET2371026.2011.0523
Gathering license deployment information...
  Discovered master server server_01.domain.com
  failed bptestbpcd to 1 of 77 clients, for details see:
  /usr/opensv/var/global/reports/20110523_175606_server_01.
domain.com/nbdeployutil-gather-20110523_175606.log
  Output for server_01.domain.com at: /usr/opensv/var/global/reports/
20110523_175606_server_01.domain.com
Gather DONE
Execution time: 9 mins 56 secs
To create a report for this master server, run one of the following:
  capacity   : nbdeployutil --report --capacity /usr/opensv/var/
global/reports/20110523_175606_server_01.domain.com
  traditional: nbdeployutil --report --traditional /usr/opensv/var/
global/reports/20110523_175606_server_01.domain.com
```

The utility generates a log file named `nbdeployutil-gather-timestamp.log` during the gathering operation. By default, the log file is created in the directory where the gathered data resides.

Example 2 - Gather information for a remote master server

```
# nbdeployutil --gather --master=server_02.example.com
```

Example 3 - Gather information for a subset of clients that the local master server protects

```
# nbdeployutil --gather --client=client_01,client_02,client_03
```

or

```
# nbdeployutil --gather --clientlist=filename.txt
```

Note: When you use the `--client` or the `--clientlist` option, some media servers may show up as not connectable in the report even though the utility can connect to them. This problem should not affect the summary information.

Analyzing the gathered data

The `nbdeployutil` utility contains the following options for generating a traditional report:

```
nbdeployutil --report [--capacity|--traditional]
[dir1 dir2 dir# | --dirsfile=filename | --parentdir=directory]
[--log=filename] [--nolog] [--runtimestats]
```

Refer to the *NetBackup Commands Reference Guide* for a complete description of the parameters.

You can generate a report for:

- A single master server
- Several master servers

Example 1 - Generate a report using data that is collected for the local master server

This example is a continuation of Example 1 from the previous topic.

```
root@server_01> admincmd/nbdeployutil --report traditional
/usr/opensv/var/global/reports/20110523_175606_server_01.domain.com
NetBackup Deployment Utility, version
7.1.0.1_EEB1_PET2326556_SET2371026.2011.0523
Analyzing license deployment ...
  Master server_01.domain.com
  Report created at: /usr/opensv/var/global/reports/
  20110523_175606_server_01.domain.com/report-capacity-server_01.
  domain.com-20110523_180636.xls
Analysis DONE
Execution time: 13 secs
```

The utility generates a log file named `nbdeployutil-report-timestamp.log` during the analysis and the report generating operation. By default, the log file is created in the directory where the gathered data resides.

Example 2 - Generate a roll-up report for several master servers

This example assumes that you have gathered the respective master server's data in directories `master1dir`, `master2dir`, `master3dir`. These directories all reside within a parent directory named `EMEA-domains`. The output (report and log file) is saved to the `EMEA-domains` directory.

```
# nbdeployutil --report traditional
--parentdir=EMEA-domains
```

This variation creates a report for a smaller set of master servers and specifies a different directory for the output.

```
# mkdir UK-masters
# nbdeployutil --report --traditional EMEA-domains/master1dir
EMEA-domains/master2dir --output=UK-masters
```

About the traditional licensing report

This topic provides a brief explanation of how to interpret the traditional license report. The utility examines the image headers in the NetBackup catalog to determine the servers and clients in the NetBackup environment. The data that is retrieved during the data collection phase can also affect the results.

Much of the report information does not affect the final values on the **Summary** tab. The information is for information purposes only. This information is useful for reaching a better understanding of your environment.

The traditional license report is a Microsoft Excel spreadsheet with seven tabs:

- **Summary**

This tab shows the final details about master servers, media servers, and clients. This tab lists the source data for generating the report. The number of media servers and the number of clients is provided, as well as capacity information.

- **Hosts**

This tab provides a listing of host names, along with associated computer information. The associated information includes information such as: platform, computer type, database software installed, SAN media server, and NDMP.

- **NDMP**

This tab shows the computers that the utility has determined are NDMP servers and the corresponding tier number of the client. When you reconcile the report, you need to address the clients that are found on this tab.

- **Virtual Servers**

This tab shows the number of the virtual servers or the virtual hosts that were detected in the environment.

- **Drives**

This tab details the type of drives as well as the host or the library where the drive resides. The tab provides the host names that are associated with each drive as well as information about virtual tape libraries, shared drives, and vaulted drives.

- **Interpreting the results**

This tab provides a general overview of how to reconcile the information in the report which your actual environment.

- **Disclaimer**

This tab shows text explaining the limits of the report's calculations and proper use of the data. For example, the figures should not be used to audit compliance.

How to reconcile the traditional licensing report

This topic reviews the different tabs in the report and provides an overview on the process of reconciling the report with the actual NetBackup environment. The utility generates a report in a Microsoft Excel format.

Reconciling the traditional licensing report output is a five step process.

Reconciling the report

- 1 Examine the **Summary** tab and confirm the correct information is displayed.
See "[Verify the Summary tab](#)" on page 18.
- 2 Review the **Hosts** tab and resolve any missing information.
See "[Complete the Hosts tab](#)" on page 19.
- 3 Resolve any missing or any incomplete information on the **NDMP** tab.
See "[Resolve the NDMP tab](#)" on page 20.
- 4 Update the **Virtual Servers** tab with any missing information.
See "[Update the Virtual Servers tab](#)" on page 20.
- 5 Confirm all information on the **Drives** tab is accurate.
See "[Confirm the Drives tab](#)" on page 20.

Verify the Summary tab

The top of the report's Summary tab details the basis for the report's information. Review the **Period Analyzed** for the source of the information for the report. The **Period Analyzed** section includes:

- Start date for the gather for each master server.
- End date for the gather for each master server.
- The total number of days gathered for each master server.
- The input directory for each master server that is associated with the report.

The start and the end dates are not necessarily the dates that are specified for the `gather` command. These are the dates within the time period that you specified where images exist. If images do not exist for a specified start or end day, the day is not listed. The nearest date with backup images is included and listed.

The **Input Directory** column displays the path to the gathered data. Within that directory is the `nbdeployutil-gather-timestamp.log` file. If non-default inputs were used in the collection of catalog data, the log file displays this information.

Under the **Options** section, confirm the list of master servers is correct. If there are missing or extra master servers, you need to rerun the report.

When you finish your review of the entire report, all the values in the **Unknown** row under **Tiering** should be zero. As you reconcile the other tabs in the report, these values should automatically update to zero.

Complete the Hosts tab

The **Hosts** tab provides a listing of all media servers and client servers that are included in the report. The tab includes master servers if they are either a media server or a client server. You need to review five areas to complete the review of this tab.

Completing the Hosts tab

- 1 Scan the **Connectable** column and see how many hosts the utility was unable to connect to for its calculations. Be aware the utility cannot connect to NDMP filers. If there is a large number of non-NDMP filer hosts the utility could not connect to, consider rerunning the utility with the `--retry` option. Use the following command to retry the connections

```
nbdeployutil --retry <path_to_the_gathered_data>
```

When that finishes, use the following command to recreate the report.

```
nbdeployutil --report <all_previously_specified_options>  
<all_previously_specified_gather_directories>
```

- 2 Check the **Tier** column for any hosts that are listed as **UNKNOWN**. You must replace these with the appropriate tier number between one and four. Please work with your Symantec Sales Engineer to determine the correct tier information.

The **Platform** and **Processors** values help determine the host's tier. These columns do not calculate the tier, but knowing this information helps you determine the appropriate value to enter in the **Tier** column.

- 3 Review the **MSEO Key Server** column and verify all the listed information is correct. **Yes** indicates the host is an MSEO key server. **No** indicates the host is not an MSEO key server. The **N/A** value indicates the host is not a media server.
- 4 Check the **Enterprise Client** column and verify that the information is correct. **Yes** indicates the host is an enterprise client and was backed up. **No** indicates the host is not an enterprise client. The **N/A** value indicates no backups were performed on the host during the report period.
- 5 Review the **SAN Media Server** column, correct any hosts where the value is **UNKNOWN**, and confirm all other values are correct. A value of **N/A** for a host indicates the host is either a client server or a master server.

Be aware the only column which contributes to the final information on the **Summary** tab is the **Tier** column. So values of **UNKNOWN** in other columns other than **Tier** indicate unknown information. All data aside from the **Tier** column is for informational purposes only.

Resolve the NDMP tab

The **NDMP** tab shows hosts the utility has determined to be NDMP servers. If there are servers listed which are not NDMP servers, delete these servers from the list. Add any missing NDMP servers to the list. For all servers, review the Tier column and confirm the information is correct. Any **Tier** values of **UNKNOWN** should be replaced with the correct tier number between one and four. Please work with your Symantec Sales Engineer and the *NetBackup Pricing and Licensing Guide* to determine the correct tier information.

Consult with your Symantec Sales team if any of the listed NDMP servers are NearStore servers.

Update the Virtual Servers tab

Complete the **Virtual Servers** tab. Replace any **UNKNOWN** values under the **Used** column with **Yes** or **No**. **Yes** indicates the host uses NetBackup's ESX-specific feature and **No** indicates it does not use the feature. Add missing virtual servers to the list and indicate **Yes** in the **Used** column.

Confirm the Drives tab

On the **Drives** tab, review the information in the **VTL** column. Verify that all virtual tape libraries are correctly listed as **Yes**. If a virtual tape library has **No** for a value in the VTL column, change that to **Yes**. Change the value for **VTL** to **No** for any drives that are incorrectly marked as a virtual tape library.

Final steps

Once you reconcile the report, correct the errors, and enter the missing information, compare the results to the install base report. The install base report is provided to you by Symantec or your reseller. Confirm everything in the report matches up with what is contained in the install base report. If there are discrepancies, consult with your Symantec sales representative to correct problems.

Capacity licensing

This chapter includes the following topics:

- [About capacity licensing](#)
- [About using nbdeployutil for capacity licensing](#)
- [About the capacity licensing report](#)
- [How to reconcile the capacity licensing report results](#)

About capacity licensing

Capacity licensing is based on the total amount of data that is protected by NetBackup. This model differs from other NetBackup license models which are based on total clients or on total storage capacity. The total amount of protected data is calculated based on the backup image header information in the NetBackup catalog. Capacity information is gathered and a report is generated. The information in the report is then reconciled with actual capacity in use. This information then forms the basis for license fees.

Requirements

To run the capacity licensing utility, the master server must meet the following requirements

- A NetBackup master server running NetBackup 6.5.6 or later. This licensing model does not apply to NetBackup versions earlier than 6.5.6. You can run this utility from any master or any media server in the environment.
- A tool for reading `.xls` files. Symantec tested the utility with Microsoft Excel, but any tool for reading and editing `.xls` files should work.

About Front-end Terabytes

The licensing fees for the use of NetBackup are based on the total number of Front-End Terabytes (FETBs) protected by NetBackup. Front-End Terabyte Calculation is a way of determining the total terabytes of data NetBackup protects. A Front-End Terabyte (FETB) is one terabyte of protected data. The data can either be on clients or devices where the software is installed or where the software is used to provide backup functionality.

The utility examines the image headers in the NetBackup catalog to determine the terabytes of data that NetBackup protects. Any partial terabytes of data are rounded up to the next whole terabyte. The final total is the sum of the FETBs for each client/policy combination that the analyzer examines. The utility measures the actual data protected. It does not measure the capacity of the storage where the data resides or the total amount of data that is stored on the device.

Consider the following:

- Assume a device with 100 TB of total storage capacity.
- A total of 65 TB of the total capacity is in use.
- NetBackup protects a total of 60 TB of the used data through multiple backup storage units.
- That is measured as 60 TB of front-end capacity.

The total terabytes of front-end capacity are independent of the number of copies NetBackup makes. A backup of 200 TB to basic disk with two copies to tape is still only 200TB of front-end capacity.

About capacity usage calculation tools

NetBackup provides three methods to calculate capacity usage.

OpsCenter	Provides a GUI interface useful for multi-server environments.
<code>nbdeployutil</code>	Provides a command-line access to capacity usage. It provides a richer set of input parameters and is highly customizable. <code>nbdeployutil</code> can also be used for business unit reporting. The utility generates a Microsoft Excel spreadsheet which you can review and modify if capacity is over counted.

PureDisk reports The `nbdeployutil` utility calculates capacity usage for some types of PureDisk backups. When PureDisk is used as a disk storage unit for NetBackup (the PureDisk Deduplication Option), the utility calculates the capacity used.

When PureDisk clients back up to a PureDisk storage pool authority and NetBackup is not involved, the `nbdeployutil` binary cannot be used. The binary cannot be used because no NetBackup master server is present. In the PureDisk only environment, refer to the **Capacity Usage Report**. Refer to the *Reports* chapter of the *Symantec NetBackup PureDisk Administrator's Guide* for more information about the **Capacity Usage Report**.

Symantec has setup a Web site for updates and the most recent information about the `nbdeployutil` utility.

<http://www.symantec.com/docs/TECH145972>

About using nbdeployutil for capacity licensing

The utility performs two steps. Data is gathered in the first step and analyzed in the second.

Table 2-1 describes the tasks to prepare a capacity deployment analysis report.

Table 2-1 Process overview to prepare a capacity deployment analysis report

Task Number	Description
Task 1	<p>Gather catalog data from one or more master servers.</p> <p>The <code>nbdeployutil</code> utility can gather data remotely for multiple master servers from a central location, provided the remote master servers have granted the initiating server access. The utility supports remotely collecting capacity data from back-level master servers (NetBackup 6.5.6 and later).</p> <p>See “Gathering capacity data” on page 26.</p>
Task 2	<p>Report on the gathered data.</p> <p>The <code>nbdeployutil</code> utility can create three different types of reports.</p> <ul style="list-style-type: none"> ■ A roll-up report for all gathered data ■ A report per master server ■ A report for a specific set of clients (e.g., a business unit level report) <p>See “Reporting on the gathered capacity data” on page 27.</p>

Table 2-1 Process overview to prepare a capacity deployment analysis report
(continued)

Task Number	Description
Task 3	Examine the results and make adjustments. See “ About the capacity licensing report ” on page 28.

Gathering capacity data

The `nbdeployutil` utility contains the following options for collecting capacity data:

```
nbdeployutil --gather [--bpimagelist=options] [--capacity]
[--client hostname1, [hostname2, hostname#] | --clientlist=filename]
[--hoursago=number] [--log=filename] [--master=hostname] [--nolog]
[--output=directory] [--runtimestats] [--start date [--end date]]
[--traditional]
```

Refer to the *NetBackup Commands Reference Guide* for a complete description of the parameters.

You can gather capacity data for:

- A single master server
- A remote master server
- A specific set of clients

Example 1 - Gather capacity information for the local master server

```
# nbdeployutil --gather
NetBackup Deployment Utility, version 7.1.0000.0000
Gathering license deployment information...
  Discovered master server marybl2g1
  Output for marybl2g1 at: D:\Program Files\VERITAS\netbackup\
  var\global\reports\20101029_170534_marybl2g1
Gather DONE
Execution time: 1 min
To create a report for this master server, run the following:
nbdeployutil.exe --report "D:\Program Files\VERITAS\netbackup\
var\global\reports\20101029_170534_marybl2g1"
```

The utility generates a log file named `nbdeployutil-gather-timestamp.log` during the gathering operation. By default, the log file is created in the directory where the gathered data resides.

Example 2 - Gather capacity information for a remote master server

```
# nbdeployutil --gather --master=sidon.example.com
```

Example 3 - Gather capacity information for a subset of clients that the local master server protects

```
# nbdeployutil --gather --client=dynamo,lettuce,marble2
```

or

```
# nbdeployutil --gather --clientlist=filename.txt
```

Reporting on the gathered capacity data

The `nbdeployutil` utility contains the following options for generating a capacity report:

```
nbdeployutil --report [--capacity]
[dir1 dir2 dir# | --dirsfile=filename | --parentdir=directory]
[--log=filename] [--nolog] [--runtimestats] [--traditional]
```

Refer to the *NetBackup Commands Reference Guide* for a complete description of the parameters.

You can generate a report for:

- A single master server
- Several master servers
- A specific subset of clients. For example, a report that contains capacity usage for business unit billing.

More information about this option is available.

See [“Business unit reporting”](#) on page 28.

Example 1 - Generate a report using data that is collected for the local master server

This example is a continuation of Example 1 from the previous topic.

```
D:\>nbdeployutil.exe --report "D:\Program Files\VERITAS\netbackup\
var\global\reports\20101029_170534_maryb12g1"
NetBackup Deployment Utility, version 7.1.0000.0000
Analyzing license deployment for master maryb12g1 ...
```

```
Report created at: D:\Program Files\VERITAS\netbackup\var\global\
reports\20101029_170534_marybl2g1\report-20101029_170705.xls
Analysis DONE
Execution time: 27 secs
```

The utility generates a log file named `nbdeployutil-report-timestamp.log` during the analysis and the report generating operation. By default, the log file is created in the directory where the gathered data resides.

Example 2 – Generate a roll-up report for several master servers

This example assumes that you have gathered the respective master server's data in directories `master1dir`, `master2dir`, `master3dir`. These directories all reside within a parent directory named `EMEA-domains`. The output (report and log file) is saved to the `EMEA-domains` directory.

```
# nbdeployutil --report --parentdir=EMEA-domains
```

This variation creates a report for a smaller set of master servers and specifies a different directory for the output.

```
# mkdir UK-masters
# nbdeployutil --report EMEA-domains/master1dir EMEA-domains/master2dir
--output=UK-masters
```

Business unit reporting

The utility can be used to examine a specific set of clients in detail.

Example - Gather data for a subset of clients for a time frame different than the default.

```
nbdeployutil.exe --gather --output BusinessUnitFinance --start "11/01/10
06:00:00" --end "11/02/10 01:00:00" --clients marybl2g1,marybl7g1
--verbose
```

To create a report for these clients, run the following:

```
nbdeployutil.exe --report "BusinessUnitFinance\20101102_155246_marybl2g1"
```

About the capacity licensing report

This topic provides a brief explanation of how to interpret the capacity license report. This topic also details how to make the corrections that reflect your backup environment configuration. The utility examines the image headers in the NetBackup catalog to determine the amount of data NetBackup protects. How you

configure your client policy and schedule settings can affect the results. The data that is retrieved during the data collection phase can also affect the results.

The capacity license deployment report is an Excel spreadsheet with four tabs:

- **Summary**

This tab shows the final figures, an overview of the basis for the report (data source), and a breakdown of the source of the capacity. The capacity breakdown includes a reporting by policy type and largest clients.

See [“Verify the completeness of the report inputs”](#) on page 30.

- **Itemization**

This tab shows a table similar to the line itemization you see in your credit card bill. Each line is a charge that contributes to the final total. Each line lists the capacity that is calculated for a client/policy combination.

See [“Examine the Itemization tab for flagged conditions in the Accuracy column”](#) on page 31.

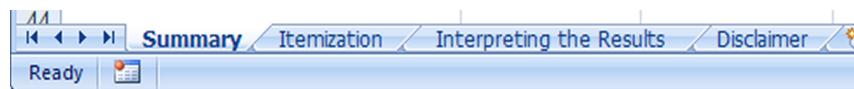
- **Interpreting the Results**

This tab shows descriptive text. The tab contains an explanation for how to examine the report and make adjustments as needed based on the unique properties of the configuration.

See [“Examine the Itemization tab for flagged conditions in the Accuracy column”](#) on page 31.

- **Disclaimer**

This tab shows text explaining the limits of the report’s calculations and proper use of the data. For example, the figures should not be used to audit compliance.



Factors influencing performance

The performance of the `nbdeployutil` utility is dependent on the system running it as well as the size of the NetBackup catalog. The gather command only executes as quickly as the `bpimagelist` command can run for 30 days worth of images.

The speed of report generation is dependent on the number of images and fragments. The operating system running the command also affects the utility’s performance. Preliminary testing at Symantec indicates this utility runs faster on Linux computers than on Windows computers.

About the capacity licensing report

The utility generates a report in a Microsoft Excel format. This topic reviews the different tabs in the report and provides an overview on the process of reconciling the report with the actual NetBackup environment.

Examining the results

Examining the deployment analysis results is a four step process.

Examining the report

- 1 Verify the completeness of the report inputs.
See [“Verify the completeness of the report inputs”](#) on page 30.
- 2 Eliminate redundant data due to client aliases and multiple IP addresses.
See [“Eliminate redundant data due to client aliases and multiple IP addresses”](#) on page 31.
- 3 Examine the **Itemization** tab for flagged conditions in the **Accuracy** column.
See [“Examine the Itemization tab for flagged conditions in the Accuracy column”](#) on page 31.
- 4 Verify correct grouping and summation of multistreamed backup images.
See [“Verify correct grouping and summation of multistreamed backup images”](#) on page 33.

Verify the completeness of the report inputs

The top of the report’s **Summary** tab shows the basis for the report's information. Examine the section marked **Analyzed** to verify the completeness of the gathered data upon which the report is based.

The **Analyzed** section displays the following:

- The master server(s) included in the report.
- The date range for catalog data.
- The number of clients and policies that are seen in the catalog output.

If the client and the policy counts are low, the report may be based on the data that was gathered with narrower, non-default inputs. The analyzer gathers 30 days worth of catalog data for all clients by default.

The **Input Directory** column displays the path to the gathered data. Within that directory is the `nbdeployutil-gather-timestamp.log` file. If non-default inputs were used in the collection of catalog data, the log file displays this information.

1	Capacity Licensing Report						
2	NetBackup Deployment Analyzer	Version 7.1.0000.0000					
3							
4	Analyzed:						
5	Master Server	Start Date	End Date	Number of Days	Total Clients	Total Policies	Input Directory
6	master1.example.com	09/27/2010	10/05/2010	7	382	373	/home/jbarr/FETB_te
7	master2.example.com	09/28/2010	10/05/2010	7	307	372	/home/jbarr/FETB_te
8	master3.example.com	09/28/2010	10/05/2010	7	10	70	/home/jbarr/FETB_te
9	master4.example.com	09/28/2010	10/03/2010	5	30	37	/home/jbarr/FETB_te

Eliminate redundant data due to client aliases and multiple IP addresses

The analyzer performs calculations based on the client name as stored in the catalog. Clients that are backed up by multiple aliases or multiple IP addresses are not collapsed into a single entry. For ease of accountability, the **Itemization** tab lists all client aliases and IP addresses used for backup separately. In some jurisdictions, the collection of the system IP address may be subject to regulation as personal data.

Determine where multiple client/policy lines refer to the same data set backed up through different interfaces. Make adjustments to the Charged Size value for all but one of the client/policy lines. We recommend retaining the value that is most recent. Annotate the duplicate client itemizations with a comment within the adjacent **Reason** cell. Indicate that the client's value is already counted under a different hostname. Please reference the hostname.

See [“Eliminate redundant counting”](#) on page 34.

Examine the Itemization tab for flagged conditions in the Accuracy column

The report’s **Itemization** tab shows the calculated capacity for each client/policy combination. The report flags conditions that have the potential to over count or to under count capacity. These conditions are identified in the **Accuracy** and **Accuracy Comment** columns.

1	Master Server	Client Name	Policy Name	Policy Type	Backup Image	Backup Date	Accuracy	Accuracy Comment
26	master1	RSVB3AVSPAP00	APP14_SMS_AV	MS-Windows-NT	RSVB3AVSPAP00_1285926056	10/01/2010	OK	
27	master1	RSVB3AVSPAP02	APP14_SMS_AV	MS-Windows-NT	RSVB3AVSPAP02_1285938853	10/01/2010	OK	
28	master1	RSVB3CFPAD01	APP4-DMS-OVR-SIM-RMS	Windows-NT	RSVB3CFPAD01_1286021066	10/02/2010	OK	
29	master1	RSVB3CFPADD01	RSVB3CFPADD01_RMA	Oracle	multiple	10/01/2010	Database Estimation (oracle)	DB size estimated via backup summation
30	master1	RSVB3CFPADP1	RSVB3CFPADP1_RMA	Oracle	multiple	10/03/2010	Database Estimation (oracle)	DB size estimated via backup summation,c
31	master1	RSVB3CFPADP1	RSVB3CFPADP1	MS-Windows-NT	RSVB3CFPADP1_1286045955	10/02/2010	Possible Overlap	Client appears in other policies
32	master1	RSVB3CFPAP01	APP4-DMS-SIM-RM	MS-Windows-NT	RSVB3CFPAP01_1286068882	10/02/2010	OK	

- Possible overlap - Client appears in multiple policies
A client in multiple backup policies has the potential to have the same data backed up more than once. Compare the policy types and names to determine if the case warrants a detailed examination of the respective policies' backup selections.
See [“Eliminate redundant counting”](#) on page 34.
- Database estimation - database size estimated via UBAK summation
The size of databases that a NetBackup database agent protects cannot be determined with certainty. Third party components external to NetBackup (e.g., RMAN) govern the composition of database backups.
The third-party component determines the number of backup streams and the contents of each stream. These backups are recorded as user-initiated backup images, i.e., UBAKs. NetBackup does not initiate backup streams, nor does it know each stream's relationship to the underlying database. Therefore the information in the catalog does not provide a single, clear, undisputable figure for the total size.
In these cases, the analyzer calculates an estimation upon which to base follow-on examinations. The analyzer uses the image header information to determine the total terabytes of data that were backed up each day within the date range examined. A day is defined as the 24 hour period from midnight to midnight. The analyzer sums all full and user-initiated backups that started within that period. The day with the largest total volume of protected data during the range that is examined is assumed to be the day when a full backup of the database was performed. This figure that is returned is an estimate of the approximate size of active data under protection for the client and policy.
See [“Confirm the accuracy of any database backups”](#) on page 35.
- Undiscoverable - No full backup found within range analyzed
The catalog has only incremental backups for the range analyzed. That error may indicate that a full backup falls outside the report's range or that a full backup does not exist.
See [“Locate policy full backup”](#) on page 34.
- Compressed Image
The client's data was sent to NetBackup in compressed form. The actual size cannot be determined with certainty. For all compressed backup images, the analyzer multiplies the final backup image size by a fixed value (the compression ratio). The value of the compression ratio is listed on the **Summary** tab.
See [“Review compressed image information”](#) on page 34.
- Size unavailable – Only snapshot is present

The catalog has only snapshots for the range analyzed. The analyzer requires a backup image of the snapshot in order to have an accurate figure for the client's protected capacity.

See [“Locate full backup for snapshot image”](#) on page 35.

- Possible multi-stream backup detected
The size of clients protected by multi-stream backups is the total of all backup images created by all streams.
See [“Determine affect of multistreamed backups”](#) on page 35.

Verify correct grouping and summation of multistreamed backup images

When a client is backed up by multiple streams, the client's size is equal to the total of all backup images created by all streams. Job throttles on the policy, the client, and the storage unit hinder the utility's ability to group the streams with certainty. For example, instead of starting within minutes of one another a subset of the backup streams may start in a different day than the rest of the backup streams. Because the utility sums only the backup images from streams that originate within the same 24 hour period (midnight to midnight), these streams are counted in separate days. Manually initiating a second full backup within the same day also skews the results. Streams from both backups are counted together as a group.

See [“Determine affect of multistreamed backups”](#) on page 35.

How to reconcile the capacity licensing report results

After you use the utility with the `--report` option, it generates a spreadsheet. After reviewing the resulting spreadsheet you can either:

- Accept the generated information without changes as the basis for license charges.
- Make changes and note the reason for the change.

As you make changes to the spreadsheet it is important to assess when any additional changes are no longer meaningful. Since licensing charges are assessed on a per terabyte basis, it may not be beneficial to dispute charges for a few gigabytes of information. You may wish to sort the clients by their backup size and focus on the largest backups first. Sorting by backup size provides two benefits. First, your efforts are initially focused on the largest clients. Second, if there are clients backing up only a few kilobytes, these backups may not capture the correct information. You may have important data which is unprotected.

Verify completeness of the report

On the **Summary** tab, look at the information under **Analyzed**. Confirm the master server or servers is correct, as well as the date, client, and policy information.

Locate policy full backup

On the **Itemization** tab, sort the list by **Accuracy Column**. For all lines with **Undiscoverable**, manually query the NetBackup catalog to determine if a full backup can be found. A full backup may exist in a time period that precedes the period the analyzer examined. Rerun the utility with specific options to restrict the collection and reporting to the specific client and a specific date range within which the full backup(s) fall. Alternatively, manually examine the client system to determine the size of data that would be backed up with the backup policy's selections and settings.

Review compressed image information

On the **Itemization** tab, sort the list by **Accuracy Comment**. For any compressed images, review the **Charged Size** column and confirm the correct information is displayed. If the information is inaccurate, change the **Charged Size** column, and add a note to the **Enter a Reason here when modifying the Charged Size** column explaining the change.

Eliminate redundant counting

On the **Itemization** tab, sort the list by **Client Name** and search for the use of hostname aliases. Look for instances where the itemization table lists the same client multiple times under the same policy but with a different hostname alias. If that occurs, zero out the **Charged Size** column for the lines with an earlier backup date. Then add a note to the **Enter a Reason here when modifying the Charged Size** column explaining why the **Charged Size** value is zero.

For some Oracle RAC backups, the presence of itemizations under different aliases can reflect the backup of different data sets. If you zero out the **Charged Size** the protected data is under counted.

If a client is found in more than one policy, confirm those policies do not have overlapping backup selections. If the backup selections overlap, find the redundant backup policies in the **Itemization** tab. Then make adjustments to the **Charged Size** value. Decrement the size by the value of the redundant backup selection and add a comment within the adjacent **Reason** cell.

Determine affect of multistreamed backups

On the **Itemization** tab, sort the list by **Accuracy Comment**. Find all backups that list **Possible multi-stream backup detected** under **Accuracy Comment** and make note of the policy name under the **Policy Name** column. Then open the log file that was generated when the `nbdeployutil --report` command ran. By default, the log file is in the directory where the gathered report is located.

Note: If OpsCenter generated the report, the log file is found on the OpsCenter server. The email with the report results contains a link to the log file location. The log file name is in the format `nbdeployutil-report-timestamp-log`.

In the log file, find the policy name for the policy in question and look at the corresponding **MAX** value. The excerpt from a log file that is shown highlights the information discussed.

