

NetBackup™ for Red Hat Virtualization Administrator's Guide

Release 11.0

Last updated: 2025-03-06

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

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

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

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

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

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

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

Cohesity, Inc.
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Cohesity Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Managing Red Hat Virtualization servers	7
	Quick configuration checklist to protect Red Hat Virtualization virtual machines	7
	Configuring secure communication between the Red Hat Virtualization server and NetBackup host	10
	ECA_TRUST_STORE_PATH for NetBackup servers and clients	13
	ECA_CRL_PATH for NetBackup servers and clients	14
	VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED for servers and clients	16
	VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients	17
	About the ports that NetBackup uses to communicate with Red Hat Virtualization	18
	Add or browse an Red Hat Virtualization manager	18
	Remove an Red Hat Virtualization manager	20
	Add a backup host to the NetBackup primary server	20
	Configure autodiscovery of the Red Hat Virtualization virtual machines	21
	Create an intelligent VM group	22
	Delete an intelligent VM group from Red Hat Virtualization	26
	Setting global limits on the use of Red Hat Virtualization resources	26
Chapter 2	Protecting RHV virtual machines	29
	Things to know before you protect Red Hat Virtualization virtual machines	29
	Protect Red Hat Virtualization VMs or intelligent VM groups	30
	Customize protection settings for a RHV asset	31
	Schedules	31
	Backup options	32
	Remove protection from VMs or intelligent VM groups	32
	View the protection status of VMs or intelligent VM groups	33

Chapter 3	Recovering RHV virtual machines	35
	Things to consider before you recover the Red Hat Virtualization virtual machines	35
	About the pre-recovery check	35
	Recover a Red Hat Virtualization virtual machine	36
	About the supported virtual disk formats and disk provisioning during VM recovery	38
Chapter 4	Troubleshooting RHV VM protection and recovery	41
	Troubleshooting tips for NetBackup for Red Hat Virtualization	41
	Error during the Red Hat Virtualization virtual machines discovery phase	42
	Error run into while backing up Red Hat Virtualization virtual machines	44
	Error while restoring Red Hat Virtualization virtual machines	45
Chapter 5	API and command line options for RHV	47
	Using APIs and command line options to manage, protect, or recover RHV VMs	47
	Additional information about the rename file	52
	Additional NetBackup options for Red Hat Virtualization configuration	53
	OVIRT_IMAGEIO_INACTIVITY_TIMEOUT option for NetBackup servers	53
	RHV_CREATEDISK_TIMEOUT option for NetBackup servers	53
	RHV_AUTODISCOVERY_INTERVAL option for NetBackup servers	54

Managing Red Hat Virtualization servers

This chapter includes the following topics:

- Quick configuration checklist to protect Red Hat Virtualization virtual machines
- Configuring secure communication between the Red Hat Virtualization server and NetBackup host
- About the ports that NetBackup uses to communicate with Red Hat Virtualization
- Add or browse an Red Hat Virtualization manager
- Remove an Red Hat Virtualization manager
- Add a backup host to the NetBackup primary server
- Configure autodiscovery of the Red Hat Virtualization virtual machines
- Create an intelligent VM group
- Delete an intelligent VM group from Red Hat Virtualization
- Setting global limits on the use of Red Hat Virtualization resources

Quick configuration checklist to protect Red Hat Virtualization virtual machines

Use NetBackup Web UI to protect the virtual machines that are created on the Red Hat Virtualization (RHV) platform.

You can also use APIs and command line options to protect and recover the Red Hat Virtualization VMs.

See “Using APIs and command line options to manage, protect, or recover RHV VMs” on page 47.

The following table describes the high-level steps or a checklist to protect the Red Hat Virtualization virtual machines:

Table 1-1 Configure and protect Red Hat Virtualization virtual machines using NetBackup

Step overview	Description and reference
Deploy NetBackup to protect Red Hat Virtualization VMs	<p>On a very high-level, to protect Red Hat Virtualization VMs, you need:</p> <ul style="list-style-type: none">■ NetBackup primary server■ NetBackup media server■ NetBackup client that can act as a backup host <p>NetBackup primary and media servers are supported on any supported server platform of NetBackup, whereas NetBackup client is supported on RHEL, SUSE, or a Windows host.</p> <p>NetBackup appliance including Flex appliance is also supported as a NetBackup primary, media server, or as a client that can act as a backup host.</p> <p>NetBackup uses an agentless architecture to protect the Red Hat Virtualization VMs. The communication between NetBackup and Red Hat Virtualization Manager happens through APIs.</p>
Configure an Red Hat Virtualization access host for backup and recovery	<p>An Red Hat Virtualization access host acts as a backup host and a recovery host during backup and recovery respectively. The access host is involved in the data movement during the backup and restore operations.</p> <p>If you plan to use a backup host that is not a NetBackup media server or an appliance, add the backup host to the NetBackup Red Hat Virtualization Access Hosts list.</p> <p>See “Add a backup host to the NetBackup primary server” on page 20.</p>
Enable secure communication between NetBackup and Red Hat Virtualization	<p>The following sections contain more information about setting up a secure communication between NetBackup and Red Hat Virtualization:</p> <ul style="list-style-type: none">■ Secure communication See “Configuring secure communication between the Red Hat Virtualization server and NetBackup host” on page 10.■ Communication ports See “About the ports that NetBackup uses to communicate with Red Hat Virtualization” on page 18.

Table 1-1 Configure and protect Red Hat Virtualization virtual machines using NetBackup (*continued*)

Step overview	Description and reference
Manage Red Hat Virtualization servers and intelligent VM groups	<ul style="list-style-type: none"> ■ Prerequisite: Adding an Red Hat Virtualization manager requires the Default Red Hat Virtualization Administrator role. ■ Managing Red Hat Virtualization servers See “Add or browse an Red Hat Virtualization manager” on page 18. ■ Managing intelligent VM groups See “Create an intelligent VM group” on page 22. See “Delete an intelligent VM group from Red Hat Virtualization” on page 26.
Protect the Red Hat Virtualization VMs	<ul style="list-style-type: none"> ■ Prerequisite: Adding an Red Hat Virtualization manager requires the Default Red Hat Virtualization Administrator role. ■ Best practices See “Things to know before you protect Red Hat Virtualization virtual machines” on page 29. ■ Protecting virtual machines See “Protect Red Hat Virtualization VMs or intelligent VM groups” on page 30.
Consider setting global limits on the use of Red Hat Virtualization resources	<p>When you protect VMs automatically when they are created, over a period of time the number of VMs protected concurrently can grow large. The large number of concurrent backups can affect the Red Hat Virtualization performance as well as backup performance.</p> <p>You can set the global limits to manage the Red Hat Virtualization resources efficiently.</p> <p>See “Setting global limits on the use of Red Hat Virtualization resources” on page 26.</p>

Additional references

The following table describes the high-level steps or a checklist to recover the Red Hat Virtualization virtual machines and additional information:

Table 1-2 Red Hat Virtualization VM recovery and additional information

Step overview	Description and reference
Removing protection from VMs	See “Remove protection from VMs or intelligent VM groups” on page 32.

