

NetBackup™ Vault Administrator's Guide

UNIX, Windows, and Linux

Release 11.0

NetBackup Vault Administrator's Guide

Last updated: 2025-03-06

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

Cohesity, Veritas, the Cohesity Logo, Veritas Logo, Veritas Alta, Cohesity Alta, and NetBackup are trademarks or registered trademarks of Cohesity, Inc. 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 Cohesity 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 Cohesity 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 Cohesity, Inc. 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. Cohesity, Inc. 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 Cohesity 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.

Cohesity, Inc.
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 Cohesity 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 Cohesity 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 Cohesity community site:

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

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity 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	About Vault	13
	About Vault	13
	About Vault and NetBackup functions	13
	About accessing NetBackup Vault	14
	About Vaulting original or duplicate images	15
	About the Vault process	15
	About choosing backup images	16
	About duplicating backup images	16
	About backing up the NetBackup catalog	16
	About ejecting media	17
	About generating reports	17
	About how Vault uses volume groups and pools	17
	About NetBackup and Vault configuration	18
	About Vault management procedures	18
 Chapter 2	 Installing Vault	 21
	About supported systems	21
	About supported robots	21
	About installing and configuring Vault on UNIX and Linux systems	22
	Adding a valid license key and configuring Vault for UNIX and Linux systems	22
	About upgrading NetBackup Vault on UNIX and Linux systems	23
	Deleting a Vault license key from UNIX and Linux systems	23
	About installing and configuring Vault on Microsoft Windows systems	24
	Adding a Vault license on a Windows system	25
	About upgrading NetBackup Vault on a Windows system	25
	Deleting the Vault license from a Windows system	25
 Chapter 3	 Best Practices	 27
	About best practices	27
	About vaulting paradigms	28

About preferred vaulting strategies	28
About Vault original backups	29
About disk staging	29
About how to ensure that data is vaulted	30
About overlapping the time window in the profile	30
About resolving multiple names for a single server	31
About specifying a robotic volume group when configuring a Vault	32
About not Vaulting more than necessary	32
About sending only the intended backups off-site	32
About avoiding vaulting partial images	32
About vaulting original backups in a 24x7 environment	34
About preparing for efficient recovery	34
About Vault NetBackup catalog requirements and guidelines	34
About naming conventions for volume pools and groups	35
About matching volume pools to data usage	35
About the primary copy	36
About suspending vaulted media	36
About revaulting unexpired media	37
About media ejection recommendations	37
About avoiding resource contention during duplication	38
About two processes trying to use the same drive	38
About when the read drive is not in the Vault robot	42
About sharing resources with backup jobs	42
About load balancing	42
About specifying different volume pools for source and destination	43
About using a separate volume pool for each Vault	44
About how to avoid sending duplicates over the network	44
About creating originals concurrently	44
About using alternate read server	45
About using advanced duplication configuration	45
About using storage units that specify a media server	47
About increasing duplication throughput	47
About basic multiple-drive configuration	47
About multiple-drive configuration that does not send data over network	48
About maximizing drive utilization during duplication	49
About scratch volume pools	50
About organizing reports	50
About organizing reports by robot	50
About organizing reports by Vault	51
About organizing reports by profile	51

	About the consequences of sharing an off-site volume group across multiple robots	51
	About generating the lost media report regularly	51
Chapter 4	Configuring NetBackup Vault	53
	About configuring NetBackup Vault	53
	About off-site volume pools	53
	Creating a volume pool	54
	About creating catalog backup schedules for Vault	55
	Creating a Vault catalog backup schedule in an existing policy	56
	About setting master server properties for Vault	58
	Setting the maximum number of Vault jobs	58
Chapter 5	Configuring Vault	61
	About configuring Vault	61
	About Vault configuration	62
	Configuration information about master servers, media servers, and storage units	62
	Robot information	63
	About configuration methods	63
	About configuring Vault Management Properties	64
	General tab (Vault Management Properties)	64
	Alternate Media Server Names tab (Vault Management Properties)	66
	Retention Mappings tab (Vault Management Properties)	68
	Reports tab (Vault Management Properties)	69
	Configuring robots in Vault	70
	Vault Robot dialog box options	70
	About creating a vault	71
	Creating a Vault	71
	About configuring Vault dialog box attributes	72
	Media access ports dialog box	74
	Creating retention mappings	74
	About creating profiles	75
	Profile dialog box	75
	About the number of profiles required	76
	Creating a profile	76
	Configuring a profile	77
	About configuring a profile using the Choose Backups tab	78
	Choose Backups tab configuration options	79
	Duplication tab	83

Catalog backup tab (Profile dialog box)	98
Eject tab (Profile dialog box)	101
Reports tab (Profile dialog box)	107

Chapter 6 Vaulting and managing media 111

About Vault sessions	112
About scheduling a Vault session	112
About running a session manually	115
About running multiple sessions simultaneously	116
About previewing a Vault session	117
Stopping a Vault session	118
About resuming a Vault session	118
About monitoring a Vault session	119
About detailed Vault job status	120
About extended error codes	121
About the list of images to be vaulted	122
About duplication exclusions	122
About ejection exclusions	123
About Vault resiliency	123
About ejecting media	124
Previewing media to be ejected	124
Ejecting media by using the NetBackup Administration Console	125
Ejecting media by using the Vault operator menu	126
Ejecting media by using the vteject command	127
Ejecting media by using a Vault policy	128
Consolidating ejects and reports	129
About injecting media	130
Injecting media for libraries with and without barcode readers	131
Injecting media by using the Vault Operator Menu	132
Injecting media by using the vltinject command	133
About using containers	134
About vaulting media in containers	135
About managing containers and media	138
Generating a Container Inventory Report	140
Assigning multiple retentions with one profile	141
About vaulting additional volumes	145
Duplicating a volume manually	145
Duplicating a volume by using Vault	146
Revaulting unexpired media	146
About tracking volumes not ejected by Vault	148

	Vaulting non-NetBackup media managed by Media Manager	150
	About notifying a tape operator when an eject begins	151
	About using notify scripts	151
	About notify script for a specific robot	153
	About notify script for a specific Vault	153
	About notify script for a specific profile	153
	Notify script order of execution	154
	About clearing the media description field	154
	Restoring data from vaulted media	154
	Replacing damaged media	155
Chapter 7	Creating originals or copies concurrently	161
	About concurrent copies	161
	About the continue or fail for concurrent copies	162
	About continue copies	163
	About fail all copies	163
	Creating multiple original images concurrently	164
	About creating duplicate images concurrently	165
	Creating concurrent copies through the catalog node	167
	Creating concurrent copies using the basic duplication tab	168
	Creating concurrent multiple copies using the advanced duplication options	171
Chapter 8	Reporting	175
	About reports	175
	About generating reports	175
	Generating reports by using the Vault operator menu	176
	Generating reports by using the <code>vlteject</code> command	177
	Creating a Vault policy to generate reports	177
	About consolidating reports	179
	About consolidated reports in previous Vault releases	180
	Viewing Vault reports	181
	Vault report types	181
	Reports for media going off site	181
	Reports for media coming on-site	186
	Inventory reports	188
	Lost Media report	193
	Non-vaulted Images Exception report	194
	About the Iron Mountain FTP file	195

Chapter 9	Administering Vault	197
	About setting up email	197
	About administering access to Vault	198
	About the Vault Operator user group permissions	198
	About printing Vault and profile information	201
	Copying a profile	202
	About moving a vault to a different robot	202
	About changing volume pools and groups	202
	About NetBackup Vault session files	203
	About setting up Vault session log files	205
	Setting the duration of Vault session files	205
	Operational issue with disk-only option on Duplication tab	206
	Operational issues with the scope of the source volume group	206
Chapter 10	Using the menu user interface	207
	About using the menu interfaces	207
	About the Vault administration interface	207
	Vault operator menu interface	209
	Vault fields in bpdjobs output	210
Chapter 11	Troubleshooting	213
	About troubleshooting Vault	213
	About printing problems	214
	About errors returned by the Vault session	214
	About media that are not ejected	214
	About media that is missing in robot	215
	Reduplicating a bad or missing duplicate tape	215
	About the tape drive or robot offline	216
	No duplicate progress message	216
	About stopping bpvault	217
	About ejecting tapes that are in use	218
	About tapes not removed from the MAP	218
	Revaulting unexpired tapes	218
	Debug logs	219
	Setting the duration and level for log files	220
	Logs to accompany problem reports	221
Appendix A	Recovering from disasters	223
	About disaster recovery	223
	About what defines a disaster	224
	About the disaster recovery process	225

- About disaster recovery plans 225
- About recovery priorities 225
- About developing disaster recovery plans 226
- About testing disaster recovery plans 227
- About disaster recovery in the NetBackup Vault context 227
- About preparing for recovery 228
- About recovering NetBackup 230
- Recovering data and restoring backup images 230
- Archiving and recovering from a specific point in time 233

Appendix B Vault file and directory structure 235

- UNIX files and directories 235
- Windows files and directories 239

Index 245

About Vault

This chapter includes the following topics:

- About Vault
- About Vault and NetBackup functions
- About accessing NetBackup Vault
- About Vaulting original or duplicate images
- About the Vault process
- About how Vault uses volume groups and pools
- About NetBackup and Vault configuration
- About Vault management procedures

About Vault

Vault is an extension to NetBackup that automates selection and duplication of backup images. It also automates ejection of media for transfer to and from a separate, off-site storage facility. NetBackup Vault also generates reports to track the location and content of the media. Vault functionality does not have to be used only for disaster recovery. You can use Vault to manage the data and the media that you store off-site for regulatory archival purposes.

About Vault and NetBackup functions

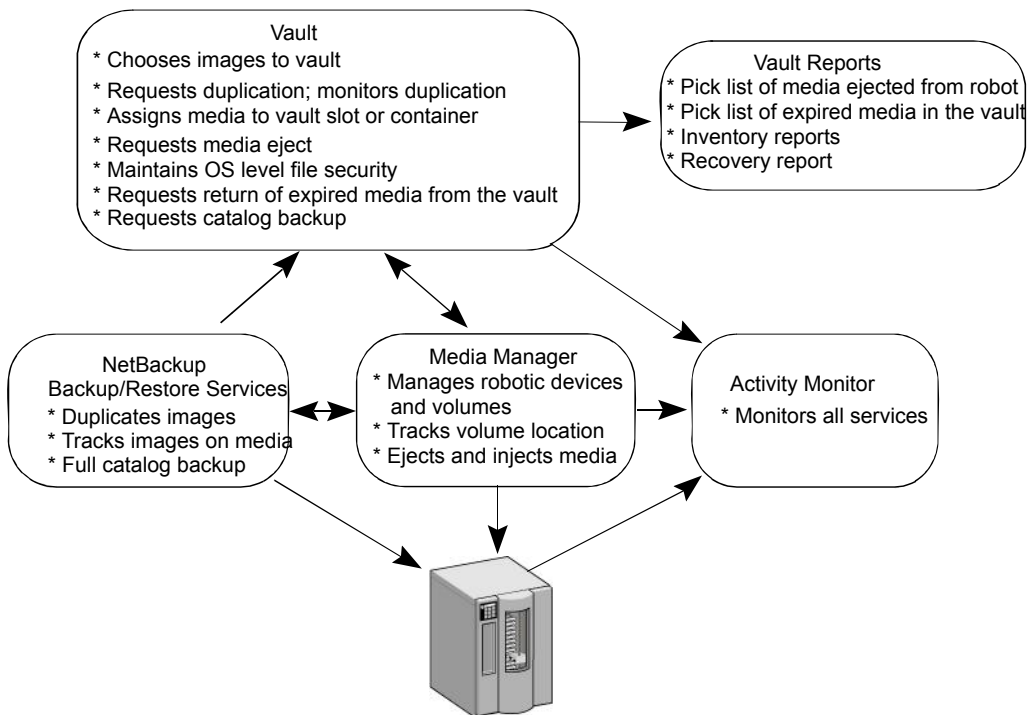
NetBackup Vault uses existing NetBackup functions for all operations, such as duplication of images, media control, reporting, and ejecting and injecting of tapes. Information from Vault is integrated with other NetBackup components and appears in the NetBackup Activity Monitor.

NetBackup Vault interacts with the following NetBackup services and catalogs:

- Media Manager, which manages robots and media
- The NetBackup catalog and the Media Manager database record of the images that have been vaulted
- The Media Manager database information which determines when expired media can be returned to the robot for reuse
- The Activity Monitor which displays the status of the Vault job

Figure 1-1 shows the NetBackup, Media Manager, and Vault relationships.

Figure 1-1 Services interaction diagram



About accessing NetBackup Vault

NetBackup Vault is installed on a NetBackup master server. The **Vault Management** node is seen on NetBackup Administration Console if you add the appropriate license key. You can add license keys during the installation or by using the **Help > License Keys**. You can use NetBackup Administration Console to configure and manage Vault.

Alternatively, you can manage Vault by using the following methods:

- Menu-based user interfaces
- Command line utilities

About Vaulting original or duplicate images

One of your most important choices is whether to send original or duplicate images off site. If you send original images off site, you do not have to duplicate images and therefore do not have to configure duplication.

Vault distinguishes between original images and duplicate images as follows:

- NetBackup creates original images during a backup job, including all of the copies that are created concurrently by a backup policy. NetBackup can create up to four copies of an image concurrently during the backup process; all are considered originals.
If you create multiple original backups in a NetBackup policy, assign the primary copy to the volume pool for the media that remains on site. And send a copy off site.
- Duplicate images are the copies that Vault creates. A Vault job reads the primary backup image and writes one or more duplicate images concurrently. The Vault job is separate from the NetBackup policy job.

About the Vault process

Vaulting is the process of sending backup images off site to a protected storage location.

For more information about the specific steps in a Vault process, see the following:

- See “About choosing backup images” on page 16.
- See “About duplicating backup images” on page 16.
- See “About backing up the NetBackup catalog” on page 16.
- See “About ejecting media” on page 17.
- See “About generating reports” on page 17.

A Vault job must select images (Choose Backups). The other steps are optional. You can separate the Vault tasks into separate jobs if desired that use different jobs to accomplish different tasks. For example, you can use one job to select and duplicate images daily, and another job to eject media and generate reports weekly.

Note: Injecting returned media back into the robot is a manual operation. The Vault reports include the media that should be recalled from the off-site location and injected into the robot.

The term *vault* refers to a logical entity that is associated with a particular robot. It also refers to the off-site storage location of a set of tapes.

About choosing backup images

The first step of the Vault process is to choose the backup images that are candidates to be transferred off site. You must configure the image selection for every Vault job. Vault uses the criteria in a Vault profile to determine which backup images are candidates to send off site. (A Vault profile is a set of rules for selecting images, duplicating images, and ejecting media.)

If you create multiple original images concurrently during a backup job, Vault can send original images off site (depending on the profile rules). If you duplicate images, Vault uses the primary backup image as the source image for the duplication operation. However, Vault duplicates from a nonprimary copy on disk if one exists. That duplication process improves performance.

Note: Vault does not select the SLP-managed images that are not lifecycle complete.

About duplicating backup images

The backup images that are candidates to be transferred off site are duplicated in the second step of the Vault process. This is the step for image duplication. Vault writes copies of the backup images on the media that you can eject and transfer off site.

Image duplication is optional. If you send your only backup image off site or create multiple original backup images and send one or more of those images off site, you do not have to duplicate images in Vault. Therefore, you do not have to configure the duplication step. However, you must write the original image to media in the off-site volume pool so it is ejected and transferred off site.

About backing up the NetBackup catalog

The NetBackup catalog is backed up in the third step of the Vault process. The NetBackup catalog consists of databases of information about the NetBackup configuration and any backups that have been performed. The information about backups includes records of the files and the media on which the files are stored, including information about media sent off-site. The catalogs also have information

about the media and the storage devices that are under the control of Media Manager.

Backing up the catalog is optional. However, vaulting a catalog backup with your data can help you recover from a disaster more efficiently. Vault creates its own catalog backup with up-to-date information.

Note: Vault does not duplicate the NetBackup catalog.

About ejecting media

The media that you then transfer to secure storage is ejected, often at a separate facility, in the fourth step of the Vault process. The Vault reporting facilities track the media that are ejected and are recalled from off-site storage for reuse after the images expire. The media can be ejected automatically by a scheduled Vault job or manually after the job has completed. The media can be ejected for each job individually or can be consolidated into a single eject operation for multiple Vault jobs.

About generating reports

Vault reports are generated in the fifth step of the Vault process. The reports track the media that Vault manages. You and your off-site storage vendor can use the reports to determine which media should be moved between your site and the off-site storage location. You can also determine when the moves should occur.

Reports can be generated as part of the Vault job or manually after the job is finished. Reports can be generated for each job individually or consolidated with a consolidated eject operation. Generating reports is optional.

About how Vault uses volume groups and pools

The volume groups identify where volumes reside. They are used as a tracking mechanism by Vault to determine where a volume is located. Volumes in a robotic volume group reside in a robot. During a Vault job, Vault searches the robotic volume group for the media that matches a profile's criteria. If media is found, Vault ejects that media and then moves it logically to an off-site volume group. A logical move means to change the volume attributes to show the new location. When a volume in off-site storage expires and is injected back into the robot, Vault moves it back into the robotic volume group.

Volume pools identify logical sets of volumes by usage. Vault uses them to determine if a volume should be ejected. Volume pools for images to be transferred off site

are known as off-site volume pools. When you create the images that you want to send off site, write them to media in an off-site volume pool. During a Vault job, Vault searches a robot for the images that match the selection criteria. If the media the images reside on are in an off-site volume pool, Vault ejects that media.

About NetBackup and Vault configuration

- Before you begin to use Vault, you must set up and configure NetBackup so that volume pools and policies are available to support Vault operations.
 - See “About configuring NetBackup Vault” on page 53.
- After configuring NetBackup for use with Vault, you must configure Vault robots and profiles.
 - See “About configuring Vault” on page 61.
- Read the best practices information to help determine how to set up and configure Vault most effectively for your environment.
 - See “About best practices” on page 27.

About Vault management procedures

Table 1-1 summarizes the operational procedures for Vault. The *NetBackup Vault Operator’s Guide* provides detailed information on day-to-day procedures.

Table 1-1 Vault management procedures

Operational procedure	Staff responsibilities
Configuration: review backup procedures and determine duplication capacity needed. Determine appropriate servers to run duplications and determine the appropriate time windows to run duplication. Configure Vault profiles. Review duplication windows for performance and throughput.	Determine the levels of duplication service on a per policy basis. Ensure sufficient hardware, software, and network capacity is available for duplication of backup images.
Choosing Backups: Vault incorporates the new criteria when choosing a backup for vaulting.	
Duplication: Set up Vault policies to run vault sessions on a schedule.	Monitor jobs to ensure that they start when scheduled.

Table 1-1 Vault management procedures (*continued*)

Operational procedure	Staff responsibilities
Determine media requirements and setup initial volume pool for duplication. Monitor volume pool usage.	Ensure that sufficient media are available for duplicates to run.
Monitoring: use the NetBackup Activity Monitor to determine progress. Set up links between log files and the email monitoring system and the paging notification.	Ensure that duplication jobs complete successfully. Ensure that errors are reported to appropriate personnel.
Vault reports: generate reports regularly to ensure that images are duplicated correctly. Compare report output with ejected and returned media.	Review production duplication cycle for thoroughness. Ensure which media goes to off-site and returning on-site matches reports.
Check duplication volume pools and catalog backup pools for available media.	Ensure sufficient media available for duplication.
Use Media Manager to expire or freeze tapes manually when needed for retrieval from the vault.	Recall the media that was not recalled during normal Vault operations.
Back up the catalog: Set up a schedule for Vault backup of image catalog. Ensure media available to store catalog.	Ensure that Vault catalog backup occurs.
Duplication capacity review: determine the capacity planning cycle, that includes the reaction time, costing factors, and new requirements.	Assist production support to help determine system, robotic, and network utilization rates (for example, disk capacity). Assist in defining requirements for the system infrastructure to use Vault effectively with other computing environment resources.
Recovery review: test recovery procedures regularly to ensure recovery of essential data from off-site storage.	Know procedures for restoring duplicated images. Know how to restore database catalog, backup software, and so on in case of disaster on NetBackup server(s).

Installing Vault

This chapter includes the following topics:

- About supported systems
- About supported robots
- About installing and configuring Vault on UNIX and Linux systems
- About installing and configuring Vault on Microsoft Windows systems

About supported systems

Vault runs on the same operating systems and versions and in the same clustering environments as NetBackup except as noted in the *NetBackup Release Notes*. The NetBackup restrictions and the limitations that relate to systems, clusters, and peripherals also apply to Vault. One exception is that Vault does not support standalone drives.

For information about supported systems and supported upgrade paths, see the *NetBackup Release Notes*.

About supported robots

Vault supports all Media Manager-supported robot types.

Vault also supports the robots that do not have media access ports and barcode readers. For best performance, and to avoid errors when you enter media IDs, Cohesity recommends that you use robots that have media access ports and barcode readers.

You can use Vault with the robots that do not have media access ports.

See “About media ejection” on page 102.

About installing and configuring Vault on UNIX and Linux systems

NetBackup Vault installs on a UNIX and Linux system when the NetBackup master server is installed or upgraded. No separate installation or upgrade procedure is required. However, you must enter a valid license key to use Vault and configure Vault.

See *NetBackup Installation Guide for UNIX and Linux*.

Adding a valid license key and configuring Vault for UNIX and Linux systems

To license NetBackup Vault on a UNIX or Linux system, the NetBackup master server must be installed and running on the UNIX or Linux computer. In addition, you must have a valid NetBackup Vault license key.

To add a valid license key and configure Vault for UNIX and Linux systems

- 1 To make sure a valid license key for NetBackup Vault is registered on the master server, enter the following command to list and add keys:

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

In a cluster environment, license NetBackup Vault on all nodes in the NetBackup cluster.

- 2 If you add the NetBackup Vault key after the NetBackup master server installation is complete, you must start the NetBackup Vault daemon. Perform the following command:

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

- 3 Use the following commands to start the NetBackup Administration Console and configure Vault.

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

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

To complete the Vault configuration, you must configure the appropriate NetBackup attributes that Vault uses and identify which NetBackup policies you want to use with Vault (or create new attributes to use with Vault). Please read the following chapters to develop an understanding of how Vault works and how best to configure Vault for your operations. You should configure the email address for notification of session status and enter alternate media server names, if appropriate.

See “About configuring Vault Management Properties” on page 64.

In a cluster environment, you can use the NetBackup Administration Console that is connected through the NetBackup virtual server name to configure Vault. You can do that regardless of which cluster server is currently active.

About upgrading NetBackup Vault on UNIX and Linux systems

On UNIX and Linux systems, NetBackup Vault is upgraded at the same time NetBackup is upgraded. To upgrade NetBackup Vault, follow the upgrade installation procedures for NetBackup.

See the *NetBackup Installation Guide*.

Deleting a Vault license key from UNIX and Linux systems

The NetBackup Vault software is not uninstalled. Rather, you deactivate NetBackup Vault by deleting the license key from the list of current NetBackup licenses. When the license key is deleted, Vault is no longer available for use. You can delete the

NetBackup Vault license key only if Vault was licensed with its own key, separate from the base NetBackup product license key.

Before you delete the NetBackup Vault license key, you should delete all Vault-specific items from NetBackup, such as volume pools, Vault policies, and so on. You can remove all Vault-specific configuration items by using the NetBackup Administration Console to delete them. Deleting the Vault configuration ensures that NetBackup does not include anything that was configured for Vault, such as volume pools.

To delete a Vault license key from UNIX and Linux systems

- 1** From the NetBackup Administration Console, select **Help > License Keys**
- 2** Select the Vault license key from the list of keys that is displayed in the **NetBackup License Keys** dialog box.

If NetBackup Vault was included as part of the base product key, performing the following step deletes your base key and you cannot use NetBackup. If you do not want to delete the NetBackup license key, do not continue.

- 3** Click **Delete**

About installing and configuring Vault on Microsoft Windows systems

NetBackup Vault is installed on a Windows system when NetBackup is installed; no separate installation procedure is required.

To use Vault, you must enter a license key as follows:

- Your license key may be a single key for the base NetBackup product and all NetBackup add-ons that you install, including Vault. If you have already installed NetBackup and entered the license key, Vault is already licensed.
- You may have a separate license specifically for the Vault option. If so, you must enter the Vault license key before you can use Vault.
See “Adding a Vault license on a Windows system” on page 25.

If you install a Vault in a cluster environment, you must license Vault on all systems in the cluster on which NetBackup master servers are installed.

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

Adding a Vault license on a Windows system

When preparing to license NetBackup Vault on a Windows system, the NetBackup master server must be installed and running on the Windows computer. Vault cannot be installed on a NetBackup media server or on a NetBackup client. In addition, you must have a valid NetBackup Vault license.

Note: If the license for NetBackup Vault was included in the license for the base NetBackup product, you do not have to perform this procedure.

To add a Vault license key on a Windows system

- 1 From the NetBackup Administration Console, choose **Help > License Keys**.
- 2 In the **NetBackup License Keys** dialog box, click **New**.
- 3 In the **License Key** dialog box, enter the NetBackup Vault license key.
- 4 Click **Add**.
- 5 Click **Close** to close the **NetBackup License Keys** dialog box.

To complete a new installation, you must configure the NetBackup attributes that Vault uses. You must also identify which NetBackup policies you want to use with Vault (or create new ones to use with Vault). Be sure to configure an email address for notification of sessions status and enter alternate media server names, if appropriate.

See “About configuring Vault Management Properties” on page 64.

About upgrading NetBackup Vault on a Windows system

On Windows systems, NetBackup Vault is upgraded at the same time NetBackup is upgraded. To upgrade NetBackup Vault, follow the upgrade installation procedures for NetBackup.

See the *NetBackup Installation Guide*.

Deleting the Vault license from a Windows system

The NetBackup Vault software is not uninstalled; rather, you deactivate Vault by deleting the license from the list of current NetBackup licenses. When the license is deleted, NetBackup Vault is no longer available for use. You can delete the Vault license only if Vault was licensed with its own license, separate from the base NetBackup product license.

Before you delete the Vault license, remove all Vault-specific configuration items by using the NetBackup Administration Console to delete them. Deleting the Vault

configuration ensures that NetBackup does not include anything that was configured for Vault, such as volume pools.

To delete the Vault license key from a Windows system

- 1** From the NetBackup Administration Console, select **Help > License Keys**.
- 2** Select the Vault license key from the list of keys that is displayed in the **NetBackup License Keys** dialog box.

If NetBackup Vault was included as part of the base product key, performing the following step deletes your base key and you cannot use NetBackup. If you do not want to delete the NetBackup license key, do not continue.

- 3** Click **Delete**.

Best Practices

This chapter includes the following topics:

- About best practices
- About vaulting paradigms
- About preferred vaulting strategies
- About how to ensure that data is vaulted
- About not Vaulting more than necessary
- About preparing for efficient recovery
- About media ejection recommendations
- About avoiding resource contention during duplication
- About how to avoid sending duplicates over the network
- About increasing duplication throughput
- About maximizing drive utilization during duplication
- About scratch volume pools
- About organizing reports
- About generating the lost media report regularly

About best practices

You can configure Vault to support how your computing environment or datacenter environment is set up and how it operates. A best-practice recommendation that may provide benefit for one environment may not provide the same benefit for

another. You should evaluate and implement any recommendations based on the benefit to your environment.

About vaulting paradigms

How you set up and name your vaults and profiles depends on your operations. For example, if you maintain a customer database and a payroll database, you may choose to organize your vaults by data usage and your profiles by time periods.

Example vaults and profiles that are organized by location and data type are as follows:

CustomerDB vault	Weekly
	Monthly
Payroll vault	Biweekly
	Monthly
	Yearly

Alternatively, if your operations are organized geographically, you can set up your vaults by location and your profiles by data type.

Example vaults and profiles that are organized by location and data type are as follows:

London vault	CustomerDB
	Payroll
Tokyo vault	CustomerDB
	Payroll

About preferred vaulting strategies

Several strategies can help you reduce resource and time contention when you back up your data and when you vault your backup media. Although these strategies may not be advantageous for all situations, they can be beneficial in many environments.

Cohesity recommends that you use one of the following strategies:

- Vault the original NetBackup backup media. Because you can produce multiple copies of images during a NetBackup policy job, fewer drives and less time may be required to create multiple original copies than duplicating media.
- Use disk staging. Send your backups to disk and then copy the data from disk to removable media. This strategy reduces the time that the backup process uses.

About Vault original backups

Cohesity recommends that you use a NetBackup policy to produce multiple, original backup images. You should then use a Vault profile to eject and transfer one or more of the original images off site.

In most situations, vaulting originals has the following advantages:

- Uses less drive-time than duplicating backup images from the original tapes. For example, a backup job that creates two originals of a backup image uses two drives: two units of drive time. Conversely, a backup job that creates one original image uses one drive, and a vault job that creates one duplicate of the original uses two drives: three units of drive time. Over time, duplicating backup images consumes more drive time than writing multiple originals during a backup job.
- Avoids configuring for duplication. In complex environments (such as with multiple media servers, multiple robots, or multiple retention period requirements), it can be difficult to configure the duplication steps of Vault profiles. It is possible to send large amounts of data over the network without careful configuration, although in storage area network (SAN) environments network traffic may not be an issue.
- Use VTL tapes. You can send your backups to VTL tapes with minimum retention period required for your recovery. You can then duplicate to physical tapes from VTL tapes and expire the VTL tape copy using the **Expire original tape backup images after xxx hour(s)** setting.

See the following information before you configure Vault, if you decide to create and vault original backups:

See “About vaulting original backups in a 24x7 environment” on page 34.

See “About avoiding vaulting partial images” on page 32.

About disk staging

Using disk staging for your backup jobs can help avoid resource contention between backup operations and Vault duplication operations. Disk staging is the process of

first writing the backup images to a disk storage unit during a NetBackup policy job and then writing the images to removable media during a Vault job.

Note: This topic is about using a disk storage unit as a destination for backup images, not about using a disk staging storage unit.

Some of the advantages of disk staging over tape-to-tape duplication are as follows:

Shortens backup time	Writing to disk is faster than writing to tape, so less time is needed for backing up.
Minimizes tape drive usage	Sending the original copy to tape and then duplicating to a second tape requires one drive to make the first copy and two drives (a read drive and a write drive) to make the second copy.
Reduces expense	Because disk access is fast and disk space is less expensive than tape drives, it is often advantageous to send your backups to disk.

You can schedule your Vault sessions to duplicate the original disk backup images to two (or more) media: one on-site volume and one off-site volume. Also, you can configure the Vault profile to free up the disk space automatically for the next round of backups.

About how to ensure that data is vaulted

When you set up NetBackup Vault, make sure that you configure it to vault all of the information that you want transferred off-site.

About overlapping the time window in the profile

To ensure that all data is vaulted, overlap the time window in the profile.

A Vault profile uses a time range as one of the criteria for choosing the backup images to be vaulted. Vault does not duplicate or eject a backup image that already has a copy in the off-site Volume Group; therefore, Vault does not process images that are already vaulted by a previous session. Perhaps more importantly, backups that were not processed if a previous session failed are processed when the profile runs again if the time window is long enough.

Therefore, configure the time window to be the sum of the following:

- The longest expected downtime for a server or robot
- Twice the length of the frequency at which the profile runs

For example, if you have a profile that duplicates images daily and your longest expected downtime is three days, configure the time window to be at least five days. If a robot fails and requires three days to repair, the next time the profile runs, it selects backup images that were not vaulted during the three-day downtime. Configuring the window to be longer, such as seven days, provides even more resiliency. A longer time window forces Vault to search a larger list of images for vault candidates. Although that consumes more processing time, the extra time may not be a problem in your environment because Vault is a batch process that does not demand immediate system response.

About the consequences of not overlapping the time window

When a vault session is delayed, some backup images may be missed if the time window does not allow Vault to select images from a wider time range. For example, suppose the time window for your daily profile extends from one day ago to zero days ago. On Tuesday, the robot has mechanical problems and the Vault profile fails. Consequently, Monday night's backups are not vaulted. On Wednesday, the robot is fixed. When the next Vault session begins on Wednesday, it only selects backup images that were created during the previous 24 hours. So Monday night's backups still are not vaulted. If the profile's time window had spanned more than one day, the session would have picked up both Monday night's and Tuesday night's backups.

About resolving multiple names for a single server

For every media server, you should add an entry on the **Alternate Media Server Names** tab of the **Vault Management Properties** dialog box. At a minimum, there should be, for each media server, an entry that contains both the abbreviated name and the fully qualified name. Also add any other names by which a media server has ever been known. Taking this action avoids a number of problems. For example, if you do not list alternate names for media servers, some images may not be recognized as a match for the criteria entered in the **Choose Backups** tab of the **Profile** dialog box and may therefore not get vaulted.

If you have multiple NIC cards in your server, make sure that the server name associated with each NIC card is listed in the **Alternate Media Server Names** tab when you configure a profile.

See "Alternate Media Server Names tab (Vault Management Properties)" on page 66.

Note: Alternate media servers apply to NetBackup Enterprise Server only.

About specifying a robotic volume group when configuring a Vault

Volumes are ejected only if they are in a robotic volume group and in one of the off-site volume pools that are specified on the profile **Eject** tab. If you want a volume to be ejected, ensure that it is in a robotic volume group and in one of the off-site volume pools that are specified on the profile **Eject** tab.

About multiple volume groups (Multiple robots)

A profile only ejects media from the robotic volume group of the vault to which the profile belongs, and a volume group cannot span robots (typically, a volume group identifies a specific robot). However, a profile can select images to duplicate that are in a different robot's volume group and in multiple volume groups. This ability is useful if you have backup images on multiple robots and want to duplicate those images on media in a robot from which the media is ejected.

If you use this configuration, configure it with care.

See "Alternative A: dedicated robot for Vault processing" on page 39.

About not Vaulting more than necessary

When you set up NetBackup Vault, do not select and transfer off-site more data than is necessary.

About sending only the intended backups off-site


When configuring your backup policies, do not assign backup images that are not intended to be moved off-site to volumes in an off-site volume pool. In some circumstances, Vault ejects a volume if it contains images that are not intended for off-site storage. For example, if volume ABC123 has three images from policy1 and three images from policy2, and policy1 is specified on the profile **Eject** tab, volume ABC123 is ejected even though it contains images from policy 2.

Use different volume pools for backup images you want to keep on site and for backup images you want to send to the vault. If you use the same volume pool for both, you vault the backup images that should remain on-site. Also, if you use the same volume pool for both, a deadlock situation may result if your Vault profile is duplicating images. It may attempt to read a backup image from the same tape to which it tries to write the image.

About avoiding vaulting partial images

Figure 3-1 shows how original backup tapes often begin and end with partial images.

Figure 3-1 Original backup media




If you eject and vault original backup media, that media may contain partial images.

To avoid vaulting partial images, use one of the following methods:

- Stop backup activity long enough to run Vault.
- If backup jobs are running, use the **Suspend Media for the Next Session** option on the profile **Eject** tab to suspend all media on which backups were written within the last day and then vault only those backups that are older than one day. No more backup images are written to that media, and that media is ready to be ejected.

Figure 3-2 shows what occurs during NetBackup and Vault operations when **Suspend Media for the Next Session** is used.

Figure 3-2 Suspending original backup media



Only use the **Suspend Media for the Next Session** option if you eject original backup media and want to avoid vaulting partial images. You should carefully consider whether to use the **Suspend Media for the Next Session** option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again, prolonging the length of time that is required to suspend the media. Therefore, some partial images on vaulted media may be acceptable. If you use this option, it is possible that the original backup media vaulted will not be full.

This option does not suspend media that is in use, such as media to which NetBackup is writing backup images.

Note: Vault only suspends media in off-site volume pools that are specified on the profile **Eject** tab.

About vaulting original backups in a 24x7 environment

If you use Vault in an environment in which backups can occur 24 hours a day, seven days a week, a profile may try to eject media to which backups are written. Because Vault cannot suspend media on which backups are currently written, an error occurs and partial images may be vaulted. The rest of the image is vaulted the next time the profile runs if that tape is not busy and the choose backups time window is large enough to select the image again.

To avoid such problems when vaulting originals, choose backups that were created a day or more ago and suspend the media to prevent writing to the media. (This recommendation assumes that your backups are complete by the time the Vault session runs.)

About preparing for efficient recovery

Preparing in advance can help you restore your data more quickly and easily. The following can help you prepare for recovery.

About Vault NetBackup catalog requirements and guidelines

Use Vault to vault the NetBackup catalogs. A current catalog backup is a critical component of an effective disaster recovery plan. Although you can rebuild the catalog by importing all of your backup media manually, it is a time-consuming process.

Vault requirements and guidelines are as follows:

- Perform the catalog backup step in Vault. Vault creates a new catalog backup with up-to-date information, it does not duplicate an existing NetBackup catalog backup. A NetBackup catalog backup is not a substitute for a Vault catalog backup. It does not include the latest information about duplicated media and media location.
- Use only one vault to do Vault catalog backup.
- Use a dedicated volume pool for Vault catalog backups.
- If you have a robot attached to the master server, use it for the Vault catalog backup. In most circumstances that master server creates the NetBackup catalog that remains on site.

See the discussion of NetBackup catalog backups in the *NetBackup Administrator's Guide, Volume I*.

- Retain the three most recent catalog backups. In most circumstances, you do not need to retain vaulted catalog backups for the same length of time that you retain other vaulted backup media. Although you only need one catalog backup in your off-site vault, maintaining the three most recent catalog backups in your off-site vault is a good practice for extra protection. The Recovery Report for Vault lists only the three most recent catalog backups in the off-site vault regardless of how many actually reside in the vault.
- To retain only the three most recent catalog backups, specify an appropriate retention level. Specify a level that dictates that older catalog backups expire and are recalled from off-site storage and only the three most recent catalog backups remain in off-site storage.

About naming conventions for volume pools and groups

How you name pools and groups can help you (and others) organize and more easily identify media if you have to recover data after a disaster.

Use the conventions as follows:

- For volume pools, try to identify the purpose or data in the pools. For example, Vaulted_Payroll, Vaulted_CustomerDB, 1_month_vault, and 7_year_vault are descriptive volume pool names.
- For Vault catalog backups, use an easily identified name for the catalog volume pool (such as Vault_Catalog_Backups).
- For off-site volume groups, use names that indicate the physical location of the data, such as offsite_volume_group.

About matching volume pools to data usage

Volumes are assigned to volume pools. To assist with recovery, create and use off-site volume pools that match your data usage (that is, the type of data). For example, if you maintain a customer database, you may want to restore all of your customer database at the same time when you recover from a disaster. Assign all of your customer database backup data to an off-site volume pool specifically for that data. Assign only backup images of the customer database to that off-site volume pool.

This volume pool (for example, Vaulted_CustomerDB) can correspond to all profiles within a logical vault or to a single profile, depending on your Vault environment configuration.

About the primary copy

The first (or only) original backup image is the primary backup copy. NetBackup always uses the primary copy for restore operations. Vault uses the primary copy for duplication operations (unless the primary copy is on removable media and another copy exists on disk). Ensure that primary copies on removable media remain on site in your robot. If the primary copy is off site, a user-initiated restore operation waits indefinitely for a mount of the off-site media.

If you create multiple original backups during a NetBackup policy job, do not assign the primary copy to an off-site volume pool unless you intend to send it off site. If you assign the primary copy to an off-site volume pool, it is ejected and is not available for restore or duplication operations.

If your Vault profile duplicates media and you send the first original off site, configure Vault to designate one of the duplicate images that remain on site as the primary copy.

About suspending vaulted media

Suspend unexpired media that is recalled and injected back into the robot so NetBackup does not write images to it. Suspending media before it is ejected also helps to prevent errors from ejecting media that is in use. Vault profile **Eject** tab options let you suspend the media that is ejected so you do not have to suspend it if it is recalled. You also can choose to suspend media before it is ejected so that partial images are not written to that media.

Suspend options available on the Vault profile **Eject** tab are as follows:

Suspend this Session's Media	To suspend media in the profile eject list for the current session. If you select Immediately , no more images are written to the media. If you select At Time of Eject , images may be written to the media until the media are ejected. Select At Time of Eject if you want the media sent off-site to be full.
-------------------------------------	--

Because **Suspend this Session's Media** operates on media in the eject list, it does not use more CPU cycles selecting media to suspend.

Suspend Media for the Next Session

To prevent partial images from being written onto media that contains images to be vaulted. Use this option only if you vault original images and want to avoid vaulting partial images on backup media.

You should carefully consider whether to use this option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again. Also, this option does not suspend media that is in use, such as media to which NetBackup is writing backup images. However, Vault does enable you to narrow the image search by entering a time window on **Eject** tab.

This option suspends duplicate media created by Vault; however, the **Suspend this Session's Media** option is a better choice for suspending duplicate media because it does not use CPU cycles to select media to suspend.

See "About avoiding vaulting partial images" on page 32.

Note: Vault only suspends media in off-site volume pools that are specified on the **Eject** tab.

About revaulting unexpired media

You should always revault media that was recalled from off-site storage and injected into the robot (for example, if you recall a volume to use for a restore operation). If you do not eject the media and transfer it to your off-site vault location, it is not available if media at your site is damaged.

About media ejection recommendations

Although Vault queues jobs, you can reduce the chances of error conditions from busy robots and reduce potential problems by doing the following:

- Eject media during a dedicated time period when no other inject or eject operations occur.
- Do not inject or eject other media while Vault is ejecting media.
- Do not inventory a robot while Vault is ejecting media.

About avoiding resource contention during duplication

Following are the resources that you should consider when you configure duplication in Vault:

- Time (that is, when the operations occur)
- Media used
- Robots and drives
- Bandwidth

Note: If you vault original backups, you do not have to use practices that avoid or reduce resource contention in Vault.

Various configurations can help you avoid resource contention. Also, a general principle that can help avoid resource contention is to wait until backups are completed before using Vault to duplicate or eject media.

About two processes trying to use the same drive

Careful configuration of your environment can help avoid resource contention during Vault duplication. Resource contention can occur when two processes try to use the same drive at the same time.

To avoid resource contention, use one of the following alternatives:

- Alternative A: dedicated robot for Vault processing
See “Alternative A: dedicated robot for Vault processing” on page 39.
- Alternative B: each robot as a Vault robot
See “Alternative B: each robot as a Vault robot” on page 40.
- Alternative C: one robot as both a backup and Vault robot
See “Alternative C: one robot as both a backup and Vault robot” on page 41.

These alternative configurations work well for multi-robot environments. They use available resources efficiently and are unlikely to cause resource allocation problems.

