

NetBackup™ Web UI VMware Administrator's Guide

Release 10.1



Last updated: 2022-10-28

Legal Notice

Copyright © 2022 Veritas Technologies LLC. All rights reserved.

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

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

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

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

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

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Managing VMware servers	7
	Add VMware servers	7
	Validate and update VMware server credentials	8
	Browse VMware servers	9
	Remove VMware servers	9
	Create an intelligent VM group	10
	Remove an intelligent VM group	17
	Add a VMware access host	17
	Remove a VMware access host	18
	Change resource limits for VMware resource types	18
	About VMware discovery	19
	Change the autodiscovery frequency of VMware assets	20
	Discover VMware server assets manually	20
Chapter 2	Protecting VMs	22
	Working with VMware policies in the web UI	22
	Protect VMs or intelligent VM groups	23
	Schedules	24
	Backup options and Advanced options	25
	Exclude disks from backups	26
	Snapshot retry options	27
	Customize protection settings for a VMware asset	28
	Remove protection from VMs or intelligent VM groups	28
	View the protection status of VMs or intelligent VM groups	29
Chapter 3	Malware scan	31
	Scan a VMware asset for malware	31
Chapter 4	Instant access	33
	Create an instant access VM	33
	Restore files and folders from a VM backup image	35
	Download files and folders from a VM backup image	37
	Things to consider before you use the instant access feature	38
	Instant access Build Your Own (BYO)	40

	Prerequisites of Instant Access Build Your Own (BYO)	41
	Hardware configuration requirement of Instant Access Build Your Own (BYO)	42
	Frequently asked questions	42
	VM malware scan	44
Chapter 5	Instant rollback	46
	Prerequisites of instant rollback	46
	Things to consider before you use the instant rollback feature	46
	Instant rollback from a VM backup image	47
Chapter 6	Continuous data protection	50
	CDP terminology	51
	CDP architecture	52
	About continuous data protection	52
	Prerequisites	53
	Capacity-based licensing for CDP	54
	Steps to configure CDP	55
	Removing VMs from the CDP gateway	55
	Defining the CDP gateway	56
	Sizing considerations	57
	Limiting concurrent CDP backup jobs	59
	Controlling full sync	61
	Monitoring CDP jobs	62
	Using accelerators with CDP	64
	Recovering CDP protected VMs	65
	Some limitations of CDP	65
	Troubleshooting for CDP	66
Chapter 7	VM recovery	70
	Recover a VM	70
	Recovery options	72
	Advanced recovery options	72
	Advanced recovery options: Format of restored virtual disks	73
	Advanced recovery options: Transport mode	74
	About VMware agentless restore	74
	Prerequisites and limitations of VMware agentless restores	75
	Recover files and folders with VMware agentless restore	77
	About restricted restore mode	77
	Recover VMware Cloud Director virtual machines	79

Chapter 8	Protecting VMs using hardware snapshot and replication	80
	About virtual machines and hardware snapshots	80
	Deployment and architecture	81
	Features and applications supported	81
	Prerequisites for hardware snapshot and replication	82
	Operations supported with hardware snapshot	84
	Configuring a VMware policy to use hardware snapshot	85
	Configuring a VMware policy to use NetBackup snapshot manager replication	89
	Jobs in the Activity Monitor that use hardware snapshot for VMs	90
	Notes and limitations	91
	Troubleshooting with VMware hardware snapshot and replication operations	92
Chapter 9	Troubleshooting VMware operations	97
	Errors when adding VMware servers	98
	Errors when browsing VMware servers	98
	Errors for the status for a newly discovered VM	99
	Error when downloading files from an instant access VM	100
	Troubleshooting backups and restores of excluded virtual disks	101
	Restore fails for a virtual machine with multiple datastores	103

Managing VMware servers

This chapter includes the following topics:

- [Add VMware servers](#)
- [Validate and update VMware server credentials](#)
- [Browse VMware servers](#)
- [Remove VMware servers](#)
- [Create an intelligent VM group](#)
- [Remove an intelligent VM group](#)
- [Add a VMware access host](#)
- [Remove a VMware access host](#)
- [Change resource limits for VMware resource types](#)
- [About VMware discovery](#)
- [Change the autodiscovery frequency of VMware assets](#)
- [Discover VMware server assets manually](#)

Add VMware servers

Use this procedure to add VMware servers and their credentials.

To add VMware servers and their credentials

- 1 On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
 The tab shows the vCenters, ESXi servers, and VMware Cloud Director servers that you can access.
- 2 Click **Add** to add a server.
- 3 Select the server type and enter its host name, and its credentials.
- 4 Choose a **Backup host for validation**.
- 5 Indicate a **Port** number for connection.
 If the default port number has not been changed on the VMware server, no port specification is required. If the VMware server has been configured to use a different port, specify that port number.
- 6 Click **Save**.
 VMs and other objects appear after the discovery process for the VMware server completes.

Validate and update VMware server credentials

After a VMware server is added, you can validate or update the credentials for the server.

To validate VMware credentials

- 1 On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
- 2 Select one or more VMware servers, then click **Validate**.
 NetBackup verifies the current credentials for the selected VMware servers.
 If the credentials are not valid, NetBackup indicates **Invalid** under **Credentials**.

To update VMware server credentials

- 1 On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
- 2 Locate the VMware server.
- 3 Select **Actions > Manage credentials**.
- 4 Update the credentials as needed.
- 5 Click **Save**.

Browse VMware servers

You can browse vCenter servers, standalone ESXi servers, and VMware Cloud Director servers to locate VMs and view their details. VM details include their protection plans and recovery points.

To browse VMware servers

- 1 On the left, click **Workloads > VMware**.
- 2 Click **VMware servers** to begin searching.

The list includes: the names and types of vCenters, standalone ESXi servers, and VMware Cloud Director servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine whether the server's VMs and other objects have been successfully discovered.

To locate a server, you can enter a string in the search field.

- 3 Click on a server to begin drilling into it.
You can navigate back to a higher level by clicking the up-arrow.
- 4 Click on a VM to view its protection status, recovery points, and restore activity.
- 5 Click **Add protection** to subscribe the VM to a plan.

Remove VMware servers

Use this procedure to remove VMware servers from NetBackup.

Note: If you delete a server, all virtual machines that are associated with the deleted VMware server are no longer protected. You can still recover existing backup images, but backups of VMs on this server fail.

To remove a VMware server

- 1 On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
The tab lists: the names and types of vCenters, standalone ESXi servers, and VMware Cloud Director servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine when the server's VMs and other objects were last discovered.
- 2 Locate the VMware server.
- 3 Select **Actions > Delete**.
- 4 If you are sure that you want to delete the VMware server, click **Delete**.

Create an intelligent VM group

You can create an intelligent VM group based on a set of filters called queries. NetBackup automatically selects virtual machines based on the queries and adds them to the group. You can then apply protection to the group. Note that an intelligent group automatically reflects changes in the VM environment and eliminates the need to manually revise the list of VMs in the group.

Note: The web UI must discover the VMs on each server before the query can select from them. If a VMware server was recently added in the web UI, its VMs may not have been discovered.

See [“Change the autodiscovery frequency of VMware assets”](#) on page 20.

To discover the VMs immediately:

See [“Discover VMware server assets manually”](#) on page 20.

Note: Intelligent VM groups are not supported for VMware Cloud Director VMs.

To create an intelligent VM group

- 1 On the left, click **Workloads > VMware**.
- 2 Click the **Intelligent VM groups** tab and then click **Add**.
- 3 Enter a name and description for the group.
- 4 Select the appropriate VMware server.
- 5 Perform one of the following:
 - Select **Include all VMs**.
This option uses a default query to select all VMs that currently reside in the vCenter or ESXi for backup when the protection plan runs.
 - To select only the VMs that meet specific conditions, create your own query: Click **Add condition**.

- 6 To add a condition, use the drop-downs to select a keyword and operator and then enter a value.

The options are described after this procedure: [Query options for creating intelligent VM groups](#).

Examples are also available: [Example queries](#)

To change the effect of the query, click **Condition** and click **AND** or **OR**, then select the keyword, operator, and value for the condition. For example:

The screenshot shows a query builder interface. At the top, there are two tabs: "AND" (selected) and "OR". To the right of the tabs are two buttons: "+ Condition" and "+ Sub-query". Below the tabs, there is a list of conditions. The first condition is selected and expanded, showing three fields: "displayName" (keyword), "Contains" (operator), and "prod" (value). Below this, there is a second condition with "tag" (keyword), "=" (operator), and "eng" (value). Each condition has a trash icon to its right.

You can also add sub-queries to a condition, if necessary. Click **Sub-query** and click **AND** or **OR**, then select the keyword, operator, and value for the sub-query condition. For example:

The screenshot shows the same query builder interface as before, but with an additional sub-query added. The first condition is still selected and expanded, showing "displayName" (keyword), "Contains" (operator), and "prod" (value). Below this, there is a second condition with "tag" (keyword), "=" (operator), and "eng" (value). Below the second condition, there is a sub-query. The sub-query has its own tabs: "AND" (selected) and "OR". To the right of the sub-query tabs are two buttons: "+ Condition" and "+ Sub-query". The sub-query contains one condition with "cluster" (keyword), "Starts with" (operator), and "clust" (value). Each condition and sub-query has a trash icon to its right.

- 7 To test the query, click **Preview**.

The query-based selection process is dynamic. Changes in the virtual environment can affect which VMs the query selects when the protection plan runs. As a result, the VMs that the query selects later when the protection plan runs may not be identical to those currently listed in the preview.

- 8 To save the group without adding it to a protection plan, click **Add**.

To save and add it to a protection plan, click **Add and protect**, select the plan, and click **Protect**.

Note: When you click **Preview** or you save the group, the query options are treated as case-sensitive when the VMs are selected for the group. Under **Virtual machines**, if you click on a VM that was not selected for the group, the **Member of virtual machine groups** field reads `none`.

However, when you add the group to a protection plan, some of the query options are treated as case-insensitive when the protection plan's backup runs. As a result, the same VM may now be included in the group and is backed up.

For the case behavior of each option, see [Query options for creating intelligent VM groups](#).

Query options for creating intelligent VM groups

Note the following for intelligent VM groups

- When using queries in **Intelligent VM groups**, the NetBackup web UI might not display an accurate list of VMs that match the query if the query condition has non-English characters. However, during the backup, the correct VMs are selected even though the VM attributes are non-English.
- Using the `not equals` filter condition on any attribute returns assets including those that have no value (null) present for the attribute. For multi-value attributes such as `tag`, the assets that do not match at least one of the values of the attribute are not returned
- When the server of an Intelligent VM group is updated, all existing access definitions configured for that Intelligent group are removed because the intelligent group is now registered with the new server namespace. You need to add new access definitions for the updated Intelligent group.

Table 1-1 Query keywords

Keyword	Description	Case-sensitive when protection plan runs
annotation	The text that is added to VM annotations in a vSphere client.	Yes
connectionState	The status of the VM connection to the ESX server. For example, if a virtual machine's ESX server is down, that virtual machine is not connected.	No
cluster	The name of the cluster (group of ESXi servers) where the VMs reside.	No
datacenter	The name of the datacenter.	No
datacenterPath	The folder structure that defines the path to a datacenter. Use this option if the datacenter name that you want to filter on is not unique in your environment.	Yes
datastore	The name of the datastore.	Yes
displayName	The VM's display name.	Yes
host	The name of the ESXi server. The ESXi host name must match the name as defined in the vCenter server.	No
dnsName	The VM's DNS name in vSphere Client.	No
guestOS	The VM guest OS type that is recorded in the vSphere client.	Yes
hostName	The VM name that is derived from a reverse lookup of its IP address.	No
instanceUuid	The VM's instance UUID. For example: 501b13c3-52de-9a06-cd9a-ecb23aa975d1	No
networkName	The name of the network switch (on an ESX server) or distributed switch.	No
powerState	The power state of the VM.	No
tag	The name of the VM's tag.	Yes
template	Indicates if the VM is a virtual machine template.	No
version	The VMware version of the virtual machine. For example, vmx-04, vmx-07, vmx-08.	Yes

Table 1-1 Query keywords (*continued*)

Keyword	Description	Case-sensitive when protection plan runs
vmFolder	The name of the VM folder (within a datacenter), which includes the path to the folder that contains the VMs. See the section called “VMFolder examples” on page 16.	No
vmxDatastore	The name of the VMX datastore (sometimes called the vmx directory or configuration datastore).	Yes
vmxDatastoreType	The type of the VMX datastore. Values are NFS or VMFS.	No

Query operators

Table 1-2 Query operators

Operator	Description
Starts with	Matches the value when it occurs at the start of a string. For example: If the value you enter is <code>box</code> , this option matches the string <code>box_car</code> but not <code>flatbox</code> .
Ends with	Matches the value when it occurs at the end of a string. For example: If the value you enter is <code>dev</code> , this option matches the string <code>01dev</code> but not <code>01dev99</code> or <code>devOP</code> .
Contains	Matches the value you enter wherever that value occurs in the string. For example: If the value you enter is <code>dev</code> , this option matches strings such as <code>01dev</code> , <code>01dev99</code> , <code>devOP</code> , and <code>development_machine</code> .
=	Matches only the value that you enter. For example: If the value you enter is <code>VMtest27</code> , this option matches <code>VMTest27</code> (same case), but not <code>vmtest27</code> , <code>vmTEST27</code> , or <code>VMtest28</code> .
!=	Matches any value that is not equal to the value that you enter.

Example queries

In this example, the query adds to the group any VM that has `prod` in its display name.

A screenshot of a query builder interface. At the top right, there are two buttons: "+ Condition" and "+ Sub-query". Below these, there is a single condition row. The first column contains the text "displayName", the second column contains a dropdown menu with "Contains" selected, and the third column contains the text "prod". To the right of the text "prod" is a trash icon.

To change the effect of the query, click **Condition** and click **AND** or **OR**, then select the keyword, operator, and value for the condition. For example:

A screenshot of a query builder interface. At the top, there are two buttons: "AND" and "OR", with "AND" selected. To the right, there are two buttons: "+ Condition" and "+ Sub-query". Below these, there are two condition rows. The first row has "displayName" in the first column, "Contains" in the second column, and "prod" in the third column. The second row has "tag" in the first column, "=" in the second column, and "eng" in the third column. Each row has a trash icon to its right.

This example uses **AND** to narrow the scope of the query: it selects only the VMs that have `prod` in their display name and that also have a tag named `eng`. If a VM does not have `prod` in its display name as well as a tag named `eng`, that VM is not added to the group.


To broaden the scope of the query, use **OR**:

A screenshot of a query builder interface. At the top, there are two buttons: "AND" and "OR", with "OR" selected. To the right, there are two buttons: "+ Condition" and "+ Sub-query". Below these, there are two condition rows. The first row has "displayName" in the first column, "Contains" in the second column, and "prod" in the third column. The second row has "tag" in the first column, "=" in the second column, and "eng" in the third column. Each row has a trash icon to its right.

In this example, **OR** causes the query to add the following to the group:

- The VMs that have `prod` in their display name (regardless of any tags).
- The VMs that have a tag named `eng` (regardless of the display name).

You can also add sub-queries to a condition, if necessary. Click **Sub-query** and click **AND** or **OR**, then select the keyword, operator, and value for the sub-query condition. For example:



In this example, the sub-query causes the query to narrow the scope further. From the VMs that have both `prod` in their display name and a tag named `eng`, only the VMs in clusters that start with `clust` are selected.

