NetBackup™ Web UI VMware Administrator’s Guide

Release 10.0

VERITAS™

Last updated: 2022-03-04

Legal Notice

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 Introducing the NetBackup web user interface 7

About the NetBackup web UI 7
Terminology 9
Sign in to the NetBackup web UI 10
Sign out of the NetBackup web UI 12

Chapter 2 Monitoring NetBackup 13

The NetBackup dashboard 13
Jobs: cancel, suspend, restart, resume, delete 13
Job monitoring 14
Search for or filter jobs in the jobs list 15

Chapter 3 Managing VMware servers 17

Add VMware servers 17
Validate and update VMware server credentials 18
Browse VMware servers 18
Remove VMware servers 19
Create an intelligent VM group 19
Remove an intelligent VM group 27
Add a VMware access host 27
Remove a VMware access host 28
Change resource limits for VMware resource types 28
About VMware discovery 29
Change the autodiscovery frequency of VMware assets 30
Discover VMware server assets manually 30

Chapter 4 Protecting VMs 32

Working with VMware policies in the web UI 32
Protect VMs or intelligent VM groups 33
    Schedules 34
    Backup options and Advanced options 35
    Exclude disks from backups 36
    Snapshot retry options 36

Customize protection settings for a VMware asset 37
Remove protection from VMs or intelligent VM groups 38
View the protection status of VMs or intelligent VM groups 38

Chapter 5 Instant access 40

Create an instant access VM 40
Restore files and folders from a VM backup image 42
Download files and folders from a VM backup image 43
Things to consider before you use the instant access feature 44
Instant access Build Your Own (BYO) 46
Prerequisites of Instant Access Build Your Own (BYO) 47
Hardware configuration requirement of Instant Access Build Your Own (BYO) 48
Frequently asked questions 48

Chapter 6 Instant rollback 51

Prerequisites of instant rollback 51
Things to consider before you use the instant rollback feature 51
Instant rollback from a VM backup image 52

Chapter 7 Continuous data protection 54

CDP terminology 54
CDP architecture 55
About continuous data protection 56
Prerequisites 57
Capacity-based licensing for CDP 58
Steps to configure CDP 58
Defining the CDP gateway 59
Sizing considerations 60
Limiting concurrent CDP backup jobs 62
Controlling full sync 64
Monitoring CDP jobs 65
Using accelerators with CDP 67
Recovering CDP protected VMs 68
Some limitations of CDP 68
Troubleshooting for CDP 69

Chapter 8 VM recovery 73

Recover a VM 73
Recovery options 75
Advanced recovery options 75

Advanced recovery options: Format of restored virtual disks 76
Advanced recovery options: Transport mode 77
About VMware agentless restore 77
Prerequisites and limitations of VMware agentless restores 78
Recover files and folders with VMware agentless restore 80
About restricted restore mode 80

Chapter 9 Troubleshooting VMware operations 82

Errors when adding VMware servers 83
Errors when browsing VMware servers 83
Errors for the status for a newly discovered VM 84
Error when downloading files from an instant access VM 85
Troubleshooting backups and restores of excluded virtual disks 86
Restore fails for a virtual machine with multiple datastores 88

Introducing the NetBackup web user interface

This chapter includes the following topics:

- About the NetBackup web UI
- Terminology
- Sign in to the NetBackup web UI
- Sign out of the NetBackup web UI

About the NetBackup web UI

The NetBackup web user interface provides the following features:

- Ability to access the primary server from a web browser, including Chrome and Firefox. For details on supported browsers for the web UI, see the NetBackup Software Compatibility List.
  Note that the NetBackup web UI may behave differently for different browsers. Some functionality, for example a date picker, may not be available on all browsers. These inconsistencies are due to the capabilities of the browser and not because of a limitation with NetBackup.
- A dashboard that displays a quick overview of the information that is important to you.
- Role-based access control (RBAC) that lets the administrator configure user access to NetBackup and to delegate the tasks for workload protection.
- Protection of assets is achieved through protection plans, job management, and visibility of the protection status of assets.

- Workload administrators can create protection plans, manage credentials, subscribe assets to the protection plans that meet the SLO, monitor protection status, and perform self-service recovery of assets.