Table 1-2

Red Hat Virtualization VM recovery and additional information
(continued)

Step overview	Description and reference
Recover the protected Red Hat Virtualization VMs	<ul style="list-style-type: none">■ Best practices See “Things to consider before you recover the Red Hat Virtualization virtual machines” on page 35.■ Supported disk format and disk provisioning See “About the supported virtual disk formats and disk provisioning during VM recovery” on page 38.■ Recover Red Hat Virtualization VMs See “Recover a Red Hat Virtualization virtual machine” on page 36.
API and command line options to protect Red Hat Virtualization VMs	<p>You can use NetBackup APIs or command line options to protect and recover Red Hat Virtualization VMs</p> <ul style="list-style-type: none">■ See “Using APIs and command line options to manage, protect, or recover RHV VMs” on page 47.■ See “Additional NetBackup options for Red Hat Virtualization configuration” on page 53.
Troubleshooting information	<ul style="list-style-type: none">■ Use the following information to troubleshoot issues regarding Red Hat Virtualization protection or recovery

Configuring secure communication between the Red Hat Virtualization server and NetBackup host

NetBackup can now validate Red Hat Virtualization server certificates using their root or intermediate certificate authority (CA) certificates.

Only PEM certificate format is supported for virtualization servers.

See “VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED for servers and clients” on page 16.

The following procedure is applicable for the NetBackup primary server and all Red Hat Virtualization access hosts.

To configure secure communication between Red Hat Virtualization server and Red Hat Virtualization access host

- 1** Configure an external certificate authority trust store on the Red Hat Virtualization access host.
- 2** Add CA certificates of the required Red Hat Virtualization server in the trust store on the access host.

In case of Windows certificate store, add the CA certificate to the Windows Trusted Root Certification Authorities.

Use the following command:

```
certutil.exe -addstore -f "Root" certificate filename
```

- 3 Use the `nbsetconfig` command to configure the following NetBackup configuration options on the access host:

For more information on the configuration options, refer to the NetBackup Administrator's Guide, Volume I.

`ECA_TRUST_STORE_PATH`

Specifies the file path to the certificate bundle file that contains all trusted root CA certificates.

This option is specific to file-based certificates. You should not configure this option if Windows certificate store is used.

If you have already configured this external CA option, append the Red Hat Virtualization CA certificates to the existing external certificate trust store.

If you have not configured the option, add all the required Red Hat Virtualization server CA certificates to the trust store and set the option.

See “`ECA_TRUST_STORE_PATH` for NetBackup servers and clients” on page 13.

`ECA_CRL_PATH`

Specifies the path to the directory where the certificate revocation lists (CRL) of the external CA are located.

If you have already configured this external CA option, append the Red Hat Virtualization server CRLs to the CRL cache.

If you have not configured the option, add all the required CRLs to the CRL cache and then set the option.

See “`ECA_CRL_PATH` for NetBackup servers and clients” on page 14.

`VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED`

This option affects Nutanix AHV, Red Hat Virtualization, and VMware secure communication. Without this option, the secure or insecure communication with workload is decided by each workload and plug-in separately.

For more information, refer to the respective workload Administrator's Guide.

For Red Hat Virtualization, secure communication is enabled by default.

This option lets you skip the security certificate validation.

See

“`VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` for servers and clients” on page 16.

`VIRTUALIZATION_CRL_CHECK`

Lets you validate the revocation status of the virtualization server certificate against the CRLs.

By default, the option is disabled.

See "VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients" on page 17.

For more information on external CA support, refer to the NetBackup Security and Encryption Guide.

ECA_TRUST_STORE_PATH for NetBackup servers and clients

The `ECA_TRUST_STORE_PATH` option specifies the file path to the certificate bundle file that contains all trusted root CA certificates.

This certificate file should have one or more certificates in PEM format.

Do not specify the `ECA_TRUST_STORE_PATH` option if you use the Windows certificate store.

The trust store supports certificates in the following formats:

- PKCS #7 or P7B file having certificates of the trusted root certificate authorities that are bundled together. This file may either be PEM or DER encoded.
- A file containing the PEM encoded certificates of the trusted root certificate authorities that are concatenated together.

This option is mandatory for file-based certificates.

The root CA certificate in Cloudera distribution can be obtained from the Cloudera administrator. It may have a manual TLS configuration or an Auto-TLS enabled for the Hadoop cluster. For both cases, NetBackup needs a root CA certificate from the administrator.

The root CA certificate from the Hadoop cluster can validate the certificates for all nodes and allow NetBackup to run the backup and restore process in case of the secure (SSL) cluster. This root CA certificate is a bundle of certificates that has been issued to all such nodes.

Certificate from root CA must be configured under `ECA_TRUST_STORE_PATH` in case of self-signed, third party CA or Local/Intermediate CA environments. For example: In case of AUTO-TLS enabled Cloudera environments, you can typically find the root CA file named with `cm-auto-global_cacerts.pem` at path `/var/lib/cloudera-scm-agent/agent-cert`. For more details, refer Cloudera documentation.

Table 1-3 ECA_TRUST_STORE_PATH information

Usage	Description
Where to use	<p>On NetBackup servers or clients.</p> <p>If certificate validation is required for VMware, Red Hat Virtualization servers, or Nutanix AHV, this option must be set on the NetBackup primary server and respective access hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA).</p>
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <p><i>ECA_TRUST_STORE_PATH = Path to the external CA certificate</i></p> <p>For example: <code>c:\rootCA.pem</code></p> <p>If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/</code>.</p>
Equivalent UI property	No equivalent exists.

ECA_CRL_PATH for NetBackup servers and clients

The `ECA_CRL_PATH` option specifies the path to the directory where the Certificate Revocation Lists (CRL) of the external certificate authority (ECA) are located.

These CRLs are copied to NetBackup CRL cache. Revocation status of the external certificate is validated against the CRLs from the CRL cache.

CRL in the CRL cache is periodically updated with the CRL on the location that is specified for `ECA_CRL_PATH` based on the `ECA_CRL_PATH_SYNC_HOURS` option.

If the `ECA_CRL_CHECK` or `HADOOP_CRL_CHECK` option is not set to `DISABLE` (or 0) and the `ECA_CRL_PATH` option is not specified, NetBackup downloads the CRLs from the URLs that are specified in the CRL distribution point (CDP) and uses them to verify revocation status of the peer host's certificate.

Note: For validating the revocation status of a virtualization server certificate, the `VIRTUALIZATION_CRL_CHECK` option is used.

For validating the revocation status of a Hadoop server certificate, the `HADOOP_CRL_CHECK` option is used.

