

NetBackup™ Administrator's Guide, Volume II

UNIX, Windows, and Linux

Release 10.3

VERITAS™

NetBackup Administrator's Guide, Volume II

Last updated: 2023-10-23

Legal Notice

Copyright © 2023 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas Alta, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas 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. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. Veritas Technologies LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	NetBackup licensing models and usage reporting	
		11
	About NetBackup licensing models	12
	Tools for creating and viewing licensing reports	13
	Setting the licensing type for scheduled reports	14
	How capacity licensing works	16
	How capacity licensing detects overlap when multiple policies protect data	17
	Backup policies and agents that support accurate licensing	17
	Requirements before gathering data for multiple primary servers in capacity licensing	19
	How NEVC licensing works	19
	Scheduling reports for capacity licensing	20
	Scheduling reports for traditional licensing	23
	Scheduling reports for NetBackup Enterprise Virtual Client (NEVC) licensing	26
	Other configuration for incremental reporting	28
	Troubleshooting failures for usage reporting and incremental reporting	30
	Manually generating licensing reports	30
	Creating and viewing the licensing report	37
	Report tab descriptions	38
	Reviewing a capacity licensing report	40
	Summary tab	41
	Client aliases and multiple IP addresses	42
	Itemization tab	42
	Clients backed up with multiple streams	45
	Data backed up by multiple policies	45
	BigData plug-ins for NetBackup	47
	NetBackup for Exchange agent	48
	NetBackup for NDMP agent	49
	NetBackup for Oracle server agent	50
	NetBackup for SQL Server agent	51
	NetBackup for VMware agent	53
	NetBackup for Hyper-V agent	54

NetBackup for RHV agent	55
NetBackup for Nutanix-AHV agent	56
NetBackup for SQLite Agent	56
NetBackup for MariaDB Agent	57
NetBackup for PostgreSQL Agent	57
NetBackup for MySQL Agent	58
NetBackup for SAP HANA	58
NetBackup for FlashBackup	59
NetBackup for Enterprise Vault	59
NetBackup for Lotus Notes	59
NetBackup for Sybase	60
NetBackup for SAP MaxDB	60
NetBackup for DB2	60
NetBackup for Informix	60
NetBackup for Microsoft SharePoint	61
NetBackup for SAP Oracle	61
NAS-Data-Protection policy	62
Cloud policy	62
NetBackup for Kubernetes	63
Windows file system backups	63
UNIX file system backups	64
Reconciling the capacity licensing report results	65
Locate full backups for clients	65
Review compressed image information	65
Eliminate redundant counting of clients	66
Determine the effect of multistreamed backups	66
Confirm the accuracy of any database backups	67
Locate full backups for snapshot images	67
Reviewing a traditional licensing report	67
Summary tab	68
Complete the Hosts tab	68
Update the Itemization tab	69
Resolve the NDMP tab	70
Update the Virtual Servers tab	70
Confirm the Drives tab	71
Final steps	71
Reviewing an NEVC licensing report	71
Summary tab	71
Hosts tab	72
Using the NetBackup Storage API to get the total backup size information	72

Chapter 2	Additional configuration	76
	About multiple NetBackup primary servers	76
	About multiple media servers with one primary server	77
	About direct I/O for backups on Windows	80
	About dynamic host name and IP addressing	82
	About setting up dynamic IP addresses and host names	83
	Configuring the NetBackup primary server	84
	bpclient commands that control client entries	87
	Configuring dynamic NetBackup clients	88
	About busy file processing on UNIX clients	91
	Configuring busy file processing on UNIX	92
	Modifying bp.conf to configure busy file processing on UNIX	93
	bp.conf file entries on UNIX	93
	How NetBackup creates and uses action files on UNIX	95
	About the logs directory on UNIX	97
	Recommended changes for modifying bpend_notify_busy on UNIX	98
	About specifying the locale of the NetBackup installation	98
	About the Shared Storage Option	100
	About Shared Storage Option components	100
	About reserving or releasing shared devices	104
	How to share robotic libraries without using the Shared Storage Option	105
	Shared Storage Option terms and concepts	105
	About the Shared Storage Option license	106
	About Shared Storage Option prerequisites	106
	About hardware configuration guidelines	107
	About installing and configuring drivers	108
	Verifying the connectivity	109
	About configuring the Shared Storage Option in NetBackup	109
	Verifying your Shared Storage Option configuration	111
	Device Monitor and Shared Storage Option	116
	Viewing SSO summary reports	117
	Operating system assistance	118
	Common configuration issues with Shared Storage Option	118
	Frequently asked questions about Shared Storage Option	120
	About the vm.conf configuration file	120
	ACS_mediatype entry in vm.conf	121
	ACS_SEL_SOCKET entry in vm.conf	121
	ACS_CSI_HOSTPORT entry in vm.conf (on UNIX)	121
	ACS_SSI_HOSTNAME entry in vm.conf	122
	ACS_SSI_INET_PORT entry in vm.conf (on UNIX)	122

ACS_SSI_SOCKET entry in vm.conf	123
ACS_TCP_RPCSERVICE / ACS_UDP_RPCSERVICE entry in vm.conf (on UNIX)	124
ADJ_LSM entry in vm.conf	124
API_BARCODE_RULES entry in vm.conf	125
AUTHORIZATION_REQUIRED entry in vm.conf in NetBackup versions 8.0 and earlier	126
AUTO_PATH_CORRECTION entry in vm.conf	127
AUTO_UPDATE_ROBOT entry in vm.conf	127
AVRD_PEND_DELAY entry in vm.conf	127
AVRD_SCAN_DELAY entry in vm.conf	128
CLEAN_REQUEST_TIMEOUT entry in vm.conf	128
CLIENT_PORT_WINDOW entry in vm.conf	129
CLUSTER_NAME entry in vm.conf	129
DAYS_TO_KEEP_LOGS entry in vm.conf	129
EMM_RETRY_COUNT entry in vm.conf	129
EMM_CONNECT_TIMEOUT entry in vm.conf	130
EMM_REQUEST_TIMEOUT entry in vm.conf	130
INVENTORY_FILTER entry in vm.conf	130
MAP_ID entry in vm.conf	131
MAP_CONTINUE_TIMEOUT entry in vm.conf	131
MEDIA_ID_BARCODE_CHARS entry in vm.conf	132
MEDIA_ID_PREFIX entry in vm.conf	133
MM_SERVER_NAME entry in vm.conf	133
RANDOM_PORTS entry in vm.conf	133
REQUIRED_INTERFACE entry in vm.conf	134
SERVER entry in vm.conf in NetBackup versions 8.0 and earlier	134
SSO_DA_REREGISTER_INTERVAL entry in vm.conf	135
SSO_DA_RETRY_TIMEOUT entry in vm.conf	135
SSO_HOST_NAME entry in vm.conf	136
VERBOSE entry in vm.conf	136
Example vm.conf file	136
How to access media and devices on other hosts	136
Host name precedence in the vm.conf file	137

Chapter 3	Holds Management	139
	About Holds Management	139
	Creating a hold	140
	Viewing hold details	140
	Adding a backup image to an existing hold	141
	Releasing a hold	141

Chapter 4	Menu user interfaces on UNIX	143
	About menu user interfaces	143
	About the tpconfig device configuration utility	144
	About the tpconfig utility menu	145
	Starting the tpconfig device configuration utility	146
	Adding robots	147
	Adding drives	147
	Updating a robot configuration	148
	Updating a drive configuration	149
	Deleting a robot	149
	Deleting a drive	150
	Configuring drive paths	150
	Configuring host credentials	151
	Displaying and writing the device configuration	151
	About the NetBackup Disk Configuration Utility	152
	Managing OpenStorage servers and disk pools	152
	Managing global disk attributes	153
Chapter 5	Reference topics	155
	Host name rules	156
	How NetBackup uses host names	156
	Updating NetBackup after changing the host name	159
	Special considerations for Name Service	160
	About reading backup images with nbtar or tar32.exe	162
	Restoring files with non-NetBackup restore utilities (on UNIX)	163
	Considerations for file restoration with non-NetBackup restore utilities (on UNIX)	164
	About the files that restores generate	165
	Factors that affect backup time	165
	Total amount of data to back up	166
	Transfer rate	166
	Methods for determining the NetBackup transfer rate	167
	NetBackup notify scripts	169
	backup_notify script	170
	backup_exit_notify script	171
	bpstart_notify script (UNIX clients)	172
	bpstart_notify.bat script (Windows clients)	174
	bpend_notify script (UNIX clients)	177
	bpend_notify.bat script (Windows clients)	179
	bpend_notify_busy script (UNIX clients)	181
	child_end_deployment_notify	181

child_start_deployment_notify	183
diskfull_notify script	184
drive_mount_notify script (on UNIX)	185
drive_unmount_notify script (on UNIX)	186
mail_dr_info script	186
media_deassign_notify script	187
nbmail.cmd script (on Windows)	187
parent_end_deployment_notify	188
parent_end_notify script	189
parent_start_deployment_notify	190
parent_start_notify script	191
pending_request_notify script	192
restore_notify script	192
session_notify script	193
session_start_notify script	193
shared_drive_notify script	193
userreq_notify script	194
Media and device management best practices	195
Media management best practices	196
Device management best practices	197
Media and device performance and troubleshooting	197
About TapeAlert	198
About TapeAlert cleaning (reactive cleaning)	198
About TapeAlert and frequency-based cleaning	199
About TapeAlert requirements	199
TapeAlert logs and codes	199
About tape drive cleaning	202
About library-based cleaning	203
About frequency-based cleaning	203
About operator-initiated cleaning	204
About using a cleaning tape	204
How NetBackup selects drives	205
How NetBackup reserves drives	206
About SCSI persistent reserve	207
About the SPC-2 SCSI reserve process	208
About SCSI reserve requirements	212
About SCSI reserve limitations	212
About SCSI reservation logging	212
About SCSI reserve operating system limitations on Windows	213
About checking for data loss	213
About checking for tape and driver configuration errors	214
About configuring SCSI reserve	214

How NetBackup selects media	215
About selecting media in robots	215
About selecting media in standalone drives	217
Volume pool and volume group examples	219
Media formats	222
Media and device management processes	225
About Tape I/O commands on UNIX	226
About requesting tapes	226
About reading and writing tape files	227
About removing tape files	227
Index	229

NetBackup licensing models and usage reporting

This chapter includes the following topics:

- [About NetBackup licensing models](#)
- [Tools for creating and viewing licensing reports](#)
- [Setting the licensing type for scheduled reports](#)
- [How capacity licensing works](#)
- [How NEVC licensing works](#)
- [Scheduling reports for capacity licensing](#)
- [Scheduling reports for traditional licensing](#)
- [Scheduling reports for NetBackup Enterprise Virtual Client \(NEVC\) licensing](#)
- [Other configuration for incremental reporting](#)
- [Troubleshooting failures for usage reporting and incremental reporting](#)
- [Manually generating licensing reports](#)
- [Creating and viewing the licensing report](#)
- [Reviewing a capacity licensing report](#)
- [Reconciling the capacity licensing report results](#)
- [Reviewing a traditional licensing report](#)

- [Reviewing an NEVC licensing report](#)
- [Using the NetBackup Storage API to get the total backup size information](#)

About NetBackup licensing models

Table 1-1 describes the NetBackup licensing models.

Table 1-1 NetBackup licensing models

Licensing model	Description
Capacity licensing	Capacity licensing is based on the total amount of data that NetBackup protects on the client or agent. The usage reporting tool securely communicates with the primary server to gather the protected data sizes and generate reports. The report includes details for the last 90 days per the license agreement and only includes details for full backups and user-directed backups (including expired backups). The capacity licensing report includes details about the mechanism that is used to calculate data size, based on the policy type. When this model is used, NetBackup automatically gathers the information through the accurate licensing method or obtains information from the backup image header.
Traditional licensing	The traditional licensing model is based on the total number of protected clients in a NetBackup environment or on the total storage capacity. This model counts the number of clients and servers, and then compares this information against licensed options. The report includes details for the last 90 days per the license agreement. Traditional usage reporting supports a single primary server.
NetBackup Enterprise Virtual Client (NEVC)	The NetBackup Enterprise Virtual Client (NEVC) licensing model is based on the total number of CPU sockets of a hypervisor. A hypervisor whose virtual machines NetBackup protects is considered for the measurement of CPU sockets. If you have a cluster of hypervisors, CPU sockets of the hypervisors that belong to a cluster are measured. Workloads that NEVC supports: ■ VMware ■ Hyper-V ■ Red Hat Virtualization (RHV) ■ Nutanix-AHV ■ Azure Stack ■ OpenStack

Tools for creating and viewing licensing reports

NetBackup provides the following tools to gather backup data information and create licensing reports.

Usage reporting widget in the NetBackup web UI

NetBackup 8.1.2 and later automatically collects and includes **Usage reporting** information in the NetBackup web UI. The usage report displays trends and details on the NetBackup data that is protected including data size, primary server, license type, and policy type.

For more information, see the *NetBackup Web UI Administrator's Guide*.

Veritas NetInsights Console

Veritas NetInsights Console is a new SaaS-based unified platform with Veritas products and features. It helps you to manage your usage and your license entitlements as well as leverages product telemetry and support data to offer software and appliance insights.

The NetInsights Console delivers a cohesive experience and eliminates the need to switch between multiple products.

To connect to Veritas NetInsights Console, use the following URL:

<https://netinsights.veritas.com>

Veritas Usage Insights

Usage information is used with Veritas Usage Insights that is part of the Veritas NetInsights Console to compare backup size data with your capacity, traditional, and NEVC licensing.

NetBackup supports the following licensing types:

- NetBackup Platform Base Complete Edition*
- NetBackup Platform Base Limited Edition*
- NetBackup Platform Base Big Data Workload Edition*
- NetBackup Platform Base NDMP Edition*
- NetBackup Platform Base Complete Edition with Flexible Licensing*
- NetBackup Enterprise Virtual Client Edition
- NetBackup Traditional Licensing Model Edition
- Alta Data Protection for NetBackup Cloud

Note: License types marked with an asterisk (*) use capacity licensing.

The usage reporting tool

NetBackup Deployment Insights is a command-line tool that can gather protected data size and create a licensing report. The `netbackup_deployment_insights` command generates a Microsoft Excel spreadsheet to review. In NetBackup 8.1.2 and later, the tool automatically gathers backup size data.

Note: Starting from 10.3, the Microsoft Excel spreadsheet report extension is changed from xls to xlsx.

It is recommended to use the NetBackup Deployment Insights tool when you manually generate the license reports. You can continue using the `nbdeployutil` command with the same options. You can then use this utility to:

- Run traditional and NEVC licensing reports on an automated schedule.
- Run capacity licensing reports on an incremental schedule.
- Manually generate reports for traditional or for capacity licensing.
- Create custom capacity license reports. For example, for a set of clients or for a business unit.

Setting the licensing type for scheduled reports

After install or upgrade, or if you change the license type, you must set the license type for a scheduled run of the `nbdeployutil` utility. Use the NetBackup web UI to set a license type. You must set the license type so that the `nbdeployutil` can generate a JSON file that NetBackup telemetry uses to send usage information to Usage Insights.

Note: You can no longer use the `bpsetconfig` command to set the licensing type.

For more information about setting a license type from the NetBackup web UI, refer to the NetBackup Web UI Administrator's Guide.

The following license types are available:

- NETBACKUP_PLATFORM_BASE_COMPLETE_EDITION
- NETBACKUP_PLATFORM_BASE_LIMITED_EDITION
- NETBACKUP_PLATFORM_BASE_BIG_DATA_WORKLOAD_EDITION

- `NETBACKUP_PLATFORM_BASE_NDMP_EDITION`
- `NETBACKUP_PLATFORM_BASE_COMPLETE_EDITION_FLEX`
- `NETBACKUP_ENTERPRISE_VIRTUAL_CLIENT_EDITION`
- `NETBACKUP_TRADITIONAL_LICENSING_MODEL_EDITION`

For NetBackup 8.3 and later, you have to set the licensing type using **NETBACKUP_PLATFORM_BASE_COMPLETE_EDITION_FLEX** to avail the benefit of Flexible licensing. For NetBackup 8.2 and earlier, Flexible licensing is available when you set Complete Edition as the license type using **NETBACKUP_PLATFORM_BASE_COMPLETE_EDITION**.

If you do not set a licensing type, the following type is selected:
NETBACKUP_PLATFORM_BASE_COMPLETE_EDITION

For more information about Flexible licensing, refer to the NetBackup Licensing Guide.

To set NetBackup Platform Base Complete Edition Flex as the license type

- 1 Open a web browser and go to the following URL:
- 2 `https://primaryserver/webui/login`
The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.
- 3 Enter your credentials and click **Sign in**.
- 4 On the left, click **Usage**.
- 5 Click **Usage reporting settings**.
- 6 The default license type or the license types that you set are displayed.
- 7 Click **Edit** and select **NetBackup Platform Base Complete Edition Flex**.
- 8 Click **Save**.

Note: In a multi-primary scenario, it is recommended that the proxy servers always use the NetBackup 8.3 version to get the maximum benefit of flexible licensing.

If you have NetBackup 8.3 and earlier clients, ensure that you install the client-side Emergency Engineering Binary (EEB) from the Veritas Support site to get the enhanced benefit of flexible licensing. You can also see the Veritas Operations Readiness Tools (SORT) website and refer to the NetBackup Emergency Engineering Binary Guide.

How capacity licensing works

NetBackup automatically uses this data-gathering method if the policy type supports accurate licensing. 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. One FETB is 1 TB 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 data is gathered and then stored in the NetBackup database. If multiple backup copies are created, only the first copy is considered.

The `nbdeployutil` utility uses accurate licensing or image headers in the NetBackup catalog to determine the terabytes of data that NetBackup protects. Any partial terabyte of data is rounded up to the next whole terabyte. The final total is the sum of the FETBs for each client or each policy combination that the analyzer examines. The utility measures the actual data protected.

The accurate licensing method provides faster performance of the `nbdeployutil` tool to gather and generate capacity licensing report.

For backup image header, if a policy type is not yet supported for the accurate licensing method, this method is used. This method is also used for NetBackup 8.0 and earlier clients.

To report on protected data with accurate licensing, a NetBackup host or client must have a valid certificate to securely connect with the primary server so that `nbdeployutil` can successfully generate a report.

For information on security certificates, see the [NetBackup Security and Encryption Guide](#).

Note: In a multi-primary scenario, only capacity based license types are supported.

The following factors affect capacity licensing:

- Multiple policies of the same type that protect the same data
- The agent that is used to perform the backup
 - See “[BigData plug-ins for NetBackup](#)” on page 47.
 - See “[NetBackup for Exchange agent](#)” on page 48.
 - See “[NetBackup for Oracle server agent](#)” on page 50.
 - See “[NetBackup for NDMP agent](#)” on page 49.
 - See “[NetBackup for SQL Server agent](#)” on page 51.
 - See “[NetBackup for VMware agent](#)” on page 53.

See [“Windows file system backups”](#) on page 63.
 See [“UNIX file system backups”](#) on page 64.
 See [“NetBackup for RHV agent”](#) on page 55.
 See [“NetBackup for Hyper-V agent”](#) on page 54.
 See [“NetBackup for Nutanix-AHV agent”](#) on page 56.
 See [“Cloud policy”](#) on page 62.
 See [“NAS-Data-Protection policy”](#) on page 62.
 See [“NetBackup for Kubernetes”](#) on page 63.
 See [“NetBackup for SQLite Agent”](#) on page 56.
 See [“NetBackup for MariaDB Agent”](#) on page 57.
 See [“NetBackup for PostgreSQL Agent”](#) on page 57.
 See [“NetBackup for MySQL Agent”](#) on page 58.

How capacity licensing detects overlap when multiple policies protect data

Users are charged based on the protected data size. If multiple backup copies are created using a single policy, only the first copy is considered for calculation. For example, if a copy created and stored on disk. The other copies are not charged to the user.

If multiple policies of the same type protect the same data, the capacity report identifies the overlap. The user needs to calculate the actual data overlap and then configure the policies considering this overlap. The user can then change the report based on the overlap. The overlap is only applicable to data that is reported with accurate licensing.

Backup policies and agents that support accurate licensing

Table 1-2 Support for accurate licensing

Policy type	NetBackup client version start of support	Primary server version
BigData (for Hadoop HDFS and Nutanix Acropolis Hypervisor)	8.1	8.1
MS-Exchange-Server	8.0	8.1
MS-SQL-Server	8.0	8.1
MS-Windows	8.0	8.1

Table 1-2 Support for accurate licensing (*continued*)

Policy type	NetBackup client version start of support	Primary server version
NDMP	8.1	8.1
Oracle	8.1	8.1
Standard	8.0	8.1
VMware	8.1	8.1
Hypervisor (for RHV)	8.2	8.2
Hypervisor (for Nutanix-AHV)	8.3	8.3
NAS-Data-Protection	8.3	8.3
Cloud	8.3	8.3
Kubernetes	NA	9.1
DataStore (for MySQL)	10.1	10.1
DataStore (for SQLite)	10.1	10.1
DataStore (for MariaDB)	10.1	10.1
DataStore (for PostgreSQL)	10.1	10.1
SAP HANA	10.2	10.2
Lotus-Notes	10.2	10.2
Informix-On-BAR	10.2	10.2
Sybase	10.2	10.2
MS-Sharepoint	10.2	10.2
FlashBackup	10.2	10.2
SAP Oracle	10.2	10.2
SAP MaxDB	10.2	10.2
Enterprise-Vault	10.2	10.2
DB2	10.2	10.2
Hyper-V	10.2	10.2

Requirements before gathering data for multiple primary servers in capacity licensing

Before you run `nbdeployutil` in an environment that has multiple primary servers, complete the following configuration using the NetBackup Administration Console. These steps only need to be completed one time.

- Create a trust relationship between the primary servers.
See the [NetBackup Administrator's Guide, Volume I](#)
- For each remote primary server in the environment, add the primary server where you plan to run `nbdeployutil` to the **Additional Servers** list (in **Host Properties**).
 - Open **Host Properties > Master Servers**.
 - Right-click on the remote primary server and click **Properties**.
 - Select **Servers**.
 - On the **Additional Servers** tab, add the primary server.
See the [NetBackup Administrator's Guide, Volume I](#)

How NEVC licensing works

The NetBackup Enterprise Virtual Client (NEVC) model enables you to license NetBackup based on how many CPU sockets are used. When you set the license type as NEVC, `nbdeployutil` captures CPU socket usage information.

Scenarios to report CPU socket count:

- If the backup of a virtual machine is from a standalone host, the usage is the CPU socket count on the standalone host.
- If the backup of a virtual machine is from a host that is part of a cluster, the usage is the CPU socket count on all hosts that are part of the cluster.
- Virtual machine migration: If a virtual machine that is part of a host is migrated to another host, the usage is the CPU socket count on both hosts. The backup from both hosts must be completed in last 90 days.


Note: Multi-primary support is not applicable for NEVC licensing.

Scheduling reports for capacity licensing

By default, NetBackup triggers `nbdeployutil` to run on a specified schedule to incrementally gather data and to generate licensing reports. For the first run, the duration of the report uses the frequency that is specified in the configuration file.

For capacity licensing, the report duration is always for the last 90 days based on the availability of the gathered data. Any data older than 90 days is not considered in the report. Each time `nbdeployutil` runs, it gathers information for the time between the latest run of `nbdeployutil` and the previous successful run.

Figure 1-1 Generating incremental capacity licensing reports



Licensing report location

The current capacity licensing report resides in the following directory:

On Windows: `install_path\NetBackup\var\global\incremental`

On UNIX: `/usr/openv/var/global/incremental`

It contains the following files:

- The generated report for the latest `nbdeployutil` result.
- Folders containing incrementally gathered data.

- The archive folder that contains the older generated reports.
- `nbdeployutil` log files.

The older reports are placed in the archive folder. Veritas recommends that you retain at least 90 days of reporting data. Data can be kept longer than 90 days, depending on the requirements of your environment. Older reports can help to show how the capacity usage has changed over time. Delete the reports or the folder when they are no longer required.

Use Case I: Using the default values for the licensing report

The `nbdeployutilconfig.txt` file is not required when you use the default parameters. `nbdeployutil` uses the following default values for capacity licensing:

- `FREQUENCY_IN_DAYS=7`
- `MASTER_SERVERS=local_server`
- `PARENTDIR=folder_name`
For Windows: `install_path\NetBackup\var\global\incremental`
For UNIX: `/usr/opensv/var/global/incremental`
- `PURGE_INTERVAL=120` (number of days)
- `MACHINE_TYPE_REQUERY_INTERVAL = 90` (number of days)

Use Case II: Using custom values for the licensing report

If the file `nbdeployutilconfig.txt` is not present, create a file using the following format:

```
[NBDEPLOYUTIL_INCREMENTAL]
MASTER_SERVERS=<server_names>
FREQUENCY_IN_DAYS=7
PARENTDIR=<folder_name_with_path>
PURGE_INTERVAL=120
MACHINE_TYPE_REQUERY_INTERVAL=90
```

To use custom values for the licensing report

- 1 Copy the `nbdeployutilconfig.txt` file to the following location:
For Windows: `install_path\NetBackup\var\global`
For UNIX: `/usr/opensv/var/global`
- 2 Open the `nbdeployutilconfig.txt` file.

- 3 Edit the `FREQUENCY_IN_DAYS` value to reflect how often you want the report to be created.

Default 7
 (recommended)

Minimum 1

Parameter deleted `nbdeployutil` uses the default value.