Note: If you vault original backups, you do not have to use practices that avoid or reduce Vault resource contention.

Alternative A: dedicated robot for Vault processing

In a multi-robot environment, dedicate one robot strictly for vault processing. The media in this robot contains only the duplicate backup copies that are to be ejected and sent to the off-site vault. This configuration works best in a storage area network (SAN) environment where all media servers have direct access to the vault robot. That way the duplication step does not send data over the network.

Two ways to achieve this configuration are as follows:

- Use a NetBackup policy to create multiple original backup images concurrently. Write the first backup image (the primary backup) to a storage unit that is not in the Vault robot. Write one of the other originals to the Vault robot and assign it to the off-site volume pool. Configure a Vault profile to eject all media in that vault's off-site volume pool. This configuration requires that all robots that are used be connected to the same NetBackup media server.
- Use Vault to duplicate images. Backup images are duplicated from all other robots to the Vault robot. Use one of the following alternatives to configure Vault to perform duplication:
 - On the **Duplication** tab of the **Profile** dialog box, do not select **Advanced Configuration** or **Alternate Read Server**. For each backup image, the media server that performed the backup also performs the duplication. All media servers send duplication data to the Destination Storage Unit media server. If the Destination Storage Unit media server is not the same as the media server that performed the backup, the data is sent over the network.
 - On the **Duplication** tab of the **Vault Profile** dialog box, specify the destination storage unit's media server as the Alternate Read Server but do not select **Advanced Configuration**. If the alternate read server also has access to all of the backup robots, no data is sent over the network.
 - On the **Choose Backups** tab of the **Profile** dialog box, specify All Media Servers in the **Media Servers** list. On the **Duplication** tab, select **Advanced Configuration**, select **Alternate Read Server**, then create an entry for each media server in your environment. To avoid sending duplication data over the network, for each media server entry specify the destination storage unit's media server as the alternate read server. That server must have access to all the robots that hold the source images so they will be duplicated. Ensure that the total number of write drives that are specified in the **Write Drives** column for each entry does not exceed the number of drives in the Vault robot.

If you use this alternative, do not use **Any Available** storage unit in your backup policies unless only your Vault storage units are set to **On Demand Only**. Using **Any Available** for other storage units may cause images not intended for off-site storage to be written to the Vault robot. You can achieve the same behavior provided

by **Any Available** storage unit by configuring your backup policies to use a storage unit group that includes all storage units except for the vault robot's. Note: if you use storage unit groups you cannot make multiple copies simultaneously.

Note: Alternate read servers apply to NetBackup Enterprise Server only.

Advantage	This configuration is most convenient for the operator, who can eject and inject tapes from only one robot, simplifying the tape rotation process.
Disadvantage	In a complex environment, this alternative can be difficult to configure if you want to avoid sending duplication data over the network.

Alternative B: each robot as a Vault robot

In a multi-robot environment, configure each backup robot to be a Vault robot. Each robot duplicates or ejects only backup images that were originally written to it.

You can do so in several ways, as follows:

- Use a NetBackup policy to create multiple original backups, assigning the copy to be vaulted to an off-site volume pool in any of the robots. For each robot, configure one vault and one profile that ejects the backups that were assigned to the off-site volume pool in that robot. Only backups on media in the off-site volume pools that are specified on the **Eject** tab and that meet the rest of the criteria specified in the profile are ejected.
- Use Vault to duplicate images. On the **Choose Backups** tab of the **Profile** dialog box, specify the robot to which the profile belongs in the **Source Volume Group** field. This limits the profile so that it duplicates only backup images that have their primary copy on media in this robot. Specify half of the available drives in the robot as read drives so that an equal number of read and write drives are available. Configure one such vault and profile for each robot. To avoid sending duplication data over the network, specify the media server of the destination storage unit as the Alternate Read Server.

Note: Alternate read servers apply to NetBackup Enterprise Server only.

Note: The destination storage unit must have at least two drives if that robot is used for both read and write functions.

These methods work well with backup policies that use **Any Available** storage unit. Using Vault to duplicate images also works well with storage unit groups if you make one copy only.

This configuration avoids resource contention when one profile attempts to duplicate images in multiple robots.

Alternative C: one robot as both a backup and Vault robot

In a multi-robot environment, configure all of the robots as backup robots and configure one of the backup robots as a Vault robot also. (One of the robots functions as both a backup robot and a Vault robot.) Configure one vault for the Vault robot, and in that vault configure one profile for each of the backup robots. In each profile, specify the backup robot in the **Source Volume Group** field of the **Choose Backups** tab and specify a destination storage unit that is in the Vault robot.

For example, if you have three robots that each have four drives, configure the three profiles as follows:

- In the profile for robot one (a backup robot only), specify the volume group in robot one as the Source Volume Group, specify four read drives, and specify a destination storage unit in robot three (robot three is the Vault robot). Images in robot one are read by four drives and written to four drives in robot three. Four duplication jobs run simultaneously.
- In the profile for robot two (a backup robot only), specify the volume group in robot two as the Source Volume Group, specify four read drives, and specify a destination storage unit in robot three. Images in robot two are read by four drives and written to four drives in robot three. Four duplication jobs run simultaneously.
- In the profile for robot three (a backup and Vault robot), specify the volume group in robot three as the Source Volume Group, specify two read drives, and specify a destination storage unit in robot three. Images in robot three are read by two drives and written to two drives. Two duplication jobs run simultaneously.

All images are duplicated to robot three and ejected from robot three.

This method works well with backup policies that use **Any Available** storage unit. Using **Any Available** storage unit in your backup policies sends backup images to media in any storage unit available. This configuration selects backup images on all the robots and duplicates them to the Vault robot.

Note: The destination robot must have at least two drives if that robot is used for both read and write functions.

About when the read drive is not in the Vault robot

The read drive does not have to be in the vault's robot. For configurations that include multiple media servers and multiple robots, we recommend that you seek advice from Cohesity Consulting Services.

About sharing resources with backup jobs

Vault duplication jobs compete with other process in NetBackup (such as regularly scheduled backups) for resources, including tape drives. If you want your Vault duplication jobs to obtain resources before other processes, assign a higher priority to the Vault jobs than is assigned to other NetBackup processes.

Vault duplication job priority is assigned for each profile in the **Duplication** tab.

Vault catalog backup jobs run at the priority assigned in the catalog backup policy unless you assign a different priority in the Vault catalog backup schedule **Multiple Copies** dialog box.

Priority for NetBackup jobs is assigned in the master server **Global Properties**.

In addition, using the Any Available storage unit for backup jobs can send some original backup images to the Vault robot. Subsequently, when Vault tries to duplicate those images, it requires a read drive and a write drive in the vault robot. If not enough drives are available, a deadlock condition can occur.

Cohesity recommends that you preview the images you want to duplicate before you run the Vault job, which shows you where the images are located and what kind of resources are required to duplicate them.

See "About previewing a Vault session" on page 117.

About load balancing

If it is feasible, Cohesity strongly recommends that you create multiple original backup images concurrently in your backup policies rather than using Vault duplication. The vault process is simpler and easier if you do not duplicate images.

If you cannot vault originals, several strategies can help you balance the load on your computing environment.

See "About profiles for both originals and duplicates" on page 42.

See "About spreading the workload" on page 43.

About profiles for both originals and duplicates

Vault can eject both original backups and duplicate images, so you can spread the load between backup jobs and Vault duplication jobs. For example, if your backup

window is too small to create multiple simultaneous copies of all backups, you can create multiple copies of some of the backups and only one copy of the other backups. Then configure a Vault profile to duplicate from the single original backups and eject both the original images and the duplicate images.

Some examples are as follows:

- NetBackup policy A creates multiple original copies and assigns one of the copies to an off-site volume pool.
- NetBackup policy B creates one copy and assigns it to an on-site volume pool.
- Your Vault profile is configured to copy backup images and assign the duplicate images to an off-site volume pool.

When you run that Vault profile, Vault copies backup images from NetBackup policy B only and does not duplicate images from policy A because an original already exists in the off-site volume pool. If you configured the profile for eject, it ejects both the copy of the original media from policy A and the duplicate media from policy B.

About spreading the workload

You can use Vault to duplicate backup images daily and eject volumes weekly. Duplication occurs every day rather than one day only, spreading the workload evenly throughout the week. The media remains in the robot until it is due to be collected by the vault vendor.

For example, if the vault vendor picks up the media every Friday, you can do the following:

- Configure a Vault profile to do duplication only, and configure a vault policy to run this profile every day of the week.
- Configure a second Vault profile to do the catalog backup and eject steps. This profile should use the same image selection criteria as the profile that duplicates images. Configure a Vault policy to run this profile before the vault vendor arrives on Friday.

This method for duplicating and ejecting media provides the added benefit of consolidated reports that are not organized by session.

About specifying different volume pools for source and destination

You should never configure a profile for duplication so that the source volume and the destination volume are in the same volume pool. That configuration will result in deadlock when NetBackup chooses the same tape as the source and the destination of the duplication operation. (This issue is a NetBackup limitation.)

About using a separate volume pool for each Vault

Jobs within the same vault are queued and then run when resources are available. However, if multiple profiles from different vaults run simultaneously and use the same off-site volume pool for duplication, those jobs could all pick the same target media. This selection would circumvent the queuing mechanism and cause undesirable results. (For example, it can cause a deadlock condition when multiple jobs try to use the same drive at the same time).

Therefore, you should configure Vault so that every vault has its own off-site volume pool.

About how to avoid sending duplicates over the network

Sending duplicate images over the network is not a problem if there is sufficient bandwidth. However, even a fiber optic storage area network (SAN) has only enough bandwidth for two or three duplication jobs at a time.

To avoid sending data over the network, use one of the following strategies:

- See “About creating originals concurrently” on page 44.
- See “About using alternate read server” on page 45.
- See “About using advanced duplication configuration” on page 45.
- See “About using storage units that specify a media server” on page 47.

About creating originals concurrently

One way to avoid sending data over the network with your Vault job is to create multiple original backup images concurrently during your scheduled backup jobs. This concurrent creation avoids the need for your Vault session to do duplication. In this scenario, Vault needs to eject the backup tapes only. Vault takes no significant resource time, except for the Catalog Backup step. (Catalog Backup is necessary to capture the changed volume database information for each vaulted tape.)

Suppose you want the on-site copy of your backups to go to one robot, and the offsite copy to go to another robot. If you create multiple backup images concurrently, all destination storage units must be on the same media server. Therefore, your media server needs a storage unit on both robots (one storage unit for your on-site copy and one for the offsite copy).

About using alternate read server

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

You can avoid sending data over the network during duplication by specifying an alternate read server if the alternate read server is as follows:

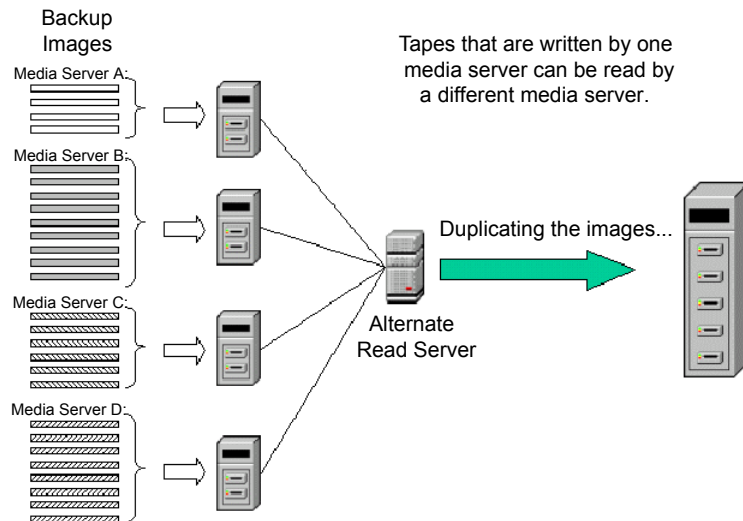
- Connected to the robot that has the original backups (the source volumes).
- Connected to the robot that contains the destination storage units.

Note: Alternate read servers apply to NetBackup Enterprise Server only.

Note: If the destination storage unit is not connected to the alternate read server, you send data over the network.

Figure 3-3 shows non-disk images that are written by media servers A, B, C, and D and are read by the alternate read server.

Figure 3-3 Tapes written by one media server can be read by another



About using advanced duplication configuration


If each media server has access to at least one unique drive in the destination robot, you can use advanced duplication to process each media server independently

and concurrently. (Note: all media from a single profile are ejected from the same robot.) You can do the same thing by configuring a separate profile for each media server rather than using advanced duplication configuration. However, multiple profiles within a single vault must run consecutively, so this may not allow you sufficient bandwidth.

Note: More than one media server applies to NetBackup Enterprise Server only.

Figure 3-4 shows that no alternate read server is used and each media server reads and duplicates its own backup images.

Figure 3-4 Each media server reads and duplicates its own backup image



You should use caution when you are specifying **All Media Servers**. For example, if you specify **All Media Servers** on the **Choose Backups** tab of a profile and also use **Advanced Configuration** on the **Duplication** tab, create an entry for each media server on the **Duplication** tab advanced configuration view.

If you list more media servers on the **Choose Backups** tab than on the **Duplication** tab, Vault assigns the images that are written by media servers that are not listed in the advanced view to the first media server that finishes its duplication job. If the first available media server is across the network, a large amount of data is sent over the network.

Another possible, though less problematic, consequence is that backup images from the media servers that are not configured for duplication may be duplicated by a different media server each time the profile is run.

About using storage units that specify a media server

NetBackup lets you create a storage unit without specifying a media server for that storage unit (that is, you can specify **Any Available** for the media server for the storage unit). When a job uses such a storage unit, NetBackup determines the media server to be used with the storage unit when a job runs.

If you specify a destination storage unit that uses **Any Available** media server for Vault duplication, NetBackup may choose a different media server for the duplication job than the source media server. If so, data is sent over the network.

Therefore, to avoid sending duplicates over the network, use storage units that specify a media server (that is, do not use storage units that are configured to use **Any Available** media server).

About increasing duplication throughput

Adding drives enables Vault to run multiple duplicate sessions concurrently. For each write drive, a separate duplication job (`bpduplicate`) is started.

The following topics provide information about multiple drive environments:

See “About basic multiple-drive configuration” on page 47.

See “About multiple-drive configuration that does not send data over network” on page 48.

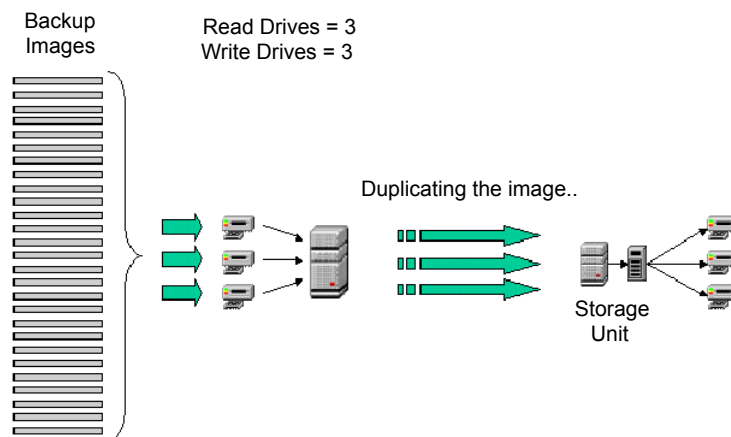
About basic multiple-drive configuration

In a basic multiple-drive configuration, there are an equal number of read and write drives, one master server, and one media server. The storage units are attached to the host on which the media server resides. A duplication process runs for each read and write drive pair. If the master server and media server reside on different hosts, media duplication data is transferred over a network.

Figure 3-5 shows a multiple-drive configuration performing a duplication process.

Note: Only NetBackup Enterprise Server allows a master server and media server to reside on different hosts.

Figure 3-5 Increasing throughput using multiple drives



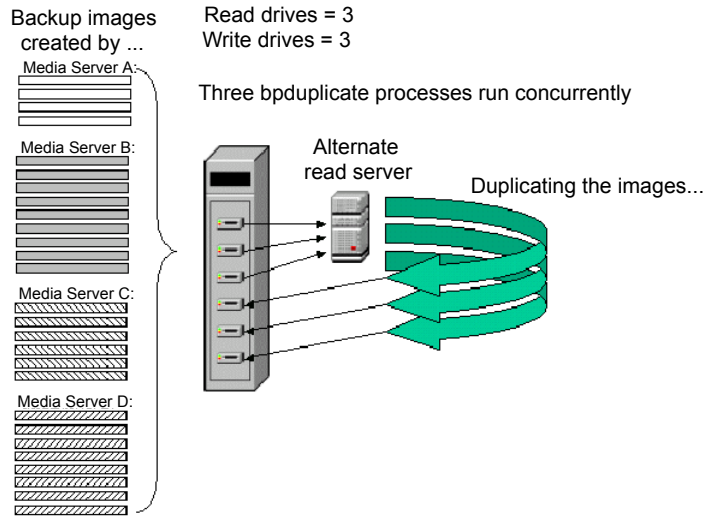
About multiple-drive configuration that does not send data over network

In a multiple-drive configuration that does not send data over the network, the configuration has an equal number of read and write drives, one master server, and multiple media servers. A separate duplication process runs for each read and write drive pair during a duplication operation. If you designate an alternate read server (media server A) for reading the images to duplicate and if the destination storage unit also resides on the alternate read server (media server A), no data is sent over the network.

Figure 3-6 shows a configuration where data is not sent over the network.

Note: Alternate read servers apply to NetBackup Enterprise Server only.

Figure 3-6 Multiple-drive configuration that does not send data over network



About maximizing drive utilization during duplication

To maximize drive utilization, Cohesity recommends that you do your duplication with as few Vault jobs as possible.

The more profiles you use, the less efficient the duplication process becomes. Drives are idle between the duplication steps of consecutive Vault jobs while Vault does all of its other processing (selecting images, backing up the catalog, and generating reports). It is much more efficient to use as few Vault profiles as possible for duplication. Therefore, if you can configure one Vault profile to duplicate all of your data, you reduce idle time and get the maximum utilization of your drives.

In Vault 5.0 and later, you can configure one Vault profile to create offsite copies with multiple, different retention. With this configuration, a single Vault profile can do all of your duplication, which keeps your drives spinning from the time of the first image to the last with no pause.

See “Assigning multiple retentions with one profile” on page 141.

About scratch volume pools

A scratch pool is an optional volume pool that you can use to ensure that volumes are allocated to the volume pools that need them. Media Manager moves volumes from a scratch pool to other pools that do not have volumes available, including Vault pools. Expired volumes are returned to the scratch pool automatically.

You can set up a scratch pool in two ways, as follows:

- Create a scratch pool and add all your volumes to it. Then create all the other volume pools but do not allocate any volumes to them. Media Manager then moves volumes from the scratch pool to the other volume pools as needed and returns the expired volumes to the scratch pool.
- Create your volume pools and allocate volumes to them. Then create a scratch pool and allocate volumes to it. Media Manager moves volumes between the scratch pool and the other volume pools as needed and returns the expired volumes to the scratch pool. This method may be the best option if you decide to add a scratch pool to an existing NetBackup configuration.

The scratch pool feature can affect reports for media coming on site. If you use a scratch pool, the Picking List for Vault, the off-site Inventory, and the All Media Inventory reports may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even though the reports may be for a specific Vault profile or session.

For information about configuring scratch pools, see the *NetBackup Administrator's Guide*.

About organizing reports

You should determine whether you want your Vault reports to group media by robot, by vault, or by profile. Your decision affects how you use your volume groups and volume pools.

Vault searches the off-site volume group for the media to include in the reports. It also uses the off-site volume pools for the same purpose. Therefore, you can use either the off-site volume group or the off-site volume pool(s) to organize media for each robot, vault, or profile.

About organizing reports by robot

To ensure that reports are organized by robot, all the vaults within each robot should use the same off-site volume group. That is, each robot has its own off-site volume group. This arrangement organizes reports by robot and maximizes the reuse of tapes. Media from one robot do not appear on the reports for another robot.

Reports do not seem consistent for an individual logical vault, but this strategy maximizes the frequency with which tapes are returned for reuse. Every time the Picking List for Vault report is generated for any profile within any vault for the robot, tapes from all profiles and logical vaults for that robot can be recalled for reuse (depending on how profiles share off-site volume pools).

About organizing reports by Vault

To ensure that the Vault reports include media for each vault, specify a separate off-site volume group for each vault within a robot (that is, each vault has its own off-site volume group) and a common off-site volume pool for all profiles within each vault (that is, all profiles in the vault use the same off-site volume pool). Doing so ensures that each report contains media from one vault.

About organizing reports by profile

If you want the reports to include only media for a single profile, use a separate off-site volume pool for each profile.

About the consequences of sharing an off-site volume group across multiple robots

If profiles from multiple robots share both an off-site volume group and one or more off-site volume pools, your vault vendor returns a group of tapes (for a single Picking List for Vault report) that were ejected from multiple robots. The operator needs to identify which tapes should be injected into each of the robots. If mistakes are made identifying and injecting tapes, you can inject the incorrect media and possibly the incorrect number of media into your robots.

About generating the lost media report regularly

You should generate the Lost Media Report regularly so you can recall media that has not been returned from the off-site vault vendor but which should have been returned.

Media can get stranded at the off-site vault for the following reasons:

- Frozen backup tapes never expire. A backup tape that does not expire does not appear on the Picking List for Vault and is not recalled from the vault.
- A backup tape appears on the Picking List for Vault and Distribution List for Robot only once. If a tape from that report is missed and is not returned to the robot, it never again is listed for recall.

About generating the lost media report regularly

- You change off-site volume group or pool names. For example, if you begin to use a new media type, you have to use a new volume pool name. If you change names, media may be stranded off-site because the Picking List for Vault is based on off-site volume pools and off-site volume groups. Media associated with the old names are not listed.

Cohesity recommends that you do not change or rename group or pool names. See “About changing volume pools and groups” on page 202.

How often you generate the Lost Media Report depends on your operations. Weekly or monthly may be often enough.

Configuring NetBackup Vault

This chapter includes the following topics:

- About configuring NetBackup Vault
- About off-site volume pools
- About creating catalog backup schedules for Vault
- About setting master server properties for Vault

About configuring NetBackup Vault

Before you configure Vault, you must do the following in NetBackup:

- Create off-site volume pools.
- Create a Vault catalog backup schedule.

You should review the best practices information. It can help you determine how to set up and configure Vault.

See “About best practices” on page 27.

You should be familiar with basic NetBackup concepts, such as volume pools and groups, policies, and storage units.

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

About off-site volume pools

Volume pools identify logical sets of volumes by usage. Vault uses volume pools to determine if a volume should be ejected. Volume pools for images to be

transferred off site are known as off-site volume pools. When you create the images that you want to send off site, write them to media in an off-site volume pool. During a Vault job, Vault searches a robot for the images that match the selection criteria. If the media that the images reside on are in an off-site volume pool, Vault ejects that media.

You need at least two dedicated volume pools such as the following:

Off-site Volume Pool	Vault ejects media from off-site volume pools. Assign the data that you want to transfer off site to media in an off-site volume pool. You can assign either original backup images that are created as part of a NetBackup policy job or duplicate images that are created by a vault job to the off-site volume pool. How many off-site volume pools you use depends on your operations.
Vault Catalog Backup Volume Pool	If you write the Vault catalog to removable media, you should use a volume pool that is dedicated for Vault catalog backups. When you configure the volume pool in Media Manager, ensure that the Catalog Backup attribute is set. You should use only one Vault catalog backup volume pool. Vault does not require a dedicated volume pool for its catalog backups. However, if you do not use one, regular NetBackup media or catalog media may be ejected.

Do not use the NetBackup volume pool for Vault media. Because the NetBackup volume pool is the default volume pool, if you use it for Vault you can send more data off-site than you intend.

See “Creating a volume pool” on page 54.

See “About naming conventions for volume pools and groups” on page 35.

See “About matching volume pools to data usage” on page 35.

See “About scratch volume pools” on page 50.

Creating a volume pool

Volume pools are configured in the **Media and Device Management > Media** node of the NetBackup Administration Console.

Ensure that the volume pools you create have sufficient volumes allocated to them (or to a scratch pool if one exists).

After a volume is assigned to an off-site volume pool, it remains in that pool and is used for rotation within that same pool (unless a scratch pool exists, in which case it is returned to the scratch pool).

See the *NetBackup Administrator's Guide, Volume I* and the NetBackup Administration help for more information about volume pools and allocating volumes to them.

To create a volume pool

- 1** In the NetBackup Administration Console, click **Media and Device Management > Media**.
- 2** Click **Actions > New > Volume Pool**.
- 3** From the **Add a New Volume Pool** dialog box, in the **Pool name** text box, enter a name for the new volume pool.

The name must be 20 characters or less and cannot contain any spaces or special characters.

- 4** In the **Description** text box, enter a brief description for the pool.
- 5** Specify the **Maximum number of partially full media**.

This option lets you specify the number of partially full media in the volume pool for each of the unique combinations of the following in that pool:

- Robot
- Drive Type
- Retention Level

This option does not apply to the None pool, catalog backup pools, or scratch volume pools.

The default value is zero, which does not limit the number of full media that are allowed in the pool.

- 6** Select the **Catalog backup pool** check box if you plan to use this volume pool to back up the NetBackup catalog. This check box creates a dedicated catalog backup pool to be used for catalog policies. A dedicated catalog volume pool reduces the number of tapes that are needed during catalog restores since catalog backup media are not mixed with other backup media.

About creating catalog backup schedules for Vault

NetBackup uses a special backup policy of type NBU-Catalog to perform catalog backups. To perform a Vault catalog backup, Vault uses a special schedule of type Vault Catalog Backup in an NBU-Catalog policy.

Before you can configure the catalog backup step in Vault, you must create a Vault Catalog Backup schedule in an NBU-Catalog policy.

See "Creating a Vault catalog backup schedule in an existing policy" on page 56.

See “Catalog backup tab (Profile dialog box)” on page 98.

See “About Vault NetBackup catalog requirements and guidelines” on page 34.

Creating a Vault catalog backup schedule in an existing policy

You can create a schedule in an existing catalog backup policy or create a new catalog backup policy and schedule. You can create more than one Vault Catalog Backup schedule per policy, and your NetBackup environment may have more than one NBU-Catalog policy.

You must specify the storage unit and, if the storage unit uses removable media, a volume pool for the Vault catalog backup.

If creating one copy of the catalog, you must do the following:

- Override the policy’s storage unit and select a storage unit.
- Override the policy’s volume pool and select the dedicated Vault catalog volume pool (for removable media only).

If making multiple copies, select the storage unit and the dedicated Vault catalog volume pool (removable media only).

You do not have to specify a volume pool for disk storage units.

Prerequisites are as follows:

- Create an NBU-Catalog type of backup policy.
See the *NetBackup Administrator’s Guide, Volume I*.

To create a Vault catalog backup schedule in an existing policy

- 1 In the NetBackup Administration Console, expand **NetBackup Management > Policies**.
- 2 From the middle-pane, double-click the policy name.
- 3 Select the **Schedules** tab.
- 4 Click **New**.
- 5 See the following topics for the options you can configure.
See “Vault catalog backup schedule configuration options” on page 56.

Vault catalog backup schedule configuration options

Table 4-1 shows the configuration options for the NetBackup Catalog policy **Schedules Attributes** tab.

Table 4-1 Schedules Attributes tab configuration options

Property	Description
Name	Enter the name of the schedule. Use a name that identifies it as a Vault catalog backup schedule.
Type of backup	Select Vault Catalog Backup.
Multiple copies	<p>To create multiple copies of the catalog, select Multiple copies, click Configure, and then select the appropriate attributes for each copy in the Configure Multiple Copies dialog box.</p> <p>The Configure Multiple Copies dialog box appears only if you select the Multiple Copies check box on the NetBackup Catalog Policy's Schedule Attributes tab, and then click Configure.</p>
Override policy storage unit	Select this option and then select the storage unit to use from the drop-down list.
Override policy volume pool	If the storage unit is on removable media, select this option. Then select the volume pool for off-site catalog backups (does not apply to disk storage units). If you use Media Manager storage units, use a dedicated off-site volume pool for Vault catalogs.
Retention	<p>Select the length of time before the catalog backup expires and the volumes are recalled from the off-site vault.</p> <p>After the Retention has passed, catalog backup media appear on the Picking List for Vault or Distribution List for Robot. Vault recalls that media so it is available to reuse as catalog backup media.</p>

The **Configure Multiple Copies** dialog box appears only if you select the **Multiple Copies** check box on an NBU-Catalog policy **Schedule Attributes** tab and then click **Configure**.

Use this dialog box to create multiple copies of a Vault catalog backup. For Media Manager storage units, all storage units must be connected to the same media server.

Table 4-2 describes the configuration options for the **Configure Multiple Copies** dialog box for catalog backup.

Table 4-2 Configure Multiple Copies dialog box options

Property	Description
Copies	Select or enter the number of copies to create concurrently. You can create up to four or the number of copies that are specified in the Maximum Backup Copies field for the NetBackup master server (if less than four). (Configured in NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes .) By default, the value is two: one original backup and one copy.
If This Copy Fails	The action to perform if a copy fails is Continue or Fail All Copies . If you choose Fail All Copies , the entire backup job fails and no copies are made. NetBackup automatically retries the job if time permits. The next time the backup window for the policy opens, NetBackup tries again to run the backup (regardless of the frequency of the schedule). NetBackup tries until the backup succeeds, although one or more backup windows may pass before the backup is successful.
Priority of Duplication Job	Specify the priority of the duplication jobs for the catalog copies, from 0 (lowest) to 99,999 (highest) priority. The job for each copy runs using this priority.
Retention	Select the length of time before the catalog backup expires and the volumes are recalled from the off-site vault. After the Retention passes, catalog backup media appears on the Picking List for Vault or Distribution List for Robot. Vault recalls that media so it is available to reuse as catalog backup media.
Storage Unit	Select the storage unit that contains the resources to which the catalog backup is written.
Volume Pool	If the storage unit is on removable media, select the volume pool for off-site catalog backups (does not apply to disk storage units). If you use Media Manager storage units, use a dedicated off-site volume pool for Vault catalogs.

About setting master server properties for Vault

Several NetBackup master server properties control some aspects of Vault.

Setting the maximum number of Vault jobs

Vault uses the **Maximum Vault Jobs** property as a threshold for queueing jobs.

The **Maximum Vault Jobs** property is configured on the NetBackup master server.

See the **Maximum Vault Jobs** property in the *NetBackup Administrator's Guide, Volume I*.

See "About running multiple sessions simultaneously" on page 116.

To set the maximum number of Vault jobs

- 1** In the NetBackup Administration Console, expand **NetBackup Management**.
- 2** Expand **Host Properties**.
- 3** Select **Master Server**.
- 4** In the right pane, select the master server and then select **Actions > Properties**.
- 5** Select **Global Attributes Properties**.
- 6** Specify the maximum number of vault jobs that can be active on the master server.

The greater the maximum number of vault jobs, the more system resources are used.

Configuring Vault

This chapter includes the following topics:

- About configuring Vault
- About Vault configuration
- About configuration methods
- About configuring Vault Management Properties
- Configuring robots in Vault
- Vault Robot dialog box options
- About creating a vault
- Media access ports dialog box
- Creating retention mappings
- About creating profiles
- Creating a profile
- Configuring a profile

About configuring Vault

When you configure Vault, you configure robots, vaults, and profiles. However, before you can configure Vault, you must first configure volume pools and a catalog backup schedule that Vault uses.

See “About configuring NetBackup Vault” on page 53.

Then, after you configure Vault profiles, you configure the policies to schedule when a Vault job runs.

See “About scheduling a Vault session” on page 112.

Before configuring NetBackup and Vault, review the best practices information. It can help you determine how to set up and configure Vault.

See “About best practices” on page 27.

About Vault configuration

Information about the general configuration for NetBackup is required so you can set up and use NetBackup Vault. Collect and record the appropriate information about the master servers and media servers, storage units, and robot information. That way the information is available when you begin to configure Vault.

Configuration information about master servers, media servers, and storage units

Collect the following information about master servers, media servers, and robotic devices, which are used in various configuration options in Vault:

Table 5-1 Configuration information

Property	Description
Master Server Host Name	The name of the host server on which the NetBackup master server and Vault are installed.
Operating System Level of Master Server	The release of the operating system residing on the system on which the NetBackup master server is installed.
Number of Media Servers	The number of media servers that are associated with the master server.
Media Server Name	<p>The name of each media server that controls the drives you want to use for the vault process. This server should also be bound to a storage unit within the NetBackup configuration. For NetBackup, all drives (of a given media type) that are attached to a server are defined as one storage unit, which is the recommended configuration for NetBackup.</p> <p>For every media server, configure alternate media server names.</p> <p>See “Adding alternate media server names” on page 67.</p>
Operating System Level of Media Servers	The release of the operating system on the host machines on which the NetBackup media server or servers are installed.
Types of Robotic Devices	The robotic devices that are associated with each media server. Use the appropriate NetBackup terminology to identify the devices (for example, TLD, ACS, TL8) or specify the actual hardware manufacturer and model names for each device.

Table 5-1 Configuration information (*continued*)

Property	Description
Storage Unit Name	The NetBackup storage units that are associated with each media server. You can use the <code>bpstulist -U</code> command to generate a list of existing storage units. Consider how many drives in each storage unit you want to use for vault sessions. You may choose to keep some drives available for restores or backups while duplication is running.
Number of Drives	The number of drives in each storage unit. Tape to tape duplication requires drives in pairs: one to read and one to write.

Robot information

Collect the following information for each robot. Although the following information is not required to configure a robot for Vault, it may help you plan your configuration so that you use resources efficiently.

The robot properties are as follows:

Table 5-2 Robot information

Property	Description
ACSLS Server	The name of the ACSLS server. StorageTek only.
ACS Number	The corresponding ACS number for this robot. You can obtain this information by using the Media Manager <code>tpconfig</code> or by using the ACSLS console commands such as <code>query acs all</code> or <code>query lsm all</code> . StorageTek only.
LSM Number	The corresponding LSM number for this robot. You can obtain this information by using the Media Manager <code>tpconfig</code> command or by using the ACSLS console commands such as <code>query acs all</code> or <code>query lsm all</code> . StorageTek only.
MAP Capacity	The capacity of the media access port (also known as cartridge access port). On StorageTek systems, you can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console.
MAP Numbers	The identifiers for the media access port. On StorageTek systems, you can obtain this information by using the ACSLS command <code>query cap all</code> from the ACSLS console.

About configuration methods

You can use the NetBackup Administration Console to configure Vault. Alternatively, you can use the Vault Administration menu user interface on UNIX systems (initiated

by the `vltaadm` command from a terminal window). These instructions describe using the NetBackup Administration Console to configure Vault.

In some circumstances, you may have to use the Vault Administration menu interface to configure Vault, as follows:

- You have to connect to the UNIX system on which the NetBackup master server is installed from a remote system that does not have the NetBackup Administration Console. For example, if you have to connect to your network by using a dial-up connection over a telephone line, you may have to use a terminal window and use the Vault Administration interface.

See “About the Vault administration interface” on page 207.

The NetBackup Vault Manager (`nbvault`) manages Vault activity and arbitrates access to the Vault robot, vault, and profile configuration information. The NetBackup Vault Manager must be running at all times so Vault functions correctly. Because NetBackup Vault Manager arbitrates access, you can run more than one instance of the NetBackup Administration Console. If one instance of an administration interface or Vault command tries to change configuration information while another instance is changing information, Vault prompts the user to reload the information by using the **Refresh** option.

Versions of Vault earlier than 6.0 do not use the NetBackup Vault Manager to arbitrate access to the Vault configuration and are not supported with this version of Vault.

About configuring Vault Management Properties

Vault Management Properties specify email addresses for event notification, alternate media server names, report properties for all vaults, and retention level mappings for all vaults.

Configure **Vault Management Properties** using the following tabs in the dialog box:

- See “General tab (Vault Management Properties)” on page 64.
- See “Alternate Media Server Names tab (Vault Management Properties)” on page 66.
- See “Retention Mappings tab (Vault Management Properties)” on page 68.
- See “Reports tab (Vault Management Properties)” on page 69.

General tab (Vault Management Properties)

The options on the **General** tab are described as follows:

Table 5-3 General tab configuration options

Property	Description
Email address for notification of session status	<p>An email notification is sent at the end of each vault session. It provides a summary of the vault session, in the form of a <code>summary.log</code> file, and the status of the operation. The subject line of the email message is formatted as follow:</p> <p>Vault Status <i>status_code</i> [<i>robot_number/vault/profile</i>] (<i>MasterServer</i>)</p> <p>By default, the email is sent to the root or administrator user account on the system on which the NetBackup master server is installed. If you enter email addresses in the E-mail address for notification of session status field, email is sent to those addresses rather than to the root user. You cannot disable notification of session status.</p> <p>To enter more than one address, separate the addresses with commas.</p>
Email address for eject notification	<p>An eject notification is sent to the email addresses entered in the Email address for eject notification field when the eject begins (includes a list of the media to be ejected) and when the eject is completed.</p> <p>Eject notification is configured for each profile on the Eject tab, for each robot on the Vault Robot dialog box, and globally for Vault on the Vault Management Properties dialog box General tab. Vault sends the notification to the first email addresses that are found in that order. You can configure different addresses in each place.</p> <p>To enter more than one address, separate the addresses with commas.</p>
Eject media, sort by	<p>You can select whether to eject media alphabetically or by expiration date. By default, Vault ejects media alphabetically.</p>
Lookback days for media going offsite reports	<p>You can enter the number of days before the day a profile runs to search for images to include in media going off-site reports. This setting reduces the amount of time to generate reports because Vault searches fewer image database records to determine which images are on the ejected media. By default, Vault searches the entire image database.</p> <p>Specifying a value does not affect whether media are ejected and vaulted. However, if a volume is ejected that has an image on it older than the period you specify, that image is not listed on the media going off-site reports.</p>

For procedures related to configuring the **General** tab options on the **Vault Management Properties** dialog see the following topic.

See “Configuring Vault Management Properties on the General tab” on page 65.

Configuring Vault Management Properties on the General tab

In the **Vault Management Properties** dialog, use the **General** tab to configure the following:

- The email address for session status notification.

- The email address for eject notification for all profiles.
- The sort order for ejected media.
- The reporting period for media going off-site reports.
See “About setting up email” on page 197.

To configure general Vault management properties on the General tab

- 1 From within Vault Management, select **Vault Management Properties** on the **Actions** menu.
- 2 In the **Vault Management Properties** dialog, select the **General** tab.
- 3 Enter information or select options as appropriate.

See “General tab (Vault Management Properties)” on page 64.

Alternate Media Server Names tab (Vault Management Properties)

Use the **Alternate Media Server Names** tab (Vault Management Properties) to add alternate names of NetBackup media servers.

Adding alternate names for media servers simplifies configuration and helps ensure that all images eligible to be vaulted are chosen. Vault expands any occurrence of one of the names in a server name group to include all of the names in the group.

For every media server, add the fully qualified name, the short name, every name that is used by storage units that refer to it, and any other names by which a media server has been known. If you have multiple network interface cards (NICs) in the server, all server names associated with each NIC.

You also can create a server name group that includes different servers. Then, in the **Media Servers** field in the **Choose Backups** tab of the **Profile** dialog box, you only have to specify the server name group rather than the individual servers. This use of the **Alternate Media Server Names** dialog box lets you use one name to specify more than one server. This name is useful if you want to duplicate images from multiple servers.

If you use the default, all media servers, for all of your vaults, you do not have to specify alternate media server names.

About alternative media server names

A media server can have more than one name. For example, a server can have a fully qualified name, a short name, and more than one network interface card, each of which has its own name. If a media server has more than one storage unit, each storage unit can use a different name for that media server.

If a media server has more than one name, images backed up by it can be identified by an alternate name. If you specify only one of the names of that media server, images that are identified by the other names are not vaulted.

The Choose backups configuration is simplified if you specify media servers (that is, specify something other than the default: all media servers). If you add alternate media server names, you only have to specify one of those names in the **Media Servers** field of the **Profile** dialog box **Choose Backups** tab. If you do not add alternate media server names, you must specify all the names that are associated with each media server on the **Choose Backups** tab.

About alternate media server names considerations

Be aware of the following concerns associated with alternate media server names:

- You must have enough drives in the specified destination storage unit to keep up with the demand for duplication. If you do not, you risk a deadlock situation.
- The specified media servers must have access to the destination storage unit. If not, you risk a deadlock situation and your Vault job fails. To prevent this situation, use the **Media Servers** criterion on the **Choose Backups** tab to ensure that only backups from certain media servers are selected.
- If multiple duplication rules use different media server names that are part of a server name group, Vault processes only the first duplication rule. Successive rules do not get processed. Also, because the media server name for the duplication rule is expanded to include all media server names in the group, all images written by all storage units that use those media server names are processed by the first duplication rule that uses any name from the group. All images are processed, but by the first duplication rule only.
- Your configuration can send data over the network, depending on the media server(s) in use.

Cohesity recommends that you specify only one destination storage unit per server. If you specify more than one, you may create a problem because Vault does not have a mechanism to choose to which destination storage unit to send the duplicate images.

Adding alternate media server names

Use the following procedure to add alternate media server names.

To add alternate media server names

- 1 From within Vault Management, select **Vault Management Properties** from the **Actions** menu.
- 2 In the **Vault Management Properties** dialog, select the **Alternate Media Server Names** tab.
- 3 In the field below the **Media Server Names** window, enter all the alternate names for the media server, separated by commas. Then click **Add**.
 - To remove a media server name group you previously added, select it and click **Delete**.
 - To change a name group you previously added, select it and click **Change**.Each server name group should occupy one line in the **Media Server Names** window.
- 4 Click **OK**.

Retention Mappings tab (Vault Management Properties)

Global retention mappings.

In the **Vault Management Properties** dialog box, use the **Retention Mappings** tab to configure alternative retentions for all vaults.

See “Configuring global retention mappings” on page 68.

See “About retention mapping” on page 69.

Configuring global retention mappings

Use the following procedure to configure global retention mappings using the **Retention Mappings** tab.

To configure global retention mappings

- 1 Left-click in the field you want to change in the **Vault Retention Level** column.
- 2 Select a retention level from the drop-down list.
- 3 Repeat for each level you want to change.
- 4 Click **OK**.

See “About retention mapping” on page 69.

See “Assigning multiple retentions with one profile” on page 141.