Note: The NetBackup web UI is best viewed at a 1280x1024 or higher screen resolution.

Access control in the NetBackup web UI

NetBackup uses role-based access control to grant access to the web UI. Access control is accomplished through roles.

- A role defines the operations that a user can perform and the features that the user can access in the web UI. For example, access to any workload assets, protection plans, or credentials.
- RBAC is only available for the web UI and the APIs.

  Other access control methods for NetBackup are not supported for the web UI and APIs, except for the Enhanced Auditing (EA).

Monitor NetBackup jobs

The NetBackup web UI lets administrators more easily monitor NetBackup job operations and identify any issues that need attention.

Protection plans: One place to configure schedules, storage, and storage options

Protection plans offer the following benefits:

- A default workload administrator can select the protection plans to use to protect assets.
- With the necessary RBAC permissions, a workload administrator can create and manage protection plans, including the backup schedules and storage that is used.
- When you select from your available storage, you can see any additional features available for that storage.

Self-service recovery

The NetBackup web UI makes it easy for a workload administrator to recover VMs, databases, or other asset types applicable to that workload.

For the workloads that support the instant access feature, users can mount a snapshot for immediate access to a VM’s files or to a database.

Terminology

The following table describes the concepts and terms in web user interface.

Table 1-1 Web user interface terminology and concepts

Term	Definition
Asset group	See *intelligent group*.
Asset	The data to be protected, such as physical clients, virtual machines, and database applications.
Backup now	An immediate backup of an asset. NetBackup performs a one-time, full backup of an asset using the selected protection plan. This backup does not affect any scheduled backups.
Intelligent group	Allows NetBackup to automatically select assets for protection based on the criteria (queries) that you specify. An intelligent group automatically stays up-to-date with changes in the production environment. These groups are also referred to as asset groups. These groups appear under the tab **Intelligent VM groups** or **Intelligent groups**.
Instant access	**Note:** Instant access is supported only a select number of workloads. An instant access VM or database that is created from a NetBackup backup image is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the snapshot directly on the backup storage device and the snapshot is treated as a normal VM or database.
Protection plan	A protection plan defines when backups are performed, how long the backups are retained, and the type of storage to use. Once a protection plan is set up, assets can be subscribed to the protection plan.
RBAC	Role-based access control. The role administrator can delegate or limit access to the NetBackup web UI through the roles that are configured in RBAC. **Note:** The roles that you configure in RBAC do not control access to the NetBackup Administration Console.

Table 1-1 Web user interface terminology and concepts *(continued)*

Term	Definition
Role	For RBAC, defines the operations that a user can perform and the assets or objects that they can access. For example, you can configure a role to manage recovery of specific databases and the credentials that are needed for backups and restores.
Storage	The storage to which the data is backed up, replicated, or duplicated (for long-term retention).
Subscribe, to a protection plan	The action of selecting an asset or an asset group to subscribe to a protection plan. The asset is then protected according to the schedule in the plan. The web UI also refers to *Subscribe* as *Add protection*.
Unsubscribe, from a protection plan	*Unsubscribe* refers to the action of removing protection or removing an asset or asset group from a plan.
Workload	The type of asset. For example: VMware, RHV, AHV, Microsoft SQL, Oracle, Cloud, or Kubernetes.

Sign in to the NetBackup web UI

Authorized users can sign in to a NetBackup primary server from a web browser, using the NetBackup web UI.

For more information, refer to the *Authorized users* section in the *NetBackup™ Web UI Administrator's Guide*.

The following sign-in options are available:

- Sign in with a username and password
- Sign in with a certificate or smart card
- Sign in with single sign-on (SSO)

Sign in with a username and password

Only authorized users can sign in to NetBackup web UI. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup primary server using a username and password

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

2. Depending on the sign-in options that are available, choose from the following:

   - Enter your credentials and click **Sign in**.
   - If the default method is not username and password, click **Sign in with username and password**. Then enter your credentials.