Note: If you enter a value that is less than 1, `nbdeployutil` automatically uses 7, which is the default value.

- 4 Edit the `MASTER_SERVERS` value to include a comma-separated list of the primary servers you want to include in the report.

Note: Veritas Usage Insights requires that primary servers be at NetBackup 8.1.2 or later.

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

For example:

- `MASTER_SERVERS=newserver, oldserver`
- `MASTER_SERVERS=newserver, oldserver.domain.com`
- `MASTER_SERVERS=myserver1.somedomain.com, newserver.domain.com`

- 5 Edit the `PARENTDIR` value to include the full path for location where the data is gathered and reported.

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

- 6** Edit the `PURGE_INTERVAL` to indicate the interval (in days) for how often you want to delete the report data. Data that is older than 120 days is automatically purged.

Default	120
Minimum	90
No value	<code>nbdeployutil</code> uses the default value.
Parameter deleted	<code>nbdeployutil</code> uses the default value.

- 7** Edit the `MACHINE_TYPE_QUERY_INTERVAL` to indicate how often to scan physical clients for updates to the machine type.

Default	90
Minimum	1
No value	<code>nbdeployutil</code> uses the default value.
Parameter deleted	<code>nbdeployutil</code> uses the default value.

Scheduling reports for traditional licensing

By default, NetBackup triggers `nbdeployutil` to run on a specified schedule to incrementally gather data and to generate licensing reports. For the first run, the duration of the report uses the frequency that is specified in the configuration file.

For traditional licensing the report duration is always for the last 90 days.

Note: Multi-primary support is not applicable for traditional licensing.

Licensing report location

The current traditional licensing report resides in the following directory:

On Windows: `install_path\NetBackup\var\global\incremental\traditional`

On UNIX: `/usr/openv/var/global/incremental/traditional`

It contains the following files:

- The generated report for the latest `nbdeployutil` result.
- Folders containing incrementally gathered data.

- The archive folder that contains the older generated reports.
- `nbdeployutil` log files.

Use Case I: Using the default values for the licensing report

The `nbdeployutilconfig.txt` file is not required when you use the default parameters. `nbdeployutil` uses the following default values for traditional licensing:

- `FREQUENCY_IN_DAYS=30`
- `PARENTDIR=folder_name`
For Windows:
`Iinstall_path\NetBackup\var\global\incremental\traditional`
For UNIX: `/usr/opensv/var/global/incremental/traditional`
- `PURGE_INTERVAL=120` (number of days).
- `MACHINE_TYPE_REQUERY_INTERVAL = 90` (number of days)

Use Case II: Using custom values for the licensing report

If the file `nbdeployutilconfig.txt` is not present, create a file using the following format:

```
[NBDEPLOYUTIL_INCREMENTAL_TRADITIONAL]
FREQUENCY_IN_DAYS=30
```

Traditional licensing uses the `PARENTDIR` and `PURGE_INTERVAL` parameters of the `[NBDEPLOYUTIL_INCREMENTAL]` tag.

```
[NBDEPLOYUTIL_INCREMENTAL]
PARENTDIR=<folder_name_with_path>
PURGE_INTERVAL=120
MACHINE_TYPE_REQUERY_INTERVAL=90
```

To use custom values for the licensing report

- 1 Copy the `nbdeployutilconfig.txt` file to the following location:
For Windows: `install_path\NetBackup\var\global`
For UNIX: `/usr/opensv/var/global`
- 2 Open the `nbdeployutilconfig.txt` file.

- 3** Edit the `FREQUENCY_IN_DAYS` value to reflect how often you want the report to be created.

Default 30
(recommended)

Minimum 1

Value of 0 Disables the incremental reporting and no licensing information is captured.

Parameter deleted `nbdeployutil` uses the default value.

- 4** Edit the `PARENTDIR` value to include the full path for location where the data is gathered and reported.

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

- 5** Edit the `PURGE_INTERVAL` to indicate the interval (in days) for how often you want to delete the report data. Data that is older than 120 days is automatically purged.

Default 120

Minimum 90

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

- 6** Edit the `MACHINE_TYPE_QUERY_INTERVAL` to indicate how often to scan physical clients for updates to the machine type.

Default 90

Minimum 1

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

Scheduling reports for NetBackup Enterprise Virtual Client (NEVC) licensing

By default, NetBackup triggers `nbdeployutil` to run on a specified schedule to incrementally gather data and to generate licensing reports. For the first run, the duration of the report uses the frequency that is specified in the configuration file.

For NEVC licensing, the report duration is always for the last 90 days.

Licensing report location

The current NEVC licensing report resides in the following directory:

On Windows: `install_path\NetBackup\var\global\incremental\NEVC`

On UNIX: `/usr/opensv/var/global/incremental/NEVC`

It contains the following files:

- The generated report for the latest `nbdeployutil` result.
- Folders containing incrementally gathered data.
- The archive folder that contains the older generated reports.
- `nbdeployutil` log files.

Use Case I: Using the default values for the licensing report

The `nbdeployutilconfig.txt` file is not required when you use the default parameters. `nbdeployutil` uses the following default values for NEVC licensing:

- `FREQUENCY_IN_DAYS=30`
- `PARENTDIR=folder_name`
For Windows: `install_path\NetBackup\var\global\incremental\NEVC`
For UNIX: `/usr/opensv/var/global/incremental/NEVC`
- `PURGE_INTERVAL=120` (number of days).

Use Case II: Using custom values for the licensing report

If the file `nbdeployutilconfig.txt` is not present, create a file using the following format:

```
[NBDEPLOYUTIL_INCREMENTAL_NEVC]
FREQUENCY_IN_DAYS=30
```

NEVC licensing uses the `PARENTDIR` and `PURGE_INTERVAL` parameters of the `[NBDEPLOYUTIL_INCREMENTAL]` tag.

```
[NBDEPLOYUTIL_INCREMENTAL]
PARENTDIR=<folder_name_with_path>
PURGE_INTERVAL=120
```

To use custom values for the licensing report

- 1 Copy the `nbdeployutilconfig.txt` file to the following location:

For Windows: `install_path\NetBackup\var\global`

For UNIX: `/usr/opensv/var/global`

- 2 Open the `nbdeployutilconfig.txt` file.
- 3 Edit the `FREQUENCY_IN_DAYS` value to reflect how often you want the report to be created.

Default	30
(recommended)	

Minimum	1
---------	---

Value of 0	Disables the incremental reporting and no licensing information is captured.
------------	--

Parameter deleted `nbdeployutil` uses the default value.

- 4 Edit the `PARENTDIR` value to include the full path for location where the data is gathered and reported.

No value	<code>nbdeployutil</code> uses the default value.
----------	---

Parameter deleted `nbdeployutil` uses the default value.

- 5 Edit the `PURGE_INTERVAL` to indicate the interval (in days) for how often you want to delete the report data. Data that is older than 120 days is automatically purged.

Default	120
---------	-----

Minimum	90
---------	----

No value	<code>nbdeployutil</code> uses the default value.
----------	---

Parameter deleted `nbdeployutil` uses the default value.

Other configuration for incremental reporting

To change the directory of the gathered data and capacity licensing report

- 1 If you have older gathered data and licensing reports, copy the complete directory to the new location.
- 2 Edit `nbdeployutilconfig.txt` and change the location of the gathered data and licensing report in the `PARENTDIR=folder_name` field.

To use the data that was gathered previously to generate a capacity licensing report

- 1 Locate the folder that was generated for the gathered data after the previous run of `nbdeployutil` and copy it to the following location:

On Windows: `install_path\NetBackup\var\global\incremental`

On UNIX: `/usr/openv/var/global/incremental`

- 2 Create the `gather_end.json` file inside the copied folder and add the following text:

```
{"success":0}
```

The next incremental run considers the data inside the copied folder to generate a capacity licensing report.

Note: Delete any other gather folders inside the copied folder to avoid gaps for the period in which data is gathered. The missing data is automatically generated during the next incremental run.

To create a custom interval report using existing gathered data for capacity licensing

- ◆ To create a report for a time interval that is different than the default interval of 90 days, run the following command:

On Windows:

```
nbdeployutil.exe --capacity --incremental --report --inc-settings
"install_dir\netbackup\var\global\nbdeployutilconfig.txt"
--hoursago <custom-time-interval>
```

On UNIX:

```
nbdeployutil.exe --capacity --incremental --report --inc-settings
"/usr/openv/var/global/nbdeployutilconfig.txt"
--hoursago <custom-time-interval>
```

The number of hours specified in `--hoursago` must be fewer than the `purge-interval` that is specified in the `nbdeployutilconfig.txt` file.

You can also use `--start` or `--end` options in the in the `nbdeployutilconfig.txt` file.

```
--start="mm/dd/yyyy HH:MM:SS"
```

```
--end="mm/dd/yyyy HH:MM:SS"
```

If the latest gather operation fails to retrieve front-end data size (FEDS) data, the custom report fails because the required backup information is not available. Let the next scheduled incremental gather run successfully and then try to generate the custom report.

Note: `nbdeployutil` uses existing gathered data to generate the custom interval report. You are not required to use the `--gather` option.

To change the directory of the gathered data and traditional and NEVC licensing report

- ◆ Edit `nbdeployutilconfig.txt` and change the location of the gathered data and licensing report in the `PARENTDIR=folder_name` field.

Troubleshooting failures for usage reporting and incremental reporting

- For incremental runs of `nbdeployutil`, notifications are sent to the NetBackup web UI. The notification details include, status of the run, duration, start time, and end time.
- `nbdeployutil` fails to gather data and generate the report for your environment. Refer to the logs to understand when the task failed and the reason for the failure.
- `nbdeployutil` fails with a `bpimagelist` error with status 37 after you run the utility manually. Ensure that you added the primary servers to the additional servers list.
 See [“Requirements before gathering data for multiple primary servers in capacity licensing”](#) on page 19.
- The following error displays because of internal web service communication failures:
 Internal Web API error occurred for primary server `SERVER_NAME`. Run `nbdeployutil` again with the gather option on primary server `SERVER_NAME`.
- For VMware or NDMP, when the backup agent fails to post licensing information to the database, a status code 5930 or 26 displays in the Activity Monitor. For more information, see the [NetBackup Status Codes Reference Guide](#).
- `nbdeployutil` may fail with errors related to loading the Perl modules. In such a scenario, it is recommended to refer the Perl documentation related to the reported error.

You can use `netbackup_deployment_insights` with the same troubleshooting points.

Manually generating licensing reports

You can manually generate reports for either the capacity or the traditional reporting model. The tool is located in the following directory:

On Windows: `install_path\NetBackup\bin\admincmd\`

On UNIX: `/usr/opensv/netbackup/bin/admincmd/`

You cannot manually generate the NetBackup Enterprise Virtual Client Edition (NEVC) licensing report.

Generating reports based on the licensing model

It is recommended to run the `netbackup_deployment_insights` command. You can continue using the `nbdeployutil` command with the same options.

If you use capacity-based licensing, run the following commands:

```
netbackup_deployment_insights
```

- For gather: `netbackup_deployment_insights --gather --capacity`
- For reports: `netbackup_deployment_insights --report --capacity`

If you use traditional based licensing, run the following commands:

```
netbackup_deployment_insights
```

- For gather: `netbackup_deployment_insights --gather --traditional`
- For reports: `netbackup_deployment_insights --report --traditional`

The tool generates the report in multiple steps. Data is gathered in the first step, and then analyzed and presented.

Table 1-3 Steps to gather and report licensing information

Step	Action	Description
Step 1	Complete prerequisites.	If you want to gather data from one primary for other remote primary servers, ensure that you have granted the necessary access across all servers: See “Requirements before gathering data for multiple primary servers in capacity licensing” on page 19. Multiple primary servers are only supported in capacity licensing. For back-level primary servers, load the engineering binary that is associated with <code>netbackup_deployment_insights</code> onto all primary servers for which you want to gather information.
	Gather the backup size data	The tool uses the following options to gather data from one or more primary servers. <pre>netbackup_deployment_insights --gather [--output=DIRECTORY] <--capacity --traditional> [--hoursago=N] [--start="mm/dd/yyyy HH:MM:SS" [--end="mm/dd/yyyy HH:MM:SS"]] [--master=HOSTNAME[,...]] [--log=FILENAME] [--runtimestats] [--nolog] [--bpimagelist=OPTIONS] [--use-bpflist] [--apikey-file=<APIKEYFILE path>]</pre> You can continue using the <code>nbdeployutil</code> command with the same options.

Table 1-3 Steps to gather and report licensing information (*continued*)

Step	Action	Description
Step 2	Analyze the data and prepare the report	The <code>netbackup_deployment_insights</code> commands uses the following options to analyze the gathered data and prepare the report: <pre>netbackup_deployment_insights --report <--capacity --traditional> <directory> ... [--dirlist=FILENAME --parentdir=DIRECTORY] [--capacity] [--debug-inputs] [--log=FILENAME] [--day-boundary=TIME] [--runtimestats] [--nolog] [--overlap-details] [--anonymize --anonymize -anon-master=SOME_NAME]</pre> For a traditional report, run: <code>netbackup_deployment_insights --report --traditional</code> For a capacity report, run: <code>netbackup_deployment_insights --report --capacity</code> You can continue using the <code>nbdeployutil</code> command with the same options. Note: For NetBackup 8.0 and later, this step requires that you provide the web service credentials.
Step 3	Examine the report and make adjustments.	

Authentication using web service credentials

If you do not use `--apikey-file=<APIKEYFILE path>`, then you must manually enter the credential information.

See [the section called “Providing web service credentials for the capacity or traditional licensing reports”](#) on page 33.

If you use `--apikey-file=<APIKEYFILE path>`, use the `apikey` file for authentication.

See [the section called “Providing API Key as credentials”](#) on page 33.

If you have enforced multi-factor authentication or have configured multi-factor authentication, you must use the `apikey` file - `--apikey-file=<APIKEYFILE path>`

For more information, refer to the Multi-factor authentication chapter from the NetBackup Web UI Administrator's Guide.

Note: In a multi-primary server scenario, if you use the `apikey` file, ensure that you add API keys of all the primary servers.

Providing web service credentials for the capacity or traditional licensing reports

When you manually run the tool for the gather operation in capacity or traditional licensing, you must enter the following credential information for authentication of the NetBackup web management service:

Table 1-4 Required credential information

Option	Description
Domain Type	Enter a domain type value from NIS, NISPLUS, WINDOWS, vx, unixpwd, ldap. This value is case-sensitive.
Domain Name	Name of the domain that the primary server host belongs to. If the primary server does not belong to any domain, enter the name of the primary server.
User name	Name of the user that has administrator privileges. For a custom user, ensure that you select the View permission for the Licensing namespace within the NetBackup Web UI. Go to Security > RBAC . For more information about RBAC, refer to the <i>NetBackup Web UI Administrator's Guide</i> .
Password	Password of the same user that has administrator privileges. When you enter the password, characters are intentionally not displayed in the command line.

Note: In a multi-primary server scenario, you must enter the credentials for all the primary servers that are mentioned with the gather command.

Providing API Key as credentials

When you manually run the tool for gather operation, you can use the `apikey` file for authentication of the NetBackup web management service.

The `apikey` file must be in the following format:

- For a single primary server:

```
<server1>:<apikey1>
```

For example,

```
Primary.netbackup.server.com:
AwBfDO9xbd5RDuu-foifdxE_rRBjEIzbtfxdU2O6B0gP4H5rUbK-UTYfXXglji9ix6s
```

- For multiple primary servers:

```
<server1>:<apikey1>
<server2>:<apikey2>
.
.
.
<serverN>:<apikeyN>
```

For example,

```
PrimaryOne.netbackup.server.com:
A9FA7cuygCW6TKxj7xB0JXOQf5CVwUMXOaZoOfjUpbLtmYyaSpulS5G0-ANnHE4C7d

PrimaryTwo.netbackup.server.com:
A9ciKsdtjhfd60GGFOZK5425nXpFJigOXUfLtszA1IHLxh2hL-OV7NEiPGG01Ewe_eK

PrimaryThree.netbackup.server.com:
A4sJ5RlXnhuSIGCfpjvRUlg2WguiYKusbbVJnPgUqS_PdyCt4Dlw9nfXlo2wblppm
```

For more information, refer to the [Using access keys, API keys, and access codes](#) chapter from the *NetBackup Web UI Administrator's Guide*.

Here is an example of gathering information for the local primary server using the `apikey` file.

```
C:\Program
Files\Veritas\NetBackup\bin\admincmd>netbackup_deployment_insights
--gather --traditional
--apikey-file="C:\Users\Administrator\Documents\apikey.txt"
NetBackup Deployment Insights, version 10.3Beta1
Gathering license deployment information...
Discovered master server xxxyyyyzzz.host1.com
Output for xxxyyyyzzz.host1.com at: C:\Program
Files\Veritas\netbackup\var\global\reports\20230821_135252_
xxxxyyyyzzz.host1.com
Gather DONE
Execution time: 6 secs
To create a report for this master server, run the following:
netbackup_deployment_insights --report --traditional "C:\Program
Files\Veritas\netbackup\var\global\reports\20230821_135252_
xxxxyyyyzzz.host1.com "
```

Performance of the usage reporting tool

The performance of the tool is dependent on the system running it as well as the size of the NetBackup catalog. The `--gather` option only executes as quickly as the `bpimagelist` command can run on 90 days 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 tool's performance.

Depending on the environment, the tool can take from several seconds to several minutes to run the `--gather` or the `--report` options.

Scanning a large number of clients increases the execution time of the usage reporting tool run accordingly. The clients that are scanned in the current run (for which the machine type is collected) are not scanned again in the subsequent runs. This reduces the execution time of `netbackup_deployment_insights`.

Overlap detection for any data that NetBackup protects with multiple policies

For capacity licensing reports for any agents that support the accurate licensing method, use the overlap option to detect any data that NetBackup protects with multiple policies. To enable overlap detection, run the `netbackup_deployment_insights report` with the `--overlap-details` option.

You can continue using the `nbdeployutil` command with the same options.

Table 1-5 Description of overlap detection option

Command line option	Description
<code>--overlap-details</code>	This option searches for any duplicate backup selections in policies of the same type. The results are noted in the Duplicate Selections column of the report. Note: The backup selection data must contain only ASCII or English-only characters. This option displays the actual backup selection data in the report.

Using scripts to automate data gathering in NetBackup 8.0 or later

If you used scripts to automate data gathering before NetBackup 8.0, these scripts must be updated when you upgrade. The data gathering process requires that you enter the credentials for the NetBackup web service.

Example: Gathering information for the local primary server

In this example, the tool gathers information for the local primary server.

```
C:\Program Files\Veritas\NetBackup\bin\
admincmd>netbackup_deployment_insights --gather
NetBackup Deployment Insights, version 9.1
Gathering license deployment information...
    Discovered master server pnnbcs5b38v002

Enter credentials for Master Server(s):

Master Server:pnnbcs5b38v002
Domain Type (NIS, NISPLUS, WINDOWS, vx, unixpwd, ldap):WINDOWS
Domain Name   :pnnbcs5b38v002
User Name     :administrator
Password      :

Data gather is in progress. This process might take some time.
    Output for pnnbcs5b38v002 at:
C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_020925_pnnbcs5b38v002
Gather DONE
Execution time: 1 min 30 secs
To create a report for this master server, run one of the following:
    capacity   : netbackup_deployment_insights.exe --report --capacity
"C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_020925_pnnbcs5b38v002"
    traditional: netbackup_deployment_insights.exe --report --traditional
"C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_020925_pnnbcs5b38v002"
```

You can continue using the `nbdeployutil` command with the same options.

The tool 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: Gathering information for a remote primary server

```
netbackup_deployment_insights --gather --master=sidon.example.com
```

Creating and viewing the licensing report

After gathering data for the licensing report, use the `--report --traditional` or the `--capacity` option to generate the licensing report. You can create a report for any of the following:

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

Creating a report using data that is collected for the local primary server

The `--gather --capacity` command is run for primary server `pnnbcsm5b38v002`:

```
C:\Program Files\Veritas\NetBackup\bin\
admincmd>netbackup_deployment_insights --gather --capacity
NetBackup Deployment Insights, version 9.1
Gathering license deployment information...
    Discovered master server pnnbcsm5b38v002

Enter credentials for Master Server(s):

Master Server:pnnbcsm5b38v002
Domain Type (NIS, NISPLUS, WINDOWS, vx, unixpwd, ldap):WINDOWS
Domain Name   :pnnbcsm5b38v002
User Name     :administrator
Password      :

Data gather is in progress. This process might take some time.
    Output for pnnbcsm5b38v002 at:
C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_021553_pnnbcsm5b38v002
Gather DONE
Execution time: 36 secs
To create a report for this master server, run the following:
    netbackup_deployment_insights.exe --report --capacity
"C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_021553_pnnbcsm5b38v002"
```

To create a capacity report based on the data gathered, the tool tells you what command you need to run:

```
C:\Program Files\Veritas\NetBackup\bin\
admincmd>netbackup_deployment_insights --report --capacity
"C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_021553_pnnbcsm5b38v002"
```

```
NetBackup Deployment Insights, version 9.1
Analyzing license deployment ...
    Report created at:
C:\Program Files\Veritas\netbackup\var\global\reports\
20210319_021553_pnnbcsm5b38v002\
report-capacity-pnnbcsm5b38v002-20210319_021917.xlsx
Analysis DONE
```

```
Execution time: 0 sec
```

It is recommended to run the `netbackup_deployment_insights` command. You can continue using the `nbdeployutil` command with the same options.

The tool 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.

Creating a roll-up report for several primary servers

This example assumes that you have gathered the respective primary server's data in directories `primary1dir`, `primary2dir`, `primary3dir`. These directories all reside within a parent directory named `EMEA-domains`. The output (report and log file) is saved to the `EMEA-domains` directory.

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

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

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

Report tab descriptions

The tool examines the image headers in the NetBackup catalog to determine one of the following:

- For traditional licensing, `netbackup_deployment_insights` determines the servers and clients in the NetBackup environment.
- For capacity licensing, `netbackup_deployment_insights` determines the amount of data NetBackup protects. The way that the client policies and schedules are configured can affect the results.
- For NEVC licensing, `netbackup_deployment_insights` determines the total number of CPU sockets of a hypervisor. A hypervisor whose virtual machines NetBackup protects is considered for the measurement of CPU sockets.

You can continue using the `nbdeployutil` command with the same options.

Table 1-6 Report tab descriptions

Tab	Description	Report type
Summary	The contents of this tab differs for a traditional or a capacity report. ■ Traditional report: Shows the final details about primary 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. ■ Capacity report: 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.	Traditional, Capacity, and NEVC
Itemization	Displays a table similar to the line itemization table you may see in a 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 or policy combination.	Traditional and Capacity
Unused clients	Displays the names of clients that are registered with the primary server but are not backed up.	Capacity
Hosts	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.	Traditional and NEVC
NDMP	A list of computers that 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.	Traditional

Table 1-6 Report tab descriptions (*continued*)

Tab	Description	Report type
Virtual Servers	A list of the virtual servers or the virtual hosts that were detected in the environment.	Traditional
Drives	Details the type of drives and the host or the library where the drive resides. Lists the host names that are associated with each drive as well as information about virtual tape libraries, shared drives, and vaulted drives.	Traditional
Interpreting the Results	Explains how to examine the report and how to reconcile the information in the report with your actual environment.	Traditional and capacity

Reviewing a capacity licensing report

The following table describes the steps to examine the capacity license report results and how to reconcile the report with the actual NetBackup environment.

Table 1-7 Examining the capacity licensing report

Step	Description	Reference
1	Verify the information on which the report is based, including the primary server, clients, and policies.	See "Summary tab" on page 41.
2	Remove redundant data created from client aliases and multiple IP addresses.	See "Client aliases and multiple IP addresses" on page 42.
3	Examine the Itemization tab for flagged conditions in the Accuracy column.	See "Itemization tab" on page 42.
4	For multistreamed backup images, verify how the images are grouped and the size of the backups.	See "Clients backed up with multiple streams" on page 45.
5	Examine the details related to overlap detection.	See "Data backed up by multiple policies" on page 45.

Table 1-7 Examining the capacity licensing report (*continued*)

Step	Description	Reference
6	Review the details for specific policy types.	See “BigData plug-ins for NetBackup” on page 47. See “NetBackup for Exchange agent” on page 48. See “NetBackup for Oracle server agent” on page 50. See “NetBackup for NDMP agent” on page 49. See “NetBackup for SQL Server agent” on page 51. See “NetBackup for VMware agent” on page 53. See “Windows file system backups” on page 63. See “UNIX file system backups” on page 64. See “NAS-Data-Protection policy” on page 62. See “Cloud policy” on page 62.

Summary tab

The top of the **Summary** tab shows the basis for the report information. Examine the section marked **Analyzed** to verify the gathered data.

The **Analyzed** section displays the following information:

- The primary servers that are included in the report.
- The date range for catalog data.
- The number of clients and policies that are included 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 values. The analyzer gathers 90 days of catalog data for all clients by default.

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

For the agents that support accurate licensing the **Overlap** column summarizes the charged sizes for all overlapping policies. These are policies for which the **Overlap** keyword appears in the **Accuracy** column, per primary server. The overlap is calculated only for the same policy type. For example, if both an **MS-Windows** and a **MS-Exchange-Server** policy back up an Exchange database, accurate licensing does not consider this policy as an overlap.