Table 1-4 ECA_CRL_PATH information

Usage	Description
Where to use	<p>On NetBackup servers or clients.</p> <p>If certificate validation is required for VMware, Red Hat Virtualization servers, Nutanix AHV, or Hadoop, this option must be set on the NetBackup primary server and respective access or backup hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA).</p> <p>If certificate validation is required for VMware, Red Hat Virtualization servers, or Hadoop, this option must be set on the NetBackup primary server and respective access or backup hosts, irrespective of the certificate authority that NetBackup uses for host communication (NetBackup CA or external CA).</p>
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format to specify a path to the CRL directory:</p> <pre>ECA_CRL_PATH = Path to the CRL directory</pre> <p>For example:</p> <pre>ECA_CRL_PATH = /usr/eca/crl/eca_crl_file.crl</pre> <p>If you use this option on a Flex Appliance application instance, the path must be <code>/mnt/nbdata/hostcert/crl</code>.</p>
Equivalent UI property	No equivalent exists.

VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED for servers and clients

The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option enables the validation of virtualization server certificates using its root or intermediate certificate authority (CA) certificates.

Before you enable the option, review the steps from the 'Validating VMware virtualization server certificates in NetBackup ' section in the NetBackup for VMware Administrator's Guide.

By default, the `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option is set to `UNDEFINED`.

The security certificate validation is enabled for Red Hat Virtualization and Nutanix AHV servers, but is disabled for VMware servers.

Note: In a scenario where an external CA can be configured for one virtualization server, but not for the other, two separate backup hosts must be used. The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` option must be set to `YES` for the backup host where the external CA can be configured. The `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` must be set to `YES` for the backup host where the external CA can be configured. The option must be set to `NO` for the other backup host.

Table 1-5 `VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED` information

Usage	Description
Where to use	On NetBackup primary server or all access hosts.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format to enable certificate validation for the Red Hat Virtualization, VMware, or Nutanix AHV servers:</p> <pre>VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED = YES</pre>
Equivalent NetBackup web UI property	No equivalent exists.

VIRTUALIZATION_CRL_CHECK for NetBackup servers and clients

The `VIRTUALIZATION_CRL_CHECK` option lets you specify the revocation check level for external certificates of the virtualization server. Based on the check, revocation status of the virtualization server certificate is validated against the certificate revocation list (CRL) during host communication.

By default, the `VIRTUALIZATION_CRL_CHECK` option is disabled. If you want to validate the revocation status of the virtualization server certificate against certificate revocation list (CRL), set the option to a different value.

You can choose to use the CRLs from the directory that is specified for the `ECA_CRL_PATH` configuration option or the CRL distribution point (CDP).

See “`ECA_CRL_PATH` for NetBackup servers and clients” on page 14.

Table 1-6 `VIRTUALIZATION_CRL_CHECK` information

Usage	Description
Where to use	On NetBackup primary server or all access hosts.
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>For information about these commands, see the NetBackup Commands Reference Guide.</p> <p>Use the following format:</p> <pre>VIRTUALIZATION_CRL_CHECK = CRL check</pre> <p>You can specify one of the following:</p> <ul style="list-style-type: none"> ■ DISABLE (or 0) - Revocation check is disabled. Revocation status of the certificate is not validated against the CRL during host communication. This is the default value. ■ LEAF (or 1) - Revocation status of the leaf certificate is validated against the CRL. ■ CHAIN (or 2) - Revocation status of all certificates from the certificate chain are validated against the CRL.
Equivalent NetBackup web UI property	No equivalent exists.

About the ports that NetBackup uses to communicate with Red Hat Virtualization

The following table describes the ports that NetBackup requires to communicate with Red Hat Virtualization:

Table 1-7

Ports required by NetBackup to communicate with Red Hat Virtualization

Port	Protocol	Destination	Purpose
80, 443	TCP	Red Hat Virtualization Manager	Provides HTTP and HTTPS access to the Red Hat Virtualization Manager
54322	TCP	Red Hat Virtualization Hosts (Red Hat Enterprise Linux hosts)	Required for communication with the ImageIO daemon (ovirtimageio-daemon)
54323	TCP	Red Hat Virtualization Manager (ImageIO Proxy server)	Required for communication with the ImageIO Proxy (ovirtimageio-proxy)

Add or browse an Red Hat Virtualization manager

You can add and browse Red Hat Virtualization managers and their credentials.

To add Red Hat Virtualization managers and their credentials

- 1

On the left, click **Workloads > Red Hat Virtualization** then click the **Red Hat Virtualization managers** tab.
- 2

Click **Add** to add an Red Hat Virtualization manager and enter the following:
 - Red Hat Virtualization manager name
 - Access credentials
 - Select a backup host using the **Backup host for validation**
 - Port number (optional)

Note: NetBackup recommends that you use the FQDN to add the Red Hat Virtualization Manager. Using an IP address or short name to add the Red Hat Virtualization Manager can create duplicate entries and cause issues during RBAC enforcement.

See “Add a backup host to the NetBackup primary server” on page 20.

3 Click **Save**.

4 To add another Red Hat Virtualization Manager credentials, click **Add**.

Inline actions on Red Hat Virtualization manager

You can run the following inline actions on an Red Hat Virtualization manager:

- **Discover:** Manually discovers the VM assets that belong to the selected Red Hat Virtualization manager.
- **Edit:** Modify the Red Hat Virtualization manager credentials.
- **Delete:** Removes the Red Hat Virtualization manager.

Bulk actions on Red Hat Virtualization managers

You can select one or more Red Hat Virtualization managers and run the following bulk actions:

- **Validate credentials:** Validates the credentials of the Red Hat Virtualization manager.
- **Delete:** Removes the Red Hat Virtualization managers.

Browse an Red Hat Virtualization manager

You can browse the Red Hat Virtualization managers and clusters to locate VMs and view their details such as their protection plans and recovery points.

To browse Red Hat Virtualization managers

- 1 On the left, click **Red Hat Virtualization**.
- 2 Click **Workloads > Red Hat Virtualization** to begin searching.

The list includes the Red Hat Virtualization managers and clusters that you have access to.

The tab shows the Red Hat Virtualization managers and clusters that you can access in the following hierarchy:

```
All
RHV_Managers
  RHV_Manager1
    Cluster1
    Cluster2
  RHV_Manager2
    Cluster3
    Cluster4
```

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

- 3 Click the Red Hat Virtualization manager to view details.
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 an Red Hat Virtualization manager

Red Hat Virtualization managers can be removed by means of the bulk or the inline actions from the **Red Hat Virtualization manager** tab.

When you remove an Red Hat Virtualization manager, you can no longer protect the Red Hat Virtualization VMs from the NetBackup.

See “Add or browse an Red Hat Virtualization manager” on page 18.

Add a backup host to the NetBackup primary server

The backup host or appliance acts as a channel to establish an indirect communication between the NetBackup primary server and the Red Hat Virtualization manager. The backup host is a NetBackup client that performs backups or restores on behalf of the virtual machines.

A NetBackup primary and media server can also be configured as the backup host. However, you do not need to add the primary or the media server acting as a backup host to the access hosts list for Red Hat Virtualization.

The secure communication happens by APIs and uses SSL.

Note: SSL requires that all backup hosts have ECA certificates.

Communication between Red Hat Virtualization and backup host requires an open port.

The following operating systems are supported for your backup host:

- Windows
- Red Hat Linux
- SUSE

When the backup host is not a NetBackup media server or an appliance, you need to add the backup host to the access hosts list for Red Hat Virtualization.