About retention mapping

Retention mapping lets you assign a retention level to a duplicate image that is based on the retention level of the original image. For example, if the retention of an original image is two weeks, you can configure the mapping so that the duplicate image that is transferred off-site has a retention level of seven years.

You configure retention mappings in the following places:

- The **Vault Management Properties** dialog box **Retention Mappings** tab (global)
- The **Vault** dialog box **Retention Mappings** tab (vault specific)

By default, each retention level maps to itself (that is, retention level 0 maps to 0, 1 maps to 1, and so on).

To use the retention mappings, you must specify **Use Mappings** for the retention level during duplication. You can specify normal retention calculation for some duplication rules and alternative retention mappings for other duplication rules. Vault uses the retention mappings in specific-to-global order. If vault-specific retention mappings do not exist, Vault uses the global retention mappings.

The retention level for a duplicate image is based on the retention level of the primary backup image. The retention period begins on the date the primary backup image was created, not on the date the duplicate image was created.

If the backup policy that created the primary backup image no longer exists, duplication of that image fails and the job continues but reports a status 306 (vault duplication partially succeeded).

Retention mapping applies to duplication only; it does not apply if you vault original NetBackup images.

Reports tab (Vault Management Properties)

Use the **Reports** tab (Vault Management Properties) to configure the following items:

- Each report that you want generated when a profile runs.
- Customize report titles.
- The destinations for each report (email, printer, and location to save).

The values that are configured on this tab are propagated to the **Reports** tab of each profile. You can override the values you configure on this tab for any report on any profile.

See “Changing report properties” on page 70.

Changing report properties

Use the following procedure to change report properties on the **Change Report Properties** dialog.

To change report properties

- 1 Double-click a report.
- 2 In the **Change Report Properties** dialog, select options and enter information as necessary.

If you change a title, the new title appears on the **Reports** tab and in the **Report Type** list box when you view Vault reports in the Administration Console.

If you consolidate your reports and also change titles, use the same title for all profiles whose reports are consolidated. The title is printed on the reports and appears in the email subject line if you email the reports.

See “About organizing reports” on page 50.

See “About reports” on page 175.

Configuring robots in Vault

Use the **Vault Robot** dialog box to specify and configure the robots from which Vault ejects media. Vault robots contain the media that has images to be stored off-site. That media is ejected so it can be transferred to the vault. The images can be original images that are created during a backup job or duplicate images that are created by a Vault duplication job.

You can select any robots that are recognized by NetBackup and that have storage units associated with them. NetBackup assigns a number to each robot that it recognizes, and eligible robots are recognized by Vault.

To configure a robot in Vault

- 1 In the NetBackup Administration Console, select **Vault Management**.
- 2 Open the **Actions** menu and select **New > New Vault Robot**.
- 3 In the **New Vault Robot** dialog box, enter information or select values as appropriate.

See “Vault Robot dialog box options” on page 70.

Vault Robot dialog box options

Table 5-4 contains the field descriptions of the **New Vault Robot** dialog box.

Table 5-4 Robot configuration information

Property	Description
Robot Number	The robot number that is assigned by Media Manager. Media Manager assigns a number to each robot that it recognizes, and eligible robots are recognized by Vault. Based on the robot number that you select, the other fields are filled in automatically.
Robot Name	The name of the robot. The name is configured in Media Manager, and Vault uses that information to populate the Robot Name field.
Robot Type	The robot type as configured in Media Manager. Vault uses that information to populate the Robot Type field.
Robot Control Host	The name of the host that controls the robot. Enter the name of the media server that controls the robot.
Use email address from Vault Management Properties for eject notification	<p>Select to use the global eject notification email address or enter email addressees, separated by commas, semicolons, or spaces, that receive notification when eject begins and ends.</p> <p>Eject notification is configured for each profile on the Eject tab, for each robot on the Vault Robot dialog box, and globally for Vault on the Vault Management Properties dialog box General tab. Vault sends the notification to the first email addresses found in that order. You can configure different addresses in each place.</p>

See “Configuring robots in Vault” on page 70.

About creating a vault

After you configure robots, you can create and configure vaults with the **Vault** dialog box.

- See “Creating a Vault” on page 71.
- See “About configuring Vault dialog box attributes” on page 72.

See “About how Vault uses volume groups and pools” on page 17.

Creating a Vault

The following are the requirements for creating a vault:

- At least one robot must be configured already in Vault.

- A robot may contain multiple vaults, but a vault cannot span robots. Therefore, if you configured three TLD robots for Vault (not connected with pass-through devices), you must define at least three logical vaults, one for each TLD robot.
- Volumes in a vault must have the same density. If a robot has volumes of different density and you want to use all of those volumes for Vault, that robot must have a separate vault for each volume density.

Use the following procedure to create a vault.

To create a vault

- 1 In the NetBackup Administration Console, expand **Vault Management**.
- 2 Select a robot in the **Vault Management** tree.
- 3 From the **Actions** menu, choose **New > New Vault**.
- 4 In the **Configuring Vault** dialog box attributes, on the **Vault Attributes** tab, enter or select values for each field.

See “About configuring Vault dialog box attributes” on page 72.
- 5 On the **Retention Mappings** tab, enter or select values for each field.
- 6 Click **OK**.

About configuring Vault dialog box attributes

A vault is a logical entity that refers to a collection of removable media drives (usually tape drives) within a robot. You can use vaults to organize the data that is going off-site. For example, you can use one vault for payroll data and another vault for customer data.

You can configure a vault by using the tabs of the Vault dialog box:

- See “Vault Attributes tab options (Vault dialog box)” on page 72.
- See “Creating retention mappings” on page 74.

Vault Attributes tab options (Vault dialog box)

Use the **Vault Attributes** tab in the **Vault** dialog box to configure the attributes of a vault.

If you are configuring a vault in an ACS robot, you also can configure the media access ports (MAPs) to use for eject operations.

See “About ACS MAP” on page 103.

See “About vaulting paradigms” on page 28.

See “About preferred vaulting strategies” on page 28.

See “About naming conventions for volume pools and groups” on page 35.

Table 5-5 lists the **Vault** dialog box configuration options that you can configure in the **Vault** dialog box.

Table 5-5 Vault dialog box configuration options

Property	Description
Change	For ACS robots only, the option that is used to configure media access ports for eject operations. If you click Change , the Media Access Ports dialog box appears, in which you can add or remove MAPs from the Media Access Ports to Use list.
Containers of Many Media	Select if your media is stored in containers at your off-site storage location.
Customer ID	Your customer identification if you selected Iron Mountain as your vault vendor. You may have a separate customer ID for each logical vault.
First Off-site Slot ID	<p>The ID of the first slot in the off-site vault. This usually is provided by your vault vendor. Off-site slot IDs are often used by the vault vendor to track media. If your vendor does not use these identifiers, you can use the default first off-site slot ID of 1. Off-site slot IDs are unique only within a given vault.</p> <p>Slot IDs are assigned consecutively from the starting slot number. Ensure that the number of media in the vault does not exceed the range of slot IDs assigned by the vault vendor. With every session, Vault starts with the off-site slot ID and counts upward, looking for slots that are no longer in use. Vault always fills in the gaps with newly vaulted media.</p> <p>If multiple vaults are defined for the same vault vendor, divide the range of assigned slots between the various vaults. For example, if the vault vendor assigned the range 1-2000 and you defined three vaults for this vault vendor, then you can assign range 1-499 to vault 1, 500-999 to vault 2, and 1000-2000 to vault 3, assuming that vault 3 has the maximum number of tapes to vault.</p>
Media Access Ports to Use	<p>For ACS robots only, the media access ports (MAPs) to use for media ejection for the current vault. To select or change MAPs to use, click Change.</p> <p>In the About the media access ports dialog box, select the MAPs to use.</p>
Off-site Volume Group	<p>The name of the off-site volume group. The Off-site Volume Group indicates that media are in off-site storage. The name should describe the data, the vault vendor, the vault location, or a combination thereof so you can easily identify the volume group. Vault moves each piece of ejected media from the Robotic Volume Group into a standalone volume group (that is, a volume group that is not under the control of the robot). If the Off-site Volume Group does not exist, it is created during the vault session. The off-site volume group name may contain up to 25 characters.</p> <p>If the off-site volume group does not exist, it is created during the vault session.</p>

Table 5-5 Vault dialog box configuration options (*continued*)

Property	Description
Robotic Volume Group	The name of the volume group associated with the robot for this vault. The Robotic Volume Group is the group that indicates that media resides in a robot. Usually, NetBackup creates a robotic volume group when media are added to a robot. A robotic library can contain volumes from more than one volume group, so a robot can have more than one robotic volume group name associated with it.
Slots for Individual Media	Select if your media is stored in slots at your off-site storage location. If you select slots, you must complete the First Off-site Slot ID field.
Vault Name	<p>The name of the vault. The name should reflect its purpose. For example, if you are creating a vault primarily to duplicate and vault records from the finance department, you might call the vault Finance. The vault name may contain up to 25 characters.</p> <p>Vault names are case sensitive.</p> <p>Note: Directory names are not case sensitive on Microsoft Windows systems. Therefore, session directories are created in the same <code>vault\sessions\vault_name</code> directory for two or more vaults that have names that differ only in case.</p>
Vault Vendor	<p>The name of your off-site vault vendor (for example, Iron Mountain). If you select Iron Mountain, you also can configure Vault to put media lists into a file formatted in compliance with Iron Mountain's electronic processing specification. You can then send this file to Iron Mountain for electronic processing of the media lists.</p> <p>See "Reports tab (Profile dialog box)" on page 107.</p>

Media access ports dialog box

This dialog box applies to ACS robots in NetBackup Enterprise Server only. Use it to configure the media access ports (MAPs) to use for eject operations.

Creating retention mappings

For Vault-specific retention mappings, use the **Retention Mappings** tab in the **Vault** dialog box to configure alternative retentions for a specific vault.

To configure retention mappings for all vaults, see the **Configuring retention mappings** tab on the **Vault Management Properties** dialog box.

See "Retention Mappings tab (Vault Management Properties)" on page 68.

When you open the **Retention Mappings** tab on the **Vault** dialog box, the **Use retention mappings from Vault Management Properties** option is selected by

default, which populates this tab with the values from the **Retention Mappings** tab on the **Vault Management Properties** dialog box.

To configure retention mappings

- 1 If **Use retention mappings from Vault Management Properties** is selected, clear the check box by clicking it.
- 2 Left-click in the field you want to change in the **Vault Retention Level** column.
- 3 Select a retention level from the drop-down list.
- 4 Repeat for each level you want to change.
- 5 Click **OK**.

See “About retention mapping” on page 69.

See “Assigning multiple retentions with one profile” on page 141.

About creating profiles

After you configure vaults, you can create and configure profiles with the **Profile** dialog box.

See the following to create and configure Vault profiles:

- See “Profile dialog box” on page 75.
- See “About the number of profiles required” on page 76.
- See “Creating a profile” on page 76.

Profile dialog box

A Vault profile is a template for a vault job; it contains the rules for selecting, duplicating, and ejecting media. A profile is associated with a specific vault, and at least one profile must exist for every vault. A vault can contain multiple profiles, although two profiles within the same vault cannot run simultaneously. Two different profiles can run simultaneously if each profile is in a different vault and if each profile uses a different off-site volume pool.

All profiles select images (that is, Choose Backups).

You must select at least one of the following profile options when you create a new Vault profile:

- **Duplication**
- **Catalog Backup**
- **Eject**

The other options are optional so you can separate the Vault tasks into separate jobs if desired, using different jobs to accomplish different tasks. For example, you can use one job to select and duplicate images daily, and another job to eject media and generate reports weekly.

You can select or deselect any of these profile options at any time during the configuration process.

After you create a profile, use a Vault policy to schedule when it should run.

See “About scheduling a Vault session” on page 112.

About the number of profiles required

The number of profiles that are required depends on your operations. If you have more than one Vault, you have more than one profile.

To determine the number of profiles that you need, use the following criteria:

- If you duplicate and eject media on a regular schedule (such as daily or weekly), you may require only one profile.
- If you duplicate images daily and eject weekly, you require two profiles, one to duplicate and one to eject media and generate reports.
- If you vault original images, you may require only one profile to choose backups, eject media, and generate reports.

Cohesity recommends that you do your duplication with as few Vault jobs as possible.

See “About maximizing drive utilization during duplication” on page 49.

Creating a profile

Once you have determined the number of profiles needed to accommodate the vault operations, you are ready to create the first profile using the **New Profile** dialog box.

To create a profile

- 1 Select a vault in the NetBackup Administration Console. From the **Actions** menu, choose **New > New Profile**.
- 2 In the **New Profile** dialog box, in the **Name** field, type a name for the profile. Cohesity recommends that you use descriptive names. Profile names are case-sensitive.

- 3 Select the steps you want this profile to perform.

You must select at least one step. However, you can change the selections when you configure the profile. Because you must always configure the choose backups step, it is not displayed on this dialog box.

- 4 Click **OK**.

The **New Profile: profile name** dialog box appears indicating that you are ready to configure the profile.

Configuring a profile

After you create a profile, the **New Profile: profile name** dialog box appears.

The **New Profile** dialog box includes the following five tabs:

Table 5-6 New Profile dialog box options

Property	Description
Choose backups tab	Enables you to specify the criteria for selecting backup images.
Duplication tab	Enables you to configure duplication of the selected backup images.
Catalog backup tab	Enables you to choose which catalog backup policy and schedule to use for creating a Vault catalog backup. For efficient disaster recovery, vault a new catalog backup each time you vault data.
Eject tab	Enables you to choose in which off-site volume pools Vault should look for the media you want to eject.
Reports tab	Enables you to choose which reports to generate.

A profile must select images (Choose Backups). The other steps are optional so you can separate the tasks into separate jobs if you want, using different jobs to accomplish different tasks. For example, you can use one profile to select and duplicate images daily, and another profile to eject media and generate reports weekly.

To configure a profile

- 1 If the **Profile** dialog box does not appear, select a profile in the NetBackup Administration Console window and select the **Change** icon in the toolbar.
- 2 Select the tab for each step that you are configuring and complete the fields.
- 3 Click **OK**.

See “About configuring a profile using the Choose Backups tab” on page 78.

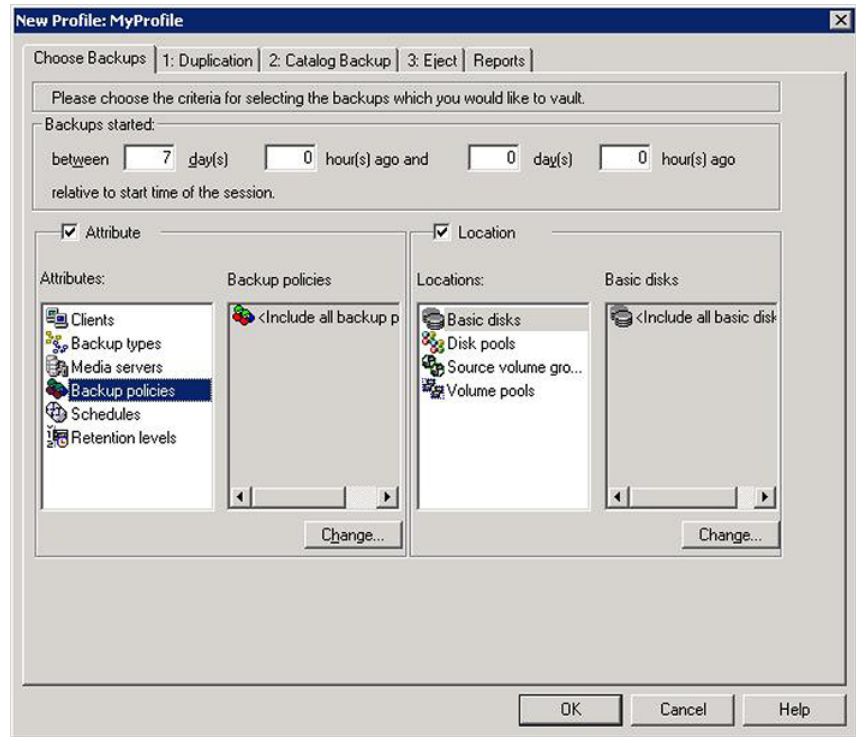
About configuring a profile using the Choose Backups tab

Use the **Choose Backups** tab to configure the search for images to be vaulted. The most basic criterion you can set is the time frame. To refine the search for images to vault, you can select **Attribute** and **Location** to configure the advanced options. The default setting of these two criteria is for the check boxes to be unselected, which means all criteria in that particular field is included in the search for an image to vault. The criteria in the **Attribute** field are logical criteria to help you refine your search. The criteria in the **Location** field represent physical locations of the images to backup.

Vault compares images in the NetBackup database with the criteria defined in the **Choose Backups** tab and generates a list of images that match the criteria. The image selection process chooses all images in the NetBackup catalog that match the criteria that you select under **Attribute** and **Location**, even images that are in a different vault. The criteria that you specify on the other tabs in the **Profile** dialog box determine whether Vault includes or excludes the selected images.

The **Choose Backups** tab enables you to quickly configure how you select criteria for your profile. For the broadest search coverage, you should leave the **Attribute** and **Location** check boxes empty and in their default state. This includes all criteria in your profile. Or, you can refine your search by using only criteria from the **Attribute** field, or search physical locations by using the criteria in the **Location** field. Finally, you can restrict your search to cover very specific areas by utilizing various criteria in the **Attribute** and the **Location** fields.

The image selection process may select catalog backup images. However, if you are duplicating images, Vault does not duplicate existing catalog images. Vault ejects the media on which those images reside if that media is assigned to a volume pool that is listed in the **Off-site Volume Pools** list on the **Eject** tab.



See “Choose Backups tab configuration options” on page 79.

See “About the list of images to be vaulted” on page 122.

See “About overlapping the time window in the profile” on page 30.

Choose Backups tab configuration options

Table 5-7 describes the options that you can use in the **Choose Backups** tab to configure a profile.

Table 5-7 Choose Backups tab configuration options

Property	Description
Attribute	A filter enabling you to choose which logical attributes you want to use when searching for an image to vault. By default, if the check box is not checked, all criteria in this field are included in the search.

Table 5-7 Choose Backups tab configuration options (*continued*)

Property	Description
Backup Policies	<p>A list of policies to use to select backup images. Enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked.</p> <p>To change the backup policies, click Change and then choose the backup policies you want to include in the profile. Policies are based on the storage unit that is used for backups. Because storage units are related to a specific robot number, choose the policies by robotic device.</p>
Backups Started	<p>The period of time from which the profile selects backups relative to the start time of the session. Time is expressed in terms of days and hours, relative to the time of the session. For example, assume the following settings:</p> <ul style="list-style-type: none"> ■ between 8 day(s) 0 hour(s) ago ■ and 1 day(s) 0 hour(s) ago <p>If the session is started on October 12 at 1:00 P.M., count backward from October 12. The vaulted backups are those started between October 4 at 1:00 P.M. (eight days before) and October 11 at 1:00 P.M. (one day before).</p> <p>If you select original backup images to send off site, the default time range is between eight days and one day before the session runs. If you duplicate images, the default time range is between seven days and zero days.</p>
Basic disks	<p>A list of basic disk paths, in the form of <server>:<path>. Selecting Basic disks refines your search to those paths that are selected. An image is selected for duplication or vaulting if its primary copy resides in any of the selected basic disk paths.</p> <p>Note: An individual basic disk from this list would also display the storage unit name in parenthesis, for example, <server>:<path> (storage unit).</p> <p>This option is enabled if you select this criteria in the Location field, or by default if the Location check box is not checked. To change the basic disks, click Change and then choose the basic disks that you want to include in this profile. By default, the Include all basic disks is selected.</p> <p>An image is selected for duplication or vaulting if its primary copy resides in any of the specified basic disks.</p>
Change	Option used to display a dialog box to change any of the criteria selected in the Attribute or Location fields.
Clients	<p>The clients for which to select backup images. Enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked.</p> <p>To change the clients, click Change and then choose the clients you want to include in this profile.</p>

Table 5-7 Choose Backups tab configuration options (*continued*)

Property	Description
Disk pools	<p>A list of disk pools that you can choose search for an image to duplicate or vault. Enabled if you select this criteria in the Location field, or by default if the Location check box is not checked. To change the disk pools, click Change and then choose disk pools that you want to include in this profile. By default, the Include all disk pools is selected.</p> <p>An image is selected for duplication or vaulting if its primary copy resides in any of the specified disk pools.</p> <p>Note: An individual disk pool from this list would also display the storage unit name in parenthesis, for example, disk pool (storage unit).</p>
Location	<p>A filter enabling you to choose which physical location you want to search for an image to vault. By default, if the check box is not checked, all criteria in this field are included in the search.</p>
Media servers	<p>The media servers from which to select backup images. This option is enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked. (Applies to NetBackup Enterprise Server only.) An image is selected for duplication or vaulting if the media server of the primary copy matches any of the values selected in the list.</p> <p>To change the media servers, click Change and then choose the media servers you want to include in this profile.</p>
Retention levels	<p>A list of retention levels that you can use to further refine your search criteria. An image is selected for duplication or vaulting if the retention level of the primary copy matches any of the values selected in the list.</p> <p>Enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked.</p> <p>To change the retention level, click Change and then choose the level or levels that you want to include in this profile. By default, the Include all retention levels is selected.</p>
Schedules	<p>A list of schedules to use to select backups. Enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked.</p> <p>To change the default, click Change and then choose the schedules you want to include in this profile. Schedules are based on the storage unit used for backups. Because storage units are related to a specific robot number, choose the schedules by robotic device.</p>

Table 5-7 Choose Backups tab configuration options (*continued*)

Property	Description
Source Volume Groups	<p>This selection criterion contains a list of Volume Groups from which to select backup images. Enabled if you select this criteria in the Location field, or by default if the Location check box is not checked. To change the default, click Change and then choose the volume groups you want to include in this profile.</p> <p>Selecting Source Volume Groups restricts the search for images to those in either all volume groups or the specific volume groups that you choose to include in your search. Usually, a Source Volume Group is specified if your master server has access to multiple robots and you want to duplicate the images that reside on media in one robot to media in another robot. The images that are read are in the Source Volume Group in one robot. The images are written to media in the Robotic Volume Group in another robot.</p> <p>An image is selected for duplication or vaulting if any fragment of its primary copy is found in a media that is from any of the selected volume groups.</p> <p>Note: If you want to exclude all tape images from this profile, select the Exclude All check box. This means there is no effect of the selection from Volume Pools criterion as both of them apply to tape images. So, if you select Exclude All for "source volume groups," it automatically means Exclude All for "volume pools" as well.</p>
Backup Types	<p>The types of backups (full, incremental, and so on) the profile captures. Enabled if you select this criteria in the Attribute field, or by default if the Attribute check box is not checked.</p> <p>To change the default, click Change and then choose the backup types you want to include in this profile. Depending on the different types of backups you configured in NetBackup policy management, you can choose the backup type. Only those types for which you configured policies are available for selection. If you want to vault all types of backups, accept the default. This criterion is optional.</p>
Volume pools	<p>A list of volume pools that you can choose to include in your search for an image to duplicate or vault. Enabled if you select this criteria in the Location field, or by default if the Location check box is not checked. To change the volume pools, click Change and then choose the volume pool or pools that you want to include in this profile. By default, the Include all Volume pools option is selected. To change the default, click Change and then choose the volume pools you want to include in this profile.</p> <p>Selecting Volume pools restricts the search for images to those in either all volume pools or just the volume pools that you choose to include in your search. An image is selected for duplication or vaulting if any fragment of its primary copy is found in a media that is from any of the selected volume pools.</p> <p>Note: To exclude all tape images from this profile, select Exclude All check box. This means there is no effect of the selection from "Source Volume Groups" criterion as both of them apply to tape images. So, if you select Exclude All for "volume pools" it automatically means Exclude All for "Source Volume Groups" as well.</p>

Duplication tab

Use the **Duplication** tab of the **Profile** dialog box to configure the rules that are used to duplicate images and to configure other duplication attributes. A duplication rule specifies the number of copies to create, a storage unit, off-site volume pool, retention period, media server (advanced configuration only), and what to do if an image copy fails (multiple copies only).

Duplication is optional. If you create multiple original backup copies concurrently during a backup job and vault one of the originals, you do not need to duplicate images in Vault.

Note: If duplication is enabled, Vault rejects the images that contain snapshots. You can use SLP to duplicate snapshot images and create multiple copies of snapshot images. Once SLP is complete for a snapshot image, use the Vault eject profile to eject the tape copy for off-site storage.

The screenshot shows the 'New Profile: MyProfile' dialog box with the 'Duplication' tab selected. The dialog has a title bar with a close button. Below the title bar are tabs: 'Choose Backups', '1: Duplication', '2: Catalog Backup', '3: Eject', and 'Reports'. The 'Duplication' tab is active. The main area contains several sections: 'Source' with options for 'Disk only' (radio button) and 'Removable media and/or disk' (radio button, selected), a 'Number of read drives' spinner set to 1, and an 'Alternate read server' checkbox (unchecked). 'Destination' section includes a 'Multiple copies' checkbox (unchecked), a 'Configure...' button, and dropdowns for 'Storage unit' (vltwin1-hcart-rob), 'Write drives' (1), 'Volume pool' (NetBackup), 'Retention level' (No change), and 'Media owner'. There is also a 'Make this copy primary' checkbox (unchecked). Below this is an 'Override default priority' section with a 'Duplication job priority' spinner set to -1. A 'Preserve multiplexing' checkbox is unchecked, with a note: '(Note: This option may slow the disaster recovery process.)'. At the bottom, there are three checked checkboxes: 'Duplicate smaller images first (applies only to disk backup images)', 'Expire original disk backup images after' (10 hours), and 'Expire original tape backup images after' (20 hours). At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

For more information on how to create one or more dedicated off-site volume pools, refer to the following:

- See “About off-site volume pools” on page 53.
- See “About the primary backup copy” on page 84.
- See “About basic duplication” on page 85.
- See “About advanced duplication” on page 85.
- See “Duplication tab configuration options” on page 87.
- See “Multiple Copies options” on page 91.
- See “Duplication Rule configurations” on page 93.
- See “About the treatment of images without a corresponding duplication rule” on page 96.
- See “About avoiding resource contention during duplication” on page 38.
- See “About how to avoid sending duplicates over the network” on page 44.
- See “About increasing duplication throughput” on page 47.
- See “About maximizing drive utilization during duplication” on page 49.

About the primary backup copy

NetBackup assigns an ordinal number to each copy of a backup image that is written by a backup policy. That number designates its sequence of creation. NetBackup also designates one of the backup images as the primary backup copy. The first backup image that is created successfully by a NetBackup policy is the primary backup; if only one copy of a backup image is created, it is the primary copy. NetBackup uses the primary copy to satisfy restore requests.

Usually, Vault duplicates from the primary copy, whether it exists on disk or removable media. Exception: for improved performance, Vault duplicates from a nonprimary copy on disk if one exists and the primary copy is on removable media.

Because both NetBackup and Vault use the primary copy, in most circumstances if a primary copy is on removable media it should remain in a robot. If the primary copy is off site, you cannot duplicate the image until the media is injected into the robot or a local copy (if available) is promoted to primary. (Exception: Vault duplicates from a nonprimary copy on disk if one exists.)

If you send the primary copy off site and you duplicate images in Vault, you can designate one of the copies that remains in the robot as the primary copy.

When the primary copy expires, NetBackup automatically promotes the backup copy that has the lowest number to primary.

About basic duplication

In basic duplication, you specify only one duplication rule. All backups are duplicated according to the same rule, and all selected images that are controlled by the specified master server are duplicated. You can create multiple copies of each backup image concurrently, but they are created using the same duplication rule.

About advanced duplication

Advanced duplication lets you specify more than one duplication rule. Vault determines which media server wrote each backup image and then applies the duplication rule corresponding to that media server to that image. In this context, the media server does not have any effect other than to identify which rule to apply to each image.

Note: Alternate read servers and multiple media servers apply to NetBackup Enterprise Server only.

If a duplication rule does not specify an alternate read server, the media server that originally wrote the backup image is used to read the original backup image during the duplication process.

Use advanced configuration only if you need to control exactly how to assign the backup images to be duplicated.

The following statements may help you understand why you should use advanced configuration:

Note: You do not need to configure advanced options if your profile duplicates images that are backed up by a single media server.

- Your robot has different types of drives or media so that you have different storage units to use as destinations for the duplication process. In this case, you may want to balance the duplication job between multiple storage units. For example, you may want to send the duplicate copies of all backup images that are written by one media server to a storage unit of one density and all backup images that are written by another media server to a storage unit of another density.
- Your profile is duplicating backup images to different media servers, each writing different types of data that require different retention periods. For example, if media server A backs up your customer database and media server B backs up warehouse inventory data, you may want to keep your customer database

in off-site storage for a longer period of time (a different retention) than your warehouse inventory data.

- You have one media server that you need reserved for other operations. For example, you use multiple media servers for duplication but dedicate one media server for backups. For that one media server you specify an alternate read server, and you let the rest of the media servers handle their own duplication.

Note: You do not need to configure advanced options if your profile duplicates images that are backed up by a single media server.

To avoid sending data over the network, do the following:

- For each duplication rule that does not specify an alternate read server, ensure that the media server controls both the source volumes and the destination storage units.
- For each duplication rule that specifies an alternate read server, ensure that
 - The alternate read server is connected to all robots that have backup images written by the media server specified for this rule.
 - The alternate read server is the same server as the media server of the destination storage unit.

New Profile: MyProfile

Choose Backups: 1: Duplication | 2: Catalog Backup | 3: Eject | Reports

☐ Skip the Duplication step ☒ Advanced configuration

☐ Alternate read server: Read original backups using media servers that are different from the media server that wrote the backups. (Note: This may send data over the network.)

SOURCE		DESTINATION		
Media Server	Read Drives	Storage Unit	Write Drives	Volume Pool
vltwin1	1	vltwin1-heart-ro...	1	NetBackup

New... Delete Change...

☐ Override default priority
Duplication job priority: -1
(higher number is greater priority)

☐ Preserve multiplexing (Note: This option may slow the disaster recovery process.)

☒ Duplicate smaller images first (applies only to disk backup images)

☒ Expire original disk backup images after 10 hours.

☒ Expire original tape backup images after 20 hours.

OK Cancel Help

Duplication tab configuration options

Table 5-8 describes configuration options for the **Duplication** tab.

Table 5-8 Duplication tab configuration options

Property	Description
Alternate Read Server	<p>This option applies to NetBackup Enterprise Server only.</p> <p>The name of an alternate read server. If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups. Using an alternate read server may transfer data over your network, affecting your site's computing environment. The Source Media Server and Alternate Read Server may be the same.</p> <p>By default this option is disabled. To configure an alternate read server, select Alternate Read Server. Then select a media server from the drop-down menu. For advanced duplication, click New to configure duplication rules.</p>

Table 5-8 Duplication tab configuration options (*continued*)

Property	Description
Change	<p>For advanced configuration only: the option used to display the Duplication Rule dialog box so you can change a destination media server and duplication rules for that server.</p> <p>If you select Alternate Read Server on the Duplication tab, the Duplication Rule dialog box has fields for both Source Media Server and Alternate Read Server. If you did not select Alternate Read Server, only a Source Backup Server field appears.</p>
Configure	For basic duplication only: the option used to display the Multiple copies dialog box.
Delete	For advanced configuration only: the option used to delete the selected destination media server and duplication rules for that server.
Duplicate Smaller Images First (applies only to disk backup images)	<p>Select to duplicate images in smallest to largest order. This capability applies only when duplicating disk backup images.</p> <p>By default, Vault duplicates images from largest to smallest, which improves tape drive utilization during duplication and duplicates more data sooner. If you know that your most important data is in smaller backup images, you can select this option so that those images are duplicated before the larger images.</p> <p>This choice does not affect the total time that is required to duplicate the images.</p> <p>Note: By default, Vault duplicates tape images using the time the backup was created. It duplicates images from the oldest to the newest.</p>
Duplication Job Priority	<p>The priority to assign to the Vault duplication jobs, from 0 to 99999. A larger number is higher priority. All duplication jobs for the profile run at the same priority.</p> <p>Vault duplication jobs compete with other process in NetBackup (such as regularly scheduled backups) for resources, including tape drives. If you want your Vault duplication jobs to obtain resources before other processes, assign a higher priority to the Vault jobs than to other NetBackup processes. Priority for backups, restores, and synthetic backups is assigned in the master server Global Properties.</p>
Expire Original Disk Backup Images	<p>The delay (in hours) until duplicated backup images become eligible to expire after the Vault session runs (applies only if the backup images are on disk and the duplication job has selected that copy as a source copy for duplication).</p> <p>You can use this option to set an earlier time for the images to become eligible to expire, however, the imageDB cleanup process performs the expiration of eligible images as a separate operation. The imageDB cleanup process is run every 12 hours by default. You can change this default value using the Image DB Cleanup Interval option on the Cleanup node of Master Server Host Properties on the NetBackup Administration Console or the <code>bpconfig -cleanup_int</code> command. Refer to the NetBackup Commands document for more information about this command.</p> <p>If the duplication of a disk image is not successful, the disk image does not expire. In addition, if the number of hours (X) equals zero, then the images expire immediately after a successful Vault duplication occurs.</p>

Table 5-8 Duplication tab configuration options (*continued*)

Property	Description
Expire Original Tape Backup Images	<p>The delay (in hours) until duplicated backup images become eligible to expire after the Vault session runs (applies only if the backup images are on disk and the duplication job has selected that copy as a source copy for duplication).</p> <p>You may choose to set this option to expire original VTL images. You can use this option to set an earlier time for the images to become eligible to expire, however, the imageDB cleanup process performs the expiration of eligible images as a separate operation. The imageDB cleanup process is run every 12 hours by default. You can change this default value using the Image DB Cleanup Interval option on the Cleanup node of Master Server Host Properties on the NetBackup Administration Console or the <code>bpconfig -cleanup_int</code> command. Refer to the NetBackup Commands document for more information about this command.</p> <p>If the duplication of a tape backup image is not successful, the tape image does not expire. In addition, if the number of hours (X) equals zero, then the images expire immediately after a successful Vault duplication occurs.</p>
Make This Copy Primary	<p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image that is created during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you have configured Continue as the fail option, the first successful copy is the primary copy.</p>
Media Owner	<p>The name of the owner of the media onto which you are duplicating images. Specify the media owner from the drop-down list box, as follows:</p> <ul style="list-style-type: none"> ■ The media owner option Any lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured). ■ The media owner option None specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ The media owner option Server Group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the drop-down list.
Multiple Copies	<p>Whether to create multiple copies concurrently. You can select Multiple Copies if the master server properties allow it.</p> <p>If you select Multiple Copies, click Configure to display the About the Multiple Copies options. See “Multiple Copies options” on page 91.</p> <p>If you configure multiple copies, you cannot configure a Storage Unit, Volume Pool, Retention Level, or Primary Copy on the basic Duplication tab.</p>

Table 5-8 Duplication tab configuration options (*continued*)

Property	Description
New	<p>For advanced configuration only, the option that is used to display the Duplication Rule dialog box in which you can add a destination media server and duplication rules for that server.</p> <p>If you select Alternate Read Server on the Duplication tab, the Duplication Rule dialog box has fields for both Source Media Server and Alternate Read Server. If you did not select Alternate Read Server, only a Source Backup Server field appears.</p>
Number of Read Drives	<p>The number of drives to use for reading backup images. When you enter a number of read drives, the same number is entered into the Destination Write Drives field. You must have an equivalent number of read and write drives available.</p>
Preserve Multiplexing	<p>Whether to preserve multiplexing. Multiplexing is the process of sending concurrent, multiple backup images from one or more clients to the same piece of media. This process speeds up duplication, but slows down restores and disaster recovery processes. If the option to preserve multiplexing is selected, the multiplexed duplication process occurs for all multiplexed images that are selected for duplication during a given Vault session.</p> <p>If the source image is multiplexed and the Preserve Multiplexing option is selected, ensure that the destination storage unit that is configured for each copy has multiplexing enabled. Multiplexing is configured in NetBackup Management > Storage Units.</p> <p>Multiplexing does not apply to disk storage units or disk staging storage units as destinations. However, if the source is a multiplexed tape and the destination is a disk storage unit or disk staging storage unit, selecting Preserve Multiplexing ensures that the tape is read on one pass rather than multiple passes.</p>
Retention Level	<p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date is the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1.</p> <p>See “Assigning multiple retentions with one profile” on page 141.</p> <p>When the retention period expires, information about the expired backup is deleted from the NetBackup and Media Manager catalog, the volume is recalled from off-site storage, and the backup image is unavailable for a restore.</p>
Skip the Duplication Step	<p>Select if you do not want to configure duplication.</p>
Source Backups Reside On	<p>The location of the backup images: disk or removable media or both. Vault duplicates images from the primary backup images on removable media or from backup images on disk.</p>

Table 5-8 Duplication tab configuration options (*continued*)

Property	Description
Storage Unit	<p>The name of a storage unit that contains the resources to which the copies of the backup images are written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source storage unit and destination storage unit can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, it is recommended that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p>
Volume Pool	<p>The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool are ejected for transfer off-site. Do not use the volume pool that was used for the original backup. NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool.</p>
Write Drives	<p>The number of write drives. This value is the same as the number of read drives.</p>

Multiple Copies options

The **Multiple Copies** dialog box appears only if you select the **Multiple Copies** checkbox on the basic **Duplication** tab and then click **Configure**. You use this dialog box to create multiple copies of a backup image concurrently.

Table 5-9 describes configuration options for the **Multiple Copies** dialog box.

Table 5-9 Multiple Copies dialog box configuration options

Property	Description
Copies	<p>The number of copies to create concurrently. You can create up to four or the number of copies that are specified in the Maximum Backup Copies field for the NetBackup master server (if less than four). (Configured in NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes.) By default, the value is two: one original backup and one copy.</p>

Table 5-9 Multiple Copies dialog box configuration options (*continued*)

Property	Description
For Each Image, If This Copy Fails	<p>The action to perform if a copy fails: Continue or Fail All Copies.</p> <p>In Vault, if you choose Fail All Copies, all copies of that image fail, independent of the success or failure of other image copy operations.</p> <p>The next time the Vault profile runs, Vault again tries to duplicate the image if the following conditions are true:</p> <ul style="list-style-type: none"> ■ The image is selected. ■ The Vault profile did not eject the primary backup. <p>By default, the option is configured to Fail All Copies in Vault.</p> <p>If you choose Continue for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted. It is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.</p> <p>See “About the continue or fail for concurrent copies” on page 162.</p>
Media Owner	<p>The name of the owner of the media onto which you are duplicating images.</p> <p>Specify the media owner from the drop-down list box, as follows:</p> <ul style="list-style-type: none"> ■ The media owner option Any lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured). ■ The media owner option None specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ The media owner option Server Group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the drop-down list.
Primary	<p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image that is created during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you configured continue as the fail option, the first successful copy is the primary copy.</p>

Table 5-9 Multiple Copies dialog box configuration options (*continued*)

Property	Description
Retention	<p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date is the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1.</p> <p>See “Assigning multiple retentions with one profile” on page 141.</p> <p>When the retention period expires, information about the expired backup is deleted from the NetBackup and Media Manager catalogs, the volume is recalled from off-site storage, and the backup image is unavailable for a restore.</p>
Storage Unit	<p>The name of a storage unit that contains the resources to which the copies of the backup images are written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source storage unit and destination storage unit can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, it is recommended that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p>
Volume Pool	<p>The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool is ejected for transfer off-site. Do not use the volume pool that was used for the original backup. NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool.</p>
Write Drives	<p>The number of write drives. This value is the same as the number of read drives.</p>

Duplication Rule configurations

The **Duplication Rule** dialog box appears if you select **New** or **Change** on the **Advanced Configuration** options of the **Duplication** tab. If you selected **Alternate Read Server** on the **Duplication** tab, an **Alternate Read Server** option appears on the dialog box.

Use the **Duplication Rule** dialog box to create multiple copies of an image and to select different media servers and read servers for the copies.

Table 5-10 describes configuration options for the **Duplication Rule** dialog box.