26	Flagged Capacity Figures (TB)						
27	Master Server	Confirmed (TB)	Overlap	Database Estimation	Possible Overlap	Multistream Estimation	
28	master	0.00	48.00	64.00	16.00	64.00	
29	newmaster	0.00	70.00	70.00	0.00	70.00	
30	oldmaster	0.00	0.00	25.45	25.45	25.45	
31	upgradedmaster	0.00	90.50	105.70	15.20	105.70	
32	Total	0.00	208.50	265.15	56.65	265.15	

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.

The report's **Itemization** tab shows the calculated capacity for each client or policy combination. The report flags any conditions that have the potential to over count

or to under count capacity. These conditions are identified in the **Accuracy** and **Accuracy Comment** columns.

- **OK - Precise data is reported**

Data that is displayed in the **Charged Size** column is protected data for a policy. A user can verify that the data is precise by referring to the policy type.

See [“Eliminate redundant counting of clients”](#) on page 66.

Overlap scenarios:

Overlap is identified within the same policy type when the accurate licensing method is used. This means that if the same data is backed up more than once by different policies of the same type (within the same client or across clients in the same primary server), the overlap is identified.

Once the overlap is identified, the overlap size is displayed in the **Overlap Size (KB)** column. After identification of the overlap, **Charged Size(KB)** is updated after reducing the calculated overlap size. The **Accuracy** column displays **OK** in such cases. For a policy where the overlap was detected is deducted from charged size, the following message is displayed:

Overlap detected for the policy and deducted from the Charged Size.

- If identical policies of the same type exist, the policy with the largest backup size is charged to the user. The **Charged Size** column displays zero for one of the identical policies.
- If a policy is a subset of another policy (consumed policy), the **Charged Size** column displays zero for the consumed policy. The user is charged for the superset policy.

- **Overlap (NetBackup 8.0 or earlier clients)**

When the **Accuracy** column displays **Overlap** for a policy, it means that the overlap exists for the given policy. The overlap size is calculated for the policy and is displayed in the **Overlap Size** column, but is not deducted from charged size.

For example, for the **MS-Windows** policy on a NetBackup client 7.7.3, if the policy has **Compression** attribute enabled, the Overlap keyword is displayed in the **Accuracy** column. Compressed size is not the correct size (being compressed) and cannot be deducted from the charged size.

- **Database estimation - database size estimated using UBAK summation**

The size of databases that a NetBackup database agent protects cannot be determined with certainty. Third party components external to NetBackup (for example, 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, or 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.

- **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.

- **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.

- **Size unavailable – Only snapshot is present**

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

Note: For Kubernetes workloads, if the snapshots do not exist in a backup, the column condition appears as "Size unavailable".

- **Possible multistream backup detected**

The size of the clients that are protected by multistream backups is the total of all backup images that are created by all streams.

The report's Itemization tab shows Teradata job name in the **Client Name** and **Policy Name** columns in case of a Teradata policy type. The backup images are grouped using Teradata job name and Teradata job index. The charged size is equal to the total size of all backup images in that group.

Master Server	Client Name	Policy Name	Policy Type	Backup Image	Charged Size (KB)
master1	Job_4N-D4	Job_4N-D4	Teradata	client1_1643613836	11914038

Clients backed up with multiple streams

When a client is backed up by multiple streams, the client's size is equal to the total of all backup images that were 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.

Data backed up by multiple policies

The `nbdeployutil` utility can use overlap detection to ensure that you are not charged for backup select data that is backed up by multiple policies of the same type. For NetBackup 8.1 and later clients, the report includes the names of the duplicate backup selections. The adjustment for any overlapping data is reflected in the **Charged Size** column of the report.

For example, if you have the following backup policies for policy type Microsoft SQL Server:

Table 1-8 Policies

Policy1	Policy2
Backup Selection:	Backup Selection:
■ "MSSQLSERVER"/"DB1"/ [Size 1.4GB] ■ "MSSQLSERVER"/"DB2"/ [Size 1.4GB]	■ "MSSQLSERVER"/"DB2"/ [Size 1.4GB] ■ "MSSQLSERVER"/"DB3"/ [Size 1.4GB]

Here MSSQLSERVER is the name of instance and DB1, DB2, DB3 are the databases. The DB file DB2 is common in both policies Policy1 and Policy2. `nbdeployutil` detects the overlap and displays in the report.

Table 1-9 Sample report columns

Policy Name	Accuracy	Accuracy Comment	Overlap Size (GB)	Duplicate Selection	Total (Readable) Size (GB)	Charged Size (GB)
Policy1	OK		0		2.8	2.8

Table 1-9 Sample report columns (*continued*)

Policy Name	Accuracy	Accuracy Comment	Overlap Size (GB)	Duplicate Selection	Total (Readable) Size (GB)	Charged Size (GB)
Policy2	OK	Overlap detected for the policy and deducted from the Charged Size	1.4	"MSSQLSERVER/" <PLACE_HOLDER> "DB2"/	2.8	1.4

In case of Availability Group (AG) the <PLACE_HOLDER> contains the AG name.

For a NetBackup client with version 8.0 or earlier, the following report is displayed policy that uses the backup image header method to gather information:

Table 1-10 Policies

Policy1: DB_File_Overlap1	Policy2: DB_File_Overlap2
Backup Selection: - Instance=INST_R DB file=DB_01 [Size 1.4GB] - Instance=INST_R DB file=DB_02 [Size 1.4GB]	Backup Selection: - Instance=INST_R DB file=DB_02 [Size 1.4GB] - Instance=INST_R DB file=DB_03 [Size 1.4GB]

In this example, the DB file DB_02 is common in both policies DB_File_Overlap1 and DB_File_Overlap2. nbdeployutil detects the overlap and displays in the report. An extract of the report specific to this example is as follows:

Table 1-11 Sample report columns

Policy Name	Accuracy	Accuracy Comment	Overlap Size (GB)	Total (Readable) Size (GB)	Charged Size (GB)
DB_File_Overlap1	Database Estimation, Multistream Estimation, OK	DB size estimated via backup summation, Multi-stream backup estimation	0	2.8	2.8

Table 1-11 Sample report columns (*continued*)

Policy Name	Accuracy	Accuracy Comment	Overlap Size (GB)	Total (Readable) Size (GB)	Charged Size (GB)
DB_File_Overlap2	Database Estimation, Multistream Estimation, OK	DB size estimated via backup summation, Multi-stream backup estimation, Overlap detected for the policy and deducted from the Charged Size	1.4	2.8	1.4

For policy DB_File_Overlap2 the customer is charged for only DB file DB_03.

BigData plug-ins for NetBackup

Capacity licensing for BigData plug-ins is specific to the **BigData** policy type for the following plug-ins:

- Hadoop plug-in for NetBackup
- Nutanix plug-in for NetBackup

After the `nbdeployutil` utility runs using the capacity licensing option, the report displays the policy type in the **Itemization** sheet as follows:

- **BigData:hadoop**
- **BigData:Nutanix-AHV**

Hadoop plug-in for NetBackup

Using the **BigData** policy, the Hadoop (HDFS) data is backed up for the directory or the backup selection that is defined in the policy. The definition of protected data for the policy is the size of the defined directories that have the `allow snapshot` option enabled on an HDFS file system. To view the defined directory size, you can browse the file system using the Hadoop web console.

Administrators can also run the following HDFS command to verify the size that capacity licensing reports. The following HDFS commands are available by default:

```
hdfs dfs -ls -R -h /<name_of_the_directory>
```

Nutanix plug-in for NetBackup

The front-end data size that is reported for Nutanix Acropolis Hypervisor is the consumed storage size. You can verify the consumed storage size on the Nutanix AHV console by navigating to **VM > Overview > Tables** under the **Storage** column. The Storage column displays the consumed storage size as compared to the total allocated storage size.

The Nutanix-AHV VM can be backed up using the BigData policy and the Hypervisor policy (all drives Included). If the same VM is backed up using both BigData and Hypervisor policies, you are charged once for the VM backup using the Hypervisor policy.

NetBackup for Exchange agent

Licensing for Exchange is specific to an **MS-Exchange-Server** policy and does not support an Exchange server that is backed up using any virtualization policy, such as **VMware**. Administrators can use any of the following ways to verify if the Exchange database size that is reported by the accurate licensing method is correct.

- Use the Microsoft Exchange Management shell command that is available by default on the Exchange server.

```
Get-MailboxDatabase -Status | select Name,DatabaseSize
```

Use the `help Get-MailboxDatabase` command to see more information about this command.
- Check the properties of Exchange database file (.edb) that was backed up on the Exchange server to view the size of the file.

Table 1-12 Reporting for Exchange backup environments

Exchange environment	Description
Database Availability Group (DAG)	A user can choose a DAG directive to back up all Exchange databases or back up an individual database of the DAG as a standalone database backup. Overlap is reported for a DAG. Regardless of which node you use to back up the Exchange DAG database, the capacity licensing report identifies the database uniquely across DAG nodes, and identifies overlap. (You can back up an Exchange DAG database from any node based on the server preference list that is configured in the backup policy.) The protected data size is calculated from the node that is used to backup the Exchange DAG database.

Table 1-12 Reporting for Exchange backup environments (*continued*)

Exchange environment	Description
Standalone Exchange server	For a standalone Exchange environment, protected data is reported. For multiple policies that have common databases, overlap is identified.

NetBackup for NDMP agent

Accurate licensing for NDMP is specific to an **NDMP** policy type. This type of licensing collects the front-end data size (FEDS) based on the backup selections that are defined in the **NDMP** backup policy types.

The `nbdeployutil` reports actual data usage without overlap by calculating the accurate data size using the related backup files and policies. The following rules are applied to guarantee data size accuracy:

- If the data size backed up by a policy with **ALL_FILESYSTEMS** directive is greater than the sum of other policies, the data size of **ALL_FILESYSTEMS** directive is counted.
- If identical policies are taken, then the policy with the higher size is counted.
- The uncounted policies are displayed with 0 data size in the **NBDeployUtil** report.
- If the data size backed up by policy P1 is consumed in another policy, P2, then only the data size of policy P2 is counted.
- If the data size is overlapped in multiple policies, then it is counted only once.

System administrators can use the following methods to verify the correctness of the data size that is reported by accurate licensing.

To verify the size of volumes with **NetApp ONTAP 9** filer, log on to **NetApp ONTAP** filer and run the following command:

```
volume show -fields size,used,available,percent-used.
```

Alternatively, the administrator can mount the volume of the **NDMP** filer (provided that the NFS protocol is enabled for that volume) on any UNIX-based client and run the following commands:

```
o du -sh
o ls -lh
```

NetBackup for Oracle server agent

Accurate licensing for Oracle is specific to an **Oracle** policy. The size of the data that is reported does not include the NetBackup for Oracle XML Archiver. This type of licensing collects the front-end data size (FEDS) for any Oracle backup that can be restored, not including transaction logs. Oracle Intelligent Policies define the backup selection on the **Instances and Databases** tab. The backup selection for script-based Oracle policies is defined based on what the script protects.

Note: The data size collection does not work properly if OS authentication is disabled.

Licensing data is collected for each database that is protected even if there are multiple databases on a single host or cluster. Licensing uses physical data file characteristics the Oracle database reports, not logical or segment sizes. The reason NetBackup collects data this way is that during a disaster recovery, RMAN needs to restore the full physical data file and not just its logical pieces.

Oracle Data Guard configurations are licensed on a per database basis. Each of the primary or the standby databases needs to be restored individually and FEDS licensing is used for any Oracle backup that can be restored. Each of the primary or the standby databases reports their FEDS data whenever NetBackup protects it during a backup operation.

The following Oracle queries are used to gather file size information:

- **Get size of database files being backed up**

This query retrieves the list of database files and their file sizes (in MB) for an instance:

```
select NAME, BYTES/1024/1024 from v$datafile;
```

This query shows the sum of the database file sizes for an instance:

```
select sum(BYTES/1024/1024) from v$datafile;
```

Note: The preceding queries do not have information about the transaction log.

- **Get the size of the control file**

This query retrieves the list of control files and their sizes (in MB), not including the transaction log:

```
select NAME, BLOCK_SIZE*FILE_SIZE_BKBS/1024/1024 controlfile_size  
from v$controlfile;
```

NetBackup for SQL Server agent

Licensing for SQL Server is specific to the **MS-SQL-Server** policy type. NetBackup collects the data for any SQL Server backup that can be restored, not including transaction logs.

The following SQL Server queries are used to gather file size information.

- **Get size of entire database**

Given the database name, this query gets the file size in MB, not including the transaction log:

```
USE <dbname>;
SELECT CAST(SUM(dbfile.size) AS FLOAT) / 128.0 AS FileSizeInMB
FROM sys.database_files AS dbfile
WHERE dbfile.drop_lsn IS NULL
AND dbfile.type <> 1;
```

- **Get size of entire database for skip read-only file groups option**

Given the database name, this query gets the file size in MB for skip ReadOnly file groups option:

```
USE <database_name>;
SELECT
sysFG.name AS FileGroupName,
SUM(CAST(dbfile.size AS float) / CAST(128 AS float)) AS FileSizeInMB
FROM
sys.database_files AS dbfile
INNER JOIN
sys.filegroups AS sysFG
ON
dbfile.data_space_id = sysFG.data_space_id
WHERE
sysFG.is_read_only = 0 and drop_lsn is null
GROUP BY
sysFG.name;
```

- **Get the size of file groups for a partial backup**

Given the database name, this query gets the file size in MB for file groups that are specified in the partial list:

```
USE <database_name>;
SELECT
sysFG.name AS FileGroupName,
SUM(CAST(dbfile.size AS float) / 128.0) AS FileSizeInMB
```

```

FROM
sys.database_files AS dbfile
INNER JOIN
sys.filegroups AS sysFG
ON
dbfile.data_space_id = sysFG.data_space_id
WHERE
drop_lsn is null
and sysFG.name in (<delimited fg name>, ...)
GROUP BY
sysFG.name;

```

- **Get size of file and corresponding file group name**

Given the file name (object name) and database name, this query gets the file size in MB and the corresponding file group name:

```

USE <database name>;
SELECT
sysFG.name AS FileGroupName,
(CAST(dbfile.size AS float) / 128.0) AS FileSizeInMB
FROM
sys.database_files AS dbfile
INNER JOIN
sys.filegroups AS sysFG
ON
dbfile.data_space_id = sysFG.data_space_id
WHERE
dbfile.name = N'<file name>' and drop_lsn is null

```

- **Get size of file group**

Given the file name (object name) and database name, this query gets the file group size in MB:

```

USE <database name>;
SELECT
SUM(CAST(dbfile.size AS float) / 128.0) AS FileSizeInMB
FROM
sys.database_files AS dbfile
INNER JOIN
sys.filegroups AS sysFG
ON
dbfile.data_space_id = sysFG.data_space_id

```

```
WHERE
```

```
sysFG.name = N'<filegroup name>' and drop_lsn is null
```

NetBackup for VMware agent

Accurate licensing for virtual machine (VM) is specific to a VMware policy. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The `nbdeployutil` utility reports actual data usage by calculating the accurate data size using the related backup files and policies. The following rules are applied to guarantee data size accuracy:

- If the data size backed up by the policy **Include all disks** is greater than the sum of other two policies, the data size of **Include all disks** is counted.
- If the data size backed up by the policies **Exclude data disks** and **Exclude boot disk** are greater than the **Include all disks**, the data size of **Exclude data disks** and **Exclude boot disk** is counted.
- If identical policies are taken, the policy with higher size is counted.
If different policies use the same virtual machine (VM) identifiers, they are detected as identical policies.

The uncounted policies are displayed with 0 data size in the FEDS table generated by the `nbdeployutil` utility.

System administrators can use the following ways to verify the data size reported by accurate licensing.

- If all disks are included in the backup, verify the size of the ESX datastore.
- If a specific disk is excluded during the backup, verify its size on the guest operating system (OS) of the virtual machine.

Backup of a virtual machine (VM) with an agent and also with a VADP policy

VMware virtual machine (VM) is backed up by VADP policy (all drives Included) and the NetBackup client installed inside the guest is backing non filesystem workloads (using policy types other than Standard/MS-Windows). You are only charged for the virtual machine (VM) backup.

The `nbdeployutil` report does not display a row for agent backup. There is only one row displayed corresponding to the VADP backup for the virtual machine (VM).

As VADP supports single file restore, file system backup using agent inside the guest is charged separately and corresponding rows are displayed in the `nbdeployutil` report. `nbdeployutil` uses the virtual machine (VM) DNS name to

correlate backup entries corresponding to VADP and agent backup. If the virtual machine (VM) DNS name is not recorded as part of the VADP backup, this correlation does not work. A virtual machine (VM) backup must have all drives included. If drives are excluded as part of the VADP backup, the agent backup is charged separately.

NetBackup for Hyper-V agent

Accurate licensing for virtual machine (VM) is specific to a Hyper-V policy. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup. The following rules are applied to guarantee data size accuracy:

- For Hyper-V backup with Windows Management Instrumentation (WMI) backup method: Charged size is the size that is used by configured disks.
- For Hyper-V backup with Volume Shadow Copy Service (VSS) backup method: Charged size is the total size of all configured disks.

Backup of a virtual machine (VM) with an agent and also with a Hyper-V policy

Note: This section is applicable for NetBackup primary server 8.3 or later and NetBackup client 8.3 or later.

Hyper-V virtual machine (VM) is backed up by a Hyper-V policy (all drives). The NetBackup client that is installed inside the guest is backed up with non-file system workloads (policy types other than Standard/MS-Windows). You are only charged for the virtual machine (VM) backup.

The `nbdeployutil` report does not display a row for agent backup. Only one row is displayed corresponding to the Hyper-V backup for the virtual machine (VM).

As Hyper-V supports single file restore, the file system backup with the agent inside the guest is charged separately. Corresponding rows are displayed in the `nbdeployutil` report. The `nbdeployutil` utility uses the virtual machine (VM) DNS name to correlate backup entries corresponding to the Hyper-V and the agent backup. If the virtual machine (VM) DNS name is not recorded as part of the Hyper-V backup, this correlation does not work. A virtual machine (VM) backup must have all drives included. If drives are excluded as part of the Hyper-V backup, the agent backup is charged separately.

NetBackup for RHV agent

Note: The following section is applicable for NetBackup primary server 8.2 or later and NetBackup client 8.2 or later.

The front-end data size that is reported for the Red Hat Virtualization (RHV) is same as consumed storage size. But this data is based on supported file systems. The value is accurate for NTFS, FAT, ext3,ext4 and inaccurate for ReFS, xfs file systems, or encrypted file systems.

Accurate licensing for Red Hat Virtualization (RHV) virtual machine (VM) collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The `nbdeployutil` utility reports actual data usage by calculating the accurate data size using the related backup VM size and policies. The following rules are applied to guarantee data size accuracy:

- If identical policies are taken, the policy with higher size is counted.
If different policies use the same virtual machine (VM) identifiers, they are detected as identical policies.

Backup of a virtual machine (VM) with an agent and also with a Hypervisor policy

This information is applicable for NetBackup primary server 8.3 or later and NetBackup client 8.3 or later.

RHV virtual machine (VM) is backed up by a Hypervisor policy (all drives). The NetBackup client that is installed inside the guest is backed up with non-file system workloads (policy types other than Standard/MS-Windows). You are only charged for the virtual machine (VM) backup.

The `nbdeployutil` report does not display a row for agent backup. Only one row is displayed corresponding to the RHV backup for the virtual machine (VM).

As RHV does not support single file restore, the file system backup with the agent inside the guest is not charged separately. Corresponding rows are not displayed in the `nbdeployutil` report. The `nbdeployutil` utility uses the virtual machine (VM) DNS name to correlate backup entries corresponding to RHV and agent backup. If the virtual machine (VM) DNS name is not recorded as part of the RHV backup, this correlation does not work.

NetBackup for Nutanix-AHV agent

Note: The following section is applicable for NetBackup primary server 8.3 or later and NetBackup client 8.3 or later.

Backup of a virtual machine (VM) using the Hypervisor policy

The front-end data size that is reported for Nutanix Acropolis Hypervisor is the consumed storage size. You can verify the consumed storage size on the Nutanix AHV console by navigating to **VM > Overview > Tables** under the **Storage** column. The **Storage** column displays the consumed storage size as compared to the total allocated storage size.

Accurate licensing for Nutanix-AHV virtual machine (VM) collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The `nbdeployutil` utility reports actual data usage by calculating the accurate data size using the related backup VM size and policies. The following rules are applied to guarantee data size accuracy:

- If identical policies are taken, the policy with higher size is counted.
If different policies use the same virtual machine (VM) identifiers, they are detected as identical policies.
- The Nutanix-AHV VM can be backed up using BigData policy and by Hypervisor policy (all drives Included). If the same VM is backed up using both BigData and Hypervisor policies, you are charged once for the VM backup using Hypervisor policy.

Note: File system backup using the agent inside the Nutanix-AHV VM is charged separately and corresponding separate rows are displayed in the `nbdeployutil` report.

NetBackup for SQLite Agent

Licensing for SQLite workload is specific to the DataStore policy. NetBackup collects the data for any SQLite database backup that can be restored. Only Full backup schedule is supported with SQLite database protection using the DataStore policy.

Use any native file system commands to get the size of the SQLite database file.

The following commands are used to verify the size that the accurate licensing method reports:

- Windows platform

- Lists the files on the system and their sizes

```
Dir <database directory>
```

- Linux platform

- Lists the files on the system and their sizes

```
ls <database directory>
```

NetBackup for MariaDB Agent

Licensing for MariaDB is specific to the DataStore policy for protecting MariaDB Instances. NetBackup collects the data for any MariaDB backup that can be restored.

In the case of MariaDB, the sizing information is gathered by listing the size of the Instance data directory of MariaDB.

The following commands are used to verify the size that the accurate licensing method reports:

- Windows platform

- Lists the files on the system and their sizes

```
Dir <Instance data directory>
```

- Linux platform

- Lists the files on the system and their sizes

```
ls <Instance data directory>
```

NetBackup for PostgreSQL Agent

Licensing for PostgreSQL is specific to the DataStore policy for protecting PostgreSQL Instances. NetBackup collects the data for any PostgreSQL backup that can be restored.

In the case of PostgreSQL, the sizing information is gathered by listing the size of the Instance data directory of PostgreSQL.

The following commands and queries are used to verify the size that the accurate licensing method reports:

- Windows platform

- Lists the files on the system and their sizes.

```
Dir <Instance data directory>
```

- Get the database size using PostgreSQL query.

```
SELECT pg_database_size ('DBNAME');
```

- Linux Platform

- Lists the files on the system and their sizes.

```
ls <Instance data directory>
```

- Get the database size using PostgreSQL query.

```
SELECT pg_database_size ('DBNAME');
```

NetBackup for MySQL Agent

Licensing for MySQL is specific to the DataStore policy for protecting MySQL Instances. NetBackup collects the data for any MySQL backup that can be restored.

In the case of MySQL, the sizing information is gathered by listing the size of the Instance data directory of MySQL.

The following commands and queries are used to verify the size that the accurate licensing method reports:

- Windows platform

- Lists the files on the system and their sizes.

```
Dir <Instance data directory>
```

- Get the database size using MySQL query.

```
SELECT ROUND(SUM(data_length + index_length), 1) FROM
information_schema.tables WHERE table_schema = 'DBNAME' GROUP
BY table_schema;
```

- Linux Platform

- Lists the files on the system and their sizes.

```
ls <Instance data directory>
```

- Get the database size using MySQL query.

```
SELECT ROUND(SUM(data_length + index_length), 1) FROM
information_schema.tables WHERE table_schema = 'DBNAME' GROUP
BY table_schema;
```

NetBackup for SAP HANA

Licensing for SAP HANA is specific to the SAP policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

NetBackup calculates the size by summing up the size of all files in the HANA Instance data directory.

For example, if you have the following SAP HANA instance directory:

`/usr/sap/InstanceName/SYS/global/hdb/data/`, the sum of the size of all files available in this directory is calculated.

The following command is used for each subdirectory to verify the size that the accurate licensing method reports:

```
ls -lh <Subdirectory path>
```

NetBackup for FlashBackup

Licensing for FlashBackup is specific to the FlashBackup policy type for UNIX clients and FlashBackup-Windows policy type for Windows clients. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following rules are applied to guarantee data size accuracy for FlashBackup:

- For a FlashBackup-Windows policy type, size of the mounted Windows disk volume is used to gather size information.
- For a FlashBackup policy type (UNIX clients), size of the raw disk partition is used to gather size information.

NetBackup for Enterprise Vault

Licensing for Enterprise Vault is specific to the Enterprise-Vault policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following Enterprise Vault queries are used to gather file size information for SQL databases:

```
SELECT SUM (CAST (COALESCE(FILEPROPERTY(dbfile.name, 'SpaceUsed'),
dbfile.size) AS FLOAT(53))) / 128.0 AS FileSizeInMB
FROM sys.database_files AS dbfile
WHERE dbfile.drop_lsn IS NULL AND dbfile.type <> 1;
```

The following rule are applied to guarantee data size accuracy for file partition:

Administrator can verify the size of the file partition protected by Enterprise-Vault policy.

NetBackup for Lotus Notes

Licensing for Lotus Notes is specific to the Lotus-Notes policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following rules are applied to guarantee data size accuracy for Lotus Notes:

- For Windows clients, an administrator can verify the size of the supported files that are protected.

- For Linux clients, an administrator can run the following UNIX commands to verify the size:
 - `df -l`
 - `ls -lh`

NetBackup for Sybase

Licensing for Sybase is specific to the Sybase policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

Use the `sp_spaceused` Sybase query to validate database size reported in usage report.