To add an access host

- 1 On the left, click **Workloads > Red Hat Virtualization**.
- 2 At the top right, click **RHV settings > Access hosts**. Then click **Add**.
- 3 Enter the name of the access host.
- 4 Click **Save**.

Configure autodiscovery of the Red Hat Virtualization virtual machines

You can edit the frequency by which NetBackup performs discovery of the Red Hat Virtualization virtual machines.

To enable autodiscovery

- 1 On the left, click **Workloads > Red Hat Virtualization** and then click the **Virtual machines** tab.
- 2 On the top right, click **RHV settings > Autodiscovery**.
- 3 Click **Edit** to set the frequency for autodiscovery. Select the frequency interval and click **Save**. The default frequency is 8 hours.

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: A background task adds the newly discovered VMs that match the query to the intelligent VM group. This background task runs 5 minutes after the start of the NetBackup Web Management service. Then the task runs every 30 minutes.

To create an intelligent VM group

- 1 On the left, click **Workloads > Red Hat Virtualization**.
- 2 Click the **Intelligent VM groups** tab and then click **Add**.
- 3 Enter a name and description for the group.

The intelligent VM group display name length must be between 1 to 256 characters.

- 4 In the **Select virtual machines** pane, select the appropriate **Red Hat Virtualization manager**.

Note: The web UI lists the servers that you can access based on your role and permissions (RBAC).

- Select the default query: **Include all VMs**.
When the protection plan runs, all VMs that are part of the Red Hat Virtualization manager are selected for backup.
 - Create your own query: Click **Add condition**.
- 5 To add a condition, use the dropdown lists to select a keyword and operator and then enter a value.

See the section called “Query options for creating intelligent VM groups” on page 24.

The following is an example query:

[+ Condition](#)

<div style="display: flex; align-items: center;"> <div style="border-bottom: 1px solid #ccc; flex-grow: 1; margin-right: 5px;">displayName</div> <div style="margin: 0 5px;">▼</div> <div style="border-bottom: 1px solid #ccc; flex-grow: 1; margin-right: 5px;">Contains</div> <div style="margin: 0 5px;">▼</div> <div style="border-bottom: 1px solid #ccc; flex-grow: 1; margin-right: 5px;">prod</div> <div style="margin-left: 5px;"> </div> </div>
--

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

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

The screenshot shows a query builder interface. At the top, there are two buttons: "AND" (selected) and "OR". To the right of these buttons are two links: "+ Condition" and "+ Sub-query". Below the buttons, there are two rows of conditions. The first row has the keyword "displayName", the operator "Contains", and the value "prod". The second row has the keyword "tagName", the operator "=", and the value "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 that is named `eng`, that VM is not added to the group.

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

The screenshot shows the same query builder interface as the previous one, but with the "OR" button selected instead of "AND". The two conditions remain the same: "displayName" Contains "prod" and "tagName" equals "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 that is named `eng` (regardless of the display name).

6 To test the query, click **Preview**.

Note: 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.

7 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 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. See the section called “Query options for creating intelligent VM groups” on page 24.

Note: If you use filters in an intelligent group, NetBackup may not display the accurate list of VMs that match the filter if the VM or the Red Hat Virtualization server has non-English characters. However, during the backup the correct VMs are selected even though the VM attributes are non-English. This behavior is only in viewing the VMs in NetBackup Web UI.

Query options for creating intelligent VM groups

Table 1-8 Query keywords

Keyword	Description
<code>cluster</code>	The name of the cluster where the VMs reside. Not case-sensitive when the protection plan runs.
<code>datacenter</code>	The name of the datacenter. Not case-sensitive when the protection plan runs.

Table 1-8 Query keywords (*continued*)

Keyword	Description
<code>displayName</code>	The VM's display name. Case-sensitive when the protection plan runs.
<code>storageDomainName</code>	The name of the storage domain. Case-sensitive when the protection plan runs.
<code>tagName</code>	The name of the VM's tag. Case-sensitive when the protection plan runs.
<code>templateName</code>	The name of the VM template. Case-sensitive when the protection plan runs.
<code>vmUuid</code>	The instance UUID of the VM. For example: 501b13c3-52de-9a06-cd9a-ecb23aa975d1 Not case-sensitive when the protection plan runs.

Table 1-9 Query operators

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

Delete an intelligent VM group from Red Hat Virtualization

Use the following procedure to remove an intelligent VM group from the Red Hat Virtualization workload.

To delete an intelligent VM group

- 1 On the left, click **Workloads > Red Hat Virtualization**.
- 2 Click on the **Intelligent VM groups** tab.
- 3 (Conditional) If the group is protected, you first need to remove it from the protection plan. Click on the group. Then click **Remove protection**.
- 4 On the **Intelligent VM groups** tab, locate and select the group. Then click **Delete**.

Setting global limits on the use of Red Hat Virtualization resources

You can control the number of simultaneous backups that can be performed on an Red Hat Virtualization resource type. The settings apply to all NetBackup policies for the currently selected primary server.

For example, to avoid overloading the overall Red Hat Virtualization cluster, you can place a limit on the number of concurrent backup jobs per Red Hat Virtualization Cluster. To control input output overhead on the storage domain array, you can limit the number of concurrent backups per storage domain.

The following resource limits are available for Red Hat Virtualization:

- **Backup Jobs per Cluster**
- **Backup Jobs per DataCenter**
- **Backup Jobs per StorageDomain**

To set limits on the use of Red Hat Virtualization resources

- 1 On the left, click **Workloads > Red Hat Virtualization**.
- 2 On the top right, click **RHV settings > Resource limits**.

For each resource type, the default is 0, (No Limit).

- 3 Select the resource type you want to change and then click **Edit**.
- 4 Choose from the following options.

Set a global limit for a RHV resource type.	<p>Locate the Global setting and select the Limits value that you want to apply.</p> <p>This value limits the number of simultaneous backups that are performed for the resource type.</p>
Set a limit for a specific RHV resource.	<p>Click Add.</p> <p>Provide a name for the resource limit.</p> <p>Select the Limits value that you want to apply.</p> <p>This value limits the number of simultaneous backups that are performed for the specific resource.</p>

Example

The following example illustrates how these limits control simultaneous backups. The settings must be done according to Red Hat Virtualization configuration in your environment.

When NetBackup connects to the Red Hat Virtualization environment for backup, it makes 1 connection per disk present on the VM. So, if a VM has 2 disks then NetBackup would make 2 connections with the Red Hat Virtualization node.

So, consider a case where Red Hat Virtualization manager is managing 2 clusters with 2 nodes in each cluster. Let's consider every node is hosting 20 VMs with 2 disks on each VM.

When the job runs, there would be 80 concurrent jobs start if no Resource Limit is set which is default behavior. Red Hat Virtualization recommends up to 10 disk connections concurrently per node in the cluster. In the example of each VM with two disks, 5 VMs per node can be ideally backed up concurrently. Hence, in this case of a cluster with two nodes, it is recommended to backup upto 10 VMs concurrently. When the **Backup Jobs per Cluster** is set to 10, this limit is enforced.

