

NetBackup™ Web UI Nutanix AHV Administrator's Guide

Release 9.1



Last updated: 2021-06-08

Legal Notice

Copyright © 2021 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

NB.docs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the NetBackup web user interface	7
		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
	Monitoring jobs	13
	Filter jobs in the job list	14
Chapter 3	Configure and protect AHV assets from WebUI	15
		15
	Configure and protect the AHV assets from the NetBackup web UI	15
		15
Chapter 4	Managing AHV clusters	17
	Quick configuration checklist to protect AHV virtual machines	18
	Configure secure communication between the AHV cluster and	
	NetBackup host	22
	Enable the iSCSI initiator service on windows backup host	25
	Install the iSCSI initiator package on Linux backup host	25
	Migrate Java GUI/CLI added clusters into WebUI	25
	Configure Nutanix AHV cluster	26
	Configure CHAP settings for iSCSI secure communication with AHV	
	clusters	27
	About the ports that NetBackup uses to communicate with AHV	27
	Add or browse an AHV cluster	28
	Add new cluster credentials	31
	Update and validate AHV cluster credentials	31
	Remove AHV Clusters	32
	Create an intelligent VM group	33

	Assign permissions to the intelligent VM group	36
	Update the intelligent VM group	37
	Remove the intelligent VM group	37
	Set CHAP for iSCSI	38
	Add an AHV access host	38
	Remove an AHV access host	39
	Change resource limits for AHV resource types	39
	Change the autodiscovery frequency of AHV assets	43
Chapter 5	Protecting AHV virtual machines	44
	Things to know before you protect AHV virtual machines	44
	Protect AHV VMs or intelligent VM groups	45
	Edit protection settings for an AHV asset	45
	Schedules and Retention	46
	Backup options	46
	Prerequisite to Enable virtual machine quiescing	47
	Remove protection from VMs or intelligent VM groups	48
	View the protection status of VMs or intelligent VM groups	48
Chapter 6	Recovering AHV virtual machines	50
	Things to consider before you recover the AHV virtual machines	51
	About the pre-recovery check	51
	Recover an AHV virtual machine	51
	About Nutanix AHV agentless files and folders restore	53
	Prerequisites for agentless files and folder recovery	54
	SSH key fingerprint	65
	Recover files and folders with Nutanix AHV agentless restore	66
	Recovery target options	67
	Pre-recovery checks	71
	About Nutanix-AHV agent-based files and folders restore	73
	Prerequisites for agent based files and folder recovery	73
	Recover files and folders with Nutanix AHV agent based restore	74
	Limitations	76
Chapter 7	Troubleshooting AHV operations	78
	Troubleshooting tips for NetBackup for AHV	78
	Error during AHV credential addition	79
	Error during the AHV virtual machines discovery phase	79
	Errors for the Status for a newly discovered VM	80
	Error run into while backing up AHV virtual machines	81
	Error while restoring AHV virtual machines	84

Chapter 8	API and command line options for AHV	94
	Using APIs and command line options to manage, protect, or recover	
	AHV virtual machines	94
	Additional NetBackup options for AHV configuration	101
	Additional information about the rename file	102

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.

Alternatively, policy management is also available for a limited number of policy types.

- Workload administrators can create protection plans, 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, with the exception of 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.
- In addition to schedules for backups, a protection plan can also include a schedule for replication and long-term retention.
- 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.

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 <i>intelligent group</i> .
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.
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 or the CLIs.
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.

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

Term	Definition
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 <i>Subscribe</i> as <i>Add protection</i> .
Unsubscribe, from a protection plan	<i>Unsubscribe</i> 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, or Cloud.
CHAP for iSCSI	The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators that are NetBackup backup or recovery hosts and iSCSI target that is Nutanix AHV cluster.

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.

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
- Enter your credentials and click **Sign in**.

For example:

For this type of user	Use this format	Example
Local user	<i>username</i>	jane_doe
Windows user	<i>DOMAINusername</i>	WINDOVS\jane_doe
UNIX user	<i>username@domain</i>	john_doe@unix

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**.