Refer to the Sybase documentation for more details about using the `sp_spaceused` command.

NetBackup for SAP MaxDB

Licensing for SAP MaxDB is specific to the SAP policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

NetBackup calculates the size considering total number of bytes read from MaxDB.

The following SAP MaxDB query is used to validate reported size in the usage report:

```
dbmcli -d <dbname> -u <username>,<password> info state
```

Refer to the MaxDB documentation for more details about using the query.

NetBackup for DB2

Licensing for DB2 is specific to the DB2 policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following DB2 query can be used to validate database size:

```
db2 call GET_DBSIZE_INFO(?,?,?,0)
```

Refer to the DB2 documentation for more details about using the query.

NetBackup for Informix

Licensing for Informix is specific to the Informix-On-Bar policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following Informix query is used to gather file size information of all the databases in all the DB spaces:

```
DATABASE sysmaster;
SELECT
    SUM(i.ti_npused * i.ti_pagesize) USED_SIZE
FROM
    systabnames n, systabinfo i, sysdatabases d
WHERE
    n.partnum = i.ti_partnum
AND
    n.dbsname = d.name;
CLOSE DATABASE;
```

NetBackup for Microsoft SharePoint

Licensing for Microsoft SharePoint is specific to the MS-SharePoint policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following Microsoft SharePoint query is used to gather file size information for each SQL databases representing the web application:

```
SELECT SUM (CAST (COALESCE(FILEPROPERTY(dbfile.name, 'SpaceUsed'),
dbfile.size) AS FLOAT(53))) / 128.0 AS FileSizeInMB
FROM sys.database_files AS dbfile
WHERE dbfile.drop_lsn IS NULL AND dbfile.type <> 1;
```

In an overlap scenario, where one policy protects **Microsoft SharePoint Resources** and another policy selects an individual web application (for example, WebApp1), `nbdeployutil` detects the overlapping selection and you are charged for Microsoft SharePoint Resources. If you add any additional data to the individual web application (WebApp1), it is updated in the usage report only after the next successful backup of the Microsoft SharePoint Resources.

NetBackup for SAP Oracle

Licensing for SAP Oracle is specific to the SAP policy type. This type of licensing collects the total number of Front-End Terabytes (FETBs) protected by NetBackup.

The following SAP Oracle queries are used to gather file size information:

For SAP Oracle backup, charged size is equal to the sum of **Data file size** and **Control file size**.

- For Data files: `select sum(BYTES/1024/1024) from v$datafile;`
- For Control files: `select BLOCK_SIZE*FILE_SIZE_BKLS/1024/1024
controlfile_size from v$controlfile;`

NAS-Data-Protection policy

Capacity licensing for all supported CloudPoint plug-ins is specific to the NAS-Data-Protection policy type.

The **Client Name** column in the itemization sheet of the licensing report displays the backup host selected at the time of backup.

The **Policy Name** column in the itemization sheet of the licensing report displays the policy name in the SLP_<Policy Name>_<NAS Volume Name format.

The entries in the report are based on per volume. See [“Clients backed up with multiple streams”](#) on page 45.

Mount NFS share on the UNIX system and you can run the following commands to verify the size reported by the accurate licensing method.

- `df -l`
- `ls -lh`

Cloud policy

Capacity licensing for all supported CloudPoint plug-ins is specific to the Cloud policy type.

- The **Client Name** column in the itemization sheet of the licensing report displays the backup host selected at the time of backup. If you back up an application or a volume then the **Client Name** column shows the host name on which the application resides or the volume is mounted.
- The **Policy Name** column in the itemization sheet of the licensing report displays the policy name in the Policy Name+<id> format.
- Usage is calculated based on disk size.
- Usage is calculated based on the accurate snapshot size.
NbDeployutil Capacity Report might report the total volume size as the snapshot size instead of the actual used size, if the CloudPoint host does not have the required permissions to read the accurate snapshot size.
See [“Troubleshooting failures for usage reporting and incremental reporting”](#) on page 30.
For information about the required permissions, refer to the *NetBackup CloudPoint Install and Upgrade Guide*.
- The entries in the report are based on each asset selected for backup.

After the `nbdeployutil` utility runs using the capacity licensing option, the report displays the policy type in the Itemization sheet as follows:

- **Cloud:aws**
- **Cloud:azure**
- **Cloud:gcp**
- **Cloud:azurestack**

NetBackup for Kubernetes

NetBackup supports capacity based licensing for Kubernetes. Here are a few points to note:

- NetBackup charges only for the persistent volumes snapped as part of snapshot operation.
- NetBackup considers persistent volumes contained within the namespace for calculation when you perform snapshot for a namespace.
- Kubernetes does not publish the actual allocated or used blocks within a volume. Hence, for license usage reporting, NetBackup considers the provisioned size.
- The client name reported under licensing reports is of the form: *<Cluster Name>:<Namespace Name>*. If no persistent volumes were snapped as part of backup, and for the policy and client combination not a single backup image exists with non-zero size reported, nbdeployutil adds a comment "size unavailable". In this case, no charges are levied.

Windows file system backups

The protected data for Windows file system backups is defined as the size of the data for the pathnames or the directives that are defined in the policy. Licensing for the Windows backup is specific to an **MS-Windows** policy. It does not support the Windows clients or servers that are backed up using any virtualization policy, such as **VMware** and **HyperV**.

Snapshot-based backups are supported on all the versions of Windows servers and clients with various features of NetBackup such as multistreaming, Accelerator, compression, and so on.

Administrators can use the following to verify the size that the accurate licensing method reports.

- `Dir`
Lists the files on the system folders and their sizes.
- `- Get-ChildItem C:\test1 | Measure-Object -property length -sum`
PowerShell command to get folder size.
- Size information from Windows Explorer.

To compare the data size that is reported in the capacity license report:

- **ALL_LOCAL_DRIVES**
Use the file system commands to calculate size of each drive. In the capacity license report, the files that are mentioned in the exclude list are not used to calculate protected data.
- **System State**
NetBackup creates the xml files of backup data under `logs\BEDS` folder and lists the files that are backed up and excluded from backup.
- **Shadow Copy Components**
NetBackup creates the xml files of backup data under `logs\BEDS` folder and lists the files that are backed up and excluded from backup.
- **Folders and files**
Use the file system commands.

UNIX file system backups

Licensing for UNIX is specific to **Standard** policy type and does not support UNIX clients or the servers that are backed up using any virtualization policy, such as **VMware**.

Using the **Standard** policy, the UNIX file system data is backed up for the directive or the backup selection that is defined in the policy. The definition of protected data for the policy is the size of **ALL_LOCAL_DRIVE** or specific mount points\partitions or files.

- **Specific mount points\partition or files**
Backs up the files and folders that are specified in the backup selection. Run the file system command and compare the size that is reported in the capacity licensing report.
- **ALL_LOCAL_DRIVE**
NetBackup posts all data under **ALL_LOCAL_DRIVE** backup directive. Run the file system commands, calculate the size of each mount point\partition, and compare the size that is reported by the capacity licensing report. The files that are mentioned in the exclude list are not included in the calculation.

The snapshot based backups are supported on all the versions of UNIX server and clients with various features of the NetBackup such as multistreaming, Accelerator, compression, and so on.

Administrators can run UNIX commands to verify the size reported by the accurate licensing method.

- `df -l`

- `ls -lh`

For more information on path names and directives and mount points and partitions, see the *NetBackup Administrator's Guide, Volume I*.

For more information on supported Windows and UNIX file systems, see the *NetBackup Software Compatibility List*:

<http://www.netbackup.com/compatibility>

Reconciling the capacity licensing report results

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 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 want 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.

Locate full backups for clients

On the **Itemization** tab, sort the list by the **Accuracy** column. For all rows that indicate 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. Run the utility again 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** column. For any compressed images, review the **Charged Size** column and confirm that 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 of clients

On the **Itemization** tab, sort the list by **Client Name** and search for the use of host name aliases. Look for instances where the itemization table lists the same client multiple times under the same policy but with a different host name alias. If that occurs, zero out the **Charged Size** column for the lines with an earlier backup date. 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 that those policies do not have overlapping backup selections. If the backup selections overlap, find the redundant backup policies in the **Itemization** tab. 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.

Note: For the agents that support accurate licensing, multiple host name aliases do not exist.

Determine the effect of multistreamed backups

On the **Itemization** tab, sort the list by the **Accuracy** column. Find all backups that list **Possible multi-stream backup detected** and make a note of the policy name under the **Policy Name** column. 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.

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

Confirm the accuracy of any database backups

Reconcile database backups in the same way that 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 backups for snapshot images

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.

Reviewing a traditional licensing report

Follow these steps to examine and reconcile the report with the actual NetBackup environment..

Table 1-13 Examining the traditional licensing report

Step	Description	Reference
1	Examine the Summary tab and confirm that the correct information is displayed.	See "Summary tab" on page 68.
2	Review the Hosts tab and resolve any missing information.	See "Complete the Hosts tab" on page 68.
3	Review the Itemization tab and resolve any missing information.	See "Update the Itemization tab" on page 69.

Table 1-13 Examining the traditional licensing report (*continued*)

Step	Description	Reference
4	Review the NDMP tab and resolve any missing information.	See “Resolve the NDMP tab” on page 70.
5	Review the Virtual Servers tab and resolve any missing information.	See “Update the Virtual Servers tab” on page 70.
6	Review the Drives tab and resolve any missing information.	See “Confirm the Drives tab” on page 71.

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

The start and the end dates are not necessarily the dates that are specified for the `--gather` option. 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 the **Input 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 that the list of primary servers is correct. If there are missing or extra primary servers, rerun the report.

When the review of the entire report is complete, all the values in the **Unknown** row under **Tiering** should be zero. As you reconcile the other tabs in the report, these values 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 primary servers if they are either a media server or a client server. Review five areas to complete the review of this tab.

To complete the Hosts tab

- 1 Scan the **Connectable** column to see if the utility was unable to connect to any hosts 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:

```
netbackup_deployment_insights --retry path_to_the_gathered_data
```

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

```
netbackup_deployment_insights --report  
all_previously_specified_options  
all_previously_specified_gather_directories
```

You can still continue using the `nbdeployutil` command with the same options.

- 2 Check the **Tier** column for any hosts that are listed as UNKNOWN. Replace these with the appropriate tier number between one and four. Work with your Veritas 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 by knowing this information you can determine the appropriate value to enter in the Tier column.
- 3 Review the **MSEO Key Server** column and verify that all the listed information is correct. Yes indicates that the host is an MSEO key server. No indicates that the host is not an MSEO key server. The N/A value indicates that the host is not a media server.
- 4 Check the **Enterprise Client** column and verify that the information is correct. Yes indicates that the host is an enterprise client and was backed up. No indicates that the host is not an enterprise client. The N/A value indicates that no backups were performed on the host during the report period.
- 5 Review the **SAN Media Server** column and correct any hosts where the value is **UNKNOWN**. Confirm that all other values are correct. A value of N/A for a host indicates that the host is either a client server or a primary server.

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

Update the Itemization tab

The report's **Itemization** tab shows the calculated capacity for each client or policy combination. The report flags any conditions that have the potential to over count

or to under count capacity. These conditions are identified in the **Accuracy** and **Accuracy Comment** columns.

- **Machine Type** column:
Indicates whether the client is a physical or virtual machine.
- **VM Host** column:
Displays the host name of the client. It may contain “UNKNOWN”, if nbdeployutil cannot find the corresponding host name for that client. This occurs:
 - If the VMware tools are not installed on the client.
 - If the client name does not match with any of the Display Name, INSTANCE_UUID, HOST_NAME, BIOS_UUID in case of VMware policy type.

Resolve the NDMP tab

The **NDMP** tab lists the hosts that 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 that the information is correct. Any **Tier** values of **UNKNOWN** should be replaced with the correct tier number between one and four. Work with your Veritas Sales Engineer and the NetBackup Pricing and Licensing Guide to determine the correct tier information.

Update the Virtual Servers tab

Complete the **Virtual Servers** tab. **CpuCount** displays the socket count for each server. It may contain UNKNOWN if the utility cannot identify the correct count. **Cluster** displays name of the cluster, if the server is part of some cluster. **Hypervisor_type** displays the type of Hypervisor under which the server belongs, For example, VMware, Hyper-V, and so on.

BackedUp indicates if a server is part of any backup through any policy that is specific to a hypervisor. The following values are displayed:

- **No:** No VM is backed up from the hypervisor.
- **Yes, Inside Guest:** VM on the hypervisor is backed up by deploying the NetBackup Agent inside the VM.
- **Yes, Outside Guest:** VM on the hypervisor is backed up using a policy that is specific to the hypervisor. For example, using VMware policy to back up a VM on an ESX server.

If a VM is backed up using both inside and outside backup, nbdeployutil displays **Yes, Outside Guest** in the Virtual Server sheet.

Confirm the Drives tab

On the **Drives** tab, review the information in the **VL** 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 the value 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 Veritas or your reseller. Confirm that everything in the report matches with the content in the install base report. If there are discrepancies, Consult with your Veritas sales representative to correct any discrepancies.

Reviewing an NEVC licensing report

The following table describes the steps to examine the NEVC license report results and how to reconcile the report with the actual NetBackup environment.

Table 1-14 Examining the NEVC licensing report

Step	Description	Reference
1	Verify the information on which the report is based, which includes the primary server, hosts, and CPU count.	See "Summary tab" on page 71.
2	Review the Hosts tab and resolve any missing information.	See "Hosts tab" on page 72.

Summary tab

The top of the **Summary** tab shows the basis for the report information. Examine the section marked **Analyzed** to verify the gathered data.

The **Analyzed** section displays the following information:

- The primary servers that are included in the report.
- The total number of hosts.
- The total number of CPU sockets for all hosts.

The analyzer gathers 90 days of CPU socket count for all hosts by default.

Hosts tab

The **Hosts** tab provides a list of all hypervisors and CPU count.

Table 1-15 Hosts tab

Option	Description
Master Server	Displays the primary server from where <code>nbdeployutil</code> is run and backups are generated.
Host Name	Displays the name of the hypervisor.
Datacenter	Displays the datacenter where the hypervisor is hosted.
Cluster	Displays the name of the cluster if hypervisor is part of a cluster.
Policy Type	Displays the policy type.
Host CPU Count	Displays the number of CPU sockets on the host.

Master Server	Host Name	Datacenter	Cluster	Policy Type	Host CPU Count
master_server_1	esx_host_1	/datacenter1	esx_cluster_1	VMWare	2
master_server_1	esx_host_2	/datacenter1	esx_cluster_1	VMWare	2
master_server_1	esx_host_3	/datacenter1	esx_cluster_1	VMWare	2
master_server_1	esx_host_4	/datacenter1	esx_cluster_1	VMWare	2
master_server_1	esx_host_5	/datacenter1	esx_cluster_1	VMWare	2
master_server_1	nutanix_hypervisor_1		nutanix_cluster_1	BigData:Nutanix-AHV	2
master_server_1	nutanix_hypervisor_2		nutanix_cluster_1	BigData:Nutanix-AHV	2
master_server_1	nutanix_hypervisor_3		nutanix_cluster_1	BigData:Nutanix-AHV	2
master_server_1	nutanix_hypervisor_4		nutanix_cluster_1	BigData:Nutanix-AHV	2

Using the NetBackup Storage API to get the total backup size information

By default, the total backup size information gather is disabled to reduce the load on the system resources. Beginning with NetBackup 8.1.2, you can use the NetBackup Storage API to get the total backup size.

The following table briefly explains the process to use the NetBackup Storage API to gather total backup size information.

Table 1-16 Process to use the NetBackup Storage API to gather total backup size information

Step	Procedure	Reference topics
1	Enable the total backup size information gather.	
2	Use the default values for backup size information gather or configure custom values.	
3	Use the NetBackup Storage API to gather total backup size information	

Enabling the total backup size information gather

By default a scheduled run of the `nbdeployutil` utility does not gather the total backup size information.

To enable the total backup size information gather

In the `nbdeployutilconfig.txt` file, add the `BETB_ENABLE` parameter:

```
[NBDEPLOYUTIL_BETB]
BETB_ENABLE=1
```

If the option is not enabled, the NetBackup Storage API displays the 404 Not Found error and records the following response:

```
{
  "errorCode": 227,
  "errorMessage": "no entity was found",
  "details": {}
}
```

Configuring the parameter values for backup size information gather

Use the following information to update the backup size information gather parameters in the `nbdeployutilconfig.txt` file:

1. Open the `nbdeployutilconfig.txt` file from the following location:

- Windows:


```
install_path\netbackup\var\global\
```
- UNIX:


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

2. Update the data gather parameters as described in the following table:

Parameter	Description
MASTER_SERVERS=<server names>	Use this option to gather backup size information from other primary servers. You can add multiple primary servers as comma-separated values. By default, the local server is selected as the primary server.
BETB_ENABLE	Enable or disables the total backup size information gather. The default value is 0. Set the value to 1 to enable the total backup size information gather.
BETB_PARENTDIR	Location where the backup size information is gathered and analyzed. A gather directory is created for every primary server. You can specify a custom directory or location for data gather. The default gather location is: ■ Windows: <pre>install_path\netbackup\var\global\</pre> ■ UNIX: <pre>/usr/opensv/var/global/</pre> The location of the gather directory is: ■ Windows: <pre><gather_location>\storage\staging\ <timestamp_primary_server_name></pre> ■ UNIX: <pre>/gather_location/storage/staging/ <timestamp_primary_server_name></pre>
BETB_FREQUENCY_IN_DAYS	The frequency in days to gather the backup size information. The default value is 1, which means that the gather runs daily. If you want a weekly report, set the value to 7.
BETB_KEEP_CMD_OUT_FILE	Option to keep or remove the <code>bpimagerlist.out</code> and <code>bpstulist.out</code> files. The default value is 0, which means the files are deleted after <code>nbdeployutil</code> runs. If you want to keep the file, set the value to 1. The sample file shows where to put the option.
BETB_LOG_KEEP	Duration in days to keep the gather folder or folders. Logs are located in the gather directory. The default value is 7, which means that the gathered data from the last seven runs is kept in the default folder or in the directory that you set in the <code>BETB_PARENTDIR</code> parameter.

Sample nbdeployutilconfig.txt file

Here is a sample `nbdeployutilconfig.txt` file that you can refer to:

```
[NBDEPLOYUTIL_INCREMENTAL]
MASTER_SERVERS=nbu.primaryserverone.com,nbu.primaryservertwo.com
[NBDEPLOYUTIL_BETB]
BETB_ENABLE=1
BETB_PARENTDIR=install_path\netbackup\var\global\>
BETB_FREQUENCY_IN_DAYS=1
BETB_KEEP_CMD_OUT_FILE=0
BETB_LOG_KEEP=7
```

About the NetBackup Storage API

The NetBackup Storage API provides access to the backup storage information of the NetBackup primary servers.

To view the NetBackup API Reference document

1. Go to the following URL:
<https://sort.veritas.com/documents>
2. From the product list, select NetBackup.
3. Find NetBackup 10.3 in the list of NetBackup product versions and click **Documentation**.
4. Locate **API Reference** and click on the document icon.

Additional configuration

This chapter includes the following topics:


- [About multiple NetBackup primary servers](#)
- [About multiple media servers with one primary server](#)
- [About direct I/O for backups on Windows](#)
- [About dynamic host name and IP addressing](#)
- [About busy file processing on UNIX clients](#)
- [About specifying the locale of the NetBackup installation](#)
- [About the Shared Storage Option](#)
- [About the vm.conf configuration file](#)

About multiple NetBackup primary servers

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

[Figure 2-1](#) 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 2-1 Multiple primary server scenario



In this environment, the two NetBackup server configurations are completely independent. You can also create a configuration where one server is the primary and the other is a media server.

About multiple media servers with one primary server

A protection domain refers collectively to the NetBackup primary 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 primary 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 primary 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 primary and the media servers.

The following are important factors to remember about primary and media servers:

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

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

Figure 2-2 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 the 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 primary and media server installation.

About NetBackup catalogs

The primary server is the default location for the NetBackup catalogs. The catalogs include the media and the volume database. The volume database contains the media usage information and the volume information that are used during the backups.

Selecting a different media server to fix storage server and disk volume connectivity issues

Storage servers and disk volumes are monitored for connectivity and capacity using one of the media servers that is credentialed to access the storage device. Resource connectivity state changes are recorded in `bpererror -disk` reports. If storage resources become unavailable, new jobs are not scheduled until this connectivity path is re-confirmed through the previously selected media server.

If the problem exists only between the selected media server and storage, a new path is not automatically selected. You need to change the NetBackup configuration to use a different media server.

To force the storage server to use a different media server, the following commands can be executed from the problematic media server:

```
tpconfig -delete -storage_server <storage server>  
-stype <storage type> -sts_user_id <userid>  
  
bpstsinfo -resyncREM -servername <primary server name>
```

To perform the change using the NetBackup Administration Console

- 1 Select **Credentials > Storage Servers**.
- 2 Right-click on the storage server name and select **Change**.
- 3 Deselect the problematic media server.
- 4 Execute the following command on the storage server:

```
bpstsinfo -resyncREM -servername <primary server name>
```

About direct I/O for backups on Windows

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:

- A Windows media server must own the storage unit.
- 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, Veritas recommends that you use `SIZE_DATA_BUFFERS_DISK`.

[Table 2-1](#) shows the possible values to include in `SIZE_DATA_BUFFERS_DISK` or `SIZE_DATA_BUFFERS`.

Table 2-1 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 on Windows

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 computer. This assumption makes it difficult to back up any clients that have network host names that might change. For example, a computer 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 2-2 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.	NetBackup requires that IP addresses of clients have a network host name. (On Windows) 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. (On UNIX) Be sure to define network host names for the range of dynamic IP addresses in the <code>hosts</code> file, NIS, and (or) DNS on the network.
Determine the NetBackup client names for the computers that have dynamic IP addresses and network host names.	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.

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

Action	Process details/requirements
Make changes on the primary server, as described.	■ 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.	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 computer in the Client Name text box.
On the primary server, enable the Announce DHCP Interval option, as described.	In the NetBackup Administration Console , in the left pane, expand NetBackup Management > Host Properties > Clients . Double-click on 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 check box.
Make changes on each dynamic NetBackup UNIX clients, as described.	■ Modify the <code>bp.conf</code> file to include a <code>CLIENT_NAME</code> entry with the correct NetBackup client name for the computer. ■ Configure the system to notify the primary server of the computer's NetBackup client name and current network host name during startup. The <code>bpdynamicclient</code> command is used to notify the primary server. ■ Configure the system to notify periodically the primary server of the computer'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, ten 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 computers share the IP addresses as previously defined.

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

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

Configuring the NetBackup primary server

Use the following procedure to configure the NetBackup primary server.

To configure the NetBackup primary server

- 1** On the primary 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 primary server.

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

On Windows:

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

On UNIX:

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

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

The `bpclient` command is in the following directory:

On Windows:

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

On UNIX:

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

See [“bpclient commands that control client entries”](#) on page 87.

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

On Windows:

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

On UNIX:

```
cd /usr/opensv/netbackup/bin/admincmd
bpclient -add -client nbclient01 -dynamic_address 1
bpclient -add -client nbclient02 -dynamic_address 1
bpclient -add -client nbclient03 -dynamic_address 1
bpclient -add -client nbclient04 -dynamic_address 1
.
.
.
bpclient -add -client nbclient20 -dynamic_address 1
```

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

On Windows:

```
install_path\NetBackup\bin\admincmd\bpclient -L -All
```

On UNIX:

```
/usr/openv/netbackup/bin/admincmd/bpclient -L -All
```

The output is similar to the following:

```
Client Name: nbclient01
Current Host:
Hostname: *NULL*
IP Address: 0.0.0.0
Connect on non-reserved port: no
Dynamic Address: yes

Client Name: nbclient02
Current Host:
Hostname: *NULL*
IP Address: 0.0.0.0
Connect on non-reserved port: no
Dynamic Address: yes
.
.
.
Client Name: nbclient20
Current Host:
Hostname: *NULL*
IP Address: 0.0.0.0
Connect on non-reserved port: no
Dynamic Address: yes
```

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 2-3 `bpclient` commands that control client entries

Action	Command
Create a dynamic client entry	On Windows: <code>bpclient.exe -add -client <i>client_name</i> -dynamic_address 1</code> On UNIX: <code>bpclient -add -client <i>client_name</i> -dynamic_address 1</code> 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's possible to create entries with <code>-dynamic_address 0</code> for static IP addressing. However, to do so is unnecessary and adversely affects performance.
Delete a client entry	On Windows: <code>bpclient.exe -delete -client <i>client_name</i></code> On UNIX: <code>bpclient -delete -client <i>client_name</i></code>
List a client entry	On Windows: <code>bpclient.exe -L -client <i>client_name</i></code> On UNIX: <code>bpclient -L -client <i>client_name</i></code>
List all client entries	On Windows: <code>bpclient.exe -L -All</code> On UNIX: <code>bpclient -L -All</code>

Configuring dynamic NetBackup clients

Configuring a dynamic Windows client

Use the following procedure to configure a dynamic Windows client.

To configure a dynamic Windows client

- 1 If it's 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 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** check box.

Additional information is available for **Announce DHCP Interval** in the [NetBackup 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 computer, as follows:

```
CLIENT_NAME = nbclient00
```

- 3 Run the `bpdynamicclient` command once when the system first starts up. `bpdynamicclient` notifies the NetBackup server of the computer'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 primary server. If the host names do not match, `bpdynamicclient` connects to the primary 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 computer 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 busy file processing on UNIX clients

Busy file processing applies only to UNIX clients.