Table 5-10 Duplication Rule dialog box configuration options

Property	Description
Alternate Read Server	<p>The name of an alternate read server. (This option applies to NetBackup Enterprise Server only.)</p> <p>If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups. Using an alternate read server may transfer data over your network, affecting your site's computing environment. The Media Server and Alternate Read Server may be the same.</p> <p>To configure an alternate read server, select a media server from the drop-down menu.</p>
Backup Server	<p>Appears if Alternate Read Server was not selected on the Duplication tab. (Applies to NetBackup Enterprise Server only.)</p> <p>The name of the media server on which the backup images reside.</p>
Copies	<p>The number of copies to create concurrently. You can create up to four or the number of copies specified in the Maximum Backup Copies field for the NetBackup master server (if less than four). (Configured in NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes.) By default, the value is two: one original backup and one copy.</p>
For Each Image, If This Copy Fails	<p>The action to perform if a copy fails: Continue or Fail All Copies.</p> <p>In Vault, if you choose Fail All Copies, all copies of that image fail, independent of the success or failure of other image copy operations.</p> <p>The next time the Vault profile runs, Vault again tries to duplicate the image if the following conditions are true:</p> <ul style="list-style-type: none"> ■ The image is selected. ■ The Vault profile did not eject the primary backup. <p>By default, the option is configured to Fail All Copies in Vault.</p> <p>If you choose Continue for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted.</p> <p>It is probable that at least one copy will succeed, but it may not be the copy assigned to the off-site volume pool.</p> <p>See "About the continue or fail for concurrent copies" on page 162.</p>

Table 5-10 Duplication Rule dialog box configuration options (*continued*)

Property	Description
Media Owner	<p>The name of the owner of the media onto which you are duplicating images.</p> <p>Specify the media owner from the drop-down list box, as follows:</p> <ul style="list-style-type: none"> ■ The media owner option Any lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured). ■ The media owner option None specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media. ■ The media owner option Server Group allows only those servers in the group to write to the media on which backup images for this policy are written. All server groups that are configured in the NetBackup environment appear in the drop-down list.
Media Server	<p>Appears only if Alternate Read Server was selected on the Duplication tab. (Applies to NetBackup Enterprise Server only.)</p> <p>The name of the media server on which the backup images reside. The Media Server and Alternate Read Server may be the same.</p>
Number of Read Drives	<p>The number of drives to use for reading backup images. When you enter a number of read drives, the same number is entered into the Destination Write Drives field. You must have an equivalent number of read and write drives available.</p> <p>Note: Vault does not let you create multiple duplication rules per media server.</p>
Primary	<p>Whether the copy should be designated the primary backup. Only designate a duplicate as the primary if the primary backup is ejected and transferred off site.</p> <p>NetBackup restores from the primary backup, and Vault duplicates from the primary backup. By default, the original backup image that is created during a NetBackup policy job is the primary copy. If the copy that you indicate as primary fails, and you configured continue as the fail option, the first successful copy is the primary copy.</p>
Retention	<p>The retention level for the copy. Each copy has a separate expiration date. If a retention level is not specified, the expiration date is the same as the original. If you specify a numeric retention level, the expiration date for the duplicate media is calculated by adding the specified retention period to the date the original backup was created. If you specify Use Mappings for the retention level, the retention period is based on the retention period of backup image copy 1.</p> <p>See “Assigning multiple retentions with one profile” on page 141.</p> <p>When the retention period expires, information about the expired backup is deleted from the NetBackup and Media Manager catalog, the volume is recalled from off-site storage, and the backup image is unavailable for a restore.</p>
Source Backups Reside On	<p>The location of the backup images: disk or removable media or both. Vault duplicates images from the primary backup images on removable media or from backup images on disk.</p>

Table 5-10 Duplication Rule dialog box configuration options (continued)

Property	Description
Storage Unit	<p>The name of a storage unit that contains the resources to which the copies of the backup images are written.</p> <p>Storage units can be Media Manager storage units, disk storage units, disk staging storage units, or Network Data Management Protocol (NDMP) storage units.</p> <p>If the Media Manager or NDMP storage unit has more than one drive, the source storage unit and destination storage unit can be the same. NDMP storage units are supported only when one copy is created per duplication rule. Because of potential NDMP performance limitations, it is recommended that you duplicate between drives that are directly attached to the same NDMP host.</p> <p>If the duplicated backup images are to be vaulted, the media in the destination storage unit must be in the Robotic Volume Group.</p> <p>All storage units must be connected to the same media server.</p>
Volume Pool	<p>The name of the off-site volume pool to which Vault assigns the duplicate media. Images on media in the off-site volume pool are ejected for transfer off-site. Do not use the volume pool that was used for the original backup. NetBackup does not verify in advance that the media ID selected for the duplicate copy is different than the media that contains the original backup. To ensure that two processes do not try to use the same volume at the same time, specify a different volume pool.</p>
Write Drives	<p>The number of write drives. This value is the same as the number of read drives.</p>

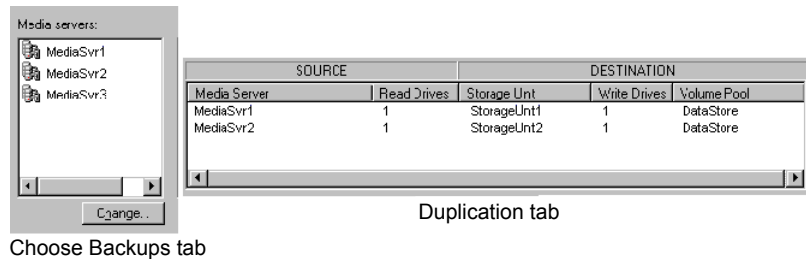
About the treatment of images without a corresponding duplication rule

In some cases, the profile may list more media servers in the **Media Servers** list on the **Choose Backups** tab than in the advanced configuration view on the **Duplication** tab.

Figure 5-1 shows that the number of media servers that appear in the **Choose Backup** tab can differ from the number of media servers that appear in the **Duplication** tab.

Note: More than one media server applies to NetBackup Enterprise Server only.

Figure 5-1 Media servers listed on Choose Backup and Duplication tabs



If this happens, images that are written by media servers that have no corresponding duplication rule must also be duplicated. Vault duplicates those images but tries to minimize total duplication time by keeping as many drives as possible busy writing data until all images are duplicated.

This situation is handled as follows:

- All images that are written by media servers that have a duplication rule are assigned to the appropriate duplication rule.
- As soon as one duplication rule finishes processing the images that are assigned to it, Vault begins to assign images written by other media servers (media servers that have no rule of their own) to the duplication rule that finished processing.
- As other rules complete the duplication of their assigned images, they too are assigned images written by other media servers that have no rule of their own.
- Eventually all images that are written by all media servers that are listed on the **Choose Backups** tab are duplicated and the duplication step is complete. If you have more media servers listed on the **Choose Backups** tab than on the **Duplication** tab, there is only one way to ensure that large amounts of duplication data do not get sent over the network.
 - Every duplication rule must specify an alternate read server. For each duplication rule, the alternate read server must be the same as the media server of the destination storage unit(s).
 - All alternate read servers must be connected to all robots that have images written by any media server listed on the **Choose Backups** tab but not on the **Duplication** tab.

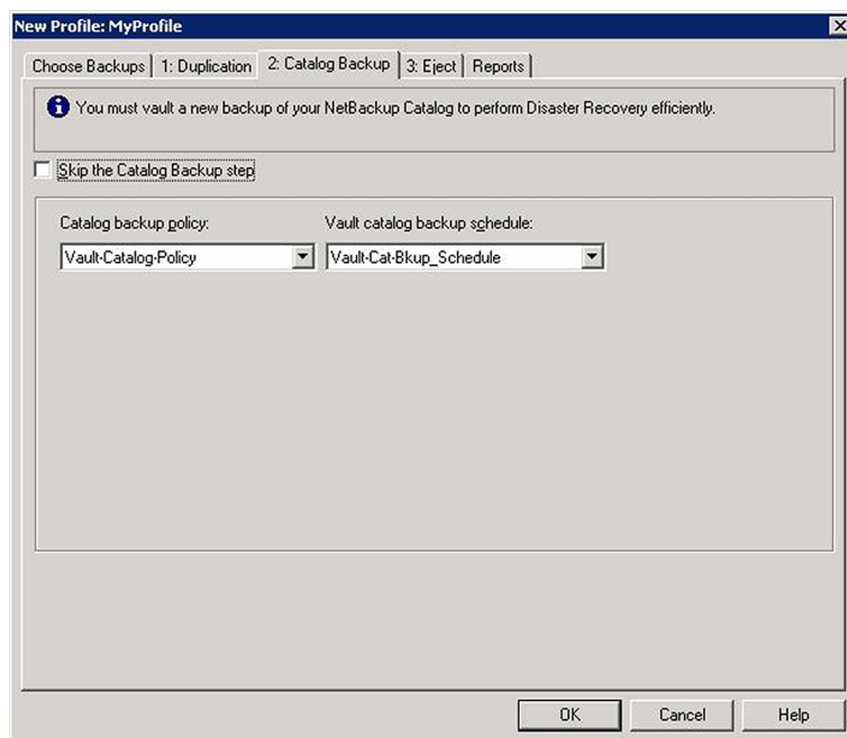
The previous configurations are best suited for a SAN environment where all media servers are visible to all robots.

Catalog backup tab (Profile dialog box)

Use the **Profile** dialog box **Catalog Backup** tab to specify the catalog backup policy and schedule that performs the Vault catalog backup. At least one NBU-Catalog policy that includes a Vault Catalog Backup schedule must exist so you can specify them on the **Catalog Backup** tab.

Vault uses the online, hot catalog backup method, which can back up the catalogs even when other NetBackup or Vault activity is occurring. (NetBackup provides two types of catalog backup; the other type is an offline, cold catalog backup that cannot occur when regular backup activity is occurring.)

You must add the Vault catalog backup volume pool to the eject list on the **Eject** tab.



See “About off-site volume pools” on page 53.

See “About creating catalog backup schedules for Vault” on page 55.

The following topics provide additional information about Vault catalog backups:

- See “About Vault catalog backups” on page 99.

- See “Catalog backup policy settings” on page 99.
- See “About critical policies ” on page 100.
- See “Catalog backup tab configuration options” on page 101.

About Vault catalog backups

NetBackup catalogs are databases that contain information about the NetBackup configuration and backups, including the files that are backed up and the media on which they are stored. Vault creates a new catalog backup with up-to-date information. It does not duplicate an existing NetBackup catalog backup. A NetBackup catalog backup is not a substitute for a Vault catalog backup because it does not include the latest information about duplicated media and media location. Therefore, you should perform a catalog backup in Vault.

To perform a Vault catalog backup, Vault uses a special schedule of type Vault Catalog Backup in an NBU-Catalog policy. (NetBackup uses a special backup policy of type NBU-Catalog to perform catalog backups.)

A Vault catalog backup occurs when a profile that performs catalog backup runs. It does not occur on the schedule that is defined in the NBU-Catalog policy. A Vault Catalog Backup schedule always performs a full backup of the entire NetBackup catalog.

You can create multiple copies concurrently of a Vault catalog backup.

Vault does not duplicate existing catalog images, but it ejects the media on which those images are stored if both of the following are true:

- Those images are selected during the choose backups step.
- The media is assigned to the dedicated Vault catalog volume pool.

If the catalog backup fails but the remainder of the Vault job succeeds, the session ends with a status 294 (vault catalog backup failed). Data is vaulted with no associated catalog backup. Cohesity believes that it is better to vault the data without a catalog backup than to fail the job and vault nothing at all for that session.

The Recovery Report for Vault shows the three most recent Vault catalog backups. If you vault your regular NetBackup catalog backups, they do not appear on the Recovery Report but do appear on other reports.

See the *NetBackup Administrator's Guide, Volume I* for information about the location of the NetBackup catalog and the files that are included in a catalog backup.

Catalog backup policy settings

Although you configure only two options on the **Catalog Backup** tab, settings in the catalog backup policy affect the Vault catalog as follows:

Table 5-11 Catalog backup tab configuration options

Property	Description
Job priority	The catalog backup job competes for resources with other backup jobs. You can specify the priority for the job, either on the policy Attributes tab (single catalog backup) or on the Configure Multiple Copies dialog box of the Vault Catalog Backup schedule.
Destination	You can send the catalog backup to any storage unit, including disk and removable media storage units. Specify the destination in the Vault Catalog Backup schedule.
Volume pool	If you use removable media storage units, you must specify a volume pool for the catalog backup in the Vault Catalog Backup schedule. You also must specify the same volume pool on the Eject tab of the profile.
Number of copies	You can create multiple copies of the catalog, and you can send them to any storage unit that is attached to the destination media server. Specify multiple copies on the Vault Catalog Backup schedule Attributes tab and then specify the number of copies on the Configure Multiple Copies dialog box.
Critical policies	Beginning with NetBackup 6.0, you do not specify the catalog files to include in the Vault catalog backup, and you cannot add other files to the catalog backup. A new NetBackup catalog policy option, critical policies, lets you select policies that should be recovered before backups from other policies are recovered. Specify critical policies on the policy Disaster Recovery tab.

About critical policies

A new NetBackup catalog policy option, critical policies, lets you select policies that should be recovered before backups from other policies are recovered. The media that contains the backup images from those critical policies are listed in the new NetBackup disaster recovery report. (Vault does not include critical policy data in its catalog backups.)

The media from those critical policies also is listed in the Vault Recovery Report so you can recall the media from those policies and recover that data before you recover data that is not as critical.

You can even recover the critical policies data before you recover the entire catalog. During catalog recovery, partial recovery of the catalog is possible. Partial recovery means you can choose which parts of the catalog backup to restore and in what order. Partial recovery lets you restore part of the catalog, then restore data from critical policies, and finally restore the remainder of the catalog.

If you want to include certain critical policies in your Vault catalog backup but not your regular NetBackup catalog backup, you need a separate NBU-Catalog policy for Vault.

For information about specifying critical policies in an NBU-Catalog policy, see the *NetBackup Administrator's Guide, Volume I*.

Catalog backup tab configuration options

The following are the configuration options for the **Catalog Backup** tab:

Table 5-12 Catalog backup tab configuration options

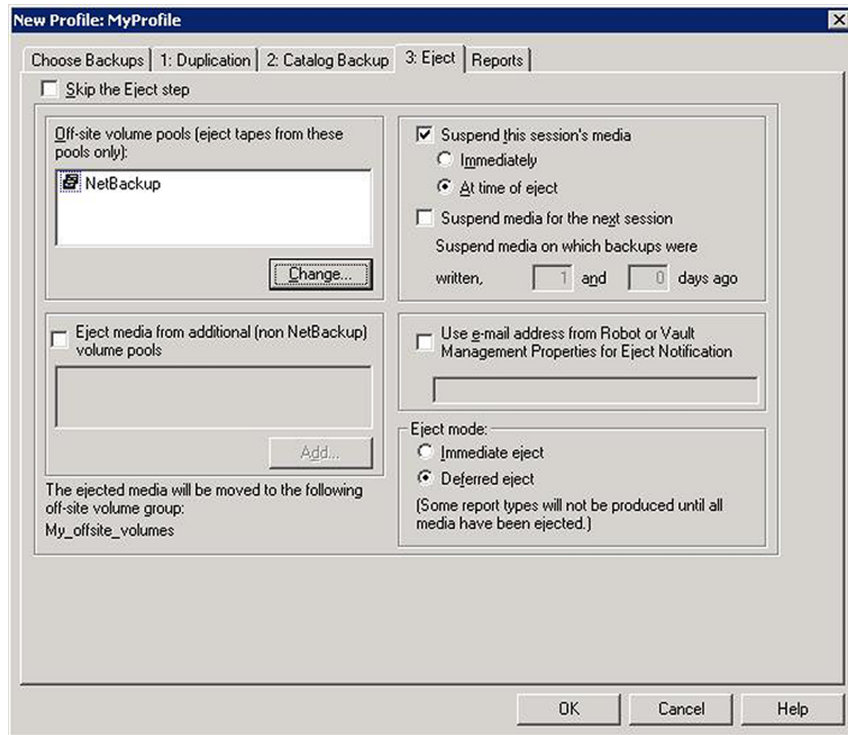
Property	Description
Catalog Backup Policy	The catalog backup policy to use. The drop-down list includes all catalog backup policies that have a Vault Catalog Backup schedule.
Vault Catalog Backup Schedule	The schedules available for the catalog backup policy. The drop-down list includes all Vault Catalog Backup schedules for the policy selected in Catalog Backup Policy field.
Skip the Catalog Backup Step	Select if you do not want to back up and vault the NetBackup and Media Manager catalogs.

Eject tab (Profile dialog box)

Use the **Profile** dialog box **Eject** tab to specify the following:

- Volume pools from which to eject media.
- Suspend option. By default, Vault suspends media when it is ejected. Vault only suspends the volumes in volume pools specified in the off-site volume pools list.
- When to eject the media (immediately when the profile runs or later).
- Email addresses for eject notification.

If you create catalog backup media in a profile in which you eject media, you must add the appropriate catalog volume pool to the off-site volume pools eject list.



About media ejection

During media ejection operations, Vault moves the media to be ejected into the default media access port (MAP) of the robotic library. You must extend the MAP, remove the media, and then retract the MAP. If more media is ejected, Vault continues to fill the MAP until all media are ejected. For libraries that have separate MAP doors such as libraries connected by pass-through mechanisms, NetBackup treats all doors as one continuous MAP. In other words, each time you are prompted by NetBackup, open all the doors, empty all the MAPs, and then close all the doors.

For ACS robots that have multiple MAPs, you can specify the MAPs to eject media to when you configure the robot in Vault.

If you use a library that does not have a MAP, you must remove the media from the library slots manually. You also have to perform the eject operation in Vault so that the appropriate database entries are completed. Although you can use automatic eject, Cohesity recommends that you use deferred eject to avoid resource contention with other NetBackup activity and you do not neglect to remove the media from the robot. The manual eject operation serves as a reminder to remove the media.

To use deferred eject for a library that does not have a MAP, do the following:

- Configure the profiles for deferred eject.
- Eject the media manually.
See "About ejecting media" on page 124.
- Remove the media from the library slots.

Do not inventory the robot until you remove the media from the MAP or library slots. If you do, you have to revault the media.

About ACS MAP

The following applies to NetBackup Enterprise Server only.

Automated cartridge system (ACS) robots can have multiple library storage modules (LSMs), each with multiple media access ports (MAPs). When you configure a vault that uses an ACS robot, you can specify any MAP or a subset of MAPs to use for media ejection. Vault ejects media to as few of the configured MAPs as possible based on a nearest MAP algorithm. The algorithm considers the volumes to be ejected, the MAPs configured for ejecting in the vault, and the configuration of the LSMs. The algorithm assumes that the LSMs are connected in a line. If your LSMs are connected in a configuration other than a line, see the following two sections within the NetBackup Administrator's Guide:

See "Adjacent LSM Specification for ACS Robots" in the *NetBackup Administrator's Guide*.

See "Media Access Port Default for ACS Robots" in the *NetBackup Administrator's Guide*.

The **Any MAP** option does not mean all MAPs; media are not ejected to all MAPs, media are ejected to the nearest MAP in each LSM.

If you specify any MAP:

- MAPs that have only one element are not used.
- Vault selects from MAPs that are on-line when the eject begins; MAPs that are offline are not considered for eject operations.
- If only a subset of MAPs are used during ejection, all MAPs are busy and unavailable. For example, if the media are ejected to only two MAPs in one LSM, all MAPs are still busy.

For all other robot types that have MAPs, media are ejected to the default MAP. NetBackup does not support ejecting to multiple MAPs for other robot types.

About eject mode (immediate or deferred)

You can eject media immediately when the profile runs or defer ejection. If you use one profile to choose and duplicate images daily and another profile to eject the media, specify deferred eject for the profile that selects and duplicates images and immediate eject for the profile that ejects media. If you defer eject, you also should defer reports.

If you select deferred eject, other actions are required to eject the media for the session. After the session ends, you can eject media for one session, for multiple sessions, or for all sessions.

Eject for one session only, as follows:

- Use the NetBackup Administration Console to eject media for the session.
- Use the Vault Operator Menu to eject media for the session.
- Use the `vlteject` command to eject media for the session.
- Create a Vault policy and enter the appropriate `vlteject` command and options in the file list.

Eject for multiple sessions for a specific profile, as follows. (This method for duplicating and ejecting media provides the added benefit of consolidated reports that are not organized by session.)

- Configure a Vault profile to duplicate only, and configure a Vault policy to run this profile on the days you want to select and duplicate images.
- Configure a second Vault profile to do the catalog backup and eject steps. This profile should use the same image selection criteria as the profile that duplicates images. Configure a Vault policy to run this profile on the day you want the media ejected.

Eject for all sessions for a specific vault or for all sessions for all vaults (consolidated eject) by doing one of the following:

- Use the NetBackup Administration Console.
- Use the Vault Operator Menu.
- Use the `vlteject` command.
- Create a Vault policy and enter the appropriate `vlteject` command and options in the file list.

If you defer eject operations, you also should defer reports until you eject media.

About media ejection timeout impact

The media ejection timeout period is the amount of time the eject process waits for the removal of the ejected media from the media access port (MAP) before an error condition occurs. The timeout period varies depending on the capability of each robot.

Table 5-13 shows the ejection timeout periods for robots.

Table 5-13 Media ejection timeout period for Vault

Robot	Timeout period	Note
Automated Cartridge System (ACS)	One week	Applies to NetBackup Enterprise Server only
Tape Library 8mm (TL8)	One week	
Tape Library DLT (TLD)	One week	
Tape Library Half-inch (TLH)	None	Applies to NetBackup Enterprise Server only
Tape Library Multimedia (TLM)	One week	Applies to NetBackup Enterprise Server only
Robots that do not have MAPs	None	

For robots that do not have timeout periods or do not have MAPs, select deferred eject and then eject the media manually. When you eject the media, ensure that the media are removed from the library in a timely manner so that other operations can occur.

Status messages that are displayed by the NetBackup Administration Console or on the command line provide information about the status of inject, eject, or inventory operations.

Note: If media are not removed and the timeout period expires, the Vault reports do not accurately show the status of the media. To recover, you should use the Vault Operator Menu (`vltopmenu`) or `vlt eject` to eject the media that was not removed from the library and generate the reports.

Eject tab configuration options

Table 5-14 lists the configuration options for the **Eject** tab.

Table 5-14 Eject tab configuration options

Property	Description
Add	The option that is used to add a volume pool to the eject list. If you click Add , the Volume Pools dialog box appears, in which you can add or remove volume pools from the eject volume pool list.
At Time of Eject (Suspend this Session's Media)	Suspend the media when it is ejected. If you also select deferred eject, images can be written to the media until it is ejected. Select this option if you want the media sent off-site to be full. Suspend at time of eject is the default.
Deferred (Eject Mode)	Defer media ejection until a later time. The reports that are marked with an asterisk (*) on the Reports tab are generated only when all media selected by the profile have been ejected.
Eject Media from additional (non-NetBackup) Volume Pools	Eject non-NetBackup media that is managed by NetBackup Media Manager. See "Vaulting non-NetBackup media managed by Media Manager" on page 150.
Immediate (Eject Mode)	Eject media immediately as part of the current Vault job. The reports that are marked with an asterisk (*) on the Reports tab are generated only when all media selected by the profile have been ejected.
Immediately (Suspend this Session's Media)	Suspend the media during the current session. No more images are written to the media even if ejection is deferred.
Use Email Address from Robot or Vault Management Properties for Eject Notification	Select to send the eject notification email to the email addresses configured on the Vault Robot dialog box or Vault Management Properties dialog box. Eject notification is sent when the eject process begins and ends. To send the notification email to different address, enter the email addresses in the field below the Use Email Address... checkbox. Eject notification is configured for each profile on the Eject tab, for each robot on the Vault Robot dialog box, and globally for Vault on the Vault Management Properties dialog box General tab. Vault sends the notification to the first email addresses found in that order. You can configure different addresses in each place.
Off-site Volume Pools	The names of the volume pools from which to eject media. Only the media in the pools that contain images that meet the selection criteria are ejected. If you create catalog backup media in a profile in which you eject media, you must add the appropriate catalog volume pool to the off-site volume pools eject list. If you use a <code>vlt_ejectlist_notify</code> script to eject media that is not created by NetBackup or Vault, you must add the volume pool in which that media resides to the Off-site Volume Pools list of the profile that you run to eject that media. See "About using notify scripts" on page 151.

Table 5-14 Eject tab configuration options (*continued*)

Property	Description
Skip the Eject Step	Select if you do not want to eject media with this profile.
Suspend Media for the Next Session (Suspend Media on Which Backups Were Written)	<p>Select to suspend original backup media. Then enter the number of days before the Vault job to suspend media.</p> <p>Use this option only if you vault original images and want to prevent NetBackup from writing partial images on backup media.</p> <p>Carefully consider whether to use this option. It uses extra CPU cycles because it queries all of the databases again and applies all of the Choose Backups filters again to select media to suspend. Also, this option does not suspend media that is in use, such as media to which NetBackup is writing backup images.</p> <p>This option suspends duplicate media that is created by Vault. However, the Suspend this Session's Media option is a better choice for suspending duplicate media because it does not use CPU cycles to select media to suspend.</p> <p>See "About avoiding vaulting partial images" on page 32.</p>
Suspend this Session's Media	<p>Select to suspend media in the eject list, then select either Immediately or At Time of Eject.</p> <p>Suspend at time of eject is the default.</p>

Reports tab (Profile dialog box)

Use the **Reports** tab to select which reports to generate for the profile, where to distribute them, and when to generate them (immediately when the profile runs or deferred until later).

You and your off-site storage vendor can use the reports to determine which media should be moved between your site and the off-site storage location and the timing of the moves.

Reports can be generated for one session or for multiple sessions (known as consolidating your reports and ejections).

Note: You must specify a destination so that reports are generated.

By default, the report files are stored in the following path:

- **UNIX**

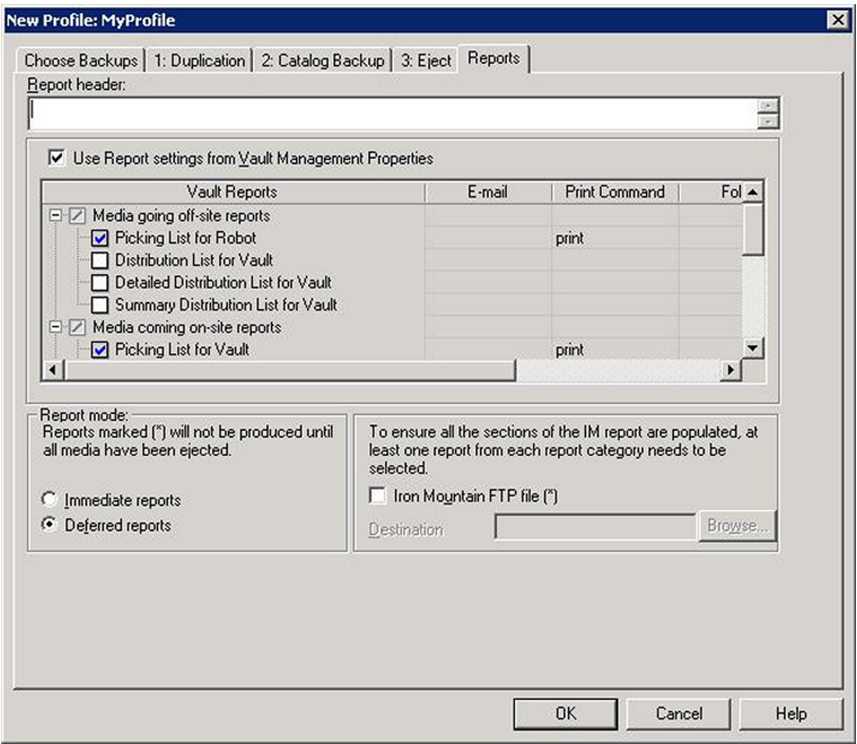
```
/usr/opensv/netbackup/logs/user_ops/vault/vault-xxx
```

- **Windows**

```
install_path\NetBackup\logs\user_ops\vault\vault-xxx
```

If you specify a different directory other than the default directory for saving the report, the report files are saved as per the following applicable naming rules:

- picklist_robot_sidXX_XXXXXXXXX.rpt
- picklist_vault_sidXX_XXXXXXXXX.rpt



See “Changing report properties” on page 110.

Generating reports is optional.

See “About report mode (immediate or deferred)” on page 109.

See “Reports that depend on eject” on page 109.

See “Reports tab configuration options” on page 109.

See “Changing report properties” on page 110.

About report mode (immediate or deferred)

Similar to the eject mode, you can specify whether reports should be generated immediately when the profile runs or deferred until later. If you defer eject, you should also defer reports. If you defer reports, you must perform or schedule another action to generate the reports.

Because some reports are generated only when media are ejected, you may choose to defer reports until the media are ejected. For example, if you duplicate images daily and eject media weekly, you can defer reports for the profile that duplicates images and use the profile that ejects media to generate reports.

If the corresponding eject process has completed when you generate reports, all pending reports are generated and distributed. Those reports are not regenerated if you run deferred reports again. If eject has not completed, the subset of reports that do not depend on completion of eject are generated. Those reports are generated again if deferred reports are ran again.

If you defer reports from multiple sessions and then generate them together, the progress is known as consolidating reports.

Reports that depend on eject

Reports can be generated from on the **Reports** tab of the **Profile** dialog box. Reports marked with an asterisk (*) on the **Reports** tab are generated only when all media that is selected by the profile are ejected. The reports for media coming on-site are not dependent on the media being ejected before they are generated.

Reports that you can generate from the **Profile** dialog box **Reports** tab are as follows:

- Reports for media coming on-site
 - Picking List for Vault
 - Distribution List for Robot
- Inventory reports
 - Vault Inventory
 - Off-site Inventory
 - All Media Inventory)
- Recovery Report for Vault

Reports tab configuration options

The following are the **Reports** tab configuration options.

Table 5-15 Reports tab configuration options

Property	Description
Deferred Reports	<p>Defer generating the reports until after the session has completed (for example, if you run Vault sessions daily and eject media weekly). Deferred is the default.</p> <p>Reports marked with an asterisk (*) are generated only when all media that is selected by the profile are ejected.</p>
Immediate Reports	<p>Generate the reports immediately as part of the current vault session. Reports that are marked with an asterisk (*) are generated only when all media that is selected by the profile are ejected.</p>
Iron Mountain FTP File	<p>If you selected Iron Mountain as your Vault vendor (in the New Vault dialog box), Iron Mountain FTP file and Destination folder items appear.</p> <p>To generate a file that you can send by FTP to Iron Mountain, select Iron Mountain FTP file and enter the name or browse to choose the Destination folder to which the file is written.</p> <p>Sending the file to Iron Mountain is not part of the vault process.</p>
Report Header	<p>If you want certain text to appear at the top of every report, enter it in the Report Header box. The header appears on all reports.</p>
Use Report Settings from Vault Management Properties	<p>Select to use the report settings configured in the Vault Management Properties Reports tab.</p>

Changing report properties

Use the following procedure to change report properties (title, email destination, printer to use, and directory to save it to).

To change report properties

- 1 Double-click a report.
- 2 In the **Change Report Properties** dialog box, select options and enter information as necessary.

If you change a title, the new title appears on the **Reports** tab and in the **Report Type** list box when you view Vault reports in the Administration Console.

If you consolidate your reports and also change titles, use the same title for all profiles whose reports are consolidated. The title prints on the reports and appears in the email subject line if you email the reports.

Vaulting and managing media

This chapter includes the following topics:

- About Vault sessions
- About previewing a Vault session
- Stopping a Vault session
- About resuming a Vault session
- About monitoring a Vault session
- About the list of images to be vaulted
- About ejecting media
- About injecting media
- About using containers
- Assigning multiple retentions with one profile
- About vaulting additional volumes
- Revaulting unexpired media
- About tracking volumes not ejected by Vault
- Vaulting non-NetBackup media managed by Media Manager
- About notifying a tape operator when an eject begins
- About using notify scripts

- About clearing the media description field
- Restoring data from vaulted media
- Replacing damaged media

About Vault sessions

A Vault session or vaulting job is the process of running the steps that are specified in a Vault profile. Before you can run a Vault session, you must configure at least one robot, one vault, and one profile.

For additional information about Vault sessions, refer to the following:

- See “About scheduling a Vault session” on page 112.
- See “About running a session manually” on page 115.
- See “About running multiple sessions simultaneously” on page 116.

You also can run a vault session by using the Vault Administration menu interface (UNIX systems only).

See “About the Vault administration interface” on page 207.

About scheduling a Vault session

To run a vault session automatically at a specific day and time, use a Vault policy. A Vault policy is a NetBackup policy that is configured to run Vault jobs; a Vault policy does not back up client systems. The policy includes the schedule for when the Vault session should run (day or date and time window) and the command to run a Vault profile.

How you schedule your sessions depends on how you conduct operations as follows:

- A Vault policy can run a profile that ejects media containing the original images that were created during a backup job. If you create multiple original backup images concurrently, you can assign one or more of the original images to an off-site volume pool. A separate Vault policy can run a Vault job that ejects the media on which those images are stored.
- A Vault policy can run a profile that selects images, duplicates those images, and ejects the media on which those images are stored. That policy can perform both operations daily or at some other interval that meets your requirements. If your vault vendor arrives daily to pick up media or you remove the off-site media from your robot immediately, you may need only one policy for that vault.

- One Vault policy can run a profile that duplicates images, and another policy can run a profile that ejects media. For example, if you create backup media daily and transfer it off site weekly, you can use one policy to create the backups daily and another policy to eject media weekly. If your vault vendor transfers your media weekly, you may prefer to eject media on the day the vault vendor arrives.

See “About media ejection recommendations” on page 37.

See “Creating a Vault policy” on page 113.

See “Vault policy configuration information” on page 114.

Creating a Vault policy

Setting up a Vault policy differs from setting up a regular policy in NetBackup, as follows:

- First, you must specify Vault as the policy type.
- Second, you do not specify clients for Vault policies.
- Third, rather than specifying files to back up on the **Backup Selections** tab, you specify one of two Vault commands to run.
 - Use the `vltrun` command to run a Vault session. You specify the robot, vault name, and profile for the job. The `vltrun` command accomplishes all the steps necessary to vault media. If the profile is configured for immediate eject, media are ejected and reports are generated. If the vault profile name is unique, use the following format:

```
vltrun profile_name
```

If the vault profile name is not unique, use the following format:

```
vltrun robot_number/vault_name/profile_name
```

- Use the `vlteject` command to eject media or generate reports for sessions that are completed already and for which media have not been ejected. The `vlteject` command can process the pending ejects or reports for all sessions, for a specific robot, for a specific vault, or for a specific profile. For example:

```
vlteject -vault vault_name -eject -report
```

Note: Include one Vault command only in a Vault policy. If you use more than one command, the first command is initiated and the successive commands are interpreted as parameters to the first command. Failure may occur and images may not be duplicated or vaulted.

See “Ejecting media by using a Vault policy” on page 128.

For more information about the `vlteject` and the `vltrun` commands, see the *NetBackup Commands Reference Guide*.

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

Note: If you create a vault policy by copying a regular NetBackup policy that has a client list configured, delete all the clients in the client list before you run the policy. If you do not, Vault creates one vault job for every client in the list even though the client is not used by the Vault job. The first vault job operates as a normal vault session; the rest terminate with a status 275 (a session is already running for this vault).

To create a Vault policy

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.
- 4 On the **Attributes** tab, select **Vault** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule.
The type of backup defaults to Automatic Vault.
- 6 Complete the schedule.
- 7 On the **Backup Selections** tab, enter the appropriate Vault command for the policy.
- 8 Click **OK**.

Vault policy configuration information

Collect and record the following information for each new schedule or policy pair you create or existing pair you consider for off-site rotation. The information you record is used when you configure Vault. It can help you determine if an existing policy can be used to create backup media that Vault can select for ejection.

Collect and record the following policy configuration information:

Policy Names	The names of all policies that are used for off-site rotation. To get information about existing policies, you can use the <code>bppllist</code> command.
Schedule Names	The names of the schedule or schedules that are associated with each policy.
Off-site Original, Duplicate, or Both	Record whether the policy selects original backup media, creates duplicate backup media, or both.
Storage Unit	The storage units for each policy.
Retention Period	The retention period for each schedule so that you have an idea of when to expect the media to return from off-site.

About running a session manually

You can run a Vault session manually either by using the NetBackup Administration Console to initiate the session or by invoking the veteran command from a command line.

See “Running a Vault session from the NetBackup Administration Console” on page 115.

See “Running a session from a command line” on page 116.

Running a Vault session from the NetBackup Administration Console

To run a Vault session from the NetBackup Administration Console, you either run the policy manually or run the profile manually.

To run a Vault policy from the NetBackup Administration Console

- 1 Expand **NetBackup Management > Policies** in the left pane of the Administration Console window.
- 2 Select the policy you want to run.
- 3 Select **Actions > Manual Backup**.

To run a Vault profile from the NetBackup Administration Console

- 1 Expand **Vault Management** in the left pane of the Administration Console window.
- 2 In the left pane, expand the robot that contains the vault and profile you want to run.

- 3 In the left pane, select the vault that you want to run.
- 4 In the Details (right) pane, click the profile that you want to run.
- 5 Select **Actions > Start Session**.

Start Session remains selected until the session begins. When the session starts, the Console displays a message similar to the following:

```
Manual vault session for profile has been started.  
Use the Activity Monitor to view progress.
```

By default, the Details pane of the Administration Console window does not show the **Volume Pools (Ejected)** and **Report Destination** columns. You can add, delete, or rearrange the columns that appear in the Details pane by selecting **View > Columns > Layout**.

Running a session from a command line

Use the following procedure to run a vault session from a command line.

To run a vault session from a command line

- 1 Add the path in which the NetBackup executable files are installed to your `PATH` environment variable
- 2 Run `vltrun` from a command line, specifying the robot number, vault, and profile as in the following example:

```
vltrun robot_number/vault_name/profile_name
```

Alternatively, you can specify only the profile if it has a unique name.

See the *NetBackup Commands Reference Guide* or the online Help in the NetBackup Administration Console.

About running multiple sessions simultaneously

Multiple Vault sessions can run at the same time. Vault uses a global setting **Maximum Vault Jobs** as a threshold for queuing jobs from any vault. A Vault job runs if resources are available and the maximum number of Vault jobs has not been reached. If resources are not available to run jobs, Vault queues jobs until the resources are available.

Vault also uses locks on the duplication and eject steps of a job to enforce queueing for the jobs that contend for those resources.

Vault queues jobs as follows:

- If **Maximum Vault Jobs** is reached, any subsequent vault job is queued and its status is shown as Queued in the Activity Monitor.
- If a job needs to duplicate images and another job from the same vault is duplicating images, the job is queued and shown as Active in the Activity Monitor. More detailed information about the status of these Active jobs appears in the **Detailed Status** tab of the **Job Details** dialog box.
- If a job needs to eject media and another job from any vault is ejecting media in the same robot, the job is queued and shown as Active in the Activity Monitor.

Queued jobs are scheduled or run when the resources that are required for them become available.

For jobs that run simultaneously, the following restrictions exist:

- Vaults should not share on-site and off-site volume pools. Profiles within the same vault can use the same volume pools, but profiles from one vault cannot use the same volume pools as profiles from another vault.
- Vaults should not share off-site volume groups. Profiles within the same vault can use the same off-site volume groups, but profiles from one vault cannot use the same off-site volume groups as profiles from another vault.

See “About detailed Vault job status” on page 120.

About previewing a Vault session

Before you run a Vault session, you can preview the session to verify that the profile selects the appropriate images for off-site storage. To preview a session, use the `vltrun` command with the `-preview` option, specifying the robot number, vault, and profile as in the following example:

```
vltrun robot_number/vault_name/profile_name -preview
```

Alternatively, you can specify only the profile if it has a unique name.

The `vltrun -preview` option starts a new vault job and performs a search on the image catalog based on the criteria that is specified on the profile **Choose Backups** tab. Then `vltrun -preview` writes the names of the images to a `preview.list` file and exits. Vault does not act on the images selected.

After you run the preview option, examine the `preview.list` file, which is located in:

- UNIX:

```
/usr/opensv/netbackup/vault/sessions/vault_name/sidxxx
```

- Windows:

```
install_path\NetBackup\Vault\sessions\vault_name\sidxxx
```

Under certain circumstances, the `preview.list` file may contain more backup images than are vaulted:

- If the profile is configured to duplicate only disk images, selected images on removable media are not vaulted.
- If images in the list do not have a copy on media in one of the Off-site Volume Pools listed for the eject step, they are not vaulted.

Stopping a Vault session

You can use Activity Monitor to stop a Vault session. The Activity Monitor must be configured to show the Vault fields.

To stop a vault session

- 1 In the Activity Monitor, select the vault session you want to stop.
- 2 From the **Action** menu, select **Cancel Job**.

If a vault session fails, you cannot run a new session until the old session ends. Use **Cancel Job** to end the failed session.

About resuming a Vault session

If a vault job fails, first consult the NetBackup Administration Console Activity Monitor or the notification of session status (the session's `summary.log`). If they do not provide enough information to determine the cause of the problem, examine the other log files.

See “Debug logs” on page 219.

After you determine the cause, you can do one of the following:

- If the session reached the Eject step or attempted to generate reports before encountering problems, you can use `vltopmenu` (or `vlteject`) to finish the eject or reporting for the session.
- Start a new session for your profile. If you are doing duplication, Vault does not duplicate any images it already duplicated, but it does eject those images if the profile is configured to eject. (This behavior is similar to checkpoint and restart.)

About monitoring a Vault session

If you configure the NetBackup Administration Console Activity Monitor to display the Vault fields, you can use the Activity Monitor to monitor the progress of Vault jobs. For a Vault job that is initiated by the NetBackup scheduler, the **Policy** field displays the policy name. If the Vault job is run by any means other than the NetBackup scheduler, the Policy field is empty.

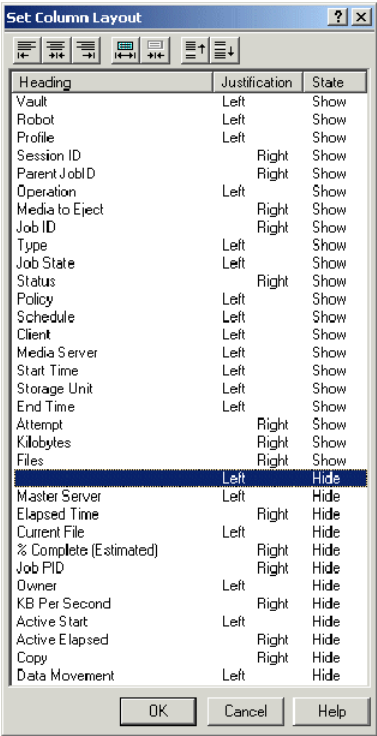
For information about configuring the Activity Monitor to display fields other than the default, see the "Monitoring NetBackup Activity" section in the *NetBackup Administrator's Guide, Volume I*.

The following are the fields that display Vault job attributes in the Activity Monitor:

Vault	The name of the vault under which this session is running.
Profile	The name of the profile that holds the processing information for a vault session.
Robot	The name of the robot the vault is associated with.
Session ID	The unique numeric value that identifies the vault session. Session ID assignment starts at 1 the first time a vault session runs after vault is installed. The value increments by one every time a new vault session runs.
Parent JobID	A Vault job that duplicates images starts one or more <code>bpduplicate</code> processes. Each of these child jobs refers to the job ID of the parent Vault job that started it.
Media to Eject	The number of tapes to be ejected for a vault session. If the profile is configured for deferred eject, the tapes are not ejected when the profile runs.
Operation	<div>The following values progress from the first value to the last as the vault job progresses:<ul style="list-style-type: none">■ Choosing Images.■ Duplicating Images.■ Choosing Media.■ Catalog Backup.■ Eject and Report.</div>

Figure 6-1 shows the Activity Monitor **Set Column Layout** dialog box showing the Vault fields at the top of the window:

Figure 6-1 Activity Monitor column layout dialog box



See “About detailed Vault job status” on page 120.

See “About extended error codes” on page 121.

About detailed Vault job status

If a job needs to duplicate images and another job from the same vault is duplicating images, the job is queued and shown as Active in the Activity Monitor. The **Detailed Status** tab of the **Job Details** dialog box shows information about such jobs.

The following are the possible messages written to the **Detailed Status** tab.