Monitoring NetBackup

This chapter includes the following topics:

- [The NetBackup dashboard](#)
- [Monitoring jobs](#)
- [Filter jobs in the job 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.

Monitoring jobs

Use the **Jobs** node to monitor the jobs in your NetBackup environment and view the details for a specific job.

To monitor a job

- 1 On the left, click **Activity monitor > Jobs**.
- 2 Click on a job name that you want to view.
On the **Overview** tab you can view information about a job.
 - The **File List** contains the files that are included in the backup image.

- The **Status** section shows the status and the status codes that are related to the job. Click the status code number to view information about this status code in the Veritas Knowledge Base.
See the [NetBackup Status Codes Reference Guide](#).
- 3 Click the **Details** tab to view the logged details about a job. You can filter the logs by error type using the drop-down menu.
See [“Filter jobs in the job list”](#) on page 14.

Filter jobs in the job list

You can filter the jobs to display the jobs in a specific state. For example, you can display all of the active jobs or all of the suspended jobs.

To filter the job list

- 1 Click **Jobs**.
- 2 Above the job list, click the **Filter** option.
- 3 In the **Filter** window, select a filter option to dynamically change the jobs that are displayed. The filter options are as follows:
 - **All**
 - **Active**
 - **Done**
 - **Failed**
 - **Incomplete**
 - **Partially Successful**
 - **Queued**
 - **Successful**
 - **Suspended**
 - **Waiting for Retry**
- 4 Click **Apply Filters**.
- 5 To remove the selected filters, click **Clear All**.

Configure and protect AHV assets from WebUI

This chapter includes the following topics:

- [Configure and protect the AHV assets from the NetBackup web UI](#)

Configure and protect the AHV assets from the NetBackup web UI

- Login to NetBackup web UI as the Administrator and to configure a role, select **Default AHV Administrator** and assign required permissions. See *Default AHV Administrator* in *NetBackup Web UI Administrator's Guide*.

Note: To perform the AHV administrator tasks, the **Default AHV Administrator** role has minimum required permissions.

- Pre-requisite configurations for AHV cluster:
 - Configure secure communication between the AHV cluster and NetBackup host.
See [“Configure secure communication between the AHV cluster and NetBackup host”](#) on page 22.
 - Enable iSCSI on the NetBackup host which you want to use as backup/restore host.
See [“Install the iSCSI initiator package on Linux backup host”](#) on page 25.
See [“Enable the iSCSI initiator service on windows backup host”](#) on page 25.
 - (Optional) Whitelist the backup host in Nutanix Prism console .

Note: To use the NFS protocol on the Linux backup or recovery host, NFS allowed listing of the host on the Nutanix AHV Prism Console is required. For details, click [here](#).

- Sign in with the Administrator role or the Default AHV Administrator role and do the following:
 - Configure and manage AHV cluster.
See [“Configure Nutanix AHV cluster”](#) on page 26.
 - Add and manage credentials.
See [“Add new cluster credentials”](#) on page 31.
 - Configure an AHV protection plan
Create a protection plan link in NetBackup™ Web UI Administrator's Guide
 - Configure an Intelligent VM group.
See [“Create an intelligent VM group”](#) on page 33.
 - Protect AHV VMs or intelligent VM groups.
See [“Protect AHV VMs or intelligent VM groups”](#) on page 45.
 - Recover V.
See [“Recover an AHV virtual machine”](#) on page 51.