VMFolder examples

For example, assume the following VM folders containing a total of 65 VMs:

`vm\VM_backup_prod1` (contains 5 VMs)

`vm\VM_backup_prod1\cluster1` (contains 10 VMs)

`vm\VM_backup_prod2` (contains 50 VMs)

To include the VMs in `vm\VM_backup_prod1` but not the VMs in `cluster1` or in any other folder:

```
VMFolder Equal "vm\VM_backup_prod1"
```

To include the VMs in `vm\VM_backup_prod1` and in its subfolder `cluster1`:

```
VMFolder Equal "vm\VM_backup_prod1"
```

OR

```
VMFolder StartsWith "vm\VM_backup_prod1"
```

Note: The first backslash is an escape character that causes the following backslash to be interpreted as a literal character.

To include all 65 VMs: `VMFolder StartsWith "vm\VM_backup_prod"`

Note: Any VM that is in a path that begins with `vm\VM_backup_prod` is included.

Remove an intelligent VM group

Use the following procedure to remove an intelligent VM group.

To delete an intelligent VM group

- 1 On the left, click **Workloads > VMware**.
- 2 Locate the group under the **Intelligent VM groups** tab.
- 3 If the group is not protected, select it and then click **Delete**.
- 4 If the group is protected, click on the group, scroll down and click the lock symbol, and click **Unsubscribe**.
- 5 Click **Remove**.

Add a VMware access host

NetBackup uses a special host that is called a VMware access host. It is a NetBackup client that performs backups on behalf of the virtual machines. The access host is the only host on which NetBackup media server or client software is installed. No NetBackup client software is required on the virtual machines. However, the access host must have access to the datastores of the virtual machines. The access host reads the data from the datastore and sends it over the network to the media server.

The VMware access host was formerly called the VMware backup host or the VMware backup proxy server. The access host is referred to as the recovery host when it performs a restore.

Note: Make sure that NetBackup media server software or client software is installed on any access host that you add.

To add a VMware access host

- 1 On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.
- 2 On the right, select **VMware settings > Access hosts**.
NetBackup lists any access hosts that were previously added.
- 3 Click **Add**.
- 4 Enter the name of the access host and then click **Add**.

Remove a VMware access host

To remove a VMware access host

- 1 On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.
- 2 On the right, select **VMware settings > Access hosts**.
NetBackup lists any access hosts that were previously added.
- 3 Locate the VMware access host and then click the delete icon.
- 4 To confirm the deletion, click **Delete**.

Change resource limits for VMware resource types

VMware resource limits control the number of backups that can be performed simultaneously on a VMware resource type. The settings apply to all NetBackup policies for the currently selected primary server.

To change the resource limits for VMware resource types

- 1 On the left, click **Workloads > VMware**.
- 2 On the top right, select **VMware settings > Resource limits**.
For each resource, the default value is **0** (No limit).
- 3 Select the VMware resource type you want to change and then **Edit**.

Note: The **Snapshot** resource limit is different from the other resource types. It sets a limit for the number of simultaneous snapshot-only operations within a vCenter domain, such as create snapshot and delete snapshot. This limit applies only during the snapshot creation and snapshot deletion phases of a backup. It does not control the number of simultaneous backup jobs. This **Snapshot** limit can be useful for controlling the effect that multiple snapshot operations have on the vCenter server. Add a specific vCenter to override the global snapshot setting for that vCenter.

4 Choose from the following options.

Set a global limit for a VMware resource type.

Locate the **Global** setting and select the **Limits** value that you want to apply.

This value limits the number of simultaneous backups that are performed for the resource type.

Set a limit for a specific VMware resource.

Click **Add**.

From the list, select the resource.

Select the **Limits** value that you want to apply.

This value limits the number of simultaneous backups that are performed for the selected resource.

5 Click **Save**.

Limits indicates the number of simultaneous backups that can be performed for the resource type. This value is the global limit. The **Override** value indicates how many resources have any limits that are different from the global limit.

Reset the resource limits for all VMware resources

To reset the resource limits for all VMware resources

- ◆ Click **Reset default values** to remove all the overrides and set all global VMware resource limits to their default values.

About VMware discovery

NetBackup automatically starts the discovery of the VMware server when you add a VMware server or update credentials. The backup host information is used to validate the credentials and perform the discovery.

To serve as a backup host, a media server or client must be at NetBackup 8.1.2 or later. For older versions, the backup host credential validation succeeds, but the discovery of the VMware servers fails. Discovery occurs at set intervals. (The default interval is every 8 hours.)

See [“Change the autodiscovery frequency of VMware assets”](#) on page 20.

To discover the VMs immediately:

See [“Discover VMware server assets manually”](#) on page 20.

Change the autodiscovery frequency of VMware assets

Automatic discovery of VMware assets occurs at regular intervals. The default frequency is every 8 hours. Use this procedure to change the autodiscovery frequency.

To change the frequency of autodiscovery of VM assets

- 1 On the left, click **Workloads > VMware**, then click the **Virtual machines** tab.
- 2 On the right, select **VMware settings > Autodiscovery**.
- 3 Select **Frequency > Edit**.
- 4 Use the up or down arrows to choose how often you want NetBackup to perform autodiscovery of VMware assets. Then click **Save**.

The range from which you may choose is 1 hour to 24 hours. To set the autodiscovery frequency in minutes or seconds or to disable autodiscovery, you must use the VMware autodiscovery API.

Discover VMware server assets manually

Use this procedure to manually discover any VMware server so that you can view and protect recently added assets.

Note: Automatic discovery of VMs and other objects in the vCenter, ESXi server, or VMware Cloud Director server begins: when server credentials are added or updated through the web UI or an API. However, the server's VMs and other objects might not appear in the UI immediately. They appear after the discovery process for the VMware server completes. Discovery also occurs at set intervals according to the `VMWARE_AUTODISCOVERY_INTERVAL` option. (The default interval is every 8 hours.) More information about this option is available:

See [“Change the autodiscovery frequency of VMware assets”](#) on page 20.

To manually discover VMware server assets

- 1** On the left, click **Workloads > VMware**, then click the **VMware servers** tab.
 The tab lists: the names and types of vCenters, standalone ESXi servers, and VMware Cloud Director servers that you have access to. You can also review the **Discovery Status** and **Last discovery attempt** to determine when the server's VMs and other objects were last discovered.

- 2** Locate and select the VMware server.

- 3** Select **Actions > Discover**.

The discovery operation may fail if the VMware server credentials are invalid. To validate and update the credentials:

See [“Validate and update VMware server credentials”](#) on page 8.

For more information about the protection status of VMs and intelligent VM groups:

See [“View the protection status of VMs or intelligent VM groups”](#) on page 29.

See [“Errors for the status for a newly discovered VM”](#) on page 99.

Protecting VMs

This chapter includes the following topics:

- [Working with VMware policies in the web UI](#)
- [Protect VMs or intelligent VM groups](#)
- [Customize protection settings for a VMware asset](#)
- [Remove protection from VMs or intelligent VM groups](#)
- [View the protection status of VMs or intelligent VM groups](#)

Working with VMware policies in the web UI

You can add new VMware policies and manage existing VMware policies in both the NetBackup Administration Console and the NetBackup web UI.

Functionality for VMware policies is the same in both interfaces, with a few differences:

In the NetBackup web UI

Query Builder keywords You must use OData keywords in the Query Builder.

In the NetBackup Administration Console

You must use VIP keywords in the Query Builder.

	In the NetBackup web UI	In the NetBackup Administration Console
VM discovery and refresh	The way you manually discover or refresh VMware server assets is different in the web UI. See “Discover VMware server assets manually” on page 20. Also use this procedure to refresh discovery of VMware server assets, for example, before you browse for VMware servers.	You use the Refresh button to manually discover VMware server assets when browsing for VMs.
Last update information	Last update information about VMware server discovery is listed for each VMware server on the Browse for virtual machines dialog. Discovery status and Last discovery attempt is also listed for each server when you click Workloads > VMware, then VMware servers to view VMware servers.	You view Last Update from the VMware policy's Client tab's Browse for virtual machines dialog. This information pertains to the last update of a VMware server within the selected vCenter, not each VMware server.

Note: Details about VMware policies are not included with this guide. You can find complete documentation for VMware policies in the [NetBackup for VMware Administrator's Guide](#).

To add or change a VMware policy in the web UI

- 1 On the left, click **Protection > Policies**.
- 2 To change a VMware policy, select it from the list.
To add a policy, click **Add**, enter a **Policy name**, and select **VMware** from the **Policy type** drop-down list.
- 3 Complete all required fields as you would in the NetBackup Administration Console.
- 4 Click **Create** to save a new policy.
Click **Save** to save changes to an existing policy.

Protect VMs or intelligent VM groups

Use the following procedure to subscribe an asset (VMs or intelligent VM groups) to a protection plan. When you subscribe an asset to a protection plan, you assign predefined backup settings to the asset.

Note: Protection plans are not supported for VMware Cloud Director VMs.

To protect VMs or VM groups

- 1 On the left, click **Workloads > VMware**.
 - 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, click the box for the VM or the VM group and click **Add protection**.
 - 3 Select a protection plan and click **Next**.
 - 4 Adjust any settings as necessary.
 - Change the backup start window.
See “[Schedules](#)” on page 24.
 - **Backup options** and **Advanced** options.
See “[Backup options and Advanced options](#)” on page 25.
 - 5 Click **Protect**.
- The results of your choices appear under **Virtual machines** or **Intelligent VM groups**.

Schedules

The following schedule settings are included in a protection plan.

Note that when you customize a protection plan for an asset, you can only edit the following schedule settings:

- Start window

Table 2-1 Schedule options for protection plans

Option	Description
Backup type	The type of backup that the schedule controls.
Recurrence (frequency)	How frequently or when to run the backup.
Keep for (retention)	How long to keep the files that were backed up by the schedule.
Replicate this backup	Replicates the snapshot to another volume.
Duplicate a copy immediately to long-term retention	Immediately after the schedule is created, a copy is duplicated to the media that is selected for long-term storage.

Table 2-1 Schedule options for protection plans (*continued*)

Option	Description
Start window	On this tab, set the window during which a backup can start.

Backup options and Advanced options

The user can adjust the following settings when subscribing to a protection plan.

Backup options

Table 2-2 Backup options for protection plans

Option	Description
Select server or host to use for backups	The host that performs backups on behalf of the virtual machines. Users can choose Automatic to have NetBackup pick the media server, based on the storage unit. Or, the user can select another host from the list. These hosts are other media servers in the environment or hosts that are configured as an access host.
If a snapshot exists, perform the following action	Specifies the action that NetBackup takes when a snapshot is discovered before NetBackup creates a new snapshot for the virtual machine backup. For example, users can choose to stop a backup if any snapshots exist. If snapshots are not automatically deleted, the performance of the virtual machine may eventually decline. Undeleted snapshots can cause restore failures due to lack of disk space.
Exclude selected virtual disks from backups	Specifies the virtual disks to exclude from backups. See "Exclude disks from backups" on page 26.

Advanced options

Table 2-3 Advanced options for protection plans

Option	Description
Enable virtual machine quiesce	By default, I/O on the virtual machine is quiesced before NetBackup creates the snapshot. In the majority of cases, you should use this default. Without quiescing file activity, data consistency in the snapshot cannot be guaranteed. If you disable the quiesce, you must analyze the backup data for consistency.

Table 2-3 Advanced options for protection plans (*continued*)

Option	Description
Allow the restore of application data from virtual machine backups	This option allows users to restore application data from full backups of the virtual machine. Note that in NetBackup 8.3 or earlier, application data for Microsoft Exchange Server or Microsoft SharePoint Server must be restored with the NetBackup Backup, Archive, and Restore interface. Data for Microsoft SQL Server must be restored with the NetBackup MS SQL Client. See the documentation for your NetBackup database agent for more details.
Transport mode	Specifies the transport mode to use for backups or how to read the data from the datastore. For more information on transport modes, see the vendor documentation for your virtualization environment.
Snapshot retry options	See “Snapshot retry options” on page 27.

Exclude disks from backups

Excluding virtual disks can reduce the size of the backup, but use these options carefully. They are intended only for the virtual machines that have multiple virtual disks.

Table 2-4 Options for excluding virtual disks

Exclude option	Description
All boot disks	Consider this option if you have another means of recreating the boot disk. The virtual machine's boot disk is not included in the backup. Any other disks are backed up. Note: Data files are available in the restored data disks. However, you cannot start a virtual machine that is restored from this backup.
All data disks	Consider this option only if you have a separate protection plan that backs up the data disks. The virtual machine's data disks are not included in the backup. Only the boot disk is backed up. Note: When the virtual machine is restored from the backup, the virtual machine data for the data disk may be missing or incomplete.

Table 2-4 Options for excluding virtual disks (*continued*)

Exclude option	Description
Exclude disks based on a custom attribute	Use this option to allow the VMware administrator to use a custom attribute to control which disks are excluded from backups. The attribute must have comma-separated values of device controllers for the disks to be excluded. For example: <code>scsi0-0, ide0-0, sata0-0, nvme0-0</code>. The default value for this attribute is <code>NB_DISK_EXCLUDE_DISK</code>. Or, you can choose your own value. If you add disks to the custom attribute value between any differential backups, those disks are excluded from the next backup. The VMware administrator must use a VMware interface to apply the attribute to the disks to exclude. See the NetBackup Plug-in for VMware vSphere Web Client Guide or the NetBackup Plug-in for VMware vSphere Client (HTML5) Guide.
Specific disks to be excluded	Use this option to exclude a specific disk by the disk type, controller, and LUN that represent the virtual device node of the disk. Click Add to specify additional disks. If you add controllers between any differential backups, their disks are excluded from the next backup.

Snapshot retry options

For most environments, the default values for the snapshot retry options are appropriate. It may be helpful to adjust these settings based on the size of the virtual machine and the processing load on the VMware server.

Table 2-5 Snapshot retry options

Option	Description
Maximum number of times to retry a snapshot	The number of times the snapshot is retried.
Maximum length of time to complete a snapshot	The time, in minutes, to allow the snapshot operation to complete. If snapshots do not complete, set this option to a specific period to force a time-out. Use the Maximum length of time to wait before a snapshot is retried setting to retry the snapshot at a later time.
Maximum length of time to wait before a snapshot is retried	The time to wait (in seconds) before the snapshot is retried.

Customize protection settings for a VMware asset

You can customize certain settings for a protection plan, including the schedule backup window and other options.

- See [“Schedules”](#) on page 24.
- See [“Backup options and Advanced options”](#) on page 25.

To customize protection settings for a VMware asset

- 1 On the left, click **Workloads > VMware**.
- 2 Do one of the following:
 - Edit the settings for a VM
 - On the **Virtual machines** tab, click on the VM that you want to edit.
 - Edit the settings for an intelligent group
 - On the **Intelligent VM groups** tab, click on the group that you want to edit.
- 3 Click **Customize protection > Continue**.
- 4 Adjust any of the following settings:
 - The backup start window.
See [“Schedules”](#) on page 24.
 - **Backup options** and **Advanced options**.
See [“Backup options and Advanced options”](#) on page 25.
- 5 Click **Protect**.

Remove protection from VMs or intelligent VM groups

You can unsubscribe VMs or intelligent VM groups from a protection plan. When the asset is unsubscribed, backups are no longer performed.

Note: When you unsubscribe an asset from a protection plan, there is a possibility that the asset displays **Classic policy** in the web UI. This situation can happen when an asset is subscribed to a protection plan and a backup runs for that asset. Then the asset is unsubscribed from the protection plan while it has a valid backup image. The web UI displays **Classic policy**, but there may or may not be an active policy protecting the asset.