```
Analyzing backups for policy <policy_name>, client <client_name>
Analyzing schedule Full
MAX 2010-09-01  14.6 T  (multiple backups      )
                   21.7 G  (client_name_1283295642) 09:00:42
                   1.0 T   (client_name_1283295643) 09:00:43
                   793.1 G (client_name_1283295644) 09:00:45
                   1.2 T   (client_name_1283295645) 09:00:48
                   1.5 T   (client_name_1283295647) 09:00:49
```

Confirm this information is correct for the policy. If the information is inaccurate, change the **Charged Size** column, and add a note to the **Enter a Reason here when modifying the Charged Size** column explaining the change.

Confirm the accuracy of any database backups

You reconcile database backups the same way you reconcile multistream backups. Find the policy name in the spreadsheet and locate the analyzed information in the `nbdeployutil-report-timestamp.log` file. Does the chosen day appear to correspond to a day upon which the complete database was backed up? If the information is inaccurate, change the **Charged Size** column, and add a note to the **Enter a Reason here when modifying the Charged Size** column explaining the change.

Locate full backup for snapshot image

Examine the backup policy attributes to determine if a backup image is ever created from the snapshot. If it is, rerun the analyzer with specific options to restrict the

collection and reporting to the specific client with a longer date range to find a full backup of the snapshot. If a backup image is never created from the snapshot, manually examine the snapshot or the client system to determine the size of the data.

Note: The log file that is associated with this report shows snapshot information.

Additional configuration

This chapter includes the following topics:

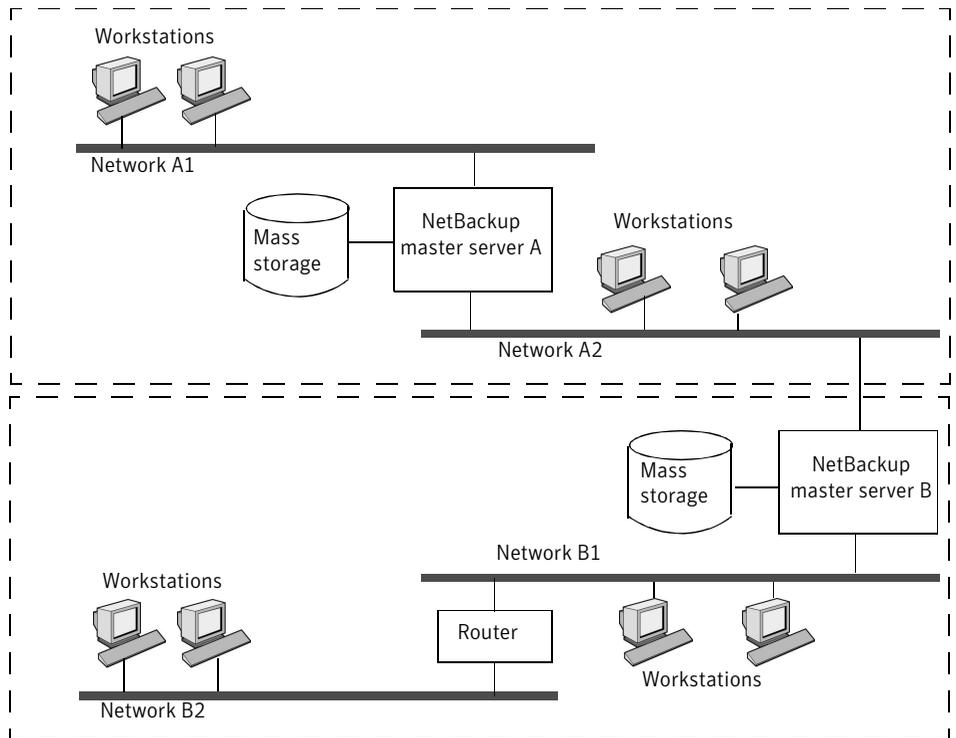
- [About multiple NetBackup master servers](#)
- [About multiple media servers with one master server](#)
- [About direct I/O for backups](#)
- [About dynamic host name and IP addressing](#)
- [About specifying the locale of the NetBackup installation](#)
- [Exporting PureDisk data to NetBackup](#)
- [About Shared Storage Option configuration](#)
- [About the vm.conf configuration file](#)

About multiple NetBackup master servers

For a large site, use multiple NetBackup master servers to optimize the backup loads. Divide the clients between the servers as necessary.

[Figure 3-1](#) shows a multiple-server configuration where the two sets of networks (A1/A2 and B1/B2) each have enough clients to justify separate servers.

Figure 3-1 Multiple master server scenario



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.

About multiple media servers with one master server

A protection domain refers collectively to the NetBackup master server, its NetBackup media servers, and its NetBackup clients. In a group of NetBackup servers, a client can have backups directed to any device on any server in the group.

Set up a NetBackup protection domain as follows:

- One master server, which controls all backup scheduling.
- Multiple media servers, which write the backup images to disk or removable media. They can have peripheral devices to provide additional storage.
- Multiple protected NetBackup clients, which send their data to the media servers.

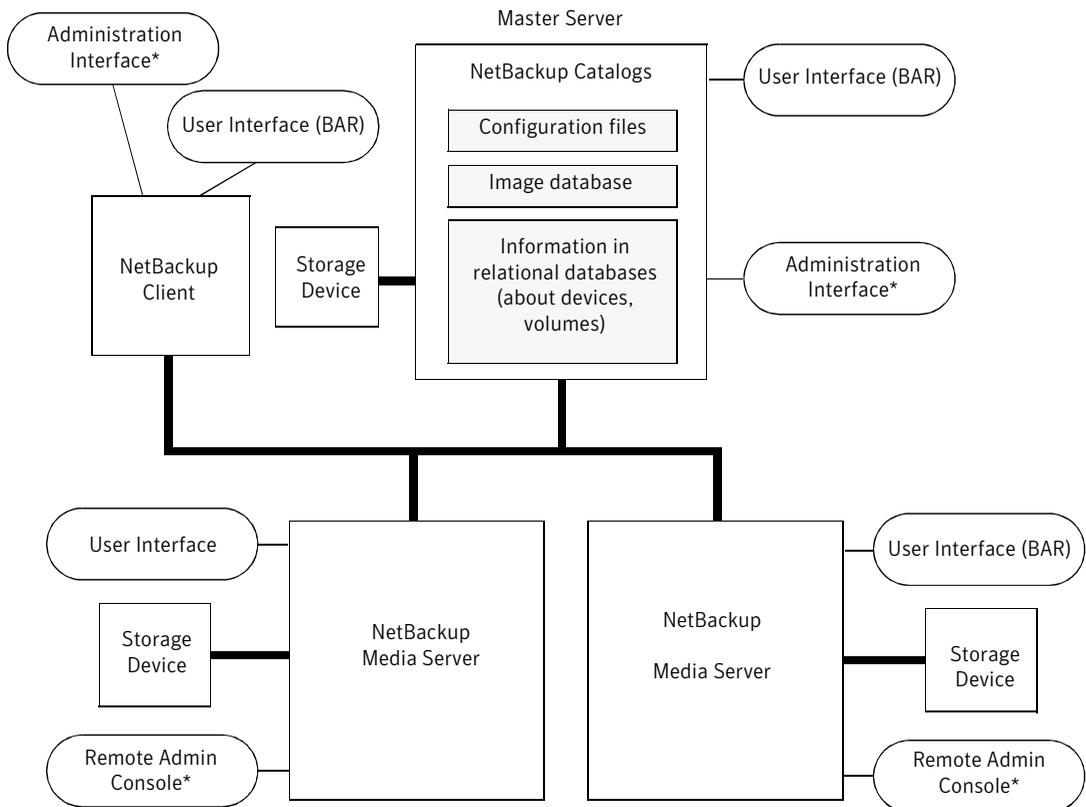
A common alternative strategy is to install extra peripherals on the clients that produce large amounts of data. The master server directs the data from the client to the client’s peripherals, which reduces network traffic because the data does not traverse the network. This strategy also distributes the backup load between the master and the media servers.

Important factors to remember about master and media servers are as follows:

- There can be only one master server in a group.
- A NetBackup master server is a media server for itself but cannot be a media server for another master server.

Figure 3-2 shows where software is installed and where the NetBackup catalogs are located (by default).

Figure 3-2 Catalog location using multiple media servers



* You can also use the Backup, Archive, and Restore user interface from a Windows client that has the Remote Administration Console installed.

About software on each server

Install NetBackup server software on each NetBackup server that has a peripheral that you want to include in a storage unit. The NetBackup installation program has choices for master and media server installation.

About NetBackup catalogs

The master server is the default location for the NetBackup catalogs. The catalogs include the media and the volume database (`emm_data.db`). The volume database contains the media usage information and the volume information that are used during the backups.

About direct I/O for backups

By default, the buffer size for disk storage units is 256 KB. If the buffer size is set to a value greater than 256 KB, backups written to that storage unit automatically use direct I/O. An increased buffer size can improve backup speed.

To increase the buffer size, the following conditions must be met:

- The storage unit must be owned by a Windows media server.
- The storage unit must be either a BasicDisk or an Array Disk storage unit.
- The backup to be stored cannot be multiplexed.
- The touch file that disables direct I/O must not be present.

`(install_path\VERITAS\NetBackup\bin\DISABLE_DIRECT_IO)`

To increase the buffer size, create one of the following touch files on the media server that owns the storage unit:

- For backups to disk

```
install_path\VERITAS\NetBackup\db\config\  
SIZE_DATA_BUFFERS_DISK
```

- For backups to disk or tape

```
install_path\VERITAS\NetBackup\db\config\  
SIZE_DATA_BUFFERS
```

If both touch files are present, `SIZE_DATA_BUFFERS_DISK` overrides the value in `SIZE_DATA_BUFFERS`. At this time, Symantec recommends that you use `SIZE_DATA_BUFFERS_DISK`.

Table 3-1 shows the possible values to include in `SIZE_DATA_BUFFERS_DISK` or `SIZE_DATA_BUFFERS ..`

Table 3-1 Absolute byte values for `SIZE_DATA_BUFFERS_DISK`, `SIZE_DATA_BUFFERS`

For a data buffer of this size (kilobytes),	enter this touch file value
32	32768
64	65536
96	98304
128	131072
160	163840
192	196608
224	229376
256	262144

Data buffer sizes continue in multiples of 32. Multiply the buffer size by 1024 for the touch file value.

A direct I/O backup triggers the following message: "Enabling direct I/O. Buffer size: *<buffer size>*."

Disabling direct I/O

Use the following procedure to disable direct I/O.

To disable direct I/O

- ◆ Create the following touch file on the media server that owns the storage unit:

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

About dynamic host name and IP addressing

Before making changes to a configuration, read this entire topic.

By default, a NetBackup server assumes that a NetBackup client name is the same as the network host name of the client machine. This assumption makes it difficult to back up any clients that have network host names that might change. For

example, a portable machine that plugs into a LAN and obtains IP addresses from a DHCP server. Or, a remote machine that dials into a PPP server. Use dynamic host name and IP addressing to define NetBackup clients that do not have fixed IP addresses and host names.

If dynamic addressing is used, remember that the NetBackup servers still require fixed IP addresses and host names.

All clients that are configured to use dynamic addressing and host names must trust each other, similar to the NetBackup altnames feature.

The following process is required to support the configurations that use dynamic IP addressing for NetBackup.

Table 3-2 Process to support the configurations that use dynamic IP addressing for NetBackup

Action	Process details/requirements
Configure the network to use a dynamic IP addressing protocol like DHCP.	<p>NetBackup requires that IP addresses of clients have a network host name.</p> <p>Be sure to define network host names for the range of dynamic IP addresses in the <code>hosts</code> file and (or) DNS on the network.</p>
Determine the NetBackup client names for the machines that have dynamic IP addresses and network host names.	<p>These NetBackup client names are used in other steps. Each NetBackup client must have a unique NetBackup client name. The NetBackup client name that is assigned to a client is permanent.</p>
Make changes on the master server, as described.	<ul style="list-style-type: none"> ■ Create NetBackup policies with client lists that include the new names. ■ Create entries in the NetBackup client database for the new client names. Use the <code>bpclient</code> command to create the entries.
Make changes on each dynamic NetBackup Windows client, as described.	<p>In the NetBackup Administration Console, in the left pane, click NetBackup Management. On the File menu, click Backup, Archive, and Restore. On the File menu, click NetBackup Client Properties. In the NetBackup Client Properties dialog box, select the General tab. Enter the correct NetBackup client name for the machine in the Client Name text box.</p>
On the master server, enable the Announce DHCP Interval option, as described.	<p>In the NetBackup Administration Console, in the left pane, expand NetBackup Management > Host Properties > Clients. Double-click on the the Windows client(s) in the right pane to open the Client Properties window. In the Client Properties window, in the left pane, expand Windows Client > Network. In the right pane, check the Announce DHCP Interval checkbox.</p>

Table 3-2 Process to support the configurations that use dynamic IP addressing for NetBackup (*continued*)

Action	Process details/requirements
Make changes on each dynamic NetBackup UNIX clients, as described.	<ul style="list-style-type: none"> ■ Modify the <code>bp.conf</code> file to include a <code>CLIENT_NAME</code> entry with the correct NetBackup client name for the machine. ■ Configure the system to notify the master server of the machine's NetBackup client name and current network host name during startup. The <code>bpdynamicclient</code> command is used to notify the master server. ■ Configure the system to notify periodically the master server of the machine's NetBackup client name and current network host name.

About setting up dynamic IP addresses and host names

Configure the network to use a dynamic IP addressing protocol. A protocol like DHCP has 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 can be assigned to NetBackup clients has a network host name. The host name should be defined in the host file, NIS, and DNS on the network.

For example, 10 dynamic IP addresses and host names are available.

The dynamic IP addresses and host names might be as follows:

```
123.123.123.70 dynamic00
123.123.123.71 dynamic01
123.123.123.72 dynamic02
123.123.123.73 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 that is assigned to a client is permanent and should not be changed. The client name that is assigned to NetBackup clients with dynamic IP addressing must not be the same as any network host names on the network. If the NetBackup client names are changed or are not unique, backup and restore results are unpredictable.

For example, 20 machines share the IP addresses as previously defined.

To make these machines NetBackup clients, assign them the following NetBackup client names:

```
nbclient01
nbclient02
nbclient03
nbclient04
.
.
.
nbclient20
```

Configuring the NetBackup master server

Use the following procedure to configure the NetBackup master server.

To configure the NetBackup master server

- 1 On the master server, create the NetBackup backup policies. For client name lists, use the NetBackup client names (for example, *nbclient01*) rather than the dynamic network host names (for example, *dynamic01*).
- 2 Create the client database on the master server.

The client database consists of directories and files in the following directory:

```
install_path\NetBackup\db\client
```

3 Create, update, list, and delete client entries with the `bpclient` command.

The `bpclient` command is in the following directory:

```
install_path\NetBackup\bin\admincmd
```

See “[bpclient commands that control client entries](#)” on page 46.

In the example, enter the following commands to create the 20 clients:

```
cd install_path\NetBackup\bin\admincmd
```

4 To see what is currently in the client database, run `bpclient` as follows:

```
install_path\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
```

The NetBackup client notifies the NetBackup server of its NetBackup client name and network host name. Then the Current Host, Hostname, and IP address fields display the values for that NetBackup client.

bpclient commands that control client entries

The `bpclient` command creates, updates, lists, and deletes client entries. The following table shows the `bpclient` commands that control client entries.

Table 3-3 `bpclient` commands that control client entries

Action	Command
Create a dynamic client entry	<pre>bpclient.exe -add -client <i>client_name</i> -dynamic_address 1</pre> <p>Where <i>client_name</i> is the NetBackup client name. The <code>-dynamic_address 1</code> argument indicates that the client uses dynamic IP addressing. It is possible to create entries with <code>-dynamic_address 0</code> for static IP addressing. However, to do so is unnecessary and adversely affects performance.</p>
Delete a client entry	<pre>bpclient.exe -delete -client <i>client_name</i></pre>
List a client entry	<pre>bpclient.exe -L -client <i>client_name</i></pre>
List all client entries	<pre>bpclient.exe -L -All</pre>

Configuring a dynamic Microsoft Windows client

Use the following procedure to configure a dynamic Microsoft Windows client.

To configure a dynamic Microsoft Windows client

- 1 If it is not already installed, install NetBackup on the Windows client.
- 2 In the **NetBackup Administration Console**, in the left pane, click **NetBackup Management**. On the menu bar, expand **File > Backup, Archive, and Restore**.
- 3 On the menu bar of the **Backup, Archive, and Restore** dialog box, expand **File > NetBackup Client Properties**.
- 4 In the **NetBackup Client Properties** dialog box, select the **General** tab. Change the **Client Name** to specify the NetBackup client name for the Windows client. Click **OK**.

- 5 In the **NetBackup Administration Console**, set **Announce DHCP Interval**. This value specifies how many minutes the client waits before it announces that it will use a different IP address.

To set the **Announce DHCP Interval**, return to the **NetBackup Administration Console**. In the left pane, expand **NetBackup Management > Host Properties > Clients**. Double-click on the the Windows client(s) in the right pane to open the **Client Properties** window. In the **Client Properties** window, in the left pane, expand **Windows Client > Network**. In the right pane, check the **Announce DHCP Interval** checkbox.

Additional information is available for **Announce DHCP Interval** in the *Administrator's Guide, Volume I*.

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.

- 6 On the client, stop and restart the NetBackup Client service to have the changes take effect.

Configuring a dynamic UNIX NetBackup client

Use the following procedure to configure a dynamic UNIX NetBackup client.

To configure a dynamic UNIX NetBackup client

- 1 If not already installed, install the NetBackup client software.
- 2 Edit the `/usr/opensv/netbackup/bp.conf` file. Use the `CLIENT_NAME` entry to specify the NetBackup client name for the machine, as follows:

```
CLIENT_NAME = nbclient00
```

- 3 Run 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 as follows:

```
bpdynamicclient -last_successful_hostname file_name
```

When `bpdynamicclient` starts up, it checks for the existence of *file_name*. If *file_name* exists, `bpdynamicclient` determines if the host name that is written in the file is the same as the current network host name. 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, create the following script in the `/etc/rc2.d` directory on a Solaris system:

```
# 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.

- 4 You must also create a root `crontab` entry to call the `bpdynamicclient` command periodically.

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
```

For DHCP, an acceptable interval to use between calls to `bpdynamicclient` is one-half of the lease period.

About specifying the locale of the NetBackup installation

The `/user/opensv/msg/.conf` file (UNIX and Linux) and the `install_path\VERITAS\msg\LC.CONF` file (Windows) contain information on the supported locales. These files define the date and the time formats for each supported locale. The `.conf` file and the `LC.CONF` file contain very specific instructions on how to add or modify the list of supported locales and formats.

The `.conf` file and the `LC.CONF` file are 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 the fifth fields define the date and the time fields and associated separators for that supported locale.

Modify the existing formats to change the default output.

For example, the TL line for the C locale is the following:

```
TL 1 C :hh:mn:ss/mm/dd/yyyy
```

An alternate specification to the order of months, days, and years is as follows:

```
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.

If the `.conf` file is not accessible, the default locales (TL lines) are:

```
TL 1 C :hh:mm:ss /mm/dd/yyyy  
TL 2 ov :hh:mm: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 defines the unrecognized locale. The fifth field defines the supported equivalent that is identified in the `TL` lines.

For example, use the following `TM` line to map the unrecognized locale French to the supported locale `fr`, the `TM` line is:

```
TM 6 french 2 fr
```

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, no default `TM` lines exist as the default locale is `C` (`ov`).

Exporting PureDisk data to NetBackup

NetBackup allows the export of PureDisk backups of Files and Folders data selections to NetBackup. NetBackup then can create copies of the data in NetBackup file format on NetBackup-supported media for possible disaster recovery purposes.

Required software and licenses

The PureDisk export capability is supported jointly by NetBackup and PureDisk.

- PureDisk requires release 6.1 MP1 or later.
For more information about software versions and licensing, see the *PureDisk Administrator's Guide*.
- The NetBackup DataStore license is required to provide the DataStore policy type selection.

Exporting PureDisk data

To export PureDisk backup data to NetBackup requires the creation of the following two active policies:

- An **Export to NetBackup** policy in PureDisk.

For information about creating an export policy, see the *NetBackup PureDisk Remote Office Edition Administrator's Guide*

- A DataStore type policy in NetBackup.
See “[Creating a NetBackup policy for a PureDisk backup data export](#)” on page 51.

The policies can be created in any order. The export takes place when the PureDisk policy runs.

Standard debugging techniques apply to both PureDisk and NetBackup. The VxBSA debug log is written to the `pdexport` directory.

Creating a NetBackup policy for a PureDisk backup data export

Use the following procedure to create a DataStore type policy for a PureDisk backup data export.

To configure a NetBackup policy for a PureDisk backup data export

- 1 Open the **NetBackup Administration Console**.
- 2 Select the master server that controls the media server to perform the export.
- 3 Select the **Policies** utility, then **Actions > New > New Policy**.
- 4 Enter a name for the policy. This name is entered in the **Parameters** tab in the PureDisk export policy.
- 5 Complete the following tabs in the **Add New Policy** dialog box:
 - **Attributes** tab
Select DataStore as the policy type. (The DataStore policy type selection appears if the DataStore license key is installed.) The compression and multiple data streams attributes are not supported for export because they are not supported upon restore. To run multiple streams, multiple export agents are required.
 - **Schedules** tab
By default, a DataStore policy type uses an Application Backup schedule. The start window for an Application Backup type is open every day for 24 hours.
You can adjust the default schedule or create a new schedule, but the start windows must coincide with the PureDisk Export policy start window.
 - **Clients** tab
In the **Clients** tab, add the name of the PureDisk export agent(s). (Multiple PureDisk export agents can indicate the same NetBackup DataStore policy. Add the export agents in the **Clients** tab as needed.) Do not include the name of the originating PureDisk clients.

■ **Backup Selections** tab

No entries are required on the **Backup Selections** tab.

6 Save and close the policy.

See “[Exporting PureDisk data](#)” on page 50.

Restoring PureDisk export data

Use the NetBackup client interface, **Backup, Archive, and Restore**, to restore the PureDisk export data to a PureDisk export agent. The system to which the data selections are restored must contain the NetBackup client software and the NetBackup engineering binary to support the export engine.

After the data is restored to the export agent, use a network transfer method to move the files to individual PureDisk clients.

To restore PureDisk export data

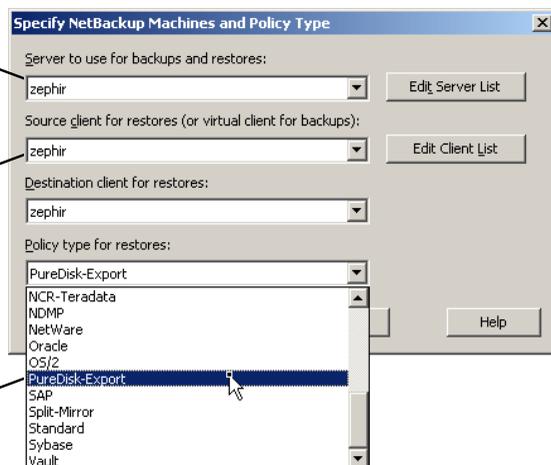
- 1 Open the NetBackup **Backup, Archive, and Restore** interface on the export agent.
- 2 Select **File > Specify NetBackup Machines and Policy Types**.
- 3 In the **Specify NetBackup Machines and Policy Types** dialog box, select the PureDisk-Export policy type to display the export data available for restore.

Although the NetBackup job runs as a DataStore policy type, the job is cataloged as a PureDisk-Export policy type under the name of the PureDisk agent.

Select the NetBackup server where the PureDisk backup data was exported to.

Select the PureDisk agent where the data was exported from. To restore the data, the agent must contain NetBackup client software.

Select PureDisk-Export as the policy type



- 4 Select a backup to restore from the NetBackup History.
- 5 Restore the files to the selected client as you would restore from a user-directed backup.

About restore support

NetBackup can restore only what PureDisk supports as part of its backups. For example, PureDisk does not provide access control list (ACL) support beyond UNIX file or directory permissions, so NetBackup cannot restore ACLs. See the PureDisk documentation for complete details.

Additional comments on restores includes the following:

- While Windows and UNIX security information can be restored, one limitation exists regarding restores to an alternate client for UNIX files. NetBackup backs up both the user ID and user name, but PureDisk backs up only the user ID. In non-PureDisk export backups during a restore to an alternate client, a user name can map to a different user ID. NetBackup performs a search for the user name and changes the user ID appropriately. For PureDisk export backups this ability is lost since the user name is not available. Files that are restored can belong to a different user.
- Windows files can be restored to UNIX systems and UNIX files can be restored to Windows systems. However, security information is lost when Windows files are restored to UNIX.

About Shared Storage Option configuration

The Shared Storage Option allows multiple NetBackup media servers to share individual tape drives (standalone drives or drives in a robotic library). NetBackup automatically allocates and unallocates the drives as backup and restore operations require.

The Shared Storage Option is a separately licensed and a separately purchased NetBackup software option that allows tape drive sharing. The license key is the Shared Storage Option key.

The Shared Storage Option is required only if multiple hosts share drives. For example, multiple NDMP hosts may share one or more drives.

The Shared Storage Option requires appropriate hardware connectivity, such as Fibre Channel hubs or switches, SCSI multiplexors, or SCSI-to-fibre bridges.