   The following are example credentials:

For this type of user	Use this format	Example
Local user	*username*	**jane_doe**
Windows user	*DOMAIN\username*	**WINDOWS\jane_doe**
UNIX user	*username*	john_doe

Sign in with a certificate or smart card

You can sign in to NetBackup web UI with a smart card or digital certificate if you are an authorized user. Contact your NetBackup security administrator for more information.

To use a digital certificate that is not on a smart card, you must first upload the certificate to the browser's certificate manager. See the browser documentation for instructions or contact your certificate administrator for more information.

To sign in with a certificate or smart card

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

2. Click **Sign in with certificate or smart card**.

3. When your browser prompts you, select the certificate.

Sign in with single sign-on (SSO)

You can sign in to NetBackup web UI with the single sign-on (SSO) option if SAML is configured as an identity provider in your NetBackup environment. Contact your NetBackup security administrator for more information.

To sign in to a NetBackup primary server using SSO

1. Open a web browser and go to the following URL.

   https://*primaryserver*/webui/login

   The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

2. Click **Sign in with single sign-on**.

3. Follow the steps as provided by your administrator.

   On subsequent logons, NetBackup signs you in automatically to the primary server.

Sign out of the NetBackup web UI

Note that NetBackup automatically signs you out of the web UI after 24 hours, which is the maximum time that is allowed for a user session. After that time NetBackup requires that you sign in again. You can also sign out if you want to change the sign-in option that you want to use (username and password, smart card, or single sign-on (SSO)).

To sign out of the NetBackup web UI

On the top right, click the profile icon and click **Sign out**.

Chapter 2

Monitoring NetBackup

This chapter includes the following topics:

- The NetBackup dashboard
- Jobs: cancel, suspend, restart, resume, delete
- Job monitoring
- Search for or filter jobs in the jobs list

The NetBackup dashboard

The NetBackup dashboard provides a quick view of the details that are related to your role in your organization.

Table 2-1 The NetBackup dashboard

Dashboard widget	Description
Jobs	Lists job information, including the number of active and queued jobs and the status of attempted and completed jobs.

Jobs: cancel, suspend, restart, resume, delete

Depending on the state of a job, you can perform certain actions on that job.

To manage a job

1. Click **Activity monitor > Jobs**.
2. Select one or more jobs.
3. The top menu shows the actions that you can perform for the selected jobs.

Cancel You can cancel the jobs that have not yet completed. They can be in one of the following states: queued, re-queued, active, incomplete, or suspended.

When a parent job is cancelled, any child jobs are also cancelled.

Suspend You can suspend backup and restore any jobs that contain checkpoints.

Restart You can restart the jobs that have completed, failed, or that have been cancelled or suspended.

A new job ID is created for the new job.

Resume You can resume the jobs that have been suspended or are in an incomplete state.

Delete You can delete the jobs that have completed. When a parent job is deleted, any child jobs are also deleted.

Note: NetBackup Kubernetes 10.0 release do not support suspend, restart, and resume operations for running the **Backup from snapshot** job.

Job monitoring

Use the **Jobs** node in the Activity monitor to monitor the jobs in your NetBackup environment. The default view for jobs is the **List view** that contains the non-hierarchical list of all the jobs. You can also use the **Hierarchical view** to see the hierarchy of parent and child jobs.

List view

Hierarchy view

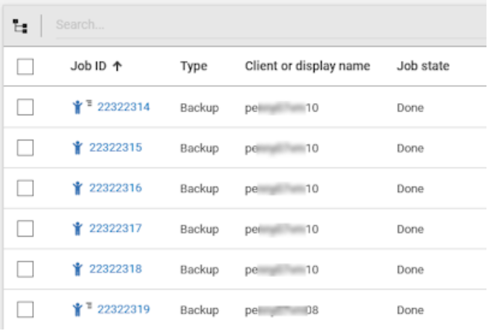

Search...