To remove protection from a VM or intelligent VM group

- 1 On the left, click **Workloads > VMware**.
- 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, click the VM or the intelligent VM group.
- 3 Click **Remove protection > Yes**.
Under **Virtual machines** or **Intelligent VM groups**, the asset is listed as Not protected.

View the protection status of VMs or intelligent VM groups

You can view the protections plans that are used to protect VMs or intelligent VM groups.

To view the protection status of VMs or intelligent VM groups

- 1 On the left, click **Workloads > VMware**.
- 2 Select the **Virtual machines** tab or **Intelligent VM groups** tab, as appropriate.

Note: Sorting on assets across asset types, that is, without the Asset Type filter, returns results grouped by asset types (Virtual Machine and Intelligent VM groups) and sorted within each asset type.

- 3 Click the VM or the intelligent VM group.

The **Protection** tab shows the details of the plans that the asset is subscribed to.

Note: If the asset has been backed up, but Status indicates it has not, see the following information.

See [“Errors for the status for a newly discovered VM”](#) on page 99.

- 4 If the asset is not protected, click **Add protection** to select a protection plan.

See [“Protect VMs or intelligent VM groups”](#) on page 23.

Malware scan

This chapter includes the following topics:

- [Scan a VMware asset for malware](#)

Scan a VMware asset for malware

The following requirements exist before you can scan for malware:

- The primary server must be at NetBackup 10.0.1 or later.
- The backups were performed with a storage server at NetBackup 10.1 or later.
- Backup images are stored on MSDP storage with instant access capability, for the supported policy type only.
- The scan host pool must be configured with scan hosts.
- The last backup must be successful.
- You must have an RBAC role with permissions to perform malware scans.

To scan a VMware asset for malware

- 1 On left, click **VMware > Virtual machines**.
- 2 Locate and click on the VM.
- 3 Select **Actions > Scan for malware**.
- 4 On the **Malware scan** page, do the following:
 - Select the date range for the scan by selecting **Start date/time** and **End date/time**.
 - Select **Scanner host pool**
 - From the **Select current status of malware scan** list select one of the following:

- **Not scanned**
- **Not infected**
- **Infected**
- **All**

5 Click **Scan for malware**.

Note: The malware scanner host can initiate a scan of three images at the same time.

6 After the scan starts, you can see the **Malware Scan Progress** on **Malware Detection**, you can see the following fields:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**

Note: Any backup images that fail validation are ignored.

- **In progress**
- **Pending**

Instant access

This chapter includes the following topics:

- [Create an instant access VM](#)
- [Restore files and folders from a VM backup image](#)
- [Download files and folders from a VM backup image](#)
- [Things to consider before you use the instant access feature](#)
- [Instant access Build Your Own \(BYO\)](#)
- [VM malware scan](#)

Create an instant access VM

You can create an instant access VM from a NetBackup backup image. The VM is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the VM's snapshot directly on the backup storage device to allow your ESXi host or cluster to treat the snapshot as a normal VM.

The mounted VM snapshot can be used for a variety of purposes. For example:

- Recovering files from the VM, or copying a vmdk file.
- Running tests on the VM, such as testing a patch.
- Troubleshooting or disaster recovery.
- Verifying an application.

Note: This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance, and Build Your Own (BYO) server. This feature requires that the NetBackup backup image is stored on a Media Server Deduplication Pool (MSDP) storage device. More information on using instance access VMs is available:

See [“Things to consider before you use the instant access feature”](#) on page 38.

To create an instant access VM

- 1 On the left, click **VMware**.
- 2 Locate the VM and click on it.
- 3 Click the **Recovery points** tab, then click the date on which the backup occurred.

The available images appear in rows with the backup timestamp for each image.

- 4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.

- 5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Create instant access virtual machine**.
- 6 Review the recovery settings and make changes if needed.

Note the **Recovery options**:

Allow overwrite of existing virtual machine	If a VM with the same display name exists at the destination, that VM must be deleted before the recovery begins. Otherwise, the recovery fails.
Power on after provisioning	Automatically powers on the VM when the recovery is complete.
Enable vMotion	Starts the migration of the VM after it is created and then displays progress of the VM migration. Note: For a NetBackup 8.1.2 storage server, the vMotion option is not used even if it is enabled.

7 Click Create.

NetBackup makes a snapshot of the VM backup image and creates an instant access mount point. The snapshot of the image appears on the **Instant access virtual machines** tab. You can now use the VM like any other VM on the ESXi server.

8 For details on the restored VM, click on the VM under the Instant access virtual machines tab and click View details.**9 When you are finished with the VM, you can click Delete to remove the mounted VM snapshot. The VM is removed from the ESXi server.**

Note: If vMotion is enabled and completed successfully, deleting a VM only removes the mounted share. The VM is still available on the ESXi server as this VM is migrated to another datastore.

Restore files and folders from a VM backup image

You can browse an instant access image of the VM to restore files and folders.

Note: More information on using instance access VMs is available:

See [“Things to consider before you use the instant access feature”](#) on page 38.

To restore files and folders from a VM backup image

1 On the left, click VMware.**2 Locate and click on the VM.****3 Click the Recovery points tab. In the calendar view, click the date on which the backup occurred.**

The available images are listed in rows with the backup timestamp for each image.

4 Select the Allow the selection of recovery of points that are malware-affected to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This option is enabled only for users with the necessary RBAC role or related RBAC permissions.

- 5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Restore files and folders**.

NetBackup creates an instant access mount point in the background.

- 6 Select the files and click **Add to restore list**.

Click on a folder to drill into it. Use the folder path to navigate back to higher levels in the hierarchy.

`yygvvm004-win10 / C / $WINDOWS.~BT / Drivers`

Enter a file name to search for files.

The restore list displays the selected files and folders with the location and size of each file.

- 7 Select the restore options:
 - **Restore everything to the original directory**
 - Enter the name of the target VM (the default is the original VM) and the username and password for the target VM.
 - **Restore everything to a different directory**
 - In **Directory for restore**, enter the destination path for restore.

Note: If the storage server is NetBackup 8.1.2, enter the `Single File Full Path` and not the `Parent Folder Path`.

- Select the **Flatten existing directory structure** check box to restore all files to a single directory.

Note: If the storage server is NetBackup 8.1.2, this option is automatically used during restore.

- Enter the name of the target VM (the default is the original VM) and the username and password for the target VM.

- 8 Select the **Overwrite existing files** check box to overwrite all the existing files.

Note: If the storage server is NetBackup 8.1.2, this option is automatically used during restore.

A summary of your selections is displayed.

- 9 Click **Start recovery** to restore the files.

The **Activity** tab displays the status of the recovery.

Download files and folders from a VM backup image

You can browse an instant access image of the VM to download files and folders.

Note: More information on using instance access VMs is available:

See [“Things to consider before you use the instant access feature”](#) on page 38.

To download files and folders from a VM backup image

- 1 On the left, click **VMware**.
- 2 Locate and click on the VM.
- 3 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.

The available images are listed in rows with the backup timestamp for each image.

- 4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.

- 5 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Download files and folders**.

6 Select the files and click **Add to download list**.

Click on a folder to drill into it. Use the folder path to navigate back to higher levels in the hierarchy.

`yygvm004-win10 / C / $WINDOWS.~BT / Drivers`

Enter a file name to search for files.

The download list displays the selected files and folders with the location and size of each file.

7 After the download package is created, click **Download**.

The **Activity** tab displays the status of the recovery.

Things to consider before you use the instant access feature

Note the following about the **Instant access virtual machines** feature:

- This feature is supported with backup copies that are created from the local or cloud LSU (logical storage unit) using the NetBackup web UI or Instant Access APIs.

For more information about limitations of instant access for cloud LSU (logical storage unit), refer to the [NetBackup Deduplication Guide](#).

- This feature is supported with backup copies that are created from protection plans using the web UI or from classic policies that are created with the NetBackup Administration Console.
- This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance, and Build Your Own (BYO) server.

Instant access on Flex WORM storage requires the following services:

- NGINX, NFS, SAMBA, WINBIND (if Active directory is required), SPWS, VPFS
- This feature is limited to 50 concurrent mount points from a Media Server Deduplication Pool (MSDP) media server or from a WORM storage server. If you have a Flex appliance, this feature is limited to 50 concurrent mount points from each node.
- By default, vSphere allows a maximum of eight NFS mounts per ESXi server. Note that NetBackup requires an NFS mount for each instant access VM you

create. To remove the NFS mount, remove the instant access VM when you are done with it.

If the NFS limit for an ESXi host has been reached and you try to create another instant access VM, the attempt fails. To increase the maximum NFS mounts per ESXi server, see the following VMware article:

<https://kb.vmware.com/s/article/2239>

- This feature does not support backups of VMs that have independent disks. VMware does not support snapshots of independent disks in a VM, either persistent disks or non-persistent disks. As a result, independent disks are not backed up.

For more information on independent disks and NetBackup, see the following article:

<https://www.veritas.com/docs/000081966>

- This feature does not support VMs that have disks that were excluded from the backup. In the Administration Console, on the NetBackup policy's **Exclude Disks** tab, select **No disks excluded**. Or, in the NetBackup web UI, in the protection plan, clear the option **Exclude selected virtual disks from backups**.
- This feature does not support VMs that have a disk in raw device mapping mode (RDM) or that have a disk in Persistent mode.
- For Windows restore, the ReFS file system is not supported.
- The version of the ESXi server that is used to create a VM using **Instant access virtual machines** must be equal to or newer than the version of the ESXi server that contains the VM backup images.
- For file or folder download with the **Download** option, the NetBackup web UI must be able to access the media server with the same name or IP address that the primary server uses to connect to that media server.
See ["Error when downloading files from an instant access VM"](#) on page 100.
- If the media server appliance uses a third-party certificate, you need to create certain configurations on the NetBackup primary server before you use this feature.
For more information, refer to the "Third-party certificates" and "Implementing third-party SSL certificates" sections in the [NetBackup Appliance Security Guide](#).
- This feature does not support restore of multiple files or folders, which are located in different volumes, partitions, or disks.
- Use the Windows administrator account credentials when you restore multiple files or folders to a Windows VM. You must be logged on to the target Windows VM with these account credentials.

- Some ACL entries are not in the restored file because ACL entries for these users or groups cannot be restored. For example, TrustedInstallers, All Application Packages.
- The Instant Access feature does not support a Windows 10 compact operating system. To verify if your operating system is compressed, run `compact "/compactos:query"` on the command prompt before backing up your VM. To disable the compression, run `compact /compactos:never` on the command prompt before backing up your VM. You can then use the Instant Access feature for your VM backups.
- To restore files and folders, the target VM must be in a normal state, and not in a sleep or hibernate mode.
- A 5-minutes-alive-session threshold is defined in Appliance and BYO web server NGINX. The files and folders that are selected for download must be compressed and downloaded within this threshold.
- To create an instant access virtual machine, you must have read and write access to the VMware data center where the virtual machine is created.
- To ensure that Instant Access works effectively after the storage server and primary server are upgraded from an earlier NetBackup version, restart the NetBackup Web Service on the upgraded primary server with the following commands:
 - `/usr/opensv/netbackup/bin/nbwmc stop`
 - `/usr/opensv/netbackup/bin/nbwmc start`
- If you have to download or restore files or folders from a Windows VM, ensure that the number of Windows registry hives are less than 10000. More information is available about [registry hives](#).
- An image cannot be deleted if an instant access VM is created from it. The instant access feature uses data from a backup image. If the image is expired, the data might be unavailable and the instant access VM may face data loss. After the instance access VM is deleted, the image can be expired.
- The instant access feature does not support hard links. If you create a universal share from an image and the image has hard link files, `vpfsd` shows these hard link files as having 0 bytes size.

Instant access Build Your Own (BYO)

You can build your own VMs (with Red Hat enterprise operating system) to support VMware instant access. You can use the following features:

- Create instant access VMs.
- VMware vMotion.
- Download files and folders.
- Restore files and folders.

To use instant access with a BYO VM created with an earlier NetBackup release, you must upgrade to NetBackup 8.3.

Prerequisites of Instant Access Build Your Own (BYO)

Prerequisites (fresh install and upgrade):

- The BYO storage server with Red Hat Enterprise Linux 7.6 and later, same as the NetBackup Appliance operating system version.
- The BYO storage server with docker installed.
 - The docker version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (RHEL extra).
 - The docker application is included in the environment path.
- The BYO storage server with NFS service installed.
- The BYO storage server with NGINX version installed.
 - The NGINX version must be same as the one in the corresponding official RHEL version release. You need to install it from the corresponding RHEL yum source (epel).
 - Ensure that the `polycoreutils` and `polycoreutils-python` packages are installed from the same RHEL yum source (RHEL server) and then run the following commands:
 - `semanage port -a -t http_port_t -p tcp 10087`
 - `setsebool -P httpd_can_network_connect 1`
 - Ensure that the `/mnt` folder on the storage server is not mounted by any mount points directly. Mount points should be mounted to its subfolders.
 - Enable the logrotate permission in selinux using the following command:
`semanage permissive -a logrotate_t`
- For BYO, docker container is used to browse VMDK files. Data related to the container is stored at the following location: `/var/lib/` and requires minimum 20 GB free space.

Hardware configuration requirement of Instant Access Build Your Own (BYO)

Table 4-1 Hardware configuration requirement

CPU	Memory	Disk
■ Minimum 2.2-GHz clock rate. ■ 64-bit processor. ■ Minimum 4 cores; 8 cores recommended. For 64 TBs of storage, the Intel x86-64 architecture requires eight cores. ■ Enable the VT-X option in the CPU configuration.	■ 16 GB (For 8 TBs to 32 TBs of storage - 1GB RAM for 1TB of storage). ■ 32 GBs of RAM for more than 32 TBs storage. ■ An additional 500MB of RAM for each live mount.	Disk size depends on the size of your backup. Refer to the hardware requirements for NetBackup and Media Server Deduplication Pool (MSDP).

Frequently asked questions

Here are some frequently asked questions for instant access Build Your Own (BYO).