You can use Shared Storage Option in the following environments:

- Fibre Channel SANs

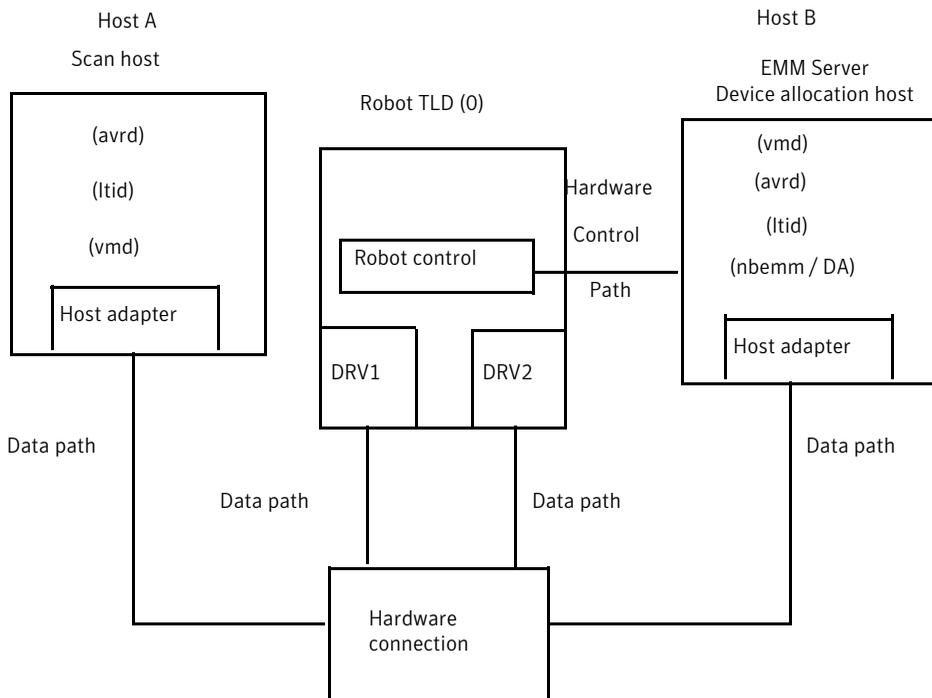
- Environments that do not use Fibre Channel, such as SCSI switches or multi-initiator configurations

About Shared Storage Option components

The NetBackup Enterprise Media Manager (EMM) manages media information. The Enterprise Media Manager also is the device allocator (DA) for shared drives.

Figure 3-3 shows an example of a shared drive configuration.

Figure 3-3 Shared Storage Option example



In this figure, the following describes Host A:

- It is connected to drives DRV1 and DRV2 through SAN hardware.
- Host A is the first host in the environment to come online with a non-zero scan ability factor. Therefore, it is the initial scan host for its drives. See “About scan hosts” on page 55.

In this figure, the following describes Host B:

- It is connected to drives DRV1 and DRV2 through SAN hardware.

- It is configured to be the EMM server, which is also the device allocation host. See [“About the device allocation host”](#) on page 57.
- It controls the robotics. Except for ACS or TLM robot types, only one robot control host exists for each robot.
- It can be configured (optionally) as a highly available (HA) server.

For a process flow diagram of Shared Storage Option components, see the *NetBackup Troubleshooting Guide*.

About SAN media servers

SAN media servers are NetBackup media servers that back up their own data. SAN media servers cannot back up the data that resides on other clients.

SAN media servers are useful for certain situations. For example, a SAN media server is useful if the data volume consumes so much network bandwidth that it affects your network negatively.

When you define a backup policy for a SAN media server, add only the SAN media server as the client.

The NetBackup Shared Storage Option can use NetBackup SAN media servers.

About SSO and the NetBackup EMM server

To coordinate network-wide allocation of tape drives, the NetBackup Enterprise Media Manager (EMM) manages all shared tape requests in a SAN. EMM responds to requests from multiple instances of NetBackup master servers, media servers, and NetBackup SAN media servers.

For shared drive configurations, the host that is configured as the EMM server is also known as the device allocation host.

See [“About the device allocation host”](#) on page 57.

EMM maintains shared drive and host information. Information includes a list of hosts that are online and available to share a drive and which host currently has the drive reserved. The Media Manager device service (`ltid`) requests shared drive information changes.

About scan hosts

Scan hosts are a component of the NetBackup Shared Storage Option.

Each shared drive has a host that is identified as the scan host. A scan host is the host from which the automatic volume recognition process (`avrd`) scans unassigned

drives. (The robotic daemons scan assigned drives.) A scan host must have data path access to the drive.

The EMM database contains the shared drive information; that information includes the scan host. Media servers receive drive status information from the EMM server.

How the scan host is determined

EMM determines scan hosts; a scan host may be different for each shared drive. The first host in the environment to come online with a non-zero scan ability factor is the initial scan host for its drives.

To configure the scan ability factor of media servers, use the `nbemmcmd` command. For more information, see *NetBackup Commands Reference Guide*.)

The scan host can change

A scan host is assigned for a shared drive until some interruption occurs.

For example, if one of the following occurs, EMM chooses a new scan host:

- The socket connection, the host, the drive, the drive path, or the network goes down.
- The drive is logically placed in the Down mode.

The scan host temporarily changes to hosts that request tape mounts while the mount is in progress. Scan host changes occur so only one host at a time has access to the drive path.

Drive paths for the scan host

If a drive has multiple paths that are configured on the selected scan host, EMM selects a scan path as follows:

- The first local device path it finds in its database in the UP state.
- The first NDMP-attached drive path it finds in its database in the UP state.

Shared tape drive polling For shared tape drives, only the scan host polls drives until a mount request is received from NetBackup. During a mount request, NetBackup uses the host that requests the mount to poll the shared drive.

This design enables NetBackup to support Dynamic Loop Switching or SAN zones. Each tape drive needs to be detected only from a single host. Each tape drive can potentially have its own scan host that switches dynamically to process errors and continue availability. A central device arbitrating component manages scan host assignments for shared drives. The arbitrating component also provides a network drive reservation system so that multiple NetBackup media servers can share a drive.

Polling a shared tape drive allows dynamic loop switching and reduces the number of device accesses and reduces CPU time. However, it cannot detect connectivity breaks (for example, discontinuity in the Fibre Channel fabric) until I/O occurs.

About the device allocation host

The device allocation host is another name for the EMM server, when the EMM server performs device allocation tasks for Shared Storage Option.

About reserving or releasing shared devices

The Shared Storage Option does not load firmware in SAN devices or communicate with hub or switch APIs. The Shared Storage Option can communicate with hub or switch APIs if you use the NetBackup `shared_drive_notify` script.

NetBackup runs the `shared_drive_notify` script when a shared drive is reserved or released.

The script requires the following parameters:

- The name of the shared drive.
- The name of the current scan host.
- The operation, which is one of the following:

RESERVED	The host on which the script is executed needs SCSI access to the drive until it is released.
----------	---

ASSIGNED	Informational only. It does not change the fact that the host that reserved the drive needs SCSI access.
----------	--

RELEASED	Only the scan host needs SCSI access to the drive.
SCANHOST	The host that executes the script has become the scan host. A host should not become a scan host while the drive is RESERVED. The scan host may change between a RESERVED operation and a RELEASED operation.

The script resides in the following directory:

```
Install_path\VERITAS\Volmgr\bin
```

Note: The script must be executable by the root user.

The script exits with status 0 upon successful completion.

How to share robotic libraries without using the Shared Storage Option

You can share robotic tape libraries among multiple NetBackup media servers by using any of the following methods:

- **Shared library support**
NetBackup allows different drives within the same robotic library to be configured on different media servers. This capability is termed shared library support. Robot types that support shared library are ACS, TL8, TLD, TLH, TLM.
- **Partitioned libraries**
Some robot vendors also let you partition libraries. One partitioned view of the robotic library includes one set of drives, while the other view has a different set of drives in the library. Partitions let two robotic control daemons on different control hosts manage the robotic library – possibly each for a different NetBackup master and media server environment.
- **Multiple master servers**
Use multiple NetBackup master servers that share a common media and device management domain. This means that the master servers use the same EMM server.

These capabilities are not related to Shared Storage Option and should not be confused with Shared Storage Option.

Shared Storage Option terms and concepts

[Table 3-4](#) describes the terms and the concepts relevant to understanding the Shared Storage Option.

Table 3-4 Shared Storage Option terms and concepts

Term	Definition
Backup Exec Shared Storage Option	The NetBackup Shared Storage Option is not the same as the Symantec Backup Exec Shared Storage Option. The Backup Exec SSO does not include support for UNIX servers and uses a different method for drive arbitration.
SAN media servers	A NetBackup SAN media server backs up its own data to shared drives. It cannot back up data on other NetBackup hosts or clients. Symantec licenses NetBackup SAN media servers.
Shared drive	When the Shared Storage Option is installed, a tape drive that is shared among hosts is termed a shared drive. For the drives that are attached to NDMP hosts, each NDMP attach host is considered an additional host.

About the Shared Storage Option license key

The Shared Storage Option is a feature that is licensed separately from base NetBackup. The NetBackup Shared Storage Option license key is based on the number of physical tape drives to share. The key activates NetBackup to share the specific number of physical drives for which you are licensed.

See [“Licensing the Shared Storage Option”](#) on page 59.

Licensing the Shared Storage Option

No special installation is required for the Shared Storage Option. When NetBackup software is installed, the Shared Storage Option software also is installed. However, you must activate the feature by entering the Shared Storage Option license key.

Note: Enter the license key on the NetBackup master server. Also enter the license key on each NetBackup media server that you use for the Shared Storage Option.

To license Shared Storage Option

- 1 To add a license to a specific server, on the **File** menu, click **Change Server** and then select the server.
- 2 In the **NetBackup License Keys** dialog box, click **New**.
- 3 In the **Add a New License Key** dialog box, enter the license key and click **Add** or **OK**.

- 4 Click **Close**.
- 5 Restart all the NetBackup services and daemons.

See [“About the Shared Storage Option license key”](#) on page 59.

About Shared Storage Option prerequisites

To configure your hardware for use with Shared Storage Option, you must ensure that the following prerequisites are satisfied:

- Configure your SAN environment.
- Attach robots and drives.
- Ensure that all of the servers recognize the shared devices. Device recognition may depend on operating system configuration, as follows:
On Windows servers, Windows recognizes devices automatically. However, in some instances you may have to install device drivers.

Some of the following tasks may be optional depending on your hardware:

- Determine the physical location of each drive within the robot. Location usually is shown on the connectors to the drives or in the vendor documentation.
This task may not be required if NetBackup device discovery accurately determines drive location within the robot.
- Connect all drives and all robots.
- Install SAN connecting hardware (for example, bridges, switches, or hubs).
- If fiber is part of your configuration and you use a SCSI-to-fiber bridge, determine the SCSI-to-Fibre Channel mapping for your tape devices.
Hard-wired SCSI IDs are converted to Fibre Channel logical unit numbers (LUNs) that the hosts read. To ensure correct drive assignments, you should know which LUNs map to which physical SCSI IDs. Use persistent LUN mapping if possible.
Familiarity with the hardware and various vendor configuration tools help you accomplish this task. See the vendor documentation for your bridge.
- Record the physical configuration.
When you set up a Shared Storage Option configuration, record your hardware information. Record the adapter, SCSI addresses, World Wide Names (WWNs), and Fibre Channel LUNs to which you connected each drive. Also, record the version levels of firmware and drivers.
- Install and configure the appropriate drivers. See your vendor documentation for instructions.

- For instructions on how to configure the HBA on Windows servers, see the HBA documentation from the vendor.
- Use any available hardware configuration interface to configure and ensure that the configuration is what you expect. For example, on Windows servers you can use the Hyperterminal interface to configure SCSI-to-fibre bridges.
Use the following order when you configure and verify the hardware:
 - Robot and shared drives
 - Bridges
 - Hub or switches
 - Hosts
- If errors occur and you suspect the operating system, refer to the operating system logs as described in your operating system documentation.

About hardware configuration guidelines

The following are hardware configuration guidelines:

- If you use SAN hardware from multiple vendors, problems may occur. Always use a SAN configuration and use firmware levels that the hardware vendor supports.
- Consult SAN device, HBA, and operating system documentation to determine how to configure operating system tape drivers and pass-through drivers to detect your SAN devices.
- Check your hub timer settings.
- Use hard arbitrated loop physical addresses rather than soft addresses. Consult with hardware suppliers to verify the recommended usage of their products.
- Check the firmware levels of all your Fibre Channel hardware (for example, bridges). Use the most recent firmware level that is known to operate with other SAN hardware devices.
- Try to duplicate SAN issues and problems using commands and utilities on the host operating system.
- Test both backup and restore capabilities. Backup jobs may complete successfully, but the data may be corrupted. For example, incorrect switch settings may cause problems.
- Ensure that your hardware and SAN configuration are operational and stable before adding Shared Storage Option software.

- Test backup and restore capabilities with dedicated tape drives before you configure them as shared drives.
- For large configurations, begin drive sharing with a few tape drives and two or three media servers (or NetBackup SAN media servers).
- Configuration and troubleshooting processes are easier on smaller configurations. If possible, create multiple and independent Shared Storage Option configurations with subsets of servers sharing subsets of SAN-attached drives.
- Use the correct start order for your Fibre Channel hardware, as follows:
 - Robots or drives
 - Bridges
 - Hubs or switches
 - Hosts
- The start sequence is longer for some devices than others. To verify that the hardware starts completely, examine indicator lights. A green light often indicates a completed start sequence.

About installing and configuring drivers

On the media server systems, install and configure drivers and modify the appropriate system configuration files.

Guidance about the NetBackup requirements is available.

See the *NetBackup Device Configuration Guide*.

Verifying the connectivity

Test your hardware configuration before you configure Shared Storage Option in NetBackup. This task is very important and is often overlooked.

Note the following points:

- Verify that all of your servers (master and media) can communicate with one another. To do so, use the `ping` command from each server to every other server. Be sure to `ping` by host name to verify that the name resolution methods function properly.
- Use the NetBackup `bpc1ntcmd` utility to resolve IP addresses into host names. For more information, see the *NetBackup Troubleshooting Guide* and the *NetBackup Commands Reference Guide*.

- Use operating system and NetBackup commands and tools to verify that the devices are configured correctly. Make sure that the operating system detects the devices on the SAN before you configure the Shared Storage Option. If the configuration does not work in the operating system, it does not work for the Shared Storage Option.
For example, on Solaris systems you can use the `mt -f tapename status` command to determine tape drive status.
- For more information and examples, see the appropriate operating system chapter in the *NetBackup Device Configuration Guide*.

About configuring the Shared Storage Option in NetBackup

You must configure your shared drives, storage units, and backup policies.

About configuring SSO in NetBackup	See “About configuring SSO in NetBackup” on page 63.
Configuring Shared Storage Option devices in NetBackup	See “Configuring Shared Storage Option devices in NetBackup” on page 63.
About adding Shared Storage Option configuration options	See “About adding Shared Storage Option configuration options” on page 64.
About configuring NetBackup storage units and backup policies	See “About configuring NetBackup storage units and backup policies” on page 64.

About configuring SSO in NetBackup

Symantec recommends that you use the Device Configuration Wizard to configure Shared Storage Option in NetBackup. Identifying devices when you configure shared devices is difficult, and the wizard increases the likelihood of a successful configuration.

With the Device Configuration Wizard, you should configure all shared drives from one host (usually the master server). Launch the wizard only one time with the current host set to the master server. You then indicate a list of media servers or NetBackup SAN media servers (in the Device Hosts screen). The wizard configures devices on all of the media servers you selected, and these hosts read the shared configuration information.

Configuring Shared Storage Option devices in NetBackup

Symantec recommends that you use the **Device Configuration Wizard** to configure shared drives. The wizard guides you through the steps to configure shared drives.

Be sure to review the limitations of the wizard in the wizard help.

To start the Device Configuration Wizard

- 1 In the **NetBackup Administration Console**, expand **Media and Device Management**.
- 2 Click **Configure Storage Devices**.

About adding Shared Storage Option configuration options

You can fine-tune your configuration by adding Shared Storage Option options to the `vm.conf` Media Manager configuration file.

See “[About the vm.conf configuration file](#)” on page 74.

About configuring NetBackup storage units and backup policies

You must configure storage units and policies for your shared drives. If you used the **Device Configuration Wizard** to configure the shared drives, you may have configured storage units and policies already.

For more information, see the following:

- *NetBackup Administrator’s Guide for UNIX and Linux, Volume I*
- *NetBackup Administrator’s Guide for Windows, Volume I*.

Configure storage units and backup policies as follows:

Configuring storage units for each media server In each storage unit definition, logically define the robot and the shared drives for that media server. For the **Maximum concurrent drives** used for backup, specify the total number of all shared drives in the robot. When you configure storage units, select a single media server. Alternatively, you can allow NetBackup to select the media server to use at the time of the backup. For example, you can configure a single storage unit that any media server that shares the storage unit can use.

Configuring a backup policy for each media server How you define a policy for a media server depends on your media server license, as follows:

- For a media server that is licensed for Shared Storage Option, the policy can back up the media server and any other NetBackup clients.
- For a NetBackup SAN media server, only the SAN media server can be backed up.

A license for a regular media server provides the greatest flexibility; a license for a NetBackup SAN media server is more restrictive.

For a policy for the clients that you want to back up anywhere in your configuration, you can choose any available storage unit. Alternatively, you can use storage unit groups (prioritized storage units).

Verifying your Shared Storage Option configuration

In a Shared Storage Option configuration, a shared drive must have the same logical name (drive name) on all of the NetBackup media servers. If the drive resides in a robotic library, it also must use the same drive number in the library. This section describes some tools you can use to verify your configuration.

How you verify that your configuration is set up correctly depends on your devices and how you configured Shared Storage Option, as follows:

- If you have serialized devices, Symantec recommends that you use the Device Configuration Wizard. The wizard verifies your configuration.
- If you have non-serialized devices, see the Symantec support site for a tech note with instructions about how to verify your configuration. The tech note title is "Verifying a Shared Storage Option (SSO) Configuration with Non-Serialized Devices."
- If you have serialized devices but you did not use the Device Configuration Wizard, use the following procedure to verify your configuration.

The verification procedures use the following NetBackup commands:

- `install_path\VERITAS\Volmgr\bin\scan`
- `install_path\VERITAS\Volmgr\bin\tpconfig`

In the following example the ADIC robotic library has six drives, but only drives 5 and 6 are configured on this particular host.

About Shared Storage Option configuration

Perform the verification on all of the NetBackup servers in your configuration. Ensure that each shared drive has the same logical drive name and same drive number ID on each media server that shares the drive.

To verify a manually-configured Shared Storage Option configuration

- 1** Execute `tpconfig -d` or `tpconfig -dl`. For NDMP devices, use `tpautoconf -probe -ndmp_host_name host_list`.

The output from `tpconfig` shows the logical names NetBackup assigns to tape drives. The following example shows drive number 5 is named QUANTUM.DLT7000.000 and drive number 6 is named QUANTUM.DLT7000.001:

Id	DriveName	Type	Residence	Status
	Drive Path			

0	QUANTUM.DLT7000.000	dlt	TLD(0) DRIVE=5	
	/dev/st/nh3c0t510			UP
1	QUANTUM.DLT.7000.001	dlt	TLD(0) DRIVE=6	
	/dev/st/nh3c0t110			UP

Currently defined robotics are:

```
TLD(0)    robotic path = /dev/sg/h3c0t010
EMM server = norway
```

- 2 Execute the `scan` command. The `scan` output shows the robot and the drive properties.

The following is example output:

```
*****
***** SDT_TAPE *****
***** SDT_CHANGER *****
*****
Device Name : "/dev/sg/h3c0t010"
Passthru Name: "/dev/sg/h3c0t010"
Volume Header: ""
Port: -1; Bus: -1; Target: -1; LUN: -1
Inquiry : "ADIC Scalar 100 3.10"
Vendor ID : "ADIC "
Product ID : "Scalar 100 "
Product Rev: "3.10"
Serial Number: "ADIC009K0340314"
WWN : ""
WWN Id Type : 0
Device Identifier: ""
Device Type : SDT_CHANGER
NetBackup Robot Type: 6
Removable : Yes
Device Supports: SCSI-2
Number of Drives : 6
Number of Slots : 50
Number of Media Access Ports: 10
Drive 1 Serial Number : "PXB03S0979"
Drive 2 Serial Number : "PXB03S0913"
Drive 3 Serial Number : "CXA04S2051"
Drive 4 Serial Number : "PXA31S1787"
Drive 5 Serial Number : "PXA37S3261"
Drive 6 Serial Number : "PXA50S2276"
Flags : 0x0
Reason: 0x0
-----
Device Name : "/dev/st/nh3c0t510"
Passthru Name: "/dev/sg/h3c0t510"
Volume Header: ""
Port: -1; Bus: -1; Target: -1; LUN: -1
Inquiry : "QUANTUM DLT7000 2561"
Vendor ID : "QUANTUM "
Product ID : "DLT7000 "
```

```

Product Rev: "2561"
Serial Number: "PXA37S3261"
WWN          : ""
WWN Id Type  : 0
Device Identifier: ""
Device Type   : SDT_TAPE
NetBackup Drive Type: 9
Removable     : Yes
Device Supports: SCSI-2
Flags : 0x4
Reason: 0x0
-----
Device Name   : "/dev/st/nh3c0t110"
Passthru Name: "/dev/sg/h3c0t110"
Volume Header: ""
Port: -1; Bus: -1; Target: -1; LUN: -1
Inquiry      : "QUANTUM DLT7000          296B"
Vendor ID    : "QUANTUM "
Product ID   : "DLT7000          "
Product Rev  : "296B"
Serial Number: "PXA50S2276"
WWN          : ""
WWN Id Type  : 0
Device Identifier: ""
Device Type   : SDT_TAPE
NetBackup Drive Type: 9
Removable     : Yes
Device Supports: SCSI-2
Flags : 0x4
Reason: 0x0

```

- 3 For each tape drive in the `tpconfig` output, do the following:**
- Use the device file name from the `tpconfig` output to locate the tape drive in the `scan` output.
 - Step 1 shows device file pathnames `/dev/st/nh3c0t5l0` and `/dev/st/nh3c0t1l0`.
 - Determine the serial number of the drive in the scan output. "Tape" in the device type field identifies a tape drive.
 - Step 2 shows example `scan` output shows the following:
 - The drive `/dev/st/nh3c0t5l0` serial number is PXA37S3261.
 - The drive `/dev/st/nh3c0t1l0` serial number is PXA50S2276.

- Verify that the serial number for the drive matches the serial number in the output from the robot section of scan. "Changer" in the device type field identifies a robot.
 In the previous examples, the serial numbers match.

Device Monitor and Shared Storage Option

You can use the NetBackup Administration Console Device Monitor to obtain information about your Shared Storage Option configuration and manage your shared drives. See the following:

For more information about the Device Monitor, see the *NetBackup Administrator's Guide, Volume I*.

[Table 3-5](#) describes information you can glean from the NetBackup Administration Console Device Monitor.

Table 3-5 Device Monitor information

Action	Information
Drive Status pane	The Control and Device Host columns contain shared drive information.
Changing the operating mode for a shared drive	For a shared drive, the Change Mode dialog contains a list of all paths to the selected drive. You can choose any number of paths to which the mode change applies.
Adding or changing a comment for a shared drive	For a shared drive, the Change Drive Comment dialog box contains the following: <ul style="list-style-type: none"> ■ A list of all paths to the selected drive ■ The current drive comment for each combination. You can choose any number of paths to which the changes apply.
Performing drive cleaning functions for a shared drive	The three available drive cleaning functions are used with shared drives are as follows: <ul style="list-style-type: none"> ■ Clean Now In the list of hosts that share the drive, you can choose only one host on which the function applies. ■ Reset Mount Time In the list of hosts that share the drive, you can choose any number of hosts on which the function applies. ■ Set Cleaning Frequency Supported for shared drives.

Viewing SSO summary reports

You can view Shared Storage Option Summary reports.

See “[Shared Storage Option summary reports](#)” on page 71.

To view SSO summary reports

- 1 In the **NetBackup Administration Console**, in the left pane, expand **Media and Device Management > Device Monitor**.
- 2 On the **Actions** menu, select **View Status of Shared Drives**.
- 3 In the **Status of Shared Drives** dialog box, select a device allocation host (or hosts) from the list.
- 4 Use **Add** to move the host to the list of hosts to scan.
- 5 Click **OK**.

The **Shared Drive Summary** and **Device Allocation Host Summary** appear in the two lower panes of the dialog.

Shared Storage Option summary reports

The following topic applies only to NetBackup Enterprise Server.

The following two reports contain the following information about the drives and hosts:

- The Shared Drive Summary shows the following:
 - Drive name
 - Device allocation host
 - Number of registered hosts
 - Drive reservation status
 - Hosts that reserve the drive
 - Current scan host
- The Device Allocation Host Summary shows the following:
 - Device allocation host
 - Host name of the registered host
 - Number of registered and reserved drives
 - Availability status
 - Scan ability factor

- Scan status (if the host is scan host for at least one SSO drive)

Operating system assistance

If errors occur during the installation or configuration of the shared devices and you suspect problems with the operating system, refer to the following:

- Operating system logs, as described in the operating system documents.
- NetBackup logs.
- Operating system man pages (UNIX or Linux servers only).
- The *NetBackup Device Configuration Guide*.

Common configuration issues with Shared Storage Option

If you cannot obtain full functionality after you configure SSO, consider the following:

- Verify that the SAN hardware uses current firmware or drivers. Hardware includes hubs, switches, HBAs, and bridges.
- Verify that the JNI HBA failover value was set to zero to avoid I/O hangs. This value applies to bridges and HBAs.
- Verify that the HBAs with the SCSI-3 protocols are compatible with the operating system drivers.
- Verify that your cluster configuration is supported.
For more information about cluster configuration, see the *NetBackup Release Notes*
- Verify that all of your Fibre Channel devices support your Fibre Channel topology. For example, in a switched fabric topology, ensure that all devices supported switched fabric.
- Verify that Shared Storage Option is licensed on each server. To do so, select Help > License keys from the NetBackup Administration Console on each server. To enable Shared Storage Option, enter the Shared Storage Option license key on each server.
- Verify that you configured Shared Storage Option from the master server. You must configure SSO from the master server not from a media server (or SAN media server).
- Verify that you configured the same robot control host on every host. Remember that except for ACS and TLM robot types, only one host controls the robot.

- Verify that you used the Device Configuration Wizard rather than the `tpconfig` utility to configure Shared Storage Option. The wizard coordinates configuration with all hosts that share the drives. The `tpconfig` utility may create inconsistent configurations.
- Verify that you selected the appropriate device hosts in the Device Configuration Wizard, including the host with robotic control.
- Fibre Channel connections to the drives and the robots cause increased complexity in a NetBackup device configuration. On some operating systems, SCSI-to-fibre bridges may result in inconsistencies in the device paths when you restart a host. After a restart of the host, the device configuration should be verified.
- Verify that names across all systems that share the drives are consistent.
- Test the drive paths on every media server.
- Define NetBackup storage units for each media server. Do not select any available media server in the storage units.
- Verify that you did not interrupt a data path during a backup. If you do, the NetBackup job fails. It can fail with media write errors or it may hang and have to be terminated manually.
- Verify that you do not use Berkeley-style close on the tape path (UNIX or Linux servers only).
- On Solaris systems, verify the following:
 - That you added tape configuration list entries in `/kernel/drv/st.conf` (if needed).
 - That you defined configuration entries for expanded targets and LUNs in `sg.links` and `sg.conf` files. If you see problems with the entries in the `/etc/devlink.tab` file (created from `sg.links`), verify the following:
The first entry uses hexadecimal notation for the target and LUN. The second entry uses decimal notation for the target and LUN.
Use a single tab character between the entries; do not use a space or a space and a tab character.
 - That you configured the operating system to force load the `sg/st/fcaw` drivers.

For more information, see the Solaris chapter of the *NetBackup Device Configuration Guide*.

Frequently asked questions about Shared Storage Option

Q. What combinations of SAN hardware components are supported for Shared Storage Option?

A. Shared Storage Option works with many hardware combinations. Symantec has an open policy on hardware support for Shared Storage Option. Consult your hardware suppliers to verify the interoperability of their products.