□	Job ID ↑	Type	Client or display name	Job state
□	22322314	Backup	pe[redacted]10	Done
□	22322315	Backup	pe[redacted]10	Done
□	22322316	Backup	pe[redacted]10	Done
□	22322317	Backup	pe[redacted]10	Done
□	22322318	Backup	pe[redacted]10	Done
□	22322319	Backup	pe[redacted]08	Done

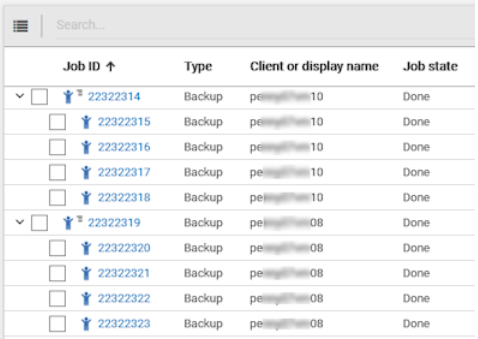

Search...

	Job ID ↑	Type	Client or display name	Job state
⌄ □	22322314	Backup	pe[redacted]10	Done
□	22322315	Backup	pe[redacted]10	Done
□	22322316	Backup	pe[redacted]10	Done
□	22322317	Backup	pe[redacted]10	Done
□	22322318	Backup	pe[redacted]10	Done
⌄ □	22322319	Backup	pe[redacted]08	Done
□	22322320	Backup	pe[redacted]08	Done
□	22322321	Backup	pe[redacted]08	Done
□	22322322	Backup	pe[redacted]08	Done
□	22322323	Backup	pe[redacted]08	Done

The type of jobs that you can view and manage depend on the RBAC role that you have. For example, a workload administrator (such as the Default VMware Administrator role) can view and manage only jobs for that workload. In contrast, the Administrator role lets you view and manage all NetBackup jobs.

If you have an RBAC role that allows access to jobs, you can see the jobs list in the job hierarchy view. For example, the Default VMware Administrator role lets you see VMware jobs in the hierarchy view. However, if you only have access to one or more VMs (asset-level access), no jobs display in the job hierarchy view.

Search for or filter jobs in the jobs list

You can search for jobs in the Activity monitor or create filters to customize the jobs that are displayed.

Search for jobs in the jobs list

The search feature lets you search for the following job information: status code (complete status code #); policy name; client or display name; client; job ID (complete job ID #), or the job's parent ID.

Search for jobs in the jobs list

1 Click **Activity monitor > Jobs**.

2 In the **Search** box, type the keyword you want to find. For example, a client name or a status code number.

Filter the job list

To filter the job list

1. Click **Activity monitor > Jobs**.
2. In the toolbar, click the **Filter** icon.
3. Click on a filter that you created. Or, click **All jobs** to display all of the available jobs.

Chapter 3

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 and ESXi 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 and standalone ESXi servers to locate VMs and view their details such as 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 and standalone ESXi 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 will 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 and standalone ESXi 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 30.

To discover the VMs immediately:

See “Discover VMware server assets manually” on page 30.

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:

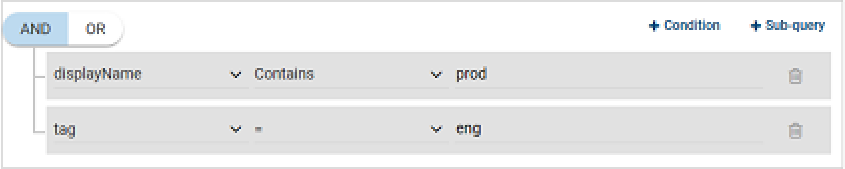

displayName Contains prod

tag = eng

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:

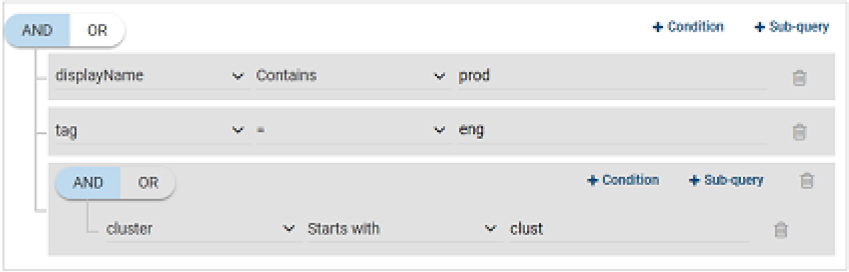

displayName Contains prod

tag = eng

cluster Starts with clust

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 3-1 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 3-1 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 26.	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 3-2 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 box, this option matches the string box_car but not flatbox.
Ends with	Matches the value when it occurs at the end of a string. For example: If the value you enter is dev, this option matches the string 01dev but not 01dev99 or devOP.
Contains	Matches the value you enter wherever that value occurs in the string. For example: If the value you enter is dev, this option matches strings such as 01dev, 01dev99, devOP, and development_machine.
=	Matches only the value that you enter. For example: If the value you enter is VMtest27, this option matches VMTest27 (same case), but not vmtest27, vmTEST27, or VMtest28.
!=	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.