Table 4-2 Frequently asked questions

Frequently asked question	Answer
How can I enable instant access file browsing (for file download and restore) on BYO after the storage is configured or upgraded without the docker installed?	Perform the steps in the following order: 1 Install the required docker version. 2 Start using the Instant Access feature. For example, you can download files, restore files, and so on.
How can I enable the VMware instant access feature on BYO after storage is configured or upgraded without the nginx service installed?	Perform the steps in the following order: 1 Install the required nginx service version. 2 Ensure that the new BYO nginx configuration entry: <code>/etc/nginx/conf.d/byo.conf</code> is part of the HTTP section of the original: <code>/etc/nginx/nginx.conf</code> file. 3 Run the command: <code>/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Table 4-2 Frequently asked questions (*continued*)

Frequently asked question	Answer
How can I resolve the following issue in the <code>vpfs-config.log</code> file that is raised from: Verifying that the MSDP REST API is available via https on port 10087	Perform the steps in the following order: 1 Install the <code>polycoreutils</code> and <code>polycoreutils-python</code> packages through yum tool. 2 Add the following rules that SELinux requires for Nginx to bind on the 10087 port. ■ <code>semanage port -a -t http_port_t -p tcp 10087</code> ■ <code>setsebool -P httpd_can_network_connect 1</code> 3 Run the following command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>
Instant Access for BYO uses a self-signed certificate by default and only supports *.pem external certificate. How do I replace it with a certificate signed by external CA (*.pem certificate), if required?	To configure the external certificate, perform the following steps. If the new certificate is already generated (the certificate must contain long and short host names for the media server), go to step 4. 1 Create the RSA public or private key pair. 2 Create a certificate signing request (CSR). The certificate must contain long and short host names for the media server. 3 The External Certificate Authority creates the certificate. 4 Replace <code><PDDE Storage Path>/spws/var/keys/spws.cert</code> with the certificate and replace <code><PDDE Storage Path>/spws/var/keys/spws.key</code> with the private key. 5 Run the following command to reload the certificate: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Table 4-2 Frequently asked questions (*continued*)

Frequently asked question	Answer
How can I disable media automount for the instant access livemount share in gnome? If the automount is enabled, the source folder is mounted from the livemount share in gnome and smaller disks appear. In this scenario, the instant access feature does not work properly. The mounted disk content source is from the <code>.../meta_bdev_dir/...</code> folder under livemount share, while the mount target is in the <code>/run/media/...</code> folder.	Follow the guideline to disable the gnome automount: https://access.redhat.com/solutions/20107
How can I resolve the following issue in the <code>/var/log/vpfs/vpfs-config.log</code> file? <pre>**** Asking the NetBackup Webservice to trust the MSDP webserver (spws) **** /usr/openv/netbackup/bin/nblistcurlcmd failed (1):</pre>	Perform the steps in the following order: 1 Ensure that your NetBackup primary server is up and there is no firewall blocking the connection between the NetBackup primary server and storage server. 2 Run the following command on storage server to verify the connection status: <code>/usr/openv/netbackup/bin/bpclntcmd -pn</code> 3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command: <code>/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

VM malware scan

You can create a malware scan livemount from a NetBackup image with NetBackup Recovery API. The livemount exports all VM files and folders via NFS or SMB protocol instantly, which allows NFS or SMB client mount the export path and do malware scan on the exported VM files and folders.

This feature provides the following malware scan APIs:

- POST
`/recovery/workloads/vmware/malware-scan-mounts`
- GET

/recovery/workloads/vmware/malware-scan-mounts

- GET
/recovery/workloads/vmware/malware-scan-mounts/{mountId}
- DELETE
/recovery/workloads/vmware/malware-scan-mounts/{mountId}.

For more details, refer NetBackup 10.0.1 API Reference on SORT

Instant rollback

This chapter includes the following topics:

- [Prerequisites of instant rollback](#)
- [Things to consider before you use the instant rollback feature](#)
- [Instant rollback from a VM backup image](#)

Prerequisites of instant rollback

The prerequisites for Instant Access Build Your Own (BYO) are also applicable to the Instant Rollback feature.

See [“Prerequisites of Instant Access Build Your Own \(BYO\)”](#) on page 41.

For NetBackup FlexScale, the software packages that instant rollback requires are included with the NetBackup FlexScale deployment. For more information, see the *Veritas NetBackup Flex Scale Administrator's Guide*.

Things to consider before you use the instant rollback feature

Note the following about the instant roll back virtual machines feature:

- This feature is supported with backup copies. These copies are created with protection plans (web UI) or classic policies (NetBackup Administration Console).
- This feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Build Your Own (BYO) server, and NetBackup FlexScale.
- This feature does not support backups of VMs that have independent disks. VMware does not support snapshots of independent disks in a VM, either

persistent disks or non-persistent disks. As a result, independent disks are not backed up.

For more information, see the following:

<https://www.veritas.com/docs/000081966>

- This feature does not support VMs that have the disks that were excluded from the backup. In the NetBackup Administration Console, on the policy's **Exclude Disks** tab, select **No disks excluded**. Or, in the NetBackup web UI, in the protection plan, clear the **Exclude selected virtual disks from backups** check box.
- This feature does not support VMs that have a disk in raw device-mapping mode (RDM).
- This feature lets you select a maximum of 100 VMs for rollback at a time. If you select more than 100 VMs the **Roll back instantly** option is not displayed. For example, if you want to rollback 180 VMs, you need create two rollback requests for the same job. One for 100 VMs and the second for 80 VMs.
- In this feature, one instant rollback VM requires one livemount. Each livemount can be retained for one day. So the number of VMs that can support roll back depend on the total number of livemounts available. By default, the livemounts value is set to 200.

You can change this default value from the following location: `storage`

`path/spws/etc/spws.cfg`

MaxAllowedLivemounts=200

For NetBackup FlexScale, the livemounts value is set to 100 by default on each MSDP engine in MSDP cluster.

You can change this default value from the following location for MSDP engine:

`/msdp/data/dp1/pdvol/spws/etc/spws.cfg`

Note: The total livemount number configured in instant rollback, VMware instant access, MSSQL instant access, and universal share must not exceed the **MaxAllowedLivemounts** value.

Instant rollback from a VM backup image

NetBackup 9.1 and later lets you roll back a VM instantly from a backup image. Only backup images that support instant access can support instant rollback.

You can perform instant rollback for multiple VMs. You can also roll back a VM multiple times to any recovery point.

For example, if you have three backup images, B1, B2, and B3, you can first roll back the VM to B1, then to B3, then to B2, and so on.

After the rollback is completed, all data after the selected recovery point is no longer available.

To instantly roll back from a VM backup image

- 1 On the left, click **VMware**.
- 2 To select the backup image, do one of the following:

Click the VM

- 1 Locate the VM and click on it.
- 2 Click the **Recovery points** tab, then click the date on which the backup occurred.

The available images appear in rows with the backup timestamp for each image.
- 3 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.
- 4 On the image or a copy of the image, click **Recover > Roll back instantly**.

Select the check box

- 1 Select the check box corresponding to the VM that you want to roll back and click **Roll back instantly**.

You can select multiple VMs to perform instant rollback.
- 2 Select any one of the roll back options:
 - **Roll back to: Most recent**
NetBackup displays the most recent instant access recovery points in a month.
 - **Roll back to: Before specific date and time**
Select the date and time.
NetBackup displays the most recent instant access recovery points going a month before the selected date and time.

Note: NetBackup displays a warning about malware affected images.
 - Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.
- 3 Click **Roll back**.

Use the **Actions** menu

- 1 Click **Actions > Roll back instantly** corresponding to the VM that you want to roll back.
- 2 Select any one of the roll back options:
 - **Roll back to: Most recent**
NetBackup displays the most recent instant access recovery points in a month.
 - **Roll back to: Before specific date and time**
Select the date and time.
NetBackup displays the most recent instant access recovery points going a month before the selected date and time.
 - Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.
- 3 Click **Roll back**.

- 3 Select the wanted options and then click **Roll back**.

The **Activity monitor** tab displays the status of the rollback.

Continuous data protection

This chapter includes the following topics:

- [CDP terminology](#)
- [CDP architecture](#)
- [About continuous data protection](#)
- [Prerequisites](#)
- [Capacity-based licensing for CDP](#)
- [Steps to configure CDP](#)
- [Removing VMs from the CDP gateway](#)
- [Defining the CDP gateway](#)
- [Sizing considerations](#)
- [Limiting concurrent CDP backup jobs](#)
- [Controlling full sync](#)
- [Monitoring CDP jobs](#)
- [Using accelerators with CDP](#)
- [Recovering CDP protected VMs](#)
- [Some limitations of CDP](#)
- [Troubleshooting for CDP](#)

CDP terminology

The following table describes the concepts and terms that are used in Continuous Data Protection (CDP).


Table 6-1 CDP terminology

Term	Explanation
CDP gateway	CDP configured media server.
VAIO	VMware framework consisting of vSphere APIs for I/O filtering. This framework enables CDP to run filters on ESXi servers and intercept any I/O requests from a guest operating system to a virtual disk.
Full sync	NetBackup fetches a VM's entire data from the ESXi.
OST	Open Storage Technology is a STU supported by NetBackup.
MSDP	Media Server Deduplication Storage Pool is a NetBackup dedupe technology engine to optimize backup storage.
Storage policy	A feature of VMware vSphere that allows administrators to create storage profiles so that the VMs do not need to be individually provisioned and so that management can be automated.
VIB	vSphere Installation Bundle. At a conceptual level a VIB is somewhat similar to a tarball or compressed archive. It is a collection of files packaged into a single archive to facilitate distribution.
nbctd	CDP service (daemon) running on the CDP gateway.
Staging area	A storage location on the CDP gateway where NetBackup temporarily stores IOs received from the ESXi.
Storage quota	Allocated limited storage size for VMs using CDP protection.
Reserved quota	Shared storage between all VMs registered to a CDP gateway.
VADP	VMware VADP is a VMware vStorage API that backs up and restores vSphere virtual machines (VMs).

CDP architecture

The CDP gateway is configured on a NetBackup media server. Once configuration is done, NetBackup starts the `nbcctd` daemon on the CDP gateway. This process services all IOs from ESX and enables other NetBackup components on the gateway to take backup. To backup this data, you also need to configure an MSDP or OST accelerator-based STU. You can configure multiple CDP gateways and MSDP/OST accelerator-based STUs as required. NetBackup REST APIs for CDP are a web API interface to leverage this feature. Refer NetBackup REST APIs Swagger documentation for more information.

Figure 6-1 CDP architecture



About continuous data protection

Continuous data protection (CDP) is a smart way to capture fast copies of backups for the VMware VMs, without stuning the VMs. Using CDP, you can rapidly make recent copies of backups and use NetBackup to retain and restore the backups as required.

Here are some salient features of CDP:

- Completely web UI based protection and recovery of VMware VMs.

- Versatile API-based protection.
- Bring Your Own Device (BYOD): You can use a Red Hat Linux-based NetBackup media server as a CDP gateway.
- Support for ESXi and various datastore types. Refer to the [Software compatibility list](#) for the latest information.
- Accelerator-based backup. Support for accelerator enabled storage like MSDP and OST.
- Support for Instant access. You can boot the VMs from MSDP storage.
- Agentless single file restore from MSDP.
- RBAC support for entire protection and restore workflow.
- Traditional and capacity-based licensing.
- CDP uses Veritas IO filter that is fully compatible with the Veritas Resiliency Platform.

Prerequisites

Prerequisites for using CDP

- CDP for VMware exclusively supports accelerator-based backup. So, CDP needs accelerator-compliant storage units based on MSDP or OST-based storage.
- CDP uses file system as staging area on CDP gateway. See the Software compatibility list for the supported file systems.
- The media server that is associated with MSDP should have NetBackup version 9.1 or higher.
- Capacity based and traditional license for enabling the feature.
- The port 33056 on the CDP gateway, must be open for ESXi server to communicate to CDP gateway.
- VMware server credentials need privileges for NetBackup to start, stop, restart, and refresh the Common Information Model (CIM) service on the ESXi host.
- You can configure a CDP gateway on RHEL-based NetBackup media server platform.
- Create a VMware storage policy for replication using the VAIO component. Attach the storage policy to each disk of the VMs that you want to protect using CDP. For details, see Veritas support knowledge base article on [How to create vtstap storage policy in VMware vCenter](#).

Veritas IO filter for VAIO requirement

You can download and deploy the VAIO drivers package, version 4.0.0, to use with your CDP deployment. Refer to the [Software compatibility list](#) for the latest version and information on how to download it.

You must install the vSphere Installation Bundle (VIB) on the vCenter cluster before configuring protection in NetBackup. Note that you do not need to deploy VIB on vCenter for restore purpose. See Veritas support knowledge base article on *Deploying an IO Filter solution to a cluster using VMware MOB*.

Storage Policy requirements

Before you can deploy CDP, you need to create a VM storage policy. The storage policy must have a component chosen as "Replication" and provider as "vtstap". This policy must be attached to each disk of VM to be protected. Otherwise backup jobs fail. For details, see Veritas support knowledge base article on [How to create vtstap storage policy in VMware vCenter](#)

Note: Detaching the storage policy results in loss of protection for the VM. If you detach the Veritas IO filter storage policy from a VM, IO tapping for the VM is stopped, so the data from this VM does not get saved in the CDP gateway. Hence, the consequent backup jobs remains blank backup jobs, even after all the data from the CDP staging area is moved to the backup storage. So, we recommend to remove protection of the VM(s) from NetBackup protection plan once you detach the vtstap policy from the VM(s).

Capacity-based licensing for CDP

Licensing collects the total number of front-end terabytes protected by NetBackup. The front-end data size for CDP backup is nearly same as consumed storage size on ESX datastore by the VMs.

The `nbdeployutil` utility reports data usage for the VMs. Following rules are applied to report data size:

- Calculate the total number of bytes written during backup (X) and the VM size from ESX datastore (Y). The reported size is the smaller value of X and Y.
- If different policies use the same virtual machine, the policy with higher data size is accounted.
- If VADP and CDP policy protects the same VM then you are charged only once, with the higher size.

Administrator can use the following steps to verify the data size reported by licensing:

- Verify the size occupied by the VMs on ESX datastore on the vCenter. Navigate to **Datastore > Files > VM**, the **Size** column shows the size occupied on datastore.
- Verify the bytes written during backup for same VM.
- Calculate the minimum of the above two values.

Steps to configure CDP

To configure CDP for your workload you must perform the following tasks.

Operations on the VMware vCenter

1. Install the I/O filters by Veritas. See Veritas support knowledge base article on [Deploying an IO Filter solution to a cluster using VMware MOB](#).
2. Attach the storage policy to ESXi. For details, see Veritas support knowledge base article on [How to create vfstap storage policy in VMware vCenter](#)

Operations on the NetBackup console

1. Create an MSDP or OST-based storage for the backup destination. See the information on how to configure storage in the *NetBackup Web UI Administrator's Guide*.
2. Create a CDP gateway.
3. Create a CDP-based protection plan for your VMware workload. See the *Managing protection plans* chapter of the *NetBackup Web UI Administrator's Guide*.
4. Protect the required VMs with the protection plan.
5. Monitor jobs.

Removing VMs from the CDP gateway

When CDP protection is no longer required for a VM, you can remove protections from that VM, or switch the VM to classic policy.

To remove CDP protection from a VM

- 1 Go to vCenter and change the VM's storage policy from `vtstap` to `Datastore default`.
- 2 In the NetBackup web UI, on the left, click **VMware** under **Workloads**, you can see a list of VMs with protection details.
- 3 Click the name of the VM, from which you want to remove protection, in the subsequent page, **Remove protection**.

You can see a confirmation message when the VM is removed.

If you remove protection from a VM without removing the `vtstap` policy from the VM, you can see a partially successful removal message in the UI. These partially removed VMs are not included in the **Total VMs subscribed** count in the **Continuous data protection gateway** tab.

Note: The partially removed VMs are neither protected by CPD nor by classic policy. Also, you cannot re-subscribe the VM to a CDP gateway. Hence, it is recommended to detach the `vtstap` storage policy from the VM, and fully unsubscribe the VM from the CDP gateway.

Defining the CDP gateway

You need to define a gateway for your CDP deployment, before you can protect any VMs. You can define the CDP gateway in a VM that is a NetBackup media or primary server.

Note: Before defining the CDP gateway ensure that your system time is synchronized with the network time.

To define a CDP gateway

- 1 On the left, click **VMware** under **Workloads**.
- 2 On the top right, click **VMware settings**, click **Continuous data protection gateway**.
- 3 Click **Add**. Enter a **Host name** and **Storage path**. The storage path should have independent file system, other than root. Do not share this same location with other applications like MSDP.

- 4
- On the next page, if your gateway version is 9.1. specify the parameter **Maximum number of concurrent jobs.** as described in the table below, and click **Save** to save the gateway.

If your gateway version is 10.0, click **Advanced** to specify the advanced parameters to configure and fine tune your CDP gateway. You can also use this set of parameters to estimate how many VMs you can support using CDP protection for a particular configuration of the gateway.