Since one data center can manage multiple clusters, **Backup Jobs per DataCenter** resource limit could be higher than **Backup Jobs per Cluster**. You need to determine the values by evaluating the effect of backup on overall environment.

Storage domain in Red Hat Virtualization can server multiple clusters in a data center. It serves to protection VM as well as backup. Its performance is dependent on type of storage technology – FC, iSCSI, NFS, gluster etc. Hence limit on storage domain using **Backup Jobs per StorageDomain** can be set based on characteristics of storage domain technology and limit can be higher than **Backup Jobs per Cluster**.

Protecting RHV virtual machines

This chapter includes the following topics:

- Things to know before you protect Red Hat Virtualization virtual machines
- Protect Red Hat Virtualization VMs or intelligent VM groups
- Customize protection settings for a RHV asset
- Remove protection from VMs or intelligent VM groups
- View the protection status of VMs or intelligent VM groups

Things to know before you protect Red Hat Virtualization virtual machines

- You cannot backup the same Red Hat Virtualization VM concurrently.
- The VMs without virtual disks cannot be protected.
- The following QCOW2 image attributes are not supported:
 - Compressed cluster
 - Encrypted disks
 - Virtual disks with internal snapshots
- If the VM virtual disks are locked when the NetBackup services shutdown or crash during a backup, use Red Hat Virtualization's `unlock_entity` command to unlock the disks. If the disks are not unlocked, the subsequent backups might fail.

See “Error run into while backing up Red Hat Virtualization virtual machines” on page 44.

- On a file storage (NFS), a QCOW2 disk gets restored as raw disk (thin provision) because of an Red Hat Virtualization limitation.
- A thin dependent cloned VM is restored as an independent cloned VM.
- If you want to use a storage that is not available through the NetBackup Web UI like a tape or basic disk based storage unit, you can use APIs or command line options to protect the VMs.
- For the minimum permissions required to perform VM backup and restore, see https://www.veritas.com/content/support/en_US/article.100050733

Protect Red Hat Virtualization VMs or intelligent VM groups

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

Note: The RBAC role that is assigned to you must give you access to the assets that you want to manage and to the protection plans that you want to use.

To protect Red Hat Virtualization VMs or VM groups

- 1 On the left, click **Workloads > Red Hat Virtualization**.
 - 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 You can adjust one or more of the following settings:
 - **Schedules and retention**
Change the backup start window.
 - **Backup options**
Adjust the server or host to use for backups.
 - 5 Click **Protect**.
- The results of your choices appear under **Virtual machines** or **Intelligent VM groups**.

Customize protection settings for a RHV asset

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

- See “Schedules” on page 31.

To customize protection settings for a RHV asset

- 1 On the left, click **Workloads > RHV**.
- 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 31.
- 5 Click **Protect**.

Schedules

The following schedule settings are included in a protection plan.

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

- Start window

Table 2-1 Schedule options for protection plans

Option	Description
Backup type	The type of backup that the schedule controls.
Recurrence (frequency)	How frequently or when to run the backup.
Keep for (retention)	How long to keep the files that were backed up by the schedule.

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

Option	Description
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

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

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.

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.

To remove protection from a VM or intelligent VM group

- 1
- On the left, click **Red Hat Virtualization**.
- 2
- On the **Virtual machines** tab or **Intelligent VM groups** tab, select the VM or the intelligent VM group.
 - For a VM, scroll down and click **Remove protection**.
 - For an intelligent VM group, scroll down and click the lock symbol and then click **Remove protection**.

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 protection 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 **Red Hat Virtualization**.
- 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, click the VM or the intelligent VM group.

The **Protection** tab shows the details of the plans that the asset is subscribed to.
- 3 If the asset is not protected, click **Add protection** to select a protection plan.
See “Protect Red Hat Virtualization VMs or intelligent VM groups” on page 30.

Recovering RHV virtual machines

This chapter includes the following topics:

- Things to consider before you recover the Red Hat Virtualization virtual machines
- About the pre-recovery check
- Recover a Red Hat Virtualization virtual machine
- About the supported virtual disk formats and disk provisioning during VM recovery

Things to consider before you recover the Red Hat Virtualization virtual machines

- Ensure that the recovery or backup host, that is added to the Red Hat Virtualization access hosts, can communicate with the Red Hat Virtualization Manager through a port .
- An Red Hat Virtualization VM that has chain of disks due to user snapshot or dependency on template, cannot retain the disk chain after a restore.
- Compressed virtual disks are not protected and cannot be recovered.

About the pre-recovery check

The pre-recovery check verifies the following:

- Usage of supported characters and the length in the display name.
- Existence of a VM with the same display name.

- Connectivity with the Red Hat Virtualization server and Red Hat Virtualization credential validation.
- Availability of the Red Hat Virtualization cluster.
- Available space with the storage domain.

Recover a Red Hat Virtualization virtual machine

You can recover a VM to the original location where it existed when it was backed up or to a different location.

To recover a VM

- 1 On the left, click **Workloads > Red Hat Virtualization**.
- 2 On the **Virtual machines** tab, locate and select the VM. Then click **Recover**.
- 3 In the calendar view on the left, click the date on which the backup occurred.
The available images are listed in rows with the backup timestamp for each image. The dates that have a green dot have a recovery point for that VM.
- 4 On the image you want to recover, click **Recover**.
- 5 To recover to the original location, do not modify the **Recovery targets**.

To recover to a different location:

Modify the **Display name**. Select the **Red Hat Virtualization manager** and the **Red Hat Virtualization cluster** where you want to recover the VM.

If you cannot change the **Red Hat Virtualization cluster**, see the following topic.

See “Error while restoring Red Hat Virtualization virtual machines” on page 45.

- 6 Click **Add** to add a storage domain and select the appropriate storage domain.
Select different storage domain for the virtual disks or select **Use the same storage domain for all virtual disks** to use the same storage domain for all the virtual disks. Click **Next**.
- 7 Review or change the following options:

Recovery options:

Overwrite existing virtual machine	If a VM with the same UUID or same name exists at the destination and if this option is selected then that VM is deleted. If a VM with the same UUID or same name exists at the destination and if this option is not selected then the restore fails and an error is displayed.
---	---

Power on after recovery	Automatically turns on the VM when the recovery is complete.
Recovery host	A backup host that you can use during recovery. By default, the recovery host is the backup host that was used during a backup.

Advanced settings:

Retain original network configuration	<p>The restored VM automatically connects to the original network using the retained NICs.</p> <p>Do not enable this option if:</p> <ul style="list-style-type: none"> ■ 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.
Create a new VM UUID	Restores the VM with a new UUID instead of the original UUID.
Remove tag associations	Removes the tags which were associated with the VM at the time of backup.

Format of restored virtual disks:

Original provision	Restores the VM's virtual disks with their original provisioning.
Thick provision	Configures the restored virtual disks in the thick format. The virtual disk space is pre-allocated when the disk is created.
Thin provision	Configures the restored virtual disks in the thin format. Restores only the populated blocks and the new blocks are allocated as required.