Please note that xxx represents a numeric value.

```
before eject, waiting for media to be unmounted; sleeping for
XXX seconds eject operation is waiting for available MAP elements
of robot duplication batch XXX started. Job ID: XXX
failed to eject XXX media. Reason: MEDIA_IN_USE
starting eject operation
suspend media for this session: failed to suspend XXX of XXX media
```



```
after duplication
suspend media for this session: failed to suspend XXX of XXX media
after catalog backup
suspend media for this session: failed to suspend XXX of XXX media
suspend media for this session: failed to suspend XXX of XXX media
at eject time
suspend media for next session: failed to suspend XXX of XXX media
vault global lock acquired
vault global lock released
vault session ID lock acquired
vault session ID lock released
vault duplication lock acquired
vault duplication lock released
vault assign slot lock acquired
vault assign slot lock released
vault eject lock acquired
vault eject lock released
vault waiting for global lock
vault waiting for session ID lock
vault waiting for duplication lock
vault waiting for assign slot lock
vault waiting for eject lock
vault lock acquisition failed
vault lock release failed
```

About extended error codes

Vault jobs may exit with exit status values greater than 255. These values are called extended error codes because they extend beyond the standard 255 NetBackup error codes. If a vault job exits with an extended error code, the exit status that is returned to the shell is 252. NetBackup adopted the convention that the exit status 252 means that an extended error code is returned by using stderr, in the following message:

```
EXIT status = extended error code
```

The Activity Monitor displays the extended error code rather than the value 252 returned to the shell in this case.

See “About errors returned by the Vault session” on page 214.

About the list of images to be vaulted

During a Vault session, Vault builds a list of images that are candidates for duplication or ejection.

The `preview.list` file, which resides in the session directory for the current Vault session, includes all images that match the criteria specified on the profile **Choose Backups** tab except for the following:

- If a copy of an image already is in the Off-site Volume Group, that image is not be included in the `preview.list` file. Because the images that have a copy in an Off-site Volume Group are already vaulted, Vault does not select them as candidates for vaulting.
- If the Source Volume Groups criteria in the **Locations** field on the **Choose Backups** tab is set to a specific volume group and if no copy of that image exists in that volume group, the image is not added to the `preview.list` file.

After the `preview.list` file is generated, Vault evaluates the images in it to determine if they should be duplicated or ejected. Because several filters (other profile configuration options) can exclude an image from duplication and ejection, the `preview.list` file can be a superset of the images that are eventually duplicated by the session.

About duplication exclusions

The following can eliminate an image from duplication:

- Catalog backup images are not duplicated.
- If Disk Only is specified on the **Duplication** tab, an image that has no disk copy is not duplicated.
- If Vault determines that an image is already duplicated, Vault does not duplicate the image again. Vault uses the following criteria to determine if an image is already duplicated:
 - For One Copy Only. If the image exists in the Off-site Volume Pool, Vault does not duplicate it. Conversely, if a copy of the image is not in the Off-site Volume Pool, Vault duplicates it.
 - For Concurrent Copies. Vault uses the **For Each Image If This Copy Fails** setting (**Continue** or **Fail All Copies**) to decide whether or not to duplicate an image. Each of the copies has its own **...If This Copy Fails** setting. Vault interprets the user's intent as follows:

Continue	If the setting for the copy is Continue , that copy is not critically important to the user. The duplication job ends with a partially successful (1) status if at least one of the other copies succeeds. If the current copy is the only one that fails, Vault does not re-duplicate the image the next time the profile runs. If all copies are set to Continue , at least one of those copies must exist or Vault duplicates the image.
Fail All Copies	If the setting for the copy is Fail All Copies , that copy is critically important to the user, and none of the copies are successful. This failure forces Vault to retry the duplication the next time the profile runs if that image is selected for duplication (if the time window of the profile allows that image to be selected again). However, if an unscheduled (and unlikely) event creates copies of the image, more than one copy of the image may be assigned to the destination volume pools. If the duplication operation results in more than the Maximum Backup Copies, the duplication step fails. (Maximum Backup Copies are configured in NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes .)

About ejection exclusions

Vault ejects media that are listed in the `eject.list` file. If the profile skips the duplication step and an image in `preview.list` has no copy in an off-site volume pool (configured on the **Eject** tab), it is not ejected.

About Vault resiliency

The functionality that Vault uses to build the list of images to be duplicated and ejected allows Vault to do the following:

- Duplicate or eject images that were not processed because of a problem during the previous run of the profile. By configuring the image selection period to be a sufficient length of time, the Vault profile duplicates an image if the duplication of that image failed during the previous run of that profile.
See “About overlapping the time window in the profile” on page 30.
- Not duplicate images that were successfully duplicated by a previous job. You can rerun a Vault session that was only partially successful, and Vault does not duplicate an image that was duplicated by a previous job. This feature provides for maximum flexibility or resiliency by letting you configure a longer image selection period without reduplicating images.

One Vault profile can vault original backup images from some NetBackup backup policies and also duplicate and vault images from other backup policies.

About ejecting media

If you configure a profile to defer ejection, you must perform or schedule another action to eject media.

You can use one of the following actions to eject media that was not ejected by a profile that selected or duplicated images:

- Manually by using the Vault Management node in the NetBackup Administration Console
- Manually by using the Vault Operator Menu
- Manually by using the `vlteject` command
- Automatically by creating and scheduling a Vault policy and entering the appropriate `vlteject` command and options in the file list

Note: You must use one of the Vault methods to eject media. If you use a NetBackup or Media Manager option to eject media, the correct database entries are not made and the Vault reports are not accurate.

For information about Vault methods to eject media, refer to the following:

- See “Previewing media to be ejected” on page 124.
- See “Ejecting media by using the NetBackup Administration Console” on page 125.
- See “Ejecting media by using the Vault operator menu” on page 126.
- See “Ejecting media by using the `vlteject` command” on page 127.
- See “Ejecting media by using a Vault policy” on page 128.
- See “Consolidating ejects and reports” on page 129.

For other related information see the following:

- See “About media ejection” on page 102.
- See “About ACS MAP” on page 103.
- See “About eject mode (immediate or deferred)” on page 104.
- See “About media ejection timeout impact” on page 105.

Previewing media to be ejected

Before you eject media, you can preview the media to be ejected. To preview that media, you can use the administration console or the `vlteject` command.

To preview media to be ejected using the administration console

- 1 Select the vault or profile for which you want to eject media.
- 2 Select **Actions > Deferred Eject**.
- 3 In the **Deferred Eject** dialog box, if necessary, select a vault, profile, or session ID. The options that are selected in the dialog box depend on whether you are ejecting for all vaults, for a single vault, or for a profile.
- 4 Click **Get Preview** and then select one or more of the profiles in the **Eject Preview** window.

To preview media to be ejected using the vlteject command

- ◆ From a command prompt, enter the `vlteject` command in the following format, specifying the robot, vault, or session for which you want to preview ejected media:

```
vlteject -preview [-vault vault_name [-profile profile_name]]  
               [-profile robot_no/vault_name/profile_name]  
               [-robot robot_no]  
               [-sessionid id]
```

Ejecting media by using the NetBackup Administration Console

Use the NetBackup Administration Console to eject media and generate reports for all vaults, for a single vault, or for a profile for which media have not yet been ejected.

When you select **Deferred Eject**, the default selections on the **Deferred Eject** dialog box depend on whether you are ejecting for all vaults, for a single vault, or for a profile. From the dialog box, you can initiate the eject operation or preview the media to be ejected. The preview shows the session IDs for which the deferred eject occurs and the media IDs for each session selected. You also can select whether to generate the reports after the ejection.

To eject media by using the NetBackup Administration Console

- 1 Select the vault or profile for which you want to eject media.
- 2 Select **Actions > Deferred Eject**.

- 3 In the **Deferred Eject** dialog box, select or change any of the options. The options that are selected in the dialog box depend on whether you are ejecting for all vaults, for a single vault, or for a profile.

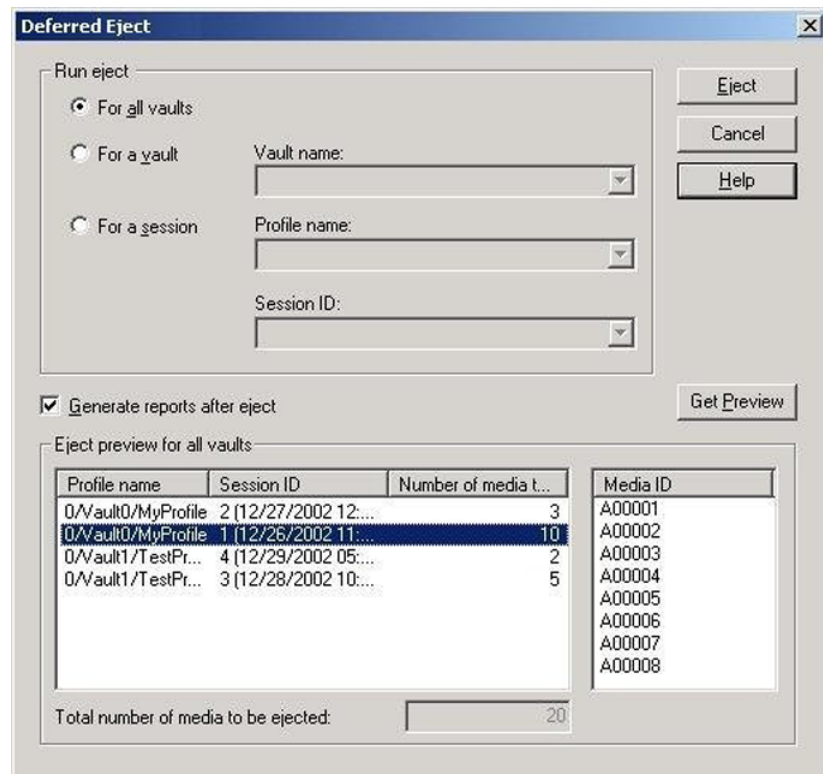
- 4 Click **Eject**.

To preview the media to be ejected, click **Get Preview**, and then select one or more of the profiles in the **Eject Preview** window.

To monitor the progress of or cancel the eject operation, use the NetBackup Administration Console Activity Monitor.

Figure 6-2 shows the Deferred Eject dialog box with all vaults selected and previewing the media that is ejected for the selected session.

Figure 6-2 Deferred Eject dialog box



Ejecting media by using the Vault operator menu

Use the **Vault Operator Menu** to eject media and generate reports for Vault sessions for which media are not yet ejected (the reports must be configured in the profiles).

The **Vault Operator Menu** calls the `vlteject` command to accomplish the media ejection. You also can use the **Vault Operator Menu** to preview the media to be ejected.

See “Vault operator menu interface ” on page 209.

To eject media by using the Vault operator menu

- 1 Start the **Vault Operator Menu** by running the `vltopmenu` command.
- 2 If necessary, select a profile.
- 3 Select one of the following options:
 - **Eject Media for This Session**
See “Ejecting media by using the `vlteject` command” on page 127.
See “Ejecting media by using a Vault policy” on page 128.
 - **Consolidate All Ejects**
 - **Consolidate All Reports and Ejects**
See “Consolidating ejects and reports” on page 129.

Ejecting media by using the `vlteject` command

Use the `vlteject` command to eject media and generate reports for Vault sessions for which media are not yet ejected (the reports must be configured in the profiles). The `vlteject` command can process the pending ejects or reports for all robots (all sessions for all vaults), for all sessions for a single vault, or for a specific Vault session.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-vault vault_name]
[-vault vault_name [-profile profile_name]]
[-profile robot_no/vault_name/profile_name]
[-legacy][-robot robot_no] [-auto y|n]
[-eject_delay seconds][-sessionid id]
```

The `vlteject` command resides in the following directory:

- **UNIX**

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

- **Windows**

```
install_path\NetBackup\bin
```

The following is a `vlteject` command example that ejects media for all robots that have sessions for which media has not yet ejected and that generates the reports:

```
vlteject -eject -report
```

The following example ejects all media that has not yet ejected for all sessions for the CustomerDB vault and generates reports:

```
vlteject -vault CustomerDB -eject -report
```

The following is a `vlteject` command example that previews the media to be ejected for the CustomerDB vault:

```
vlteject -vault CustomerDB -preview
```

See the *NetBackup Commands Reference Guide* for more information about the `vlteject` command.

See also "Using NetBackup Commands" in the NetBackup Administration Console help.

To eject media by using the `vlteject` command

- 1 In a terminal window or command window, change to the directory in which the `vlteject` command resides.
- 2 Run the command, using the appropriate options and parameters.

Ejecting media by using a Vault policy

Use a Vault policy to eject media or generate reports for the Vault sessions that have completed already and for which media have not ejected. In the Vault policy, specify Vault as the policy type, do not specify clients, and specify the `vlteject` command on the **Backup Selections** tab.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name]  

[-profile robot_no/vault_name/profile_name]  

[-robot robot_no] [-vault vault_name [-sessionid id]]  

[-auto y|n] [-eject_delay seconds] [-legacy]
```

The `vlteject` command resides in the following directory:

- UNIX

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

- Windows


```
install_path\NetBackup\bin
```

The following is an `vlteject` command example that ejects media for all robots that have sessions for which media has not yet ejected and generates the reports:

```
vlteject -eject -report
```

The following example ejects all media that has not yet ejected for all sessions for the CustomerDB vault and generates reports:

```
vlteject -vault CustomerDB -eject -report
```

See the *NetBackup Administrator's Guide, Volume I*, for more information about creating NetBackup policies.

See the *NetBackup Commands Reference Guide* for more information about the `vlteject` command.

See also "Using NetBackup Commands" in the NetBackup Administration Console help.

To create a Vault policy that ejects media

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.
- 4 On the **Attributes** tab, select **Vault** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule.
The type of backup defaults to Automatic Vault.
- 6 Complete the schedule.
- 7 Bypass the **Client** tab (clients are not specified for Vault jobs).
- 8 On the **Backup Selections** tab, enter the `vlteject` command and the appropriate options for the policy.
- 9 Click **OK**.

Consolidating ejects and reports

You can eject media from more than one vault session, which is known as consolidating ejects. For example, you may use one vault policy to duplicate media daily but eject media only at the end of the week.

If you consolidate ejects, you should also consolidate reports. Some restrictions may apply if you consolidate reports. By default, you cannot consolidate reports between vaults that use slots and vaults that use containers.

See “About consolidating reports” on page 179.

See “Eject tab (Profile dialog box)” on page 101.

See “Reports tab (Profile dialog box)” on page 107.

To consolidate ejects and reports for a profile

- 1 Select **Deferred Eject** on the profile **Eject** tab.

This action ensures that tapes are not ejected automatically for each Vault session.

- 2 Select **Deferred Reports** on the profile **Reports** tab.

This action ensures that reports are not generated automatically for each Vault session.

- 3 Eject media and generate reports by using one of the methods that are described in Ejecting media.

See “About ejecting media” on page 124.

About injecting media

In a normal volume rotation, you have to inject media back into a robot after media expires and is returned from your off-site storage location so that it is available for reuse. You also may need to inject unexpired media for restore or disaster recovery operations.

Injecting media updates the NetBackup and Media Manager catalogs so that the correct location of the media is recorded. If the robot does not have a barcode reader to identify the media being injected, you still must use an inject option so the location of the media is updated in the databases.

How you accomplish the process of injecting the media depends on the robot library as follows:

- If your library has a media access port (MAP), you insert the media to be injected into the MAP. Then use one of the injecting options that is discussed in this topic to move that media from the MAP to the library slots. If the library has a barcode reader, the appropriate database changes are made automatically.
- If the library does not have a MAP, you insert the media into the library slots or into a cartridge which is then placed into the robot. If the library has a barcode reader, the appropriate database changes are made automatically.

- If your library does not have a barcode reader, you must use the **Move media** option of the NetBackup Administration Console so the databases are updated.

You can inject media as follows:

- See “Injecting media for libraries with and without barcode readers” on page 131.
- See “Injecting media by using the Vault Operator Menu” on page 132.
- See “Injecting media by using the vltinject command” on page 133.

The vault fields in the Media Manager database are cleared when the media are unassigned while in a robotic volume group or moved into a robotic volume group and then unassigned (that is, injected back into the robot).

The following are the Media Manager database fields dedicated to Vault information:

vltcid	The ID of the container (container vaulting only).
vltname	The name of the vault.
vltreturn	The date the volume or container should be returned from the off-site vault.
vltsent	The date the volume or container was sent off-site.
vltsid	The ID of the session that vaulted the volume or container.
vltslot	The ID of the slot in which the volume resides in the off-site vault (slot vaulting only).

Volume pool, volume group, and media description fields are used for all volumes, not only the volumes that Vault uses. So, they are not cleared when media are injected back into a robot. You can, however, configure NetBackup so that the media description field is cleared.

See “About clearing the media description field” on page 154.

Injecting media for libraries with and without barcode readers

Use the NetBackup Administration Console to inject media for libraries that have barcode readers and libraries that do not have barcode readers.

To inject media for libraries with barcode readers

- 1 Insert the media into the robotic library slots or media access port.
- 2 In the NetBackup Administration Console, click **Media and Device Management > Media > Robots**.
- 3 Select the robotic library where you inserted the volume.

- 4 Click **Actions > Inventory Robot**.
- 5 In the Inventory operation section, select **Update volume configuration**.
- 6 If your robot has a media access port into which you placed the media, select **Empty** media access port before update in the Inventory operation section.
- 7 To configure advanced options, click **Advanced Options**.
- 8 To clear any previous display in the Results section, click **Clear Results**.
- 9 Click **Start** to start the update.
- 10 Repeat as necessary until all media are injected.

To inject media for libraries without barcode readers

- 1 Insert the media into the robotic library slots (or into the cartridge and then inject the cartridge into the robot).
- 2 In the NetBackup Administration Console, click **Media and Device Management > Media**.
- 3 Select the volume to be injected into the library.
- 4 Click **Actions > Move**.
- 5 In the **Move Volumes** dialog box, select or enter the robot, volume group, and slot number.

Use the **First Slot Number** field to enter the slot into which you placed the volume.
- 6 Click **OK** to move the volume.
- 7 Repeat as necessary until all media are injected.

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

Injecting media by using the Vault Operator Menu

You can use the Vault Operator Menu to inject media. The Vault Operator Menu calls the `vinject` application programming interface (API) to inject the media.

Note: Ensure that all the media in the MAP are from the current vault (that is, the vault for the currently selected profile in `vltopmenu`). If they are not, inject fails.

To inject media by using the Vault Operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 If necessary, select a profile.

- 3 If necessary, select a session.
 - 4 Load the media in the robot's media access port.
 - 5 Select **Inject Media into Robot**.
 - 6 Repeat until all media are injected into the robot.
- See "Vault operator menu interface " on page 209.

Injecting media by using the `vltinject` command

The `vltinject` command injects volumes into a robot and updates the Media Manager volume database. It requires as an option the name of a profile (if unique) or a robot number, vault, and profile name.

The following is the format of the `vltinject` command:

```
vltinject profile|robot/vault/profile
```

The following example command injects volumes that were vaulted by the Payroll profile and that were returned from the off-site vault:

```
vltinject Payroll
```

Note: If you use the same profile name across vaults in your vault configuration, then you should specify the `robot/vault/profile` attributes.

The following example injects volumes that were vaulted by the Weekly profile in the Finance vault and that were returned from the off-site vault:

```
vltinject 8/Finance/Weekly
```

See the *NetBackup Commands Reference Guide* for more information about the `vltinject` command.

See also "Using NetBackup Commands" in the NetBackup Administration Console help.

To inject media by using the `vltinject` command

- 1 In a terminal window or command window, change to the directory in which the `vltinject` command resides, as follows:

- UNIX

```
/usr/openv/netbackup/bin
```

- Windows

```
install_path\NetBackup\bin
```

- 2 Load the media to be injected into the robot's media access port.
- 3 Run the command, using the appropriate options and parameters.
- 4 Repeat until all media are injected.

About using containers

A container is a box in which you can place media and then transfer that box to your off-site storage location. When you configure a vault, you select whether the media are stored in containers or slots at your off-site storage location. Vault tracks, reports, and recalls your media regardless of how the media are transferred and stored off site.

After the media are ejected from your robot, you must add the media logically to containers by using either the Vault Operator Menu or the `vltcontainers` command.

The options available for adding media to containers are as follows:

- Enter the container ID and media IDs by typing them in with the keyboard. Using this method, you can add media to more than one container.
- Scan the container ID and media IDs by using a keyboard interface barcode reader. (Keyboard interface readers are also known as keyboard wedge readers because they connect, or wedge, between the keyboard and the keyboard port on your computer.) Using this method, you can add media to more than one container.
- Read an input file that contains the IDs or numeric equivalents of barcodes of all the media to be added to one container. If you have a barcode reader that can write to a file, you can scan media barcodes and use that output file as input for the `vltcontainers` command.
- Add all the media that is ejected by a specific session to one container.

The default return date of a container is the date of the volume in the container that is returned the latest. You can change the return date during the container ID and media ID entry process or at any time thereafter before a container is recalled.

You also can delete a container from the NetBackup and Media Manager databases. The following describes when you should delete an empty container:

- If a container becomes empty due to moving media to other containers, Vault deletes that empty container.
- If a container becomes empty due to a media inject, the empty container remains and the user will have to manually delete it using the `vltcontainers` command.

The following example shows you how to delete container ABC123 from the NetBackup and Media Manager catalogs.

```
vltcontainers -delete -vltcid ABC123
```

Note: The container must be empty before it can be deleted.

If you use containers, Vault reports on the containers and media outside the context of a profile or session.

See “About vaulting media in containers” on page 135.

See “About managing containers and media” on page 138.

See “Generating a Container Inventory Report” on page 140.

About vaulting media in containers

You can use either the Vault Operator Menu or the `vltcontainers` command to add media IDs to containers.

See “Vaulting container media by using the Vault operator menu” on page 135.

See “Vaulting container media by using the `vltcontainers` command” on page 136.

Vaulting container media by using the Vault operator menu

After the media eject from your robot, you can use the Vault Operator Menu to enter the container ID and media IDs.

See “Vault operator menu interface ” on page 209.

To vault media in containers by using the Vault operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 Eject the media to be added to the containers.
- 3 Select **Container Management**.

4 Select one of the following options:

- | | |
|---|--|
| Move media into one or more containers | Select if you want to use the keyboard to enter media IDs and container IDs or use a barcode scanner to scan the barcodes on the volumes and containers. |
| Move all media ejected by this session, into one container | Select if you want to add all media that are ejected by a session to a container. |
| Move all media listed in a file, into one container | Select if you want to add all media that are listed in an input file to a container. |

5 Follow the prompts to complete the process of logically moving media into containers.

Vaulting container media by using the `vltcontainers` command

After the media eject from your robot, you can use the `vltcontainers` command to enter the container ID and media IDs. The following is the format of the `vltcontainers` command:

```
vltcontainers -run [-rn robot_number]
vltcontainers -run -usingbarcodes [-rn robot_number]
vltcontainers -run -vltcid container_id -vault vault_name -sessionid session_id
vltcontainers -run -vltcid container_id -f file_name [-rn robot_number] [-usingbarcodes]
vltcontainers -view [-vltcid container_id]
vltcontainers -change -vltcid container_id -rd return_date
vltcontainers -delete -vltcid container_id
vltcontainers -version
```

The following examples show how to use the `vltcontainers` command to add media to a container:

- To add the volumes that were ejected from robot number 0 to containers and enter the IDs by typing them in, use the following command:

```
vltcontainers -run -rn 0
```


- To add the volumes that were ejected from robot number 0 to containers and use a barcode reader to scan the container ID and media IDs, use the following command:

```
vltcontainers -run -usingbarcodes -rn 0
```

- To change the return date of container ABC123 to December 07, 2004, use the following command:

```
vltcontainers -change -vltcid ABC123 -rd 12/07/2004
```

- To delete container ABC123 from the NetBackup and Media Manager catalogs, use the following command:

```
vltcontainers -delete -vltcid ABC123
```

- To add all media that were ejected by session 4 of vault MyVault_Cntrs to container ABC123, use the following command:

```
vltcontainers -run -vltcid ABC123 -vault MyVault_Cntrs  
-sessionid 4
```

- To add media that were listed in file C:\home\jack\medialist that are ejected from robot number 0 to container ABC123, use the following command:

```
vltcontainers -run -vltcid ABC123 -f C:\home\jack\medialist -rn  
0
```

- To add media to container ABC123 that was ejected from a robot that is attached to the master server and read the barcodes for that media from file C:\home\jack\medialist, use the following command:

```
vltcontainers -run -vltcid ABC123 -f C:\home\jack\medialist  
-usingbarcodes
```

See the *NetBackup Commands Reference Guide* for more information about the `vltcontainers` command.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

To vault container media by using the `vltcontainers` command

- 1 Eject the media to be added to the containers.
- 2 Run the `vltcontainers` command, using the appropriate options and parameters.
- 3 Follow the prompts to move the media logically into containers.

About managing containers and media

After the media and containers are sent to your off-site storage location, you can still perform tasks to manage the containers and media. You can view and change return dates of containers. If a container is recalled and is empty of media, you can delete the information about a container from the NetBackup and Media Manager databases.

See “Managing container media by using the Vault operator menu” on page 138.

See “Managing container media by using the `vltcontainers` command” on page 139.

Managing container media by using the Vault operator menu

You can use the Vault Operator Menu to change a container return date and to delete the information about a container from the NetBackup databases.

See “Vault operator menu interface ” on page 209.

Note: (This note applies to Iron Mountain users.) To change a container return date, change the date using the Vault Operator Menu or the `vltcontainers` command. Then resend the Container Inventory Report or the Iron Mountain FTP file to Iron Mountain. Do not use the Iron Mountain account management facilities to change a container return date. If you do, the Vault reports do not match the report information that is maintained by Iron Mountain.

To view a container return date by using the Vault operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 Select **Container Management > View a container’s return date**.
- 3 Follow the prompts to enter a container name.

To change a container return date by using the Vault operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 Select **Container Management > Change a container's return date**.
- 3 Follow the prompts to enter container names and change dates.

To delete a container by using the Vault operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 Select **Container Management > Delete a container**.
- 3 Follow the prompts to enter a container name and delete the records of a container.

If a container becomes empty after it is recalled and all media that reside in it are either injected back into the robot or assigned to another container, it is deleted from the NetBackup and Media Manager databases.

Managing container media by using the `vltcontainers` command

You can use the `vltcontainers` command to view and change a container return date and to delete the information about a container from the NetBackup and Media Manager databases.

See the *NetBackup Commands Reference Guide* for more information about the `vltcontainers` command.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

Note: (This note applies to Iron Mountain users.) To change a container return date, change the date using the Vault Operator Menu or the `vltcontainers` command. Then resend the Container Inventory Report or the Iron Mountain FTP file to Iron Mountain. Do not use the Iron Mountain account management facilities to change a container return date. If you do, the Vault reports do not match the report information that is maintained by Iron Mountain.

To view a container return date by using the `vltcontainers` command

- ◆ Run the `vltcontainers` command by using the `-view` option. For example, to view the return date of container ABC123, use the following command:

```
vltcontainers -view -vltcid ABC123
```

To change a container return date by using the `vltcontainers` command

- ◆ Run the `vltcontainers` command by using the `-change` option and specifying the `-vltcid` parameter and argument and `-rd` parameter and argument. For example, to change the return date of container ABC123 to December 07, 2004, use the following command:

```
vltcontainers -change -vltcid ABC123 -rd 12/07/2004
```

To delete a container by using the `vltcontainers` command

- ◆ Run the `vltcontainers` command by using the `-delete` option and specifying the `-vltcid` parameter and argument. For example, to delete container ABC123 from the NetBackup and Media Manager catalogs, use the following command:

```
vltcontainers -delete -vltcid ABC123
```

To be deleted, a container must be empty.

Generating a Container Inventory Report

The Container Inventory Report shows all the containers that are configured in your vaulting environment, the return date of each container, and the media that are in each container. Alternatively, you can specify a container ID to generate a report of the media in a specific container.

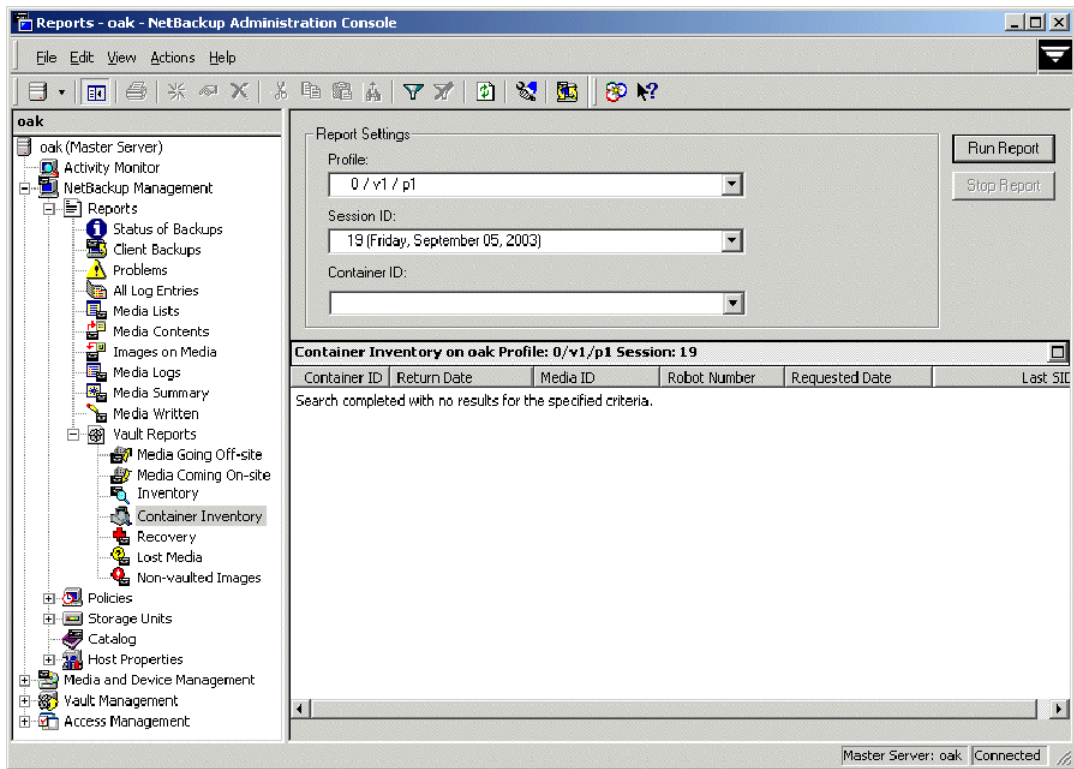
If you are using containers, all of the other Vault reports list the ID of the container in which the volume resides rather than a slot number. Reports do not show container information until after you add container and media IDs in Vault. Media are removed logically from a container when they are injected back into the robot.

You also can generate the Container Inventory Report by using the Vault Operator Menu (Run **Individual Reports > Container Inventory**).

Note: You must specify a directory path where the report is generated.

See “Vault operator menu interface ” on page 209.

Figure 6-3 provides an example of the **Container Inventory Report** window:

Figure 6-3 Container Inventory Report window**To generate a container inventory report**

- 1 In the NetBackup Administration Console, select **Reports > Vault Reports > Container Inventory**.
- 2 In the **Container ID** field, select **All Containers** or the ID of the container for which you want a report.
- 3 Click **Run Report**.

Assigning multiple retentions with one profile

Different types of data often are retained for different lengths of time. For example, you may want to vault your finance data for seven years and your customer data for 20 years. To do this, the off-site copy of your backups will have different retentions based on the type of data. Vault can process different types of data individually if your backups are organized based on the type of data being protected.

For example, if you have separate backup policies based on the data type, such as a Finance backup policy and a CustomerDB backup policy.

When a Vault session creates a duplicate image, Vault usually assigns the same retention to all of the duplicates created by using one of the following duplication options:

- Specifying **No Change** keeps the same expiration date as the original copy.
- Specifying a numeric retention level applies that retention, calculated from the backup time of the original image.

Alternatively, you can configure Vault to calculate a retention for the duplicate copy based on the type of data. During duplication, specifying Use Mappings instructs the profile to use the alternative retention mappings. The retention for a duplicate copy of a particular type of data is based on the retention level that is assigned by the backup policy for that type of data. The retention mapping converts the original backup image's retention to a new retention for the duplicated copy.

For example, suppose you want to retain the on-site copy of all your data for two weeks, the off-site copy of your Finance data for seven years, and the off-site copy of your CustomerDB data for 20 years.

To assign multiple retentions with one profile

- 1** Using Host Properties in the NetBackup Administration Console, configure retention levels 1 and 11 to be two weeks, retention level 12 to be seven years, and retention level 13 to be 20 years.
- 2** In your Finance backup policy, assign retention level 1 (two weeks) to the first (or only) copy configured in the schedule.
- 3** In your CustomerDB policy, assign retention level 11 (two weeks) to the first (or only) copy configured in the schedule.

- 4 In your Vault profile, on the **Duplication** tab configure **Retention Level** to be **Use Mappings**.

5 Configure the retention mappings as follows:

0	0
1	12
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	13
12	12

With this mapping, duplicated images from the Finance policy or schedule are assigned a retention level of 12 (seven years). Duplicated images from the CustomerDB policy or schedule are assigned a retention level of 13 (20 years). The duplicated images are written to different media if the **Allow Multiple Retentions Per Media** property is not set (**NetBackup Management > Host Properties > Master Server > Media**).

You can configure retention mappings globally for all vaults by using the **Retention Mappings** tab on the **Vault Management Properties** dialog box or for each vault by using the **Retention Mappings** tab on the **Vault** dialog box.

When you configure a profile, you can specify normal retention calculation for some duplication rules and alternative retention mappings for other duplication rules.

The values for the retention levels are configured in **NetBackup Management > Host Properties > Master Server > Retention Periods**.

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

See "Retention Mappings tab (Vault Management Properties)" on page 68.

About vaulting additional volumes

Usually, you create the necessary copies of backup media during a NetBackup policy job or a Vault profile duplication job. Then the Vault profile ejects the media for transfer off site. After the Vault profile is run, you cannot run the profile again to create additional copies of media that was already sent off site.

However, you can use other methods to create and eject additional copies of backup media after the NetBackup policy and Vault profile have run. You can duplicate the volume manually or you can configure Vault to duplicate the volume. If you want to duplicate and eject one or several additional volumes one time only, the easiest solution is to duplicate the volume manually.

To duplicate an additional volume, the primary copy of the volume must be in the robot. If the primary copy is not in the robot but a duplicate copy is, you can use the `bpchangeprimary` command to change the duplicate to primary before you create an additional volume.

See the *NetBackup Commands Reference Guide* for information about the `bpchangeprimary` command.

See also "Using NetBackup Commands" in the NetBackup Administration Console help.

Duplicating a volume manually

If you use the following instructions to duplicate a volume manually, the tape appears on the Picking List for Vault report when it expires. It is recalled from the off-site vault as part of your normal operations.

To duplicate a volume manually

- 1 Duplicate the volume manually by using the `bpduplicate` command. When duplicating the volume, specify the same off-site volume pool that is used for the volume already vaulted.
- 2 Assign the vault vendor's slot number for the volume by using the `vltoffsitemedia` command. The slot number is assigned in the `vltslot` field.

You can assign values to other vault fields if desired.

Do not assign a value to the `vltreturn` field. If you assign a value, the volume never appears on the Picking List for Vault report.

- 3 Move the volume into the off-site volume group by using the `vmchange` command. Use the same off-site volume group as the first vaulted copy.

If the volume is in the same off-site volume group and the same off-site volume pool that is used by the regularly scheduled Vault profile, this volume appears on the Picking List for Vault report when the first vaulted copy expires (if you did not assign a value to the `vlreturn` field).
- 4 Eject the volume.
- 5 Edit the file of the Picking List for Robot report to insert this volume into the list. Then print the report and give it to the vault vendor.

See the *NetBackup Commands Reference Guide* for information about the `bpduplicate` and `vltoffsitemedia` commands.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

Duplicating a volume by using Vault

If you use Vault to create and eject an additional copy of a volume that is already in the off-site vault, you must use a different vault, off-site volume group, and off-site volume pool than the first volume. The additional volume does not appear in the Picking List for Vault, so you must use the Lost Media Report to recall the volume after it expires.

To duplicate a volume by using Vault

- 1 Create and configure a new vault. Use a different off-site volume group than the first volume.
- 2 Create and configure a new profile to duplicate and eject the volume. Assign the volume to a different off-site volume pool than the first volume.
- 3 Configure the eject step of the profile to search the off-site volume pool in which the additional volume was assigned.
- 4 Run the profile.
- 5 To recall the volume after it expires, run the Lost Media Report.

If you run the Lost Media Report as part of your normal operations, the volume appears on the report after it expires.

Revaulting unexpired media

When you inject unexpired media from off-site storage back into a robot (for example, to perform a restore), you should revault the media. If you have to revault many

tapes, you should create a new profile to revault them. If only a few tapes are to be revaulted, revaulting them manually may be the easiest and fastest option.

To revault unexpired media by creating a new profile

- 1 Copy the original Vault profile that was used to eject the media.
- 2 In this new profile, change the **Choose Backups** time window to shift it far enough back in time so that it selects the images on the volumes that you want to revault.

- 3 Start a session using this new vault profile.

Vault recognizes that copies of images eligible to be vaulted exist and does not duplicate the images even if the duplication step is configured. The profile ejects the volumes to be revaulted.

- 4 If you are vaulting containers, logically add the volumes to containers. The container ID field is cleared when media that was vaulted in containers is injected back into the robot, so you must add the media to containers.

See "About using containers" on page 134.

If you are vaulting media in slots, Vault assumes that the media are returned to the same slots in off-site storage from which they were recalled.

- 5 Delete the new profile you created to do the revaulting.
- 6 If you froze your media during the data restore process, use the `bpmedia` command to unfreeze it.

If you froze the media, you have to unfreeze it so it is recalled and returned to volume pool rotation when it expires. Vaulted media that are suspended are unsuspended automatically when they expire and are recalled.
- 7 Return the media to your vault vendor so that all backups on that media will be available for future disaster recovery.
- 8 Run the Recovery Report to ensure that media are available in off-site storage for future use.

See the *NetBackup Commands Reference Guide* for information about the `bpmedia` command.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

To revault media manually

- 1 Manually eject the media, using one of the following methods:
 - Use the `vmchange` command.

- In the NetBackup Administration Console, select the media ID and then select the **Eject Volumes from Robot....** operation on the **Actions** menu.

`vlteject` and `vltopmenu` do not work for this purpose.

- 2 Manually transfer the media to the off-site volume group, using one of the following methods:
 - Use the `vmchange` command.
 - In the NetBackup Administration Console, select the media ID and then select the **Change Volume Group....** operation on the **Actions** menu.
- 3 If you are vaulting containers, logically add the volumes to containers. The container ID field is cleared when media vaulted in containers is injected back into the robot, so you must add the media to containers.

See "About using containers" on page 134.

If you are vaulting media in slots, Vault assumes that the media are returned to the same slots in off-site storage from which they were recalled.

- 4 If you froze your media during the data restore process, use the `bpmedia` command to unfreeze it.

If you froze the media, you have to unfreeze it so that it is recalled and returned to volume pool rotation when it expires. Vaulted media that are suspended are unsuspended automatically when they expire and are recalled.

- 5 Return the media to your vault vendor so that all backups on that media are available for future disaster recovery.
- 6 Run the Recovery report to ensure that the media are available for future disaster recovery operations.

See the *NetBackup Commands Reference Guide* for information about the `vmchange` command.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

About tracking volumes not ejected by Vault

Eject is the event that Vault uses to update the NetBackup database for the location of volumes and to recall volumes. In normal operation, Vault must eject volumes so they are tracked and recalled from off-site storage after they expire.

If you ejected volumes or removed them from your robot without using a Vault eject method, it is still possible to use Vault to track them if they are in a volume pool you are using for off-site media. Using commands, you can change the attributes for

volumes that were not ejected by Vault so they appear in the Vault reports and are recalled when they expire.

Note: This process works only if the volumes already are in a volume pool on the **Eject** tab of a profile. You cannot change the pool of an assigned volume.

To track volumes not ejected by Vault do the following:

- Use the `vmchange` command to change the volume group of the volumes, which Vault uses to track their location. For example, the following `vmchange` command changes the volume group of volume A00001:

```
vmchange -new_v offsite_volgrp -m A00001
```

- Use the `vltoffsitemedia` command to change the Vault-specific attributes. The following `vltoffsitemedia` example changes the vault attributes of volume A00001:

```
vltoffsitemedia -change -m A00001 -vltname offsite_vault -vltsent  
07/03/2004 -vltreturn 0 -vltslot 99 -vltsession 33
```

A catalog volume is processed the same except the return date is set to the date the volume should be recalled.

If you are adding the volumes to slots at your off-site storage location, you can use the `vltoffsitemedia` command with the `-list` option to find empty slots into which you can add the volumes.

If you are placing the volumes in containers, use the `vltcontainers` command to add the volumes logically to containers after you specify the off-site volume group and the vault attributes.

See “About using containers” on page 134.

The default return date of a container is the date of the volume in the container that is returned the latest. You may have to change the container return date if the volume you add expires later than any volume already in the container.

The following are the `vltoffsitemedia` options you can use to set the necessary volume attributes:

<code>-vltname</code> <code>vault_name</code>	The name of the vault.
<code>-vltreturn</code> <code>date</code>	Set the return date to 0. Vault uses the latest expiration date of the images on the volume as the return date. Exception: if the volume is a NetBackup catalog backup volume, set the date the volume should be recalled from off-site.

<code>-vltsent date</code>	Set the sent date to the date the volume was ejected. The format of the date depends on your locale setting. For the C locale, the date syntax is mm/dd/ yyyy [hh[:mm[: ss]]].
<code>-vltsession session_id</code>	The ID of the session that vaulted the volume or container. Set it to a nonzero number that is different from existing session IDs
<code>-vltslot slot_id</code>	The ID of the slot in which the volume resides in the off-site vault. Ensure that this is an empty slot at your off-site storage location. If you are placing the volume in a container, do not specify this option.

Vaulting non-NetBackup media managed by Media Manager

Vault can eject and track media that was not created by NetBackup if the media are managed by Media Manager. Vault uses an eject notify script to add valid media IDs to the eject list. Vault ejects that media if you add the volume pool in which that media resides to the **Offsite Volume Pools** list on the **Eject** tab of the **Profile** dialog box.

Vault assigns vendor slot or container IDs to the media. Those media appear on the Picking List for Robot and Picking List for Vault reports. Vault ejects notify script media even if no other media are selected for ejection by the Vault profile.

Notify script templates are provided with Vault. The following procedure documents how to copy and modify the `vlt_ejectlist_notify` script. The scripts include information about how to modify and test them.

See “About using notify scripts” on page 151.

To vault non-NetBackup media managed by Media Manager

- 1 Copy the `vlt_ejectlist_notify` script and name it appropriately (that is, add the appropriate extension to the name).
- 2 Edit the script as follows:
 - Add the media IDs of the non-NetBackup media that you want to eject.
 - Add a `vltoffsitemedia` command and use the `vltreturn` option to set a date to recall the media from the vault.