Managing AHV clusters

This chapter includes the following topics:

- [Quick configuration checklist to protect AHV virtual machines](#)
- [Configure secure communication between the AHV cluster and NetBackup host](#)
- [Enable the iSCSI initiator service on windows backup host](#)
- [Install the iSCSI initiator package on Linux backup host](#)
- [Migrate Java GUI/CLI added clusters into WebUI](#)
- [Configure Nutanix AHV cluster](#)
- [Configure CHAP settings for iSCSI secure communication with AHV clusters](#)
- [About the ports that NetBackup uses to communicate with AHV](#)
- [Add or browse an AHV cluster](#)
- [Add new cluster credentials](#)
- [Update and validate AHV cluster credentials](#)
- [Remove AHV Clusters](#)
- [Create an intelligent VM group](#)
- [Assign permissions to the intelligent VM group](#)
- [Update the intelligent VM group](#)
- [Remove the intelligent VM group](#)
- [Set CHAP for iSCSI](#)
- [Add an AHV access host](#)

- [Remove an AHV access host](#)
- [Change resource limits for AHV resource types](#)
- [Change the autodiscovery frequency of AHV assets](#)

Quick configuration checklist to protect AHV virtual machines

Use NetBackup web UI to protect and recover the virtual machines that are created on the AHV platform. You can also use APIs and command line options for the same.

See [“Using APIs and command line options to manage, protect, or recover AHV virtual machines”](#) on page 94.

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

Table 4-1 Configure and protect AHV virtual machines using NetBackup

Step overview	Description and reference
Deploy NetBackup to protect AHV VMs	On a very high level to protect AHV VMs you need: ■ NetBackup primary server ■ NetBackup media server (Recommended) ■ NetBackup client that can act as a backup host The operating system of the backup host must be a Linux RHEL, SUSE, or Windows. The backup host can be a NetBackup media server or a client, or an NetBackup Appliance. NetBackup appliance including Flex appliance and Flex scale appliance is also supported as a NetBackup media server that can act as a backup host. NetBackup uses an agentless architecture to protect the AHV VMs. The communication between NetBackup and AHV cluster happens through Nutanix AHV APIs.

Table 4-1

Configure and protect AHV virtual machines using NetBackup
(continued)

Step overview	Description and reference
Configure an AHV access host for backup and recovery	An AHV 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. 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 AHV Access Hosts list. Note: A backup host which is not a media server or an appliance needs to have the NetBackup client installed on it. See “Add an AHV access host” on page 38.
Enable secure communication between NetBackup and AHV	The following sections contain more information about setting up a secure communication between NetBackup and AHV: - Secure communication See “Configure secure communication between the AHV cluster and NetBackup host” on page 22. - Communication ports See “About the ports that NetBackup uses to communicate with AHV” on page 27.
Manage AHV clusters and intelligent VM groups	- Managing AHV clusters See “Add or browse an AHV cluster” on page 28. - Managing intelligent VM groups See “Create an intelligent VM group” on page 33. See “Remove the intelligent VM group” on page 37.
Protect the AHV VMs	- Prerequisite: Adding an AHV cluster requires the Default AHV Administrator role. - Best practices See “Things to know before you protect AHV virtual machines” on page 44. - Protecting virtual machines See “Protect AHV VMs or intelligent VM groups” on page 45.

Table 4-1

Configure and protect AHV virtual machines using NetBackup
(continued)

Step overview	Description and reference
iSCSI Transport for Windows backup hosts	Prerequisite For Windows 2012 or later, iSCSI client initiator is present on Windows. By default, the iSCSI initiator service is stopped or disabled on Windows. See "Enable the iSCSI initiator service on windows backup host" on page 25. Note: If a backup or a recovery host that is selected is on Windows, ensure that iSCSI service is running on windows computer to avoid failure of backup or restore jobs.
iSCSI Transport for Linux backup hosts	Prerequisite To use iSCSI, <code>scsi-initiator-utils</code> package must be installed. By default, it is installed on the RHEL/SUSE. See "Install the iSCSI initiator package on Linux backup host" on page 25. Note: To use the NFS protocol on the Linux backup or recovery host, NFS allowed listing of the host on the Nutanix AHV Prism Console is required. For details, refer https://www.veritas.com/content/support/en_US/doc/127664414-132725336-0/v127698742-132725336. If the <code>iscsi-initiator-utils</code> package is already installed on the backup or the recovery host, ensure that the iSCSI daemon is running. ■ To check the status of the daemon use the command <code>systemctl status iscsid</code>. ■ If the daemon is disabled then run the command <code>systemctl enable iscsid</code> and then run the command to start the iSCSI daemon <code>systemctl start iscsid</code>.