8 Click **Next** to run the **Pre-recovery check**.

The **Pre-recovery check** validates all the recovery parameters and displays errors, if any. You can fix the errors before starting the recovery.

9 Click **Start recovery**.

If you refresh the display, the **Restore activity** tab shows the job progress.

For information on the recovery status codes, see the NetBackup Status Codes Reference Guide.

About the supported virtual disk formats and disk provisioning during VM recovery

Red Hat Virtualization supported allocation methods for virtual disks

Red Hat Virtualization supports the following allocation methods for virtual disks:

- Pre-allocated (Thick provision)
Pre-allocated indicates fully-allocated raw disks.
- Thin provision
Thin provisioned disks are of the following types:
 - Raw sparse (default on file storage such as NFS)
 - QCOW2 (default on block storage such as FC, SAN, iSCSI)The thin provisioned virtual disk that is created on block storage is always in QCoW2 format.

Virtual disk provisioning for Red Hat Virtualization VM recovery

Based on the disk provisioning option that you select in NetBackup, the virtual disks are created as described in the following table:

Table 3-1 Virtual disk provisioning for Red Hat Virtualization VM recovery

The disk provisioning option that is selected during restore	Original disk format during backup		
	RAW sparse	RAW pre-allocated	QCOW2
Original or default	RAW sparse (QCOW2 on block storage)	RAW pre-allocated	RAW sparse (QCOW2 on block storage)
Thin	RAW sparse (QCOW2 on block storage)	RAW sparse (QCOW2 on block storage)	RAW sparse (QCOW2 on block storage)
Thick	RAW pre-allocated	RAW pre-allocated	RAW pre-allocated

Disk formats for VM templates

- VM templates can have the disks that have RAW or QCOW2 format.
- Storage allocation can be thin (dependent) or clone (independent).
In Clone (independent) allocation, the contents of template disks are copied to the VM disk.

About the supported virtual disk formats and disk provisioning during VM recovery

In Thin (dependent) allocation, the template disks are referred as base disks for the VM.

- When multiple VMs are deployed from the same template with Thin allocation, then the VMs share the template disks.

Troubleshooting RHV VM protection and recovery

This chapter includes the following topics:

- Troubleshooting tips for NetBackup for Red Hat Virtualization
- Error during the Red Hat Virtualization virtual machines discovery phase
- Error run into while backing up Red Hat Virtualization virtual machines
- Error while restoring Red Hat Virtualization virtual machines

Troubleshooting tips for NetBackup for Red Hat Virtualization

For more information about Red Hat Virtualization troubleshooting, check the following details:

- For discovery job failures:
 - Check the **Job details** section for the job in Activity monitor.
 - Check the `ncfnbcs` log.
- For snapshot job failures:
 - Check the **Job details** section for the job in Activity monitor.
 - Check the `bpfis` log.
 - For Red Hat Virtualization-related errors, check **Events** section on Red Hat Virtualization manager console.
- For backup job failures:

- Check the **Job details** section for the job in Activity monitor.
- Check the `bpbkar` and `VxMS` logs.
- For Red Hat Virtualization-related errors, check **Events** section on Red Hat Virtualization manager console.
- For restore job failures:
 - Restore job fails with error 2822 (Hypervisor policy restore error)
 - Check the **Job details** section for the job in Activity monitor.
 - Check the `bprd`, `bpVMutil`, `VxMS`, or `ncfnbrestore` logs.
 - For Red Hat Virtualization-related errors, check **Events** section on Red Hat Virtualization manager console.
- For upgrades:

When you upgrade NetBackup primary server, the assets (VM and Intelligent VM groups) are migrated from the older database tables to the new database tables.

A background migration task runs when the NetBackup Web Management Console starts. During this migration, the NetBackup Web UI displays the following message:

The migration process is in progress. The provided data might be inconsistent and incomplete.

Error during the Red Hat Virtualization virtual machines discovery phase

The following table describes the problem that might occur when you try to discover Red Hat Virtualization virtual machines.

Table 4-1 Error run into during the Red Hat Virtualization virtual machines discovery phase

Error message or cause	Explanation and recommended action
The Red Hat Virtualization assets are not discovered after the correct Red Hat Virtualization manager credentials are added. The VM discovery operation fails.	<p>The maximum allowed length of the Red Hat Virtualization manager name is 255 characters, however, if the characters exceed 95, the asset discovery fails.</p> <p>Workaround:</p> <ul style="list-style-type: none">■ Ensure that the Red Hat Virtualization manager name has 95 or fewer characters.

Table 4-1 Error run into during the Red Hat Virtualization virtual machines discovery phase (*continued*)

Error message or cause	Explanation and recommended action
The discovery job fails with error 200. (Scheduler found no backups or clients to deploy NetBackup)	<p>Ensure that the query specified in the policy or intelligent VM group is correct. The VMs that need protection are added recently to Red Hat Virtualization manager or the VM configuration has changed and the autodiscovery or discover now was not triggered.</p> <ul style="list-style-type: none"> Run discover now and retry the backup. The maximum allowed length of the Red Hat Virtualization manager name is 255 characters, however, if the characters exceed 95, the asset discovery fails. Workaround: Ensure that the Red Hat Virtualization manager name has 95 or fewer characters. The asset discovery does not work if the Red Hat Virtualization manager credentials are added using <code>tpconfig</code>. Workaround: From NetBackup web UI, run Discover for the specified Red Hat Virtualization manager. Ensure that you add the Red Hat Virtualization manager credentials using the APIs or the NetBackup web UI.
The GET asset API does not work when you use the <code>tolower</code> and <code>toupper</code> functions.	<p>NetBackup Web UI:</p> <p>For the filter in the intelligent groups, NetBackup web UI might not display accurate list of the Red Hat Virtualization VMs that match the filter if the VM or the Red Hat Virtualization server has non-English characters. However, during the backup, correct VMs are selected even though its attributes are non-English. This behavior is only in viewing the VMs in the NetBackup web UI.</p> <p>GET assets API of asset service:</p> <p>The GET Assets API does not give the desired result if the <code>tolower</code> or <code>toupper</code> functions are used together for the assets that have non-English characters.</p>
There is a delay in the API response.	For a large number of Red Hat Virtualization assets, if you add large and random offsets to an API request, there is an increase in the processing time that leads to a delay in the API response.