The script runs the `vltoffsitemedia` command(s) and assign the expiration date(s). The media appears on the Picking List for Vault on the date it expires.

- 3 Place the script in the NetBackup `bin` directory.
- 4 Configure a Vault profile so that it includes the volume pool in which the media are assigned in the **Off-site Volume Pools** list on the **Eject** tab of the **Profile** dialog box.

When the profile runs and if the script runs successfully, the media ejects.

See the *NetBackup Commands Reference Guide* for information about the `vltoffsitemedia` command.

See also, "Using NetBackup Commands" in the NetBackup Administration Console help.

About notifying a tape operator when an eject begins

Vault can send an email notification when the eject process begins. Eject notification is configured for each profile on the **Eject** tab, for each robot on the **Vault Robot** dialog box, and globally for Vault on the **Vault Management Properties** dialog box. Vault sends the notification to the first email addresses that are found in that order.

To configure eject email notification, see the following:

- See "About configuring Vault Management Properties" on page 64.
- See "Configuring robots in Vault" on page 70.
- See "Eject tab (Profile dialog box)" on page 101.
- See "About setting up email" on page 197.

About using notify scripts

A Vault job can call notify scripts at specific points during the execution of the job. Vault includes template notify scripts that you can customize for your use. You can use a script for a robot, a vault, or a profile.

The template notify scripts are in the following directory:

- UNIX

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

- Windows

```
install_path\NetBackup\bin\goodies
```

On Windows systems, the names of the scripts include a `.cmd` extension. They include instructions that can help you edit them for your needs.

To call and run a script, it must be copied to the NetBackup `bin` directory. A script must return a normal status (0) for the Vault job to continue processing. In case of failure, the script must return a nonzero status code to cause the Vault job to stop. On UNIX systems, the return status is communicated to the Vault job through the `exit` call. On Windows systems, the scripts communicate the return status in a file that is defined by the `EXIT_STATUS` environment variable, which is set by Vault.

The following scripts are provided with Vault:

<code>vlt_start_notify</code>	Called by the Vault session after it starts. For example, you can use it to send notification when the Vault job begins.
<code>vlt_ejectlist_notify</code>	Called by the Vault session before the list of media to be ejected (the <code>eject.list</code>) is built. Use this script to add media managed by Media Manager but not created by NetBackup or Vault to the eject list. The script writes media IDs to the <code>addon_medialist</code> file. Vault reads the <code>addon_medialist</code> file and ejects the media that is listed in that file during the current Vault session if the volume pool in which that media resides is in the Off-site Volume Pools list on the Eject tab of the Profile dialog box.
<code>vlt_starteject_notify</code>	Called by the Vault session after the <code>eject.list</code> file is built and before the automatic eject process begins. Use this script to send notification when the eject process begins or perhaps to suspend the media in the eject list. If the eject step is not configured for the profile, the <code>vlt_starteject_notify</code> script is not called.
<code>vlt_endeject_notify</code>	Called at the end of eject processing. Use this script to send notification when the eject process ends. If the eject step is not configured for the profile, the <code>vlt_endeject_notify</code> script is not called.
<code>vlt_end_notify</code>	Called by the Vault session immediately before it exits. One use for this script is to start another vault job. Then run Vault jobs in succession and avoid resource contention.

Before you use a notify script, ensure that your systems are set up properly for email.

See “About setting up email” on page 197.

For information on how to use the notify script, see the following:

- See “About notify script for a specific robot” on page 153.
- See “About notify script for a specific Vault” on page 153.
- See “About notify script for a specific profile” on page 153.
- See “Notify script order of execution” on page 154.

About notify script for a specific robot

You can use a notify script to create a unique, customized script for each robot in your configuration. To create a notify script for a specific robot, append the robot number to the script name and copy the script to the NetBackup `bin` directory.

For example, a `vlt_start_notify` script for a specific robot appears as follows:

```
vlt_start_notify.robot_number
```

The script is run for all profiles that are created for the robot.

Use the same methodology to create other notify scripts.

About notify script for a specific Vault

You can use a notify script to create a unique, customized script for each vault in your configuration. To create a notify script for a specific vault, append the robot number and vault name to the script name and copy the script to the NetBackup `bin` directory.

For example, a `vlt_start_notify` script for a specific robot or vault combination appears as follows:

```
vlt_start_notify.robot_number.vault_name
```

The script is run for all profiles that are created for a specific vault.

Use the same methodology to create other notify scripts.

About notify script for a specific profile

You can use a notify script to create a unique, customized script for each profile in your configuration. To create a notify script for a specific profile, append the robot

number, vault name, and profile name to the script name and copy the script to the NetBackup `bin` directory.

For example, a `vlt_start_notify` script for a specific robot or vault or profile combination appears as follows:

```
vlt_start_notify.robot_number.vault_name.profile_name
```

This script is run for a specific profile that is defined for a specific vault.

Use the same methodology to create other notify scripts.

Notify script order of execution

The notify scripts run in specific to general order, as follows:

1. `script_name.robot_number.vault_name.profile_name`
2. `script_name.robot_number.vault_name`
3. `script_name.robot_number`
4. `script_name`

About clearing the media description field

When media are returned from the off-site vault during a typical volume rotation, they are expired and ready for reuse. To avoid confusion, it may be helpful to clear the media description information when an expired volume is returned to the robot.

You can configure NetBackup so that the media description field is cleared when media are returned to the robot. To do so, use the `nbemmcmd` to set the `VAULT_CLEAR_MEDIA_DESC` parameter. The media description field clears when other Vault information clears from the Media Manager volume database.

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

See the *NetBackup Commands Reference Guide* for information about the `nbemmcmd` command.

See also "Using NetBackup Commands" in the NetBackup Administration Console help.

Restoring data from vaulted media

You may need to restore images from the media that is stored in your off-site vault. The high-level procedure in this section describes how to restore data from vaulted media.

To restore data from vaulted media

- 1 Recall the media.
- 2 Change the images to be recovered to primary (NetBackup restores from the primary image).

Use the `bpchangeprimary` command to promote a copy to primary.

See the *NetBackup Commands Reference Guide*.
- 3 If the media is not suspended or frozen, suspend the media.

Use the `bpmedia` command to suspend the media.

See the *NetBackup Commands Reference Guide*.
- 4 Inject the media into the robot.

See “About injecting media” on page 130.

Injecting the media moves it into the robot and also changes the off-site volume group attribute of the media to the robotic volume group so NetBackup knows that the media are in the robot.
- 5 Using the Backup, Archive, and Restore interface, restore the data.

See the *NetBackup Backup, Archive, and Restore Getting Started Guide*.
- 6 After restoring all the data, revault the media.

See “Revaulting unexpired media” on page 146.

Replacing damaged media

If media in your robot is damaged, you can use copies of the media (if available) from your off-site storage location to replace the damaged media. You also can use this process to recover images if the primary backup expired, the volume was overwritten, and a copy in off-site storage is still available.

Note: This image recovery process assumes that the NetBackup system and image catalog are current and up to date.

The instructions use an example to illustrate how to run the various commands that are used in the recovery process. Modify the command examples as appropriate for your purposes.

Most of the commands used to recover from damaged media are in the following directory:

- UNIX

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

■ Windows

```
install_path\netbackup\bin\admincmd
```

After you recover and restore the damaged media, you should revault the media so that it again is available for recovery.

See “Revaulting unexpired media” on page 146.

To replace damaged media

1 Identify the damaged media.

When you receive an error message during a restore, the errors are logged to the restore log and also show up on the Activity Monitor as the restore fails. You can set up a procedure using NetBackup scripts to send errors to an event management console to notify the storage administrator immediately of this type of media error.

2 Determine which backup images are on the damaged tape.

To identify all images on a specific tape, run the `bpimmedia` command. It scans the entire NetBackup image catalog, so it may take a few minutes depending on the size of that catalog. For example, the following shows that volume S05423 contains one image from client `fgolddust`. It also shows that this image was duplicated because it has (FRAG 2) entries. The full image name is `"fgolddust_0862806643"`:

```
# bpimmedia -mediaid S05423
IMAGE fgolddust 2 fgolddust_0862806643 golddust_BR1 0
Full_Weekly 0 3 19360 8654 85043 0 0
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0
*NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0
*NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0
*NULL*
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0
*NULL*
```

3 Determine which duplicate tapes were used and their host.

In step 2, the (FRAG 2) entries show that an image was duplicated: the (FRAG 2 1) entry is the duplicate copy. On copy 1 there were four fragments (usually due to multiplexing). The (FRAG 2 -1) entry is the true image restore duplicate. In this case, the image fgolddust_0862806643 uses media S04440 for duplicating all of the original fragments. This is normal because the original image was multiplexed onto four tapes, while the duplicate was de-multiplexed during image duplication, and could fit on one tape.

Also note that the host for the media is printed for each fragment, in this case nirvana. With media servers, the host could be different than the master server. Under Vault, the duplication should normally occur on the same server that made the original backup, so the host server names are the same for both copies of the image.

To confirm this information, use the bpimagelist command, as follows:

```
# bpimagelist -backupid fgolddust_0862806643
IMAGE fgolddust 0 0 2 fgolddust_0862806643 goldust_BR1 0 *NULL*
root Full_Weekly 0 3 862806643 4591 865485043 0 0 2356562 19360
2 7 1 goldust_BR1_0862806643_FULL.f *NULL* *NULL* 0 1 0 2
865830643 *NULL* 1 0 0 0 0 *NULL*
HISTO -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
FRAG 1 -1 2293 0 2 6 2 S05423 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 1 232848 0 2 6 1 S02643 nirvana 64512 2 862804830 3 0
*NULL*
FRAG 1 2 1225539 0 2 6 2 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 3 70182 0 2 6 3 S02643 nirvana 64512 0 862804830 3 0
*NULL*
FRAG 1 4 825700 0 2 6 1 S05423 nirvana 64512 2 862808446 3 0
*NULL*
FRAG 2 -1 2293 0 2 6 2 S04440 nirvana 32768 0 862927577 2 0
*NULL*
FRAG 2 1 2335584 0 2 6 1 S04440 nirvana 32768 2 862927577 2 0
*NULL*
```

To confirm which is the primary copy (the copy to use for restores), use the -L option with bpimagelist, as follows:

```
UNIX: # bpimagelist -L -backupid fgolddust_0862806643 | grep
Primary
Primary Copy: 1
Windows: bpimagelist -L -backupid fgolddust_0862806643 | find
```

```
Primary  
Primary Copy: 1
```

- 4 Tell NetBackup to use the duplicated copy rather than the original.

Run the `bpimage -npc` command and option to change the primary copy. The new primary copy is used for restoring an image:

```
# bpchangeprimary -copy 2 -id fgoldddust_0862806643 -cl fgoldddust
```

To confirm the new primary copy, use the following command:

```
UNIX: # bpimagelist -L -backupid fgoldddust_0862806643 | grep  
Primary  
Primary Copy: 2  
Windows: bpimagelist -L -backupid fgoldddust_0862806643 | find  
"Primary"  
Primary Copy: 2
```

- 5 Freeze the duplicated copy to ensure restore.

Use the command `bpmedia -freeze` to prevent NetBackup from expiring the images on the media and to ensure that the media is assigned in Media Manager. You should also use the media host for this image that was printed by `bpimmedia` in step 2. This is required when the host is different than the computer on which you are running this command.

```
bpmedia -freeze -m S04440 -host nirvana
```

- 6 Recall media from the vault.

Recall the appropriate volume from off-site storage.

To determine the media ID, slot number, or container ID of the tape to recall, use the `vmquery` command, located in the following directory:

- UNIX

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

- Windows

```
install_path\volmgr\bin.
```

In the following example, the slot number (S278) is listed in the vault slot field:

```
vmquery -m S04440  
=====
```

```

media ID:                S04440
media type:              8MM cartridge tape (4)
barcode:                 S04440
media description:       Added by Media Manager
volume pool:             Vaulted_CustomerDB (2)
robot type:              NONE - Not Robotic (0)
volume group:            DB_offsite_volumes
vault name:              Customer_DB_Vault
vault sent date:         ---
vault return date:       ---
vault slot:              S278
vault session id:        1
created:                 Tue Sep 3 10:08:32 2000
assigned:                 Tue May 6 00:11:45 2001
last mounted:            Tue May 6 11:34:25 2001
first mount:             Tue Sep 3 18:20:48 2000
expiration date:         ---
number of mounts:        21
max mounts allowed:      ---
=====

```

7 Inject recalled media back into the robot.

When the tape is returned from the off-site vendor, inject it into the appropriate robotic library. First, insert the tape into the robot media access port. Then, from the NetBackup Administration Console, choose **Media and Devices Management**. Choose the **Inventory Robot...** option. Select the **Empty Media Access Port Prior to Update** checkbox.

You can also perform this function using the `vltinject` command.

8 Perform a normal restore operation.

The restore should read from the new primary copy. The restore log should show a mount request for the duplicate media.

9 Unfreeze media that is used for duplicates.

After the restore is successful, unfreeze the duplicate media to allow the normal expiration process. If you want to send the tape off-site again, either remove it from the robot or leave it in the robot as the primary copy. Cohesity recommends that you suspend the media so that no images are written to it.

```
bpmmedia -unfreeze -m S04440 -host nirvana
```

10 Create new duplicate images.

Optionally, you can create new duplicate images for transfer to your off-site vault vendor.

See “Reduplicating a bad or missing duplicate tape” on page 215.

11 Modify the NetBackup catalog for a large number of images.

In a disaster recovery situation in which a large number of images need their primary copy modified, run the `bpchangeprimary` command. This command changes the primary copy of all the backup images in the off-site volume pool for which the media was returned from the off-site vault.

Creating originals or copies concurrently

This chapter includes the following topics:

- About concurrent copies
- About the continue or fail for concurrent copies
- Creating multiple original images concurrently
- About creating duplicate images concurrently

About concurrent copies

You can create up to four copies of the same backup image concurrently. Those copies are created concurrently by the Inline Tape Copy feature. If the images are created during a NetBackup policy job, all are considered original images. If the images are duplicated by using the NetBackup Administration Console Catalog node or during a Vault job, they are considered duplicate images.

You must configure NetBackup to allow a sufficient number of copies in the **Maximum Backup Copies** field for the NetBackup master server. (Configured in **NetBackup Management > Host Properties > Master Server > *server_name* > Global NetBackup Attributes**.) By default, the value is two.

All storage units must be connected to the same media server. Also, you must configure the storage unit to allow a sufficient number of concurrent jobs to support the concurrent copies (**Maximum Concurrent Jobs** or **Maximum Concurrent Drives Used for Backup** setting).

You can write images concurrently to the following storage units:

- Media Manager storage units. If the Media Manager storage unit has more than one drive, the source storage unit and destination storage unit can be the same.
- Disk storage units.
- Disk staging storage units.
- Network Data Management Protocol (NDMP) storage units only during Vault duplication. Also only one copy is allowed per duplication rule (NDMP is not supported during original backup). If the NDMP storage unit has more than one drive, the source storage unit and destination storage unit can be the same. Although specifying an NDMP storage unit restricts the number of copies to one, you can use multiple duplication rules to specify other storage units for images that are created by other media servers. For example, you can use one duplication rule to read an image from one media server and write a copy to an NDMP storage unit. Then you can use another duplication rule to read an image from a different media server and write copies to other storage units. (To specify multiple duplication rules in a Vault profile, select **Advanced Configuration** on the **Profile** dialog box **Duplication** tab.) Because of potential NDMP performance limitations, it is recommended that you duplicate between the disk drives and tape drives that are directly attached to the same NDMP host.

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

You cannot create images concurrently using the following:

- Storage unit groups
- Quarter-inch cartridge (QIC) devices
- Third-party copies

About the continue or fail for concurrent copies

When making multiple copies of images concurrently, you can choose how an operation behaves if one of the copies fails. Your choice also can determine whether copies are ejected, depending on the success or failure of the copy operation. It is possible for a duplication operation to succeed but no ejection to occur.

In NetBackup, your continue or fail choice affects only the current image copy. In Vault, your choice affects all copies of that image.

By default, the option is configured to **Continue in NetBackup** and to **Fail All Copies in Vault**.

- See “About continue copies” on page 163.

- See “About fail all copies” on page 163.

About continue copies

If you choose **Continue for all copies**, the concurrent copy job is considered successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted. It is probable that at least one copy will succeed, but it may not be the copy that is assigned to the off-site volume pool for ejection.

To ensure that media are ejected even if a concurrent copy operation fails during a NetBackup policy backup, do one of the following:

- Configure a Vault profile to duplicate the image, assign the copy to the off-site volume pool, and select **Fail All Copies**. If the copy fails during the original NetBackup backup job, the Vault profile subsequently duplicates it. If the copy succeeds during the original backup job, the Vault profile does not duplicate it. Either way, a copy is ejected for transfer off site.
- Monitor the Activity Monitor for a failed status for the copy that is assigned to the off-site volume pool. If that copy fails, duplicate that image and assign it to the off-site volume pool so it is ejected. You can use the Administration Console Catalog node or the `bpduplicate` command to duplicate the copy.

About fail all copies

The behavior of the fail option and the default settings depend on whether the concurrent copies operation was configured in Vault or in NetBackup:

- In Vault, if you choose **Fail All Copies**, all copies of that image fail, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault again tries to duplicate the image if the following conditions are true:
 - The image is selected.
 - The Vault profile did not eject the primary backup.
- In NetBackup, if you choose **Fail All Copies**, the entire backup job fails and no copies are made. In this case, normal NetBackup behavior ensures that a successful backup for this policy eventually occurs. That is, NetBackup automatically retries the backup if time permits and, the next time the backup window for the policy opens, NetBackup again tries to run the backup (regardless of the frequency of the schedule). NetBackup retries until the backup succeeds, although one or more backup windows may pass before the backup is successful.

Creating multiple original images concurrently

In a NetBackup policy job, you can create multiple original backup images concurrently. Vaulting original images has many benefits, including easier configuration of Vault, fewer chances for resource contention, and possibly fewer drives required.

To create multiple backup images concurrently

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 In the middle pane, double-click an existing policy.
- 3 Select the **Schedules** tab.
- 4 Double-click an existing schedule or click **New** to create a new schedule.
- 5 In the **Schedule Attributes** tab of the **Schedule** dialog, select **Multiple Copies** and then click **Configure**.
- 6 In the **Configure Multiple Copies** dialog box, specify the number of copies to be created simultaneously.

The maximum is four. Copy 1 is the primary copy. If copy 1 fails, the first successful copy is the primary copy.
- 7 Specify the priority of the duplication job for each copy, from 0 to 99999.

A larger number is higher priority. All copies are duplicated at the same priority.
- 8 Specify the storage unit where each copy is stored.

If a Media Manager storage unit has more than one drive, it can be used for both the source and the destination. Network Data Management Protocol (NDMP) storage units are not supported when creating multiple copies during a NetBackup policy job.
- 9 Specify the volume pool to which each copy is assigned.
- 10 Select the retention level for each copy.

If you select **No Change**, the expiration date is the same for the duplicate copies and original copies.

If you select a different retention period, the expiration date of the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 200x, and its retention period is one week, the new copy's expiration date is November 21, 200x.
- 11 Select whether to **Continue** the other copies if a copy operation fails or to **Fail All Copies**.

12 Specify who should own the media onto which NetBackup writes the images:

Any	NetBackup chooses the media owner, either a media server or server group.
None	Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	All media server groups that are configured in your NetBackup environment appear in the drop-down list. Specifying a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written.

13 Click **OK**.

14 Configure other schedule criteria as appropriate.

About creating duplicate images concurrently

You can create multiple duplicate backup images concurrently either by using the NetBackup Catalog node or by configuring the **Duplication** tab of a Vault profile. Duplication is not always possible, so you must understand when you can use duplication in NetBackup.

See “Creating concurrent copies through the catalog node” on page 167.

See “Creating concurrent copies using the basic duplication tab” on page 168.

See “Creating concurrent multiple copies using the advanced duplication options” on page 171.

Table 7-1 describes when duplication is and is not possible in NetBackup.

Table 7-1 Possible circumstances for duplicating a Vault profile

Possible to duplicate backups	Not possible to duplicate backups
<ul style="list-style-type: none">■ From one storage unit to another.■ From one media density to another.■ From one server to another.■ From multiplex to nonmultiplex format.■ From multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. This process is done with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)	<ul style="list-style-type: none">■ While the backup is being created (unless you create multiple backup images concurrently during the backup job).■ While any other backup image is being written to a tape that contains the source primary backup.■ When the primary backup image is not available.■ By using the NetBackup scheduler to schedule duplications automatically (unless you use a Vault policy to schedule duplication) of the NetBackup catalogs.■ When it is a multiplexed image of the following:<ul style="list-style-type: none">■ Auspex FastBackup■ NDMP backup■ Backups to or from disk type storage units■ Nonmultiplexed backups

If you do multiplexed duplication, be aware of the following:

- When you duplicate multiplexed SQL-BackTrack backups with multiplex mode enabled, it is necessary to duplicate all of the backups in the multiplexed group. This ensures that the fragment order and size are maintained in the duplicate. Otherwise, it is possible that restores from the duplicated backups will not work. A multiplexed group is a set of backups that were multiplexed together during a single multiplexing session.
- When you duplicate multiplexed backups, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups have a multiplexing factor that is no greater than that used during the original backup.
- If all backups in a multiplexed group are duplicated to a storage unit that has the same characteristics as the one where the original backup was created, the duplicated group is identical, with the following exceptions:
 - If end of media (EOM) is encountered on either the source media or destination media.

- If any of the fragments in the source backups are zero length (which can occur if many multiplexed backups start at the same time), during duplication these zero-length fragments are removed.

Creating concurrent copies through the catalog node

Use the following procedure to create concurrent copies of backup images manually through the NetBackup Administration Console Catalog node.

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

To create concurrent images through the catalog node

- 1** In the NetBackup Administration Console, expand **Master Server > NetBackup Management > Catalog**.
- 2** In the **Action** field, select **Duplicate**.
- 3** Select the search criteria for the image you want to duplicate, and then click **Search Now**.
- 4** Right-click the image you want to duplicate and select **Duplicate** from the shortcut menu.
- 5** In the **Setup Duplication Variables** dialog, specify the number of copies to create.

If enough drives are available, the copies are created simultaneously. Otherwise, the system may require operator intervention, for example, if four copies are to be created and there are only two drives.

- 6** Specify the priority of the duplication job for each copy, from 0 to 99999. A larger number is higher priority. All copies are duplicated at the same priority.
- 7** If you want one of the duplicated copies to become the primary copy, check the appropriate box.

See "Duplication tab configuration options" on page 87.

- 8** Specify the storage unit where each copy is stored.
- 9** Specify the volume pool to which each copy is assigned.

NetBackup does not verify in advance that the media ID selected for the duplicate copy is not the same as the media ID of the volume that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure that a different volume is used.

- 10 To change the retention level for the copy, select one of the retention level options.

If **No Change** is selected for the retention period, the expiration date is the same for the duplicate copy and source copy.

If you specify a numeric retention level, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 200x and its retention period is one week, the new copy's expiration date is November 21, 200x.
- 11 Specify whether the remaining copies should continue or fail if the specified copy fails.
- 12 Specify who should own the media onto which NetBackup writes the images.

Any	NetBackup chooses the media owner, either a media server or server group.
None	Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	All media server groups that are configured in your NetBackup environment appear in the drop-down list. Specifying a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written.
- 13 If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, select **Preserve Multiplexing**.
- 14 Click **OK** to start duplicating.
- 15 Click the **Results** tab, and then select the duplication job just created to view the job results.

Creating concurrent copies using the basic duplication tab

You can create multiple duplicate images concurrently in Vault by selecting either **Multiple Copies** on the basic **Duplication** tab or **Advanced Configuration** on the basic **Duplication** tab, which displays the advanced duplication criteria.

You can use the following instructions to create multiple copies concurrently from the basic **Duplication** tab.

For instructions on how to configure duplication in Vault, see the following:

See "Duplication tab" on page 83.

See "Duplication tab configuration options" on page 87.

To create concurrent multiple copies using the basic duplication tab

- 1** Indicate whether the images you want to duplicate reside on disk storage units only or on disk or media storage units.
- 2** Enter the number of drives to use for reading backup images for duplication.

When you enter a number of read drives, the same number is entered into the destination **Write Drives** field. You must have an equivalent number of read and write drives available.
- 3** To use a media server that is different from the server that wrote the images, check **Alternate Read Server** and select the media server to use. (Alternate read servers apply to NetBackup Enterprise Server only.)

If robots (or drives) are shared by more than one media server, you can specify a different media server to read the original backups than the media server that wrote the backups.
- 4** Select **Multiple Copies**, and then click **Configure**.
- 5** In the **Multiple Copies** dialog, select the number of copies to create.

The number of copies you choose cannot exceed the number of copies that are specified in the **Maximum Backup Copies** field for the NetBackup master server. (Configured in **NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes**.) By default, the value is two, which means one original backup and one copy.
- 6** If you want one of the copies to be the primary copy, select which copy is to be primary.
- 7** Specify the storage unit to be used for the duplication. When you specify the storage unit, the following applies:
 - If the Media Manager or Network Data Management Protocol (NDMP) storage unit has more than one drive, the source storage unit and destination storage unit can be the same.
 - NDMP storage units are supported only when one copy is created.
 - All storage units must be connected to the same media server.
- 8** Specify a volume pool for each copy.
- 9** Specify the retention level for each copy.

See “Assigning multiple retentions with one profile” on page 141.

When the retention period expires, information about the expired backup is deleted from the NetBackup and Media Manager catalog, the volume is recalled from off-site storage, and the backup image is unavailable for a restore.

10 Indicate what action is to be taken if a copy fails.

In Vault, if you choose **Fail All Copies**, all copies of that image fails, independent of the success or failure of other image copy operations. The next time the Vault profile runs, Vault again tries to duplicate the image if the following conditions are true:

- The image is selected.
- The Vault profile did not eject the primary backup.

By default, the option is configured to **Fail All Copies** in Vault.

If you choose **Continue** for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted. It is probable that at least one copy will succeed, but it may not be the copy that is assigned to the off-site volume pool.

11 Specify who should own the media onto which you are duplicating images:

Any	NetBackup chooses the media owner.
None	Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	All media server groups that are configured in your NetBackup environment appear in the drop-down list. Specifying a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written.

12 Click **OK** to return to the basic **Duplication** tab.

13 Specify the priority of the Vault duplication jobs, from 0 to 99999. A larger number is higher priority. All duplication jobs for the profile run at the same priority.

14 To preserve multiplexing, select **Preserve Multiplexing**.

See “Duplication tab configuration options” on page 87.

15 Select **Duplicate Smaller Images First (applies only to disk backup images)** to duplicate images in smallest to largest order.

16 Check **Expire Original Disk Backup Images...** and then enter the number of hours after this Vault session completes to expire the disk images.

If the duplication of a disk image fails, the disk image does not expire.

17 After you complete the dialog box, click **OK**.

Creating concurrent multiple copies using the advanced duplication options

You can use the following instructions to create multiple copies concurrently from the advanced configuration criteria of the Vault profile **Duplication** tab.

For instructions on how to configure duplication in Vault, see the following:

See “Duplication tab” on page 83.

See “Duplication tab configuration options” on page 87.

To create concurrent multiple copies using the advanced configuration options

- 1** On the **Duplication** tab, select **Advanced Configuration**.
- 2** To use a server that is different from the server that wrote the images, select **Alternate Read Server**. (Alternate read servers apply to NetBackup Enterprise Server only.)

If robots (or drives) are shared by more than one media server, you can designate a different media server to read the original backups than the media server that wrote the backups.

If you select **Alternate Read Server**, an **Alternate Read Server** column heading appears in the SOURCE area.

- 3** To add a destination media server and duplication rules for that server, click **New**.

If you selected **Alternate Read Server** on the **Duplication** tab, the **Duplication Rule** dialog box has fields for both source **Media Server** and **Alternate Read Server**. If you did not select **Alternate Read Server**, only a source **Backup Server** field appears.

- 4** Select the source **Backup Server** or, if you selected **Alternate Read Server** on the **Duplication** tab, select the source **Media Server**.
- 5** If you selected **Alternate Read Server** on the **Duplication** tab, select an **Alternate Read Server**. (Alternate read servers apply to NetBackup Enterprise Server only.)

The source media server and alternate read server may be the same.

- 6** Select the number of copies to create.

You can create up to four or the number of copies that are specified in the **Maximum Backup Copies** field for the NetBackup master server (if less than four). (**Configured in NetBackup Management > Host Properties > Master Server > server_name > Global NetBackup Attributes**.) By default, the value is two: one original backup and one copy.

- 7

If you want one of the copies to be the primary copy, select which copy is to be primary.
- 8

Specify the storage unit to be used for the duplication.

In addition, all storage units must be connected to the same media server.
- 9

Specify a volume pool for each copy.
- 10

Specify the retention level for each copy.

See “Assigning multiple retentions with one profile” on page 141.

When the retention period expires, information about the expired backup is deleted from the NetBackup and Media Manager catalog, the volume is recalled from off-site storage, and the backup image is unavailable for a restore.
- 11

Indicate what action is to be taken if a copy fails.

In Vault, if you choose **Fail All Copies**, all copies of that image fails, independent of the success or failure of other image copy operations. The next time the vault profile runs, Vault again tries to duplicate the image if the following conditions are true:
 - The image is selected.
 - The Vault profile did not eject the primary backup.By default, the option is configured to **Fail All Copies** in Vault.

If you choose **Continue** for all copies, Vault considers the duplication job successful if any of the copies succeed. However, it is possible that a copy of the image may never get vaulted. It is probable that at least one copy will succeed, but it may not be the copy that is assigned to the off-site volume pool.
- 12

Specify who should own the media onto which you are duplicating images:

Any	NetBackup chooses the media owner.
None	Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	All media server groups that are configured in your NetBackup environment appear in the drop-down list. Specifying a media server group allows only those media servers in the group to write to the media on which backup images for this policy are written.
- 13

Click **OK** to return to the **Duplication** tab.

- 14** Specify the priority of the Vault duplication jobs, from 0 to 99999. A larger number is higher priority. All duplication jobs for the profile run at the same priority.
- 15** Indicate whether you want to **Preserve Multiplexing**.
See “Duplication tab configuration options” on page 87.
- 16** Select **Duplicate Smaller Images First (applies only to disk backup images)** to duplicate images in smallest to largest order.
- 17** Check **Expire Original Disk Backup Images...** and then enter the number of hours after this Vault session completes to expire the disk images.
Use this option to free up space on the disk for subsequent backup images. Be sure that you allow enough time for the duplication operation to complete.
If the duplication of a disk image fails, the disk image does not expire.
- 18** Click **OK**.

Reporting

This chapter includes the following topics:

- About reports
- About generating reports
- About consolidating reports
- Viewing Vault reports
- Vault report types

About reports

The reports for each profile are configured on the **Reports** tab in the **Profile** dialog box. When you configure a Vault profile, you specify which reports should be generated, when they should be generated, and how and to whom they should be distributed.

After reports are generated and distributed, you can view and print them until the Vault logs for that session are deleted.

To view sample Vault reports and log files for this release, refer to the following file on the Cohesity Support Web site.

About generating reports

If the reports for a profile are configured as immediate, the reports are generated when the profile runs.

If the reports for a profile are deferred, use one of the following methods to generate the reports after the profile runs:

- Select the **Deferred Eject** option in the Administration Console and then selecting **Generate Reports After Eject**.
- Use the **Vault Operator Menu** interface.
- Use the `vlteject` command.
- Use a Vault policy that runs the `vlteject` command.

When you generate reports, the reports that are selected on the **Profile** dialog box **Reports** tab are generated and distributed to the destinations specified.

Reports can be generated for one session or for multiple sessions. Generating reports and ejecting media from more than one vault session is known as consolidating your reports and ejections. For example, you may duplicate images daily but eject media and generate reports only at the end of the week.

- See “Generating reports by using the Vault operator menu” on page 176.
- See “Generating reports by using the `vlteject` command” on page 177.
- See “Creating a Vault policy to generate reports” on page 177.
- See “Ejecting media by using the NetBackup Administration Console” on page 125.
- See “Reports tab (Profile dialog box)” on page 107.
- See “Reports that depend on eject” on page 109.

Generating reports by using the Vault operator menu

You can use the Vault Operator Menu to generate reports.

To generate reports by using the Vault operator menu

- 1 Start the Vault Operator Menu by running the `vltopmenu` command.
- 2 If necessary, select a profile.
- 3 Select one of the following options:
 - **Run Reports for This Session**
 - **Run Individual Reports**
 - **Consolidate All Reports**
 - **Consolidate All Reports and Ejects**
Consolidating reports and ejects also ejects media.
- 4 Continue as prompted by the Vault Operator Menu.

Generating reports by using the `vlteject` command

You can use the `vlteject` command with the `-report` option to generate reports from the command line. The following is the syntax for the command that generates all reports that are not yet generated:

```
vlteject -report
```

You also can specify a robot, vault, profile, or session for which to generate reports.

If the corresponding eject process has completed, pending reports from the selected sessions are generated and distributed. The reports are not generated again if `vlteject` is run again. If the ejection has not been completed, the subset of reports that do not depend on completion of eject are generated. These reports are generated again if `vlteject -report` is run again after eject has been completed.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name]  
[-profile robot_no/vault_name/profile_name]  
[-robot robot_no] [-vault vault_name [-sessionid id]]  
[-auto y|n] [-eject_delay seconds] [-legacy]
```

The `vlteject` command resides in the following directory:

- UNIX: `/usr/opensv/netbackup/bin`
- Windows: `install_path\NetBackup\bin`

See the *NetBackup Commands Reference Guide*.

To generate reports by using the `vlteject` command

- 1 In a terminal window or command window, change to the directory in which the `vlteject` command resides.
- 2 Run the command, using the appropriate options and parameters.

Note: In case of non-administrative users using NBAC, a specially privileged user in **Vault Operator** group is authorized to run the command `vlteject -report`. This privileged user only can generate the reports for `detail.log` and `session.log` files, on demand.

Creating a Vault policy to generate reports

You can use a Vault policy to generate reports for the Vault sessions that have completed already and for which reports have not been generated. On the **Backup**

Selections tab, in the Vault policy, specify Vault as the policy type, do not specify clients, and specify the `vlteject` command with the `-report` option.

You also can specify a robot, vault, profile, or session for which to generate reports.

If the corresponding eject process has completed, pending reports from the selected sessions are generated and distributed. The reports are not generated again if `vlteject` is run again.

If the eject operation has not been completed, the subset of reports that do not depend on completion of eject are generated and distributed. These reports are generated again if `vlteject` is run again.

The following is the format of the `vlteject` command:

```
vlteject -eject -report -preview [-profile profile_name]
[-profile robot_no/vault_name/profile_name]
[-robot robot_no] [-vault vault_name [-sessionid id]]
[-auto y|n] [-eject_delay seconds] [-legacy]
```

The `vlteject` command resides in the following directory:

- **UNIX**