Information about VSP (Volume Snapshot Provider) is available for Windows clients.

See the [NetBackup Administrator's Guide, Volume I](#).

A busy file is a file that was detected as changed during a user or a scheduled backup. Typically, detection occurs if a process writes to a file while NetBackup attempts to back it up.

The following conditions result in the detection of busy files:

- Read error on the file
- File modification time changed
- File inode time changed
- File size changed

The backup usually completes with a status of 1, which indicates that the backup was partially successful. Busy file processing allows the user control the actions of NetBackup when busy files are detected.

Busy file processing can be configured in the **Busy File Settings** host properties for UNIX clients.

See the [NetBackup Administrator's Guide, Volume I](#).

Busy file processing can also be enabled by adding the `BUSY_FILE_PROCESSING` option to the client `/usr/opensv/netbackup/bp.conf` file. Then add other busy file options to control the processing of busy files. The options can exist in both the client `/usr/opensv/netbackup/bp.conf` file and a user's `$HOME/bp.conf`. The user's `bp.conf` file takes precedence when the options are in both places.

NetBackup creates several files and directories when it processes busy files. Initially, a working directory named `busy_files` is created under `/usr/opensv/netbackup`. NetBackup then creates the `/actions` directory under `busy_files` and places `action` files in that directory. An `action` file contains the information that NetBackup uses to control the processing of busy files.

By default, the contents of the action file are derived from the `BUSY_FILE_ACTION` options in `bp.conf`. A user can also create an action file to control a specific backup policy and schedule. NetBackup creates a logs directory under `busy_files` for storing busy file status and diagnostic information.

Configuring busy file processing on UNIX

Use the following procedure to use the `bp.conf` file to configure busy file processing.

To configure busy file processing

- 1 Modify the `bp.conf` file options.

See [“Modifying bp.conf to configure busy file processing on UNIX”](#) on page 93.

- 2 Copy the `bpend_notify_busy` script, located on the primary server:

```
/usr/opensv/netbackup/bin/goodies/bpend_notify_busy
```

to the following path on the client:

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

Be sure to set the file access permissions to allow groups and others to run `bpend_notify`.

(This step is also performed when configuring busy file processing in the **Busy File Settings** host properties.)

- 3 Configure a policy with a user backup schedule for the busy file backups.

This policy services the backup requests that the `repeat` option in the `actions` file generates. The policy name is significant. By default, NetBackup alphabetically searches (upper-case characters first) for the first available policy with a user backup schedule and an open backup window. For example, a policy name of `AAA_busy_files` is selected ahead of `B_policy`.

(This step is also performed when configuring busy file processing in the **Busy File Settings** host properties.)

Modifying bp.conf to configure busy file processing on UNIX

Use the following procedure to modify the `bp.conf` file to configure busy file processing.

To modify the `bp.conf` file to configure busy file processing

- 1 Configure busy file processing by using the **Busy File Settings** host properties for UNIX clients.

See the [NetBackup Administrator's Guide, Volume I](#).

- 2 Or, configure busy file processing by using the entries in the `bp.conf` file on the client. The user can also configure a `bp.conf` file in a home directory. The busy file options that are specified in the user's `bp.conf` file apply only to user-directed backups. NetBackup ignores busy file processing for user backups if a `bp.conf` file does not exist in their home directory.

The `bp.conf` file entries to use are as follows:

- `BUSY_FILE_PROCESSING`
- `BUSY_FILE_DIRECTORY`
- `BUSY_FILE_ACTION`

bp.conf file entries on UNIX

The following table describes the `bp.conf` file entries that are used to configure busy file processing.

Table 2-4 `bp.conf` file entries

Entry	Description
<code>BUSY_FILE_PROCESSING</code>	Enables the NetBackup busy file-processing feature. By default, this entry is not present in the client's <code>/usr/opensv/netbackup/bp.conf</code> file.
<code>BUSY_FILE_DIRECTORY</code>	Specifies an alternate path to the busy files working directory. This entry is not required. By default, this entry is not present in the client's <code>/usr/opensv/netbackup/bp.conf</code> or <code>\$HOME/bp.conf</code> file. By default, NetBackup creates the <code>busy_files</code> directory in <code>/usr/opensv/netbackup</code> or the user's home directory.

Table 2-4 bp.conf file entries (*continued*)

Entry	Description
BUSY_FILE_ACTION	Directs the action that NetBackup performs on busy files. By default, this entry is not present in the client's <code>/usr/opensv/netbackup/bp.conf</code> or <code>\$HOME/bp.conf</code> file. Multiple entries may exist in the following form: <pre>BUSY_FILE_ACTION = filename_template action_template</pre> Where ■ <i>filename_template</i> is the absolute pathname and file name of the busy file. The shell language metacharacters <code>*</code>, <code>?</code>, <code>[]</code>, <code>[-]</code> can be used to match the patterns of file names or parts of file names. ■ <i>action_template</i> is one of the following: MAIL mail Directs NetBackup to mail a busy file notification message to the user that the <code>BUSY_FILE_NOTIFY_USER</code> option specifies. REPEAT repeat [repeat_count] Directs NetBackup to retry the backup on the specified busy file. A repeat count can be specified to control the number of backup attempts. The default repeat count is 1. IGNORE ignore Directs NetBackup to exclude the busy file from busy file processing. The file is backed up and a log entry that indicates that the file was busy appears in the All Log Entries report. BUSY_FILE_NOTIFY_USER Specifies the recipient of the busy file notification message when <code>BUSY_FILE_ACTION</code> is set to <code>MAIL</code> or <code>mail</code>. By default, <code>BUSY_FILE_NOTIFY_USER</code> is not in <code>/usr/opensv/netbackup/bp.conf</code> or <code>\$HOME/bp.conf</code> file on a client. By default, the mail recipient is root. See Table 2-5 on page 95.

The following table shows examples of how a `BUSY_FILE_ACTION` entry works.

Table 2-5 Examples of how a `BUSY_FILE_ACTION` entry works

Example	Description
<pre> BUSY_FILE_PROCESSING BUSY_FILE_DIRECTORY = /tmp BUSY_FILE_NOTIFY_USER = kwc BUSY_FILE_ACTION = /usr/* mail BUSY_FILE_ACTION = /usr/local ignore </pre>	This example causes NetBackup to take the following actions when it encounters busy files: ■ Create the busy files working directory in <code>/tmp</code> ■ Send an email notification message to user <code>kwc</code> for all busy files that it finds under <code>/usr</code> except for those in <code>/usr/local</code>.
<pre> BUSY_FILE_PROCESSING BUSY_FILE_ACTION = /usr/* repeat 2 BUSY_FILE_ACTION = /usr/openv mail BUSY_FILE_ACTION = /usr/local ignore </pre>	This example causes NetBackup to take the following actions when it encounters busy files: ■ Send a busy file notification message to root for busy files in <code>/usr/openv</code>. ■ Repeat the backup up to two times for all busy files that it finds under <code>/usr</code>, except for those in <code>/usr/openv</code> and <code>/usr/local</code>. ■ Exclude the busy files in <code>/usr/local</code> from all actions.

How NetBackup creates and uses action files on UNIX

When a backup operation begins, NetBackup creates a default action file named `actions` in the `busy_files/actions` directory. The contents of the `actions` file are derived from the `BUSY_FILE_ACTION` options in the `bp.conf` file.

Normally, NetBackup refers to the default action file for all future busy file processing. To override the default, create an action file to control a specific backup policy and schedule. The following entries show the naming convention for the policy and the schedule action files:

```

actions.policy_name.schedule_name
actions.policy_name

```

Where *policy_name* and *schedule_name* correspond to a predefined backup policy and schedule.

NetBackup performs the following steps when it searches for an action file.

Table 2-6 NetBackup steps when it searches for an action file

Step	Example
Checks for a file that names a specific policy and schedule, such as:	<code>actions.policy_name.schedule_name</code>
If a file for a specific policy and schedule is not found, NetBackup searches for a less specific name, such as the following:	<code>actionpolicy_names</code>
If a less specific name does not exist, NetBackup refers to the default action file. The contents of user-created action files are similar to the default. Optional comment lines can be included. The specification is the same as for the <code>BUSY_FILE_ACTION</code> option:	<pre># comment_line filename_template action_template</pre> Example 1: The <code>bp.conf</code> file might contain the following: <pre>BUSY_FILE_ACTION = /usr/openv mail BUSY_FILE_ACTION = /usr/* repeat 2 BUSY_FILE_ACTION = /usr/local ignore</pre> If yes, the default actions file (named <code>actions</code>) contains the following lines: <pre>/usr/openv mail /usr/* repeat 2 /usr/local ignore</pre> Example 2: An action file name for a backup policy <code>production_servers</code> with a schedule name <code>full</code> follows: <pre>actions.production_servers.full</pre> The <code>actions</code> file can contain the following line: <pre>/bin/* repeat</pre> If yes, NetBackup repeats the backup for busy files in the <code>/bin</code> directory.

About the logs directory on UNIX

During busy file processing NetBackup creates a number of files under the `busy_files/logs` directory. These files contain status and diagnostic information. NetBackup derives the names of these files from the policy name, schedule name, and process ID (PID) of the backup.

NetBackup creates the following logs:

- Busy file log

NetBackup records the names of any busy files in the busy file log. The name of the busy file log has the following form:

```
policy_name.schedule_name.PID
```

- Diagnostic log file

NetBackup generates a log file that contains diagnostic information. The name of the log file has the following form:

```
log.policy_name.schedule_name.PID
```

- Retry log file

NetBackup also generates a retry file that contains diagnostic information that is recorded when the repeat option is specified. The name of the retry file has the following form:

```
policy_name.schedule_name.PID.retry.retry_count
```

Where *retry_count* starts at zero and increases by one every time a backup is repeated. Processing stops when *retry_count* is one less than the number that is specified by the `repeat` option.

Example:

To service busy file backup requests, the administrator defined a policy named `AAA_busy_files` that has a user backup schedule named `user`. A scheduled backup is initiated with the policy named `production_servers`, schedule named `full`, and PID of 1442.

If busy files are detected, NetBackup generates the following files in the `/usr/opensv/netbackup/busy_files/logs` directory:

```
production_servers.full.1442
  log.production_servers.full.1442
```

If the actions file has the repeat count set to 2, NetBackup generates the following files:

```
production_servers.full.1442.retry.0
AAA_busy_files.user.10639
log.AAA_busy_files.user.10639
```

If a repeat backup is attempted, NetBackup generates the following files:

```
production_servers.full.1442.retry.1
AAA_busy_files.user.15639
log.AAA_busy_files.user.15639
```

Recommended changes for modifying bpend_notify_busy on UNIX

The administrator can modify busy file processing by changing the `bpend_notify_busy` script.

The only recommended changes are as follows:

- Changing the `RETRY_POLICY` and `RETRY_SCHED` variables from `NONE` to the busy file backup policy name and schedule name.
- Remove the files in the logs directory after busy file processing (these logs are not removed automatically):
 - At the end of the `busy_files()` function, add the following command:

```
/bin/rm -f $LOG_FILE
```

- After the call to the `busy_files()` function in main, add the following commands:

```
/bin/rm -f $BUSYFILELOG
/bin/rm -f $RETRY_FILE
```

About specifying the locale of the NetBackup installation

NetBackup applications can display a wide range of international date and time formats as determined by the locale of the installation. To help ensure consistency among the applications, NetBackup uses a single, configurable source to define the locale conventions.

The `install_path\VERITAS\msg\LC.CONF` file (on Windows) and the `/usr/openv/msg/.conf` file (on UNIX) 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:mn:ss /mm/dd/yyyy
```

```
TL 2 ov :hh:mn:ss/mm/dd/yyyy
```

Note that `C` and `ov` are synonymous.

- **TM Lines**

The `TM` lines define a mapping from unrecognized locales to those supported by NetBackup, as defined by the `TL` lines.

The third field of the `TM` lines 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).

About the Shared Storage Option

The Shared Storage Option allows multiple NetBackup servers (both NetBackup servers and NDMP hosts) 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 requires appropriate hardware connectivity. For example, a fibre channel hub.

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


The Shared Storage Option is a separately licensed and purchased NetBackup option that lets you share a specific number of physical drives.

About Shared Storage Option components

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

[Figure 2-3](#) shows an example of a shared drive configuration.

Figure 2-3 Shared Storage Option example



The following items describe the NetBackup components for the Shared Storage Option example in [Figure 2-3](#).

- The master server hosts the Enterprise Media Manager (EMM) service. It's the device allocation host.
See [About the device allocation host](#).
- Host A:
 - Is a NetBackup media server that runs the Automatic Volume Recognition (avrd) process, the NetBackup Device Manager service (ltid), and the NetBackup Volume Manager (vmd) service.
 - Is connected to drives DRV1 and DRV2 through SAN hardware.
 - Is the first host in the environment to come online with a non-zero scan ability factor. Therefore, it's the initial scan host for its drives.
See [About scan hosts](#).
- Host B:
 - Is a NetBackup media server that runs the Automatic Volume Recognition (avrd) process, the NetBackup Device Manager service (ltid), and the NetBackup Volume Manager (vmd) service.
 - Is connected to drives DRV1 and DRV2 through SAN hardware.

- Controls the robotics. Except for ACS robot types, only one robot control host exists for each robot.

For a process flow diagram of Shared Storage Option components, see the *NetBackup Logging Reference Guide*:

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

About the device allocation host

The NetBackup Enterprise Media Manager (EMM) service allocates devices for Shared Storage Option jobs and tasks. The EMM service runs on the NetBackup master server. The host that allocates devices is also known as the device allocation host.

About SSO and the NetBackup EMM service

To coordinate network-wide allocation of tape drives, the NetBackup Enterprise Media Manager (EMM) service manages all shared tape requests in a shared drive environment. EMM responds to requests from a single NetBackup master server for its corresponding media servers, and NetBackup SAN media servers within a single NetBackup domain.

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 service.

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 `nbermmcmd` command. For more information, see the *NetBackup Commands Reference Guide*.

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

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 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 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's 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 `shared_drive_notify` script resides in the following directory:

- On Windows: `install_path\VERITAS\Volmgr\bin`
- On UNIX/Linux: `/usr/opensv/volmgr/bin/shared_drive_notify`

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 and TLD.
- **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.

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

Shared Storage Option terms and concepts

[Table 2-7](#) describes the terms and the concepts relevant to understanding the Shared Storage Option.

Table 2-7 Shared Storage Option terms and concepts

Term	Definition
Backup Exec Shared Storage Option	The NetBackup Shared Storage Option is not the same as the Veritas 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. Veritas 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

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

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 UNIX or Linux servers, you may have to modify configuration files, such as the sg driver on Solaris systems.
 - 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.
- On UNIX and Linux servers, create any device files that are needed. Depending on the operating system, a reconfiguration system start (`boot -r`) may create these files automatically.

Create the device files for each drive; use the Fibre Channel LUNs of the drives and adapters in the device file names. Add the names of the device files to your notes to complete the correlation between device files and physical drive location. Use the *NetBackup Device Configuration Guide* and the man pages that are available with the operating system.

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

- On UNIX and Linux servers, customize the operating system by modifying the appropriate system configuration files. This task requires knowledge of the system files that use the Shared Storage Option environment and their formats. For example, on Sun Solaris systems you may need to modify the `sg`, `st`, and HBA driver files.

Modify the HBA driver files to bind Fibre Channel devices (WWN) to a specific target ID. For procedures, see the operating system documentation.

- 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 the 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*.

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

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 `bpcintcmd` utility to resolve IP addresses into host names. For more information, see the *NetBackup Troubleshooting Guide* and the *NetBackup Commands Reference Guide*, available through the following URL:
<http://www.veritas.com/docs/DOC5332>

- 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*, available at the following URL:

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

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 110.

Configuring Shared Storage Option devices in NetBackup

See “Configuring Shared Storage Option devices in NetBackup” on page 110.

About adding Shared Storage Option configuration options

See [“Configuring Shared Storage Option devices in NetBackup”](#) on page 110.

About configuring NetBackup storage units and backup policies

See [“About configuring NetBackup storage units and backup policies”](#) on page 111.

About configuring SSO in NetBackup

Veritas 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

Veritas 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**.

Note: You must restart the NetBackup Device Manager (ltd) on all the servers that share tape drives whenever you perform the following actions:

- Configure the shared drives to a newly added media server.
 - Add or remove the shared drives paths.
-

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 120.

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.

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).

For more information, see the *NetBackup Administrator's Guide, Volume I*:

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

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, Veritas recommends that you use the Device Configuration Wizard. The wizard verifies your configuration.
- If you have non-serialized devices, see the Veritas support site for tech note TECH31764, "Verifying a Shared Storage Option (SSO) Configuration with Non-Serialized Devices. It describes how to verify your configuration.
- 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:

- On Windows:

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

- On UNIX/Linux:

```
usr/opensv/volmgr/bin/scan  
usr/opensv/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.

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
      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/nh3c0t510` and `/dev/st/nh3c0t110`.
 - 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/nh3c0t510` serial number is PXA37S3261.
 The drive `/dev/st/nh3c0t110` 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 **Device Monitor** in the **NetBackup Administration Console** 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*:

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

Table 2-8 describes information you can glean from the **Device Monitor**.

Table 2-8 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: ■ 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: ■ 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 117.

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 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*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

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*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>
- 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 the Shared Storage Option, enter the Shared Storage Option license 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 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*, available at the following URL:

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

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. Veritas 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 Veritas support Web site:

<http://www.netbackup.com/compatibility>

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.

Q. How can I reserve or release shared devices?

A. 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.

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.

On Windows, the pathname is `install_path\Volmgr\vm.conf`.

On UNIX, the pathname is `/usr/opensv/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 136.

ACS_mediatype entry in `vm.conf`

The following configuration entry applies to NetBackup servers:

```
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*:

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

ACS_SEL_SOCKET entry in `vm.conf`

The following configuration entry applies to NetBackup servers:

```
ACS_SEL_SOCKET = socket_name
```

By default, `acssel` 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_CSI_HOSTPORT entry in `vm.conf` (on UNIX)

The following configuration entry applies to NetBackup servers:

```
ACS_CSI_HOSTPORT = ACS_library_software_hostname socket_name
```

The 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.

The valid values for `socket_name` are 1024 - 65535 and 0. The value must match the value on the ACSLS server for the port that the CSI uses for inbound packets.

If 0 (zero), NetBackup uses the previous behavior of CSI and `acsssi` (no specific ports).

This entry specifies the port where the `acsssi` process sends its ACSLS requests on the ACSLS server. The ACSLS CSI must use this port to accept inbound ACSLS requests from `acsssi` processes.

This entry, the `ACS_SSI_INET_PORT` entry, and the `ACS_TCP_RPCSERVICE` entry are commonly used with firewall implementations. With these three entries in the `vm.conf` file, TCP connections use the designated destination ports. Note that TCP source ports are not restricted.

See [“ACS_SSI_INET_PORT entry in vm.conf \(on UNIX\)”](#) on page 122.

See [“ACS_TCP_RPCSERVICE / ACS_UDP_RPCSERVICE entry in vm.conf \(on UNIX\)”](#) on page 124.

For example, a NetBackup media server has two ACSLS servers (`ACSL_1` and `ACSL_2`) behind firewalls. Both servers listen for queries on port 30031 and the firewall allows traffic through this port.

The `vm.conf` entries are as follows:

```
ACS_TCP_RPCSERVICE
ACS_CSI_HOSTPORT = ACSLS_1 30031
ACS_CSI_HOSTPORT = ACSLS_2 30031
ACS_SSI_INET_PORT = ACSLS_1 30032
ACS_SSI_INET_PORT = ACSLS_2 30033
```

Each `acsssi` process sends queries to the respective ACSLS server's port 30031, and the ACSLS server is configured to listen for queries on this port.

ACS_SSI_HOSTNAME entry in vm.conf

The following configuration entry applies to NetBackup servers:

```
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_INET_PORT entry in vm.conf (on UNIX)

The following configuration entry applies to NetBackup servers:

```
ACS_SSI_INET_PORT = ACS_library_software_hostname socket_name
```

The 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.

The `socket_name` entry specifies the port that `acsssi` uses for incoming ACSLS responses. Valid values are 1024 - 65535 and 0. This value must be unique for each `acsssi` process.

A value between 1024 - 65535 indicates the number to be used as the TCP port on which `acsssi` accepts ACSLS responses.

0 (zero) indicates that the previous behavior (allow the port to be dynamically allocated) should remain in effect.

This entry, the `ACS_CSI_HOSTPORT` entry, and the `ACS_TCP_RPCSERVICE` entry are commonly used with firewall implementations. With these three entries in the `vm.conf` file, TCP connections use the designated destination ports. Note that TCP source ports are not restricted.

See “[ACS_CSI_HOSTPORT entry in vm.conf \(on UNIX\)](#)” on page 121.

See “[ACS_TCP_RPCSERVICE / ACS_UDP_RPCSERVICE entry in vm.conf \(on UNIX\)](#)” on page 124.

For example, a NetBackup media server has two ACSLS servers (ACSL_1 and ACSLS_2) behind firewalls. Ports 30032 and 30033 have been opened in the firewall for `acsssi` to ACSLS server communication.

The entries would be as follows:

```
ACS_TCP_RPCSERVICE
ACS_SSI_INET_PORT = ACSLS_1 30032
ACS_SSI_INET_PORT = ACSLS_2 30033
ACS_CSI_HOSTPORT = ACSLS_1 30031
ACS_CSI_HOSTPORT = ACSLS_2 30031
```

The NetBackup media server starts two `acsssi` processes. One listens for ACSLS_1 responses on port 30032, and the other listens on port 30033 for responses from ACSLS_2.

ACS_SSI_SOCKET entry in vm.conf

The following configuration entry applies to NetBackup servers:

```
ACS_SSI_SOCKET = ACS_library_software_hostname socket_name
```

The 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.

ACS_TCP_RPCSERVICE / ACS_UDP_RPCSERVICE entry in `vm.conf` (on UNIX)

The following configuration entries apply to NetBackup servers:

```
ACS_TCP_RPCSERVICE
ACS_UDP_RPCSERVICE
```

These entries specify the method over which `acsssi` communicates with ACSLS servers: TCP or UDP.

Only one entry should be entered into `vm.conf`. NetBackup uses UDP if both entries are found or neither entry is found.

For `acsssi` firewall support, `ACS_TCP_RPCSERVICE` must be entered in `vm.conf`.

See [“ACS_CSI_HOSTPORT entry in `vm.conf` \(on UNIX\)”](#) on page 121.

See [“ACS_SSI_INET_PORT entry in `vm.conf` \(on UNIX\)”](#) on page 122.

ADJ_LSM entry in `vm.conf`

The following configuration entry applies to NetBackup servers:

```
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 2-4 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 2-4 Pass-through example

Interconnections for
Robot 700



API_BARCODE_RULES entry in vm.conf

The following configuration entry applies to NetBackup servers:

`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` in NetBackup versions 8.0 and earlier

Note: This entry is not applicable for NetBackup 8.1 or later versions.

Starting with NetBackup 8.1, the Volume Manager service (`vmd`) validates all requests from remote hosts against the primary servers and the media servers for the domain, as known to `nbemm`. The `vm.conf` entries are no longer used for this determination, and requests from hosts in other NetBackup domains are no longer allowed.

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, Veritas 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 (`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. This entry is read and interpreted on the host on which the NetBackup Device Manager (`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 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 TLD robots that post a unit attention when their MAP is opened.

Veritas 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 Windows, NetBackup reports PEND if the drive reports Busy when a volume is unmounted. Use this entry to minimize the display of this misleading status.

The 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`

Note: This entry affects tape drive cleaning requests as well as tape mount and tape dismount requests.

Use this entry to specify the following:

- How long NetBackup waits for a drive to be cleaned before it removes the request from the queue.
- How long NetBackup waits for a tape to be mounted or dismounted before it removes the request from the queue.

Unprocessed requests to clean a drive or to mount or dismount a tape are removed from the queue after 30 minutes.

```
CLEAN_REQUEST_TIMEOUT = minutes
```

The *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`

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

```
CLUSTER_NAME = cluster_alias
```

See [“Host name precedence in the `vm.conf` file”](#) on page 137.

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
```

The default is 30 days. A value of zero means that the logs are not deleted. This entry does not affect the debug logs that Unified Logging creates.

Information about Unified Logging is available in the [NetBackup Logging Reference Guide](#).

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 `vmdaemon` 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.

INVENTORY_FILTER entry in `vm.conf`

The following configuration entry applies to NetBackup servers:

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

Used to filter the robot inventory results in ACS 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 only be ACS.

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

mode is `BY_ACS_POOL` for ACS.

Example

```
INVENTORY_FILTER = ACS 0 BY_ACS_POOL 4 5
```

MAP_ID entry in vm.conf

The following configuration entry applies to NetBackup servers:

```
MAP_ID = robot_num map_ID
```

Use this entry to configure the default media access port (MAP) to use to eject media from the 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 124.

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 the 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`

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

```
MM_SERVER_NAME = host_name
```

See [“Host name precedence in the `vm.conf` file”](#) on page 137.

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.

On UNIX, if random ports are not specified in the NetBackup configuration, specify `RANDOM_PORTS = NO` in the `vm.conf` file.

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

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

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** in the **NetBackup Administration Console: NetBackup Management > Host Properties > Double-click on master server > Port Ranges > Use random port assignments**.