A list of SAN components that have been tested with NetBackup is available on the Symantec support Web site:

<http://entsupport.symantec.com>

Q. If NetBackup allocates four drives to a server and it finishes with two of the drives, does NetBackup reallocate the two drives? Or does NetBackup wait until the backup schedule that uses the four drives is completely finished before it reallocates the drives?

A. The two available drives are reallocated and used. NetBackup monitors drive status and notifies the NetBackup scheduler of drive availability.

Q. Does NetBackup Shared Storage Option use the IP protocol or the SCSI protocol?

A. Both. IP protocol is used to provide coordination between servers. Shared Storage Option uses SCSI protocol (SCSI reserve) as an added layer of protection.

About the `vm.conf` configuration file

The `vm.conf` file contains configuration entries for media and device management. NetBackup can create this file, but if it does not exist, you must create it.

pathname is `install_path\Volmgr\vm.conf`.

Various NetBackup components read this configuration file on the host where the component runs. The NetBackup component is a command, daemon, process, or utility. The host can be a NetBackup administration client or a server where administration operations are requested.

See “[Example `vm.conf` file](#)” on page 90.

ACS_mediatype entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
ACS_mediatype = Media_Manager_mediatype
```

If this entry is used in `vm.conf`, the ACS media type is mapped to the specified Media Manager media type. More than one `ACS_mediatype` entry can be specified.

This entry is read and interpreted on the host on which `vmcheckxxx` and `vmupdate` run during a robot inventory operation. Use this entry on every NetBackup media server that functions as an ACS robot control host.

A list of the valid `ACS_mediatype` entries is available.

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

ACS_SEL_SOCKET entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
ACS_SEL_SOCKET = socket_name
```

By default, `acs_sel` listens on socket name 13740. If this entry is specified in `vm.conf`, the default can be changed. This entry is read and interpreted on the host on which `acsd` runs.

ACS_SSI_HOSTNAME entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
ACS_SSI_HOSTNAME = host
```

Use `ACS_SSI_HOSTNAME` to specify the host to which RPC return packets from ACS library software are routed for ACS network communications. By default, the local host name is used. This entry is read and interpreted on the host on which `acsd` and `acsssi` run. Do not use the IP address of the host for this parameter.

ACS_SSI_SOCKET entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
ACS_SSI_SOCKET = ACS_library_software_hostname socket_name
```

Valid value for `ACS_library_software_hostname` is the host name of the ACS library host. Do not use the IP address of the ACS library host for this parameter.

By default, `acsssi` listens on unique, consecutive socket names; the names begin with 13741. If this entry is specified in `vm.conf`, specify socket names on an ACS library software host basis. This entry is read and interpreted on the host where `acsd` and `acsssi` are running.

ADJ_LSM entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
ADJ_LSM = robot_num ACS_ID,LSM_ID ACS_ID,LSM_ID
```

In an ACS robot with multiple library storage modules (LSMs), pass-through mechanisms can move ejected media to the media access port (MAP). A pass-through mechanism passes media from one LSM to another. This travel time can be excessive when media must pass through several LSMs.

Use this entry to specify the physical orientation of the LSMs in an ACS robot. If this entry is specified in `vm.conf`, you do not need to know which MAP (or ACS CAP) to select for efficient ejects. NetBackup determines the appropriate MAP to complete the media eject by using a nearest-MAP algorithm.

This nearest-MAP algorithm is based on the physical orientation of the LSMs that defined with this entry. This algorithm is only for the cases where more than one MAP is requested to handle the eject. If this algorithm is used, any `MAP_ID` entries in `vm.conf` are ignored.

Note: nearest-MAP capability is only available by using the `vmchange` command with the `-map` option or the Vault administrative interface. It is not available from the NetBackup Administration Console.

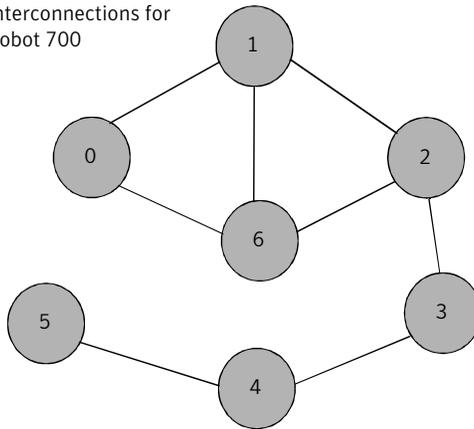
Without this entry present, NetBackup assumes that all LSMs are interconnected with pass-through ports, except for the first LSM and the last LSM. The LSMs are interconnected in a line formation.

robot_num is the robot number. *ACS_ID* and *LSM_ID* are the coordinates of the LSM.

Figure 3-4 is a diagram of LSM interconnections that are described by the following entries:

```
ADJ_LSM = 700 0,0 0,1
ADJ_LSM = 700 0,0 0,6
ADJ_LSM = 700 0,1 0,2
ADJ_LSM = 700 0,1 0,6
ADJ_LSM = 700 0,2 0,6
ADJ_LSM = 700 0,2 0,3
ADJ_LSM = 700 0,3 0,4
ADJ_LSM = 700 0,4 0,5
```

The robot has pass-through mechanisms between 7 LSMs.

Figure 3-4 Pass-through exampleInterconnections for
Robot 700

API_BARCODE_RULES entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server as follows:

```
API_BARCODE_RULES
```

If this entry is specified in `vm.conf`, barcode rule support for API robots is enabled.

NetBackup barcode rules allow default media mappings to be overridden. Barcode rules are especially useful when multiple generations of the same tape drive use the same type of media.

For example STK 9940A and STK 9940B drives use STK1R media, but write data at different densities. The drive must be configured by using different drive types such as HCART or HCART2. Specify a barcode rule for a series of bar codes to configure some of the media as HCART2. Other STK1R media not in this barcode range are configured as HCART (the default for STK1R). Without this entry, a robot inventory operation configures all media of type STK1R as either HCART or HCART2, depending on how the drive was configured.

AUTHORIZATION_REQUIRED entry in `vm.conf`

This entry specifies that NetBackup should use the `vm.conf` file `SERVER` entry to control which hosts can monitor and control devices on this host. This entry is read and interpreted on the media server on which the NetBackup `vmd` service runs, as follows:

```
AUTHORIZATION_REQUIRED
```

If this entry is specified in `vm.conf`, the `vm.conf` file also must include a `SERVER` entry for every media server that controls devices on this host.

If no `AUTHORIZATION_REQUIRED` entry exists and no `SERVER` entries exist, any NetBackup server can monitor and control devices on this host.

For maximum security, Symantec recommends that you use this entry and `SERVER` entries.

This entry is read and interpreted on media servers on which the NetBackup `vmd` service runs.

AUTO_PATH_CORRECTION entry in `vm.conf`

If this entry is specified in `vm.conf`, it specifies whether automatic device path remapping is enabled or disabled, as follows:

```
AUTO_PATH_CORRECTION = YES|NO
```

If the value is `NO`, the device configuration remains unchanged when the NetBackup Device Manager service (`ltid`) is started. Therefore, the saved device configuration may be different than the actual configuration after devices are changed and the server is restarted.

If the value is `YES`, NetBackup tries to discover attached devices and then automatically update the device configuration for any device paths that are incorrect. On Windows computers, this entry is read and interpreted on the host on which the NetBackup Device Manager service runs. On UNIX and Linux computers, this entry is read and interpreted on the host on which `ltid` runs.

Device path remapping is enabled by default on Windows and Linux servers. It is disabled by default on all other servers.

AUTO_UPDATE_ROBOT entry in `vm.conf`

Use this entry to inject media automatically from the Media Access Port (MAP) into a TL8 or TLD robot and update the EMM database. Media are injected if the robot generates a unit attention message.

```
AUTO_UPDATE_ROBOT
```

This entry only operates with the TL8 or TLD robots that post a unit attention when their MAP is opened.

Symantec recommends that this entry not be used with partitioned libraries. Most robotic libraries with multiple partitions do not post a unit attention when the MAP is opened.

AVRD_PEND_DELAY entry in `vm.conf`

If this entry is specified in `vm.conf`, `avrd` waits *number_of_seconds* before it displays a pending status (PEND) in the Device Monitor. This entry is read and interpreted on the host on which `avrd` runs.

```
AVRD_PEND_DELAY = number_of_seconds
```

On some server operating systems (Windows and HP-UX), NetBackup reports PEND if the drive reports Busy when a volume is unmounted. Use this entry to minimize the display of this misleading status.

minimum for *number_of_seconds* is zero. The maximum is 255. The default value is 180 seconds.

AVRD_SCAN_DELAY entry in `vm.conf`

If this entry is specified in `vm.conf`, `avrd` waits *number_of_seconds* between normal scan cycles. This entry is read and interpreted on the host on which `avrd` runs.

```
AVRD_SCAN_DELAY = number_of_seconds
```

Use this entry to minimize tape mount times. Without this entry, NetBackup delays mount requests by an average of 7.5 seconds.

The minimum for *number_of_seconds* is 1. The maximum is 180. A value of zero converts to one second. The default value is 15 seconds. If a value is used that is greater than the default, NetBackup delays mount requests and drive status updates in the Device Monitor.

Note: If *number_of_seconds* is set to a value that allows media to be changed within one scan cycle, NetBackup may not detect media changes. Data loss may occur.

CLEAN_REQUEST_TIMEOUT entry in `vm.conf`

Use this entry to specify how long NetBackup waits for a drive to be cleaned before it removes the cleaning request from the cleaning queue. Unprocessed requests to clean a drive are removed from the queue after 30 minutes.

```
CLEAN_REQUEST_TIMEOUT = minutes
```

minutes can be from 1 to 144000 (100 days). The default value is 30 and a value of zero converts to the default value of 30.

CLIENT_PORT_WINDOW entry in `vm.conf`

Use this entry to specify the range of non-reserved ports on this host that are used to connect to `vmd` on other hosts. This entry is read and interpreted on the host on which `vmd` runs.

```
CLIENT_PORT_WINDOW = start end
```

For example, the following entry permits ports from 4800 through 5000:

```
CLIENT_PORT_WINDOW = 4800 5000
```

The operating system determines the non-reserved port to use in the following cases:

- A `CLIENT_PORT_WINDOW` entry is not specified.
- A value of zero is specified for `start`.

CLUSTER_NAME entry in `vm.conf`

```
CLUSTER_NAME = cluster_alias
```

This entry specifies the virtual name for the media server on which the `vm.conf` file resides.

See [“Host name precedence in the `vm.conf` file”](#) on page 90.

CONNECT_OPTIONS entry in `vm.conf`

This entry only affects connections to NetBackup 7.0 and earlier. For connections to NetBackup 7.0.1 and later, the `veritas_pbx` port is used.

Add this entry in `vm.conf` to specify the options that enhance firewall efficiency with NetBackup. Server connection options can be any of the following: use `vnetd` or the daemon's port number, use only `vnetd`, or use only the daemon's port number.

```
CONNECT_OPTIONS = server_name 0 0 [0|1|2]
```

`CONNECT_OPTIONS` entries can be specified for multiple servers.

`server_name` is the name of the media server to connect to.

The first and second options currently are not used. Specify zero for these options.

The third option specifies the connection method to use to connect to `server_name` as follows:

- A value of 0 specifies to use `vnetd` to connect to a daemon on the server. If the `vnetd` service is not active, connect by using the traditional port number of the daemon.
- A value of 1 specifies to use `vnetd` only to connect to a daemon on the server.
- A value of 2 specifies to use the traditional port number of the daemon to connect to the daemon on the server. The default value is 2.

The following example entry specifies to use either `vnetd` or the daemon's port number to connect to server shark:

```
CONNECT_OPTIONS = shark 0 0 0
```

The following example entry specifies to use `vnetd` only to connect to server dolphin:

```
CONNECT_OPTIONS = dolphin 0 0 1
```

The following example entry specifies to use the daemons's port number only to connect to server perch:

```
CONNECT_OPTIONS = perch 0 0 2
```

DAS_CLIENT entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
DAS_CLIENT = client_name
```

If this entry is specified in `vm.conf`, specify the DAS client name that the TLM robot uses for communications with the DAS/SDLC server. By default this client name is the host name of the media server. This entry is read and interpreted on the host where `tlmd` is running.

DAYS_TO_KEEP_LOGS entry in `vm.conf`

If this entry is specified in `vm.conf`, specify the number of days to keep debug logs before `vmd` deletes them. This entry is read and interpreted on the hosts where `vmd` is running.

```
DAYS_TO_KEEP_LOGS = days
```

A value of zero means that the logs are not deleted. The default is zero. This entry does not affect the debug logs that Unified Logging creates.

Information about Unified Logging is available.

See the *NetBackup Troubleshooting Guide for UNIX, Windows, and Linux*.

EMM_RETRY_COUNT entry in `vm.conf`

The `vmd` daemon and the `ltid` daemon use this entry to determine how many times to retry requests to the NetBackup Enterprise Media Manager.

```
EMM_RETRY_COUNT = number_of_retries
```

The default is one retry.

Only change the value of this `vm.conf` file entry when directed to do so by a NetBackup support representative. If this entry is added to the `vm.conf` file or if this value is changed, restart the `vmd` daemon and the `ltid` daemon.

EMM_CONNECT_TIMEOUT entry in `vm.conf`

This value applies for broken connections between the NetBackup Enterprise Media Manager and the following daemons: the `vmd` daemon and the `ltid` daemon. These two daemons use this entry to determine for how long they should try to reconnect to the NetBackup Enterprise Media Manager.

```
EMM_CONNECT_TIMEOUT = number_of_seconds
```

The default is 20 seconds.

Only change the value of this `vm.conf` file entry when directed to do so by a NetBackup support representative. If this entry is added to the `vm.conf` file or if this value is changed, restart the `vmd` daemon and the `ltid` daemon.

EMM_REQUEST_TIMEOUT entry in `vm.conf`

The `vmd` daemon and the `ltid` daemon use this entry to determine how many seconds to allow a request to the NetBackup Enterprise Media Manager to complete.

```
EMM_REQUEST_TIMEOUT = number_of_seconds
```

The default is 300 seconds.

Only change the value of this `vm.conf` file entry when directed to do so by a NetBackup support representative. If this entry is added to the `vm.conf` file or if this value is changed, restart the `vmd` daemon and the `ltid` daemon.

ENABLE_ROBOT_AUTH entry in `vm.conf`

Symantec encourages the use of Symantec Product Authentication and Authorization for NetBackup Access Control (NBAC) instead of legacy security implementations.

For information about the `ENABLE_ROBOT_AUTH` configuration entry, see the NetBackup 6.0 documentation. Information on Symantec Product Authentication and Authorization is available.

See the *NetBackup Security and Encryption Guide*.

INVENTORY_FILTER entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
INVENTORY_FILTER = robot_type robot_number mode value1 [value2 ...]
```

Used to filter robot inventory results in ACS or TLH robot types. Add this entry to the configuration file (`vm.conf`) on the NetBackup server on which the inventory operation is invoked. This entry is read and interpreted on the host on which `vmcheckxxx` and `vmupdate` run.

Note: This entry may be required for an ACS robot and the ACS library software host with an STK Library Station. Newer versions of STK Library Station allow robot inventory commands to function correctly so filters are not required.

robot_type can be ACS or TLH.

robot_number is the number of the robot as was configured in NetBackup.

mode is `BY_ACS_POOL` for ACS or `BY_CATEGORY` for TLH.

See the following examples:

```
INVENTORY_FILTER = ACS 0 BY_ACS_POOL 4 5
INVENTORY_FILTER = TLH 0 BY_CATEGORY FFFA CDB0
```

MAP_ID entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
MAP_ID = robot_num map_ID
```

Use this entry to configure the default media access port (MAP) to use to eject media from Automated Cartridge System (ACS) robots. This default is selected in

the NetBackup Administration Console, but you can also select other Media Access Ports for ejects.

If the MAP is not available or the `vm.conf` file does not contain this entry, NetBackup uses the default MAP selection process. By default, NetBackup uses the smallest MAP that can hold the number of media to be ejected.

If NetBackup selects multiple MAPs, NetBackup uses the nearest-MAP algorithm rather than the MAP that is specified in the MAP ID entry.

See “[ADJ_LSM entry in vm.conf](#)” on page 75.

robot_num is the robot number. *map_ID* is in the format of an ACS CAP (cartridge access port) ID and cannot contain any spaces.

The following example specifies the MAP ID for ACS robot number 700. The ACS CAP ID of 0,1,0 is used.

```
MAP_ID = 700 0,1,0
```

MAP_CONTINUE_TIMEOUT entry in `vm.conf`

This entry applies only when the `vmchange` command is used and the `-w` option is specified.

```
MAP_CONTINUE_TIMEOUT = seconds
```

The default timeout value for *seconds* is 300 (5 minutes). *seconds* cannot be zero and values greater than 1200 (20 minutes) can cause the robotic daemon to cancel the operation.

If this entry is specified in `vm.conf`, the SCSI robotic daemons wait the specified number of seconds before they time out. A timeout can occur while the daemons wait for user reply after the user removes volumes from the media access port. If a timeout occurs, NetBackup aborts the operation.

This entry is read and interpreted on the host on which the SCSI-controlled robotic daemon or process runs.

Note: Non-mount activities such as a robotic inventory cannot occur during this timeout period.

MEDIA_ID_BARCODE_CHARS entry in `vm.conf`

If this entry is specified in `vm.conf`, it controls NetBackup media ID generation. This entry is read and interpreted on the host on which `vmcheckxxx` and `vmupdate` run as part of the robot inventory operation.

```
MEDIA_ID_BARCODE_CHARS = robot_num barcode_length media_ID_rule
```

Note: To use this entry, the robot must support bar codes and the robot type cannot be an API robot.

Choose how NetBackup creates media IDs by defining the rules that specify which characters of a barcode on tape NetBackup uses. Alphanumeric characters can be specified to be inserted in the ID.

Multiple entries can be added to the `vm.conf` file. For example, specify media ID generation for each robot or for each barcode format that has different numbers of characters. The multiple entries allow flexibility for multimedia.

If no `MEDIA_ID_BARCODE_CHARS` entries exist or the entry is invalid, NetBackup uses the rightmost six characters of the barcode to create its media ID.

robot_num is the robot number.

barcode_length is the length of the barcode.

A *media_ID_rule* consists of a maximum of six fields that colons delimit. Numbers in the fields define the positions of the characters in the barcode that NetBackup extracts (from left to right). For example, if the number 2 is in a field, NetBackup extracts the second character from the barcode. The numbers can be specified in any order.

If the pound sign (#) prefixes a character, that character is inserted in that position in the generated ID. Any alphanumeric characters must be valid for a media ID. Use rules to create media IDs of many different formats. However, if the generated media ID is different from the label on the media, media management may be more difficult.

The following is an example rule and the resulting generated media ID:

```
Barcode on the tape: 032945L1
Media ID rule:      #N:2:3:4:5:6
Generated media ID: N32945
```

MEDIA_ID_PREFIX entry in `vm.conf`

If this entry is specified in `vm.conf`, it defines the media ID prefixes to use for media without bar codes. This entry is read and interpreted on the host where `vmcheckxxx` and `vmupdate` are running as part of the robot inventory operation.

```
MEDIA_ID_PREFIX = media_id_prefix
```

The best way to add media to a robot is to use the Robot Inventory Update Volume Configuration operation.

MM_SERVER_NAME entry in `vm.conf`

```
MM_SERVER_NAME = host_name
```

This entry specifies the name other NetBackup servers and clients should use when they refer to this server.

See [“Host name precedence in the `vm.conf` file”](#) on page 90.

PREFERRED_GROUP entry in `vm.conf`

Symantec encourages the use of Symantec Product Authentication and Authorization for NetBackup Access Control (NBAC) instead of legacy security implementations.

For information about the `PREFERRED_GROUP` configuration entry, see the NetBackup 6.0 documentation. Information on Symantec Product Authentication and Authorization is available.

See the *NetBackup Security and Encryption Guide*.

PREVENT_MEDIA_REMOVAL entry in `vm.conf`

This topic applies to the TL8 robots only.

Specifying this entry changes the default operation for TL8 robots. Without this entry present, NetBackup allows the removal of media.

If this entry is specified in `vm.conf`, TL8 robots run the SCSI command `PREVENT MEDIUM REMOVAL`. The robot's main door or the MAP cannot be opened while the robotic control daemon runs.

This entry is read and interpreted on the host on which the TL8 robot control daemon or process (`tl8cd`) runs.

To override `PREVENT_MEDIA_REMOVAL`, do one of the following:

- Use the test utility and run `allow media removal`.
- Use `inject` or `eject` for access, when volumes are added or moved.

RANDOM_PORTS entry in `vm.conf`

Use this entry to specify whether NetBackup chooses port numbers randomly or sequentially for communication with other NetBackup servers. This entry is read and interpreted on hosts on which `vmd` runs.

```
RANDOM_PORTS = YES|NO
```

If `YES` or no entry exists (the default), NetBackup chooses port numbers randomly from those that are available in the allowed range.

If `NO`, NetBackup chooses numbers sequentially. NetBackup begins with the highest number in the allowed range, and then tries the next highest, and so on until a port is available.

To specify no random ports in the NetBackup configuration file, do one of the following:

- Specify `RANDOM_PORTS = NO` in the `bp.conf` file on UNIX.
- Use the NetBackup **Host Properties** on Windows.

REQUIRED_INTERFACE entry in `vm.conf`

```
REQUIRED_INTERFACE = host_name
```

This entry specifies the name of the network interface that the media server uses to connect to another media server.

A NetBackup server can have more than one network interface, and by default the operating system determines the one to use. To force NetBackup to connect through a specific network interface, use `REQUIRED_INTERFACE` and specify the name of that network interface.

See “[Host name precedence in the `vm.conf` file](#)” on page 90.

SERVER entry in `vm.conf`

This entry determines the name other NetBackup servers should use when they refer to this server.

`SERVER` entries in the `vm.conf` file are used for NetBackup media server security.

```
SERVER = host_name
```

`SERVER` entries work with the `AUTHORIZATION_REQUIRED` entry to control which hosts can monitor and control devices on this host.

If the `AUTHORIZATION_REQUIRED` entry exists, the `vm.conf` file must include a `SERVER` entry for every media server that controls devices on this host. If the `vm.conf` file contains any `SERVER` entries, it also must include a `SERVER` entry for itself or it cannot manage its own devices.

If no `AUTHORIZATION_REQUIRED` entry exists and no `SERVER` entries exist, any NetBackup server can monitor and control devices on this host.

For security, the entries that allow only specific hosts to access the devices must be added remotely.

This entry is read and interpreted on media servers on which the NetBackup `vmd` service runs.

SSO_DA_REREGISTER_INTERVAL entry in `vm.conf`

This entry determines the name other NetBackup servers should use when they refer to this server.

This configuration entry applies only to NetBackup Enterprise Server.

```
SSO_DA_REREGISTER_INTERVAL = minutes
```

This `vm.conf` entry is for the Shared Storage Option (SSO) for Tape feature only. It is read and interpreted on the host on which `ltid` runs.

`ltid` on a scan host periodically registers its shared drives with `EMM/DA` to ensure that it is still provides the drive scanning function. Only one of the hosts that share a drive scan the drive. This reregistration allows conditions such as a device allocator restart to have minimal effect on use of shared drives.

The default for the reregistration interval is 5 minutes. Use the `SSO_DA_REREGISTER_INTERVAL` entry to tune this interval. After the entry is added, stop and restart `ltid` for the change to take effect.

SSO_DA_RETRY_TIMEOUT entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
SSO_DA_RETRY_TIMEOUT = minutes
```

This `vm.conf` entry is for the Shared Storage Option (SSO) for Tape feature only. It is read and interpreted on the host on which `ltid` runs.

The Device Manager `ltid` delays before if one of the following events occurs:

- Problems during communications with `EMM/DA`.
- Failure trying to reserve a shared drive.

The default value for the delay is 3 minutes. Use the `SSO_DA_RETRY_TIMEOUT` entry to tune this delay period. After the entry is added, stop and restart `ltid` for the change to take effect.

SSO_HOST_NAME entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
SSO_HOST_NAME = host_name
```

This `vm.conf` entry is for the Shared Storage Option (SSO) for Tape feature only. It is read and interpreted on the host on which `ltid` runs.

This entry specifies the name that the current host uses to register, reserve, and release shared drives with EMM/DA. The default is the local host name.

TLH_mediatype entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
TLH_mediatype = Media_Manager_mediatype
```

If this entry is specified in `vm.conf`, IBM ATL media types in tape library Half-inch (TLH) robots are mapped to Media Manager media types. This entry is read and interpreted on the host where `vmcheckxxx` and `vmupdate` are running as part of the robot inventory operation.

TLM_mediatype entry in `vm.conf`

This configuration entry applies only to NetBackup Enterprise Server.

```
TLM_mediatype = Media_Manager_mediatype
```

If this entry is specified in `vm.conf`, DAS/SDLC media types in tape library Multimedia (TLM) robots are mapped to Media Manager media types. This entry is read and interpreted on the host where `vmcheckxxx` and `vmupdate` are running as part of the robot inventory operation.

VERBOSE entry in `vm.conf`

If this entry is specified in `vm.conf`, all Media Manager components on the host are started with verbose logging enabled.

Use this option only if problems occur or if requested by Symantec support. After the problem is resolved, remove the debug logs or add a `DAYS_TO_KEEP_LOGS` entry.

Example `vm.conf` file

The following is an example of a `vm.conf` file, on host `server1`:

```
SERVER = server1
SERVER = server2
MEDIA_ID_PREFIX = NV
MEDIA_ID_PREFIX = NETB
ACS_3490E = HCART2
```

Host name precedence in the `vm.conf` file

NetBackup identifies the media server by using the following name precedence:

- `CLUSTER_NAME` entry if present in `vm.conf`.
- `MM_SERVER_NAME` entry if present in `vm.conf`.
- `REQUIRED_INTERFACE` entry if present in `vm.conf`.
- The same name that NetBackup uses.
- `gethostname()` name.

Reference topics

This chapter includes the following topics:

- [Host name rules](#)
- [About reading backup images with tar](#)
- [Factors that affect backup time](#)
- [Methods for determining the NetBackup transfer rate](#)
- [NetBackup notify scripts](#)
- [Media and device management best practices](#)
- [About TapeAlert](#)
- [About tape drive cleaning](#)
- [How NetBackup selects drives](#)
- [How NetBackup reserves drives](#)
- [How NetBackup selects media](#)
- [Volume pool and volume group examples](#)
- [Media formats](#)
- [Media Manager commands](#)

Host name rules

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.

See [“About dynamic host name and IP addressing”](#) on page 41.

NetBackup uses TCP/IP host names to connect to NetBackup servers and clients. NetBackup validates its connections by performing a reverse host name lookup. That is, NetBackup determines the IP address of a connection and then uses the IP address to look up the host name with `gethostbyaddr()`. The host name and address resolution must be set up correctly in DNS, WINS, or the local `%Systemroot%\system32\drivers\etc\hosts` file (if necessary).

Note: Place the system host name and IP address in the `%Systemroot%\system32\drivers\etc\hosts` file to accelerate name lookups.

See [“Using the System Monitor with NetBackup”](#) on page 102.

How NetBackup uses host names

A major consideration is the extent to which you qualify host names. In many cases, the short host name of a computer is adequate. If the network environment contains multiple domains, 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 only `mercury`.

The following topics discuss how NetBackup stores and uses host names. These topics also address factors to consider when you choose host names.

Note: Do not change the host name of a NetBackup server. This practice is not recommended. You may need to import all previously used media to the server before you can use it under the new host name.

The following table discusses topics that address how NetBackup stores and uses host names.

Table 4-1 How NetBackup stores and uses host names

Topic	Description
Policy configuration	<p>The configured name for a client is the host name as it is added to a policy. This name is how the client is identified in the NetBackup configuration.</p> <p>The server uses the client's configured name to connect to the client and start the processes that satisfy client requests. Always use qualified host names to add clients to a policy so that all NetBackup servers can connect to the clients.</p> <p>When a client makes a user backup, archive, or restore request to the NetBackup server, the server uses the peer name of the client. The peer name (identified from its TCP connection) is used to determine the client's configured name.</p> <p>If you add a client to more than one policy, always use the same name in all cases. If the same name is not used, the client cannot view all the files that are backed up on its behalf. In this case, file restores become complicated because both user action and administrator action is required to restore from some of the backups.</p>
Image catalog	<p>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.</p> <p>Every backup for a client has a separate file in this subdirectory. Each of these backup records contains the host name of the server on which the backup was written.</p>
Error catalog	<p>NetBackup uses entries in the error catalog for generating reports. These entries contain the host name of the server that generates the entry and the client's configured name, if applicable. The server host name is normally the server's short host name. (For example, <i>servername</i> instead of <i>servername.null.com</i>.)</p>
Catalog backup information	<p>This topic applies to NetBackup Enterprise Server.</p> <p>If you include a media server's catalog files in the NetBackup catalog, qualify the host name of the media server in the file path. Qualified names are necessary because they allow the master server to connect to the media server.</p>

Updating NetBackup after changing the host name

Do not change the host name of a NetBackup server. A name change might require that all previously used media be imported to the server before the host can be used under the new name.

Use the following steps to update the NetBackup configuration if a client's host name is changed.

To update NetBackup after a master server name change

See ["To update NetBackup after a master server name change"](#) on page 94.

To update NetBackup after a client name change See [“To update NetBackup after a client name change”](#) on page 94.

To update NetBackup after a master server name change

- 1 On the master server, delete the client’s old name from all policies where it exists and add the client’s new name to those policies. You do not need to reinstall NetBackup software on the client. The client continues to have access to all previous backups.
- 2 Create a file named `ALTPATH` in the image catalog directory.

For example, if the client name is `client1`, the `ALTPATH` file is created in the following location:

```
Install_path\VERITAS\NetBackup\db\images\client1\  
ALTPATH
```

- 3 Create a directory for the new `client2` in the `\images` directory:

```
Install_path\VERITAS\NetBackup\db\images\client2
```

- 4 On the first line of the `client1\ALTPATH` file, specify the path to the directory for the new client. The path is the only entry in the `ALTPATH` file.