```
/usr/openv/netbackup/bin
```

- **Windows**

```
install_path\NetBackup\bin
```

See the *NetBackup Commands Reference Guide*.

To create a Vault policy to generate reports

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box.
- 4 On the **Attributes** tab, select **Vault** as the policy type.
- 5 On the **Schedules** tab, click **New** to create a new schedule.
The type of backup defaults to Automatic Vault.
- 6 Complete the schedule.
- 7 Bypass the **Client** tab (clients are not specified for Vault jobs).

- 8 On the **Backup Selections** tab, enter the `vlteject` command with the `-report` option and any other appropriate options.
- 9 Click **OK**.

About consolidating reports

You can generate reports and eject media from more than one vault session, which is known as consolidating your reports and ejections. For example, you may duplicate images daily but eject media and generate reports only at the end of the week. To do so, specify deferred reports on the **Reports** tab and deferred eject on the **Eject** tab for each profile for which you want to consolidate reports. Then, eject the media and generate the reports.

Note: If you consolidate reports, you should also consolidate ejects.

See “About ejecting media” on page 124.

When you generate the reports, you select the robot, vault, or profile sessions for which reports were deferred (that is, for the reports that are pending).

You can consolidate the following:

- All sessions for a profile
- All sessions for a vault
- All sessions for a robot
- All sessions for all vaults

A consolidated report includes information from all sessions in which that report is specified in the profile. For example, a consolidated Picking List for Robot includes the appropriate media from all sessions whose profile has Picking List for Robot selected on the **Reports** tab.

Table 8-1 defines the elements of a consolidated report.

Table 8-1 Consolidated reports elements

Report element	Description
Report Header	Includes the following: <ul style="list-style-type: none">■ Identifies the report as a consolidated report.■ The robots, vaults, or profiles that are included in the report.■ The sessions that are included in the report.

Table 8-1 Consolidated reports elements *(continued)*

Report element	Description
Report Body	Shows the media from all sessions that are included in the consolidation in which the report is selected on the profile Reports tab. If media applies to more than one session, only the information from the most recent session is included. Similarly for containers, only the container information from the most recent session appears.
Summary	Shows the same information as in a nonconsolidated report.

In the Recovery Report, the earliest date range among the consolidated sessions is the end date. The Recovery Report is generated from the current date and time to that end date. Time ranges specified in individual profiles are used to generate the reports. Each time range is calculated based on the start time of the individual sessions participating in the consolidated report.

If the ejection has not completed, the subset of reports that do not depend on completion of eject are generated. These reports are generated again if deferred reports are run again.

If you consolidate reports and also rename reports, use the same customized report title for all profiles whose reports are consolidated. The customized report title is printed on the report and appears in the email subject line if you email the reports.

Note: Reports cannot be consolidated between the vaults that use slots and the vaults that use containers.

About consolidated reports in previous Vault releases

Consolidated reports in previous releases of Vault concatenated the same reports from each session together. One advantage of that style is that you can consolidate reports from slot and container-based vaults, which you cannot do with the new style of consolidated report.

If you prefer the previous style of consolidated reports, use the `vlteject` command `-report` and `-legacy` options to consolidate the reports in that old style. If you currently do immediate reports, you have to change to deferred reports and either run the `vlteject` command manually or create a Vault policy to schedule the `vlteject -report -legacy` job. If you currently do deferred reports, you can add the `-legacy` option to the `vlteject` command you use to generate reports.

Viewing Vault reports

You can use the NetBackup Administration Console to view and print reports for Vault sessions for which reports are already generated. You can only view reports if the session directory for that vault still exists. Only some of the reports are valid. For example, the picking list reports are only valid on the date they were generated.

To view Vault reports

- 1 Select **NetBackup Management > Reports > Vault Reports**.

- 2 Select one of the Vault reports or report types.

When you select a report or a report type, the **Reports** window appears. The **Reports** window includes a **Report Settings** area and a report contents window.

- 3 Enter or select the appropriate values for the report you want to generate.

Usually, you must specify a profile and a session ID. You also may have to specify a date range or time period.

- 4 Click **Run Report**.

- 5 To print the report, click **File > Print**.

Vault report types

This topic introduces and describes a variety of Vault reports as well as Vault report types.

Reports for media going off site

The reports for media going off-site show media that have ejected from the robot and are transported off-site. They vary in the amount of detail that is included in each report. Media on these reports are from the `eject.list` file for the session or, from consolidated reports, the combined `eject.list` files of all unreported vault jobs. An unreported vault job is one that ejects media but does not generate reports.

Picking List for Robot report

The **Picking List for Robot** report shows the volumes that are ejected from the robot that should be transported off-site. This report is sorted by media ID and should be used by the operations staff as a checklist for media that has ejected from the robots. You can save the report for tracking purposes, or you can reprint it as long as the session directory still exists.

Column descriptions in the **Picking List for Robot** report are as follows:

IMAGES	The number of images on the volume. For Vault catalog backup media from releases earlier than NetBackup 6.0, the column displays zero.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EJECTED	Yes or No.
EXPIRATION	Date when the images on the volume expire. For Vault catalog backup media, displays the date that is calculated as a return date during the volume assignment.
MBYTES	The size in megabytes of images on the volume. For Vault catalog backup media from releases earlier than NetBackup 6.0, the field is empty.
CATEGORY	The type of media: <ul style="list-style-type: none"> ■ NBU. NetBackup media that contains backup images. ■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases. ■ Old NBU CTLG. Catalog backup media from NetBackup releases earlier than 6.0 ■ Add-on. Media not managed by NetBackup.
MEDIA	The ID of the media.
ROBOT	The number of the robot from which the media was ejected. (Consolidated report only.)
SLOT ID	The ID of the slot in which the volume will reside at the off-site vault. (Slot vaulting only.)

Distribution List for Vault report

The **Distribution List for Vault** report shows the volumes that have ejected from the robot and are transported off-site. This report is sorted by off-site slot number or container number and should accompany the media that is destined for the off-site vault. The vault vendor should use this report to verify that all the volumes listed were actually received.

Column descriptions in the **Distribution List for Vault** report are as follows:

IMAGES	The number of images on the volume.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EJECTED	Yes or No.

EXPIRATION	Date when the images on the volume expire. For Vault catalog backup media, displays the date that is calculated as a return date during the volume assignment.
MBYTES	The size in megabytes of images on the volume. For Vault catalog backup media from releases earlier than NetBackup 6.0, the field is empty.
CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media not managed by NetBackup.
MEDIA	The ID of the media.
RETURN DATE	The date the container should be returned from the off-site vault.
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)
IMAGES	The number of images on the volume.

Detailed Distribution List for Vault report

The **Detailed Distribution List for Vault** report shows the volumes that have ejected from the robot and are transported off-site. This report is similar to the **Picking List for Robot** and **Distribution List for Vault** reports except that it includes detailed information about the images on each volume. Because backup jobs can span volumes, fragments of a backup image may appear on more than one volume. If two or more fragments of the same image are on one volume, they are reported on one line rather than on a separate line for each fragment. That is, each image is listed once for every media its fragments reside on.

This report is useful at a disaster recovery site. Cohesity recommends that you send this report off-site.

Column descriptions in the **Detailed Distribution List for Vault** report are as follows:

BACKUP ID	Identifier that NetBackup assigns when it performs the backup.
CLIENT	Name of the client that was backed up.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)

EJECTED	Yes or No.
EXPIRATION	Date when the images on the volume expire. For Vault catalog backup media, displays the date that is calculated as a return date during the volume assignment.
IMAGES	The number of images on the volume.
KBYTES	The size in kilobytes of the complete backup image. The size of the complete image is listed even if the image is a fragment. For Vault catalog backup volumes, the field is empty.
MBYTES	The size in megabytes of all images on the volume. For Vault catalog backup media from releases earlier than NetBackup 6.0, the field is empty.
CATEGORY	The type of media: <ul style="list-style-type: none"> ■ NBU. NetBackup media that contains backup images. ■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases. ■ Add-on. Media not managed by NetBackup.
MEDIA	The ID of the media.
PARTIAL	Partial images on the volume. The field displays: <ul style="list-style-type: none"> ■ COMPLETE if all fragments reside on that volume. ■ PARTIAL if some fragments reside on other volumes. ■ EXTRA if the images do not belong to the session.
POLICY	Name of the policy that was used to back up the client.
SCHEDULE	Name of the schedule that was used to back up the client.
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)
VAULT	The name of the vault to which the volume's profile belongs. (Consolidated report only.)
WRITTEN	The date the image was written.

Summary Distribution List for Vault report

This report is similar to the **Detailed Distribution List for Vault** report, except that the entry for each piece of media lists only a unique client, policy, schedule, and date. That is, if multiple backup jobs for a given client, policy and schedule (usually seen with RDBMS backups or SAP backups) are written to the same volume on

the same date, only one line of information prints on this report. The **Detailed Distribution List** would show each of these backup jobs as a separate entry, which can generate a long report. The **Summary Distribution List for Vault** report summarizes the information and presents it in a more compact form. This report is also useful for disaster recovery situations; we recommend that you send this report off-site.

Column descriptions in the **Summary Distribution List for Vault** report are as follows:

IMAGES	The number of images on the volume.
BACKUP TIME	When the backup occurred.
CLIENT	Name of the client that was backed up.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EJECTED	Yes or No.
EXPIRATION	Date when the images on the volume expire. For Vault catalog backup media, displays the date that is calculated as a return date during the volume assignment.
MBYTES	The size in megabytes of images on the volume. For Vault catalog backup media from releases earlier than NetBackup 6.0, the field is empty.
CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media not managed by NetBackup.
MEDIA	The ID of the media.
POLICY	Name of the policy that was used to back up the client.
SCHEDULE	Name of the schedule that was used to back up the client.
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)
VAULT	The name of the vault to which the volume's profile belongs. (Consolidated report only.)

Reports for media coming on-site

The reports for media coming on-site show the volumes that are requested back from the off-site vault. These reports can be generated before or after media has ejected for the current Vault session.

To appear in these reports, media must:

- Be in a Vault off-site volume group.
- Be in an eject volume pool for a profile in the current vault or in a scratch pool.
- Be in any vault if the media is in a scratch pool.
- Have a non-null return date which has passed when the report is generated (Catalog backup only).
- Have a vault container value that is non-null (container vaulting only).
- Be unassigned (NetBackup media only).

Picking List for Vault report

The **Picking List for Vault** report shows the volumes that are requested back from the off-site vault. This report should be sent off-site to the vault vendor.

Volumes are listed on this report because Vault determined that they are in an off-site volume group and that all images have expired. When Vault identifies these volumes, it changes the **Return Date** field for the media and adds the media ID and date requested to this report.

Expired media appear on this report only once: either on the date the media expire or the next time the report is generated (if the report is not generated on the date a volume expires). If media appear on the report but are not recalled, they appear on the **Lost Media** report.

A slot at the off-site vault from which an expired volume is recalled is available for use one day after the volume has physically returned to the robot.

If you use a scratch pool, this report may include volumes from other profiles or vaults that expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

Column descriptions in the **Picking List for Vault** report are as follows:

CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
DENSITY	Density of the volume.

LAST MOUNT	The date the volume was last mounted. (Session report only; does not appear for consolidated report.)
LAST SID	The session ID of the Vault session that recalled the media.
CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media that NetBackup does not manage.
MEDIA	The ID of the media.
REQUESTED	Date when the volume is requested back from the off-site vault.
RETURN DATE	The date the container should be returned from the off-site vault. (Container vaulting only.)
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)
VAULT	The name of the vault to which the volume's profile belongs. (Consolidated report only.)

Distribution List for Robot report

The **Distribution List for Robot** report shows the volumes to be requested back from the off-site vault. This report is identical to the **Picking List for Vault**, except that it includes the robot to which the media should be returned. Retain this report to use as a checklist for the media that are returned from the off-site vault.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

Column descriptions in the **Distribution List for Robot** report are as follows:

CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
DENSITY	Density of the volume.
LAST MOUNT	The date the volume was last mounted. (Session report only; does not appear for consolidated report.)
LAST SID	The session ID of the Vault session that recalled the media.

CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media that NetBackup does not manage.
MEDIA	The ID of the media.
REQUESTED	Date when the volume is requested back from the off-site vault.
RETURN DATE	The date the container should be returned from the off-site vault. (Container vaulting only.)
ROBOT	The number of the robot in which the media resided. (Consolidated report only.)
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)

Inventory reports

The inventory reports show the location of the media. These reports are not generated until the media have been ejected.

See “Vault Inventory report” on page 188.

See “Off-site Inventory report” on page 189.

See “All Media Inventory report” on page 190.

If you use the NetBackup Administration Console to display an inventory report, you must select a profile that ejects media. Also, select the most recent session for that profile so the most recent data is reported.

Vault Inventory report

The **Vault Inventory** (or **Inventory List for Vault**) report shows media that are off-site at the vault vendor and media being sent off-site (outbound media in transit to the vault).

To appear in this report, media must be:

- In the off-site volume group.
- In the current vault or in any vault if the media is in a scratch pool.
- In an eject volume pool for a profile in the current vault, in a scratch pool, or in the catalog volume pool.

Cohesity recommends that you send this report to your vault vendor so they can verify that they have the volumes that Vault indicates are at the vault vendor.

Column descriptions in the **Vault Inventory** report are as follows:

ASSIGNED	The date that NetBackup Media Manager assigned the volume.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EXPIRATION	Date when the images on the volume expire.
CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media that are not managed by NetBackup.
MEDIA	The ID of the media.
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)
VAULT	The name of the vault to which the volume's profile belongs. (Consolidated report only.)

Off-site Inventory report

The **Offsite Inventory** (or **Full Inventory List for Vault**) report includes the information in the Vault Inventory report and also includes any volumes that were requested back from the off-site vault vendor (that is, inbound media in transit). Usually, this report is not generated on a daily basis. Rather, the **Inventory List for Vault** report is sent to the vault vendor to perform verification.

If you use a scratch pool, this report may include volumes from other profiles or vaults that have expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

Column descriptions in the **Offsite Inventory** report are as follows:

ASSIGNED	The date when the volume was assigned by NetBackup Media Manager.
CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EXPIRATION	Date when the images on the volume expire.

CATEGORY	The type of media: <ul style="list-style-type: none">■ NBU. NetBackup media that contains backup images.■ New NBU CTLG. Catalog backup media from NetBackup 6.0 and later releases.■ Add-on. Media that are not managed by NetBackup.
MEDIA	The ID of the media.
REQUESTED	The date the volume was requested to be returned from the off-site vault.
SLOT ID	The ID of the slot in which the volume resides in the vault. (Slot vaulting only.)
VAULT	The name of the vault to which the volume's profile belongs. (Consolidated report only.)

All Media Inventory report

The **All Media Inventory** (or **Complete Inventory List for Vault**) report shows all volumes in the off-site volume pool.

To appear in this report, media must be:

- In the robotic volume group or off-site volume group.
- In the current vault or in any vault if the media is in a scratch pool.
- In an eject volume pool for a profile in the current vault or in a scratch pool.

If you use a scratch pool, this report can include volumes from other profiles or vaults that expired and moved back into the scratch pool even if the report is for a specific Vault profile or session.

Note: Volumes within the off-site volume pool must belong to either the off-site volume group or the robotic volume group or they do not appear on this report.

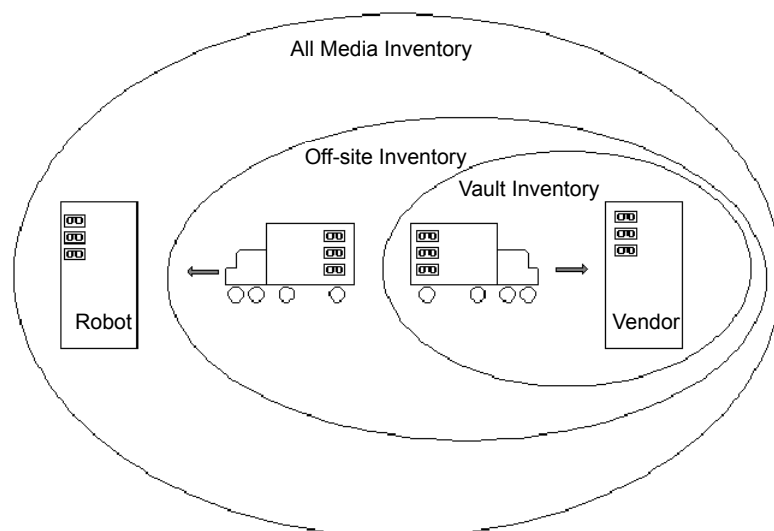
Column descriptions in the **All Media Inventory** report are as follows:

CONTAINER ID	The ID of the container in which the volume resides in the vault. (Container vaulting only.)
EXPIRATION	Date when the images on the volume expire.
LOCATION	Shows where the volume resides. For a session report, shows Robot or Vault. For a consolidated report, shows the robot number if the volume is in the robot or the vault name if the volume is off-site.

MEDIA	The ID of the media.
REQUESTED	The date the volume was requested to be returned from the off-site vault.
SID	The ID of the session that duplicated and/or ejected this volume.
SLOT ID	The ID of the slot in which the volume resides in the off-site vault. (Slot vaulting only.)

Figure 8-1 shows the different scopes of the reports:

Figure 8-1 Graphical representation of inventory reports scope



Container Inventory report

The **Container Inventory** report shows all the containers that are configured in your vaulting environment, the return date of each container, and the media that are in each container. Alternatively, you can specify a container ID to generate a report of the media in a specific container.

Generate this report only if you vault your media in containers. Reports do not show container information until after you add container and media IDs in Vault. Media are removed logically from a container when they are injected back into the robot.

Column descriptions in the **Container Inventory** report are as follows:

CONTAINER ID	The ID of the container in which the volume resides in the vault.
LAST SID	The last session ID of the profile that accessed this volume.
MEDIA ID	ID of the media that are in the container.
REQUESTED	Date when the container is requested back from the off-site vault.
RETURN DATE	The date the container should be returned from the off-site vault.
ROBOT	The robot from which the volumes were ejected.

Recovery Report for Vault report

The **Recovery Report for Vault** shows all policies that are defined on a NetBackup master server and all media that are required to restore the backups between a given set of dates. The report displays the date range to which the images on the media apply.

This report also includes the following:

- The three most recent Vault catalog backups in the vault’s off-site volume group. Only Vault catalog backups appear on this report. NetBackup catalog backups do not appear even if they are ejected and transferred off site.
- The information from the disaster recovery file that is generated by the catalog backup policy.

In a consolidated **Recovery** report, the earliest date range among the consolidated sessions is the end date. The **Recovery** report is generated from the current date and time to that end date. Time ranges specified in individual profiles are used to generate the reports. Each time range is calculated based on the start time of the individual sessions participating in the consolidated report.

Sending the **Recovery** report to the vault vendor on a regular basis helps with disaster recovery efforts. If a disaster destroys the master server, you cannot generate a **Recovery** report to determine which volumes to request from the vault vendor. Therefore, it is very important that the vault vendor has a copy of the Recovery report.

Field descriptions in the **Recovery Report for Vault** report are as follows:

POLICY	Name of the policy that was used to back up the client.
SCHEDULE	Name of the schedule that was used to back up the client.
CLIENT	Name of the client that was backed up. (Excluding catalog backup media.)

MEDIA	The ID of the media. (Vault catalog backup media only.)
MEDIA ID	The ID of the media. (Excluding catalog backup media.)
WRITTEN	The date the catalog backup was written to the volume. (Vault catalog backup media only.)
VAULT	The off-site vault at which the media are stored.
SLOT/CONTAINER	The ID of the slot or container in which the volume resides in the off-site vault.

Lost Media report

The **Lost Media** report lists expired media that has not been returned from the off-site vault vendor.

Media that appear on the Lost Media report have the following:

- An assigned date of 0 or are frozen.
- A robot number value of null.
- A vault sent date value that is non-null.
- A non-null return date.

Media can get lost for various reasons, as follows:

- A volume appears on the **Picking List for Vault** only once. If a volume from that report is missed and is not returned to the robot, it never again is listed for recall.
- Frozen backup media never expires. Media that does not expire does not appear on the **Picking List for Vault** and is not recalled from the vault.
- You change or rename your off-site volume groups and pools. For example, if you begin using a new media type, you have to use a new volume pool name. Cohesity recommends that you do not change or rename group or pool names.

You must generate the **Lost Media** report; it is not generated when you eject media. You do not have to configure your profiles for the **Lost Media** report. Usually, the media that are included in the **Lost Media** report should be returned from off-site and injected back into the appropriate vault in the robot.

A good practice is to run the **Lost Media** report periodically, such as weekly or monthly (depending on your operations).

Column descriptions in the **Lost Media** report are as follows:

DENSITY	Density of the volume.
---------	------------------------

LAST MOUNT	The date the volume was last mounted.
MEDIA ID	The ID of the media.
REQUESTED	Date when the volume is requested back from the off-site vault.
VAULT	The vault to which the volume belongs.
VOLUME GROUP	The volume group to which the volume is assigned.

Non-vaulted Images Exception report

The **Non-vaulted Images Exception** report shows images and media that were not vaulted when the Vault session was run. When a **Non-vaulted Images Exception** report for a given session is generated, the current status of images (in the preview.list of the session) and the media they reside on is checked for those that were not vaulted. This report lists the images (from the preview.list file) and the media that match the Choose backups criteria, and were not vaulted at the time the report was generated.

If you generate the report after a session runs, images that expired since the session ran are not on the report even though they were not vaulted. Therefore, to use the **Non-vaulted Images** report effectively, generate it when a session runs and save it so you can refer to it later.

You can save the report by specifying a directory to save the report to on the **Reports** tab (**Profile** dialog box).

See “Reports tab (Profile dialog box)” on page 107.

Reports also are saved in the session directory.

Column descriptions in the **Non-vaulted Images** report are as follows:

ASSIGNED	The date when NetBackup Media Manager assigned the volume.
BACKUP ID	Identifier that NetBackup assigns when it performs the backup.
CREATED	The date the volume was created (original backup or duplicated).
EXPIRATION	Date when the images on the volume expire.
MEDIA	The ID of the media.
POLICY	Name of the policy that was used to back up the client.
ROBOT	The number of the Vault robot in which the volume resides. (Consolidated report only.)
SCHEDULE	Name of the schedule that was used to back up the client.

SLOT ID	The ID of the slot in which the volume resides in the robot. (Slot vaulting only.) (Session report only; does not appear for consolidated report.)
VOLUME GROUP	The group to which the volume is assigned.
VOLUME POOL	The pool to which the volume is assigned.

About the Iron Mountain FTP file

If Iron Mountain is your vault vendor, you can configure Vault to produce an **Iron Mountain Electronic Format** report, which is a file that can include the following reports:

- Picking List for Vault (the "P" section)
- Distribution List for Vault (the "D" section)
- Vault Inventory report (if you are vaulting in slots) (the "I" section)
- Container Inventory report (if you are vaulting containers) (the "C" section)
- Recovery report (the "R" section)

The reports that are included in the file depend on your selections on the **Reports** tab of the **Profile** dialog box. You must select a report so that it appears in the Iron Mountain report file.

The report is in a format that Iron Mountain's automated vaulting mechanism can read and contains the information that they require. You can use the file transfer protocol (FTP) to send the report file to Iron Mountain, and they use it to update their vaulting mechanism automatically.

Before you send the report to Iron Mountain, verify that the volumes that were ejected match the **Distribution List for Vault**. Contact Iron Mountain to determine where and when to send the report.

Administering Vault

This chapter includes the following topics:

- About setting up email
- About administering access to Vault
- About printing Vault and profile information
- Copying a profile
- About moving a vault to a different robot
- About changing volume pools and groups
- About NetBackup Vault session files
- Operational issue with disk-only option on Duplication tab
- Operational issues with the scope of the source volume group

About setting up email

Depending on your computing environment, you may have to configure NetBackup or your computing environment so that notification email from NetBackup functions properly.

On UNIX systems, NetBackup uses the `sendmail` mail transfer agent to send email. If `sendmail` is not installed, you must install it and configure your environment so it functions correctly.

On Windows systems, NetBackup uses the `nbmail.cmd` script (in `install_path\VERITAS\NetBackup\bin\goodies`) on the NetBackup master server to send email. For email notifications, NetBackup passes the email address, subject, and message to the script. NetBackup then uses the mail program that is specified in the script to send email. For instructions on configuring the script, see the

comments in the `nbmail.cmd` script. Default NetBackup behavior: `nbmail.cmd` does not send email.

Note: If you use the Blat email client to deliver email on Windows systems, include the `-mime` option on the `blat` command in the `nbmail.cmd` script. That ensures the Vault reports are mailed correctly.

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

About administering access to Vault

NetBackup provides two mutually exclusive methods for controlling user access:

- **Access Management.** Access Management lets you control access to NetBackup by defining user groups and granting explicit permissions to these groups. Configuring user groups and assigning permissions is done using Access Management in the NetBackup Administration Console. Access Management is the newest method and is the preferred method in future NetBackup releases.
- **Enhanced Authorization and Authentication.** Enhanced authentication allows each side of a NetBackup connection to verify the host and user on the other side of the connection. By default, NetBackup runs without enhanced authentication. Enhanced authorization determines if authenticated users (or groups of users) have NetBackup administrative privileges. By default, NetBackup gives administrative privileges to UNIX root administrators or Windows system administrators on NetBackup servers.

If both Access Management and Enhanced Authorization and Authentication are configured, Access Management takes precedence.

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

About the Vault Operator user group permissions

NetBackup Access Management is used to define user groups, specify which actions each user group can perform, and assign users to those user groups. Each user group can perform only the actions explicitly granted and no others.

When Vault is installed and licensed, NetBackup includes a Vault Operator user group that has permission to perform the operator actions necessary for the Vault process.

Table 9-1 lists the permissions that the Vault Operator user group has in NetBackup Access Management terminology.

Table 9-1 Vault Operator permission sets defaults

Permission sets	Permissions	Vault Operator
Operate media	Browse media	X
	Read media	X
	Inject media	X
	Eject media	X
	Move media	X
	Assign media	X
	Deassign media	X
	Update database	X
	Update barcodes	X
	New	X
	Delete	X
	Expire	X
Read report	Browse report	X
	Read report	X
Operate robot	Browse robot	X
	Read robot	X
	Inventory robot	X
	New robot	X
	Delete robot	X
Drive	Browse drive	X
	Read drive	X
NBU_Catalog	Browse	X
	Read	X
Job	Browse job	X
	Read job	X

Table 9-1 Vault Operator permission sets defaults (continued)

Permission sets	Permissions	Vault Operator
	Suspend job	X
	Resume job	X
	Cancel job	X
	Delete job	X
	Restart job	X
	New job	X
Service	Browse service	X
	Read service	X
Host Properties	Browse Host Properties	X
	Read Host Properties	X
License	Browse license	X
	Read license	X
Volume group	Browse volume group	X
	Read volume group	X
	New volume group	X
	Delete volume group	X
Volume Pool	Browse volume pool	X
	Read volume pool	X
Dev Host	Browse device host	X
	Read device host	X
Vault	Browse vault	X
	Read vault	X
	Manage containers	X
	Run reports	X
ServerGroup	Browse	X

Table 9-1 Vault Operator permission sets defaults (continued)

Permission sets	Permissions	Vault Operator
	Read	X

These permissions are granted only in the scope of actions that are performed in Vault. For example, the Vault Operator group has permission to update databases, but only to the extent that is allowed by Vault, such as when ejecting media changes volume group information for the volume ejected. As defined in the default permission sets, the Vault Operator cannot use the NetBackup Administration Console to change database information that is not related to the operate media actions.

If you use Access Management to administer access by using the default Vault Operator group, those permission sets and permissions apply regardless of whether the actions are initiated from the Vault Operator Menu or the NetBackup Administration Console.

A NetBackup Security Administrator (a user group that is defined within NetBackup Access Management) can use Access Management to add users to the Vault Operator group and change the permission sets and permissions of the Vault Operator group. A Security Administrator also can create new user groups to define new roles.

Because you can change which actions user groups can perform, the Vault documentation cannot specify which actions are or are not allowed by Access Management. If an action cannot be performed because of access management restrictions, NetBackup Administration Console messages explain the restriction.

See the *NetBackup Security and Encryption Guide*.

Note: Giving operators access to the Vault Operator Menu also gives operators the capability to change report destinations. If you do not want your operators to view reports and change report destinations, do not give them access to the Vault Operator Menu. For example, you may not want your operators to see the Recovery Report or to be able to change to whom reports are emailed.

About printing Vault and profile information

You can print a list of the information (robots, vaults, or profiles) that currently appears in the Details pane of the Administration Console. From the **File** menu, choose **Print** or click the **Print** icon on the toolbar.

Copying a profile

If you want to create a profile that is similar to another profile, you can copy the existing profile, rename it, and then change the attributes.

Note: The new profile must belong to the same robot as the original profile.

To copy a profile

- 1 Select the profile you want to copy.
- 2 Open the **Actions** menu and select **Copy Profile**.
- 3 In the **Copy a Profile** dialog box, from the drop-down list under the **Vault** field, choose the vault in which to place the new profile.
- 4 Enter a new name for the profile.
- 5 Click **OK**.

About moving a vault to a different robot

A vault is associated with (that is, belongs to) a specific robot. However, you can change the robot to which the vault belongs. To do so, right-click the robot and select **Change**. Then complete the dialog box, specifying another robot for the vault, and click **OK**.

Note: All vaults that were associated with the previous robot are now associated with the new robot chosen in the dialog box. Some profile configurations may be invalid under the new robot. For example, if the previous robot was associated with a media server that the new robot is not associated with, the configuration is invalid.

Robots are configured in NetBackup through Media Manager. The action that is described in this topic does not change the configuration of a robot in Media Manager.

About changing volume pools and groups

If you change to a new off-site volume group or off-site volume pool(s), you can ensure that media are recalled by configuring a profile that only generates the reports needed to recall media, as follows:

- Configure a vault that uses the old off-site volume group.
- Configure a profile in that vault that does the following:

- Selects no images (that is, configure the **Choose Backups** step so that no backup images are selected).
- Skips the Duplication step and Catalog Backup step.
- Specifies the old volume pools in the volume pool list on the Eject step.
- Generates only the Picking List for Vault and Distribution List for Robot reports.
- Schedule the profile to run on a regular basis.

Media in the old off-site volume group and in the old volume pool(s) are recalled from off-site storage as they expire. After each volume is recalled and injected back into the robot, change its volume pool and group to the new ones. If a volume is returned to a scratch volume pool from which all media are allocated, you do not have to change the volume pool.

After all media in those volume pools and groups are recalled, you can delete the vault, volume group, and volume pools.

Note: If you have media in your off-site vault, Cohesity recommends that you do not change or rename your off-site volume group(s) or off-site volume pool(s). If you begin using new volume pools and groups in your Vault profiles, the reports that recall expired media do not include the old groups and pools.

About NetBackup Vault session files

The directory that is generated for each vault session collects information for the session in 2 log files. The `detail.log` file contains a record of each action that is performed for the session. Some of the information in `detail.log` is also recorded in the NetBackup log files. The `summary.log` file contains a brief description of the vault session, and the results of the session. If email notification is enabled, the information in this file is appended to the email.

The `detail.log` has information about the number of images that are selected by a particular session. In addition, it should record information (during the duplication step) about the total number of images and the number of images duplicated. If these numbers do not match, it means that some images were not duplicated. The log should contain information about which images were not duplicated, either because they were duplicated in a previous session or because the duplication failed for some reason. The actual images that are selected by the session show only if a higher debug level (level 5) is used.

The session log files are located in the following path:

■ UNIX

```
/usr/openv/netbackup/vault/sessions/vault_name/sidxxx

/usr/openv/netbackup/vault/sessions/vault_name/sidxxx/logs
```

On a NetBackup cluster, the `vault/sessions` directory is present in the shared location:

Example: `/opt/VRTSnbu/netbackup/vault/sessions/vault_name/sidxxx/`

■ Windows

```
install_path\NetBackup\vault\sessions\vault_name\sidxxx
```

On a NetBackup cluster, the `vault\sessions` directory is present in the shared location.

Example: `<shared-path>\NetBackup\vault\sessions\vault_name\sidxxx\`

The name of the vault is `vault_name` and it is used for the session. The unique session identifier `xxx` is the identifier that Vault assigns to each vault session. The session ID starts at 1 the first time that Vault runs and increments by 1 for each new session. The session identifier for a Vault session can be found by viewing the Activity Monitor entry for that session.

By default, the session log files are created using the following naming convention, on Windows and UNIX platforms.

- `/usr/openv/netbackup/logs/vault/<vault>/<session-id>/logs/<username>.detail.log`
- `/usr/openv/netbackup/logs/vault/<vault>/<session-id>/logs/<username>.session.log`

For Windows administrative users, the user name **ALL_ADMINS** gets appended in the log file name. In case of UNIX administrative users, the user name **root** gets appended in the log file name. For non-administrative users, only for the NBAC, with a specially privileged user in **Vault_Operator** group who is authorized to run `vlteject -report`, generates the session logs on demand. For more information on legacy logging, refer to the **Legacy Logging** section in *NetBackup Logging Reference Guide*.

Note: Directory names are not case-sensitive on Microsoft Windows systems. Therefore, session directories are created in the same `vault\sessions\vault_name` directory for two or more vaults that have names difference only in the letter-case.

Table 9-2 describes the Vault session logs.

Table 9-2 Vault session log names

Name	Purpose
<code>duplicate.log.nn</code>	Progress information for duplication operations; generated by the <code>-L</code> option of the <code>bpduplicate</code> command.
<code>preview.list</code>	Summary of images to be duplicated if Duplication step is configured or ejected if Eject step is configured and Duplication step is not.
<code>image.list</code>	Lists all images and partial images for a session.
<code>detail.log</code>	Summary of each action that is performed for a Vault session.
<code>summary.log</code>	Brief description of the Vault session and its results. If email notification is enabled, data in this log file is appended.

About setting up Vault session log files

Vault generates session logs and debug logs. The session logs can help you keep track of Vault processes.

See “About NetBackup Vault session files” on page 203.

See “Setting the duration of Vault session files” on page 205.

See “Debug logs” on page 219.

Setting the duration of Vault session files

Vault’s session files are stored in the following directory:

- UNIX

```
/usr/openv/netbackup/vault/sessions
```

- Windows

```
install_path\NetBackup\vault\sessions
```

You can configure the length of time NetBackup retains these files by using the NetBackup Administration Console.

To set the duration of Vault session files

- 1 In the NetBackup Administration Console, select **Host Properties**.
- 2 Select **Master Server** under **Host Properties**.
- 3 In the right pane, right-click the master server and choose **Properties**.

4 Select **Clean-up**.

5 In the **Delete vault logs** field, set the length of time after which to delete the Vault working files.

When the set time has elapsed, the entire `sidxxx` directory is deleted.

You should plan to retain each `sidxxx` directory at least as long as the period of time over which you plan to span consolidated ejects. We suggest that you keep these directories at least a week longer than the consolidation span. If the `sidxxx` directory has been deleted, Vault is unable to eject tapes or generate reports from that session.

Operational issue with disk-only option on Duplication tab

If **Disk Only** is specified on the **Duplication** tab, an image that has no disk copy is not duplicated even if a copy of that image exists on removable media and was selected during the Choose Backups step.

See “About the list of images to be vaulted” on page 122.

See “Operational issues with the scope of the source volume group” on page 206.

Operational issues with the scope of the source volume group

The Source Volume Group on the **Choose Backups** tab spans all steps of a Vault profile (most notably Duplication and Eject). However, if you do not duplicate images, you do not have to specify a source volume group (the **Source Volume Group** field is ignored). Conversely, the **Source of Backups...** field on the **Duplication** tab applies only to the Duplication step.

Even if no images are selected for duplication, images can still be ejected if they are in the Source Volume Group and in an off-site volume pool that is specified on the profile **Eject** tab.

See “Operational issue with disk-only option on Duplication tab” on page 206.

Using the menu user interface

This chapter includes the following topics:

- About using the menu interfaces
- About the Vault administration interface
- Vault operator menu interface
- Vault fields in bpdjobs output

About using the menu interfaces

Vault also includes the following two menu user interfaces (MUIs) that you can use in a terminal window:

- The Vault Administration interface, which lets you configure Vault. It provides the same functionality as Vault Management in the NetBackup Administration Console.
- The Vault Operator Menu interface, which provides a way to eject media and print reports for one or more Vault sessions.

About the Vault administration interface

The information in this section applies to UNIX and Linux systems only.

The Vault Administration interface lets you configure and run Vault from a text-based menu. You can perform the same actions in the **Vault Administration** menu as in the NetBackup Administration Console.

You can use the Vault Administration interface from any character-based terminal (or terminal emulation window) that has a termcap definition or terminfo definition. Use the `vltadm` command to start the Vault Administration interface, and run the `vltadm` command only from the UNIX system on which the NetBackup master server resides. You must have root privileges to run the `vltadm` command.

The `vltadm` command and interface are available on UNIX systems only.

The `vltadm` command resides in the following directory:

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

When you run the `vltadm` command, the following menu appears in the terminal window:

```
Vault Administration
-----
      Robot Name:  <none>
      Vault Name:  <none>
      Profile Name: <none>
r)  Browse all configured robots
v)  Browse all configured vaults for selected robot
p)  Browse all configured profiles for selected vault
n)  Robot management...
t)  Vaults for selected robot...
f)  Profiles for selected vault...
c)  Copy selected profile...
s)  Start session for selected profile...
a)  Vault properties...
h)  Help
q)  Quit
ENTER CHOICE:
```

To browse through specific robots, vaults, or profiles already configured in Vault, press `r`, `v`, or `p`. When the correct robot, vault, or profile appears, type the letter of the action you want to perform.

You can configure different criteria in the Vault Administration interface.

See “About configuring Vault” on page 61.

For help on the currently displayed menu, select the **help** option on that menu. Help includes a tutorial for learning and using the Vault Administration interface.

Vault operator menu interface

The Vault Operator Menu interface lets an authorized user eject and inject tapes and print reports for one or more Vault sessions (an authorized user is one who can run the `vltopmenu` command).

The `vltopmenu` command, which starts the Vault Operator Menu, resides in the following directory:

- UNIX: `/usr/opensv/netbackup/bin`
- Windows: `install_path\NetBackup\bin`

When you run the `vltopmenu` command, the NetBackup Vault Operator menu appears in the terminal window.

```

                                NetBxxxackup Vault Operator Menu
Current Profile: None
Current Session: 0
Current Report Destinations - Print command: /usr/ucb/lpr
                                Email:
                                Directory:

p) Select Profile                m) Modify the Report Destinations...

u) Profile Up                   r) Run Reports for This Session

d) Profile Down                 v) Run Individual Reports...

s) Select Session              cr) Consolidate All Reports

i) Inject Media into Robot      e) Eject Media for This Session

ce) Consolidate All Ejects      re) Consolidate All Reports and Ejects

c) Container Management...

q) Quit

Selection-->
```

Upon startup, the menu displays the current profile, session, and report destinations.

The `vltopmenu` command writes messages about its operations to the log file for Vault commands:

- UNIX

```
/usr/opensv/netbackup/logs/vault/log.mmddyy
```

- Windows

```
install_path\NetBackup\logs\vault\mmddyy.log
```

See the *NetBackup Vault Operator's Guide*.

Vault fields in bpdjobs output

The NetBackup activity monitor utility, `bpdjobs`, displays Vault jobs by default.

Table 10-1 shows the Vault-specific fields that you receive if you run the `bpdjobs` command and use the `-vault` option.

Table 10-1 Vault fields in bpdjobs output

Field	Description
Robot	The name of the robot with which the vault is associated.
Vault	The name of the vault under which the session is running.
Profile	The name of the profile that holds the configuration information for the vault session.
Session ID	The session ID is a unique numeric value for the vault job. Session ID assignment starts at 1 the first time a vault job is run after vault has been installed. The value increments by 1 every time a new vault job runs.
Tapes to Eject	The number of tapes to be ejected for a vault session. If the profile is configured for deferred eject, the tapes may not be ejected yet.
Operation	<div>For Vault jobs, the field contains one of the following values. These values progress from the first value to the last as the Vault job progresses as follows:</div> <ul style="list-style-type: none">■ Choosing Images■ Duplicating Images■ Choosing Media■ Catalog Backup■ Eject and Report■ Done

If a Vault job completes successfully (with exit status = 0), the **State** field and the **Operation** field both contain the value **Done**. If a vault job fails, the Operation field contains the operation occurring at the time the job failed.

Troubleshooting

This chapter includes the following topics:

- About troubleshooting Vault
- About printing problems
- About errors returned by the Vault session
- About media that are not ejected
- About media that is missing in robot
- Reduplicating a bad or missing duplicate tape
- About the tape drive or robot offline
- No duplicate progress message
- About stopping bpvault
- About ejecting tapes that are in use
- About tapes not removed from the MAP
- Revaulting unexpired tapes
- Debug logs

About troubleshooting Vault

This topic contains information about how to diagnose potential problems or errors that you may encounter. Some of these areas include printing problems, missing media issues, no duplicate messages, to how to stop bpvault.

About printing problems

Problems with printing reports that appear to be Vault problems often are problems with the print command configured on the profile **Reports** tab. Therefore, you should test print commands from a command line on the server on which Vault is installed to ensure that they work correctly.

In some rare cases with Microsoft Windows, the print command works correctly when tested from a command prompt but does not work when configured on the profile **Reports** tab. The issue may be with how Windows calls the print command. If you experience such a problem, from a command prompt on the master server on which Vault is installed, enter the following command (use the appropriate server and printer names):

```
NET USE lpt1 \\servername\printername PERSISTENT:YES
```

This problem can also occur in mixed environments of UNIX and Windows.

About errors returned by the Vault session

Every Vault session writes a detailed error status to `stderr`, as follows:

- If the error generated by the Vault session is less than or equal to 255, it returns the actual error code. Error codes less than or equal to 255 (except 252) map to standard NetBackup error codes and are documented in the *NetBackup Troubleshooting Guide*.
- If the Vault session fails with an error code greater than 255, it returns error code 252 and the actual error code is written to `stderr`. Codes greater than 255 are called NetBackup Extended Error Codes and are not supported by all operating systems.

The format of the error text written to `stderr` is:

EXIT status = *error code*

See the *NetBackup Troubleshooting Guide*.

About media that are not ejected

If no media are ejected, it may be because of the following:

- All images have already been vaulted, so no images were selected. Vault determines that a backup image has already been vaulted if a copy of the image is already on a volume in an off-site volume group.

- The media to be vaulted are in a volume group other than the robotic volume group specified for the vault to which the profile belongs.

About media that is missing in robot

Duplication may fail if NetBackup does not know that a requested piece of media is in the robot. For example, a tape may have been moved to the off-site volume group inadvertently even though it remains in the robot. To compare the tapes actually stored in the robot with the Media Manager database, use the NetBackup Administration Console Inventory Robot option.

If the tape is in the robot, use the NetBackup Administration Console to move the tape to the robotic volume group.

If the tape is not found, delete it from the NetBackup system. If the tape is missing yet is assigned and has valid duplicate images, use the command `bpexpdate` to expire the images before you delete the tape from Media Manager. This command is documented in the *NetBackup Administrator's Guide*.

Reduplicating a bad or missing duplicate tape

If a duplicate tape is lost or damaged, you can reduplicate the images that were on the tape if the primary backup images still reside in the robot.

To reduplicate a bad or missing duplicate tape

- 1 Determine which images were on the tape by running the `bpimmedia` command.
 The `bpimmedia` command scans the entire NetBackup image catalog. It may take a few minutes depending on the size of that catalog. Save the output because you need to verify that the correct images were reduplicated.
- 2 Expire the lost or damaged duplicate tapes by using the `bpexpdate` command.
- 3 Determine when the images were created by using the `bpimagelist` command.
- 4 Create a profile that has the same criteria as the profile that created the missing duplicate tape except for the following:
 - Specify policy names only for the policy names used to create the images on the missing tape.

- Set the time window so the profile selects the images on the missing tape. For example, if the original backups were made 30 days ago, set the time window between 32 days and 28 days ago.
- 5 Run the profile by selecting it in the Administration Console and then select **Actions > Start Session**.

Ensure that no other Vault sessions are running before running this new profile.

Before duplicating images, you can verify that the correct images are selected by previewing the session.

See “About previewing a Vault session” on page 117.

About the tape drive or robot offline

If you have a problem with ACSLS drives going offline, try to configure the drives in an UP state or reset the drive. If drives persistently go offline, duplication may hang.

Also, if the tape drives are listed as AVR control in the Administration Console Media and Device Management node, there may be a problem with the robotics control. List all drives as robotically controlled (that is, TLD, ACS, and so on). They are converted to AVR control if a problem occurs with the robot. To diagnose the problem, examine the system logs (for example, `/var/adm/messages` on UNIX systems) for error messages. You can also use the robotic test utilities (such as `robtest`) to further debug the problem.

No duplicate progress message

If you see a message similar to the following in the Vault `detail.log`, the Vault process has not received any new information from the `bpduplicate` process within the time frame specified (in this example, 30 minutes):

```
bpduplicate_progress_logname: no activity in 30 minutes
```

bpduplicate_progress_logname is the progress log that `bpduplicate` creates as it runs the duplication for Vault.

This file resides in the following directory:

- UNIX

```
/usr/openv/netbackup/vault/sessions/vault_name/sidxxx  
/logs/duplicate.log.n
```

- Windows


```
install_path\NetBackup\vault\sessions\vault_name\sidxxx  

\logs\duplicate.log.n
```

The *vault_name* is the name of the vault that is used for the session, *xxx* is the unique session ID, and *n* is the number of the instance of the `bpduplicate` command (1 for the first instance, 2 for the second, and so on).

This message does not necessarily indicate that an error occurred. If the image that is currently being duplicated is very large (for example, several gigabytes), this message appears only for informational purposes. To determine if a problem exists, you can determine the size of the current image. First examine the last few lines of the `details.log` file to determine backup image ID.

Then run the `bpimagelist` command and specify the image ID, as in the following example:

- UNIX

```
bpimagelist -L -backupid server2_0897273363
```

- Windows

```
bpimagelist.exe -L -backupid server2_0897273363
```

The output of this command shows you various statistics about this backup image, including the number of kilobytes written during the backup. If the number is relatively small, there may be a problem with the duplication process. Sometimes this delay is caused by a media mount (which normally does not occur in robotic devices during duplication), by hardware problems, or by the media in use. Examine the Activity Monitor to determine if there are any hardware problems and also check the system logs. If the backup image is very large, regard this message as informational.

About stopping bpvault

To stop a bpvault duplication process, you can use the `vltrun -halttdups` command from the command line interface.

See the *NetBackup Commands Reference Guide*.

This command sends a SIGUSR2 signal to the main duplication process, which is the primary vault job that is currently in the duplication process. This signal is automatically propagated to the other bpvault duplication instances without waiting for any current duplication job instance to finish. The current duplication completes and no more duplication instances run. However, the `bpvault.all` script continues to complete processing of the other bpvault commands.

To end a bpvault process immediately rather than wait a number of minutes for the current bpduplicate processes to complete, first run the `-halttdups` command. Then manually stop the bptm jobs on each server. However, do not stop the bpduplicate processes. The errors that are received by bpvault are logged and these images are not successfully duplicated. The failure of the bptm or bpduplicate means that these images are duplicated on the next attempt if the number of duplicate_days is not exceeded.

About ejecting tapes that are in use

If Vault is configured to eject original media, it is possible that a piece of media is in use during the eject process (for example, for a restore or a media verify procedure). In this case, an error message is generated by Media Manager. A similar error may be generated by non-Media Manager controlled robots if a piece of media is currently in use.

If you receive one of these errors, use the Vault Operator Menu (`vltopmenu`) to re-eject the media after the media is no longer in use. You may receive additional error messages because the rest of the media for the scheduled job has already ejected.

About tapes not removed from the MAP

If media are not removed from the robot's media access port (MAP) and a timeout condition occurs, the job fails with status code 338 (vault eject timed out). If this failure occurs, the Vault reports do not accurately reflect the status of the media.

To recover, you should use the Vault Operator Menu (`vltopmenu`) or the `vlteject` command to eject the media that was not removed from the library and generate the reports.

Revaulting unexpired tapes

If Vault tapes that have not expired are injected back into a robot, you can revault them manually.

To revault unexpired tapes

- 1 Eject the media, as follows:
 - In the NetBackup Administration Console, select the robot into which the media was injected (**Media and Device Management > Robots**).
 - Select the media ID(s) you want to eject.

- Select the **Eject Volumes from Robot....** operation on the **Actions** menu.
- 2 Transfer the media to the off-site volume group by doing the following:
 - Select the media ID(s).
 - Select **Actions > Change Volume Group**.
 - Choose the appropriate off-site volume group from the **New Volume Group Name** drop-down menu.
 - 3 Return the media to your vault vendor so that all backups on that media are available for future disaster recovery.
 - 4 Run the Recovery report to ensure that the media is available for future disaster recovery operations.

Alternatively, you can use the `vmchange` command to eject the media and transfer it to the off-site volume group.

Debug logs

Vault writes debug logs in the standard NetBackup debug logging path as follows:

- Vault commands write messages in log files in a `vault` directory. You must first create the `vault` directories so that daily log files are generated. If the directories do not exist, the log files are not created.

The log files are in the following vault directories:

- UNIX

```
/usr/opensv/netbackup/logs/vault/log.mmdyy
```

- Windows

```
install_path\NetBackup\logs\vault\mmdyy.log
```

- The NetBackup Vault Manager (`nbvault`) writes messages in Veritas Unified Logging (VxUL) log files. (VxUL creates log file names and messages in a standardized format.)
- All unified logs are written to the following locations:
 - UNIX