Parameter	Description
Maximum number of concurrent jobs.	The maximum number of CDP jobs that can run simultaneously in the gateway. A bigger number may indicate increased peak resource consumption
Maximum number of simultaneous initial sync	Number of VMs that can take full backup simultaneously during the initial phase of CDP protection. Specifying a higher value than the default, may cause increased resource consumption and affect existing protection.
Reserved memory for Continuous data protection	Reserved memory for the gateway. Enter a value in GB that is equal to or smaller than 90% of the total physical memory.
Data staging area per VM	Specify storage for each VM.
Reserved staging area	Additional storage area to handle the I/O spikes in the VMs.

- 5
- Click **Estimate the number of VMs** to calculate how many VMs this gateway will be able to support for this given configuration.
- 6
- Click **Save**, to add the gateway.

Sizing considerations

This section describes the sizing requirements of the CDP gateway, based on the workload in your environment.

Note: If the CDP gateway plans to support large number of VMs, it is recommended to deploy the CDP gateway, and the MSDP or media server hosting the storage unit, on different hosts.

Note: If CDP gateway and MSDP are co-located on the same media server, then CDP service consumes 20% of available memory (RAM) for its internal use. If the CDP gateway is standalone on media server, it consumes 50% of available memory for the same. From NetBackup version 10.0 onwards, you can configure this value in the UI.

Gateway sizing

You need to size the CDP based on the number of VMs that you want to protect. Consider the requirements described in this section, while calculating requirements for the gateway.

CDP enables you to continuously tap the IOs done by the VMs. NetBackup, by default, uses 10-GB storage space on the staging area per VM. When IO tapping starts, the CDP service starts writing the data into this 10GB storage. Once this storage limit is reached, the CDP service (nbcctd) triggers a backup job to move this data from the gateway to the backup storage.

Out of the total available space on the CDP staging path, by default, NetBackup reserves 25% for usage beyond allocated storage per VM. This storage is common for the subscribed VMs to the gateway. See “[Defining the CDP gateway](#)” on page 56. , for how to do it on version 10.0 onwards. You can reconfigure this value in the `nbcct.conf` file in NetBackup 9.1.

To configure reserved storage in NetBackup 9.1

- 1 Logon to CDP gateway.
- 2 Navigate to the `<staginglocation>/nbcct/` directory, and open the `nbcct.conf` file in a text editor.
- 3 Enter the required values against the parameters
`CCT_VM_QUOTA_SIZE_IN_MB` and
`CCT_VM_QUOTA_RESERVE_PERCENT`
- 4 Restart the `nbcctd` service.

Storage requirement for the gateway

When NetBackup receives the data from the ESXi IO daemon, it stores the data in the in-memory cache. Recommended is minimum 160 MB of data for each VM.

For example, you protect 40 VMs in a gateway. So, you need $40 \times 160 \text{ MB} = 6400 \text{ MB}$ RAM. Allocating more RAM increases the in-memory cache size when CDP service starts, ultimately increasing the IO performance of the service.

Similarly, to stage $40 \times 10\text{-GB} = 400\text{-GB}$ (75%) + 134GB (25%) reserved, that is approximately 540 GB space you need to have on the staging area.

Increasing per VM storage allows to NetBackup to backup more data per backup job. Increasing reserved storage for the CDP gateway lets you receive more data without any interruption to the protection. Note that even when the staging path is fully occupied, it does not affect the applications inside the VM. NetBackup catches up the data produced by applications during that time, moves it to the backup storage in the subsequent backup jobs.

Note: If NFS is used for the staging area, minimum required throughput is 100 MB/sec.

First 24-hours experience

When you start using the CDP feature, it is important to observe the system and tune according to your business demand, add hardware configuration to maximize the protection and performance. First, you can use default values and start subscribing the VMs according to the requirements mentioned in this section. You should check the following:

- Number of immediate backup jobs that the CDP service triggers due to staging storage full condition.
- You can check the CDP backup engine notifications on NetBackup web UI.
- Underlying provisioned storage performance. Like the NetBackup installation disk, CDP staging area, and MSDP storage disks.
- Network utilization and available bandwidth.
- CPU and memory consumption when receiving data from the ESXi, and when the backup jobs are running.

Note: If you observe slow IOs from the IO daemon, check network bandwidth and system RAM. See [“Defining the CDP gateway”](#) on page 56. , for how to increase the in-memory cache size in NetBackup 10.0 onwards. For NetBackup 9.1, you can do it using the `CCT_POOL_SIZE_QUOTA_PERCENTAGE` parameter in the `nbcct.conf` file.

Limiting concurrent CDP backup jobs

You can set a limit for the simultaneous CDP snapshot jobs that can run in the CDP gateway at a time. For example, if you protect 20 VMs, and you have set a limit of 5, then only 5 VMs can run simultaneous backups, and 15 VMs will be in queue. This setting is required for optimized use of your system and network resources. By default, the resource limit value is 0, representing no limit.

See [“Defining the CDP gateway”](#) on page 56. for information on how to do it on NetBackup version 10.0 onwards. For NetBackup 9.1 follow the procedure described below.

To set value to resource limit, we have the following API:

```
POST /config/resource-limits

{
  "data": [
    {
      "type": "resource-limits",
      "id": "string",
      "attributes": {
        "resources": [
          {
            "resourceType": "string",
            "resourceName": "string",
            "resourceLimit": 0,
            "additionalData": "string"
          }
        ]
      }
    }
  ]
}
```

Here,

- `Id` represents the workload that is `Cdp`
- `resourceType` should be `Cdp-Backup`
- `resourceName` represents the CDP gateway host name. It should be same as specified in the protection plan. If you keep an empty string for `resourceName`, the `resourceLimit` value is set as a global limit, which is applicable to all the configured CDP gateways.
- The `resourceLimit` value sets the value of backup jobs for that gateway.

To retrieve the list of resource limits for a CDP workload type, use:

```
GET - /config/resource-limits/cdp
```

To update the value of `resourceLimit` for particular gateway, hit POST API with change in `resourceLimit` for the same record.

To delete the specified granular resource limits, use:

```
DELETE - /config/resource-limits
```

Only the resource limit set for a particular resource can be deleted. Provide both the resource type and the specific resource of that type.

Controlling full sync

When you subscribe a VM to a CDP enabled protection plan, NetBackup initiates full sync, to get the entire data of the newly protected VM. For a newly subscribed VM, NetBackup does not have any data to apply the incremental backup features, hence full sync is initiated. During a full sync, NetBackup captures the entire data of the VM, from the underlying VMDKs to the CDP staging location, and subsequently to the NetBackup STUs.

Full sync is normally triggered when you subscribe a new VM to a CDP enabled protection plan, but in certain scenarios, you can manually trigger a full sync:

- Accidental corruption or deletion: CDP maintains backed up data of the VMs at the staging location in proprietary format files. If these files for a VM are accidentally deleted or corrupted, the subsequent backup job for the VM fails citing data integrity mismatch. In this case, you can initiate a force rescan schedule backup, and subsequently a full sync of the VM takes place.
- Following a manually triggered force-rescan schedule.
- CDP service can trigger full sync to receive VM data whenever necessary.

During full sync, data flows from the ESXi to the CDP gateway. Depending on the data size of the VMs, the volume of this data can be substantially large that can consume a lot of resources like network, memory, processing power, and storage. This also affects the backup operations of the VMs subscribed earlier.

If you subscribe more than 5 VMs at a time, say 7, then, full sync is initiated for 5 VMs, and 2 are in wait state.

Therefore, it is recommended to limit the number of concurrent full sync operations to optimize system resources. The default number of concurrent full sync is 5. This allows 5 VMs to perform full sync concurrently. Other VMs needing full sync need to wait in a queue. This way, the system resources are managed optimally.

Recommendation for controlling full sync:

- Subscribe the VMs in batches of five or less.
- Once a subscribed VM completes full sync, you can see message in the UI, then you can proceed to subscribe the next batch.

Configuring full sync

See [“Defining the CDP gateway ”](#) on page 56. for information on how to configure full sync on NetBackup version 10.0 onwards.

In NetBackup 9.1, you can configure the number of concurrent full sync operations by specifying a value for the `CCT_MAX_FULL_SYNC_REQS` parameter, in the `nbccct.conf` file. For example, `CCT_MAX_FULL_SYNC_REQS=7`

Monitoring CDP jobs

More information is available on monitoring jobs in the web UI.

[NetBackup dashboard](#)

CDP follows the same job hierarchy as the traditional NetBackup agent for VMware. Protection starts with the job discovering the VM and its attributes. A child job called Preparing for Backup follows it. This child job determines the changed blocks based on previous images and current data available on gateway. A backup job to move data from CDP gateway to destination storage unit, follows the child job.

If there is not enough space for each VM, on the gateway, the backup image may not be fully recoverable. Such images are referred to as partial non-recoverable images and are not available to restore from the web UI. But the subsequent backup jobs, create recoverable backup images. If an image is non-recoverable, NetBackup triggers a backup job automatically when it receives consistent data from ESXi.

Viewing notifications

For most CDP activities, you can see notifications in the web UI. These notifications are helpful to know how the IO tapping on the gateway performs. You can see notifications when things have stopped working or any action is required from your side. The following are some important scenarios when you can see notifications:

- While backing up data. When a backup job moves data from staging area to back up storage.
- VM full sync has started/suspended/resumed/done.
- Partial image is generated.
- No space left in the staging area storage.
- When there is an error while writing in-memory data to staging area location.

Here are some notifications:

Table 6-2 Viewing notifications

Message	Scenario	Severity	Priority
Temporarily disconnecting from the IO filter to the Continuous data protection service on the gateway. Either the allocated staging area is almost full, or the memory usage is at maximum.	The staging space allocated to CDP is almost full, and CDP service temporarily disconnects from the IO filter. This may also happen, if backup jobs are not able to move data from the CDP gateway staging database to the backup storage. Check backup job failure reasons and STU's underlying storage.	Critical	High
Input/Output error occurred for the VM: <uuid>	CDP service is not able to perform IO on staging location due to myriad of reasons like, underlying disk snapped out of storage, or file-system went into read-only mode, and so on.	Error	High
Terminating the Continuous data protection service, as the staging area memory is full.	If the staging space is less than 1 GB, CDP raises this error and terminates the service.	Critical	High
Data storage quota full for the VM: <uuid>, bearing jobid: \${jobid}. Moving data to backup storage.	During VM's data transfer, if the total data crosses the configured VM quota, then a backup job is triggered to move the staging data to backup destination.	Info	Low
Cannot move data to backup storage, for the VM: <uuid>. Storage quota for the VM is full.	Data movement from the gateway to the backup location failed.	Error	High
Full sync started for the VM: <uuid>.	Initiated the full sync process for this VM.	Info	Low
Full sync resumed for the VM: <uuid>.	Full sync for the VM is resumed after some unexpected interruption.	Info	Low
Full sync completed for the VM: <uuid>.	The initial full sync for VM is complete.	Info	Low

Table 6-2 Viewing notifications (*continued*)

Message	Scenario	Severity	Priority
Full sync suspended for the VM: <uuid>.	Full sync operation fails, for some reason like, network glitch.	Info	Low
Backup image generated for the VM: <uuid> is not recoverable.	When a VM sync is in progress, if the VM quota is reached, a backup job is triggered. When the backup job is completed the image may not be recoverable, as NetBackup is moving the intermediate data generated on the guest VM.	Info	Low

Viewing jobs

CDP uses the activity monitor to display the following job information:

- Parent backup job - discovery job to discover the VM information.
- Preparing for backup - identify the point in time data for the VM.
- Backup - move data from the staging path to the backup storage.

Using accelerators with CDP

CDP for VMware exclusively supports accelerator-based backup. So, CDP needs accelerator-compliant storage units based on MSDP or OST-based storage.

Force rescan

Force rescan enhances safety, and establishes a baseline for the next accelerator backup. This feature protects against any potential damage like failure of checksum verification on the data in the staging area.

When you use accelerator-based forced rescan, it clears the data on CDP gateway staging area. So, any corrupted data is replaced with fresh data synced from the ESXi server. Note that the first backup job triggered by forced rescan may not have all data needed for a recoverable image. As data becomes available, the subsequent backups are triggered automatically making the images recoverable.

Recommendations for using forced rescan:

- Do not trigger force rescan for the VMs which are turned off.
- If the staging location memory is full, you can see a notification in the UI. Initiate the force rescan only when sufficient memory is available at the staging location.

To manually trigger the backup with force rescan execute the following command in the command prompt or the Linux terminal:

```
bpbbackup -i -p policyname -s <schedulename>
```

NetBackup creates a schedule named `ForcedRescan` for every protected VM.

Recovering CDP protected VMs

VMs protected by NetBackup CDP for VMware have same backup image format as the NetBackup agent for VMware. So, all recovery operations are same as the NetBackup agent for VMware.

Here are some minor differences:

- Agentless single file recovery is supported only if MSDP is configured for instant access.
- Recovery from the vCenter plug-in is not supported.
- Cannot restore VMs from CDP-based backup images through Java UI.

Web UI does not allow recovery of the images shown as partial and non-recoverable. You can restore them using NetBackup API. However, the VMs may not boot after the recovery.

Some limitations of CDP

Here are some limitations of CDP:

- NetBackup features like Intelligent policy and Backup now, and Roll back instantly from web UI are not supported.
- CDP for VMware and Veritas Resiliency Platform does not work together for the same VM. However, both products can protect different VMs on the same vCenter cluster.
- CDP does not support any standalone ESX, which is not managed by any VC. An ESXi which is not part of any ESXi cluster but is managed by VC, is also not supported.
- You must turn on the VMs before subscribing them to a CDP-based protection plan, and also for the first full backup.
- Do not protect the same VM using both CDP and VMware vStorage APIs for Data Protection (VADP) policies. If you protect a VM using CDP policy, then the VADP backup for that VM may fail.

- After subscribing a VM for CDP backup policy, if any disk from the VM is removed or a new disk is added, the subsequent backups fail. In such cases, unsubscribe the VM from CDP protection, and subscribe it again.
- Due to VMware limitation, if you try to protect a VM using the NetBackup agent for VMware and CDP, both at the same time, backup operation fails with error or the operation might crash with symbols from VDDK.

Troubleshooting for CDP

VAIO stops sending data to CDP gateway

Happens when the IOFilter encounters problem and hence enters into NOOP (Non-Operational) mode.

Possible reasons:

- IOfilter encountered problem with datastore.
- IOfilter encountered problem while reading from vm disk on ESXi server.

Workaround:

Remove the VTSTAP policy from all the disks of protected the VMs and reattach.

Error: Storage policy is not detached from one or more virtual disks of virtual machine.

Happens when the storage policy is not detached from all the virtual disks of the VM. The next backups fail with error code 156.

Workaround:

Remove the Veritas I/O filter based storage (vtstap) policy from all the disks the VM that CDP protected previously. You can do this operation on the vCenter.

Error: Failed to retrieve or parse the version of Veritas IO filter.

You may get this error when trying to subscribe one or more VMs to CDP protection plan. Occurs when the CIM server service on ESXi server is non-responsive.

Workaround:

Restart the CIM server service on the ESXi server and retry the VM subscription to CDP protection plan. You can find the CIM server service of ESXi server, under Configure > Services section of the ESXi.

nbcctd service goes in inconsistent state. Cannot configure the CDP gateway.

Possible reasons:

- When you mount a read-only file system and provide its path in the CDP gateway configuration, service is configured, but the gateway fails to start.
- When you try to configure the gateway again, by giving a read/write path, the service still fails to start.

Workaround: Retry the operation after you remove the `nbcct` directory from:

- `<NBU installation path>/netbackup/nbcct` in NetBackup 9.1.
- `<staginglocation>/nbcct` in NetBackup 10.0 onwards.