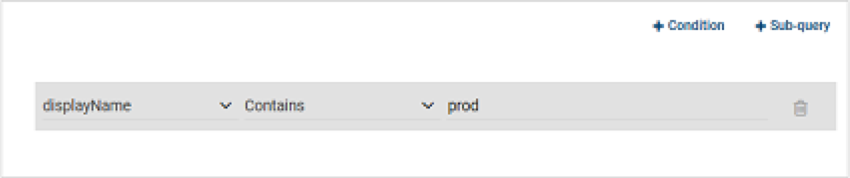

displayName Contains prod

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:

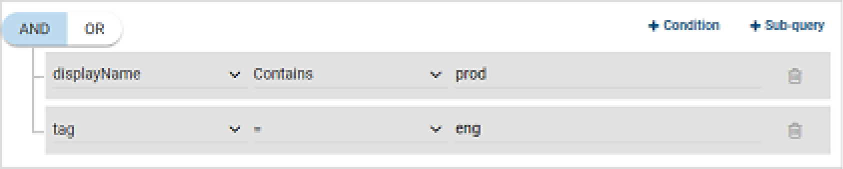

displayName Contains prod

tag = eng

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**:

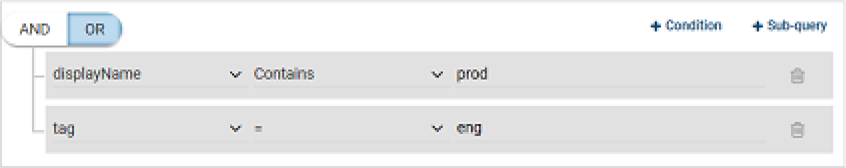

displayName Contains prod

tag = eng

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:

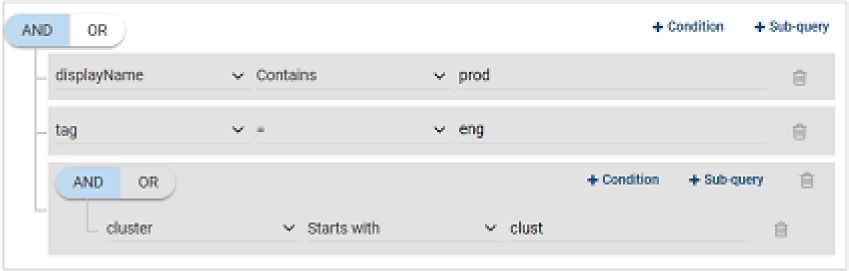

displayName Contains prod

tag = eng

AND OR

cluster Starts with clust

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 30.

To discover the VMs immediately:

See “Discover VMware server assets manually” on page 30.

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 or ESXi 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 30.

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 and standalone ESXi 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 18.

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 38.

See "Errors for the status for a newly discovered VM" on page 84.

Chapter 4

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	**In the NetBackup Administration Console**
Query Builder keywords	You must use OData keywords in the Query Builder.	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 30. 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.

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 34.
   - **Backup options** and **Advanced** options.
     See "Backup options and Advanced options" on page 35.
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 4-1 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.
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 4-2 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 36.

Advanced options

Table 4-3 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.
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 36.

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 4-4 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.
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: `scsi0-0,ide0-0,sata0-0,nvme0-0`. The default value for this attribute is `NB_DISK_EXCLUDE_DISK`. 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 4-5 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 34.
- See “Backup options and Advanced options” on page 35.

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 34.
   - **Backup options** and **Advanced** options.
     See “Backup options and Advanced options” on page 35.
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 84.