Table 4-1

Configure and protect AHV virtual machines using NetBackup

(continued)

Step overview	Description and reference
Configure CHAP settings for iSCSI secure communication with Nutanix AHV clusters	One-way CHAP: - iSCSI initiator authenticates with the target (AHV) using the random generated CHAP password/secret. Mutual CHAP - automatic: - NetBackup Credential Management Service (CMS) automatically generates a credential with the prefix AHV_ISCSI_MUTUAL_AUTO_ for the backup/recovery host CHAP password. This credential is used for mutual authentication between the iSCSI initiator that is NetBackup backup/recovery host and the target that is AHV. You can set a retention period for these auto-generated CHAP passwords. The default retention period for the auto-generated CHAP passwords is 90 days from the date of creation. Note: The default configuration is one-way CHAP. To enable the Mutual CHAP option: See “Configure CHAP settings for iSCSI secure communication with AHV clusters” on page 27.
Set global limits on the use of AHV resources	VMs are automatically protected, 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 AHV performance as well as backup performance. You can set the global limits to manage the AHV resources efficiently. See “Change resource limits for AHV resource types” on page 39.

Table 4-1

Configure and protect AHV virtual machines using NetBackup
(continued)

Step overview	Description and reference
NetBackup Automatic Backup Host selection	NetBackup Automatic backup host selection option internally uses NetBackup media server load balancing to allocate snapshot/backup jobs to an available, supported media servers. NetBackup avoids sending jobs to busy media servers. Note: Application consistent backups require NetBackup 9.1 or later on the media server. Prerequisite - Click Storage > Storage Server. Add all the supported media servers for load balancing. - Click Storage > Storage Unit > <i>storage unit name</i> > Under Media Server section and select Allow NetBackup to automatically select. - When you create an AHV Protection plan, select Automatic for the Select server or host to use for backups setting.

Configure secure communication between the AHV cluster and NetBackup host

NetBackup can now validate AHV cluster certificates using their root or intermediate certificate authority (CA) certificates.

Only PEM certificate format is supported for virtualization servers.

The following procedure is applicable for the NetBackup media servers acting as backup hosts and all AHV access hosts.

To configure secure communication between AHV cluster and AHV access host:

- 1 Use the `openssl s_client -connect <Nutanix Cluster FQDN>:9440 -showcerts < /dev/null` command from a Linux system to obtain the Nutanix certificates.
- 2 Scroll to the end of the results and copy the last certificate which starts from:

```
-----BEGIN CERTIFICATE-----
<Certificate>
-----END CERTIFICATE-----
```

Note: Ensure to copy the five dashes before and after the BEGIN and END CERTIFICATE.

- 3 Paste this in a text file and then rename it as `<certificate file name>.pem` and copy it to a path to your backup host. Recommended path is:
 - For Linux: `/usr/opensv/netbackup.`
 - For Windows: `<Install drive>\Program Files\Veritas\Netbackup.`
- 4 ■ For Linux: Enter the PEM file path
`ECA_TRUST_STORE_PATH=/usr/opensv/netbackup/<certificate file name>.pem` in the `bp.conf` on the backup host.

- For Windows: Run the command `<Install drive>\Program Files\Veritas\Netbackup\bin\nbsetconfig`.

5 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 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. 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 AHV 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.
VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED	This option affects AHV, RHV, 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. It is recommended by NetBackup that secure communication should be enabled using the <code>ECA_TRUST_STORE_PATH</code> option. Disabling this option lets you skip the security certificate validation. Lets you validate the revocation status of the virtualization server certificate against the CRLs. By default, the option is disabled.

For more information on external CA support, refer to the [NetBackup Security and Encryption Guide](#).