CDP-based protection plan fails with the error: Storage policy is not attached to one or more virtual disks of virtual machine to be registered for IO tapping.

Possible reasons:

Currently, NetBackup supports only `vtstap` policy as storage policy for CDP. If you try to subscribe a VM using hybrid storage policy (encryption + replication) it shows the error.

Workaround: Avoid using hybrid storage policy (encryption + replication) for CDP protected VMs.

CDP service is not getting started after media server restart or mount path-related changes.

Possible reasons:

The configured staging area is unmounted post reboot or having an unsupported file system. For example, if you configure the CDP gateway using a supported mount like `/mnt/stage_area` and do not configure auto mount. After a system restart, this path points to root file system, which CDP does not support, hence the CDP service (`nbcctd`) cannot start.

Workaround: Ensure that the staging area or the relevant disk mounts are remounted properly, whenever there are changes in the system related to unmount or system reboot.

VM gets unsubscribed in powered off state and having I/O tapping policies attached to the VMDK. It should give warning to remove storage policies and then unsubscribe.

Possible reasons:

While removing CDP protection, if the protected VM is powered off, CDP gateway cannot get the required information of storage policies from VAIO. Hence, though the CDP protection is removed from the VM, the I/O tapping policies are still attached to the VMDK of that VM, it continues to tap the I/O'S and affect performance.

Workaround: Always detach the storage policy of the VMs before unsubscribing the VMs, irrespective of its powered on or off state.

Subscription to NetBackup protection plan fails, but the backup jobs keep dumping data in the staging area.

Explanation

Occurs when you protect a NetBackup primary server, using the same primary server's protection plan.

Workaround: We do not recommend protecting a NetBackup primary server using a protection plan made using the same primary server. If this error occurs, detach the Veritas storage policy from the NetBackup primary server VM, and unsubscribe the VM from the protection plan.

Cannot delete CDP protection plan when the CDP gateway is unreachable.

Explanation:

CDP policy is not deleted after removing the entries in case of an unreachable host.

Workaround: The CDP protection plan subscription does not get removed as we are not deleting the CDP policy before cleaning up the CDP host. So, we need to call Delete policy API manually after calling the Delete CDP gateway API, to delete the entries of the unreachable gateway.

You can clean up an unreachable CDP gateway using the following API:

To DELETE CDP Gateway

URL : `https://netbackup/config/cdp-gateway/force`

HTTP Method : DELETE

Headers:

Authorisation: Bearer <Token>

Content-Type: `application/vnd.netbackup+json;version=7.0; charset=UTF-8`

To Delete Policy

URL : `https://netbackup/config/policies/policy_name`

HTTP Method : DELETE

Headers:

Authorisation: Bearer <Token>

After successful execution of above two APIs, the mapping for the policy and the VM is still visible in the web UI. If you try to remove protection of that VM through web UI, you can see an error message saying: **Subscription ID not found**. This is expected behavior.

CDP gateway update operation fails to restart the CDP service (nbcctd) on the gateway

Explanation: CDP gateway update operation tries to restart the service. If stopping the service takes longer time than usual, then update operation shows an error, indicating that the CDP service (nbcctd) failed to restart.

Workaround: In this case, check if the `nbcctd` service is running on the gateway. If the service is running, wait for it to shut down. To manually stop the service, use the command: `/usr/opensv/netbackup/bin/nbcctd -terminate`. When the service has stopped, start it using the command `/usr/opensv/netbackup/bin/nbcctd -X`.

VM recovery

This chapter includes the following topics:

- [Recover a VM](#)
- [About VMware agentless restore](#)
- [Prerequisites and limitations of VMware agentless restores](#)
- [Recover files and folders with VMware agentless restore](#)
- [About restricted restore mode](#)
- [Recover VMware Cloud Director virtual machines](#)

Recover a VM

You can recover a VM to its original location where it existed when it was backed up or to different location. You can choose to recover from the default copy of the backup image or from an alternate copy, if one exists. The default copy is also known as the primary copy.

To recover a VM

- 1 On the left, click **VMware**.
- 2 Locate and click on the VM.
- 3 Click the **Recovery points** tab. In the calendar view on the left, select the date on which the backup occurred.

The available images are listed in rows with the backup timestamp for each image.

- 4 Select the **Allow the selection of recovery of points that are malware-affected** to be able to recover malware affected images. This option is only available for the recovery points which contains malware-affected images.

Note: This options is enabled only for users with required permissions.

- 5 On the image that you want to recover, select one of the following image recovery options:
 - **Recover**
 Recover from the default copy of the backup image. This option is displayed if only one copy exists.
 - **Recover from default copy**
 Recover from the default copy of the backup image. This option is displayed if more than one copy exists.
 - **nn copies**
 Recover from the default copy or a different copy of the backup image. NetBackup allows up to ten copies of the same backup image. All available copies are displayed when you select this option. For each copy, the **Storage** name, **Storage Server**, and the **Storage server type** are displayed. Click **Recover** for the copy that you want to recover.
- 6 Choose the type of recovery that you want to perform:
 - **Restore virtual machine:** Recover the backup image to the original location or to an alternate location.
 - **Create instant access virtual machine:** Recovers the backup image to a new instant access virtual machine. This option is available only if the backup image has instant access capability.
 See [“Create an instant access VM”](#) on page 33.
 - **Download files and folders:** Downloads the files and folders from a VM backup image. This option is available only if the backup image has instant access capability.
 See [“Download files and folders from a VM backup image”](#) on page 37.
 - **Restore files and folders:** Restores the files and folders from a VM backup image. This option is available only if the backup image has instant access capability.
 See [“Restore files and folders from a VM backup image”](#) on page 35.
- 7 Review the **Restore to** values.

The default values come from the backup image of the VM.

- To recover to the original location, click **Next**.
- To recover to an alternate location, change the restore values. Then click **Next**.

8 Review or change the **Options**.

See [“Recovery options”](#) on page 72.

9 Review or change the **Advanced** options.

See [“Advanced recovery options”](#) on page 72.

See [“Advanced recovery options: Format of restored virtual disks”](#) on page 73.

See [“Advanced recovery options: Transport mode”](#) on page 74.

10 Click **Pre-recovery check**.

NetBackup verifies the credentials and appropriate paths and connectivity, determines whether the datastore or datastore cluster has available space, and reviews, other requirements.

11 Resolve any errors.

You can choose to ignore the errors. However, the recovery may fail.

12 Click **Start recovery**.

Click the **Restore Activity** tab to monitor a job's progress. Select a specific job to view its details.

Recovery options

Allow overwrite of existing virtual machine	NetBackup deletes any VM with the same display name that exists at the destination, before starting recovery. Note that, NetBackup deletes any VM with the same display name, it may not be the same VM, but another VM having the same display name .
Power on after recovery	Automatically powers on the VM when the recovery is complete.
Recovery host	Indicate the host that you want to use to perform the recovery. By default, the recovery host is the one that performed the backup.

Advanced recovery options

Create a new BIOS UUID	Restores the VM with a new BIOS UUID instead of the original BIOS UUID.
-------------------------------	---

Create a new instance UUID	Restores the VM with a new instance UUID instead of the original instance UUID.
Remove backing information for devices	For example, this option restores the VM without restoring any ISO file that was mounted when the VM was backed up. If this option is disabled, the recovery might fail if the backing information is not longer available for devices, such as DVD/CD-ROM drives, or serial or parallel ports.
Remove original network configuration	Removes the NIC cards from the VM. Note that for network access, the restored VM requires network configuration. Enable this option if: ■ The network connections on the destination virtual machine have changed since the backup was made. ■ The original virtual machine still exists and a duplicate VM may cause conflicts.
Retain original hardware version	Restores the VM with its original hardware version (such as 4). It retains the original version even if the target ESXi server by default uses a different hardware version (such as 7 or 8). If the target ESXi server does not support the virtual machine's hardware version, the restore may fail. If this option is disabled, the restored virtual machine is converted to the default hardware version that the ESXi server uses.

Advanced recovery options: Format of restored virtual disks

Original provisioning	Restores the VM's virtual disks with their original provisioning.
Thick provisioning lazy zeroed	Configures the restored virtual disks in the thick format. The virtual disk space is allocated when the disk is created. This option restores the populated blocks, but initializes vacant blocks with zeros later, on demand. Note: If the vmdk is completely written, VMware automatically converts a lazy-zeroed disk to Thick provisioning eager zeroed.
Thick provisioning eager zeroed	Configures the restored virtual disks in the thick format. Restores the populated blocks and immediately initializes vacant blocks with zeros (eager zeroed). Creation of the virtual disks may take more time with this option. However, if the restore occurs over a SAN, the eager zeroed feature may speed up the restore by reducing network communication with the vCenter server.

- Thin provisioning** Configures the restored virtual disks in the thin format. Restores the populated blocks but does not initialize vacant blocks or commit them. Thin provisioning saves disk space through dynamic growth of the vmdk file. The vmdk files are no larger than the space that the data on the virtual machine requires. The virtual disks automatically increase in size as needed.
- Note:** If the vmdk is completely written, VMware automatically converts a thin disk to **Thick provisioning eager zeroed**.

Advanced recovery options: Transport mode

The **Transport mode** specifies the mode to use for backups or how to read the data from the datastore. For more information on transport modes, see the vendor documentation for your virtualization environment.

Note the following when you select a transport mode:

- The SAN mode is not supported for the virtual machines that use VMware Virtual Volumes (VVols).
- For the hotadd mode, the virtual machines that use VVols and the backup host (hotadd) virtual machine must reside on same VVol datastore. For more information about the hotadd transport mode, see the [NetBackup for VMware Administrator's Guide](#).

About VMware agentless restore

NetBackup 8.2 and later supports VMware agentless restore. The agentless restore lets you restore individual files and folders to virtual machines where the NetBackup client is not installed. By using VxUpdate, NetBackup can deploy the recovery tool to the virtual machines, restore files and folders, and perform the required cleanup. NetBackup does not require a connection to the target virtual machine to recover the files. All recovery is handled through the ESX server using VMware vSphere Management APIs.

A video is available that describes NetBackup VMware agentless restore:

[VMware agentless recovery video](#)

Overview of the agentless restore process

- 1 The NetBackup primary server receives input from either the NetBackup web UI or the Agentless Recovery API. The input is the files and folders for restore along with the credentials for the target virtual machine. These credentials must have administrator or root privileges.
- 2 The primary server sends the requested data to the restore host.

- 3 The restore host confirms that it has the necessary VxUpdate recovery package to perform restore. If it's not available, the restore host downloads the required package from the primary server using VxUpdate.
- 4 The restore host pushes recovery tool to virtual machine using the vSphere management API.
- 5 The data stream containing the user-selected files and folders is staged in a vmdk that is associated with a temporary virtual machine. Veritas creates the temporary virtual machine for the agentless restore.
- 6 The vmdk that NetBackup created on the temporary virtual machine is attached to the target virtual machine.
- 7 The recovery tool is invoked and the files and folders are recovered.
- 8 NetBackup performs the necessary cleanup. All temporary files and objects that are created as part of the process are deleted or removed. Among the objects that are deleted and removed are the recovery tool, the temporary virtual machine, and the staging vmdk.
- 9 The job is finished.

Prerequisites and limitations of VMware agentless restores

Prerequisites

The following prerequisites exist for VMware agentless restores:

- You must provision VxUpdate packages for all platforms for which you have virtual machines where you want to perform agentless recovery.
- You must have credentials with administrator or root permissions for the target virtual machine.
- The target VM is where the files are recovered. It must be powered on and have the latest version of VMware Tools installed.
- The target VM should have at least one Paravirtual Controller with available LUNs. Or, available space for a Paravirtual SCSI Controller.
- The default staging location on the target VM is %TEMP% or %TMP% for Windows and the root directory (/) for Linux.
- The staging location must exist on the target VM file system.

Limitations

The following limitations exist for VMware agentless restores:

- Agentless restores to Windows target VMs can fail if you use an account other than the built-in **Administrator for Windows Guest OS** account as the **Target VM Credentials**. The restore fails because **Run all administrators in Admin Approval Mode** is enabled. More information is available:
https://www.veritas.com/content/support/en_US/article.100046138.html
- VMware agentless restores can only be used for the restore of files and folders.
- In some instances, when you perform an agentless restores, orphaned VMs starting with `NB_` are left behind. Using the ESX server credentials to perform the restore on the target VM even though the vCenter manages the ESX server can cause this condition. This condition is a known limitation of VMware. To resolve the problem, register the vCenter in NetBackup and use vCenter credentials for backups and restores. The orphaned VMs starting with `NB_` can be removed from inventory manually by logging into the vCenter using VMware vSphere Client.
- Restore job fails if NetBackup is unable to use the staging directory. This directory is specified in the `TMP` or `TEMP` environment variable.
- Restore job fails if NetBackup does not have sufficient privileges to the staging directory. Or, if there is insufficient space in the staging directory.
- If you select **Flatten existing directory structure** and **Overwrite existing files** options, you risk an incorrect restore if it contains multiple files with the same file name. In this case, the last file that is restored is the one that is present when the restore completes.
 If you select **Flatten existing directory structure** and you do not select **Overwrite existing files**, the restore succeeds. The first file that is restored is present when the restore completes. To prevent this issue, do not select **Flatten existing directory structure** when restoring multiple files with the same name.
- The **Flatten existing directory structure** and **Append string to file names** options are only applicable to files. They are not available for directories.
- Multiple restore jobs to the same VM are not supported. The user must start another job as needed for that VM once the first restore job for that VM has completed.
- If a backup and a restore occur simultaneously on the same VM, one or both jobs can have unexpected results. If a backup or a restore exits with a non-zero status code, one possible cause is simultaneous jobs occurring on the same VM.
- Veritas does not recommend VMware agentless restore if a NetBackup client already exists on the target VM. The NetBackup administrator must use the agent-based restore in such cases.

- For the current list of guest operating systems that NetBackup supports for the target VM, see *Supported guest operating systems for VMware* in the following document:
[Support for NetBackup in virtual environments](#)

Recover files and folders with VMware agentless restore

To restore VMware files and folders using agentless restore

- 1 Confirm the target VM is powered on.
- 2 On the left, click **Workloads > VMware**.
- 3 Locate and click on the VM that contains the files and folders for restore.
- 4 Click the **Recovery points** tab. In the calendar view, click the date on which the backup occurred.

The available images are listed in rows with the backup timestamp for each image.
- 5 On the image you want to recover from, click **Restore files and folders**.
- 6 Specify the files and folders you want recovered. Then click **Next**.
- 7 Specify the target VM to which you want the files and folders recovered, as well as the administrator credentials for the target VM. Click **Next**.
- 8 Click **Next**.
- 9 Specify additional recovery options for the restored files and folders. Click **Next**.

NetBackup performs a pre-recovery check using the options you specified.
- 10 **Review** displays the status of the pre-recovery check along with the options you selected for the recovery. Once you confirm that they are correct, click **Start recovery**.

About restricted restore mode

The restricted restore mode option is a form of VMware agentless restore for restricted environments such as Windows User Account Control (UAC). The user-selected files are first staged to the recovery host and then restored to the virtual machine. The recovery host must have sufficient space for staging.