```
Install_path\VERITAS\NetBackup\db\images\client2
```

To update NetBackup after a client name change

- 1 On PC clients, change the client name setting either through the user interface or in a configuration file.

See the online Help in the **Backup, Archive, and Restore** client interface.

- 2 On UNIX clients, change the `CLIENT_NAME` value in the `bp.conf` file to the new name.

If users on UNIX clients have a `bp.conf` file in the `$HOME` directory, users must change `CLIENT_NAME` in that file to the new name.

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` library function. If the name is unknown to the master server Domain Name Service, the master server may not be able to reply to client requests.

This possible situation depends on how the client and the server are configured. If `gethostname` on the client returns host names that DNS on the master server cannot resolve, problems occur.

One possible solution is to reconfigure the client or the master server DNS hosts file. Another option is to create a special file in the `altnames` directory on the master server. The file forces the translation of NetBackup client host names.

```
install_path\NetBackup\db\altnames\host.xlate
```

Each line in the `host.xlate` file contains three elements: a numeric key and two host names. Each line is left-justified, and a space character separates each element of the line:

```
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, which indicates a configured name translation.
- *hostname_from_client* is the value to translate. The client name must correspond to the name that is obtained by running the client's `gethostname`. The value must be sent to the server in the request.
- *client_as_known_by_server* is the name to substitute for *hostname_from_client* for request responses. The name must match the name in the NetBackup configuration on the master server and must also be known to the master server's network services.

Consider the following example:

```
0 xxxx xxxx.eng.aaa.com
```

The line specifies that when the master server receives a request for a configured client name (numeric key 0), the name `xxxx.eng.aaa.com` always replaces `xxxx`.

The substitution resolves the problem if the following conditions are true:

- When `gethostname` is run on the client, it returns `xxxx`.
- The master server's network services `gethostbyname` library function did not recognize the name `xxxx`.
- The client was configured and named in the NetBackup configuration as `xxxx.eng.aaa.com`. And, this name is also known to network services on the master server.

About reading backup images with tar

NetBackup for UNIX uses a modified GNU `tar` for reading backup images. By using the modified `tar`, NetBackup can understand compressed files, sparse files, long pathnames, and ACL information. It offers features similar to those in `cpio`.

Although non-NetBackup versions of `tar` can be used to restore files, they provide only limited restore capabilities.

Note: You cannot use the NetBackup modified-GNU `tar` on UNIX or `tar32.exe` on Windows to extract files from a NetBackup for Windows backup image.

Consequences of using a non-NetBackup tar

Non-NetBackup versions of `tar` do not supply all of the restore capabilities that the NetBackup `/usr/openv/netbackup/bin/tar` provides. Possible problems result.

The following is a list of consequences that can occur if using a non-NetBackup `tar`:

- Compressed backups cannot be recovered.
- Multiplexed backups cannot be recovered.
- Solaris extended attributes cannot be restored to a client.
- VxFS named data streams cannot be restored to a client.
- Backups cannot be recovered that contain raw partitions. (Includes FlashBackup images.)
- NDMP client backup images cannot be restored, though NDMP vendors may have tools or the utilities that can perform a restore directly from the media.
- Non-NetBackup versions of `tar` may have trouble with sparse files and often skip sparse files.
- HP CDFs are restored with non-NetBackup versions of `tar`. The directory is no longer hidden and the name of the directory has a `+` appended to it.
- If the backup spans more than one piece of media, you must read and combine the fragments from the media to give to `tar`. To combine the fragments, the system's `dd` command may be useful.

Another possibility is to use `tar` on the fragments. To use `tar` on fragments can allow recovery of 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` combine the `atime`, `mtime`, and `ctime` strings with the file name and create the file paths that are not desirable.

About the files that tar generates

Any version of `tar` (including NetBackup-modified `tar`) can generate a number of files depending on the circumstances of the recovery, as the following table shows.

Table 4-2 Files that tar generates

File	Description
<code>@@MaNgLeD.nnnn</code>	For backups containing pathnames longer than 100 characters, <code>tar</code> generates the files that are named <code>@@MaNgLeD.nnnn</code> that contain the actual file.
<code>@@MaNgLeD.nnnn_Rename</code>	<code>tar</code> generates another file (<code>@@MaNgLeD.nnnn_Rename</code>) that explains how to rename the <code>@@MaNgLeD.nnnn</code> files to return the files to the correct location.
<code>@@MaNgLeD.nnnn_Symlink</code>	For long names of symbolic links, <code>tar</code> generates the files that are named <code>@@MaNgLeD.nnnn_Symlink</code> . These files contain descriptions of the symbolic links that must be made to return a link to the correct file.
For cross-platform ACLs restores, <code>tar</code> creates and stores the ACLs in <code>.SeCuRiT.y.nnnn</code> files in the <code>root</code> directory	The files can either be read or deleted. Regenerate the ACLs to the corresponding files by hand.
For cross-platform VxFS extent attribute restores, <code>tar</code> creates and stores extent attributes in <code>.ExTeNt.nnnn</code> files in the <code>root</code> directory	The files can either be deleted or read and the extent attributes regenerated by hand to the corresponding files.

Factors that affect backup time

The amount of time that NetBackup requires to complete a backup is an important factor in setting up schedules. The importance of time is particularly true for the sites that handle large amounts of data. For example, the total backup time can exceed the time that is allotted to complete backups and interfere with normal network operations. Longer backup times also increase the possibility of a problem

that disrupts the backup. The time to back up files can also give an indication of how long it may take to recover the files.

Figure 4-1 shows the major factors that affect backup time.

Figure 4-1 Backup time formula

$$\text{Backup time} = \frac{\text{Total data}}{\text{Transfer rate}} + \frac{\text{Compression factor (optional)}}{1} \times \text{Device delays}$$

Total amount of data to back up

The total amount of data to back up depends on the size of the files for each client in the policy. The total amount of data also depends on whether the backup is a full backup or an incremental backup.

The implications are as follows:

- Full backups involve all the data. Therefore, a full backup usually takes longer than an incremental backup.
- Differential incremental backups include only the data that changed since the last full or incremental backup.
- Cumulative incremental backups include all the data that changed since the last full backup.

For incremental backups, the amount of data depends on the frequency with which files change. If a large number of files change frequently, incremental backups are larger.

Transfer rate

The transfer rate depends on the following factors.

Table 4-3 Transfer rate factors

Factor	Description
Speed of the backup device	Backups that are sent to tapes with a transfer rate of 800 kilobytes per second are generally faster than tapes with a transfer rate of 400 kilobytes. (Assume that other factors allow for the faster transfer rate.)
Available network bandwidth	The available bandwidth is less than the theoretical network bandwidth and depends on how much other network traffic is present. For example, multiple backups occurring on the same network compete for bandwidth.

Table 4-3 Transfer rate factors (*continued*)

Factor	Description
Speed with which the client can process the data	The speed varies with the hardware platform and depends on the other applications that run on the platform. File size is also an important factor. Clients can process larger files faster than smaller ones. A backup for 20 files, 1 megabyte each, is faster than a backup for 20,000 files that are 1 kilobyte each.
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 that run on the platform. The number of concurrent backups being performed also affects server speed.
Network configuration can affect performance	For example, when some machines run full-duplex and some run half-duplex in an Ethernet environment, the throughput is significantly reduced.
Device delays	<p>Device delays can be due to the following factors:</p> <ul style="list-style-type: none"> ■ The device may be busy or slow to load the media. ■ The device may be slow to find the location on the media at which to start writing the backup. <p>These delays can vary widely and depend on the devices and the computing environments.</p>

Methods for determining the NetBackup transfer rate

Calculate three variations of the backup transfer rate by using NetBackup report data.

Three NetBackup transfer rates and calculation methods are available.

Table 4-4 NetBackup transfer rates

Transfer rate	Description
Network transfer rate	<p>The network transfer rate is the rate provided in the All Log Entries report.</p> <p>The network transfer rate considers only the time it takes to transfer data over the network from client to server.</p> <p>This rate ignores the following:</p> <ul style="list-style-type: none"> ■ The time the device requires to load and position media before a backup. ■ The time that the tape file requires to close and write an additional NetBackup information record to the tape.
Network transfer plus end-of-backup processing rate	<p>This rate ignores the time it takes to load and position media before a backup. However, the rate does include 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:</p> <pre data-bbox="606 826 955 847">begin writing backup id xxx</pre> <p>until the message</p> <pre data-bbox="606 927 1018 947">successfully wrote backup id xxx</pre> <p>To calculate the transfer rate, divide this time (in seconds) into the total bytes that are transferred. (The total bytes that are transferred is recorded in the All Log Entries report.)</p>
Total transfer rate	<p>This transfer rate includes the time it takes to load and position the media as well as the end-of-backup processing. Use the List Client Backups report to calculate the transfer rate by dividing Kilobytes by Elapsed Time (converted to seconds).</p>

The Microsoft Windows System Monitor also displays the NetBackup transfer rate.

See [“Using the System Monitor with NetBackup”](#) on page 102.

Examples of reports that provide backup data to calculate transfer rates

Assume that the reports provide the following data.

Sample **All Log Entries** report:

```

TIME                SERVER/CLIENT  TEXT
04/28/09 23:10:37  windows giskard begin writing backup
                   id giskard_0767592458, fragment 1 to
                   media id TL8033 on device 1 . . .
04/29/09 00:35:07  windows giskard successfully wrote
                   backup id giskard_0767592458,
                   fragment 1, 1161824 Kbytes at
                   230.325 Kbytes/sec

```

Sample List Client Backups Report:

```

Client:                giskard
Backup ID:             giskard_0767592458
Policy:               production_servers
Client Type:          Standard
Sched Label:          testing_add_files
Schedule Type:        Full
Backup Retention Level:  one week (0)
Backup Time:           04/28/09 23:07:38
Elapsed Time:         001:27:32
Expiration Time:      05/05/09 23:07:38
Compressed:           no
Kilobytes:            1161824
Number of Files:      78210

```

The following three rates were compiled with the backup data from the sample reports:

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

See [“Using the System Monitor with NetBackup”](#) on page 102.

Using the System Monitor with NetBackup

NetBackup adds the NetBackup Disk/Tape performance object to the list of objects that the Windows System Monitor monitors.

Four counters are available for the NetBackup Disk/Tape performance object, as follows:

- Disk/Tape Read Bytes (GB)
- Disk/Tape Read Bytes/sec (KB)
- Disk/Tape Write Bytes (GB)
- Disk/Tape Write Bytes/sec (KB)

The NetBackup performance object supports instances in the System Monitor. The instances can be drive names or absolute paths to which NetBackup writes, or from which NetBackup is reads.

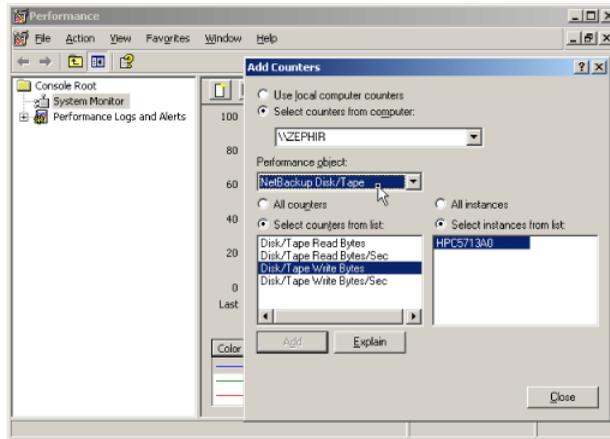
The System Monitor displays object instances when NetBackup reads or writes from the disk or tape. The read or write counters are updated depending on the type of NetBackup operation performed. The object instance is removed from the list once the NetBackup operation completes.

If the performance is monitored locally or remotely during a NetBackup read or write operation, the object instance exists after the NetBackup operation is complete. In this case, the object instance is removed when performance monitoring stops.

To monitor NetBackup counters remotely, the initiating computer attaches to the target computer's WinLogon process through RPC. To attach to the process locks the object instances. Thus, the object instances remain until the system reboots.

To use the System Monitor with NetBackup

- 1 Open the System Monitor on the Windows system.
- 2 In the **Performance** window, in the left pane, click on **System Monitor**.
- 3 To add a counter, click the plus sign (+) on the toolbar of the right pane. When the **Add Counters** window opens, select **NetBackup Disk/Tape** from the **Performance object** drop-down list.



In order for the NetBackup objects to be available for selection, the following conditions must be met:

- The drive must be connected to a Windows media server (or SAN media server). A NetBackup job must be active (a drive is in use).
 - The user must have permissions to read the Windows registry.
 - Performance data collection is enabled (select **Host Properties > Media Servers > Universal Settings > Enable performance data collection**).
- 4 Select the counter to display from the list of available counters. Available counters are as follows:
 - Disk/Tape Read Bytes (GB)
 - Disk/Tape Read Bytes/sec (KB)
 - Disk/Tape Write Bytes (GB)
 - Disk/Tape Write Bytes/sec (KB)
 - 5 Select one or more object instances from the list of instances. Instances appear when NetBackup begins to read or write from the disk or the tape drives.
 - 6 Click **Add**. The number of bytes that are read or written is updated dynamically, along with the rate.

NetBackup notify scripts

NetBackup provides scripts or batch files that can collect information and be used to notify administrators of specific events.

The `Install_path\VERITAS\NetBackup\bin\goodies\` directory contains sample shell scripts to modify. The scripts in the `\goodies` directory are not supported but are intended as examples to customize.

The following scripts are active on the master server:

`Install_path\VERITAS\NetBackup\bin\backup_notify.cmd`

See [“backup_notify.cmd on Windows”](#) on page 105.

`Install_path\VERITAS\NetBackup\bin\backup_exit_notify.cmd`

See [“backup_exit_notify.cmd on Windows”](#) on page 105.

`Install_path\VERITAS\NetBackup\bin\diskfull_notify.cmd`

See [“diskfull_notify.cmd on Windows”](#) on page 116.

`Install_path\VERITAS\NetBackup\bin\mail_dr_info.cmd`

See [“mail_dr_info.cmd”](#) on page 116.

`Install_path\VERITAS\NetBackup\bin\goodies\media_deassign_notify`

See [“media_deassign_notify”](#) on page 117.

`Install_path\VERITAS\NetBackup\bin\nbmail.cmd`

See [“nbmail.cmd”](#) on page 117.

`Install_path\VERITAS\NetBackup\bin\goodies\pending_request_notify`

See [“pending_request_notify”](#) on page 120.

`Install_path\VERITAS\NetBackup\bin\restore_notify.cmd`

See [“restore_notify.cmd on Windows”](#) on page 120.

`Install_path\VERITAS\NetBackup\bin\session_notify.cmd`

See [“session_notify.cmd on Windows”](#) on page 120.

`Install_path\VERITAS\NetBackup\bin\session_start_notify.cmd`

See [“session_start_notify.cmd on Windows”](#) on page 120.

`Install_path\VERITAS\NetBackup\bin\userreq_notify.cmd`

See [“userreq_notify.cmd on Windows”](#) on page 121.

`Install_path\VERITAS\NetBackup\bin\goodies\parent_end_notify.cmd`

See [“parent_end_notify.cmd on Windows”](#) on page 118.

`Install_path\VERITAS\NetBackup\bin\goodies\parent_start_notify.cmd`

See [“parent_start_notify.cmd on Windows”](#) on page 119.

`Install_path\VERITAS\Volmgr\bin\shared_drive_notify.cmd`

See “[shared_drive_notify.cmd on Windows](#)” on page 121.

The following scripts are run on clients:

`Install_path\VERITAS\NetBackup\bin\goodies\bpstart_notify.bat`

See “[bpstart_notify.bat \(Microsoft Windows clients only\)](#)” on page 109.

`Install_path\VERITAS\NetBackup\bin\goodies\bpnd_notify.bat`

See “[bpnd_notify.bat \(Microsoft Windows clients only\)](#)” on page 113.

To use the client scripts, first create the script on the client.

See “[bpstart_notify.bat \(Microsoft Windows clients only\)](#)” on page 109.

See “[bpnd_notify.bat \(Microsoft Windows clients only\)](#)” on page 113.

For more information, see the comments in the scripts.

Note: This note applies only to the NetBackup Enterprise Server. If you use either the `bpstart_notify` or `bpnd_notify` scripts, do not include any commands that write to `stdout`. NetBackup sends the output that is written to `stdout` to the server as part of the backup. The resulting backup can abort with an error message that pertains 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. Its use can cause data to be written to `stdout` or `stderr`.

backup_notify.cmd on Windows

The `backup_notify.cmd` script runs on the NetBackup server where the storage unit is located. It is called each time a backup is successfully written to media.

NetBackup passes the following parameters to this script:

- The name of the program doing the backup
- The backup-image name or path

For example:

```
backup_notify.cmd bptm host_0695316589
```

backup_exit_notify.cmd on Windows

The `backup_exit_notify.cmd` script runs on the master server. It is called to perform site-specific processing when an individual backup completes.

NetBackup passes the following parameters to the script:

<code>clientname</code>	Specifies the name of the client from the NetBackup catalog.
<code>policyname</code>	Specifies the policy name from the NetBackup catalog.
<code>schedname</code>	Specifies the schedule name from the NetBackup catalog.
<code>schedtype</code>	Specifies one of the following: <code>FULL</code> , <code>INCR</code> (differential incremental), <code>CINC</code> (cumulative incremental), <code>UBAK</code> , <code>UARC</code>
<code>exitstatus</code>	Specifies the exit code for the entire backup job.
<code>stream</code>	Specifies the backup stream number for a job. 0 = The backup job is not running multiple data streams. -1 = The job is a parent job.
<code>done_trying</code>	Specifies whether the job will retry. 0 = The job is not complete and will retry. 1 = The job is complete and will not retry. If the system is configured to make 3 attempts in 12 hours, the job could run this script up to 3 times. On the final attempt, the <code>done_trying</code> flag is set to 1. The job has either completed successfully or has failed and exhausted the number of tries.

For example:

```
backup_exit_notify.cmd clientname1 pol_prod sched_fulls FULL 0 -1 1
backup_exit_notify.cmd clientname2 pol_prod sched_incr INCR 73 0 1
```

bpstart_notify (UNIX clients only)

On UNIX clients, NetBackup calls the `bpstart_notify` script each time the client starts a backup or an archive.

Note: Ensure that this script can be run by others on the client before it is used. To do so, run `chmod 755 script_name`, where `script_name` is the name of the script.

To use this script, copy the following file from the server:

```
Install_path\VERITAS\NetBackup\bin\goodies\bpstart_notify.bat
```

Then place the script in the following location on the UNIX client:

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

Modify the script and ensure that you have permission to run the script.

The `bpstart_notify` script runs each time a backup or an archive starts and initialization is completed. The script runs before the tape is positioned. 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 runs in the foreground and the `bbpkar` process on the client waits for it to complete before continuing. Any commands in the script that do not end with an ampersand character (`&`) run serially.

The server expects the client to respond with a `continue` message within the time that the `BPSTART_TIMEOUT` option specifies 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:

<code>clientname</code>	Specifies the name of the client from the NetBackup catalog.
<code>policyname</code>	Specifies the policy name from the NetBackup catalog.
<code>schedname</code>	Specifies the schedule name from the NetBackup catalog.
<code>schedtype</code>	Specifies one of the following: <code>FULL</code> , <code>INCR</code> (differential incremental), <code>CINC</code> (cumulative incremental), <code>UBAK</code> , <code>UARC</code>

Note: The `bpstart_notify` script also runs for NetBackup catalog backups if a `.policyname[.schedule]` is not specified.

For example:

```
bpstart_notify client1 pol_cd4000s sched_fulls FULL
bpstart_notify client2 pol_cd4000s sched_incrementals INCR
bpstart_notify client3 pol_cd4000s sched_fulls FULL
bpstart_notify client4 pol_cd4000s sched_user_backups UBAK
bpstart_notify client5 pol_cd4000s sched_user_archive UARC
```

To create a `bpstart_notify` script for a specific policy or policy and schedule combination, create script files with a `[.policyname]` or `.policyname.schedulename`

suffix. The following are two examples of script names for a policy (production) that has a schedule (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 policy that are named production. The second script affects scheduled backups in the policy that is 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 script 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 `bpbkarr` process creates these variables. The following are examples of the strings that are available to the script to use to record information about a backup:

```
BACKUPID=client1_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 2009
```

In addition, the following environment variables can be used to support multiple data streams.

Table 4-5 Environment variables used to support multiple data streams

Environment variable	Description
STREAM_NUMBER	Specifies the stream number. The first stream from a policy, client, and schedule is 1. A 0 value indicates that multiple data streams are not enabled.
STREAM_COUNT	Specifies the total number of streams to be generated from this policy, client, and schedule.
STREAM_PID	Specifies the pid (process ID) number of <code>bpbkarr</code> .
RESTARTED	Specifies checkpointed restarts or checkpointed backup jobs. A value of 0 indicates that the job was not resumed. (For example, upon first initiation.) A value of 1 indicates that the job was resumed.

bpstart_notify.bat (Microsoft Windows clients only)

For all Windows clients, you can create batch scripts that provide notification whenever the client starts a backup or archive.

To use this script, copy the following file from the server:

```
Install_path\VERITAS\NetBackup\bin\goodies\bpstart_notify.bat
```

Then place the file on the client 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 for backups of a specific policy or schedule.

To create a script that applies to all backups, name the script `bpstart_notify.bat`.

To create a `bpstart_notify` script that applies only to a specific policy or policy and schedule combination, add a `.policyname` or `.policyname.schedule` suffix to the script name.

The following are examples of `bpstart_notify` script names:

- The following script applies only to a policy named `days`:

```
install_path\netbackup\bin\bpstart_notify.days.bat
```

- The following script applies only to a schedule that is named `fulls` in a policy named `days`:

```
install_path\netbackup\bin\bpstart_notify.days.fulls.bat
```

The `bpstart_notify` script also runs for NetBackup catalog backups if a `.policyname[.schedule]` is not specified.

The first script affects all scheduled backups in the policy named `days`. The second script affects scheduled backups in the policy 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.policy.schedule.bat  
bpstart_notify.policy.bat  
bpstart_notify.bat
```

For example, if there are both `bpstart_notify.policy.bat` and `bpstart_notify.policy.schedule.bat` scripts, NetBackup uses only the `bpstart_notify.policy.schedule.bat` script.

Note: `bpstart_notify` scripts can provide a different level of notification than the `bpstart_notify` scripts. For example, to use one of each, the script names might be `bpstart_notify.policy.bat` and `bpstart_notify.policy.schedule.bat`.

NetBackup passes the following parameters to the script:

- %1 Specifies the name of the client from the NetBackup catalog.
- %2 Specifies the policy name from the NetBackup catalog.
- %3 Specifies the schedule name from the NetBackup catalog.
- %4 Specifies one of the following: `FULL`, `INCR`, `CINC`, `UBAK`, `UARC`
- %5 Specifies that the status of the operation is always 0 for `bpstart_notify`.
- %6 Specifies the 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 policy and schedule, the results file must be named

install_path\netbackup\bin\BPSTART_RES.policy.schedule

If the script applies to a specific policy, the results file must be named

install_path\netbackup\bin\BPSTART_RES.policy

If the script applies to all backups, the results file must be named

install_path\netbackup\bin\BPSTART_RES

An `echo 0> %6` statement is one way for the script to create the file.

NetBackup deletes the existing results file before it calls the script. After the script runs, 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 time that the NetBackup `BPSTART_TIMEOUT` option specifies. 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 clients, `bpstart_notify` script can use the following environment variables for the support of multiple data streams.

Table 4-6 Environment variables that support multiple data streams

Environment variable	Description
<code>STREAM_NUMBER</code>	Specifies the stream number. The first stream from a policy, client, and schedule is 1. A 0 value indicates that multiple data streams are not enabled.
<code>STREAM_COUNT</code>	Specifies the total number of streams to be generated from this policy, client, and schedule.
<code>STREAM_PID</code>	Specifies the pid (process ID) number of <code>bpbkar</code> .

`bpend_notify` (UNIX clients only)

To receive a notification whenever a UNIX client completes a backup or an archive operation, copy the following file from the server:

```
Install_path\VERITAS\NetBackup\bin\goodies\bpend_notify
```

Then place the file in the following location on the UNIX client:

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

Modify the script and ensure that you have permission to run the script.

Note: The `bpend_notify` script is run when the client is finished sending data, but the server has not yet completed writing to media.

Note: Ensure that this script can be run by others on the client before it is used. To do so, run `chmod 755 script_name`, where *script_name* is the name of the script.

The `bpend_notify` script runs each time a backup or archive completes. For archives, it runs after the backup but before the files are removed.

If `bpend_notify` exists, it runs in the foreground and `bpbkar` on the client waits until it completes. Any commands that do not end with an ampersand character (&) run serially.

The server expects the client to respond within the time that the `BPEND_TIMEOUT` NetBackup configuration option specifies. 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 because it can delay the server from servicing other clients.

NetBackup passes the following parameters to the script:

<code>clientname</code>	Specifies the name of the client from the NetBackup catalog.
<code>policyname</code>	Specifies the policy name from the NetBackup catalog.
<code>schedname</code>	Specifies the schedule name from the NetBackup catalog.
<code>schedtype</code>	Specifies one of the following: <code>FULL</code> , <code>INCR</code> (differential incremental), <code>CINC</code> (cumulative incremental), <code>UBAK</code> , <code>UARC</code>
<code>exitstatus</code>	Specifies the exit code from <code>bpbkar</code> . The status is the client status and does not indicate that the backup is complete and successful. The client can display a status 0 when, due to a failure on the server, the All Log Entries report displays a status 84.

Note: The `bpend_notify` script also runs for NetBackup catalog backups if a `.policyname[.schedule]` is not specified.

For example:

```
bpend_notify client1 pol_1 fulls FULL 0
bpend_notify client2 pol_1 incrementals INCR 73
```

To create a `bpend_notify` script for a specific policy or policy and schedule combination, create script files with a `.policyname` or `.policyname.schedule` suffix. The following are two examples of script names for a policy that is named `production` with a schedule that is 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 policy `production`. The second script affects scheduled backups in the policy `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`.

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 the strings that are available to the script for use to record information about a backup:

```
BACKUPID=client1_0857340526
UNIXBACKUPTIME=0857340526
BACKUPTIME=Sun Mar 2 16:08:46 2011
```

The following environment variables can be used for the support of multiple data streams.

Table 4-7 Environment variables used for support of multiple data streams

Environment variable	Description
STREAM_NUMBER	Specifies the stream number. The first stream from a policy, client, and schedule is 1. A 0 value indicates that multiple data streams are not enabled.
STREAM_COUNT	Specifies the total number of streams to be generated from this policy, client, and schedule.
STREAM_PID	Specifies the pid (process ID) number of <code>bpbkar</code> .
FINISHED	Specifies the status of the checkpointed restarts of backup jobs. A value of 0 indicates that the client was not finished sending all of the data. A value of 1 indicates that the client was finished sending all the of data.

`bpend_notify.bat` (Microsoft Windows clients only)

For Windows 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\bpend_notify.bat
```

Install_path is the directory where NetBackup is installed.