REQUIRED_INTERFACE entry in vm.conf

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

```
REQUIRED_INTERFACE = host_name
```

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 137.

SERVER entry in vm.conf in NetBackup versions 8.0 and earlier

Note: This entry is not applicable for NetBackup 8.1 or later versions.

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 `vmid` 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.

The following configuration entry applies to NetBackup servers:

```
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 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`

The following configuration entry applies to NetBackup servers:

```
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`

The following configuration entry applies to NetBackup servers:

```
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.

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 Veritas 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
```

How to access media and devices on other hosts

Note: This topic is not applicable for NetBackup 8.1 or later versions.

Starting with NetBackup 8.1, the Volume Manager service (`vmc`) validates all requests from remote hosts against the primary servers and the media servers for the domain, as known to `nbemm`. The `vm.conf` entries are no longer used for this determination, and requests from hosts in other NetBackup domains are no longer allowed.

For NetBackup to access media and device management functionality on a remote NetBackup host, you may need to add a `SERVER` entry to the `vm.conf` file on the remote host.

The `SERVER` entries are used in the NetBackup `bp.conf` and `vm.conf` files for security. You can add the entries that allow only specific hosts to access those capabilities remotely.

If the `vm.conf` file on a remote host contains no `SERVER` entries, a host can manage media and devices on the remote host if it's added to the `bp.conf` file of the server you logged into. You do not need to add a `SERVER` entry to the `vm.conf` file.

If the `vm.conf` file on a remote host contains any `SERVER` entries, add a `SERVER` entry for the host on which the **NetBackup Administration Console** is running (the server you logged into) to that `vm.conf` file.

Assume that you have three hosts named `eel`, `yak`, and `shark`. You want to centralize device management on host `shark` and also permit each host to manage its own devices.

The following example scenario applies:

- The `vm.conf` file on `shark` contains the following:

```
SERVER = shark
```

The `vm.conf` file on `shark` does not require any additional `SERVER` entries, because all device management for `shark` is performed from `shark`.

- The `vm.conf` file on `eel` contains the following, which lets `eel` manage its own devices and permits `shark` to access them:

```
SERVER = eel  
SERVER = shark
```

- The `vm.conf` file on `yak` contains the following, which lets `yak` manage its own devices and permits `shark` to access them:

```
SERVER = yak  
SERVER = shark
```

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 name of the host in the Server host properties of the primary server.
- `gethostname()` `name`.

Holds Management

This chapter includes the following topics:

- [About Holds Management](#)
- [Creating a hold](#)
- [Viewing hold details](#)
- [Adding a backup image to an existing hold](#)
- [Releasing a hold](#)

About Holds Management

NetBackup provides an option to put backup images on hold. The holds mechanism lets you retain the backup images for as long as you need without altering the expiration date.

You can manage the holds by using the command-line interface. You can perform the following:

- Create a hold.
See [“Creating a hold”](#) on page 140.
- View the list of holds.
See [“Viewing hold details”](#) on page 140.
- Add one or more backup images to an existing hold.
See [“Adding a backup image to an existing hold”](#) on page 141.
- Release a hold from the backup image.
See [“Releasing a hold”](#) on page 141.

Note: All hold operations except listing are audited.

Creating a hold

You can create a hold on one or more backup images by using the `nbholdutil -create` command.

Caution: Creating a hold on backup images may disrupt new backups from completing. Storage may fill up if previous backups are not automatically expired.

Note: When you retry a failed Hold creation, an empty hold is created if the backup images have expired between the initial hold and the retry.

To create a hold

The `nbholdutil -create` command lets you create a hold for a backup image.

On a command prompt on the NetBackup master server, enter `nbholdutil -create` with appropriate options and elements. For example:

```
nbholdutil.exe -create -holdname legal_case1 -backupid  
win81.sky.com_1307425938 -allcopy
```

This command creates a hold called `legal_case1`. The backup image ID is `win81.sky.com_1307425938`. You must provide either the `-allcopy` option or the `-primarycopy` option. The `-allcopy` operation indicates that the hold includes all copies of the selected backup image. The `-primarycopy` option indicates that the hold includes only the primary copy of the selected backup image.

For more information about related command options, see the *Veritas NetBackup Commands Reference Guide*.

To display help information about the command and its options, enter `nbholdutil -help [-option]`

Viewing hold details

You can view the list of holds by using the `nbholdutil -list` command.

To view hold details

On a command prompt on the NetBackup master server, enter the `nbholdutil -list` command with appropriate options and elements. For example:

```
nbholdutil.exe -list
```

When you upgrade NetBackup to version 7.7, the legal holds are converted to user holds, which can be managed by using the `nbholdutil` command.

If the hold name of a legal hold is same as a user hold, all the hold names are renamed as follows:

- The legal hold names are suffixed with `_1`. For example, `hold_1`. The number 1 in the hold name denotes that it was a legal hold before conversion.
- The user hold names are suffixed with `_3`. For example, `hold_3`. The number 3 in the hold name denotes that it is a user hold.

For more information about related command options, see the *Veritas NetBackup Commands Reference Guide*.

To display help information about the command and its options, enter `nbholdutil -help [-option]`

Adding a backup image to an existing hold

You can add one or more backup images to an existing hold by using the `nbholdutil -add` command.

To add a backup image to an existing hold

On a command prompt on the NetBackup master server, enter the `nbholdutil -add` command with appropriate options and elements. For example:

```
nbholdutil.exe -add -holdname hold123 -reason "Reason1" -backupid
win81.sky.com_1307425938 -primarycopy
```

This command adds primary copy of the backup image `win81.sky.com_1307425938` to the existing hold with hold ID equal to `hold123`.

For more information about related command options, see the *Veritas NetBackup Commands Reference Guide*.

To display help information about the command and its options, enter `nbholdutil -help [-option]`

Releasing a hold

You can release holds by using the `nbholdutil -delete` command.

Note: A backup image expires as per the expiry date when all the holds that include that backup image are released.

To release a hold

On a command prompt on the NetBackup master server, enter the `nbholdutil -delete` command with appropriate options and elements. For example:

```
nbholdutil.exe -delete -holdname legal_case1 -force -reason
Legal_Case1 resolved
```

This command releases a hold that is called `legal_case1`. For more information about related command options, see the *Veritas NetBackup Commands Reference Guide*.

The command `nbholdutil -delete` lets you release a hold.

Menu user interfaces on UNIX

This chapter includes the following topics:

- [About menu user interfaces](#)
- [About the tpconfig device configuration utility](#)
- [About the NetBackup Disk Configuration Utility](#)

About menu user interfaces

NetBackup provides several menu user interfaces on UNIX systems to help manage some NetBackup functionality:

- See [“About the tpconfig device configuration utility”](#) on page 144.
- See [“About the NetBackup Disk Configuration Utility”](#) on page 152.

These utilities are alternatives to using the **NetBackup Administration Console**. The terminology, general concepts, and results are the same regardless of the administration method that is used.

Note: Many NetBackup processes set an upper limit on the number of concurrently open file descriptors allowed by the process. That limit is inherited by the notify scripts run by the process. In the rare event that a command invoked by a notify script requires many additional file descriptors, the script must increase the limit appropriately before invoking the command.

About the `tpconfig` device configuration utility

UNIX systems only.

The NetBackup `tpconfig` device configuration utility is a character-based, menu-driven interface to configure robots, drives, and logon credentials. It can be used at any terminal (or terminal emulation window) for which `termcap` or `terminfo` is defined.

The NetBackup command utilities are alternatives to the **NetBackup Administration Console**. The terminology, general concepts, and results are the same regardless of which method you use.

After you configure devices, you can use the **NetBackup Administration Console** to configure volumes.

The following list describes the attributes of device configuration and how to use the `tpconfig` utility to configure those attributes.

The `tpconfig` device configuration utility attributes are as follows:

- **Robot number**

You assign a robot number when you add a robot to the configuration. `tpconfig` prompts you to enter a number or accept the next available robot number that appears. This number identifies the robot in displays and listings, and it follows the robotic type in parentheses.

If you configure robots on multiple systems, robot numbers must be unique. If you connect drives from a robot to multiple systems, specify the same robot number for the robot on all systems.

- **Robotic control path**

For most robots, you or the operating system creates this path in the `/dev` directory when you add a robot to the configuration. When the `tpconfig` utility prompts you, enter the path to the robotic control as found in the `/dev` directory. If the entries do not exist, more information is available.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

The control path to a robot may be on another host. If so, enter the host name of the host instead of a path. When you define a robot that another host controls by another host, the robot number must be the same on both hosts.

Information about how to configure robotic control is available.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

- **Host name**

You must specify a host name in the following cases:

- When you add an ACS robot, enter the name of the host on which the ACS Library Software resides instead of a robotic control path.
- When you add TLD robots that have robotic control on another host, you are prompted for the name of that host.
- **No rewind on close device name**

You specify an no rewind on close device name when you add a drive. Usually the letter n precedes or follows the device name. If the device name entries do not exist, you must create them.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

In `tpconfig` displays and `tpconfig` output, the device names are shown under the heading `DrivePath`.
- **Drive status**

Drive status indicates whether NetBackup considers a drive available. You specify the initial drive status when you add a drive to the configuration. You can change the status. To do so, use the Update option of the Drive Configuration menu in `tpconfig` (ensure that the device `daemonltid` is not active). If the device `daemon ltid` is active, use the Administration Console Device Monitor or the `vmopr cmd` command.

About the tpconfig utility menu

The **Device Configuration Utility** menu contains the following information:

Device Management Configuration Utility

- 1) Drive Configuration
- 2) Robot Configuration
- 3) Credentials Configuration
- 4) Print Configuration
- 5) Help
- 6) Quit

Enter option:

[Table 4-1](#) describes the main menu selections.

Table 4-1 `tpconfig` main menu selections

Menu choice	Description
Drive Configuration	Opens a menu to add, delete, or update drive definitions; list definitions of drives and robots; or configure drive paths.
Robot Configuration	Opens a menu to add, delete, or update robot definitions or list definitions of drives and robots
Credentials Configuration	Opens a menu to add, delete, update, or list credentials for the following: ■ NDMP filer ■ Disk array ■ OpenStorage server ■ Virtual machine
Print Configuration	The List Configuration commands on subsequent menus let you display the current configuration on the screen or write it to a file. If you specify the <code>-d</code> option only on the <code>tpconfig</code> command, <code>tpconfig</code> writes the current configuration to stdout (the screen) without invoking the menus. Other command options are available. Run <code>tpconfig -help</code>. See the <i>NetBackup Commands Reference Guide</i>, available at the following URL: http://www.veritas.com/docs/DOC5332
Help	Online Help is available on the main menu and most submenus.
Quit	Terminates the utility and returns you to the UNIX prompt.

You can return to the main menu from anywhere in the utility by entering **Ctrl C** or by using the **Escape** key.

Starting the `tpconfig` device configuration utility

Several methods exist to start the `tpconfig` utility.

Note: If the Media Manager device daemon is running, stop it by using the `stopltid` command.

To start tpconfig from a UNIX shell

- ◆ Enter the following command in a UNIX shell (you must have root user privileges):

```
/usr/opensv/volmgr/bin/tpconfig
```

Adding robots

When you configure robots and drives, first add the robots by using the **Robot Configuration** menu. Then add the drives by using the **Drive Configuration** menu.

To change standalone drives to robotic, use the **Update** option of the **Drive Configuration** menu.

See [“Updating a drive configuration”](#) on page 149.

To add a robot

- 1 Select the **Robot Configuration** menu.
- 2 Select the **Add** option.
- 3 From the list of possible robot types, select the one you want to add.
- 4 Enter a robot number that you know is unused or accept the default robot number.
- 5 Indicate where the robotic control for the library is by entering the device file path or library name. The **Help** option on the **Robot Configuration** menu has examples of typical path names.
- 6
 - If robotic control is on another host, enter that host name.
For an ACS robot, enter the name of the ACS library software host.
 - If robotic control is on this host, enter the device file path or library name.
The **Help** option on the **Robot Configuration** menu has examples of typical path names.
For an ACS robot, enter the name of the ACS library software host.
- 7 If no conflicts are detected with the new configuration, a message appears to indicate that the robot was added.

Adding drives

Use the following procedure to add a drive.

To add a drive

- 1 Select the **Drive Configuration** menu.
- 2 Select the **Add** option.

- 3 From the list of possible drive types, select the one you want to add.
- 4 Enter the no rewind on close device path as shown in the `/dev` directory.
The **Help** option on the **Drive Configuration** menu has examples of typical path names.
- 5 Enter the drive status (Up or Down).
- 6 If a robot exists to which you can add the drive, specify whether to add the drive to the robot. Alternatively, you can configure the drives as a standalone drive.

If there are no robots to which you can add the drive, `tpconfig` automatically adds the drive as a standalone drive.

If you add a drive to a robot and more than one possible robot exists, enter the number of the robot that controls the drive.

Depending on the type of robot, you may also be prompted to add the robot drive number.
- 7 For a drive in an ACS robot, you are prompted for four drive identifiers.

More information on ACS robots is available.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

More information is available.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>
- 8 Type a drive name or press the **Enter** key to use the default drive name.

If you use the shared drives option, all hosts that share the same physical drive must use the same name for the drive. Descriptive drive names are recommended.

Updating a robot configuration

Use the following procedure to change the robot number or the robotic control path.

To change the robot number or the robotic control path

- 1 On the main menu, select **Robot Configuration**.

If only one robot is configured, you do not have to select **Update** or enter the robot number. If only one robot is configured, skip to step 4.
- 2 On the **Robot Configuration** menu, choose **Update**.

- 3 Enter the number of the robotic library you want to change.
- 4 Enter a new robot number to replace the existing robot number or press **Enter** to retain the current robot number.

You are prompted to enter robotic control information. The actual prompts depend on the type of robotic library you update.
- 5 Enter the appropriate robotic control path or name of the server that hosts the robot.

Updating a drive configuration

You can change information for a drive (for example, you can add it to a robot).

To change information for a drive

- 1 On the main menu, select **Drive Configuration**.
- 2 On the **Drive Configuration** menu, choose **Update**.
- 3 Enter the name of the drive you want to update.

The current drive information is displayed, followed by prompts to change each field.
- 4 Enter a new value or press **Enter** to retain the existing value.

One of the prompts asks if you want to configure the drive in a robot. If you do, `tpconfig` adds the drive immediately or gives you the opportunity to choose from any existing robot of the appropriate type.

When you have responded to all prompts, a revised Drive Information display appears, along with the following prompt:


```
Are you sure you want to UPDATE drive name xxxxx? (y/n) n:
```
- 5 Answer yes by pressing **y**.

Deleting a robot

Use the following procedure to delete a robot.

To delete a robot

- 1 On the main menu, select **Robot Configuration**.

If only one robot is configured, you do not have to select **Update** or enter the robot number. If only one robot is configured, skip to step 4.
- 2 On the **Robot Configuration** menu, choose **Delete**.

- 3 If more than one robot is configured, enter the number of the robot to delete.
- 4 Enter **y** to delete the robot.
If you respond with **n**, press any key to return to the **Drive Configuration** menu.

Deleting a drive

Use the following procedure to delete a drive.

To delete a drive

- 1 On the main menu, select **Drive Configuration**.
- 2 In the **Drive Configuration** menu, select **Delete**.
- 3 Enter the name of the drive you want to delete:
- 4 Enter **y** to delete the drive.
If you respond with **n**, press any key to return to the **Drive Configuration** menu.

Configuring drive paths

Use the following procedures to configure and manage drive paths.

To display the drive path menu

- 1 From the **Drive Configuration** menu, select **Drive Path Configuration**.
- 2 Enter a drive name.

To add a drive path

- 1 Select **Add** from the **Drive Path Configuration** menu.
- 2 Enter a valid drive path.
- 3 Specify the drive status on the path.

The drive can be Up, Down, or Disabled for the path.

To delete a drive path

- 1 Select **Delete** from the **Drive Path Configuration** menu.
- 2 Enter the drive path to delete.

To update a drive path

- 1 Select **Update** from the **Drive Path Configuration** menu.
- 2 Enter the drive path to update.

- 3 Specify a new drive path or press **Enter** to update the status of the drive path.
- 4 A prompt similar to the following is displayed:
- 5 Enter the path status.

Configuring host credentials

You can add, delete, update, or configure the following default host credentials:

- NDMP filer
You can add the credentials for a specific filer on a specific server. You also can add credentials to be used for all NDMP Filers by all NetBackup servers.
- OpenStorage server
- Virtual machine

To configure host credentials

- 1 On the main menu, select **Credentials Configuration**.
- 2 On the credentials menu, select the type of credential you want to configure.

Please select the type of host you are trying to configure:

```
1) (N)dmp Filer
2) (O)penStorage Server
3) (V)irtual Machine
```

- 3 Select an option at the specific credentials menu and follow the prompts.

Displaying and writing the device configuration

You can display or write out the current configuration from every menu in `tpconfig`.

To display the configuration from the main menu

- 1 Press 4) `Print Configuration`.
- 2 Press **Enter**.

To display the configuration from a submenu

- ◆ Select the `List Configuration` option by pressing the corresponding number.

To write the current configuration to a file

- 1 On the main menu, press 4) `Print Configuration`.
- 2 Enter the file name.

To write the current configuration to standard output

- ◆ Enter the following command in a UNIX shell:

```
tpconfig -d
```

About the NetBackup Disk Configuration Utility

The NetBackup Disk Configuration Utility is a character-based, menu-driven interface that lets you configure and manage disk storage entities. Use this utility for optional NetBackup products that use disk pool storage. It can be used at any terminal (or terminal emulation window) for which `termcap` or `terminfo` is defined.

The NetBackup command utilities are alternatives to the **NetBackup Administration Console**. The terminology, general concepts, and results are the same regardless of which method you use.

After you configure the disk storage, you also can configure a disk storage unit by using a UNIX utility.

Managing OpenStorage servers and disk pools

Use the OpenStorage Disk Management menu to configure and manage disk storage for the OpenStorage Disk Storage Unit Option.

To manage OpenStorage servers and disk pools

- 1 From the NetBackup disk configuration utility main menu, press **o** (OpenStorage Disk Management) to display the OpenStorage Disk Management menu.

The menu contains the following information:

```
OpenStorage Disk Management
```

```
-----
```

- a) Add Storage Server
- u) Update Storage Server
- r) Remove Storage Server
- v) View Storage Server
- g) Get Storage Server Configuration
- e) Engage Storage Server Configuration

- p) Preview Disk Volumes
- n) Create New Disk Pool

- t) Take Disk Pool Inventory
- m) Merge Two Disk Pools
- c) Change Disk Pool
- s) Change Disk Pool State
- w) Update Disk Pool Replication Properties From Storage Server
- k) Add Volumes To Disk Pool
- d) Delete Disk Pool
- l) List Disk Pools

- h) Help
- q) Quit Menu

```
ENTER CHOICE:
```

- 2 Select a menu option and follow the prompts to configure and manage OpenStorage.

Managing global disk attributes

Use the Global Disk Management Attributes menu to configure and manage disk storage attributes for all disk pool features.

To manage global disk attributes

- 1 From the NetBackup disk configuration utility main menu, press **g** (Global Disk Management Attributes) to display the Global Disk Management Attributes menu.

The menu contains the following information:

```
Global Disk Management Attributes
```

```
-----
```

- l) List Global Disk Management Attributes
- s) SharedDisk SCSI Persistent Reservation

- h) Help
- q) Quit Menu

```
ENTER CHOICE:
```

- 2 Select a menu option and follow the prompts to configure and manage attributes.

Reference topics

This chapter includes the following topics:

- [Host name rules](#)
- [About reading backup images with nbtar or tar32.exe](#)
- [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 and device management processes](#)
- [About Tape I/O commands on UNIX](#)

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 82.

On Windows:

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.

Windows systems do not permit NetBIOS computer names that exceed 15 characters. It is recommended that you use a computer name that is less than 15 characters. The 16th character of a NetBIOS computer name is reserved to identify the functionality that is installed on the registered network device. Note that the interactive NetBackup commands such as `bpbntat -login` truncate the host name values to use the first 15 characters. If you need to provide the host name along with the NetBackup domain name, you should consider this truncation.

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: (On Windows) Veritas recommends that you do not change the host name of a NetBackup server. 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 the topics that address how NetBackup stores and uses host names.

Table 5-1 How NetBackup stores and uses host names

Topic	Description
Server and client names on UNIX servers and clients	On both UNIX servers and clients, the <code>SERVER</code> entries in the <code>bp.conf</code> file define the NetBackup servers that are allowed access. The first <code>SERVER</code> entry identifies the primary server. The first <code>SERVER</code> entry indicates the server to which client requests are made. For this reason, the <code>SERVER</code> name must be one by which all clients can connect to the server. If more than one <code>SERVER</code> entry exists, the additional entries identify other NetBackup servers that can initiate scheduled backups on the client. The <code>bp.conf</code> file must have multiple <code>SERVER</code> entries if you configure remote media servers. The NetBackup Request daemon (<code>bprd</code>) and NetBackup Database Manager daemon (<code>bpdbm</code>) do not run on any server other than a primary. When a client makes a list or restore request to the server, the NetBackup client name is used to determine whether to allow the operation. (The client name as specified on the client.) The client name that is used is usually the <code>CLIENT_NAME</code> from the <code>bp.conf</code> file of the client. Or, the client name can be the actual host name of the client if not in the <code>bp.conf</code> file. Alternate client restores can use the name that is specified through the user interface or with a parameter on the <code>bprestore</code> command. For a successful request, the client name must match the name that is specified for the client in the NetBackup configuration on the server. The only exception to this rule is if the server is configured to allow alternate client restores.
Host names on Windows servers and PC clients	Windows NetBackup servers and clients also have <code>SERVER</code> and <code>CLIENT_NAME</code> settings. On these systems, specify server and client settings in the NetBackup Administration Console.

Table 5-1 How NetBackup stores and uses host names (*continued*)

Topic	Description
Policy configuration	(On Windows) The configured name for a client is the host name as it's added to a policy. This name is how the client is identified in the NetBackup configuration. (On UNIX) The configured name for a client is the host name as it's added to a policy. This name is how the client is identified in the NetBackup configuration. NetBackup also adds a <code>CLIENT_NAME</code> entry to a UNIX client's <code>bp.conf</code> file when software is first installed on the client. 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. 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. 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.
Image catalog	A subdirectory in the image catalog is created for a client when a backup is first created for that client. The subdirectory's name is the client's configured name. Every backup for a client has a separate file in this subdirectory. Each of these backup records contains the host name of the server on which the backup was written.
Error catalog	NetBackup uses the 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, <code>servername</code> instead of <code>servername.null.com</code>.)
Catalog backup information	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 primary server to connect to the media server.

Many NetBackup user-defined strings must not contain non-US ASCII characters, including the following:

- Host name (primary server, media server, Enterprise Media Manager (EMM) server, volume database host, media host, client)
- Policy name
- Policy keyword (Windows only)
- **Backup, Archive, and Restore** keyword (Windows only)

- Storage unit name
- Storage unit disk pathname (Windows only)
- Robot name
- Device name
- Schedule name
- Media ID
- Volume group name
- Volume pool name
- Media description
- Vault policy names
- Vault report names
- BMR Shared Resource Tree (SRT) name
- `nbcertcmd` command

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 primary server name change | See "To update NetBackup after a primary server name change" on page 160. |
| To update NetBackup after a client name change | See "To update NetBackup after a client name change" on page 160. |

To update NetBackup after a primary server name change

- 1 On the primary 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 (On UNIX) Create a symbolic link from the client's old image directory to its new image directory. For example,

```
cd /usr/opensv/netbackup/db/images ln -s  
old_client_name new_client_name
```

- 3 (On Windows) 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
```

- 4 (On Windows) Create a directory for the new `client2` in the `\images` directory:

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

- 5 (On Windows) 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 Name Service

In some requests to the primary server, the NetBackup Request Daemon/Service needs to confirm the configured name by which it knows the connecting client host. It does this by using the `getnameinfo` library function, supplied by the operating

system, to query name services. The returned peer name is then compared to the client names known to the primary server, both current clients in a policy and clients for which a backup image is available. If no match is found, then each known client is checked for hostname aliases using the `getaddrinfo` library function to query name services. The returned aliases are then compared to the peer name. This comparison is performed by the `get_ccname` function which will appear in the debug logs for the `bprd` process.

`get_ccname`: determine configured name for client2-nic2.veritas.com

`get_ccname`: configured name is: client2

If `getnameinfo` returns an unexpected peer name or the various `getaddrinfo` calls do not return a matching alias, problems occur.

One possible solution is to reconfigure name services, possibly via the primary server hosts file. Another option is to create a special file in the `altnames` directory on the primary server. The file provides a direct translation of the peer names to the client host names.

On Windows:

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

On UNIX:

```
/usr/opensv/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 peername_from_client_IP_address 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.
- *peername_from_client_IP_address* is the value to translate.
- *client_as_known_by_server* must match the name in the NetBackup configuration on the primary server and must also be known to the media server's network services. It should also be known to the primary server's network services.

Consider the following example:

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

The line specifies that when the primary server receives a request for a configured client name (numeric key 0), the client name `xxxx.eng.aaa.com` is used for peer name `xxxx`.

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

- When `getnameinfo` on the primary server resolves the source IP from the client to `xxxx`, which is not a policy client and does not have any backup images.
- When `getaddrinfo` on the primary for the known policy and backup clients does not return an alias that matches the name `xxxx`.
- The client was configured and named in the NetBackup configuration as `xxxx.eng.aaa.com`.