The default staging location on the recovery host is

`install_path\VERITAS\NetBackup\var\temp\staging`. NetBackup creates this directory with the correct permissions the first time it is accessed. You can change

the staging location with the `AGENTLESS_RHOST_STAGING_PATH` registry setting on the recovery host. This `REG_SZ` registry key does not exist by default. It must be created in

`HKEY_LOCAL_MACHINE\SOFTWARE\VERITAS\NetBackup\CurrentVersion\Config.`

If you change the staging location, Veritas recommends that you let NetBackup create the staging directory. When you let NetBackup create the directory, the permissions are set correctly. For NetBackup to create the new staging directory, the immediate parent directory must exist. If you want the restore to use `E:\recovery\staging`, then `E:\recovery` must exist. If the `E:\recovery` directory does not exist, the restore fails.

If you create the directory yourself, the **SYSTEM**, the domain administrator, and the local administrator accounts must have **Full Control** permissions. Additionally, Access Control Lists inherited from the parent directory are not secure and must be disabled.

Restricted restore mode supports alternate location restores. You can configure the alternate location in the NetBackup web UI.

Limitations of restricted restore mode:

- Restricted restore mode is currently only supported on Windows. The recovery host must also be Windows.
- The file ownership of the restored files is set to the account that was used for the NetBackup backup operation.
- Restore of ACLs is not supported.
- Restricted restore mode does not support renaming of targets for soft links.
- Restricted restore mode creates new files where hard links had previously been used.
- Irregular files such as sparse files, device files, special files, and junction points are not supported.
- A supported version of VMware Tools must be running for the restore to succeed.
- File path length with the directory cannot exceed 260 characters.

Performance considerations

File transport through the required infrastructure for this restore method is significantly slower than VMware agentless restores. As a result of performance concerns, Veritas recommends limiting the restore to fewer than 100 files and less than 1 GB of data.

Recover VMware Cloud Director virtual machines

You can only recover a virtual machine (VM) to a VMware Cloud Director if the VM was backed up from a VMware Cloud Director.

To recover a VMware Cloud Director VM

- 1 Select **Workloads > VMware >** and select virtual machine to recover.
- 2 On the **Recovery points** tab, select **Recover > Restore virtual machine**.
- 3 On the **Recovery target** page, select to restore the VM to either **VMware Cloud Director** or **vSphere**.
 - If you select **vSphere**, refer to the following information:
See “[Recover a VM](#)” on page 70. information.
 - If you select **VMware Cloud Director**, continue with this procedure.
- 4 On the **Recovery target** page, specify the VMware Cloud Director and vSphere recovery destination information.
 - The default values shown restore the VM back to its original location.
 - If you change any of the VMware Cloud Director recovery destination information, you must update the **vSphere** recovery destination information.
 - If you accept the default VMware Cloud Director recovery destination information, you can change the **vSphere** recovery destination information if necessary.

Click **Next**.

- 5 On the **vApp options** screen, specify the vApp information.
 - To restore to an existing vApp, browse the list of vApps or enter the name of a vApp that exists.
 - To restore to a new vApp, enter the name of the new vApp.
 - The **Status** shows **New** if the vApp does not exist in VMware Cloud Director. A new vApp is created.

Click **Next**.

- 6 For the **Recovery options** page, specify any recovery options for your restore and click **Next**.
- 7 The **Review** screen summarizes the selections made. A pre-recovery check attempts to determine if there are issues with any of the selected options. You can override any errors shown, however, the recovery can fail if errors are not addressed.

Protecting VMs using hardware snapshot and replication

This chapter includes the following topics:

- [About virtual machines and hardware snapshots](#)
- [Deployment and architecture](#)
- [Features and applications supported](#)
- [Prerequisites for hardware snapshot and replication](#)
- [Operations supported with hardware snapshot](#)
- [Configuring a VMware policy to use hardware snapshot](#)
- [Configuring a VMware policy to use NetBackup snapshot manager replication](#)
- [Jobs in the Activity Monitor that use hardware snapshot for VMs](#)
- [Notes and limitations](#)
- [Troubleshooting with VMware hardware snapshot and replication operations](#)


About virtual machines and hardware snapshots

Hardware snapshot-based solution for VMware uses storage array snapshots for protecting VMware virtual machines. The benefits of using hardware snapshot are the reduced stun time for virtual machine. The VM snapshot is retained only for the duration of hardware snapshot.

This solution uses the NetBackup snapshot manager for performing hardware snapshots. For more information about NetBackup snapshot manager, refer to the *NetBackup™ Snapshot Manager for Data Center Administrator's guide*.

Deployment and architecture

Following is the deployment and architecture diagram of VMware hardware snapshot-based solution.



Note: This solution supports only the VMware datastores which are created on the NAS storage. It does not support the VMware datastores created on the SAN storage.

For all the supported NAS storage arrays, refer to the *NetBackup Snapshot Manager* section, under *Snapshot Solutions* in the *NetBackup Hardware and Cloud Storage Compatibility List (HCL)*.

Features and applications supported

Hardware snapshot-based protection for VMware includes the following features for protecting the virtual machine snapshots and replicated copies:

- Creates an instantaneous hardware snapshot of virtual machines.

- Backs up the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Block level incremental backup (BLIB) of the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Accelerator enabled backups of the virtual machines from the snapshots at primary locations and from replicated snapshots at remote locations.
- Supports browsing of virtual machine snapshots.
- Restores a virtual machine from its vmdk files that are in a snapshot.
- Restores an individual vmdk that is present in a snapshot.
- Restores the individual files from the vmdk files in a snapshot.
- Supports the storage lifecycle policies (SLPs).
- Under the Application Protection, following applications are supported in the VMware policy:
 - Microsoft Exchange databases
 - Microsoft SQL server

Prerequisites for hardware snapshot and replication

Following are the prerequisites for hardware snapshot-based support explained in the table.

Table 8-1 Prerequisites for hardware snapshot support

Support parameter	Description
System	■ All supported NetBackup primary, media server platforms. ■ Backup host for VMware must be RHEL, SUSE or Windows. ■ Snapshot manager server supported to the operating system platform as follows: ■ Ubuntu 16.04 and 18.04 Server LTS ■ Red Hat Enterprise Linux (RHEL) 8.2 and 7.x

Table 8-1 Prerequisites for hardware snapshot support *(continued)*

Support parameter	Description
Configuration	■ NetBackup version 10.1 primary, media server and backup host. ■ VMware backup host can be on any of the NetBackup appliance form factor: ■ NBA ■ Flex ■ NetBackup FlexScale (NBFS) ■ NetBackup Snapshot Manager version 10.1
Permission	■ On Windows backup host, the following NetBackup services must be started using similar domain user account. ■ NetBackup Client Service ■ NetBackup Legacy Network Service ■ The domain user must be part of local administrative group.
VMware NFS datastores	VMware NFS datastores mounted on the ESX host must be version NFS 4.1. or NFS 3.0.
VMware VCenter and ESX sever hosting virtual machine	Virtual machines must reside on the NFS datastores.

Operations supported with hardware snapshot

Table 8-2 Virtual machine operations with hardware snapshot

Operation	Description and notes
Create array-based snapshots of virtual machines on the NFS datastore.	Configure a storage lifecycle policy (SLP) and a backup policy to create array snapshots of virtual machines. The snapshots remain on the array or filer and are not backed up to a NetBackup media server storage unit. Note: ■ The snapshots are created on the NFS datastores only. ■ The virtual machine or its individual files can be restored directly from the snapshots on the storage array. The snapshots can also be replicated to other locations. ■ For fast browsing of files to restore, include the Index From Snapshot option in the SLP. This option, catalogs the metadata of the virtual machine.
Back up quiesce virtual machines from a snapshot (or snapshot replica) which resides on the NFS datastore.	Configure SLP and backup policy to make a backup image from the virtual machine snapshot. NetBackup backs up only the virtual machines quiesce before the snapshot occurs. The backup image is written to NetBackup storage unit. The image is retained according to the policy's retention period. Note: The Application consistent snapshot option in the policy must be enabled (Under Options > Snapshot Client Options).
Restore a virtual machine from a snapshot (or snapshot replica) that is on the NFS datastore or from the backup image written to NetBackup storage unit.	Use the NetBackup web UI interface to restore the virtual machine. Supported restore destinations are the original (NFS) datastore or an alternate datastore (NFS or non-NFS).

Table 8-2 Virtual machine operations with hardware snapshot (*continued*)

Operation	Description and notes
Restore individual files and VMDK from a snapshot (or snapshot replica) that is on the NFS datastore or from the backup image written to NetBackup storage unit.	Use the Backup , Archive , and Restore interface to restore the files. Note: ■ To restore files from a replica of the snapshot, the replica must exist in the same NetBackup domain as the snapshot. ■ To restore files to the original virtual machine, a NetBackup client must be installed on the original virtual machine.
Index from Snapshot	The Index From Snapshot operation catalogs the metadata of the virtual machine. This allows fast browsing of files to restore. Use Index From Snapshot option in the SLP to use this feature.
Live Browse of snapshot	The live-browse function allows you to browse the content of VM snapshot that resides on the storage array.

Configuring a VMware policy to use hardware snapshot

The following procedure describes how to configure a VMware policy to create hardware snapshots of virtual machines residing on NFS datastore using NetBackup web UI.

See [“Working with VMware policies in the web UI”](#) on page 22. for more information about configuring VMware policies.

Table 8-3 Configuration steps and description

Step	Description	Reference
1	Configure the NetBackup Snapshot Manager server in NetBackup.	For more details, refer the NetBackup™ Snapshot Manager for Data Center Administrator's guide.
2	Configure the NAS storage array plug-in.	For more details, refer the NetBackup™ Snapshot Manager for Data Center Administrator's guide.

Table 8-3 Configuration steps and description (*continued*)

Step	Description	Reference
3	Add the VMware backup host to your NetBackup configuration.	See the topic on adding the VMware backup host to NetBackup, in the NetBackup for VMware Guide.
4	Configure NetBackup access credentials for the VMware vCenter server or ESX server.	See the topic on adding NetBackup credentials for VMware, in the NetBackup for VMware Guide.
5	Configure the SLP to use snapshot	For more details, refer to these chapters in the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> , topic: <i>Configuring storage lifecycle policies for snapshots and snapshot replication</i> .
6	Configure a NetBackup VMware policy to perform the operations that are specified in the SLP	For more details, refer to the NetBackup VMware Administrator's Guide.

Only those policy options that are necessary to configure VMware policy to use hardware snapshot of VMs residing on NFS Datastore are listed in the following procedure.

Note: This feature is not supported with protection plan for VMware. You need to create a VMware policy using NetBackup web UI to use this feature.

Following are the steps to create VMware policy to protect the VM residing on NFS Datastore using hardware snapshots.

To create a policy to use VM hardware snapshot in the web UI

- 1 On the left, click **Protection > Policies**.
- 2 To change a VMware policy, select it from the list.
To add a policy, click **Add**, enter a **Policy name**, and select **VMware** from the **Policy type** drop-down list.
- 3 Configure the options on the policy **Attributes** tab. The following items are specific to creating a VMware policy for hardware snapshots:
 - **Policy storage:** Select the SLP which you want to use and is configured for the snapshot-based protection.
 - **Perform snapshot backups:** Enable this option to automatically select other options which are required for the snapshot backup.

- **Perform snapshot backup options:** Click **Options** button to view the **Snapshot Options** dialog box and the default configuration **Parameters** as follows:
 - **Snapshot Type:** Select the appropriate snapshot type. By default, **Auto** option is selected which enables NetBackup to automatically determine the snapshot type to be used for array snapshot.
 - **Snapshot Manager:** Select the *NetBackup Snapshot Manager* host which communicates with the storage array to perform the snapshot operations.

Note: To view the list of configured NetBackup Snapshot Manager hosts, provide VIEW permissions to the user. `MANAGE > SNAPSHOT-MGMT-SERVER > View`

- **Maximum Snapshots :** Sets the maximum number of Snapshots to be retained at one time. When the maximum is reached, snapshot rotation occurs:
 The next snapshot causes the oldest to be deleted. Managed by SLP retention is automatically selected if the Fixed or the Expire after Copy retention is currently selected in the SLP.
- **Application Consistent Snapshot:** This option is enabled by default. In most cases, NetBackup recommends that you keep this option enabled.
 If this option is disabled, data in the virtual machine may not be in a consistent state when the snapshot occurs. The snapshot may not capture all the data in the virtual machine.

Note:

 - To allow the SLP to create a backup image from the snapshot, this option must be enabled.
 - If this option is disabled, note the following about the VMware tab:
 - Exclude deleted blocks and Exclude swap and paging files are disabled.
 - Enable block-level incremental backup is disabled.
 - The **Application Protection** option is disabled.
- **Use Accelerator:** Select this option to accelerate backup operations.

Note: To accelerate backups, Backup From Snapshot must be defined in the SLP.

The MSDP storage unit used for the **Backup From Snapshot** operation must be same as MSDP storage unit used in the Snapshot (or snapshot replication).

- 4 Click **Schedules** tab, select full and incremental schedule for backup and then click **Add**.

Note: To add Incremental schedule, **Block-level incremental backup (BLIB)** option should be enabled.

- 5 Use the **Clients** tab to create a query for the automatic selection of virtual machines. The selected VMs must reside on the NFS datastore.

Note: For instructions on creating a query, refer to the *Configuring the automatic selection of virtual machines for backup* section in the *NetBackup for VMware Administrator's Guide*.

- 6 Use **VMware** tab to select the virtual machine backup options.
Select **Enable block-level incremental backup (BLIB)** option to perform block level incremental backups.

Note: The **Transport modes** are not supported and are disabled. NetBackup uses the VMware file transport mode to move the data between the backup host and the storage array.

Note: Under the **Application Protection** options, only Microsoft Exchange and Microsoft SQL are supported.

- 7 When the policy configuration is complete, click **Create**.

Configuring a VMware policy to use NetBackup snapshot manager replication

Using the NetBackup™ Snapshot Manager for Data Center you can replicate the hardware snapshots of VMs. The replicated snapshots are accessed on VMware backup hosts to create point in time backup copies of VMs. The following procedure describes how to configure a VMware policy to use hardware snapshots and replication of VMs residing on NFS datastore, using NetBackup web UI.

See [“Working with VMware policies in the web UI”](#) on page 22. for more information about configuring VMware policies.

Table 8-4 Configuration steps with description, and reference topics

Step	Description	Reference topic
1	Register the Snapshot Manager server in NetBackup	For more details, refer the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .
2	Configure the NAS storage array plug-in	For more details, refer the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> .
3	Add the VMware access host to your NetBackup configuration	See “Add a VMware access host” on page 17.
4	Configure NetBackup access credentials for the VMware vCenter server or ESX server.	See “Add VMware servers” on page 7. See “Validate and update VMware server credentials” on page 8.
5	Configure the SLP to use snapshot and replication	For more details, refer to these chapters in the <i>NetBackup™ Snapshot Manager for Data Center Administrator's Guide</i> : ■ Storage array replication ■ Configuring storage lifecycle policies for snapshots and snapshot replication ■ Supported storage array in data center

Table 8-4

Configuration steps with description, and reference topics
(continued)

Step	Description	Reference topic
6	Configure a NetBackup VMware policy to perform the operations that are specified in the SLP.	See “Configuring a VMware policy to use hardware snapshot” on page 85., procedure <i>To create a policy to use VM hardware snapshot in the web UI</i>

Jobs in the Activity Monitor that use hardware snapshot for VMs

You can use the NetBackup Activity Monitor to keep track of virtual machines backups as they occur. The number of jobs that appear in the Activity Monitor depends on the policy's **Application Consistent Snapshot** option.

Note: By default, the **Application Consistent Snapshot** option is enabled. In most cases, NetBackup recommends that you keep this option enabled. If this option is disabled, data in the virtual machine may not be in a consistent state when the snapshot occurs.