Error run into while backing up Red Hat Virtualization virtual machines

The following table describes the problem that might occur when you back up Red Hat Virtualization virtual machines:

Table 4-2 Error while backing up Red Hat Virtualization virtual machines

Error message or cause	Explanation and recommended action
After a NetBackup backup operation, the VM snapshot on the Red Hat Virtualization manager is not deleted.	<p>If a disk attached to the VM is in an inactive state, then the Red Hat Virtualization manager does not delete the VM snapshot after a backup operation is complete.</p> <p>Workaround:</p> <ul style="list-style-type: none">Before the backup operation, verify the state of the disks that are attached to the VM and ensure that they are active.Ensure that the disks are not attached while the VM is running, thus preventing the disk to be in an inactive state.
VM backup fails with the following error: "The virtual machine has no disks or contains only Raw Device Mappings for disks: Status 25"	<p>This temporary error might occur when the VM snapshot is not available for the backup operation. The backup job is successful on the second attempt.</p>
Unable to remove the older snapshots from the Red Hat Virtualization manager when the disk is in a locked state. The following error is displayed: A NetBackup snapshot of the virtual machine exist.	<p>Workaround:</p> <ul style="list-style-type: none">Refer to the following article for steps to unlock the disk: https://access.redhat.com/solutions/396753Manually remove the older snapshots from Red Hat Virtualization manager.
Accelerator option does not work.	<p>When a backup policy is created using APIs and the use Accelerator option is enabled, the policy gets created but the NetBackup Accelerator feature does not work.</p> <p>NetBackup Accelerator is not supported for Red Hat Virtualization.</p>

Error while restoring Red Hat Virtualization virtual machines

The following table describes the problem that might occur when you restore an Red Hat Virtualization virtual machine.

Table 4-3 Error run into while restoring Red Hat Virtualization virtual machines

Error message or cause	Explanation and recommended action
VM recovery to alternate location fails on a Windows primary server.	For a Windows NetBackup primary server, ensure that the rename file ends with an empty line.
Pre-recovery check runs successfully when a VM with the same UUID exists in the Red Hat Virtualization cluster and the option to overwrite the VM is not enabled, but the VM restore fails. The following error message is seen: Info bpVMutil (pid=1196) FTL - Virtual machine exists and overwrite option not specified, can not proceed with restore. end Restore; elapsed time Hypervisor policy restore error. (2822)	Pre-recovery check compares the VM display name instead of UUID to find out if VM already exists, hence the check completes successfully. But if the overwrite option is not set, the restore job fails if a VM with the same UUID already exists. Workaround: Restore the VM with a new UUID <ol style="list-style-type: none"> 1 Start the recovery process. 2 On the Recovery Options page, click Advanced. 3 Enable Create a new VM UUID. 4 Proceed with the recovery process and click Start recovery to restore. Overwrite the existing VM that has the same UUID <ol style="list-style-type: none"> 1 Start the recovery process. 2 On Recovery Options page, enable the Overwrite existing virtual machine option. 3 Proceed with the recovery process and click Start recovery to restore.
When you try to recover an Red Hat Virtualization VM image that is imported from a different domain using the Web UI, the pre-recovery check fails and displays that by default the recovery host is the same access host that was used during back up.	During the recovery of imported Red Hat Virtualization VM images, for the recovery host, select the access host in the target domain as a recovery host or select the target primary server.

API and command line options for RHV

This chapter includes the following topics:

- Using APIs and command line options to manage, protect, or recover RHV VMs
- Additional NetBackup options for Red Hat Virtualization configuration

Using APIs and command line options to manage, protect, or recover RHV VMs

This topic lists the APIs and command line options to protect or recover the Red Hat Virtualization virtual machines. Only the important variables and options are mentioned in this topic.

Following sections are part of this topic:

- Add the RHV manager credentials
- Validate the Red Hat Virtualization manager credentials
- Create an Red Hat Virtualization VM backup policy
- Restore the Red Hat Virtualization VM at the original location
- Restore the Red Hat Virtualization VM to an alternate location

For detailed information on the APIs and command lines, use these references:

- All the NetBackup APIs are listed at the following location:
Services and Operations Readiness Tools (SORT) > Knowledge Base > Documents
- For more information about the commands, refer to the *NetBackup Commands Reference Guide*.

Add the Red Hat Virtualization manager credentials

Table 5-1 Add the Red Hat Virtualization manager credentials

API or command line options	Important variables and options
POST /netbackup/config/servers/vmservers	<ul style="list-style-type: none">serverName is the name of the Red Hat Virtualization managervmType is RED_HAT_VIRTUALIZATION_MANAGER
tpconfig command	<ul style="list-style-type: none">virtual_machine is the name of the Red Hat Virtualization manager.vm_type is 10. The number 10 stands for Red Hat Virtualization Manager.

Validate the Red Hat Virtualization manager credentials

Table 5-2 Validate the Red Hat Virtualization manager credentials

API or command line options	Important variables and options
POST /netbackup/config/servers/vmservers/ {serverName}/validate-credential	<ul style="list-style-type: none">{serverName} is the name of the Red Hat Virtualization manager.validationHost is a whitelisted Windows or Linux backup host.

Create an Red Hat Virtualization VM backup policy

Table 5-3 Create an Red Hat Virtualization VM backup policy

API or command line options	Important variables and options
POST /netbackup/config/policies/	<ul style="list-style-type: none"> ■ <code>policyType</code> is Hypervisor ■ <code>backuphost</code> is a whitelisted Windows or Linux host. ■ <code>snapshotMethodArgs</code> can have the following values to back up a VM using VM UUID: <pre>application_consistent=1 Virtual_machine_backup=1 vm_identifier=GUID (catalog uses VM UUID) file_system_optimization=1 exclude_swap=1</pre> <p><code>vm_identifier=[GUID:VM GUID]</code> is the primary VM identifier Supported values are: DISPLAYNAME and GUID.</p> ■ In <code>backupSelections > selections</code>, use the filter option as <code>"rhv:/?filter=Displayname Contains <name_filter>"</code> to filter Red Hat Virtualization VMs of a specific name. Apart from Displayname, you can use the other filter criteria mentioned for Intelligent VM groups.
admincmd command	<ul style="list-style-type: none"> ■ In <code>bpplclients -add <discoveryhost></code> Hypervisor Hypervisor, the hypervisor discovery host is a whitelisted Windows or Linux host. ■ In <code>bpplinfo</code>, the policy type (<code>-pt</code>) is Hypervisor. ■ In <code>bpplinclude</code>, use the filter option as <code>"rhv:/?filter=Displayname Contains <name_filter>"</code> to filter Red Hat Virtualization VMs of a specific name. ■ In <code>bpplinfo</code> <ul style="list-style-type: none"> ■ Value of <code>use_virtual_machine</code> is 5 for Red Hat Virtualization VMs. ■ Value of <code>snapshot_method</code> is <code>Hypervisor_snap</code>. <p>For optimized backup, you can use:</p> <pre>file_system_optimization=1 exclude_swap=1</pre>

After you create the policy, other commands like creating the schedule for the policy or triggering the policy backup remain the same. For more information about the commands, refer to the *NetBackup Commands Reference Guide*.