About reading backup images with nbtar or tar32.exe

NetBackup uses tar-formatted backup images. By using the NetBackup `tar32.exe` on Windows or `nbtar` on UNIX or Linux, NetBackup can understand compressed files, sparse files, long pathnames, and ACL information. It offers features similar to those in `cpio`.

Although non-NetBackup restore utilities that process tar-formatted images can be used to restore files, they provide only limited restore capabilities. You cannot use the NetBackup `tar32.exe` or `nbtar` to extract files from a NetBackup for Windows backup image.

Consequences of using non-NetBackup restore utilities

Non-NetBackup restore utilities do not supply all of the restore capabilities that the NetBackup `/usr/opensv/netbackup/bin/nbtar` provides. Possible problems result.

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

- 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.
- Raw partitions cannot be recovered. (This applies to FlashBackup images as well.)
- 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 restore utilities may have trouble with sparse files and often skip sparse files.
- HP CDFs are restored with non-NetBackup versions of restore utilities. 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 the restore utility. To combine the fragments, the system's `dd` command may be useful.
Another possibility is to use a restore utility on the fragments. To use a restore utility on fragments can allow recovery of any file in the backup other than the one that spanned the 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.

Restoring files with non-NetBackup restore utilities (on UNIX)

This sequence assumes that the media is known to Media Manager and that the tape drive is under Media Manager control.

Before you begin, obtain the following information:

- The media ID of the tape that contains the required backup.
- The tape file number of the backup on the tape.
See the NetBackup **Images on Media** report for this tape.
- The tape type and density.
- The tape pool.

To restore files with a non-NetBackup utility

- 1 Enter the following command:

```
tpreq -m media_id -a r -d density -p poolname -f  
/tmp/tape
```

Where the following is true:

media_id is the media ID of tape that contains the backup.

density is the density of the tape.

poolname is the volume pool to which the tape belongs

- 2 Enter the following command: `mt -f /tmp/tape rew`

- 3 Enter the following command: `mt -f /tmp/tape fsf file_#`

Where the following is true:

file_# is the tape file number of the backup on tape. Determine the tape file number by checking the NetBackup Images on Media report for the tape.

- 4 Enter the following command: `mt -f /tmp/tape fsr`

- 5 Enter the following command:

```
/bin/nbtar -tvfb /tmp/tape blocksize
```

Where the following is true:

blocksize is 64 (assume that the tape is written with 32K blocks)

- 6 Enter the following command: `tpunmount /tmp/tape`

Considerations for file restoration with non-NetBackup restore utilities (on UNIX)

When you restore files with non-NetBackup restore utilities, be aware of the following considerations:

- The file restoration procedure with non-NetBackup utilities does not apply to the encrypted backups that use NetBackup Encryption. Encrypted backups are recoverable. However, the backups cannot be decrypted.
- To determine if a backup is encrypted, run a non-NetBackup restore utility such as `tar -t` before the recovery. The output for an encrypted backup is similar to the following example:

```
erw-r--r-- root/other Nov 14 15:59 2014 .EnCryYpTiOn.388
-rw-r--r-- root/other Oct 30 11:14 2015 /etc/group.10-30
```

Where the *e* at the beginning of line one indicates that the backup is encrypted. (Additional messages appear during recovery.)

- The file restoration procedure with non-NetBackup utilities does not work on the Solaris platform. You cannot use `/usr/sbin/tar` on Solaris to read NetBackup backups. The Solaris `tar` command uses the `ctime` and the `atime` fields differently than other `tar` commands.

When `/usr/sbin/tar` is used to restore backups, directories with large numbers are created at the top level. These directories are from the `ctime` and the `atime` fields being read as pathnames.

You can use `/usr/opensv/netbackup/bin/nbtar` to read the backups on Solaris platforms.

- Steps 1 and 6 from the file restoration procedure with non-NetBackup utilities are optional in a standalone environment. If step 1 is skipped, DOWN the drive and then substitute the /dev path of the drive in place of /tmp/tape in the other steps. Remember to UP the drive when you are done.
See [“To restore files with a non-NetBackup utility”](#) on page 163.

About the files that restores generate

The `nbtar` command and any restore utility that processes tar-formatted images, can generate a number of files depending on the circumstances of the recovery, as [Table 5-2](#) shows.

Table 5-2 Files that restores generate

File	Description
@@MaNgLeD.nnnn	For backups containing pathnames longer than 100 characters, <code>nbtar</code> generates the files that are named @@MaNgLeD.nnnn that contain the actual file.
@@MaNgLeD.nnnn_Rename	<code>nbtar</code> generates another file (@@MaNgLeD.nnnn_Rename) that explains how to rename the @@MaNgLeD.nnnn files to return the files to the correct location.
@@MaNgLeD.nnnn_Symlink	For long names of symbolic links, <code>nbtar</code> generates the files that are named @@MaNgLeD.nnnn_Symlink. These files contain descriptions of the symbolic links that must be made to return a link to the correct file.
For cross-platform VxFS extent attribute restores, <code>nbtar</code> creates and stores extent attributes in .ExTeNt.nnnn files in the root 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 5-1](#) shows the major factors that affect backup time.

Figure 5-1 Backup time formula

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

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 5-3 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 5-3 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 computers run full-duplex and some run half-duplex in an Ethernet environment, the throughput is significantly reduced.
Compression (on UNIX)	Software compression often multiplies the backup time by a factor of two or three for a given set of data.
Device delays	Device delays can be due to the following factors: ■ 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. These delays can vary widely and depend on the devices and the computing environments.

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 5-4 NetBackup transfer rates

Transfer rate	Description
Network transfer rate	The network transfer rate is the rate provided in the All Log Entries report. The network transfer rate considers only the time it takes to transfer data over the network from client to server. This rate ignores the following: ■ The time the device requires to load and to 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	This rate ignores the time it takes to load and to 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: <pre>begin writing backup id xxx</pre> until the message <pre>successfully wrote backup id xxx</pre> To calculate the transfer rate, divide this time (in seconds) into the total bytes that are transferred. (The total bytes that are transferred are recorded in the All Log Entries report.)
Total transfer rate	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).

On Windows, the Microsoft Windows System Monitor also displays the NetBackup transfer rate.

Examples of the 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 TLD033 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 KB at 230.325 KB per second

Network transfer plus end-of-backup processing rate:

23:10:30 - 00:35:07 = 01:24:30 = 5070 seconds

1161824 KB/5070 = 229.157 KB per second

Total transfer rate:

Elapsed time = 01:27:32 = 5252 seconds

1161824 Kbytes/5252 = 221.216 KB per second

NetBackup notify scripts

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

All of the NetBackup notify scripts are located in the `goodies` directory, which contains sample shell scripts to modify. Three `volmgr` notify scripts are installed in

the `volmgr` `goodies` directory. The scripts in the `goodies` directories are not supported but are intended as examples to customize.

The `goodies` directories are found in the following locations:

On Windows: `Install_path\NetBackup\bin\goodies\`

`Install_path\volmgr\bin\goodies`

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

`/usr/opensv/volmgr/bin/goodies`

Notes about using scripts

- Ensure that others can run the script after modifying. To do so, run `chmod ugo+rx script_name`, where `script_name` is the name of the script.
- If you use either the `bpstart_notify` or `bpend_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`.
- Many NetBackup processes set a limit on the number of concurrently open file descriptors that are allowed. That limit is inherited by the notify scripts run by the process. In the rare event that a command invoked by a notify script requires many additional file descriptors, the script must increase the limit appropriately before invoking the command.

The following topics describe the scripts that are active on the primary server and those that are active on the client.

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

Additional comments appear in the scripts.

backup_notify script

The `backup_notify.cmd` script (on Windows) and the `backup_notify` script (on UNIX) runs on the NetBackup server where the storage unit is located. It's called each time a backup is successfully written to media.

To use this script, copy the script to the `bin` directory:

- UNIX:

```
/usr/opensv/netbackup/bin/goodies/backup_notify to
/usr/opensv/netbackup/bin
```

- **Windows:**

```
install_path\netbackup\bin\goodies\backup_notify.cmd to
install_path\netbackup\bin
```

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

NetBackup passes the following parameters to this script:

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

See the following Windows example:

```
backup_notify.cmd bptm host_0695316589
```

backup_exit_notify script

The `backup_exit_notify.cmd` script (on Windows) and the `backup_exit_notify` script (on UNIX) run on the primary server. It's called to perform site-specific processing when an individual backup completes.

To use this script, copy the script to the `bin` directory:

- **UNIX:**

```
/usr/opensv/netbackup/bin/goodies/backup_exit_notify to
/usr/opensv/netbackup/bin
```

- **Windows:**

```
install_path\netbackup\bin\goodies\backup_exit_notify.cmd to
install_path\netbackup\bin
```

Modify the script and confirm that you have permission to run 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: FULL, INCR (differential incremental), CINC (cumulative incremental), UBAK, UARC
<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.

See the following UNIX example:

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

bpstart_notify script (UNIX clients)

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

Note: Ensure that others can run this script on the client before it's used. To do so, run `chmod ugo+rx script_name`, where *script_name* is the name of the script.

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

```
/usr/opensv/netbackup/bin/goodies/bpstart_notify
```

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. The `bpbkar` process on the client waits for the script to complete before

it continues. 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 seconds. If the script needs more time than 300 seconds, increase the value to allow more time. (The `BPSTART_TIMEOUT` option corresponds to the **Backup start notify timeout** on the **Timeouts** host properties.)

Note: The **Client read timeout** (`CLIENT_READ_TIMEOUT` option) must be equal to or greater than the **Backup start notify timeout** (`BPSTART_TIMEOUT` option). If the **Client read timeout** is less than the **Backup start notify timeout**, the job can time out while the `bpstart_notify` script is running.

NetBackup passes the following parameters to the script:

<code>clientname</code>	Specifies the name of the client from the NetBackup catalog.
<code>polycyname</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 `.polycyname[.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 `.polycyname` or `.polycyname.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 `bpbkar` 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 2016
```

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

Table 5-5 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>bpbkar</code> .
RESTARTED	Specifies the 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 script (Windows clients)

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:

On Windows:

```
Install_path\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.schedulename* 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: `bpend_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 `bpend_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 `BPSTART_TIMEOUT` option specifies on the server. The default for `BPSTART_TIMEOUT` is 300 seconds. If the script needs more time than 300 seconds, increase the value to allow more time. (The `BPSTART_TIMEOUT` option corresponds to the **Backup start notify timeout** on the **Timeouts** host properties.)

Note: The **Client read timeout** (`CLIENT_READ_TIMEOUT` option) must be equal to or greater than the **Backup start notify timeout** (`BPSTART_TIMEOUT` option). If the **Client read timeout** is less than the **Backup start notify timeout**, the job can timeout while the `bpstart_notify` script is running.

bpend_notify script (UNIX clients)

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

```
/usr/opensv/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 other administrators can run the notify scripts after they are modified. To do so, run `chmod ugo+rx 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>polycname</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>bpbkcar</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.schedulename` suffix. The following are two examples of script names for a policy that is named `production` with a schedule that is named `fulls`:

```
/usr/openv/netbackup/bin/bpend_notify.production
/usr/openv/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 `bpbkcar` 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 5-6 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 bpbkar.
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 of the data.

bpend_notify.bat script (Windows clients)

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\bpnd_notify.days.fulls.bat
```

Note: The `bpnd_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 `bpnd_notify` script and checks for them in the following order:

```
bpnd_notify.policy.schedule.bat  
bpnd_notify.policy.bat  
bpnd_notify.bat
```

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

Note: `bpstart_notify` scripts can provide a different level of notification than the `bpnd_notify` scripts. For example, if you had one of each, they could be `bpstart_notify.policy.bat` and `bpnd_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 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.

%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.

bpend_notify_busy script (UNIX clients)

Use the `bpend_notify_busy` script to configure busy file processing when using the `bp.conf` file.

See [“About busy file processing on UNIX clients”](#) on page 91.

Busy file processing can also be configured in the **Busy File Settings** host properties of the **NetBackup Administration Console**.

child_end_deployment_notify

The `child_end_deployment_notify` script (on UNIX) and the `child_end_deployment_notify.cmd` script (on Windows) runs on the NetBackup primary server. NetBackup calls the script each time a deployment child job completes. The script runs after all other deployment steps have completed.

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

On UNIX: `/usr/opensv/netbackup/bin/goodies/child_end_deployment_notify`

On Windows:

`install_path\NetBackup\bin\goodies\child_end_deployment_notify.cmd`

Then place the script in the following location on the primary server:

On UNIX: `/usr/openv/netbackup/bin/`

On Windows: `install_path\NetBackup\bin\`

To run properly, the script must be executable. To make the script executable on a UNIX primary server, run `chmod ugo+rx child_end_deployment_notify`

NetBackup passes the following parameters to the script based on the platform:

Table 5-7 UNIX `child_end_deployment_notify` parameters

Parameter	Details
<code>JobID</code>	Specifies the job ID of the child job.
<code>GroupJobID</code>	Specifies the child job's parent job's job ID.
<code>PolicyName</code>	Specifies the policy name of the deployment policy.
<code>ClientName</code>	Specifies the host name of the client as it is found in the deployment policy.
<code>Operation</code>	Specifies the operations to perform: precheck, stage, and install.
<code>Package</code>	Specifies the deployment policy's package.
<code>PrecheckStatus</code>	Specifies the status of the precheck child job step.
<code>StageStatus</code>	Specifies the status of the stage step child job step, if performed.
<code>InstallStatus</code>	Specifies the status of the install child job step, if performed.
<code>JobStatus</code>	Specifies the exit status code for the child job.

Table 5-8 Windows `child_end_deployment_notify.cmd` parameters

Parameter	Details
<code>%1</code>	Specifies the job ID of the child job.
<code>%2</code>	Specifies the child job's parent job's job ID.
<code>%3</code>	Specifies the policy name of the deployment policy.
<code>%4</code>	Specifies the host name of the client as it is found in the deployment policy.

Table 5-8 Windows child_end_deployment_notify.cmd parameters
(continued)

Parameter	Details
%5	Specifies the operations to perform: precheck, stage, and install.
%6	Specifies the deployment policy's package.
%7	Specifies the status of the precheck child job step.
%8	Specifies the status of the stage step child job step, if performed.
%9	Specifies the status of the install child job step, if performed.
%10	Specifies the exit status code for the child job.

child_start_deployment_notify

The `child_start_deployment_notify` script (on UNIX) and the `child_start_deployment_notify.cmd` script (on Windows) runs on the NetBackup primary server. NetBackup calls the script each time a new deployment child job starts and initialization is completed. The script runs before all deployment steps are initiated.

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

On UNIX: `/usr/opensv/netbackup/bin/goodies/child_start_deployment_notify`

On Windows:

`install_path\NetBackup\bin\goodies\child_start_deployment_notify.cmd`

Then place the script in the following location on the primary server:

On UNIX: `/usr/opensv/netbackup/bin/`

On Windows: `install_path\NetBackup\bin\`

To run properly, the script must be executable. To make the script executable on a UNIX primary server, run `chmod ugo+rx child_start_deployment_notify`

NetBackup passes the following parameters to the script based on the platform:

Table 5-9 UNIX child_start_deployment_notify parameters

Parameter	Details
JobID	Specifies the job ID of the child job.
GroupJobID	Specifies the child job's parent job's job ID.

Table 5-9 UNIX child_start_deployment_notify parameters (*continued*)

Parameter	Details
PolicyName	Specifies the policy name of the deployment policy.
ClientName	Specifies the host name of the client as it is found in the deployment policy.
Operation	Specifies the operations to perform: precheck, stage, and install.
Package	Specifies the deployment policy's package.

Table 5-10 Windows child_start_deployment_notify.cmd parameters

Parameter	Details
%1	Specifies the job ID of the child job.
%2	Specifies the child job's parent job's job ID.
%3	Specifies the policy name of the deployment policy.
%4	Specifies the host name of the client as it is found in the deployment policy.
%5	Specifies the operations to perform: precheck, stage, and install.
%6	Specifies the deployment policy's package.

diskfull_notify script

The `diskfull_notify.cmd` script (on Windows) and the `diskfull_notify` script (on UNIX) run 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`).

To use this script, copy the script to the `bin` directory:

- UNIX:

```
/usr/opensv/netbackup/bin/goodies/diskfull_notify to  
/usr/opensv/netbackup/bin
```

- Windows:

```
install_path\netbackup\bin\goodies\diskfull_notify.cmd to  
install_path\netbackup\bin
```

Then place the script in the `/usr/opensv/netbackup/bin/` directory on the UNIX client. For a Windows client, place the script in the `install_path\netbackup\bin` directory.

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

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:

```
/disk1/images/host_08193531_cl_F1
```

See the following Windows example:

```
diskfull_notify.cmd bpdm
```

drive_mount_notify script (on UNIX)

The NetBackup `tpreq` command runs the `drive_mount_notify` script (if it exists) immediately after media is mounted in a pre-selected, robotic drive. This script is not valid for standalone drives.

Each time a tape volume is mounted, this script gathers information on the drive that is mounted. This script also lets you perform special-handling. For example, you can use the script to gather log sense or other data from a drive and place it in an output file. You can change the information that the script gathers by modifying the script.

After the script runs, control is then returned to NetBackup to resume processing.

This script is located in the following directory:

```
/usr/opensv/volmgr/bin/goodies
```

To use this script, activate it and place it into the `/usr/opensv/volmgr/bin` directory. See the script for instructions about how to activate it and how to modify it.

drive_unmount_notify script (on UNIX)

The NetBackup `tpunmount` command runs the `drive_unmount_notify` script (if it exists) after media is unmounted. This script is valid for robotic drives and standalone drives.

Each time a tape volume is unmounted, this script gathers information about the drive that is unmounted. This script also lets you perform special-handling. For example, you can use the script to gather log sense or other data from a drive and place it in an output file. You can change the information that the script gathers by modifying the script.

After the script runs, control is then returned to NetBackup to resume processing.

This script is located in the following directory:

```
/usr/opensv/volmgr/bin/goodies
```

To use this script, activate it and place it into the `/usr/opensv/volmgr/bin` directory. See the script for instructions about how to activate it and how to modify it.

mail_dr_info script

Use the `mail_dr_info.cmd` script (on Windows) and the `mail_dr_info.sh` script (on UNIX) to send NetBackup disaster recovery information to specified recipients after running an online, hot catalog backup.

By default, this script does not exist. You must create it. How you do so depends on the operating system type of your primary server.

On Windows: To create the script, copy the following script from the primary server:

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

and place it into the following location:

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

On UNIX: To create the script, touch the following file:

```
/usr/opensv/netbackup/bin/mail_dr_info.sh
```

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 either the path or a comma-separated list for the path to the recovery email and the disaster recovery (DR) package.

On Windows: 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.

On UNIX: NetBackup checks to see if `mail_dr_info.sh` is present in `/usr/opensv/netbackup/bin`. If `mail_dr_info.cmd` exists, NetBackup passes the parameters to the script. `mail_dr_info.sh` is not an installed file. Users must create the script.

media_deassign_notify script

The NetBackup 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.)

On Windows: Copy

`Install_path\NetBackup\bin\goodies\media_deassign_notify.cmd` into `Install_path\NetBackup\bin\` on the primary server.

On UNIX: Copy `/usr/opensv/netbackup/bin/goodies/media_deassign_notify` into `/usr/opensv/netbackup/bin/` on the primary 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 script (on Windows)

Use the `nbmail.cmd` script to send the 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 to accept script parameters. UNIX platforms have a built-in SMTP transfer method.

To create the script on a client, copy

`Install_path\NetBackup\bin\goodies\nbmail.cmd` from the primary 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_deployment_notify

The `parent_end_deployment_notify` script (on UNIX) and the `parent_end_deployment_notify.cmd` script (on Windows) runs on the NetBackup primary server. NetBackup calls the script each time a deployment parent job completes. The script runs after all other deployment steps have completed.

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

On UNIX: `/usr/opensv/netbackup/bin/goodies/parent_end_deployment_notify`

On Windows:

`install_path\NetBackup\bin\goodies\parent_end_deployment_notify.cmd`

Then place the script in the following location on the primary server:

On UNIX: `/usr/opensv/netbackup/bin/`

On Windows: `install_path\NetBackup\bin\`

To run properly, the script must be executable. To make the script executable on a UNIX primary server, run `chmod ugo+rx parent_end_deployment_notify`

NetBackup passes the following parameters to the script based on the platform:

Table 5-11 UNIX `parent_end_deployment_notify` parameters

Parameter	Details
JobID	Specifies the job ID of the parent job.
PolicyName	Specifies the policy name of the deployment policy.

Table 5-11 UNIX parent_end_deployment_notify parameters (*continued*)

Parameter	Details
Operation	Specifies the operations to perform: precheck, stage, and install.
Package	Specifies the deployment policy's package.
JobStatus	Specifies the exit status code of the parent job.
ClientCount	Specifies the number of child jobs that the parent job initiated.

Table 5-12 Windows parent_end_deployment_notify.cmd parameters

Parameter	Details
%1	Specifies the job ID of the parent job.
%2	Specifies the policy name of the deployment policy.
%3	Specifies the operations to perform: precheck, stage, and install.
%4	Specifies the deployment policy's package.
%5	Specifies the exit status code of the parent job.
%6	Specifies the number of child jobs that the parent job initiated.

parent_end_notify script

NetBackup calls the `parent_end_notify.cmd` script (on Windows) and the `parent_end_notify` script (on UNIX) each time a parent job ends.

To use this script, copy the script to the `bin` directory:

- UNIX:

```
/usr/opensv/netbackup/bin/goodies/parent_end_notify to
/usr/opensv/netbackup/bin
```

- Windows:

```
install_path\netbackup\bin\goodies\parent_end_notify.cmd to
install_path\netbackup\bin
```

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

NetBackup passes the following parameters to the script:

`clientname` 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>	Specifies the stream number; it's always -1.
<code>stream_count</code>	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 <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_deployment_notify

The `parent_start_deployment_notify` script (on UNIX) and the `parent_start_deployment_notify.cmd` script (on Windows) runs on the NetBackup primary server. The script runs each time a new deployment parent job starts and initialization is completed. The script runs before all deployment steps are initiated.

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

On UNIX:

`/usr/opensv/netbackup/bin/goodies/parent_start_deployment_notify`

On Windows:

`install_path\NetBackup\bin\goodies\parent_start_deployment_notify.cmd`

Then place the script in the following location on the primary server:

On UNIX: `/usr/opensv/netbackup/bin/`

On Windows: `install_path\NetBackup\bin\`

To run properly, the script must be executable. To make the script executable on a UNIX primary server, run `chmod ugo+rx parent_start_deployment_notify`

NetBackup passes the following parameters to the script based on the platform:

Table 5-13 UNIX `parent_start_deployment_notify` parameters

Parameter	Details
<code>JobID</code>	Specifies the job ID of the parent job.
<code>PolicyName</code>	Specifies the policy name of the deployment policy.

Table 5-13 UNIX parent_start_deployment_notify parameters (*continued*)

Parameter	Details
Operations	Specifies the operations to perform: precheck, stage, and install.
Package	Specifies the deployment policy's package.

Table 5-14 Windows parent_start_deployment_notify.cmd parameters

Parameter	Details
%1	Specifies the job ID of the parent job.
%2	Specifies the policy name of the deployment policy.
%3	Specifies the operations to perform: precheck, stage, and install.
%4	Specifies the deployment policy's package.

parent_start_notify script

NetBackup calls the `parent_start_notify.cmd` script (on Windows) or the `parent_start_notify` script (on UNIX) each time a parent job starts.

To use this script, copy the script to the `bin` directory:

- UNIX:

```
/usr/opensv/netbackup/bin/goodies/parent_start_notify to  
/usr/opensv/netbackup/bin
```

- Windows:

```
install_path\netbackup\bin\goodies\parent_start_notify.cmd to  
install_path\netbackup\bin
```

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

NetBackup passes the following parameters to the script:

<code>clientname</code>	Specifies the name of the client from the NetBackup catalog.
<code>polycyname</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: FULL, INCR (differential incremental), CINC (cumulative incremental), UBAK, UARC
<code>status</code>	Specifies the exit code for the entire backup job.

`streamnumber` Specifies the stream number; for a parent job it's always -1.

pending_request_notify script

The NetBackup Media Manager 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. (A root user must run the script.)

On Windows: Copy

`install_path\NetBackup\bin\goodies\pending_request_notify.cmd` into
`install_path\NetBackup\bin\` on the primary server.

On UNIX: Copy `/usr/opensv/netbackup/bin/goodies/pending_request_notify`
into `/usr/opensv/netbackup/bin/` on the primary 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 script

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

To use this script, copy the script to the `bin` directory:

- UNIX:

`/usr/opensv/netbackup/bin/goodies/restore_notify` to
`/usr/opensv/netbackup/bin`

- Windows:

`install_path\netbackup\bin\goodies\restore_notify.cmd` to
`install_path\netbackup\bin`

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

NetBackup passes the following parameters to the script:

`programname` Specifies the name of the program doing the restore or other read operation.

`pathname` Specifies the path to the backup name or path.

`operation` Specifies one of the following: `restore`, `verify`, `duplication`, `import`

session_notify script

The `session_notify.cmd` script (on Windows) and the `session_notify` script (on UNIX) run on the primary server. It's 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.

To use this script, copy the script to the `bin` directory:

- UNIX:
`/usr/opensv/netbackup/bin/goodies/session_notify` to
`/usr/opensv/netbackup/bin`
- Windows:
`install_path\netbackup\bin\goodies\session_notify.cmd` to
`install_path\netbackup\bin`

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

session_start_notify script

The `session_start_notify.cmd` script (on Windows) and the `session_start_notify` script (on UNIX) run on the primary 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.

To use this script, copy the script to the `bin` directory:

- UNIX:
`/usr/opensv/netbackup/bin/goodies/session_start_notify` to
`/usr/opensv/netbackup/bin`
- Windows:
`install_path\netbackup\bin\goodies\session_start_notify.cmd` to
`install_path\netbackup\bin`

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

shared_drive_notify script

NetBackup runs the `shared_drive_notify.cmd` script (on Windows) and the `shared_drive_notify` script (on UNIX) 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's 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. The scan host may change between a RESERVED operation and a RELEASED operation.

This script is located in the `/usr/opensv/volmgr/bin/goodies` directory on UNIX and the `install_path\Volmgr\bin\goodies` directory on Windows.

To use this script, activate it and place it into the `/usr/opensv/volmgr/bin` (UNIX) or `install_path\Volmgr\bin` (Windows) directory.

See the script for instructions about how to activate it and how to modify it.

The script must be executable by the root user.

The script exits with status 0 upon successful completion.

userreq_notify script

The `userreq_notify.cmd` script (on Windows) and the `userreq_notify` script (on UNIX) run on the primary server.

To use this script, copy the script to the `bin` directory:

- UNIX:
`/usr/opensv/netbackup/bin/goodies/userreq_notify` to
`/usr/opensv/netbackup/bin`
- Windows:
`install_path\netbackup\bin\goodies\userreq_notify.cmd` to
`install_path\netbackup\bin`

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

NetBackup calls the script 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.

See the following UNIX example:

```
userreq_notify backup mercury jdoe
userreq_notify archive mercury jdoe
userreq_notify manual_backup mercury jdoe
userreq_notify restore mercury jdoe
userreq_notify 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 website:

<http://www.netbackup.com/compatibility>

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

- Use only the NetBackup commands that Veritas 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 the 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, primary 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 barcodes are unique across the entire enterprise.
- On UNIX hosts: To use the devices that NetBackup controls but are used with other applications, do the following to avoid the potential loss of data:
 - Use the NetBackup `tpreq` command to mount media on a drive and `tpunmount` to remove media from the drive. If you use these commands, another application can control a device when NetBackup is finished with the device.
 - Down the drive, if the drive is in the UP state.
- On Windows hosts: 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). See the *NetBackup Device Configuration Guide* at the following URL:
<http://www.veritas.com/docs/DOC5332>
- Use only computers, operating systems and devices that Veritas 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 primary 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 Status Codes Reference Guide* for all error conditions:
<http://www.veritas.com/docs/DOC5332>
- Always install the latest NetBackup release updates that are available from Veritas.
- 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 203.