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

See “Protect VMs or intelligent VM groups” on page 33.

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)

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 44.

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. 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**.
5. 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.
6. 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.
7. For details on the restored VM, click on the VM under the **Instant access virtual machines** tab and click **View details**.
8. 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 44.

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. 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.
5. 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.

   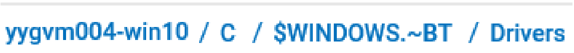

   yygvm004-win10 / C / $WINDOWS.~BT / Drivers

   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.
6. 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.

7 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.

8 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 44.

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

   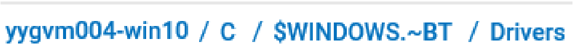

   yygvm004-win10 / C / $WINDOWS.~BT / Drivers

   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.
6. 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 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 optio **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 85.
- 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/openv/netbackup/bin/nbwmc stop
  - /usr/openv/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 show 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 `policycoreutils` and `policycoreutils-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 5-1 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 5-2 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: `/etc/nginx/conf.d/byo.conf` is part of the HTTP section of the original: `/etc/nginx/nginx.conf` file. **3** Run the command: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

Table 5-2 Frequently asked questions *(continued)*

Frequently asked question	Answer
How can I resolve the following issue in the vpfs-config.log 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 `policycoreutils` and `policycoreutils-python` packages through yum tool. **2** Add the following rules that SELinux requires for Nginx to bind on the 10087 port. ■ `semanage port -a -t http_port_t -p tcp 10087` ■ `setsebool -P httpd_can_network_connect 1` **3** Run the following command: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`
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 `<PDDE Storage Path>/spws/var/keys/spws.cert` with the certificate and replace `<PDDE Storage Path>/spws/var/keys/spws.key` with the private key. **5** Run the following command to reload the certificate: `/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

Table 5-2 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 `.../meta_bdev_dir/...` folder under livemount share, while the mount target is in the `/run/media/...` folder.	Follow the guideline to disable the gnome automount: https://access.redhat.com/solutions/20107

How can I resolve the following issue in the `/var/log/vpfs/vpfs-config.log` file?
`**** Asking the NetBackup Webservice to trust the MSDP webserver (spws) **** /usr/openv/netbackup/bin/nblibcurlcmd failed (1):`
	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:
	`/usr/openv/netbackup/bin/bpclntcmd -pn`
	3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command:
	`/usr/openv/pdde/vpfs/bin/vpfs_config.sh --configure_byo`

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 47.

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, and Build Your Own (BYO) server.
- 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

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** 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. **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. **3** Click **Roll back**.

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

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

Chapter 7

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
- 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 7-1 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.
nbcctd	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 7-1 CDP architecture