Restore the Red Hat Virtualization VM at the original location

Table 5-4 Restore the Red Hat Virtualization VM at the original location

API or command line options	Important variables and options
POST /netbackup/recovery/workloads/rhv/ scenarios/full-vm/recover	<ul style="list-style-type: none"> ■ <code>client</code> is the VM identifier of the protected VM. The VM identifier is the VM UUID. ■ <code>recoveryHost</code> is a whitelisted Windows or Linux host. ■ Set the following values: <code>defaultVmDiskProvisioning</code> <code>powerOnAfterRecovery</code> <code>overwriteExistingVm</code> <code>removeNetworkInterfaces</code> <code>retainVmGuid</code> <code>removeTagAssociations</code>
bprestore command	<ul style="list-style-type: none"> ■ <code>vmproxy</code> is a whitelisted Windows or Linux backup host. ■ <code>vmserver</code> is the name of the Red Hat Virtualization manager. ■ <code>vmhypervisor</code> specifies restore from the Hypervisor policy type ■ Use the following values to modify the VM configuration: <ul style="list-style-type: none"> ■ <code>vmst</code> to remove the VM tags. ■ <code>vmpoweron</code> to start the VM after the VM restore. ■ <code>vmsn</code> to remove the VMs network interfaces. ■ <code>vmid</code> to retain the original VM UUID of the VM. Alternatively, use the <code>-K</code> option to retain the existing VM with the same UUID and not overwrite it. ■ <code>thickdisk</code> to configure the restored virtual disks in the thick format. The virtual disk space is allocated when the disk is created. ■ <code>thindisk</code> to configure the restored virtual disks in the thin format. The populated blocks are restored but the vacant blocks are not initialized or committed.

Restore the Red Hat Virtualization VM to an alternate location

Table 5-5 Restore the Red Hat Virtualization VM to an alternate location

API or command line options	Important variables and options
<div>POST</div> <div>/netbackup/recovery/workloads/rhv/scenarios/full-vm/recover</div>	<div><ul style="list-style-type: none">■ <code>client</code> is the VM name of the protected VM. The VM name can either be the display name (<code>displayName</code>) or the UUID.■ <code>rhvServer</code> is the name of the alternate Red Hat Virtualization manager.■ <code>recoveryHost</code> is a whitelisted Windows or Linux host.■ <code>vmhypervisor</code> specifies restore from the Hypervisor policy type■ Set the following values:<div><div>defaultVmDiskProvisioning</div><div>powerOnAfterRecovery</div><div>overwriteExistingVm</div><div>removeNetworkInterfaces</div><div>retainVmGuid</div><div>removeTagAssociations</div></div></div>

Table 5-5 Restore the Red Hat Virtualization VM to an alternate location
(continued)

API or command line options	Important variables and options
<code>bprestore</code> command	<ul style="list-style-type: none"> ■ <code>vmproxy</code> is a whitelisted Windows or Linux backup host. ■ <code>vmserver</code> is the name of the Red Hat Virtualization manager. ■ Use the following values to modify the VM configuration: <ul style="list-style-type: none"> ■ <code>vmst</code> to remove the VM tags. ■ <code>vmpoweron</code> to start the VM after the VM restore. ■ <code>vmsn</code> to remove the VMs network interfaces. ■ <code>vmid</code> to retain the original VM UUID of the VM. Alternatively, use the <code>-K</code> option to retain the existing VM with the same UUID and not overwrite it. ■ The <code>-R</code> option defines the path of the rename file. Use the rename file to recover the VM to an alternate location or change the VM configuration. Sample rename file: <pre>change vmname to new_vm_name change /storage_domain_1/disk1_UUID to /storage_domain_2/ change /storage_domain_1/disk2_UUID to /storage_domain_2/ change cluster to new_cluster_name</pre> <p>Note: For a Windows NetBackup host, you must add an empty line at the end of the rename file entries.</p> <p>See “Additional information about the rename file” on page 52.</p>

Additional information about the rename file

- You can specify destination storage domain for all the disks or for some specific list of disks.
- If you do not specify a destination storage domain for one of the disks, then that disk is restored to the original location.
- If you specify a destination storage domain for a non-existing or invalid disk, the VM restore fails.
- For a Windows NetBackup host, you must add an empty line (carriage return) after all the rename file entries.

Additional NetBackup options for Red Hat Virtualization configuration

Use the following NetBackup command options for additional Red Hat Virtualization configuration:

- See “OVIRT_IMAGEIO_INACTIVITY_TIMEOUT option for NetBackup servers” on page 53.
- See “RHV_CREATEDISK_TIMEOUT option for NetBackup servers” on page 53.
- See “RHV_AUTODISCOVERY_INTERVAL option for NetBackup servers” on page 54.

OVIRT_IMAGEIO_INACTIVITY_TIMEOUT option for NetBackup servers

This option specifies the timeout period of client inactivity in seconds. When this timeout period is surpassed the oVIRT engine aborts the transfer session. The client inactivity usually occurs when there is a traversal of a disk chain. For example, while backing up a thin-dependent VM or a VM that consists of user snapshots.

Table 5-6 OVIRT_IMAGEIO_INACTIVITY_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the <code>OVIRT_IMAGEIO_INACTIVITY_TIMEOUT</code> option. The default inactive timeout period is 172800 seconds (48 hours).
Example	The following entry tells the NetBackup backup job to set the client inactivity timeout period to 172800 seconds (48 hours). <code>OVIRT_IMAGEIO_INACTIVITY_TIMEOUT = 172800</code>

RHV_CREATEDISK_TIMEOUT option for NetBackup servers

This option specifies the timeout period for creating a virtual disk during the restore of an Red Hat Virtualization VM. If the Red Hat Virtualization VM with a large pre-allocated disk was backed up and then restored on a file storage such as NFS, then the disk creation function can time out before the restored virtual disk is fully realized.

Table 5-7 RHV_CREATEDISK_TIMEOUT information

Usage	Description
Where to use	On NetBackup primary servers.
How to use	Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the <code>RHV_CREATEDISK_TIMEOUT</code> option.
Example	<p>The following entry tells the NetBackup backup job to set the create disk timeout period to 172800 seconds (48 hours).</p> <pre>RHV_CREATEDISK_TIMEOUT = 172800</pre> <p>The range for <code>RHV_CREATEDISK_TIMEOUT</code> is 0 hours to 48 hours.</p>

RHV_AUTODISCOVERY_INTERVAL option for NetBackup servers

This option controls how often NetBackup scans the Red Hat Virtualization servers to discover virtual machines to display in the NetBackup web UI.

NetBackup attempts autodiscovery first with the same host for which the last discovery attempt was successful. If autodiscovery fails with that host, NetBackup tries again with other hosts in the following order:

- The NetBackup primary server
- The access host, client, or proxy server
- The media server

Table 5-8 Red Hat Virtualization_AUTODISCOVERY_INTERVAL information

Usage	Description
Where to use	On NetBackup primary servers.

Table 5-8

Red Hat Virtualization_AUTODISCOVERY_INTERVAL
information *(continued)*

Usage	Description
How to use	<p>Use the <code>nbgetconfig</code> and the <code>nbsetconfig</code> commands to view, add, or change the option.</p> <p>The default is 28,800 seconds (8 hours). The minimum value is 300 seconds (5 minutes) and the maximum is 31,536,000 seconds (1 year).</p> <p>Use the following format:</p> <pre>RHV_AUTODISCOVERY_INTERVAL = number of seconds</pre> <p>For example:</p> <pre>RHV_AUTODISCOVERY_INTERVAL = 100000</pre> <p>This entry should appear only once in the configuration file.</p> <p>Note: After changing this option, stop and restart the NetBackup services. For VM discovery, the <code>Netbackup Discovery Framework</code> service must be running.</p>