Table 8-5

Job flow in the Activity Monitor

Application consistent snapshot option	Job flow in the Activity Monitor
Enabled	The first job discovers the virtual machines. This job is labeled Backup. Backup job starts with the following: ■ A Snapshot job for each virtual machine. ■ A Snapshot job for each datastore.
Disabled	The first job discovers the virtual machines. This job is labeled Backup. Backup job starts with the following: ■ A Snapshot job to collect all the virtual machines' configuration data. ■ A Snapshot job for each datastore.

Example 1: Virtual machine jobs with the **Application Consistent Snapshot** option enabled.

5 Jobs (0 Queued 0 Active 0 Waiting for Retry 0 Suspended 0 Incomplete 5 Done - 0 selected)										Search
Job ID	Type	Client	Job Policy	State	Start Time	State Details	Status	Job Schedule		
5	Backup From Snapshot	VMwareNAS_DemoVM	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 1:44:17 PM		0 -			
6	Backup From Snapshot	VMwareNAS_DemoVM	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 1:44:30 PM		0 Full_BK			
4	Backup	r7515-112v01.windia.veritas.com	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 1:38:20 PM		0 -			
4	Snapshot	NetAPP_SS_200	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 1:38:45 PM		0 Full_BK			
3	Snapshot	VMwareNAS_DemoVM	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 1:38:59 PM		0 -			

The jobs occurred as follows:

- The discovery (parent) Backup job for virtual machine discovery is ID 2.
- Job 3 made VMware snapshots of the virtual machine VMwareNAS_DemoVM.
- Job 4 made snapshots of datastore NetAPP_SS_200.
- Job 5 parent Backup from Snapshot export and mount the snapshot.
- Job 6 child Backup from Snapshot does the backup and creates the backup image.

Example 2: Virtual machine jobs with the **Application Consistent Snapshot** option disable.

7	Backup	r7515-112v01.windia.veritas.com	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 2:55:48 PM	0 -		
9	Snapshot	NetAPP_SS_200GB	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 2:56:00 PM	0 Full_BK		
8	Snapshot	r7515-112v01.windia.veritas.com	VMware_NAS_Demo_Pol	Done	Jun 28, 2022 2:55:57 PM	0 -		

The jobs occurred as follows:

- The discovery (parent) Backup job for virtual machine discovery is ID 7.
- Job 8 collected the configuration data of all the virtual machines selected (VM1, VM2, and so forth).
- Jobs 9 snapshots of the virtual machine datastores.

Notes and limitations

Note the following about VMware NAS hardware snapshot for virtual machines datastores:

- Index from Snapshot, Single File Restore (SFR) from snapshot and Live browse from snapshot for XFS, Btrfs file system are currently not supported.
- Agentless Single File Restore (ALVR) is not supported from the NetBackup web UI.
- GRT and Individual VMDK restore is not supported from the NetBackup web UI.
- While restoring the full VM from the incremental snapshot copy, the restore is performed only from the snapshot which is taken during the incremental backup. In the case of restore from the incremental backup image copy, the restore is performed from all the incremental images and the full backup image.

- To restore from the incremental backup images, all the primary copies must be either snapshot copies or the backup images copies.
- VMware policy which protects Microsoft Exchange using hardware snapshot-based backups, in that policy only Windows must be specified as the backup host.

Troubleshooting with VMware hardware snapshot and replication operations

About gathering information and checking logs

To create detailed log information, place a **VERBOSE** entry in the `bp.conf` file on the NetBackup primary and client. Or set the global logging level to a high value in the **Logging** dialog, under both primary server Properties and Client Properties.

These directories can eventually require a lot of disk space. Delete them when you are finished troubleshooting and remove the **VERBOSE** option from the `bp.conf` file. Or reset the Global logging level to a lower value.

Logging directories for Linux platform

To create logging directories use `/usr/opensv/netbackup/logs/mklogdir` script. You can also create the directories using an access mode of 755 so, NetBackup can write to the logs.

Table 8-6 Linux logging directories for snapshot operation

Path of log directory	Where to create the directory
<code>/usr/opensv/netbackup/logs/bprd</code>	NetBackup primary server
<code>/usr/opensv/logs/nbjm</code>	NetBackup primary server
<code>/usr/opensv/netbackup/logs/bpbm</code>	NetBackup media server
<code>/usr/opensv/netbackup/logs/bpfis</code>	Backup host client

Table 8-7 Linux logging directories for backup operation

Path of log directory	Where to create the directory
<code>/usr/opensv/netbackup/logs/bprd</code>	NetBackup primary server
<code>/usr/opensv/logs/nbjm</code>	NetBackup primary server
<code>/usr/opensv/netbackup/logs/bpdbm</code>	NetBackup primary server

Table 8-7 Linux logging directories for backup operation (*continued*)

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bptm	NetBackup media server
/usr/opensv/netbackup/logs/bpbm	NetBackup media server
/usr/opensv/netbackup/logs/bpfis	Backup host client
/usr/opensv/netbackup/logs/bppfi	Backup host client
/usr/opensv/netbackup/logs/bpbkar	Backup host client
/usr/opensv/netbackup/logs/bppfi	Backup host client
/usr/opensv/netbackup/logs/vxms	Backup host client

Table 8-8 Linux logging directories for single file restore operation

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/netbackup/logs/bpcd	Restore host client
/usr/opensv/netbackup/logs/bpbkar	Restore host client
/usr/opensv/netbackup/logs/bpfis	Restore host client
/usr/opensv/netbackup/logs/bppfi	Restore host client
/usr/opensv/logs/ncfnbhfr	Restore host client
/usr/opensv/netbackup/logs/vxms	Restore host client
/usr/opensv/netbackup/logs/tar	Destination client where the files are restored.

Table 8-9 Linux logging directories for full VM restore operation

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/netbackup/logs/bpbm	NetBackup media server
/usr/opensv/logs/ncfnbvmcopyback	Restore host client
/usr/opensv/netbackup/logs/vxms	Restore host client

Table 8-10 Linux logging directories for live browse operation

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/netbackup/logs/bpdbm	NetBackup primary server
/usr/opensv/netbackup/logs/bpcd	Backup host client
/usr/opensv/netbackup/logs/bppfi	Backup host client
/usr/opensv/logs/ncfnbbrowse	Backup host client
/usr/opensv/netbackup/logs/vxms	Backup host client

Table 8-11 Linux logging directories for index from snapshot operation

Path of log directory	Where to create the directory
/usr/opensv/logs/ncflbc	Backup host client
/usr/opensv/netbackup/logs/vxms	Backup host client

Logging folders for Windows platforms:

Table 8-12 Windows logging directories for snapshot operation

Path of log directory	Where to create the directory
/usr/opensv/netbackup/logs/bprd	NetBackup primary server
/usr/opensv/netbackup/logs/bpdbm	NetBackup primary server
/usr/opensv/netbackup/logs/bpcd	NetBackup media server
/usr/opensv/netbackup/logs/bppfi	Backup host client

Table 8-13 Windows logging directories for backup operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\nbjm	NetBackup primary server
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bptm	NetBackup media server
install_path\NetBackup\logs\bpbrm	NetBackup media server

Table 8-13 Windows logging directories for backup operation (*continued*)

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bpfis	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\bpbkar	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Table 8-14 Windows logging directories for single file restore operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpccd	Restore host client
install_path\NetBackup\logs\bpbkar	Restore host client
install_path\NetBackup\logs\bpfis	Restore host client
install_path\NetBackup\logs\bppfi	Restore host client
install_path\NetBackup\logs\ncfnbhfr	Restore host client
install_path\NetBackup\logs\vxms	Restore host client
install_path\NetBackup\logs\tar	Destination client where the files are restored.

Table 8-15 Windows logging directories for full VM restore operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server
install_path\NetBackup\logs\bpbrm	NetBackup media server
install_path\NetBackup\logs\ncfnbvmcopyback	Restore host client
install_path\NetBackup\logs\vxms	Restore host client

Table 8-16 Windows logging directories for live browse operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bprd	NetBackup primary server

Table 8-16 Windows logging directories for live browse operation (*continued*)

Path of log directory	Where to create the directory
install_path\NetBackup\logs\bpdbm	NetBackup primary server
install_path\NetBackup\logs\bpcd	Backup host client
install_path\NetBackup\logs\bppfi	Backup host client
install_path\NetBackup\logs\ncfnbbrowse	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Table 8-17 Windows logging directories for index from snapshot operation

Path of log directory	Where to create the directory
install_path\NetBackup\logs\ncflbc	Backup host client
install_path\NetBackup\logs\vxms	Backup host client

Troubleshooting VMware operations

This chapter includes the following topics:

- [Errors when adding VMware servers](#)
- [Errors when browsing VMware servers](#)
- [Errors for the status for a newly discovered VM](#)
- [Error when downloading files from an instant access VM](#)
- [Troubleshooting backups and restores of excluded virtual disks](#)
- [Restore fails for a virtual machine with multiple datastores](#)

Errors when adding VMware servers

Table 9-1 Errors adding VMware servers

Error message or cause	Explanation and recommended action
Virtualization server credential validation fails.	This error occurs when the NetBackup primary server is in a DNAT or a similar setup can access only a few specified NetBackup hosts (<code>PROXY_SERVERS</code>). The credentials validation occurs in the following order: ■ The auto-discovered discovery host is used to access the virtualization server. ■ If the autodiscovery does not find any information about the virtualization server on the discovery host, the NetBackup primary server is used. Workaround: When you add the virtualization server credentials, select the proxy server that has access to the virtualization server as the backup host for validation. Note: Adding or updating VMware credentials also automatically starts the discovery of the VMware server. When backup host information is provided in the request, it is used to perform validation of credentials as well as for performing the discovery. For discovery, NetBackup 8.1.2 is the minimum version that is supported for a NetBackup media server or client that serves as a backup host. For older versions, backup host credential validation succeeds, but the discovery of VMware servers fails.
Unable to obtain the list of trusted Certificate Authorities.	This error might occur when VMware server credentials are added, updated, or validated. It occurs if the environment is configured to enabled communication between NetBackup (primary server, media server, or client) and vCenter, ESX, or any other VMware entity using authenticated certificates. Workaround: Ensure that certificates are installed and are valid.

Errors when browsing VMware servers

The following table describes the problems that may occur when you click on a server under **VMware servers**.

Table 9-2 Errors browsing VMware servers

Error message or cause	Explanation and recommended action
No VMs or other objects were discovered for the VMware server.	■ If the server was added recently, the VM discovery process for that server may not have completed yet. Recommended action: Wait for the discovery process to finish. Note: The discovery of VMs and other objects in the vCenter, ESXi server, or VMware Cloud Director server begins: when server credentials are added or updated through the web UI or an API. However, the server's VMs and other objects might not appear in the UI immediately. They appear after the discovery process for the VMware server completes. Discovery also occurs at set intervals according to the <code>VMWARE_AUTODISCOVERY_INTERVAL</code> option. (The default interval is every 8 hours.) To perform autodiscovery of VMware server objects at a different frequency: See “Change the autodiscovery frequency of VMware assets” on page 20. ■ VMs or other objects of the VMware server may not be accessible for the added VMware server credentials. Recommended action: From the option menu on the right of the row, select Edit. Review the VMware server credentials and correct them as needed.

Errors for the status for a newly discovered VM

The following table describes a problem that may occur when you review the status of a newly discovered VM under **Virtual machines**.

Table 9-3 Errors encountered when you review Status for a newly discovered VM

Error message or cause	Explanation and recommended action
The protection status of a VM indicates that it has not been backed up. However, a backup job that includes the VM has successfully completed.	In the NetBackup web UI, the protection status for a newly discovered VM does not indicate that it is backed up until the next backup of the VM has completed. In some circumstances, a new VM is backed up before the discovery of that VM has happened, as in the following scenario: ■ By default, autodiscovery occurs every 8 hours. ■ A new VM is added to the environment. ■ A backup job completes successfully before discovery completes. For example, a backup job that uses existing policies where the new VM is included as part of the backup selection criteria. ■ Later, discovery completes. However, in the NetBackup web UI, the protection status of the VM indicates that it has not been backed up. If you encounter a similar situation, you can still browse the recovery points and recover them. However, it is only after another backup of the VM successfully completes that the protection status indicates that the VM has been backed up. To review the protection status of a newly discovered VM in the NetBackup web UI, Veritas recommends that you wait until the next successful backup has completed. Then, the protection status of the VM should correctly indicate its protection status.

Error when downloading files from an instant access VM

The following table describes the problems that may occur when you download individual files from an instant access VM.

Table 9-4 Errors in downloading files

Error message or cause	Explanation and recommended action
Chrome: This site can't be reached Firefox: Server not found Edge: Hmmmm...can't reach this page	This error can occur for any of the following reasons: ■ The web UI is unable to access the NetBackup media server with the name or IP address that the NetBackup primary server uses to connect to that media server. For example: If the primary server connects to the media server using <code>MSserver1.veritas.com</code>, the web UI must also be able to reach <code>MSserver1.veritas.com</code>. If the primary server uses a short name for the media server such as <code>MSserver1</code>, the web UI must be able to reach <code>https://MSserver1/...</code> Recommended action: Verify that the primary server and the web UI use the same name or IP address to access the media server (check the <code>hosts</code> file). For example: If the primary server uses the media server's short name, add the media server's short name and IP address to the <code>hosts</code> file of the PC or other host where the web UI is running. The hosts file location on Windows: <code>C:\Windows\System32\drivers\etc\hosts</code> The hosts file location on UNIX or Linux: <code>/etc/hosts</code> ■ The web UI is unable to access the NetBackup media server because that server is behind a firewall. Recommended action: Contact the NetBackup security administrator.

Troubleshooting backups and restores of excluded virtual disks

Refer to the following table if you encounter restore issues for a backup that was configured to exclude virtual disks.

Table 9-5 Issues with excluding virtual disks

Issue	Explanation
The boot disk was backed up even though it was excluded from the backup.	The virtual machine only has a boot disk and no other disks.
	The boot disk is part of a managed volume (Windows LDM or Linux LVM). NetBackup can only exclude a boot disk if it is fully contained on a single disk.
	The virtual machine's boot disk is an independent disk and has no other disks.
	NetBackup was not able to identify the boot disk. The boot disk must include the boot partition and the system or the boot directory.
A restored boot disk has no data.	The boot disk is an independent disk. NetBackup cannot back up the data in this type of disk.
A restored virtual machine has a disk that contains missing or incomplete data.	The disk that has missing or incomplete data was excluded from the backup.
A data disk (or disks) was backed up even though it was excluded from the backup.	The virtual machine has only one disk (such as C:). In this case, the single drive is backed up and is not excluded.
A virtual machine is restored to an unexpected state.	You added a disk to the virtual machine and changed the settings that exclude disks. However, you did not create a backup of the entire virtual machine after you made the change.
Not all files can be restored individually.	If you remove disks from the custom attribute value between the differential backups, only those files that changed since the last backup can be restored individually. Alternatively, you can restore the entire virtual disk or the VM. After the next full backup, you can restore any of the files individually.
	If you remove controllers from Specific disks to be excluded between the differential backups, only those files that changed since the last backup are available for restore. All files are available for restore after the next full backup.

Restore fails for a virtual machine with multiple datastores

Table 9-6 Issues with restores of a virtual machine with multiple datastores

Issue	Explanation
Restore fails because the datastore did not have enough space for the .vmdk files.	This issue can occur when a virtual machine is configured on multiple datastores and a leftover snapshot existed on the virtual machine when it was backed up. NetBackup tries to restore all .vmdk files to the snapshot datastore. Alternatively, you can restore the virtual machine to an alternate location.