Enable the iSCSI initiator service on windows backup host

Do one of the following:

- 1
 - Click **Server Manager> Tools> iSCSI initiator**.
 - A message window displays, **To start the service now and have the service start automatically each time the computer restarts, click the Yes button**. Click **Yes** to confirm.
- 2 Alternatively, to enable the iSCSI service from the administrative tools:
 - Open **Control Panel -> Administrative Tools -> Open Services**.
 - Find **Microsoft iSCSI Initiator Service**.
 - Right-click on it and click on **Start**.

Note: The default option for this service is **Manual**. Change the setting to **Automatic** to auto start the service running when you reboot.

Install the iSCSI initiator package on Linux backup host

To install the iSCSI initiator package use the following yum and zypper commands:

- `yum install iscsi-initiator-utils` - RedHat.
- `zypper -n install open-iscsi` - SuSE.

Migrate Java GUI/CLI added clusters into WebUI

Credential management for JAVA GUI/CLI and WebUI are separate.

- Clusters added via Java GUI/CLI won't be reflected in WebUI and vice-versa.
- If there are any existing clusters in Java GUI/ CLI, user must manually add these cluster/clusters and its credentials in WebUI.

Note: Once a cluster is added in WebUI and then if it is deleted from Java GUI/CLI, the cluster would still exist in WebUI and vice-versa.

- Once the cluster is added in WebUI and if the cluster credential needs to be updated, it must be updated from WebUI only.

Consider the following scenario:

- A cluster exists in both WebUI and Java UI.
- Cluster credential is updated in WebUI only.
- The cluster is deleted from WebUI.

Impact: Backups and restores may fail on Java GUI as the credential of cluster added in Java GUI was not updated.

Recommendation: Update the credential from Java GUI.

- After the cluster is added in WebUI, the backups with existing policies would still be successful, even if the cluster is deleted from Java GUI. However, restore jobs can not be triggered from Java GUI in this scenario; as it requires cluster to be present on Java GUI.
- If a cluster is added from Java GUI and WebUI and then it is deleted it from Java GUI, the cluster can still be seen at WebUI and vice-versa.
- If a cluster was in WebUI as well as Java GUI and its credential was updated in WebUI and then the cluster is deleted from WebUI, backup and restores may fail as the cluster added in Java UI was not updated. One may have to update the credential from Java UI to get things working.

Configure Nutanix AHV cluster

Prerequisite:

Configure iSCSI Data Services IP on the Nutanix AHV cluster

- ◆ As per **Nutanix** recommendations, data services IP for iSCSI must be configured on Nutanix AHV.

Go to the Nutanix AHV cluster Prism console at <https://<Nutanix cluster FQDN/IP>:9440>.

Click **Settings > Cluster Details > Set iSCSI Data Services IP**.

Note: If this setting is not configured, for Windows backup host, backup/restore jobs fail and for Linux backup host, jobs fall back to use NFS.

Note: The failure of the backup/restore job for windows backup hosts is displayed as failure in **Activity Monitor > Job Details**. The fallback from iSCSI to NFS in case of Linux backup hosts is mentioned as a warning in the job details.

Configure CHAP settings for iSCSI secure communication with AHV clusters

The CHAP settings apply to all AHV clusters configured to currently selected primary server.

- 1 On the left, select **Workloads > Nutanix AHV**.
- 2 On the top, click **AHV settings**.
- 3 Select **CHAP for iSCSI**.
- 4 Select the appropriate CHAP option.

About the ports that NetBackup uses to communicate with AHV

The following table describes the ports that NetBackup requires to communicate with AHV:

Table 4-2 Ports required by NetBackup to communicate with AHV

Port	Protocol	Destination	Purpose
80, 443	TCP	AHV Cluster	Provides HTTP and HTTPS access to the AHV Cluster
54322	TCP	AHV Hosts	Required for communication with the <code>ImageIO</code> daemon (<code>ovirtimageio-daemon</code>)
54323	TCP	AHV Cluster (<code>ImageIO</code> Proxy server)	Required for communication with the <code>ImageIO</code> Proxy (<code>ovirtimageio-proxy</code>)
860, 3260	iSCSI uses TCP	AHV Cluster	iSCSI provides block-level access to storage devices with the <code>SCSI</code> . iSCSI facilitates data transfers usually over the ethernet.
111	TCP		Portmapper
2049	TCP		NFS
9440			Prism console

Add or browse an AHV cluster

You can add and browse AHV cluster and their credentials.