```
/usr/opensv/logs/nbvault
```

- Windows `install_path\NetBackup\logs\nbvault`
 The originator ID (OID) for NetBackup Vault is 166.

See the "Using Logs and Reports" chapter of the *NetBackup Troubleshooting Guide* for more information on unified logging.

The NetBackup Vault Manager (`nbvault`) manages Vault activity and arbitrates access to the Vault robot, vault, and profile configuration information. The NetBackup Vault Manager must run at all times so Vault functions correctly. On Windows systems, use the `bpup.exe` and `bpdwn.exe` commands to start and stop the NetBackup services, including the NetBackup Vault Manager. On UNIX systems, use `netbackup start` and `netbackup stop` to start and stop NetBackup and Vault daemons.

See "Setting the duration and level for log files" on page 220.

See "Logs to accompany problem reports" on page 221.

Setting the duration and level for log files

The amount of information logged and how long it is retained is controlled by the following NetBackup configuration parameters:

- The length of time NetBackup retains debug logs. This setting affects all debug log files that NetBackup generates, not only the Vault debug logs.
- Vault logging level. Cohesity recommends that you use a debug level of 5 when you generate logs that you send to Cohesity for troubleshooting purposes. You can set the debug level to 5 for all Vault sessions or you can use the `-verbose` option on the `vltrun` command in the Vault policy that initiates the Vault job.

To set the duration and level for log files

- 1 In the NetBackup Administration Console, expand **NetBackup Management**.
- 2 Expand **Host Properties**.
- 3 Select **Master Server**.
- 4 In the right pane, select the master server and then click **Actions > Properties**.
- 5 Select the **Logging** tab.
- 6 Enter a value for **Keep logs for**. Enter the length of time to retain the NetBackup log files. This setting applies to all NetBackup logs, including but not limited to the Vault logs.
- 7 Enter a value for **Keep Vault logs for**. The logging level corresponds to the `bp.conf` entry `VAULT_VERBOSE = level` on UNIX systems. If this value is not specified in the `bp.conf` entry, then the default verbose level is set to 5.

Logs to accompany problem reports

To troubleshoot problems, Cohesity Customer Service requires a set of log files that NetBackup and Vault produces.

In most circumstances, you must provide log files from the following NetBackup processes:

- `admin` (on the master server); administrative commands process
- `bpcd` (on the master server); NetBackup client daemon manager
- `bpsched` (on the master server); NetBackup backup scheduler
- `bptm` (on the media server); NetBackup tape manager
- `nbvault` (on the master server); Vault Manager service or daemon
- `vault` (on the master server); Vault commands

Session log files also are useful for troubleshooting problems, and you should include the appropriate session log files with problem reports you send to Cohesity.

If you use the `vlteject` command or the Vault Operator Menu (`vltopmenu`) to perform consolidated ejects and reports, the following log file may also be useful:

- **UNIX**

```
/usr/opensv/netbackup/vault/sessions/vlteject.mstr
```

- **Windows**

```
install_path\NetBackup\vault\sessions\vlteject.mstr
```

See the *NetBackup Troubleshooting Guide*.

Recovering from disasters

This appendix includes the following topics:

- About disaster recovery
- About disaster recovery in the NetBackup Vault context
- About preparing for recovery
- About recovering NetBackup
- Recovering data and restoring backup images
- Archiving and recovering from a specific point in time

About disaster recovery

This topic provides information about recovering data by using NetBackup and Vault, when you have to recall your media from your off-site storage location. It also provides general information about preparing for a disaster recovery situation.

See the "Disaster Recovery" section in the *NetBackup Troubleshooting Guide*.

Data backup is essential to any data protection strategy, especially a strategy that is expected to assist in disaster recovery. Regularly backing up data and then being able to restore that data within a specified time frame are important components of recovery. Regardless of any other recovery provisions, backup protects against data loss from complete system failure. And off-site storage of backup images protects against damage to your on-site media or against a disaster that damages or destroys your facility or site.

To perform recovery successfully, the data must be tracked to know at what point in time it was backed up. Knowing the time allows your organization to assess the information that cannot be recovered. Configure your data backup schedules to allow your organization to achieve its recovery point objective (RPO), which is the

point in time before which you cannot accept lost data. If your organization can accept one day's data loss, your backup schedule should be at least daily so you can achieve an RPO of one day before any disaster.

Your organization also may have a recovery time objective (RTO), which is the expected recovery time or how long it takes to recover. Recovery time is a function of the type of disaster and of the methods that are used for recovery. You may have multiple RTOs, depending on which services your organization must recover and when.

High availability technologies can make the recovery point very close or even identical to the point of failure or disaster, and they also can provide very short recovery times. However, the closer to the failure that you place your RTO and RPO, the more expensive it is to build and maintain the systems that are required to achieve recovery. Your analysis of the costs and benefits of various recovery strategies should be part of your organization's recovery planning. Understanding disaster recovery planning lets you place Vault and tape-based backups that are stored off-site in the proper context within your disaster recovery objectives.

About what defines a disaster

For an organization, a disaster is an unplanned event that interrupts its ability to function. Usually, the event affects the delivery of critical business functions and results in a loss of data.

The following are generally recognized as the types of disasters possible:

- Technological disasters result in shortcomings in performance, availability, capacity, and accessibility of your IT infrastructures. Technological disasters include computer or Internet crimes, computer viruses, power failures, network or telecommunication failures, hardware or software failures, and other similar failures.
- Human disasters are caused by people, including accidents, explosions, fires, riots, terrorist activities, and other events.
- Natural disasters are caused by nature, including hurricanes, tornadoes, earthquakes, floods, and other natural events.

The effect of a disaster often depends on the scale and timing of the event. Although a disaster is an event that is beyond your control, you can control the way in which your organization reacts to a disaster. By planning and preparing for a disastrous event, you can minimize the effect of the disaster.

About the disaster recovery process

Disaster recovery is the process of responding to an interruption in the services your organization uses to operate. Disaster recovery usually is focused on information, network, and telecommunication services, often at an alternative site and using one or more data-recovery methods.

Disaster recovery is part of a larger topic that is called business recovery, which includes restoring the actual capability for employees to perform their jobs. Business recovery includes logistic related items, such as telephones, office space, living arrangements for employees, and other logistical items. Business recovery itself is part of a larger topic that is called business continuity planning, which includes plans to manage the crisis to your organization, help resume normal business operations, and other inventions.

A resilient organization uses business continuity planning to help ensure that it can survive a disaster and resume operations at an acceptable level.

About disaster recovery plans

A disaster recovery plan is a plan to resume or recover a specific essential operation, function, or process of an organization. Although disaster recovery usually is used to describe information technology and telecommunication services recovery, other services an organization uses to conduct operations can and should be considered part of a plan. For example, an organization's people also are subject to the effects of a disaster and planning should include the effect on them and the resources necessary to help them recover so they can perform their duties.

By planning how your company responds in the event of a disaster, you ensure that your company can do the following:

- Protect critical data.
- Minimize the effect of a disaster.
- Use resources most effectively.
- Maintain business continuity.

About recovery priorities

Your organization must decide between recovery cost (the infrastructure and testing) and the level of functionality that must be recovered. You may choose to recover only the most critical business functions immediately and then recover other functions later. Although all functions of an organization should be valuable and necessary for the organization to operate, it may be acceptable to operate at a reduced level for a specific period of time. The longer your organization can operate without a

function, the easier and less expensive it becomes to recover that function. Therefore, given the higher cost of rapid recovery, only those functions that are required for immediate operation need to be recovered quickly. Delaying recovery of some functions can be a good business decision.

About developing disaster recovery plans

Developing a disaster recovery plan usually begins with an impact analysis that identifies the functions an organization requires to operate and determines how long each function can be unavailable until it affects the organization to an unacceptable extent.

Understanding the effect of disaster helps you identify the objectives for the recovery plan.

The following are examples of the objectives that may be in a disaster recovery plan:

- Ensure service to customers by making critical resources available.
- Minimize economic loss.
- Secure company assets.
- Minimize decision making during the recovery process.
- Reduce reliance on key individuals.
- Ensure a safe and orderly recovery within predetermined time period.
- Maintain a sense of security and organizational stability.

The priority you assign your objectives depends on the needs of your organization. By setting clear, prioritized objectives for your disaster recovery plan, you can reduce your organization's exposure to risks and ensure that your critical systems and networks are available during disruptions.

You can use the two following approaches to create disaster recovery plans:

- A general plan that is used any time a disaster occurs. A general plan should be flexible and is often impact-driven rather than disaster driven (that is, based on the effect to your organization rather than the type of disaster). A general plan usually is based on assumptions that define the scope of each impact in the plan. A general plan is easy to maintain and convenient. However, because it may require that some decisions are made at the time of disaster (such as assessing the type of impact and determining the response), the beginning of recovery can be delayed.
- Multiple smaller plans, each used for a specific disaster that your organization has determined is most likely to occur. For example, individual plans often are

created for power outages, network outages, fires, floods, and other similar occurrences. Individual disaster-specific plans are easier to create than a general plan. It is often clear which plan should be used, so fewer decisions are required at the beginning of recovery, which can result in quicker recovery. However, which plan to use may not always be clear (for example, if a fire causes a power outage). And if a disaster occurs for which a plan does not exist, recovery may be slow to begin and difficult to achieve.

A disaster recovery plan should be easy to follow and not require interpretation. Do not include unnecessary detail. If the plan is implemented, it will be in a time of high stress and pressure to perform; therefore, the plan should be simple, specific, and well tested.

You should publicize your disaster recovery plan within your organization so that everyone knows about it, understands how it works, and understands the reasoning behind the decisions in the plan.

About testing disaster recovery plans

Developing a disaster recovery plan is a waste of time and resources if it is not tested regularly, thoroughly, and frequently (frequently depends on any changes in your organization's functions and environment). The goal of disaster recovery testing is not to pass a test but to find out what does not work. Design your tests to find problems because it is better to find them during a test than during an actual recovery situation.

Testing can be as simple as calling all of the phone numbers in an emergency notification list to verify that everyone can be reached when needed. Or testing can be as complex as actually conducting operations at a recovery site to ensure that everything works correctly. Between those extremes, variations include walkthroughs, during which everyone that is involved in the recovery process discusses their roles in a moderated recovery scenario, and simulations that initiate the recovery plan but use simulated data. Using a combination of testing scenarios to test specific parts of the plan also can be effective.

About disaster recovery in the NetBackup Vault context

In the NetBackup Vault context, disaster recovery means restoring the NetBackup application (master server, media servers, and clients) and then restoring the data that is stored at the off-site storage facility.

If you use NetBackup and Vault to backup your applications and store removable media at a secure off-site location, you also can perform application recovery so you do not have to reinstall your applications.

About preparing for recovery

Recovering data can be a difficult and time consuming process. The success of recovery often depends on how well you prepare for disaster. Your preparations for disaster and what you have to accomplish during a recovery depends on your recovery systems. For example, suppose your recovery site and systems are already operational and have NetBackup and Vault installed. You do not have to protect the NetBackup installation media and the license keys and install NetBackup during the recovery process. You only have to recover the NetBackup catalogs and data. Conversely, if your recovery systems do not have NetBackup and Vault installed and configured, you have to prepare for that and accomplish it during recovery.

You should do the following to prepare for recovery using NetBackup and Vault. (You may not have to do some of the items listed, and you may have to do more than what is listed.)

- Develop a disaster recovery plan.
- Test the disaster recovery plan.
- Back up and vault data regularly. In addition to backing up files on a regular basis, it is important to select the correct files to back up. You should back up all data that your organization's impact analysis determines is critical and store copies at a secure, off-site storage location.
- If you can recover to the same or identical hardware, back up and vault the applications that your organization's impact analysis determines are critical. You also should back up system files so you can quickly restore a system to normal operation.
 - Include all operating system files in your backups. For Microsoft Windows systems, the Windows system directories include the registry, without which it is impossible to restore a system to its original configuration. If you are using a NetBackup exclude list for a client, do not specify any Windows system files in that list.

Restoring operating system files is not helpful if you are recovering data to a different system at your original site or disaster recovery site. You can back up those files, but then not restore them if you are recovering to a different system or site.

- Back up executable and other files for applications you need to conduct critical operations. You may want to save money by using fewer tape

volumes, but backing up critical applications ensures that you can restore them to their exact configurations. For example, if you have applied software updates or patches, restoring from a backup eliminates the need to reapply them, reducing recovery time.

- Every time you vault media, store the Recovery Report securely. The same disaster that destroys your site can destroy your Recovery Report. You need the Recovery Report to identify the media you need to recall from off-site storage. Your vault vendor may let you vault your Recovery Report. If you have a recovery site equipped with computers, email the Recovery Report to that site.
- Record and protect the names of the policies that are used to backup the data you want to recover. The Recovery Report is organized by policy. You need to know which policies are used so you can identify the media you need to recover.
- Record and protect the names of the off-site volume groups for the data you want to recover. Those names are used during the recovery process. Alternatively, you can obtain the off-site volume group names after you restore the NetBackup catalog (because the catalog includes the Vault configuration).
- Document the commands and options that you need to recover data. For example, the `bpchangeprimary` command is used to promote the vaulted images to primary images so that you can restore from them. So you should have a record of the commands and options that you need during the recovery process.
- Protect the NetBackup and Vault installation media. You need the media so you can install NetBackup and Vault on the recovery system if it is not already installed.
- Record and protect the license keys for NetBackup and Vault. You need them for NetBackup and Vault on the recovery system if you have to install NetBackup. You can use temporary license keys if necessary.
- Protect the installation media and record the license keys for any other Cohesity products that must be installed on the recovery systems. For example, if you use the Veritas File System and Veritas Volume Manager on the recovery systems, you need their license keys when you install those products.
- Protect the installation media for the operating system and other applications that are required to run the systems you are using for recovery.
- Protect your DR plan. The same disaster that destroys your site can destroy your DR plan and recovery report. You should have copies stored so that they will be available when needed. Your vault vendor may let you vault a copy of the DR plan.

Note: Effective disaster recovery procedures are specific to an environment and provide detailed information about everything that should be accomplished to prepare for disaster and to recover after disaster occurs. Cohesity provides general disaster recovery information that is intended as a model only. You must evaluate the information and then develop your own disaster recovery plans and procedures.

About recovering NetBackup

See the "Disaster Recovery" section of the *NetBackup Troubleshooting Guide* for information about recovering NetBackup master servers, media servers, and client systems after a disk failure. The procedures include reinstalling NetBackup and may include reinstalling the operating system if it is required.

You must first ensure that your computer and robotic hardware and any necessary network equipment is operational before you reinstall NetBackup. You can then use an appropriate procedure in the *NetBackup Troubleshooting Guide*.

After you reinstall NetBackup, you may have to configure robots and drives. (If you reinstall NetBackup on the original system, the device configuration is restored if you recover the NetBackup catalogs).

Recovering data and restoring backup images

Recovering data can be a difficult and time consuming process. The success of recovery often depends on how well you prepare for disaster and subsequent recovery.

Note: Effective disaster recovery procedures are specific to an environment and provide detailed information about everything that should be accomplished to prepare for disaster and to recover after disaster occurs. Cohesity provides general disaster recovery information that is intended as a model only. You must evaluate the information and then develop your own disaster recovery plans and procedures.

The steps you have to perform to recover can depend on the configuration of your original system and robotic devices, your original NetBackup configuration, the configuration of your recovery system and robots, and the configuration of NetBackup on the recovery systems. Therefore, providing specific disaster recovery procedures for all situations is not possible. Rather, these procedures are intended as general guidelines from which you can create your own procedures for recovering NetBackup and the data that was transferred off-site.

Although some detail is included about restoring the NetBackup catalogs to a recovery system, these procedures do not provide detail about every step of the recovery procedure.

Information in this section assumes the following:

- The primary backup images are unavailable.
- The NetBackup master and media server software, Vault software, client software, and devices are installed and robots and drives are configured on systems to which you are recovering data.
- The NetBackup catalogs and data media have not been recalled from off-site storage.
- The Recovery Report is available.
- You know the name of the off-site volume group name to which the recovered images belong.
- You know the names of the original master and media servers.

To recover data and restore backup images

- 1** Using the Recovery Report, identify the current catalog backup media and the media that is required to restore the data.

The Recovery Report is organized by policy, so you should determine which policies were used to back up the data you want to recover.

- 2** Recall the appropriate catalog backup and data media from off site storage.

- 3** Recover the catalog.

See "Recover the Catalog from an Online, Hot Backup" in the *NetBackup Troubleshooting Guide*.

- 4** If necessary, perform device discovery in NetBackup.

By default, NetBackup catalog backups include the NetBackup device configuration information. If the recovery system has a different device configuration than the original system, the device configuration from the original system overwrites the device configuration when the catalog is recovered.

- 5 If the master and media server names on the recovery system are different than the original system, change the server names in the NetBackup catalogs by using the `bpimage` command.

The `bpimage` command and options required are as follows:

```
bpimage -newserver recovery_system -oldserver original_system
```

You can also use the `bpimage` command if the old system had separate media servers and the recovery system has a combined master and media server. Use the name of the combined master or media server for the argument to the `-newserver` option.

- 6 Specify the backup copy from which to restore by adding the copy number to the file `ALT_RESTORE_COPY_NUMBER`.

After you add the copy number to the file, all subsequent restore operations restore from that backup copy. The restore can be from the Backup, Archive, and Restore interface or from the `bprestore` command. Exception: the `bprestore -copy x` option and argument override the value in the `ALT_RESTORE_COPY_NUMBER` file.

The `ALT_RESTORE_COPY_NUMBER` file must be in the following directory on the NetBackup master server:

- UNIX

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

- Windows

```
install_path\VERITAS\NetBackup
```

See "Restoring from a specific backup copy" in the Backup, Archive, and Restore help.

- 7 If the media is not suspended or frozen, suspend the media.

Use the `bpmedia` command to suspend the media. Suspending the media prevents NetBackup from writing backup images to that media.

- 8 If the NetBackup Administration Console is not running, start it.

- 9 Inject the media into the robot.

Injecting the media moves it into the robot and also changes the off-site volume group attribute of the media to robotic volume group so NetBackup knows that the volumes are in the robot and ready for restore operations.

See "About injecting media" on page 130.

- 10** Using the Backup, Archive, and Restore interface, restore the data.

See the *NetBackup Backup, Archive, and Restore Getting Started Guide*.

- 11** After restoring all of the data, revault the media.

See “Revaulting unexpired media” on page 146.

Archiving and recovering from a specific point in time

If your data center or computing environment requires recovery to a specific point in time (not just to the most recent valid backups), you can set up a process that ensures that you can recover both the NetBackup catalog and the data for that specific time. You should retain the corresponding catalog backups for the same length of time as the corresponding data backups.

The following high-level information is intended as an overview of how to archive the catalog and data so you can recover to a specific point in time. Detailed instructions for accomplishing all the tasks necessary are not included.

Before you can recover to a specific point in time, you must archive the catalog and data.

To archive the catalog and data to a specific point in time

- 1** Use your normal procedures to vault the data and NetBackup catalogs for that data.
- 2** Use the `bpmmedia` command to freeze the data volumes and catalog volumes that you want to retain.

Freezing the volumes prevents them from becoming unassigned and from appearing on the Picking List for Vault report. Do not recall the volumes from off-site storage when they expire.

- 3** Vault a printed copy of the Recovery Report for that specific point in time.

You need the Recovery Report from the specific point in time so you can recall and restore the appropriate catalog and data volumes.

- 4** Optionally, remove the media IDs from the volume database.

This reduces the size of the database and improves performance. Depending on the number of volumes, maintaining the media IDs in the volume database may not degrade performance much.

To recover the catalog and data to a specific point in time

- 1 Retrieve the appropriate printed Recovery Report from off-site storage.
- 2 Using the Recovery Report, recall the appropriate catalog backup and data volumes from off site storage.
- 3 Recover the catalog that you recalled from off-site storage.
That version of the catalog contains information about the archived volumes and the images on them.
- 4 Use the `bpexpdate` command to reset the expiration date on the recalled volumes so they are not expired. Use the `-policy` option to change the expiration date for all media that is associated with a specific policy.
- 5 Change the images to be recovered to primary (NetBackup restores from the primary image).

To change a large number of images to primary, use `bpchangeprimary -group` option to specify all images in a specific off-site volume group.

For information about the `bpchangeprimary` command, see the *NetBackup Commands Reference Guide*.

- 6 Restore the data.

The *NetBackup Administrator's Guide, Volume I* includes an alternative procedure for archiving the catalog. That alternative procedure uses the `catarc` catalog archive policy to archive old data in the NetBackup catalog. You can then vault the archived catalog data or a copy of the archived catalog data.

See "Catalog Archiving" in the *NetBackup Administrator's Guide, Volume I*.

Vault file and directory structure

This appendix includes the following topics:

- UNIX files and directories
- Windows files and directories

UNIX files and directories

Vault is installed in `/usr/opensv/netbackup` on UNIX systems. The following table describes the files in each of the directories Vault creates and uses on UNIX systems. Also included are Vault files that reside in NetBackup directories. The files are either copied into the directories during the installation process or created as Vault sessions run. The paths are specified in relation to the NetBackup directory.

Table B-1 lists the files and directories for Vault in UNIX.

Table B-1 Files and directories for Vault in UNIX

Directory, Program, or File	Purpose
<code>bin/nbvault</code>	Vault daemon that manages Vault activity.
<code>bin/vltadm</code>	Curses-based utility to configure NetBackup Vault and monitor its operations. <code>vltadm</code> requires root (administrator) privileges.
<code>bin/vltcontainers</code>	Command used to add media logically to containers.
<code>bin/vlteject</code>	Command used to eject media from Vault sessions and run the reports that are selected in the profile.

Table B-1 Files and directories for Vault in UNIX (*continued*)

Directory, Program, or File	Purpose
<code>bin/vltinject</code>	Command used to inject media into a robot and update the Media Manager database.
<code>bin/vltoffsitemedia</code>	Command that lets the user change the off-site parameters of a given piece of media.
<code>bin/vltopmenu</code>	Utility that lets the user initiate a menu screen containing the various options that an operator of the NetBackup Vault feature can use. It allows the user to eject or inject media, print various reports individually or collectively, and consolidate reports and ejects across sessions.
<code>bin/vltrun</code>	Process that runs all the NetBackup commands that are used during a Vault session.
<code>bin/goodies/vlt_ejectlist_notify</code>	Script called by the vault session just before vault tapes are ejected.
<code>bin/goodies/vlt_end_notify</code>	Script called by the vault session just before it exits.
<code>bin/goodies/vlt_endeject_notify</code>	Script called by the vault session at the end of eject processing.
<code>bin/goodies/vlt_start_notify</code>	Script called by the vault session after it starts.
<code>bin/goodies/vlt_starteject_notify</code>	Script called by the vault session when the eject process starts.
<code>db/vault/vault.xml</code>	Vault configuration file.
<code>help/vltadm</code>	Contains help files for the Vault Administration (vltadm) interface.
<code>/logs/vault</code>	Directory where the Vault commands log messages.
<code>vault</code>	Main directory for Vault. Contains session directories.
<code>vault/sessions</code>	A subdirectory that contains working session directories and log files. In a cluster environment, the sessions directory must reside on the shared disk.
<code>vault/sessions/cntrDB</code>	The database of information for media that is vaulted in containers.
<code>vault/sessions/sidxxx</code>	Subdirectory that contains working session subdirectories. Can be manually removed to reduce disk usage if necessary.
<code>vault/sessions/vault_name/session.last</code>	Counter to show current duplication session

Table B-1 Files and directories for Vault in UNIX (*continued*)

Directory, Program, or File	Purpose
<code>allmedia_inventory.rpt</code>	All Media Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>container_inv.rpt</code>	Container Inventory report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault/sessions/vault_name/sidxxx/logs/detail.log</code>	Shows the output of every command that was run during the session.
<code>detailed_distlist_vault.rpt</code>	Detailed Distribution List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>distlist_robot.rpt</code>	Distribution List for Robot Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>distlist_vault.rpt</code>	Distribution List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault/sessions/vault_name/sidxxx/logs/duped.images</code>	List of images successfully duplicated during the session.
<code>vault/sessions/vault_name/sidxxx/logs/duplicate.log.nn</code>	Contains the output from <code>bpduplicate</code> .

Table B-1 Files and directories for Vault in UNIX *(continued)*

Directory, Program, or File	Purpose
<code>vault/sessions/vault_name/sidxxx/eject.list</code>	List of media to be ejected for the session.
<code>lostmedia.rpt</code>	Lost Media Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report will include the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>non_vaulted.rpt</code>	Non-vaulted Images Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>offsite_inventory.rpt</code>	Off-site Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>picklist_robot.rpt</code>	Picking List for Robot report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>picklist_vault.rpt</code>	Picking List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault/sessions/vault_name/sidxxx/preview.list</code>	A list of all images that to be considered for duplication or ejection by the current vault session.

Table B-1 Files and directories for Vault in UNIX (*continued*)

Directory, Program, or File	Purpose
<code>recovery.rpt</code>	Recovery Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault/sessions/vault_name/sidxxx/logs/summary.log</code>	<p>This file provides a concise view of the critical session details that occurred during a session. For example, the file can contain the following pieces of information:</p> <ul style="list-style-type: none">■ The duration of a session■ The robot, vault, profile, session ID, and Job ID■ A summary report that shows the master server, start and stopped times, and the exist status. The exit status can show if an error occurred and which session step if occurred at. <p>This log is appended for email notification.</p>
<code>summary_distlist_vault.rpt</code>	Summary Distribution List report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault_inventory.rpt</code>	Vault Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.

Windows files and directories

Vault is installed in the directory that is specified by `install_path\NetBackup` on Windows systems. The following table describes the files in each of the directories Vault creates and uses on Windows. Also included are Vault files that reside in NetBackup directories. The files are either copied into the directories during the installation process or created as Vault sessions run. The paths are specified in relation to the NetBackup directory.

Table B-2 lists the files and directories for Vault in Windows.

Table B-2 Files and Directories for Vault in Windows

Directory, Program, or File	Purpose
bin\nbvault.exe	NetBackup Vault Manager service.
bin\vlcontainers.exe	Command used to add media logically to containers.
bin\vlteject.exe	Command used to eject media from Vault sessions and run the reports that are selected in the profile.
bin\vltinject.exe	Command used to inject media into a robot and update the Media Manager database.
bin\vltoffsitemedia.exe	Command that lets the user to change the off-site parameters of a given piece of media.
bin\vltopmenu.exe	Utility that lets the user initiate a menu screen containing the various options that an operator of the NetBackup Vault feature can use. It lets the user eject or inject media, print various reports individually or collectively, and consolidate reports and ejects across sessions.
bin\vltrun.exe	Process that runs all the NetBackup commands that are used during a Vault session.
bin\goodies\slt_ejectlist_notify	Script called by the vault session just before vault tapes are ejected.
bin\goodies\slt_end_notify	Script called by the vault session just before it exits.
bin\goodies\slt_endeject_notify	Script called by the vault session at the end of eject processing.
bin\goodies\slt_start_notify	Script called by the vault session after it starts.
bin\goodies\slt_starteject_notify	Script called by the vault session when the eject process starts.
db\vault\vault.xml	Vault configuration file.
logs\nbvault	Directory where the Vault service, nbvault.exe, logs messages.
logs\vault	Directory where the Vault commands log messages.
vault	Main Vault directory. Contains the session directories.
vault\sessions	A subdirectory containing working session directories and log files. In a cluster environment, the sessions directory must reside on the shared disk.
vault\sessions\cntrDB	Database of information for media that is vaulted in containers.

Table B-2 Files and Directories for Vault in Windows (*continued*)

Directory, Program, or File	Purpose
<code>vault\sessions\sidxxx</code>	Subdirectory containing working session subdirectories. Can be manually removed to reduce disk usage if necessary.
<code>vault\sessions\vault_name\session.last</code>	Counter to show current duplication session
<code>allmedia_inventory.rpt</code>	All Media Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>container_inv.rpt</code>	Container Inventory report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault\sessions\vault_name\sidxxx\logs\detail.log</code>	Shows the output of every command that was run during the session.
<code>detailed_distlist_vault.rpt</code>	Detailed Distribution List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report in that folder includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>distlist_robot.rpt</code>	Distribution List for Robot report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>distlist_vault.rpt</code>	Distribution List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.

Table B-2 Files and Directories for Vault in Windows (*continued*)

Directory, Program, or File	Purpose
<code>vault\sessions\vault_name \sidxxx\logs\duped.images</code>	List of images successfully duplicated during the session.
<code>vault\sessions\vault_name \sidxxx\logs\duplicate.log.nn</code>	Contains the output from bpduplicate.
<code>vault\sessions\vault_name \sidxxx\eject.list</code>	List of media to be ejected for the session.
<code>lostmedia.rpt</code>	Lost Media Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>non_vaulted.rpt</code>	Non-vaulted Images Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>offsite_inventory.rpt</code>	Off-site Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>picklist_robot.rpt</code>	Picking List for Robot report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>picklist_vault.rpt</code>	Picking List for Vault report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.

Table B-2 Files and Directories for Vault in Windows (*continued*)

Directory, Program, or File	Purpose
<code>vault\sessions\vault_name \sidxxx\preview.list</code>	A list of all images to be considered for duplication or ejection by the current vault session.
<code>recovery.rpt</code>	Recovery Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault\sessions\vault_name \sidxxx\logs\summary.log</code>	<p>This file provides a concise view of the critical session details that occurred during a session. For example, the file can contain the following pieces of information:</p> <ul style="list-style-type: none">■ The duration of a session■ The robot, vault, profile, session ID, and Job ID■ A summary report that shows the master server, start and stopped times, and the exist status. The exit status can show if an error occurred and which session step if occurred at. <p>This log is appended for email notification.</p>
<code>summary_distlist_vault.rpt</code>	Summary Distribution List report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.
<code>vault_inventory.rpt</code>	Vault Inventory Report. If a report name includes a timestamp, it is part of a consolidated report operation. If you configure the profile Reports tab to write reports to a folder, the file name of the report includes the session ID. Report file names that include <code>_ff</code> are for debug purposes only. They are created when the Administration Console generates reports.

Index

A

- access management 198
- ACS number 63
- ACSLs server 63
- adding alternate media server names 31, 66
- administering access to Vault 198
- advanced duplication 85
 - to avoid sending data over the network 46
- All Media Inventory report 190
- alternate media server names
 - adding 31, 66
- Alternate Media Server Names tab 66
- alternate read server
 - to avoid sending data over the network 45
- Alternate Read Server setting 87, 94
- Any Available storage unit setting 40–41
- assigning multiple retentions with one profile 141
- At Time of Eject setting (Suspend This Session's Media) 106

B

- Backup Policies setting 80
- Backups Started setting 80
- bpdbjobs
 - changes for Vault 210

C

- CAP.. *See* media access port
- catalog backup
 - Catalog Backup Schedule setting 101
 - Skip the Catalog Backup Step setting 101
- Catalog Backup tab 98
- changing off-site volume group name 202
- changing off-site volume pool name 202
- changing volume group name 202
- changing volume pool name 202
- choose backups
 - Backup Policies setting 80
 - Backups Started setting 80
 - Clients setting 80

- choose backups (*continued*)
 - disk pools 81
 - Media Servers setting 81
 - Schedules setting 81
 - Source Volume Group setting 82
 - Type Of Backups setting 82
 - Volume pools 82
- Choose Backups tab 78
- clearing the media description field 154
- clearing the vault fields 131
- Clients setting 80
- clusters 21
 - shared disk 236, 240
- Complete Inventory List for Vault. *See* All Media Inventory report
- Complete Inventory List for Vault.. *See* All Media Inventory report
- concurrent copies
 - continue or fail 162
 - during advanced duplication 171
 - during basic duplication 168
 - during original backup 164
 - during Vault duplication 165
 - overview 161
 - through the NetBackup Catalog node 167
 - to avoid sending duplicates over the network 44
- configuration
 - Catalog Backup tab 98
 - Choose Backups tab 78
 - Duplication tab 83
 - Eject tab 101
 - methods 63
 - profile 77
 - Reports tab 107
 - robots 70
 - volume pools 54
- Container Inventory Report 191
- Containers of Many Media setting 73
- continue
 - for concurrent copies 162
- Copies setting 91, 94

- copy
 - primary backup 84
- copying a profile 202
- creating
 - a profile 75
 - a vault 71
 - a vault policy 113
- critical policies 100
- Customer ID setting 73

D

- debug logs 219
- deferred reports 110
- Deferred setting (Eject Mode) 106
- Destination Storage Unit setting 91, 93, 96
- Destination Volume Pool setting 91, 93, 96
- Detailed Distribution List for Vault 183
- disaster recovery
 - definition 225
 - definition of disaster recovery plan 225
 - developing a disaster recovery plan 226
 - preparing for disaster 223
 - priorities 226
 - testing a disaster recovery plan 227
- disk staging 29
- Distribution List for Vault 182, 187
- duplicate images 15
- duplication
 - advanced 85
 - Alternate Read Server setting 87, 94
 - basic 85
 - Copies setting 91, 94
 - Destination Storage Unit setting 91, 93, 96
 - Destination Volume Pool setting 91, 93, 96
 - Expire Original Disk Backup Images setting 88
 - Expire Original Tape Backup Images setting 89
 - Fail Copies setting 92, 94
 - increase throughput 47
 - Multiple Copies dialog 91
 - Multiple Copies setting 89
 - multiplexed 166
 - Number Of Read Drives setting 90, 95
 - Number Of Write Drives setting 91, 93, 96
 - Preserve Multiplexing setting 90
 - Primary Copy setting 89, 92, 95
 - Retention Level setting 90
 - Retention setting 93, 95
 - Source Backup Server setting 94–95
 - Source Backups Reside On setting 90, 95

- duplication *(continued)*
 - through the NetBackup Catalog node 167
- Duplication tab 83
- duplication throughput
 - configuring for multiple drives 47
 - multiple-drive scenario 48

E

- e-mail
 - notifying a tape operator when eject begins 151
- eject
 - consolidating ejects 129
 - Eject Mode setting 106
 - Suspend Media On Which Backups Were Written setting 107
- Eject Mode settings 106
- Eject tab 101
- eject.list file 123
- ejected tapes returned to robot 218
- ejecting partial images
 - use suspend to avoid 32
- error codes 214
 - extended 121
- EXIT status 214
- Expire Original Disk Backup Images setting 88
- Expire Original Tape Backup Images setting 89
- extended error codes 121

F

- fail
 - for concurrent copies 162
- Fail Copies setting 92, 94
- First Off-Site Slot ID setting 73
- Full Inventory List for Vault.. See Off-site Inventory report

G

- General tab 65

I

- images
 - duplicate 15
 - original 15
 - primary backup 84
- Immediate Eject setting 106
- immediate reports 110
- Immediate setting (Eject Mode) 106

- Immediately setting (Suspend This Session's Media) 106
- Inline Tape Copy
 - during advanced duplication 171
 - during basic duplication 168
- installation
 - on Windows systems 24
 - prerequisites for a Windows systems 25
 - prerequisites for UNIX and Linux systems 22–23
- Inventory List for Vault.. See Vault Inventory report
- Iron Mountain Electronic Format report 195

L

- license key
 - adding on Windows systems 25
- load balancing
 - by duplicating daily and ejecting weekly 43
 - profiles for both originals and duplicates 42
- log files
 - debug logs 219
 - set duration 220
 - vault session 205
- LSM number 63

M

- MAP.. See media access port
- master server
 - host name of 62
- media access port
 - capacity 63
 - number 63
 - to use for eject 73
- media description field
 - clearing 154
- media missing in robot 215
- media server names 62
 - adding alternate 31, 66
- Media Servers setting 81
- menu user interface
 - bpdjobs 210
 - overview 207
 - Vault Administration Interface 207
 - Vault Oerator Menu Interface 209
 - vltadm 207
 - vltopmenu 209
- Monitoring a vault session 119
- moving a vault to different robot 202

- multiple copies 161
 - overview 161
- Multiple Copies dialog 91
- Multiple Copies setting 89
- multiple retention mappings 141
- multiple volume groups 32
- multiplexed duplication 166

N

- network
 - avoid sending duplicates over the network 44
- New Profile dialog 76
- New Vault dialog 72
- New Vault Robot dialog 70
- notify scripts
 - for a specific profile 153
 - for a specific robot 153
 - for a specific vault 153
 - order of execution 154
 - using 151
 - vlt_ejectlist_notify 152
 - vlt_end_notify 152
 - vlt_endeject_notify 152
 - vlt_start_notify 152
 - vlt_starteject_notify 152
- notifying a tape operator when eject begins 151
- Number Of Read Drives setting 90, 95
- Number Of Write Drives setting 91, 93, 96

O

- Off-site Inventory report 189
- off-site volume group 73
 - renaming 202
- Off-Site Volume Group setting 73
- off-site volume pool 54
 - renaming 202
- Off-Site Volume Pools setting 106
- Off-site Volume Pools windows (Eject tab) 106
- organizing reports by profile 51
- organizing reports by robot 50
- organizing reports by vault 51
- original images 15

P

- partial backups
 - use suspend to avoid vaulting 32
- Picking List for Robot 181
- Picking List for Vault 186

- policy 113
 - configuration information 114
 - names 115
 - NetBackup policy for Vault 112
 - schedule names 115
- preferred vaulting strategies 28
- Preserve Multiplexing setting 90
- preview vault session 117
- preview.list file 117
- primary backup copy 84
- primary backup image 84
- Primary Copy setting 89, 92, 95
- printing
 - troubleshooting problems 214
- profile
 - configuring 77
 - copying a 202
 - creating 75
 - notify script for a specific 153
 - organizing reports by 51
 - overlap time window 30
 - printing information 201
- Profile dialog 75
- profiles
 - for both originals and duplicates 43

R

- recovering backup images 155
- recovery
 - keep primary copy on site 36
 - match volume pools to usage 35
 - of damaged media 155
 - preparing for efficient 34
 - revault unexpired media 37
 - use precise naming conventions 35
 - vault NetBackup catalogs 34
- Recovery Report for Vault 192
- renaming off-site volume group 202
- renaming off-site volume pool 202
- renaming volume group 202
- renaming volume pool 202
- reports
 - All Media Inventory 190
 - Complete Inventory List for Vault.. See All Media Inventory report
 - consolidating reports 179
 - Container Inventory Report 191
 - deferred 110
 - Detailed Distribution List for Vault 183
 - reports *(continued)*
 - Distribution List for Robot 187
 - Distribution List for Vault 182
 - Full Inventory List for Vault.. See Off-site Inventory report
 - immediate 110
 - inventory 188
 - Inventory List for Vault.. See Vault Inventory report
 - Iron Mountain Electronic Format 195
 - Lost Media Report 51, 193
 - Off-site Inventory 189
 - organizing by profile 51
 - organizing by robot 50
 - organizing by vault 51
 - Picking List for Robot 181
 - Picking List for Vault 186
 - printing 175
 - Recovery Report for Vault 192
 - Summary Distribution List for Vault 185
 - types of 181
 - Vault Inventory report 188
 - Reports tab 69, 107
 - resource contention
 - avoid by robot usage 38
 - avoiding 38
 - load balancing 42
 - sharing resources with backup jobs 42
 - specifying different volume pools for source and destination 43
 - use disk staging to avoid 29
 - vault original backups to avoid 29
 - when the read drive is not in the vault robot 42
 - when two processes try to use the same drive 38
 - resuming a vault session 118
 - Retention Level setting 90, 93, 95
 - Retention level setting 81
 - Retention Mappings tab 68
 - retention period based on copy one retention period 141
 - revault media 147
 - robot
 - configuring for vault 70
 - control host 71
 - notify script for a specific 153
 - number 71
 - organizing reports by 50
 - types of 62
 - Robot Name setting 71

- Robot Number setting 71
- Robot Type setting 71
- robotic volume group 32
- Robotic Volume Group setting 74
- running multiple sessions simultaneously 116

S

- Schedules setting 81
- scratch volume pools 50
- server name group 66
- session files
 - setting the duration of 205
- setting the duration of session files 205
- Skip the Catalog Backup Step setting 101
- Skip the Eject Step setting 107
- Slot ID 73
- Slots for Individual Media setting 74
- Source Backup Server setting 94–95
- Source Backups Reside On setting 90, 95
- source volume group 40–41
- Source Volume Group setting 82
- status codes 214
- storage unit
 - Any Available setting 40–41
 - name of 63
 - number of drives in 63
- Summary Distribution List for Vault 185
- supported clusters 21
- supported robots 21
- supported systems 21
- suspend
 - Suspend Media On Which Backups Were Written
 - setting 32
 - Suspend Option setting 32
 - use suspend to avoid partial backups 32
- Suspend Media for the Next Session setting 107
- Suspend Media on Which Backups Were Written
 - setting 107
- Suspend Option setting 107
- Suspend This Session's Media setting 107

T

- time window
 - overlapping 30
- troubleshooting
 - bad or missing duplicate media 215
 - drive or robot offline 216
 - ejected tapes returned to robot 218

- troubleshooting (*continued*)
 - ejecting tapes while in use 218
 - error codes 214
 - logs 221
 - media missing in robot 215
 - no duplicate progress message 216
 - printing problems 214
- Types Of Backups setting 82

U

- uninstall Vault
 - from UNIX systems 23
- UNIX files and directories 235

V

- Vault
 - accessing 14
- vault
 - clearing the Vault fields 131
 - creating a 71
 - moving to a different robot 202
 - name 74
 - notify script for a specific 153
 - only the intended backups 32
 - organizing reports by 51
 - original backups 29
 - paradigm 15, 28
 - policy 112
 - printing vault information 201
 - vendor 74
- Vault Administration Interface 207
- Vault dialog 72
- Vault Inventory report 188
- Vault Management Properties dialog
 - Alternate Media Server Names tab 66
 - General tab 65
 - Reports tab 69
 - Retention Mappings tab 68
- Vault Name setting 74
- Vault Operator Menu Interface 209
- vault original backups 29
- vault session
 - previewing 117
 - resuming 118
 - running 112
 - running multiple sessions simultaneously 116
 - setting the duration of session files 205
 - stopping 118

- Vault Vendor setting 74
- vaulting
 - containers 73
 - defined 15
 - original backups in a 24x7 environment 34
 - paradigm 15, 28
 - preferred strategies 28
- vltoffsitemedia command
 - adding an expiration date 150
 - assigning slot number 145
 - changing Vault attributes 149
- volume group
 - configuration 73
 - renaming 202
 - robotic 74
 - source 82
- volume pool
 - configuration 54
 - destination 91, 93, 96
 - for backup media that remains on site 54
 - for catalog media 54
 - match to data usage 35
 - naming conventions 35
 - off-site 54
 - overview 54
 - renaming 202
 - use scratch volume pools 50
- Volume pools
 - specifying different source and destination 43

W

- Windows files and directories 239
- Windows systems
 - delicensing Vault 25
- working files
 - setting the duration of 205