You can create `bpend_notify` scripts that provide notification for all backups or for backups of a specific policy or schedule.

To create a `bpend_notify` script that applies to all backups, name the script `bpend_notify.bat`

To create a script that applies only to a specific policy or policy and schedule combination, add a *.policyname* or *.policyname.schedulename* suffix to the script name as follows:

- The following script applies only to a policy named days:

```
Install_path\netbackup\bin\bpend_notify.days.bat
```

- The following script applies only to a schedule that is named fulls in a policy named days:

```
Install_path\netbackup\bin\bpend_notify.days.fulls.bat
```

Note: The `bpend_notify` script also runs for NetBackup catalog backups if a `.policyname[.schedule]` is not specified.

The first script affects all scheduled backups in the policy named days. The second script affects scheduled backups in the policy 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.policy.schedule.bat  
bpend_notify.policy.bat  
bpend_notify.bat
```

For example, if there are both `bpend_notify.policy.bat` and `bpend_notify.policy.schedule.bat` scripts, NetBackup uses only `bpend_notify.policy.schedule.bat`.

Note: `bpstart_notify` scripts 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.policy.bat` and `bpend_notify.policy.schedule.bat`.

NetBackup passes the following parameters to the script when the backup completes:

- %1 Specifies the name of the client from the NetBackup catalog.
- %2 Specifies the policy name from the NetBackup catalog.
- %3 Specifies the schedule name from the NetBackup catalog.

- %4** Specifies one of the following: `FULL`, `INCR`, `CINC`, `UBAK`, `UARC`
- %5** Specifies the status of the operation. It is the same status as is sent to the NetBackup server. The status is 0 for successful backups and 1 for partially successful backups. If an error occurs, the status is the value associated with that error.
- %6** Specifies the 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 policy and schedule, the results file must be named
- ```
Install_path\netbackup\bin\BPEND_RES.policy.schedule
```
- If the script applies to a specific policy, the results file must be named
- ```
Install_path\netbackup\bin\BPEND_RES.policy
```
- If the script applies to all backups, the results file must be named
- ```
Install_path\netbackup\bin\BPEND_RES
```
- An echo `0 > %6` statement is one way for the script to create the file.
- NetBackup deletes the existing results file before it calls the script. After the script runs, 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 time that the `BPEND_TIMEOUT` option specifies. 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 clients, the `bpend_notify` script can use the following environment variables for the support of multiple data streams.

**Table 4-8** Environment variables for support of multiple data streams

| Environment variable       | Description                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>STREAM_NUMBER</code> | Specifies the stream number. The first stream from a policy, client, and schedule is 1. A 0 value indicates that multiple data streams are not enabled. |
| <code>STREAM_COUNT</code>  | Specifies the total number of streams to be generated from this policy, client, and schedule.                                                           |
| <code>STREAM_PID</code>    | Specifies the pid (process ID) number of <code>bpbkar</code> .                                                                                          |

## diskfull\_notify.cmd on Windows

The `diskfull_notify.cmd` script runs on the NetBackup server that contains the storage unit. The disk media manager (`bpdm`) calls this script if it encounters a disk full condition while it writes a backup to a disk storage unit. The default action is to report the condition and immediately try to write the data again. (The file being written is kept open by the active `bpdm`).

The script can be modified to send a notification to an email address. Or modified to perform actions such as removing other files in the affected directory or file system.

NetBackup passes the following parameters to the script:

|                          |                                                                |
|--------------------------|----------------------------------------------------------------|
| <code>programname</code> | Specifies the name of the program (always <code>bpdm</code> ). |
| <code>pathname</code>    | Specifies the path to the file being written.                  |

For example:

```
diskfull_notify.cmd bpdm

/disk1/images/host_08193531_c1_F1
```

In previous releases, the `diskfull_notify.cmd` script default condition was to sleep for five minutes when a disk storage unit became full. To retain this behavior upon upgrade, do one of the following:

- Copy the `netbackup/bin/diskfull_notify.old_revision_number` script to `netbackup/bin/diskfull_notify`, or
- Modify the script, to change `sleep 0` to:

```
sleep 300
```

## mail\_dr\_info.cmd

Use `mail_dr_info.cmd` to send NetBackup disaster recovery information to specified recipients after running an online, hot catalog backup.

To create the script, copy the following script from the master server:

```
Install_path\VERITAS\NetBackup\bin\nbmail.cmd
```

Place it into the following location:

```
Install_path\NetBackup\bin\mail_dr_info.cmd.
```

NetBackup passes the following parameters to the script:

- %1 Specifies the recipient's address. For multiple addresses, enter *email1,email2*
- %2 Specifies the subject line.
- %3 Specifies the message file name.
- %4 Specifies the attached file name.

NetBackup checks to see if `mail_dr_info.cmd` is present in

`Install_path\NetBackup\bin`. If `mail_dr_info.cmd` exists, NetBackup passes the parameters to the script.

---

**Note:** All NetBackup email notifications require that a public domain SMTP mail client be configured. (For example, blat.) For details, see the comments in the `nbmail.cmd` script.

---

## media\_deassign\_notify

The NetBackup Enterprise Media Manager calls the `media_deassign_notify` script after media is deassigned. To send an email notification when media is deassigned, include an email address in the script where indicated. (The script must be run as the root user.)

Copy `Install_path\NetBackup\bin\goodies\media_deassign_notify.cmd` into `Install_path\NetBackup\bin\` on the EMM server. (Usually the master server.)

If the script exists in the `\bin` directory, the following parameters are passed to the script: media ID, legacy media type, barcode, robot number, and robot type.

## nbmail.cmd

Use `nbmail.cmd` to send specified recipients notifications about scheduled backups. The recipients email addresses must also be configured in the **Universal Settings** host properties.

Windows systems also require that you install the Simple Mail Transfer Protocol application to transfer messages in order to accept script parameters. UNIX platforms have a built-in SMTP transfer method.

To create the script on a client, copy

`Install_path\VERITAS\NetBackup\bin\goodies\nbmail.cmd` from the master server into `Install_path\NetBackup\bin` of each client that is to receive the notification.

NetBackup passes the following parameters to the script:

- %1 Specifies the address of the recipient. For multiple addresses, enter *email1,email2*
- %2 Specifies the contents of the subject line.
- %3 Specifies the file that is sent in the body of the email. This is generated by another script.
- %4 Specifies the attached file name.

NetBackup checks to see if `nbmail.cmd` is present in `Install_path\NetBackup\bin`. If `nbmail.cmd` exists, NetBackup passes the parameters to the script.

## parent\_end\_notify.cmd on Windows

NetBackup calls the `parent_end_notify.cmd` script each time a parent job ends.

NetBackup passes the following parameters to the script:

- `clientname` Specifies the name of the client from the NetBackup catalog.
- `policyname` Specifies the policy name from the NetBackup catalog.
- `schedname` Specifies the schedule name from the NetBackup catalog.
- `schedtype` Specifies one of the following: `FULL`, `INCR` (differential incremental), `CINC` (cumulative incremental), `UBAK`, `UARC`
- `status` Specifies the exit code for the entire backup job.
- `stream` Specifies the stream number; it is always -1.
- `stream_count` Specifies that if the job starts normally, the stream count indicates how many streams were started.  
  
Verifies the number of streams that complete and run `backup_exit_notify`. If a failure occurs that makes it impossible to start any streams, a stream count of -1 is returned.

## parent\_end\_notify

The `parent_end_notify` script runs on the master server and is located in the following directory:

**Windows:** `Install_path\VERITAS\NetBackup\bin\goodies\`

**UNIX:** `/usr/opensv/netbackup/bin/goodies`

NetBackup calls the `parent_end_notify` script each time a parent job ends. In NetBackup 6.5.6, two new parameters have been added the script.

NetBackup passes the following parameters to the script:

|                           |                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>clientname</code>   | Specifies the name of the client from the NetBackup catalog.                                                                                                                                                                                                                                                                                       |
| <code>policyname</code>   | Specifies the policy name from the NetBackup catalog.                                                                                                                                                                                                                                                                                              |
| <code>schedname</code>    | Specifies the schedule name from the NetBackup catalog.                                                                                                                                                                                                                                                                                            |
| <code>schedtype</code>    | Specifies one of the following: <code>FULL</code> , <code>INCR</code> (differential incremental), <code>CINC</code> (cumulative incremental), <code>UBAK</code> , <code>UARC</code>                                                                                                                                                                |
| <code>status</code>       | Specifies the exit code for the entire backup job.                                                                                                                                                                                                                                                                                                 |
| <code>stream</code>       | New in 6.5.6; specifies the stream number. The stream number is always -1.                                                                                                                                                                                                                                                                         |
| <code>stream_count</code> | New in 6.5.6; specifies the stream count. If the job starts normally, the stream count indicates how many streams were started.<br><br>Use this count to verify the number of streams that complete and run <code>backup_exit_notify</code> . If a failure occurs that makes it impossible to start any streams, a stream count of -1 is returned. |

## parent\_start\_notify.cmd on Windows

NetBackup calls the `parent_start_notify.cmd` script each time a parent job starts.

NetBackup passes the following parameters to the script:

|                           |                                                                                                                                                                                     |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>clientname</code>   | Specifies the name of the client from the NetBackup catalog.                                                                                                                        |
| <code>policyname</code>   | Specifies the policy name from the NetBackup catalog.                                                                                                                               |
| <code>schedname</code>    | Specifies the schedule name from the NetBackup catalog.                                                                                                                             |
| <code>schedtype</code>    | Specifies one of the following: <code>FULL</code> , <code>INCR</code> (differential incremental), <code>CINC</code> (cumulative incremental), <code>UBAK</code> , <code>UARC</code> |
| <code>status</code>       | Specifies the exit code for the entire backup job.                                                                                                                                  |
| <code>streamnumber</code> | Specifies the stream number; for a parent job it is always -1.                                                                                                                      |

## pending\_request\_notify

The NetBackup Enterprise Media Manger calls the `pending_request_notify` script after a pending request is issued for a media resource (tape volume). To send an email notification when a pending request is initiated, include an email address in the script where indicated. (The script must be run by the root user.)

Copy `Install_path\NetBackup\bin\goodies\pending_request_notify.cmd` into `Install_path\NetBackup\bin\` on the EMM server. (Usually the master server.)

If the script exists in the `\bin` directory, the following parameters are passed to the script: media ID, barcode, action code, robot type, robot number, media server, volume group, and pending time (in seconds since the UNIX epoch).

## restore\_notify.cmd on Windows

The `restore_notify.cmd` script runs on the server that contains the storage unit. The NetBackup tape or disk manager (`bptm` or `bpdm`) calls the script when it finishes sending data to the client during a restore. The script is called regardless of whether data is sent.

NetBackup passes the following parameters to the script:

|                          |                                                                                                                             |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <code>programname</code> | Specifies the name of the program doing the restore or other read operation.                                                |
| <code>pathname</code>    | Specifies the path to the backup name or path.                                                                              |
| <code>operation</code>   | Specifies one of the following: <code>restore</code> , <code>verify</code> , <code>duplication</code> , <code>import</code> |

## session\_notify.cmd on Windows

The `session_notify.cmd` script runs on the master server. It is called at the end of a backup session if at least one scheduled backup succeeded. NetBackup passes no parameters to this script. Scheduling is suspended until this script completes, so no other backups can start until that time.

## session\_start\_notify.cmd on Windows

The `session_start_notify.cmd` 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 before it starts the first backup. NetBackup passes no parameters to this script.

## shared\_drive\_notify.cmd on Windows

NetBackup runs the `shared_drive_notify.cmd` script when a shared drive is reserved or released.

- The name of the shared drive.
- The name of the current scan host.
- The operation, which is one of the following:

|          |                                                                                                                                                                                                                                     |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RESERVED | Specifies that the host on which the script is executed needs SCSI access to the drive until it is released.                                                                                                                        |
| ASSIGNED | Informational only. Specifies that the host that reserved the drive needs SCSI access.                                                                                                                                              |
| RELEASED | Specifies that only the scan host needs SCSI access to the drive.                                                                                                                                                                   |
| SCANHOST | Specifies that the host that executes the script has become the scan host. A host should not become a scan host while the drive is RESERVED.<br><br>The scan host may change between a RESERVED operation and a RELEASED operation. |

The script resides in the following directory:

```
Install_path\VERITAS\Volmgr\bin\shared_drive_notify.cmd
```

The script must be executable by the root user.

The script exits with status 0 upon successful completion.

## userreq\_notify.cmd on Windows

The `userreq_notify.cmd` script runs on the master server.

NetBackup calls it each time a request is made to either of the following:

- List files that are in backups or archives
- Start a backup, archive, or restore

You can change this script to gather information about user requests to NetBackup.

NetBackup passes the following parameters to the script:

|                     |                                                                                                                                                                            |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>action</code> | Specifies the action and can have the following values: <code>backup</code> , <code>archive</code> , <code>manual_backup</code> , <code>restore</code> , <code>list</code> |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                         |                            |
|-------------------------|----------------------------|
| <code>clientname</code> | Specifies the client name. |
| <code>userid</code>     | Specifies the user ID.     |

For example:

```
userreq_notif.cmd backup mercury jdoe
userreq_notify.cmd archive mercury jdoe
userreq_notify.cmd manual_backup mercury jdoe
userreq_notify.cmd restore mercury jdoe
userreq_notify.cmd list mercury jdoe
```

## Media and device management best practices

Use the following best practices for NetBackup media and device management. Follow these recommendations to minimize problems and to reduce the time and the effort that is required to administer the configuration.

For a list of supported devices, server platforms, and the latest device mapping file, see the NetBackup support Web site:

<http://entsupport.symantec.com>.

The following items are general best practices for media and device management:

- Use only the NetBackup commands that Symantec documents and supports.
- Refer to the NetBackup release notes for configuration and operational changes in the current release or in future releases. The release notes also contain information about all new functionality in each release.
- Use the documented methods for terminating the NetBackup Media Manager daemons and services.
- Periodically verify the backups by using **NetBackup Management > Catalog** in the **NetBackup Administration Console**. Also, periodically restore files to prove that restores work correctly.
- Always back up the NetBackup catalogs. You may also want to back up the `vm.conf` file and the `bp.conf` (UNIX system) files on the media servers.
- When you restore the NetBackup catalog (for example, master server databases and the EMM database), use backups from the same point in time.
- Ensure that all names and numbers for devices and all media IDs and bar codes are unique across the entire enterprise.
- To use the devices that NetBackup controls but are used with other applications, down the drive if the drive is in the UP state.

## Media management best practices

The following items are NetBackup media management best practices:

- Use the robot inventory update operation for media management.
- Use a scratch pool for unassigned media.
- Configure cleaning cartridges for tape drives and use TapeAlert for automatic drive cleaning if the drives support automatic cleaning.
- Replace old media according to the life-span recommendations of the manufacturer. Replace old cleaning media also.
- Use the robotic libraries that have a bar code reader and use only the bar code labels that the robot vendor recommends.
- Use bar code rules for media type assignment when you inventory multimedia libraries. Use bar code naming conventions to differentiate between data and cleaning tapes and different physical media types. A common convention is a prefix that identifies the type of media.
- Before performing inject or eject commands, ensure that the media access port is empty. Although NetBackup can handle a port that is not empty, some libraries can have problems.

## Device management best practices

The following items are device management best practices:

- Monitor the NetBackup system log for device errors encountered.
- Monitor devices by using the NetBackup Device Monitor.
- Investigate the causes of all the drives that are down.
- Do not use the robotic test utilities while running backup or restore jobs.
- Read the *NetBackup Device Configuration Guide* before configuring devices on media servers (or SAN media servers).
- Use only computers, operating systems and devices that Symantec supports. For supported devices, see the NetBackup hardware compatibility list on the NetBackup support site.
- Use only fully-serialized devices. A fully-serialized SCSI library should report a serial number for the robot and also a serial number for each drive in the robot.
- Always configure and use pass-through paths for robotic libraries and drives.
- When possible, use SCSI persistent reserve or SCSI reserve and release.

- Use persistent bindings for fibre-attached devices.
- Use the **NetBackup Device Configuration Wizard** to configure the devices.
- Download and install the latest device mapping file from the NetBackup support Web site before you use the **NetBackup Device Configuration Wizard**.
- Use consistent logical drive types for all physical drive types on all servers in the environment. For example, use the DLT drive type as the logical drive type for all DLT7000 drives.
- Do not load vendor medium-changer drivers on Microsoft Windows hosts. The default Microsoft medium-changer driver is acceptable (but is not required) for use with NetBackup.

## Media and device performance and troubleshooting

The following items are performance and troubleshooting best practices:

- Use the performance-tuning documents available on the NetBackup support Web page.
- Use only a dedicated server for the NetBackup master server and Enterprise Media Manager (EMM) server. Do not use a server that hosts other applications or one that stores data. Plan periodic maintenance for all of the backup servers.
- Consult the Troubleshooter in the **NetBackup Administration Console** or the *NetBackup Troubleshooting Guide* for all error conditions.
- Always install the latest NetBackup release updates that are available from Symantec.
- Verify all SCSI-related operating system configuration files (such as the Solaris `st.conf` file), when you install system release updates.
- For problems with devices, consult the vendor for firmware upgrades and consult the NetBackup hardware compatibility list for supported firmware levels.
- Do not use the NetBackup `DISABLE_RESOURCES_BUSY` touch file.
- Do not disable the operating system `TCP_NODELAY` functionality.

## About TapeAlert

TapeAlert is a tape drive status monitor and message utility. The TapeAlert utility can detect tape quality problems, defects in tape drive hardware, and the need to clean drives. For the tape drives that support TapeAlert, the TapeAlert firmware

monitors the drive hardware and the media. Error, warning, and informational states are logged on a TapeAlert log page.

For the drives that do not support TapeAlert, configure and use frequency-based cleaning.

See [“About frequency-based cleaning”](#) on page 130.

## About TapeAlert cleaning (reactive cleaning)

Reactive cleaning by using TapeAlert is a function of the tape drive. The drive determines and initiates the cleaning when needed. If a drive supports the TapeAlert capability and it is enabled on the drive, the NetBackup `bptm` process polls the drive for status from TapeAlert.

TapeAlert allows reactive cleaning for most drive types. Not all platforms, robots, drives, or firmware levels support tape alert reactive cleaning.

A drive with TapeAlert capability tracks how many read and write errors it has encountered within a certain time period. Although a drive can recover from these errors, the drive sets a `CLEAN_NOW` or `CLEAN_PERIODIC` flag when a threshold is reached.

If the `bptm` process detects that either of the following flags are set, it performs a cleaning at one of the following times:

- At the end of a backup or a restore to the drive.
- Before the next backup or restore to the drive.

Symantec recommends that you use reactive cleaning.

## About TapeAlert and frequency-based cleaning

Using TapeAlert with frequency-based cleaning ensures that a drive is cleaned at least every  $x$  hours, depending on the setting for the cleaning frequency. In addition, the drive can be cleaned sooner if the drive sets the `CLEAN_NOW` or `CLEAN_PERIODIC` TapeAlert flag.

When TapeAlert is used without frequency-based cleaning, a drive is cleaned only when the drive sets its `CLEAN_NOW` or `CLEAN_PERIODIC` flags.

## About TapeAlert requirements

To use TapeAlert, all of the following conditions must be true:

- The host platform, robot type, and drive support drive cleaning.

- The drive must support the TapeAlert capability, and the TapeAlert are enabled on the drive.  
 To determine if a drive supports TapeAlert, see the Symantec support Web site.
- A cleaning tape is configured and available in NetBackup for the robotic library. The cleaning cartridge is compatible with the drive that needs to be cleaned.
- The cleaning tape has not reached its end of life.
- Pass through device files are configured on UNIX and Linux media servers. See the *NetBackup Device Configuration Guide*.

## TapeAlert logs and codes

TapeAlert codes are derived from the T10 SCSI-3 Stream Commands standard (see <http://t10.org/>). For the list of codes that the device supports, see the device's documentation.

TapeAlert checks for errors of the following types:

- Recoverable read and write drive problems
- Unrecoverable read and write drive problems
- Hardware defects
- Wrong or worn-out media
- Expired cleaning tapes
- Abnormal errors

A set of TapeAlert conditions is defined that can cause the media in use to be frozen. Another set of conditions are defined that can cause a drive to be downed.

NetBackup writes TapeAlert conditions into the following logs:

- The `bptm` log
- The error log
- The job details log
- The system log on UNIX and Event Viewer on Windows

[Table 4-9](#) describes the codes.

**Table 4-9** TapeAlert log codes

| TapeAlert code | Default action | Error type    | Error message |
|----------------|----------------|---------------|---------------|
| 0x01           | None           | Warning - WRN | Read warning  |

**Table 4-9** TapeAlert log codes (*continued*)

| TapeAlert code | Default action     | Error type           | Error message                              |
|----------------|--------------------|----------------------|--------------------------------------------|
| 0x02           | None               | Warning - WRN        | Write warning                              |
| 0x03           | None               | Warning - WRN        | Hard error                                 |
| 0x04           | Freeze media - FRZ | Critical - CRT       | Media                                      |
| 0x05           | Freeze media - FRZ | Critical - CRT       | Read failure                               |
| 0x06           | Freeze media - FRZ | Critical - CRT       | Write failure                              |
| 0x07           | Freeze media - FRZ | Warning - WRN        | Media life                                 |
| 0x08           | Freeze media - FRZ | Warning - WRN        | Not data grade                             |
| 0x09           | None               | Critical - CRT       | Write protect                              |
| 0x0a           | None               | Informational - INFO | No removal                                 |
| 0x0b           | None               | Informational - INFO | Cleaning media                             |
| 0x0c           | None               | Informational - INFO | Unsupported format                         |
| 0x0d           | Freeze media - FRZ | Critical - CRT       | Recoverable mechanical cartridge failure   |
| 0x0e           | Freeze media - FRZ | Critical - CRT       | Unrecoverable mechanical cartridge failure |
| 0x0f           | Freeze media - FRZ | Warning - WRN        | Mic failure                                |
| 0x10           | None               | Critical - CRT       | Forced eject                               |
| 0x11           | None               | Warning - WRN        | Read only                                  |
| 0x12           | None               | Warning - WRN        | Directory corrupted on load                |
| 0x13           | Freeze media - FRZ | Informational - INFO | Nearing media life                         |
| 0x14           | Clean drive - CLN  | Critical - CRT       | Clean now                                  |
| 0x15           | Clean drive - CLN  | Warning - WRN        | Clean periodic                             |
| 0x16           | Freeze media - FRZ | Critical - CRT       | Expired cleaning media                     |

**Table 4-9** TapeAlert log codes (*continued*)

| TapeAlert code | Default action     | Error type           | Error message               |
|----------------|--------------------|----------------------|-----------------------------|
| 0x17           | Freeze media - FRZ | Critical - CRT       | Invalid cleaning tape       |
| 0x18           | None               | Warning - WRN        | Retension requested         |
| 0x19           | None               | Warning - WRN        | Dual-port error             |
| 0x1a           | None               | Warning - WRN        | Cooling fan failure         |
| 0x1b           | None               | Warning - WRN        | Power supply failure        |
| 0x1c           | None               | Warning - WRN        | Power consumption           |
| 0x1d           | None               | Warning - WRN        | Drive maintenance           |
| 0x1e           | Down drive - down  | Critical - CRT       | Hardware A                  |
| 0x1f           | Down drive - DOWN  | Critical - CRT       | Hardware B                  |
| 0x20           | None               | Warning - WRN        | Interface                   |
| 0x21           | None               | Critical - CRT       | Eject media                 |
| 0x22           | None               | Warning - WRN        | Download fail               |
| 0x23           | None               | Warning - WRN        | Drive humidity              |
| 0x24           | None               | Warning - WRN        | Drive temperature           |
| 0x25           | None               | Warning - WRN        | Drive voltage               |
| 0x26           | None               | Critical - CRT       | Predictive failure          |
| 0x27           | None               | Warning - WRN        | Diagnostics req.            |
| 0x28 - 0x31    | None               | Informational - INFO | Undefined                   |
| 0x32           | None               | Warning - WRN        | Lost statistics             |
| 0x33           | Freeze media - FRZ | Warning - WRN        | Directory invalid on unload |
| 0x34           | Freeze media - FRZ | Critical - CRT       | System area write failure   |
| 0x35           | Freeze media - FRZ | Critical - CRT       | System area read failure    |
| 0x36           | Freeze media - FRZ | Critical - CRT       | No start of data            |

**Table 4-9** TapeAlert log codes (*continued*)

| TapeAlert code | Default action     | Error type           | Error message                |
|----------------|--------------------|----------------------|------------------------------|
| 0x37           | Freeze media - FRZ | Critical - CRT       | Loading failure              |
| 0x38           | Freeze media - FRZ | Critical - CRT       | Unrecoverable unload failure |
| 0x39           | None               | Critical - CRT       | Automation interface failure |
| 0x3a           | None               | Warning - WRN        | Firmware failure             |
| 0x3d - 0x40    | None               | Informational - info | Undefined                    |

## About tape drive cleaning

The following types of drive cleaning are available by using NetBackup:

- Reactive cleaning  
See [“About TapeAlert cleaning \(reactive cleaning\)”](#) on page 125.  
Symantec recommends that you use reactive cleaning.
  - Library-based cleaning  
See [“About library-based cleaning”](#) on page 129.
  - Frequency-based cleaning  
See [“About frequency-based cleaning”](#) on page 130.
  - Operator-initiated cleaning  
See [“About operator-initiated cleaning”](#) on page 130.
- See [“About using a cleaning tape”](#) on page 131.

## About library-based cleaning

NetBackup does not support library-based cleaning for most robots because robotic library and operating systems vendors implement this cleaning in different ways. (Library-based cleaning also is known as robotic cleaning or auto cleaning.) These different methods often interfere with NetBackup robotic control operations.

NetBackup does not define the cleaning media that is used for library-based cleaning, and the robotic library manages the cleaning media.

Because TapeAlert provides the same type of cleaning as library-based cleaning, Symantec recommends disabling library-based cleaning when you use TapeAlert.

## About frequency-based cleaning

Frequency-based cleaning occurs when the accumulated mount time exceeds the time you specify for the cleaning frequency. NetBackup updates the mount time for the drive each time a tape is unmounted.

The cleaning frequency is configured when a drive is added to NetBackup. Change the cleaning frequency by changing the drive properties or by using the **Media and Device Management Device Monitor** in the **NetBackup Administration Console**.

If the following conditions are met, drive cleaning occurs when the accumulated mount time exceeds the time specified for the cleaning frequency:

- The drive is in a robotic library that supports drive cleaning.
- A cleaning tape is configured and available for the robotic library.
- The cleaning tape has cleanings remaining.

NetBackup cleans the drive immediately after a tape is unmounted. Drive cleaning does not unmount a drive in the middle of an active backup. The mount time is reset after the drive is cleaned. The cleaning frequency value remains the same.

A cleaning can occur within a backup if the backup spans tapes. For example, if cleaning is due after the first tape is full, NetBackup cleans the drive before it mounts the next tape.

Media can remain in a drive for extended periods. It does not affect the cleaning frequency because NetBackup increments the mount time only when NetBackup assigns the media to a process.

Frequency-based cleaning is not supported for drives in the ACS or the TLH libraries that are under API robotic control. The robotic library software controls the drive cleaning. To manage drive cleaning for these robots, use the robot vendor interfaces.

See [“About TapeAlert and frequency-based cleaning”](#) on page 125.

## About operator-initiated cleaning

A drive cleaning can be initiated regardless of the cleaning frequency or accumulated mount time of the drive. Clean standalone drives or robotic drives if a cleaning tape of the correct media type and residence for the drive was added to NetBackup.

NetBackup reports that a drive needs cleaning if either of the following conditions are true:

- The value for the mount time is greater than the cleaning frequency.

- The TapeAlert CLEAN\_NOW or CLEAN\_PERIODIC flag is set.

And either of the following conditions must be true:

- The drive is a standalone drive and a cleaning tape is not defined.
- The drive is a standalone drive and no cleaning tape has any cleanings that remain.

NetBackup displays NEEDS CLEANING as follows:

- The **Tape Cleaning Comment** column of the **Drive List** in the **Devices** node of the **NetBackup Administration Console**.
- The comment field of the output from the `tpclean -L` command.

## About using a cleaning tape

You can specify the number of cleanings that are allowed for a cleaning tape. This number is reduced with each cleaning. When the number of cleanings is zero, NetBackup stops by using the cleaning tape. Then, use a new cleaning tape or increase the number of cleanings that are allowed for the tape.

---

**Note:** NetBackup does not control the cleaning tapes that library-based cleaning uses.

---

Symantec suggests following the recommendations from cleaning tape vendors for the amount of tape usage. If you clean a tape past its recommended life, cleaning delays can occur (due to excessive tape position operations) and drives can be downed.

## How NetBackup selects drives

NetBackup stores media information and device configuration and status information in the EMM database. When a robotic mount request is issued, the NetBackup Resource Broker (`nbrb`) queries the EMM database for the media ID of the volume requested. If the volume is in the EMM database, the media request is matched with a compatible drive in the robot. The mount request is forwarded to the appropriate robotic daemon (UNIX) or process (Windows) based on the location of the media. Location is the robotic library and the storage slot number, if applicable.

A drive must meet the following criteria to be selected for the mount request:

- The drive is configured.
- The drive is in the robotic library that contains the media.

- The drive allows the requested media density.

The EMM server (`nbemm`) manages the drives and requests for locally-attached or shared drives in the EMM domain.

The EMM server manages the drives by doing the following actions:

- Determines which of the drives are currently available.

A drive is available if it is one of the following:

- Configured as UP
- Not assigned
- Compatible with the media type
- Not reserved by another host

- Picks an available drive that was least recently used.

NetBackup selects the robotic-based drives over standalone drives unless the correct media already is loaded in a standalone drive.

The first drive in the drive configuration is used first, and then the second drive, and then the next. Use the `tpconfig -d` command to see the drive order in the configuration.

If some of the drives are shared drives, NetBackup chooses a nonshared drive first (if one is available). NetBackup chooses a shared drive first so the shared drives can be used on other hosts that share the drives. Shared drives require the Shared Storage Option.

## How NetBackup reserves drives

In multiple-initiator (multiple host bus adapter) environments, device-level access protection is required to avoid unintended sharing of tape devices and possible data loss problems. (Shared Storage Option is a multiple-initiator environment.) Access protection on a tape drive prevents an HBA that is not the reservation owner from issuing commands to control the drive. SCSI access protection operates at the SCSI target level and depends on correct operation of the fiber-to-SCSI bridge or the native fiber device hardware.

The only commonly available technique for this purpose is SPC-2 SCSI reserve and release functionality. All tape drive vendors support the SPC-2 SCSI reserve method. NetBackup has used SPC-2 SCSI reserve since NetBackup 3.4.3; it is the default tape drive reservation method in NetBackup. SPC-2 SCSI reserve is effective for most NetBackup environments.

Alternatively, the new SCSI persistent reserve method may be more effective in either of the following environments because it provides device status detection and correction:

- NetBackup media servers are in a cluster environment  
NetBackup can recover and use a reserved drive after a failover (if NetBackup owns the reservation). (With SPC-2 SCSI reserve, a drive reset usually is required because the reservation owner is inoperative.)
- Environments where high drive availability is important  
NetBackup can resolve NetBackup drive reservation conflicts and maintain high drive availability. (SPC-2 SCSI reserve provides no method for drive status detection.)  
However, the SCSI persistent reserve method is not supported or not supported correctly by all device vendors. Therefore, analyze the environment to ensure that all of the hardware supports SCSI persistent reserve correctly.  
NetBackup lets you configure either SCSI persistent reserve or SPC-2 SCSI reserve.

The following table describes the protection options.

**Table 4-10** Protection options

| Option                       | Description                                                                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SCSI persistent reserve      | Provides SCSI persistent reserve protection for SCSI devices. The devices must conform to the SCSI Primary Commands - 3 (SPC-3) standard.                                        |
| SPC-2 SCSI reserve (default) | Provides SPC-2 SCSI reserve protection for SCSI devices. The devices must conform to the reserve method and release management method in the SCSI Primary Commands - 2 standard. |
| No protection                | Other HBAs can send the commands that may cause a loss of data to the tape drives.                                                                                               |

You can configure access protection for each NetBackup media server. The protection setting configures tape drive access protection for all tape drive paths from the media server on which the setting is configured. The media server setting for any drive path can be overridden.

SCSI reservations provide protection for NetBackup Shared Storage Option environments or any other multiple-initiator environment in which drives are shared.

## About SCSI persistent reserve

The NetBackup process that reads from or writes to the media in a drive (`bptm`) issues SCSI persistent reserve commands to do the following:

- Register with the tape drive's device server (the server is a logical unit within a drive that processes SCSI tasks)
- Request an exclusive access reservation

If the tape drive's device server grants the reservation, the NetBackup process has exclusive use of the device. The reservation prevents other host bus adapters (HBAs) from issuing any commands that can cause data loss.

If the reservation fails, NetBackup fails the job.

When the NetBackup process is finished with the drive, NetBackup unloads the drive and sends a persistent reserve clear command to the drive. The command removes both the reservation and the registration.

SCSI persistent reserve also provides device status detection, which NetBackup uses to resolve reservation conflicts within NetBackup.

The reservation does not prevent other applications on the host that has the reservation from using the same device and from causing data loss. For example, if a user on the same host issues a UNIX `mt` command, the `mt` command can take control of the drive.

Also, other HBAs can clear or release a SCSI persistent reservation. Therefore, an application can clear another HBA reservation (although it should not do so).

## About SCSI persistent reserve commands

When a device receives an exclusive access type SCSI persistent reservation command, it does not process commands from any other HBA. The device processes commands from another HBA only when the HBA that owns the SCSI persistent reservation clears the reservation. If an application sends a command to a reserved device, the device fails the command by returning a status of RESERVATION CONFLICT. The only exceptions to this action are several commands that cannot interfere with the reservation, such as Inquiry or Request Sense.

A device stays reserved until one of the following events occurs on the device:

- Released by the HBA that reserved it
- Power cycled (usually)
- Preempted by a SCSI persistent reserve command

## About SCSI persistent reserve conflicts

NetBackup uses unique reservation keys. Therefore, NetBackup attempts to resolve conflicts with other NetBackup reservations. If a conflict exists, NetBackup sends SCSI commands to unload the drive. Based on the drive status, NetBackup tries

to unload the drive again by using additional information to release or preempt the persistent reservation.

In cluster environments after a failover event, NetBackup on the active cluster node detects the persistent reservation and clears the reservation. NetBackup regains use of the drive without power-cycling the drive.

If NetBackup does not own the persistent reservation, NetBackup reports a pending status in the Device Monitor. The reservation owner must clear the reservation before NetBackup can use the drive. For example, NetBackup does not clear a NetApp persistent reservation.

## About the SPC-2 SCSI reserve process

The NetBackup process issues an SPC-2 SCSI reserve command to the tape drive that contains the media. (The process can be `bptm`, `bprecover`, or `bpbackupdb`.) If the device is not reserved, NetBackup acquires a reservation. The drive does not process commands from any other host bus adapters (HBAs) until NetBackup releases the reservation or the reservation is broken. If the reservation fails, NetBackup fails the job.

The reservation does not prevent other applications on the host that has the reservation from using the same device and from causing data loss. For example, if a user on the same host issues a UNIX `mt` command, the `mt` command can take control of the drive.

After the NetBackup process finishes with the media, it issues an SPC-2 SCSI command to release the reservation during the unmount operation. The release frees the device for access by another HBA.

SCSI reserve does not provide a method to determine if a device is reserved. Only the reservation owner (the host bus adapter) can release the reservation. However, these limitations do not interfere with NetBackup operations in most environments.

## About SPC-2 SCSI reserve commands

When a device receives an exclusive access type SCSI persistent reservation command, it does not process commands from any other HBA. The device processes commands from another HBA only when the HBA that owns the reservation issues the release command. If an application sends a command to a reserved device, the device fails the command by returning a status of RESERVATION CONFLICT. The only exceptions to this action are several commands that cannot interfere with the reservation, such as Inquiry or Request Sense.

A device stays reserved until one of the following events occurs on the device:

- Released by the HBA that reserved it
- Released by a TARGET or a LOGICAL UNIT RESET  
These resets are protocol dependent and differ between parallel SCSI and FCP (SCSI on Fibre Channel ). These resets can be issued from any HBA.
- Released by Fibre Channel LOGO, PLOGO, PRLI, PRLO, or TPRLO action or failed discovery (link actions)
- Power cycled

A negative consequence of SPC-2 SCSI reserve occurs if the HBA that owns the reservation fails. A device stays reserved until the reservation is removed or broken. Only the original HBA can remove the reservation, which means the system must be available. If the HBA that owns the reservation fails, it cannot remove the reservation. Therefore, the reservation must be broken.

To break a reservation, one of the following actions must break the reservation:

- SCSI reset
- Bus device reset
- LUN device reset
- Power cycle
- Fibre Channel link actions may break reservations

SPC-2 SCSI reserve commands are mandatory for all SCSI-2 and SCSI-3 devices. See the SCSI 2 standard for a detailed description of SCSI reserve command operation and behavior.

## About SCSI reservation conflicts

The NetBackup Automatic Volume Recognition process (`avrd`) manages access to tape devices. A properly configured NetBackup environment and properly configured tape devices should not receive a reservation conflict message from a tape drive. When `avrd` starts, it issues an SPC-2 SCSI release to all configured, nondisabled tape drive paths that are currently in the Up state. The command releases all devices that were SPC-2 reserved at the time of a system restart or crash. The SCSI release command returns tape devices to general availability after a system crash.

If the `avrd` process receives a reservation conflict message, it changes the status of the device to PENDING. It also writes the following message in the system log:

```
Reservation Conflict status from DRIVENAME (device NUMBER)
```

Also, the **NetBackup Administration Console Device Monitor** or the output from the `vmopr cmd` command shows PENDING in the Control column.

If a conflict occurs, a reservation problem can exist. If the HBA that reserves the drive is unavailable (for example, due to a system crash or hardware failure), it cannot release the reservation. NetBackup cannot release or break an SPC-2 SCSI reservation automatically. Force a release or break the reservation to make the drive available, even for a failover server in a cluster environment.

When the conflict is resolved, the following message is written to the log:

```
Reservation Conflict status cleared from DRIVENAME (device NUMBER)
```

## About forcing a release of an unavailable HBA's SPC-2 reservation

To force a release of an unavailable HBA's SPC-2 reservation, use the following NetBackup `vmopr cmd` command and option:

```
vmopr cmd -crawlreleasebyname drive_name
```

This option requests that all hosts that are registered to use the drive issue SPC-2 SCSI release commands to the drive.

Issue the `vmopr cmd` command on the host that is the device allocator (DA host). Alternatively, use the `-h` option of the command to specify the DA host. The DA host is also the EMM server.

---

**Note:** Use this command after a PENDING status appears in the **NetBackup Administration Console Device Monitor**. However, do not issue this command during backups.

---

More information about using the `vmopr cmd` command is available.

See *NetBackup Commands Reference Guide*.

## Breaking a reservation

If you cannot release an SPC-2 SCSI reservation, try to use an operating system command that forces a device reset. A device reset breaks a reservation. The procedure depends on the operating system type.

---

**Note:** The reset operation can reset other devices in the configuration. Loss of data is also possible. Try alternate methods first to break the reservation on a device (by using switch and bridge hardware).

---

Lastly, if the following operating system commands cannot break the reservation, power-cycle the drive. A power cycle breaks SPC-2 SCSI drive reservations (and usually breaks SCSI persistent drive reservations).

#### To break an SPC-2 reservation on Solaris

- 1 Issue `mt -f drive_path_name forcereserve`.
- 2 Issue `mt -f drive_path_name release`.

See the `mt(1)` man page for more information.

#### To break an SPC-2 reservation on HP-UX

- ◆ Issue `st -f drive_path_name -r`.

See the `st(1m)` man page for more information.

#### To break an SPC-2 reservation on AIX

- ◆ Issue `tctl -f drive_path_name reset`.

See the `tctl` man page (in the IBM AIX Commands Reference) for more information.

## About SCSI reserve requirements

To use SCSI persistent reserve or SPC-2 SCSI reserve, the following requirements must be met:

- There must be pass through driver access to all shared drives. The pass through driver must be installed and all required paths must be created. Information about how to configure and use the pass through driver for UNIX operating systems is available. See the *NetBackup Device Configuration Guide*.
- You must configure the operating systems on the NetBackup media servers so they let NetBackup control SCSI persistent reserve or SPC-2 SCSI reserve.
- On HP-UX systems, disable the operating system's use of SPC-2 SCSI reserve. See the *NetBackup Device Configuration Guide*.
- Depending on the tape drives, you may have to disable the operating system's use of SPC-2 SCSI reserve. AIX and Solaris may require such a change. See the *NetBackup Device Configuration Guide*.

## About SCSI reserve limitations

The NetBackup implementation of SCSI persistent reserve and SPC-2 reserve has the following limitations:

- SCSI persistent reserve and SPC-2 reserve do not apply to NDMP drives. The NDMP filer is responsible for providing exclusive device access.
- Third-party copy configurations must be configured correctly. To retain reservation of a tape device during a third-party copy backup, configure the NetBackup `mover.conf` file. Do not use SCSI persistent reserve on the drive paths that are used for third-party copy backups. See the *NetBackup Snapshot Client Administrator's Guide*.
- With SPC-2 SCSI reserve, devices may remain reserved after a failover in cluster environments or multipath environments with failover capability. You cannot use SPC-2 SCSI reserve if the following factors are true: The failover does not break the device reservations and those devices that were in use during the failover must be available without manual intervention. Use SCSI persistent reserve.
- If the drive path changes, the backup jobs and the restore jobs fail. Therefore, jobs fail in cluster environments or any multipath environments that share paths dynamically. If you cannot disable dynamic path sharing, you cannot use SPC-2 SCSI reserve or SCSI persistent reserve in NetBackup.

## About SCSI reservation logging

The `bptm` process logs SCSI reservation-related commands. Examine the `bptm` log on all NetBackup media servers to ensure that the SCSI operations are logged. SCSI reservation commands are labeled SCSI PERSISTENT RESERVE or SCSI RESERVE in the log.

In addition, information about the SCSI persistent reservations that are broken are also written to the NetBackup Problems report.

## About server operating system limitations

This topic applies to Windows servers.

Windows operating systems cannot distinguish between a reserved device and a busy device. Therefore, PEND appears in the Device Monitor if another application controls the tape drive. NetBackup cannot share tape devices with other applications. If you use other applications, use the NetBackup `tpreq` command or Down the drive before using the drive.

These operating systems also may report PENDING if the drive reports Busy when a volume is unmounted. Use the `AVRD_PENDING_DELAY` entry in the `vm.conf` configuration file to filter out these extraneous reports.

## About checking for data loss

To detect data loss, the `bptm` process reads the tape position and then verifies the actual position against the expected position.

If the actual position is less than the expected position at the end of the backup process, the following events occur:

- The tape is frozen.
- The backup fails.
- The following error message entry is written to the `bptm` log:

```
FREEZING media id xxxxxx, External event caused rewind during
write, all data on media is lost
```

## About possible data loss causes

If tape drive access protection is not enabled on the NetBackup media servers, the following may cause data loss: configuration errors, incorrect paths, multiple master servers, incorrect Shared Storage Option configurations, and third-party or operating system utilities.

If access protection is enabled on all NetBackup media servers, the following can cause data loss: any third-party or operating system utilities that run on the server that runs the NetBackup backup job.

Unfortunately, data loss cannot be prevented only recognized after the fact. NetBackup does not remove catalog information about the backup sessions that were lost. Use the `bpxpdate` command to expire the images for the lost backup sessions.

## About checking for tape and driver configuration errors

To detect data loss, the `bptm` process reads the tape position and then verifies the actual position against the expected position.

If a configuration problem causes the actual position to be greater than the expected position at the end of the backup process, the following events occur:

- The tape is frozen.
- The backup fails.

- The following error message entry is placed in the `bptm` log:

```
FREEZING media id xxxxxx, too many data blocks written, check
tape/driver block size configuration
```

The backup data may be usable. If so, import the image by using the NetBackup `bpimport` command so the data is available for restores.

## About common configuration problems

Identify and fix the source of the configuration problem that causes data loss. The most common configuration error is a failure to configure the driver for variable length blocks.

A less common error may be in the tape driver's configuration data, such as in the `/kernel/drv/st.conf` file on a Solaris system.

Information about tape driver configuration is available.

See the *NetBackup Device Configuration Guide*.

## About configuring SCSI reserve

The SCSI reserve protection setting configures tape drive access protection for all tape drives from the media server on which the setting is configured. You can configure the protection for each media server and override the global setting for any drive path.

To configure SCSI reserve protection on a media server: use the **NetBackup Administration Console** to set the media server host property **Enable SCSI Reserve** on the **Media** tab.

To override the media server protection setting: use the **NetBackup Administration Console** to set the drive path property **Override SCSI reserve settings** when you add a drive or change a drive's properties.

## How NetBackup selects media

How NetBackup selects media depends on whether the media is in a robot or a standalone drive.

See [“About selecting media in robots”](#) on page 142.

See [“About selecting media in standalone drives”](#) on page 145.

## About selecting media in robots

When NetBackup receives a request for a volume, it searches the EMM database for the media ID. The external media ID should correspond to the NetBackup media ID.

A request for a volume includes the following attributes:

- The media ID
- The device density
- The file name that is used to link to the device that is assigned.

NetBackup selects a volume in a robot in the following order:

NetBackup searches the 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 that the backup schedule requires. However, if the NetBackup **Media** host property **Allow multiple retentions per media** is specified for the server, NetBackup does not search by retention level.
- In the volume pool that the backup job requires.
- Not in a FULL, FROZEN, IMPORTED, or SUSPENDED state.
- Of the same density that the backup job requested, and in the robot that that the backup job requested.
- Not currently in use by another backup or a restore.
- Not written in a protected format. NetBackup detects tape format after the volume is mounted. If the volume is in a protected format, NetBackup unmounts the volume and resumes the search.

If a suitable volume is found, NetBackup uses it.

If NetBackup cannot find a mounted volume that satisfies all of the previous conditions, it checks the media catalog for any volume that is suitable.

- If a suitable volume is in a robot, NetBackup issues the commands that move the volume to a drive, position the heads to the beginning of the volume, and assign it to the request. No manual intervention is required.
- If a suitable volume is not in a robot but is in a standalone drive, NetBackup automatically mounts and assigns it. No manual intervention is required.
- If a suitable volume is not in a robot or a standalone drive and the request is media-specific, NetBackup may pend a mount request. A media-specific mount request is one for a restore, for an import, or from the `tpreq` command.
- If a suitable volume is not in a robot or a standalone drive, NetBackup may attempt to use another volume only as follows: For backup jobs for which any other media can be used.

If a suitable volume does not exist or if a suitable volume is at end of media (EOM), NetBackup assigns a new volume. NetBackup may assign a new volume even if a volume is not full (because NetBackup received an EOM message from the drive).

The new volume must meet 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 NetBackup)
- Has not exceeded the maximum number of mounts allowed

If more than one volume qualifies, NetBackup chooses the volume that was least recently used.

NetBackup then adds it to the media catalog and assigns it the specified retention level.

If there are no unassigned volumes of the requested type, the backup terminates with an error message that no media were available.

NetBackup selects a volume in a robot in the following order:

- NetBackup searches the 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 that the backup schedule requires. However, if the NetBackup **Media** host property **Allow multiple retentions per media** is specified for the server, NetBackup does not search by retention level.
    - In the volume pool that the backup job requires.
    - Not in a FULL, FROZEN, IMPORTED, or SUSPENDED state.
    - Of the same density that the backup job requested, and in the robot that that the backup job requested.
    - Not currently in use by another backup or a restore.
    - Not written in a protected format. NetBackup detects tape format after the volume is mounted. If the volume is in a protected format, NetBackup unmounts the volume and resumes the search.
- If a suitable volume is found, NetBackup uses it.
- If NetBackup cannot find a mounted volume that satisfies all of the previous conditions, it checks the media catalog for any volume that is suitable.
    - If a suitable volume is in a robot, NetBackup issues the commands that do the following: Move the volume to a drive, position the heads to the

beginning of the volume, and assign it to the request. No manual intervention is required.

- If a suitable volume is not in a robot but is in a standalone drive, NetBackup automatically mounts and assigns it. No manual intervention is required.
- If a suitable volume is not in a robot or a standalone drive and the request is media-specific, NetBackup may pend a mount request. A media-specific mount request is one for a restore, for an import, or from the `tpreq` command.
- If a suitable volume is not in a robot or a standalone drive, NetBackup may attempt to use another volume only as follows: For backup jobs for which any other media can be used.
- If a suitable volume does not exist or if a suitable volume is at end of media (EOM), NetBackup assigns a new volume. NetBackup may assign a new volume even if a volume is not full (because NetBackup received an EOM message from the drive).

The new volume must meet 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 NetBackup)
- Has not exceeded the maximum number of mounts allowed
- If more than one volume qualifies, NetBackup chooses the volume that was least recently used.  
NetBackup then adds it to the media catalog and assigns it the specified retention level.
- If there are no unassigned volumes of the requested type, the backup terminates with an error message that no media were available.

See [“About spanning media with automatic media selection”](#) on page 144.

## About spanning media with automatic media selection

After an end of media (EOM) is reached, automatic media selection depends on whether NetBackup is configured to allow backups to span media, as follows:

- NetBackup spans media if the NetBackup **Media** host property **Allow backups to span media** is specified for the server.  
In this case, 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 the media **Allow backups to span media** property is not specified.  
In this case, the backup terminates abnormally and the operation is retried according to the NetBackup Global Attributes host property, **Schedule backup attempts**.

## About selecting media in standalone drives

The following topics explain media selection and other aspects of standalone drive operations:

See [“About selecting media by using standalone drive extensions”](#) on page 145.

See [“About disabling standalone drive extensions”](#) on page 146.

See [“About spanning media”](#) on page 146.

See [“About leaving standalone drives in the ready state”](#) on page 147.

### About selecting media by using standalone drive extensions

With NetBackup standalone drive extensions, NetBackup tries to use any labeled or any unlabeled media that is in a standalone drive. This capability is enabled by default during installation.

The media selection process is as follows:

- If a backup is requested and an appropriate standalone drive contains a volume, NetBackup tries to select and use that volume.
- If an appropriate drive does not contain a volume, NetBackup selects a volume.  
See [“About selecting media in robots”](#) on page 142.  
The Device Monitor shows the mount request, and an operator must manually insert the volume and assign it to a drive.

A volume that was used previously for backups must meet the following criteria:

- Not be FULL, FROZEN, or SUSPENDED
- Contain backups at the retention level and be in the same volume pool as the backup that requires a volume.

However, if the NetBackup **Media** host property **Allow multiple retentions per media** is specified for the server, NetBackup does not require a specific retention level.

NetBackup selects unlabeled media only if the existing volumes that meet the appropriate criteria do not have available space to contain the new backup images.

If the media is unlabeled, the following actions occur:

- NetBackup labels the media.
- NetBackup adds a media ID to the volume configuration, if necessary.  
If a media ID is added, the NetBackup Media ID prefix (non-robotic) is used as the first characters of the media ID.
- If a media ID prefix is not specified, the default prefix is the letter A. For example, A00000.
- NetBackup adds the requested volume pool to the volume configuration (if the backup policy specifies a volume pool).

If the unused media is unlabeled, label it by using the `bplabel` command. Specify the `-u` parameter to force assignment of a specific drive index, which eliminates the need to assign the drive manually.

## About disabling standalone drive extensions

Disable the standalone drive extensions by clearing the NetBackup media server host property, **Enable standalone drive extensions**. If this property is cleared, NetBackup uses the same method to select media for standalone drives as it uses for robotic drives.

## About spanning media

Media selection after an end of media (EOM) condition depends on whether NetBackup is configured to allow backups to span media, as follows:

- NetBackup spans media if the **Allow backups to span media** host property is specified for the server. NetBackup selects another volume to begin the next fragment, and the resulting backup has data fragments on more than one volume.  
After an EOM condition, NetBackup attempts to use an unassigned volume rather than one that already has images on it. NetBackup checks the EMM database for a volume that is the correct media type, in the correct volume pool, and so on.  
If a suitable unassigned volume is unavailable, NetBackup selects a volume.
- NetBackup does not span media if the **Allow backups to span media** host property is not specified. The backup terminates abnormally when the end of media is reached. The operation is rescheduled according to the master server host property **Schedule backup attempts**.

You can further configure NetBackup behavior for standalone drives. Normally, when NetBackup spans media and an EOM is encountered on a standalone drive, NetBackup searches for other media or generates a pending mount request. You can configure a wait period for standalone drives. The wait period is helpful when a gravity feed tape stacker takes a long time to load the next media in the drive.

To configure NetBackup to wait, specify the **Media request delay** media server host property. This property specifies the number of seconds NetBackup waits to use a volume that is loaded in a compatible drive. After the wait period expires, NetBackup searches for another drive. NetBackup also waits to generate a pending mount request during tape span operations. The **Media request delay** property applies only when standalone drive extensions are enabled.

### About leaving standalone drives in the ready state

To leave standalone drives in a ready condition after a backup or restore completes, use the `nbemmcmd` command to enable the `-do_not_eject_standalone` option. NetBackup does not eject the tape after an operation completes. The media is still ejected if EOM is reached or an error is encountered. Also, the media is ejected if the drive needs to be used with another media or the media needs to be used with another drive.

One standalone drive may be ready and contain suitable media.

Detailed information on the `nbemmcmd` command is available.

See *NetBackup Commands Reference Guide*.

## Volume pool and volume group examples

The following three examples show the relationship between volume pools and volume groups.

**Figure 4-2** shows an example of one volume pool (named `NB_pool`) and several volume groups.

You can move volumes between the groups in the robotic library and any groups that are off site. All volumes, however, remain in the same pool.

Media in the same volume pools are in different volume groups. Note that the data is stored on separate volumes by assigning different volume pools. The volumes in a pool can be in more than one physical location and in more than one volume group.

**Figure 4-2** Volume pool with multiple volume groups

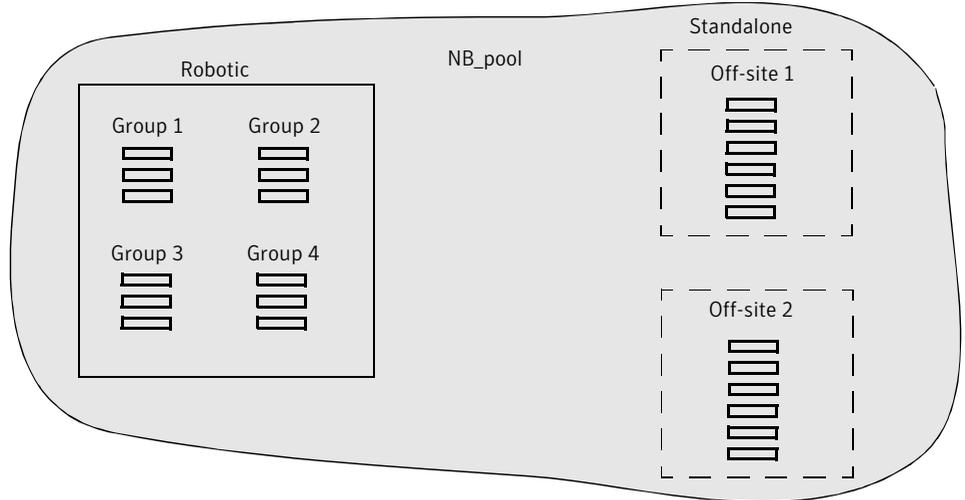
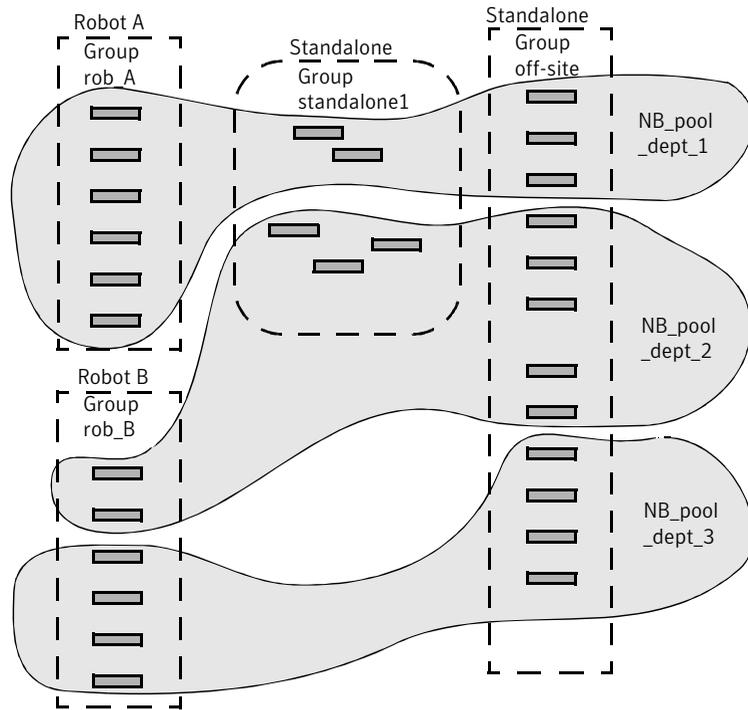


Figure 4-3 shows how the volumes in the pool NB\_pool\_dept\_1 are spread among the rob\_A, standalone1, and off-site volume groups.

These groups also have volumes from more than one pool (though the volumes in each group must all be the same type). You also can configure a scratch pool from which NetBackup can transfer volumes when a volume pool has no media available.

**Figure 4-3** Volume groups with multiple volume pools

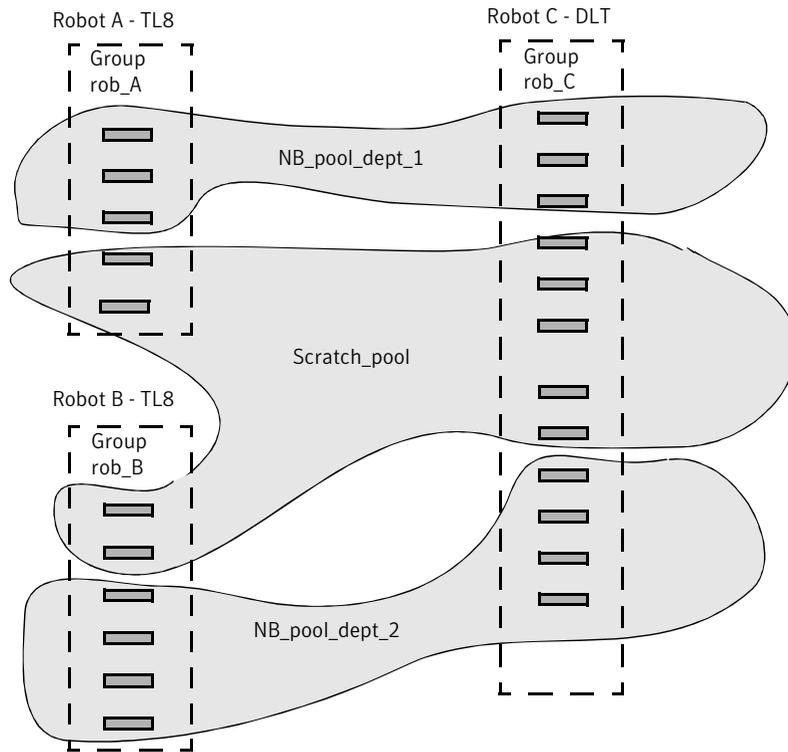


In [Figure 4-4](#), the scratch pool is named `Scratch_pool`. The three robots contain volumes from that pool in addition to those from other pools.

Assume the following sequence of events:

- A backup job requires a DLT volume, so NetBackup attempts to assign one from `NB_pool_dept_1` in Robot C.
- Robot C has no unassigned volumes available in the `NB_pool_dept_1` pool.
- NetBackup searches the scratch pool for an unassigned DLT volume in Robot C. If a volume is available, NetBackup moves it to `NB_pool_dept_1`. Otherwise, NetBackup logs a `media unavailable` status.

**Figure 4-4** Scratch pool example



## Media formats

NetBackup writes media in a format that allows the position to be verified before NetBackup appends new backups.

[Table 4-11](#) shows the symbols that are used in the media format descriptions.

**Table 4-11** Media format symbols

| Symbol              | Description                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------|
| MH                  | Media header (1024 bytes).                                                                             |
| *                   | Tape mark.                                                                                             |
| BH                  | Backup header (1024 bytes).                                                                            |
| BH1 ... BH <i>n</i> | Backup headers (1024 bytes). One for each job that is part of the set of the jobs that are multiplexed |

**Table 4-11** Media format symbols (*continued*)

| Symbol | Description                                                 |
|--------|-------------------------------------------------------------|
| Image  | Data from the backup.                                       |
| EH     | Empty backup header, which is used for position validation. |

[Table 4-12](#) provides more information about how the media formats are used in different situations.

**Table 4-12** Media format descriptions

| Format                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standard tape format     | <p>For all tape media except quarter-inch cartridge (QIC) and WORM, the format for the backups that are not multiplexed is as follows:</p> <p>MH * BH Image * BH Image * BH Image * EH *</p> <p>When a new backup image is added, the tape is positioned to the EH and the 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.</p> <p>When NetBackup encounters the end of media during a write operation, it terminates the tape with two tape marks and does not write an EH.</p> |
| QIC and WORM tape format | <p>This format is used for quarter-inch cartridge (QIC) and WORM media. Unlike the standard tape format, NetBackup does not write empty backup headers (EH). The format is as follows:</p> <p>MH * BH Image * BH Image * BH Image *</p> <p>To append backup images to QIC media, NetBackup positions to the end of data (EOD) and then starts the next backup.</p>                                                                                                                                                                                                                 |

**Table 4-12** Media format descriptions (*continued*)

| Format                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fragmented backup format | <p>For fragmented backups, the media format is similar to the standard tape format. The difference is that NetBackup breaks the backup image into fragments of the size that are specified when the storage unit is configured.</p> <p>The following is an example:</p> <pre>MH * BH Image (frag 1)* BH Image (frag 2)* BH Image (frag n) * EH *</pre> <p>Fragmentation is intended primarily for storing large backup images on a disk type storage unit.</p> <p>For multiplexed backups, image fragmentation results in faster restores because NetBackup can advance to the specific fragment before it begins a search for the file.</p> <p><b>Note:</b> If an error occurs in a backup, the entire backup is discarded and the backup restarts from the beginning. It does not restart from the fragment where the error occurred. Exception: checkpoint and restart backups resume from the last checkpoint fragment.</p> |
| Multiplexing format      | <p>The tape format for multiplexed backups is as follows:</p> <pre>MH * BH1 ... BHn Image ...</pre> <p>By default, the data image is in 64-kilobyte blocks. Each block also contains 512 bytes that are reserved for multiplexing control information and to identify the backup to which the block corresponds.</p> <p>When a job ends or a new job is added to the multiplexing set, NetBackup writes a tape mark. NetBackup then starts multiplexing the revised set of jobs.</p> <p>The following is an example:</p> <pre>MH * BH1 BH2 BH3 Image* BH2 BH3 Image* BH2 BH3 BH4 Image. .</pre>                                                                                                                                                                                                                                                                                                                                 |

**Table 4-12** Media format descriptions (*continued*)

| Format               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Spanning tape format | <p>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 described for fragmented backups. The first fragment on the next tape begins with the buffer of data where the end of media occurred.</p> <p>The following is the first tape format (NetBackup does not write an EH and terminates the tape with two tape marks):</p> <p>MH * ... *BHn Image (frag 1) **</p> <p>The following is the second tape format:</p> <p>MH * BHn Image (frag2)* ... * EH *</p> |

## Media Manager commands

Detailed information about the Media Manager commands is available. These commands are located in `install_path\VERITAS\Volmgr\bin`.

See *NetBackup Commands Reference Guide* for detailed information about most of the commands that are in the following tables.

---

**Note:** Start and stop services by using the **Services** tool available in **Administrative Tools** in the Microsoft Windows Control Panel. If they are started from the command line, some services occupy that NetBackup Console session until they are stopped.

---

[Table 4-13](#) shows the Media Manager services and processes and the commands that start each.

[Table 4-14](#) lists commands and the devices and processes that each stops.

**Table 4-13** Starting services and processes

| Command | Description                                                                                                                                                          |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| acsd    | <p>The Automated Cartridge System robotic process. The Device Manager <code>ltid</code> starts this process.</p> <p>Applies only to NetBackup Enterprise Server.</p> |
| avrd    | The Automatic Volume Recognition process. The Device Manager <code>ltid</code> starts this process.                                                                  |

**Table 4-13** Starting services and processes (*continued*)

| Command | Description                                                                                                                                                                                           |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ltid    | Starts the NetBackup Device Manager service. Starting <code>ltid</code> also starts the robotic, robotic control, Media Manager volume, and automatic volume recognition daemons.                     |
| tl4d    | The tape library 4MM robotic process. The Device Manager <code>ltid</code> starts this process.                                                                                                       |
| tl8cd   | Starts the tape library 8MM robotic-control process. The Device Manager <code>ltid</code> starts this process.                                                                                        |
| tl8d    | The tape library 8MM robotic process. The Device Manager <code>ltid</code> starts this process.                                                                                                       |
| tldec   | Starts the tape library DLT robotic-control process. The Device Manager <code>ltid</code> starts this process.                                                                                        |
| tldec   | The tape library DLT robotic process. The Device Manager <code>ltid</code> starts this process.                                                                                                       |
| tlhcd   | Starts the tape library Half-inch robotic-control process. The Device Manager <code>ltid</code> starts this process.<br><br>Applies only to NetBackup Enterprise Server.                              |
| tlhd    | The tape library Half-inch robotic process. The Device Manager <code>ltid</code> starts this process.<br><br>Applies only to NetBackup Enterprise Server.                                             |
| tlmd    | The tape library Multimedia process. The Device Manager <code>ltid</code> starts this process.<br><br>Applies only to NetBackup Enterprise Server.                                                    |
| vmd     | The NetBackup Volume Manager service. The Device Manager <code>ltid</code> starts this process.                                                                                                       |
| vmgcd   | The NetBackup Status Collection service. The <code>nbemm</code> command starts this service on the same host as the EMM server if one or more NetBackup 5.x servers are present in the configuration. |

**Table 4-14** Stopping services and processes

| Command  | Description                                              |
|----------|----------------------------------------------------------|
| stopltid | Stops the device, robotic, and robotic-control services. |

**Table 4-14** Stopping services and processes (*continued*)

| Command                | Description                                                                                               |
|------------------------|-----------------------------------------------------------------------------------------------------------|
| <code>tl1dcd -t</code> | Stops the tape library DLT robotic-control process.                                                       |
| <code>tl18cd -t</code> | Stops the tape library 8MM robotic-control process.                                                       |
| <code>tl1hcd -t</code> | Stops the tape library Half-inch robotic-control process.<br>Applies only to NetBackup Enterprise Server. |



# UNIX reference topics

This chapter includes the following topics:

- [About exclude and include lists on UNIX clients](#)
- [Schedules for user backups or archives](#)

## About exclude and include lists on UNIX clients

On UNIX clients, create the exclude and include lists in the following files on the client:

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

---

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

---

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

---

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

---

The following types of files appear in an exclude list:

- \*.o files
- core files
- a.out files
- Files that begin or end with ~ (backups for editors)
- Files and directories under /tmp, /usr/tmp

- Man pages
- Software that you can restore from original installation tapes
- Automounted directories
- CD-ROM file systems
- NetBackup automatically excludes the following file system types:
  - `mntfs` (Solaris)
  - `proc` (all UNIX platforms)
  - `cdrom` (all UNIX platforms)
  - `cacheFs` (AIX, Solaris, SGI, UnixWare)

Check with users before any files are excluded from backups.

---

**Note:** Symantec 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.

---

## Syntax rules for exclude lists

The following syntax rules apply to exclude lists:

- Blank lines or lines that begin 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, precede the character 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 each bracket with a backslash as in

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

A backslash (\) acts as an escape character only when it precedes a special or a wildcard character. NetBackup normally interprets a backslash literally because a backslash is a legal character to use in pathnames.

- If all files are excluded in the backup selections list, NetBackup backs up only what is specified by full path names in the include list. Files can be excluded by using / or \* or by using both symbols together (/\*).
- Spaces are considered legal characters. Do not include extra spaces unless they are part of the file name.

For example, if you want to exclude a file named

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

and the 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 the directory path, enter the name without a preceding slash. For example:

`test`

rather than

`/test`

This is equivalent to prefixing the file pattern with a slash:

`/`

`/*/`

`/**/*.`

`/**/*./*`

and so on.

- Do not use patterns with links in the names. For example, assume `/home` is a link to `/usr/home` and `/home/doc` is in the exclude list. The file is still backed up in this case because the actual directory path, `/usr/home/doc`, does not match the exclude list entry, `/home/doc`.

## Example of an exclude list

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

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

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

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

## Exclude lists for specific policies or schedules

NetBackup lets you create an exclude list for a specific policy or for a policy and a schedule combination. Create an `exclude_list` file with a *policyname* or *policyname.schedulename* suffix. The following two file examples use a policy that is named `wkstations`. The policy contains a schedule that is named `fulls`:

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

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

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

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

NetBackup uses only:

```
exclude_list.wkstations.fulls
```

## About creating an include list on a UNIX client

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

---

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

---

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

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

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

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

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

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

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

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

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

## Schedules for user backups or archives

To have NetBackup use a specific policy and schedule for user backups or archives of a UNIX client, add the following options to the `/usr/opensv/NetBackup/bp.conf` file:

- `BARCHIVE_POLICY`
- `BARCHIVE_SCHED`
- `BBACKUP_POLICY`
- `BBACKUP_SCHED`

These options can also be added to a user's `$HOME/bp.conf` file on the client.

# Index

## Symbols

.ExTeNt.nnnn files 97  
.SeCuRiTy.nnnn files 97  
@@MaNgLeD.nnnn files 97  
@@MaNgLeD.nnnn\_Rename files 97  
@@MaNgLeD.nnnn\_Symlink files 97

## A

Access control  
  lists (ACLs) 97  
ACS or TLM robot types 55  
ACS\_vm.conf entry 74  
ACS\_SEL\_SOCKET  
  vm.conf entry 75  
ACS\_SSI\_HOSTNAME  
  vm.conf entry 75  
ACS\_SSI\_SOCKET  
  vm.conf entry 75  
ADJ\_LSM  
  vm.conf entry 75  
AIX cachefs file system 158  
All Log Entries report 100  
Allow backups to span media 145  
Alternate client restores  
  host.xlate file 94  
Announce DHCP interval property 42  
API\_BARCODE\_RULES  
  vm.conf entry 77  
Arbitrated Loop Physical Address (ALPA) 61  
AUTHORIZATION\_REQUIRED  
  vm.conf entry 77  
AUTO\_PATH\_CORRECTION  
  vm.conf entry 78  
AUTO\_UPDATE\_ROBOT  
  vm.conf entry 78  
AVRD\_PEND\_DELAY  
  vm.conf entry 79, 140  
AVRD\_SCAN\_DELAY  
  vm.conf entry 79

## B

Backup Exec 59  
backup\_exit\_notify script 105  
backup\_notify script 105  
Backups  
  backup\_exit\_notify script 105  
  backup\_notify script 105  
  bpend\_notify script  
    UNIX client 111  
    Windows client 113  
  bpstart\_notify script  
    UNIX client 106  
    Windows client 109  
  compressed 96  
  diskfull\_notify script 116  
  estimating time required 98  
  multiplexed 96  
  session\_notify script 120  
  session\_start\_notify script 120  
bpclient commands 46  
bpclntcmd utility 62  
bpdynamicclient 48  
bpend\_notify script  
  UNIX client 111  
  Windows client 113  
bpstart\_notify script  
  UNIX client 106  
  Windows client 109

## C

Capacity licensing  
  about 23  
  and multistreamed backups 35  
  nbdeployutil 24–26  
  reconciling report results 33  
  reporting 17, 27, 29–31  
cdrom file system 158  
CLEAN\_REQUEST\_TIMEOUT  
  vm.conf entry 79  
cleaning  
  automatic 129

- cleaning (*continued*)
  - frequency-based 130
  - library-based 129
  - TapeAlert reactive 125
  - times allowed 131
- CLIENT\_PORT\_WINDOW
  - vm.conf entry 80
- Clients
  - changing host names 93
  - dynamic UNIX client 47
  - exclude files list, UNIX 157
  - include files list 160
- cluster environments 139
- CLUSTER\_NAME
  - vm.conf entry 80
- Compressed backups 96
- CONNECT\_OPTIONS
  - vm.conf entry 80
- crawlreleasebyname
  - vmopr cmd option 137

**D**

- DAS\_CLIENT
  - vm.conf entry 81
- DataStore policy type 51
- DAYS\_TO\_KEEP\_LOGS
  - vm.conf entry 81
- Device
  - drivers 60
- device
  - configuration wizard 63
  - delays 99
  - using with other applications 122
- Device allocation host 55, 57
- Devices
  - configuration wizard 63
  - configuring 60
- devices
  - management practices 123
- DHCP server 41
- diskfull\_notify script 116
- Domain Name Service (DNS) hostnames 94
- drives
  - cleaning 129
  - cleaning manually 130
- Dynamic host name and IP addressing 41, 43–44, 46–47

**E**

- EMM\_REQUEST\_TIMEOUT
  - vm.conf entry 82
- EMM\_RETRY\_COUNT
  - vm.conf entry 82
- Enable performance data collection property 103
- ENABLE\_ROBOT\_AUTH
  - vm.conf entry 83
- Escape character on UNIX 159
- Exclude files list
  - UNIX 157
- Exclude lists
  - creating 157
  - example 159
  - files on UNIX 157
  - for specific policies and schedules 160
  - syntax rules 158
  - wildcards in 158
- Extended attribute files 96
- ExTeNt.nnnn files 97

**F**

- Files
  - @@MaNgLeD.nnnn\_Symlink 97
  - goodies scripts 103
- files
  - .ExTeNt.nnnn 97
  - .SeCuRiT.y.nnnn 97
  - @@MaNgLeD.nnnn 97
  - @@MaNgLeD.nnnn\_Rename 97
- Firmware levels 60–61
- FlashBackup 96
- frequency-based drive cleaning 130
- Front-End Terabyte (FETB) Calculation 24

**G**

- GNU tar 96
- Goodies directory 103

**H**

- Host names
  - changing client name 93
  - changing server name 92–93
  - client peername 93
  - correct use 91
  - short 93
- host.xlate file and alternate client restores 94
- HyperTerminal 61

**I**

Include files list 160  
 INVENTORY\_FILTER  
   vm.conf entry 82–83

**L**

library-based cleaning 129  
 License keys  
   for Shared Storage Option 53, 59  
 Licensing  
   about 13, 23  
   analyzing gathered data 16  
   nbdeployutil 13–14, 24–26  
   reconciling report results 18, 33  
   reporting 17, 27, 29–31

**M**

mail\_dr\_info.cmd 116  
 MAP\_CONTINUE\_TIMEOUT  
   vm.conf entry 84  
 MAP\_ID, vm.conf entry 83  
 Maximum concurrent drives for backup 64  
 Media  
   using tar to read images 96  
 media  
   best practices 123  
   formats 150  
   selection algorithm 142–143, 145  
   spanning 145–146  
 media and device management  
   best practices 122  
   performance and troubleshooting 124  
 Media Manager  
   best practices 122  
   configuration file 74  
   security 87  
 media\_deassign\_notify script 117  
 MEDIA\_ID\_BARCODE\_CHARS  
   vm.conf entry 84  
 MEDIA\_ID\_PREFIX  
   vm.conf entry 85  
 MM\_SERVER\_NAME  
   vm.conf entry 86  
 mntfs file system 158  
 Multiple servers 37  
 Multiplexing (MPX)  
   recovering backups 96

  multiplexing (MPX)  
     backups 152  
     tape format 152  
 Multistreamed backups 35

**N**

Named data streams  
   VxFS 96  
 nbdeployutil 13–14, 24, 26–27  
 nbemm 54  
 nbemm/DA  
   definition 54  
 nbmail.cmd 117  
 NDMP 96, 139  
 NetBackup Access Control (NBAC)  
   use of 83, 86  
 Network transfer rate 100  
 Notification scripts 103

**P**

Peername  
   client 93  
 pending\_request\_notify script 120  
 Performance Monitor  
   using with NetBackup 103  
 PREFERRED\_GROUP  
   vm.conf entry 86  
 PREVENT\_MEDIA\_REMOVAL  
   vm.conf entry 86  
 proc file system 158  
 PureDisk  
   exporting backups to NetBackup 50–51  
   reports 25  
   required licenses 50  
   restoring export data 52–53

**R**

RANDOM\_PORTS  
   vm.conf entry 87  
 Raw partitions 96  
 reactive cleaning 125  
 REQUIRED\_INTERFACE  
   vm.conf entry 87  
 RESERVATION\_CONFLICT status 136  
 restore\_notify script 120  
 Restores  
   restore\_notify script 120  
 robotic cleaning 129

## Robots

- sharing without SSO 58

**S**

- SAN media server 55, 59

- SAN Shared Storage Option (see SSO) 53

- Scan host 54, 56

## Scripts

- backup\_exit\_notify 104

- backup\_notify 104

- bpend\_notify 105

- bpstart\_notify 106, 109

- diskfull\_notify 104

- goodies 103

- mail\_dr\_info.cmd 104

- media\_deassign\_notify 104

- nbmail.cmd 104

- notification 103

- parent\_end\_notify 104

- parent\_start\_notify 104

- pending\_request\_notify 104

- restore\_notify 104

- session\_notify 104

- session\_start\_notify 104

- shared\_drive\_notify 57, 105, 121

- userreq\_notify 104

- SCSI persistent reserve 133

- SCSI reserve and release 133

- break a reservation 136–137

- error recovery 137

- limitations 139

- PEND status 137

- requirements 138

- RESERVATION CONFLICT 135–136

- SCSI-to-fibre

- bridges 61

- SeCuRiT<sub>y</sub>.nnnn files 97

- SERVER

- vm.conf entry 87

- Servers

- changing host names 92–93

- NetBackup

- multiple 37

- multiple media servers 38

- SAN media server 55

- session\_notify script 120

- session\_start\_notify script 120

- SGI cachefs file system 158

- shared drives

- definition 59

- Shared drives (see SSO) 53

- Shared library support 58

- Shared robots

- without SSO 58

- Shared Storage Option

- license key for 59

- Shared storage option

- key 53

- shared\_drive\_notify script 57

- Simple Mail Transfer Protocol 117

- Solaris

- extended attributes 96

- file systems 158

- spanning media 144–146, 153

- Scripts

- bpstart\_notify 105

- SSO

- definition 53

- device allocation host 55, 57

- Device Allocation Host Summary 71

- hardware requirements 53

- scan host 54, 56

- Shared Drive Summary 71

- supported SAN hardware 74

- terminology 59

- vm.conf entries 89

- SSO components configuration

- examples 54

- SSO\_DA\_REREGISTER\_INTERVAL

- vm.conf entry 88

- SSO\_DA\_RETRY\_TIMEOUT

- vm.conf entry 88

- SSO\_HOST\_NAME

- vm.conf entry 89

- standalone drive

- extensions

- disabling 146

- Storage area network (SAN) 53, 59–61

- supported

- SAN hardware 74

- Symantec Backup Exec 59

- System Monitor

- using with NetBackup 102–103

- System Monitor, using with NetBackup 102–103

**T**

- tape
  - spanning 145–146
- tape formats 151
- TapeAlert 124–125
  - log codes 126
- Tar
  - GNU 96
    - to read backup images 96
  - tested SAN components 74
  - TLH\_vm.conf entry 89
  - TLM\_vm.conf entry 89
- Traditional licensing
  - about 13
  - analyzing gathered data 16
  - nbdeployutil 13–14
  - reconciling report results 18
- Transfer rate 98–99

**U**

- UnixWare cachefs file system 158
- userreq\_notify script 121
- using devices with other applications 122

**V**

- VERBOSE, vm.conf entry 89
- veritas\_pbx port 80
- vm.conf file
  - ACS\_entries 74
  - ACS\_SEL\_SOCKET entries 75
  - ACS\_SSI\_HOSTNAME entries 75
  - ACS\_SSI\_SOCKET entries 75
  - ADJ\_LSM entries 75
  - API\_BARCODE\_RULES entries 77
  - AUTHORIZATION\_REQUIRED entries 77
  - AUTO\_PATH\_CORRECTION entries 78
  - AUTO\_UPDATE\_ROBOT entries 78
  - AVRD\_PEND\_DELAY entries 79
  - AVRD\_SCAN\_DELAY entries 79
  - CLEAN\_REQUEST\_TIMEOUT entries 79
  - CLIENT\_PORT\_WINDOW entries 80
  - CLUSTER\_NAME entry 80
  - CONNECT\_OPTIONS entries 80
  - DAS\_CLIENT entries 81
  - DAYS\_TO\_KEEP\_LOGS entries 81
  - ENABLE\_ROBOT\_AUTH entries 83
  - INVENTORY\_FILTER entries 82–83
  - MAP\_CONTINUE\_TIMEOUT entries 84

**vm.conf file (continued)**

- MAP\_ID entries 83
- MEDIA\_ID\_BARCODE\_CHARS entries 84
- MEDIA\_ID\_PREFIX entries 85
- MM\_SERVER\_NAME entry 86
- overview 74
- PREFERRED\_GROUP entries 86
- PREVENT\_MEDIA\_REMOVAL entries 86
- RANDOM\_PORTS entries 87
- REQUIRED\_INTERFACE entry 87
- SERVER entries 87
- SSO\_DA\_REREGISTER\_INTERVAL entries 88
- SSO\_DA\_RETRY\_TIMEOUT entries 88
- SSO\_HOST\_NAME entries 89
- TLH\_entries 89
- TLM\_entries 89
- VERBOSE entries 89
- volume groups
  - examples 147
- volume pools
  - examples 147
- VxFS
  - extent attributes 97
  - named data streams 96

**W**

- Wildcard characters
  - in exclude lists 158
  - UNIX escape character 159
- Windows System Monitor, using with NetBackup 102
- wizard
  - shared drive configuration 63
- Wizards
  - device configuration 63