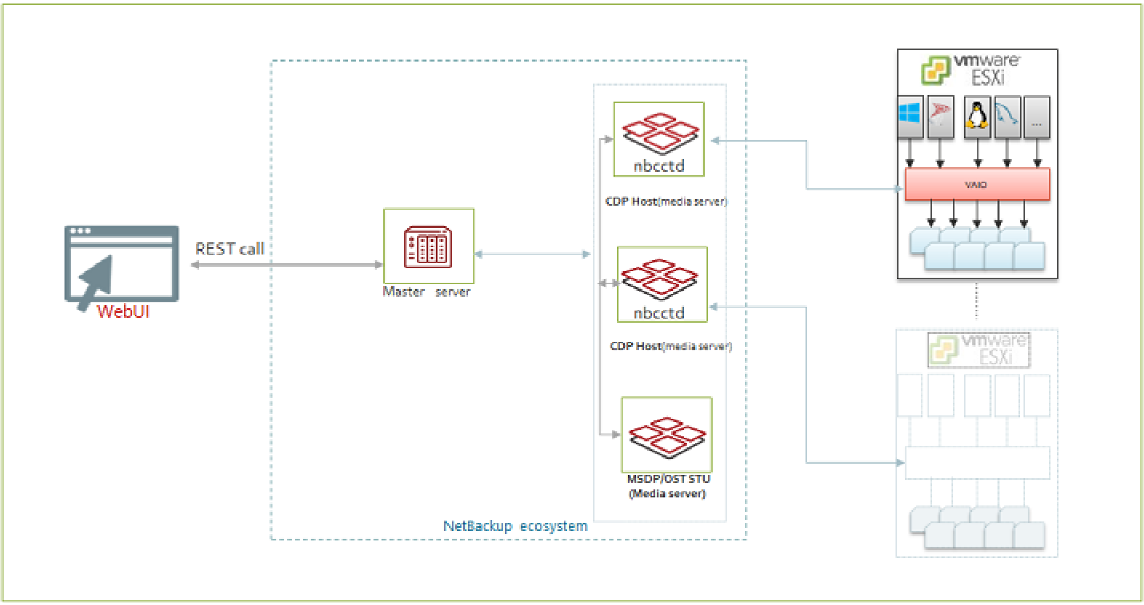

About continuous data protection

Continuous data protection (CDP) is a smart way to capture fast copies of backups for the VMware VMs, without stunning 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 RedHat Linux based NetBackup media server as 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.

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 vtstap storage policy in VMware vCenter

Operations on the NetBackup console

1. Create an MSDP or OST-based storage for the backup destination. See the *Configuring Storage* chapter of 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.

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 59. , 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. Recommenced is minimum 160 MB of data for each VM.

For example, you protect 40 VMs in a gateway. So, you need 40*160 MB = 6400 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 * 10-GB = 400-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 59. , 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 59. 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 59. 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 `nbcct.conf` file. For example, CCT_MAX_FULL_SYNC_REQS=7

Monitoring CDP jobs

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

See "The NetBackup dashboard" on page 13.

NetBackup Web UI Administrator's Guide.

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 7-2 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 7-2 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:

```
bpbackup -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 vmdk 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 IO 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 IO tapping policies are still attached to the VMDK of that VM, it continues to tap the IO'S and impact performance.

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

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 webUI. If you try to remove protection of that VM through webUI, you can see an error message saying: **Subscription ID not found**. This is expected behavior.

Chapter 8

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 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. 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.

5 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 40.
- **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 43.
- **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 42.

6 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**.

7 Review or change the **Options**.

See "Recovery options" on page 75.

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

See "Advanced recovery options" on page 75.

See "Advanced recovery options: Format of restored virtual disks" on page 76.

See "Advanced recovery options: Transport mode" on page 77.

9 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.

10 Resolve any errors.

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

11 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

Deletes any VM with the same display name that exists at the destination. That VM must be deleted before the recovery begins. Otherwise, the recovery fails.

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 VMware authorization credentials for the target virtual machine. These credentials must have administrator or superuser 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:

- You must provision VxUpdate packages for all platforms for which you have virtual machines where you want to perform agentless recovery.
- You must have an account with administrator or root permissions on the target virtual machine.
- The target VM is where the files are recovered. It must be powered on and have VMware Tools installed.
- The target VM should have at least one Paravirtual Controller with available LUNs or available space for 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.
- You must have the latest version of VMware Tools installed to perform agentless restores.

Limitations:

- 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 directory that is specified in the TMP or TEMP environment variable as the staging directory.
- 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, and 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. Under **Select files**, specify the files and folders you want recovered then click **Next**.
7. Under **Recovery target**, specify the target VM to which you want the files and folders recovered, as well as the administrator credentials for the target VM.
8. On **Recovery options**, specify additional recovery options for the restored files and folders.
9. After you 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, proceed with the restore.

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\tmp\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.

Chapter 9

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 (`PROXY_SERVERS`). 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 or ESXi 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.) To perform autodiscovery of VMware server objects at a different frequency: See "Change the autodiscovery frequency of VMware assets" on page 30. ■ 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: `Hmmm…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 `MSserver1.veritas.com`, the web UI must also be able to reach `MSserver1.veritas.com`. If the primary server uses a short name for the media server such as `MSserver1`, the web UI must be able to reach `https://MSserver1/...`
	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 `hosts` 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 `hosts` file of the PC or other host where the web UI is running.
	The hosts file location on Windows:
	`C:\Windows\System32\drivers\etc\hosts`
	The hosts file location on UNIX or Linux:
	`/etc/hosts`
	■ 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.