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 TapeAlert 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.

It is recommended that you use reactive cleaning.

See [“About TapeAlert”](#) on page 198.

See [“About tape drive cleaning”](#) on page 202.

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 Veritas Support website.

- 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 media servers.

See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

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

The following table describes the codes.

Table 5-15 TapeAlert log codes

TapeAlert code	Default action	Error type	Error message
0x01	None	Warning - WRN	Read warning
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

Table 5-15 TapeAlert log codes (*continued*)

TapeAlert code	Default action	Error type	Error message
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
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

Table 5-15 TapeAlert log codes (*continued*)

TapeAlert code	Default action	Error type	Error message
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
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 198. It is recommended that you use reactive cleaning.
Library-based cleaning	See “About library-based cleaning” on page 203.
Frequency-based cleaning	See “About frequency-based cleaning” on page 203.
Operator-initiated cleaning	See “About operator-initiated cleaning” on page 204.

See [“About using a cleaning tape”](#) on page 204.

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, it is recommended to disable 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 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 199.

See [“About tape drive cleaning”](#) on page 202.

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.

Veritas 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 service (`nbemm`) manages the drives and requests for locally-attached or shared drives in the EMM domain.

The EMM service 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 5-16 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 `bpbbackupdb`.) 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 PEND. 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 PEND 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:

```
vmoprncmd -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 `vmoprncmd` command on the primary server. Alternatively issue the command on a media server and use the `-h` option of the command to specify the primary server. The NetBackup EMM service allocates devices (that is, the DA host or device allocation host).

Note: Use this command after a PEND status appears in the **NetBackup Administration Console Device Monitor**. However, do not issue this command during backups.

More information about using the `vmoprncmd` command is available.

See the *NetBackup Commands Reference Guide*, available at the following URL:

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

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.

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*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>
- You must configure the operating systems on the NetBackup media servers so they let NetBackup control SCSI persistent reserve or SPC-2 SCSI reserve.
- Depending on the tape drives, you may have to disable the operating system's use of SPC-2 SCSI reserve. Solaris may require such a change.
See the *NetBackup Device Configuration Guide*, available at the following URL:
<http://www.veritas.com/docs/DOC5332>

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.
- With SPC-2 SCSI reserve, devices may remain reserved after a failover in cluster environments or multi-path 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 multi-path 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 SCSI reserve operating system limitations on Windows

Windows operating systems cannot distinguish between a reserved device and a busy device. Therefore, PEND appears in the **NetBackup Administration Console 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 PEND if the drive reports Busy when a volume is unmounted. Use the `AVRD_PEND_DELAY` entry in the `vm.conf` configuration file to filter out these extraneous reports.

About checking for data loss

To detect data loss, the NetBackup `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 primary 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 `bpexpdate` 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*, available at the following URL:

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

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 215.

See [“About selecting media in standalone drives”](#) on page 217.

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.

[Table 5-17](#) describes the order in which NetBackup selects a volume in a robot.

Table 5-17 How NetBackup selects a volume in a robot

Order	Description
1.	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 the backup job requested. ■ Not currently in use by another backup or a restore. ■ Not written in a protected format. NetBackup detects the 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.

Table 5-17 How NetBackup selects a volume in a robot (*continued*)

Order	Description
2.	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 <code>tpreq</code> 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.
3.	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
4.	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.
5.	If there are no unassigned volumes of the requested type, the backup terminates with an error message that no media were available. NetBackuptakes no action.

See [“About spanning media with automatic media selection”](#) on page 216.

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 217.

See [“About disabling standalone drive extensions”](#) on page 218.

See [“About spanning media”](#) on page 218.

See [“About leaving standalone drives in the ready state”](#) on page 219.

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 215.
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 primary 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 the *NetBackup Commands Reference Guide*, available at the following URL:

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


Volume pool and volume group examples

The following three examples show the relationship between volume pools and volume groups.

See [Figure 5-2](#) on page 220. for 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 5-2 Volume pool with multiple volume groups

See [Figure 5-3](#) on page 221. for examples of 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 5-3 Volume groups with multiple volume pools




See [Figure 5-4](#) on page 222. for an example where 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 5-4 Scratch pool example



Media formats

NetBackup writes media in a format that allows the position to be verified before NetBackup appends new backups.

The following table shows the symbols that are used in the media format descriptions.

Table 5-18 Media format symbols

Symbol	Description
MH	Media header (1024 bytes).
*	Tape mark.

Table 5-18 Media format symbols (*continued*)

Symbol	Description
BH	Backup header (1024 bytes).
BH1 ... BH n	Backup headers (1024 bytes). One for each job that is part of the set of the jobs that are multiplexed.
Image	Data from the backup.
EH	Empty backup header, which is used for position validation.

The following table provides more information about how the media formats are used in different situations.

Table 5-19 Media format descriptions

Format	Description
Standard tape format	For all tape media except quarter-inch cartridge (QIC) and WORM, the format for the backups that are not multiplexed is as follows: MH * BH Image * BH Image * BH Image * EH * 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 position validation. 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.
QIC and WORM tape format	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: MH * BH Image * BH Image * BH Image * To append backup images to QIC media, NetBackup positions to the end of data (EOD) and then starts the next backup.

Table 5-19 Media format descriptions (*continued*)

Format	Description
Fragmented backup format	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. The following is an example: <pre>MH * BH Image (frag 1)* BH Image (frag 2)* BH Image (frag n) * EH *</pre> Fragmentation is intended primarily for storing large backup images on a disk type storage unit. 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. Note: 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.
Multiplexing format	The tape format for multiplexed backups is as follows: <pre>MH * BH1 ... BHn Image ...</pre> 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. 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. The following is an example: <pre>MH * BH1 BH2 BH3 Image* BH2 BH3 Image* BH2 BH3 BH4 Image</pre>
Spanning tape format	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. The following is the first tape format (NetBackup does not write an EH and terminates the tape with two tape marks): <pre>MH * ... *BHn Image (frag 1) * *</pre> The following is the second tape format: <pre>MH * BHn Image (frag2)* ... * EH *</pre>

Media and device management processes

Table [Table 5-20](#) shows the NetBackup services and processes that control storage devices with removable media. NetBackup starts the processes as needed, but you can start some of them manually. The table also shows the commands that start each one.

These commands are located in the following directories:

UNIX `/usr/opensv/volmgr/bin`

Windows `install_path\VERITAS\Vlmgr\bin`

For detailed information about the commands, see the *NetBackup Commands Reference Guide*, available at the following URL:

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

Table 5-20 Starting services and processes

Command	Description
<code>acsd</code>	The Automated Cartridge System robotic process. The Device Manager <code>ltid</code> starts this process.
<code>avrd</code>	The Automatic Volume Recognition process. The Device Manager <code>ltid</code> starts this process.
<code>ltid</code>	Starts the NetBackup Device Manager service. Starting the Device Manager also starts the robotic, robotic control, Media Manager volume, and automatic volume recognition daemons. To stop the device, robotic, and robotic-control services, use the <code>stopltid</code> command.
<code>tlcdc</code>	Starts the tape library DLT robotic-control process. The Device Manager <code>ltid</code> starts this process. To stop the tape library DLT robotic-control process, use <code>tlcdc -t</code> .
<code>tldd</code>	The tape library DLT robotic process. The Device Manager <code>ltid</code> starts this process.
<code>vmd</code>	The NetBackup Volume Manager service. The Device Manager <code>ltid</code> starts this process.

On UNIX, you can use the `kill pid` command to stop the process for the daemon with the specified *pid* (process ID).

On Windows, you can 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.

For detailed information about most of the commands that are in the following tables, see the *NetBackup Commands Reference Guide*, available at the following URL:

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

About Tape I/O commands on UNIX

To troubleshoot or test Media Manager, use the commands that are described in the following sections to manipulate volumes. Similarly, if you do not use NetBackup, you can use these commands to mount and manipulate volumes.

About requesting tapes

The `tpreq` command lets you request a tape of a particular density and specify various options, such as the access mode. This command reserves a single drive and creates a file in the current working directory (unless a full path is specified). The file acts as a symbolic link to the tape and all subsequent access to the tape is through this file name. Users do not have to be concerned with the full path to a specific device file.

For all types of tapes, the tape is mounted and assigned when you enter the `tpreq` command.

By default, NetBackup assigns drives that support DLT cartridge tapes. You can use the density option on `tpreq` to request a drive that supports another density. For a list of supported densities and drive types, see the `tpreq` man page.

The density for the physical write is not selected automatically on drives. It's requested, so an operator can satisfy the correct drive. One of two methods is used to determine the drive density: the `/dev` device name that was used when the drive was configured or by how the drive is configured physically.

A `tpreq` command must include a media ID and a file name. If the tape volume is associated with a volume pool, the name of the volume pool can also be specified by using the `-p` parameter. If you specify the pool name, the name is validated against the pool name that is associated with the media in the EMM database.

The NetBackup `tpreq` command runs the `drive_mount_notify` script (if it exists) immediately after media is mounted in a pre-selected, robotic drive.

See [“drive_mount_notify script \(on UNIX\)”](#) on page 185.

See the *NetBackup Commands Reference Guide*, available at the following URL:

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

About reading and writing tape files

Reading or writing tape files involves copying the file from tape to disk or from disk to tape. To perform read or write operations, use one of the UNIX commands that performs input and output operations, for example `tar` or `mt`.

Positioning tape files

The `mt` command positions tape files by skipping forward or backward according to tape marks.

The following options are available on the `mt` command for positioning tapes:

- `eof, weof`
Writes an end-of-file tape mark at the current position on the tape according to the count option on `mt`.
- `fsf, bsf`
Spaces forward or backward the number of tape marks on the count option.
- `fsr, bsr`
Spaces forward and backward the number of records according to the count option on `mt`. `bsr` is only supported for the undefined record type.

The following example uses the `mt` command to skip forward three files on a tape:

```
mt -f tape1 fsf 3
```

Rewinding tape files

When a file is rewound, it is positioned to the beginning of the data. To rewind a tape file, you can use the `mt` command.

`tape1` is positioned to the beginning of the tape volume that is associated with the file.

The following command rewinds file `tape1`:

```
mt -f tape1 rewind
```

The count option is not used for the rewind operation. If you specify a count, `mt` ignores it.

About removing tape files

When you have completed reading or writing tape files, use the `/usr/opensv/volmgr/bin/tpunmount` command to end the assignment of the tape file. This command removes from the directory the tape file you created by with `tpreq` and removes the tape volume from the tape drive. The `tpunmount` command is required for each file that the `tpreq` command creates.

See the *NetBackup Commands Reference Guide*, available at the following URL:

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

On UNIX, the NetBackup `tpunmount` command runs the `drive_unmount_notify` script (if it exists) after media is unmounted.

See “[drive_unmount_notify script \(on UNIX\)](#)” on page 186.

Index

Symbols

.ExTeNt.nnnn files 165
@@MaNgLeD.nnnn files 165
@@MaNgLeD.nnnn_Rename files 165
@@MaNgLeD.nnnn_Symlink files 165

A

accurate licensing 42, 66
ACS robot types 102
ACS_vm.conf entry 121
ACS_CSI_HOSTPORT 121
ACS_SEL_SOCKET vm.conf entry 121
ACS_SSI_HOSTNAME vm.conf entry 122
ACS_SSI_INET_PORT
 vm.conf entry 123
ACS_SSI_INET_PORT vm.conf entry 122
ACS_SSI_SOCKET vm.conf entry 124
ADJ_LSM vm.conf entry 124
All Log Entries report 168
Allow backups to span media 217
alternate client restores
 host.xlate file 161
Announce DHCP interval property 83
API_BARCODE_RULES vm.conf entry 126
Arbitrated Loop Physical Address (ALPA) 108
AUTO_PATH_CORRECTION vm.conf entry 127
AUTO_UPDATE_ROBOT vm.conf entry 127
AVRD_PEND_DELAY
 vm.conf entry 213
AVRD_PEND_DELAY vm.conf entry 127
AVRD_SCAN_DELAY vm.conf entry 128

B

Backup Exec 105
backup_exit_notify script 171
backup_notify script 170
backups
 backup_exit_notify script 171
 backup_notify script 170

backups (*continued*)
 bpend_notify script
 UNIX client 177
 Windows client 179
 bpstart_notify script
 UNIX client 172
 Windows client 174
 compressed 162
 diskfull_notify script 184
 estimating time required 165
 multiplexed 162
 session_notify script 193
 session_start_notify script 193
blat mail 187
bpclient commands 87
bpclntcmd utility 109
bpdynamicclient 90
bpend_notify script
 UNIX client 177
 Windows client 179
bpend_notify_busy script 181
bpstart_notify script 173, 177
 UNIX client 172
 Windows client 174
BPSTART_TIMEOUT 173, 177
busy file processing
 bp.conf entries 93
 Busy file settings property 91, 93
 configuration overview 91
 configuring on UNIX 92
 creating action files 95
 logs directory 97
 modifying bpend_notify_busy 98
BUSY_FILE_ACTION bp.conf entry 94
BUSY_FILE_DIRECTORY bp.conf entry 93
BUSY_FILE_PROCESSING bp.conf entry 93

C

capacity reporting
 parameters 21
character device 148

- child_end_deployment_notify script 181
- child_start_deployment_notify script 183
- CLEAN_REQUEST_TIMEOUT vm.conf entry 128
- cleaning
 - automatic 203
 - frequency-based 203
 - library-based 203
 - TapeAlert reactive 198
 - times allowed 204
- Client read timeout property 173, 176
- CLIENT_PORT_WINDOW vm.conf entry 129
- CLIENT_READ_TIMEOUT 173, 177
- clients
 - changing host names 159
 - dynamic UNIX client 89
- cluster environments 212
- CLUSTER_NAME vm.conf entry 129
- compressed backups 162
- control path
 - robotic 144
- crawlreleasebyname
 - vmoprcmd option 211

D

- DAYS_TO_KEEP_LOGS vm.conf entry 129
- device
 - configuration wizard 110
 - delays 167
 - file
 - robotic 147
 - using with other applications 196
- device allocation host 101–102
- device configuration utility. *See* tpconfig
- devices
 - configuration wizard 110
 - configuring 106
 - files 107
 - management practices 197
- DHCP server 82
- direct I/O on Windows 80
- disk pools, creating 152
- diskfull_notify script 184
- display device configuration 151
- display, configuring locale 98
- Domain Name Service (DNS) hostnames 161
- drive_mount_notify script 185
- drive_unmount_notify script 186
- drives
 - ACS information 148

- drives (*continued*)
 - adding 147
 - character device 148
 - cleaning 202, 204
 - configuring 144
 - deleting 150
 - name 148
 - no rewind device 148
 - robot drive number 148
 - robot number that controls 148
 - standalone 148
 - type 148
 - update configuration 149
 - volume header device 148
- dynamic host name and IP addressing 82–84, 88–89

E

- EMM_CONNECT_TIMEOUT vm.conf entry 130
- EMM_REQUEST_TIMEOUT vm.conf entry 130
- EMM_RETRY_COUNT vm.conf entry 129
- encrypted backups 164
- extended attribute files 162
- ExTeNt.nnnn files 165

F

- files
 - .ExTeNt.nnnn 165
 - @@MaNgLeD.nnnn 165
 - @@MaNgLeD.nnnn_Rename 165
 - @@MaNgLeD.nnnn_Symlink 165
 - goodies scripts 169
 - name on tpreq 226
 - positioning on tape 227
- firmware levels 106, 108
- FlashBackup 162
- frequency-based drive cleaning 203

G

- goodies directory 169

H

- holds
 - creating 140
 - releasing 141
 - viewing hold details 140
- host names
 - changing client name 159
 - changing server name 156, 159

host names *(continued)*
 client peername 158
 correct use 156
 robotic control selection 144, 147
 short 158
 host.xlate file and alternate client restores 161
 HyperTerminal 107

I

INVENTORY_FILTER vm.conf entry 130

L

library-based cleaning 203
 licensing 12
 for Shared Storage Option 106
 locale, configuring 98

M

mail_dr_info.cmd 186
 mail_dr_info.sh 186
 MAP_CONTINUE_TIMEOUT vm.conf entry 131
 MAP_ID, vm.conf entry 131
 Maximum concurrent drives for backup 111
 media
 best practices 196
 formats 222
 selection algorithm 215, 217
 spanning 217–218
 media and device management
 best practices 195
 performance and troubleshooting 197
 Media Manager
 best practices 195
 configuration file 120
 security 134
 media_deassign_notify script 187
 MEDIA_ID_BARCODE_CHARS vm.conf entry 132
 MEDIA_ID_PREFIX vm.conf entry 133
 MM_SERVER_NAME vm.conf entry 133
 multiplexing (MPX)
 backups 224
 recovering backups 162
 tape format 224

N

named data streams 162
 nbemm 100

nbemm/DA
 definition 100
 nbholdutil -create 140
 nbmail.cmd 187
 nbtar 162, 165
 NDMP 212
 client backups 162
 host credentials 151
 NetBackup Disk Configuration Utility
 about 152
 network transfer rate 168
 notification scripts 169

O

Online Help for tpconfig 146
 open files. *See* busy-file processing

P

parent_end_deployment_notify script 188
 parent_end_notify script 189
 parent_start_deployment_notify script 190
 parent_start_notify script 191
 peername of client 158
 pending_request_notify script 192
 positioning tape files 227
 printing device configuration 151

R

random ports, setting on server 133
 RANDOM_PORTS vm.conf entry 133
 raw partitions 162
 reactive cleaning 198
 reading tape files 227
 release 141
 removing tape files 227
 reporting, capacity 20
 reporting, traditional 23
 requests
 user tape 226
 REQUIRED_INTERFACE vm.conf entry 134
 RESERVATION_CONFLICT status 210
 restore_notify script 192
 restores
 from a non-NetBackup tar 163
 restore_notify script 192
 rewind
 devices
 none 148

rewind (*continued*)
 tape files 227

robotic cleaning 203

robots

- adding 147
- configuring 144
- control host 144, 147
- deleting 149
- device file 147
- drive 148
- number 147–148
- sharing without SSO 105
- type 147
- update configuration 148

S

SAN media server 103, 105

SAN Shared Storage Option (see SSO) 100

scan host 101–102

scripts

- bpend_notify_busy 181
- bpstart_notify 172–174, 177
- child_end_deployment_notify 181
- child_start_deployment_notify 183
- drive_mount_notify 185
- drive_unmount_notify 186
- goodies 169
- notification 169
- parent_end_deployment_notify 188
- parent_end_notify 189
- parent_start_deployment_notify 190
- parent_start_notify 191
- shared_drive_notify 104, 120, 194

SCSI persistent reserve 207

SCSI reserve and release 207

- break a reservation 209, 211
- error recovery 210
- limitations 212–213
- PEND status 210–211
- requirements 212
- RESERVATION CONFLICT 209–210

SCSI-to-fibre

- bridges 107

SERVER vm.conf entry 134

servers

- changing host names 156, 159
- SAN media server 103

session_notify script 193

session_start_notify script 193

shared drives. *See* SSO

- definition 105

shared library support 105

shared robots

- without SSO 105

Shared Storage Option

- license for 106

Shared storage option

- license 100

shared_drive_notify script 104, 120

Simple Mail Transfer Protocol 187

Solaris

- extended attributes 162

spanning media 217–218, 224

SSO

- definition 100

- device allocation host 102

- Device Allocation Host Summary 117

- hardware requirements 100

- scan host 101–102

- Shared Drive Summary 117

- supported SAN hardware 120

- terminology 105

- vm.conf entries 136

SSO components configuration

- examples 100

SSO_DA_REREGISTER_INTERVAL vm.conf
 entry 135

SSO_DA_RETRY_TIMEOUT vm.conf entry 135

SSO_HOST_NAME vm.conf entry 136

standalone drive

- tpconfig 148

standalone drives

- disabling extensions 218

Storage area network (SAN) 100, 105–107

storage servers, creating 152

T

tape configuration utility. *See* tpconfig

tape drives, cleaning 202

tape formats 223

tape spanning 217–218

TapeAlert

- about 198

- cleaning flags 204

- frequency-based cleaning 199

- log codes 199

- reactive cleaning 198

- requirements 199

- tapes and tape files
 - density 226
 - positioning tape file 227
 - reading and writing 227
 - removing tape files 227
 - requesting tapes 226
 - rewinding 227
 - volume pool assignment 226
- tar32.exe 162
- tested SAN components 120
- Timeouts host properties 173, 176
- tpconfig
 - about 144
 - adding a drive 147
 - adding a robot 147
 - adding NDMP host credentials 151
 - deleting a drive 150
 - deleting robots 149
 - menus 145
 - Online Help 146
 - printing device configuration 151
 - starting 146
 - stopping 146
 - update drive configuration 149
 - update robot configuration 148
- tpreq, using to request tapes 226
- tpunmount, using to remove tape files 227
- transfer rate 166–167

U

- unified logging 129
- userreq_notify script 194
- using devices with other applications 196

V

- VERBOSE, vm.conf entry 136
- Veritas Backup Exec 105
- vm.conf file
 - ACS_ entries 121
 - ACS_CSI_HOSTPORT entries 121
 - ACS_SEL_SOCKET entries 121
 - ACS_SSI_HOSTNAME entries 122
 - ACS_SSI_INET_PORT entries 122–123
 - ACS_SSI_SOCKET entries 124
 - ADJ_LSM entries 124
 - API_BARCODE_RULES entries 126
 - AUTO_PATH_CORRECTION entries 127
 - AUTO_UPDATE_ROBOT entries 127

- vm.conf file *(continued)*
 - AVRD_PEND_DELAY entries 127
 - AVRD_SCAN_DELAY entries 128
 - CLEAN_REQUEST_TIMEOUT entries 128
 - CLIENT_PORT_WINDOW entries 129
 - CLUSTER_NAME entry 129
 - DAYS_TO_KEEP_LOGS entries 129
 - EMM_CONNECT_TIMEOUT entries 130
 - EMM_REQUEST_TIMEOUT entries 130
 - INVENTORY_FILTER entries 129–130
 - MAP_CONTINUE_TIMEOUT entries 131
 - MAP_ID entries 131
 - MEDIA_ID_BARCODE_CHARS entries 132
 - MEDIA_ID_PREFIX entries 133
 - MM_SERVER_NAME entry 133
 - overview 120
 - RANDOM_PORTS entries 133
 - REQUIRED_INTERFACE entry 134
 - SERVER entries 134
 - SSO_DA_REREGISTER_INTERVAL entries 135
 - SSO_DA_RETRY_TIMEOUT entries 135
 - SSO_HOST_NAME entries 136
 - VERBOSE entries 136
- volume groups examples 219
- volume header device 148
- volume pools examples 219
- VxFS
 - extent attributes 165
 - named data streams 162

W

- Windows, direct I/O 80
- wizards
 - device configuration 110
 - shared drive configuration 110
- writing tape files 227