To add AHV cluster and their credentials

- 1 On the left, click **Nutanix AHV** then click the **AHV cluster** tab.
- 2 Click **Add** to add an AHV cluster and enter the following:
See [“Error during AHV credential addition”](#) on page 79.

■ Cluster name

Note: NetBackup recommends that you use the FQDN to add the AHV cluster. The cluster name must follow the limit of 218 characters.

■ REST API port (default: 9440)

This port must remain open between backup host and AHV cluster.

See [“About the ports that NetBackup uses to communicate with AHV”](#) on page 27.

- **Select a backup host**

This backup host is used for validation and discovery.

Note: Credential validation and discovery of virtual machines is only supported by NetBackup 9.1 or later.

- **Associate credential**

Do one of the following:

- Select an existing credential see *Managing Credentials* in the [NetBackup™ Web UI Administrator's Guide](#).
- See [“Add new cluster credentials”](#) on page 31.

3 Click Add and Manage permissions.

Validations for all the inputs are performed.

Select the roles you want to have access to this cluster. See *Managing role-based access control* in the [NetBackup™ Web UI Administrator's Guide](#).

4 To add another AHV cluster credentials, click Add.

Inline actions on AHV cluster

You can run the following inline actions on an AHV cluster:

- **Discover:** Manually discovers the VM assets that belong to the selected AHV cluster.
- **Edit:** Modify the AHV cluster credentials.
- **Delete:** Removes the AHV cluster.
- **Manage Permissions:** Used to add or manage the permissions on selected cluster.

Bulk actions on AHV cluster

You can select one or more AHV cluster and run the following bulk actions:

- **Discover:** Manually discovers the VM assets that belong to the selected AHV cluster.

Note: Discovery is triggered in a sequential manner for clusters one after the other.

- **Validate credentials:** Validates the credentials of the AHV cluster.
- **Delete:** Removes the AHV cluster.

Browse an AHV cluster

You can browse the AHV clusters to locate VMs and storage containers and their details.

To browse AHV cluster

- 1 On the left, click **Nutanix AHV**.
- 2 Click the **AHV cluster** tab, and begin searching.

The list includes the AHV clusters that you have access to.

The tab shows the AHV clusters that you can access in the following hierarchy:

```
All
AHV_clusters
  cluster1
    VirtualMachine
    StorageContainer
  cluster2
    VirtualMachine
    StorageContainer
```

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

- 3 Click on an AHV cluster to view details.
- 4 Click on a virtual machine to view its protection status, recovery points, and restore activity.
- 5 Click **Add protection** to subscribe the selected VM to a protection plan. You can also select **Backup now**, **Recover**, and **Manage Permission** options.

Note: Ensure that the protection plan for AHV exists to add protection to an asset.

- 6 Click on a storage container to view free space, and last discovered time.

Note: Once the data exceeds the advertising capacity, the additional data is shown as negative and for such value. NetBackup WebUI displays an empty field and corresponding API shows -ve value for the free space field for particular storage container.

- 7 For storage container, you can **Manage Permission**.

Note: **Manage Permission** is enabled only when you select the storage container.

Add new cluster credentials

- 1 On left, click **Nutanix AHV**, then click the **AHV cluster** tab.
- 2 Click **+Add** to add a new cluster.
- 3 On the **Add AHV cluster > Associate credentials** page, click **Add a new credential**.
- 4 On the **Add credential** page, enter the details such as **Credential name**, **Username**, and **Password**.
- 5 Click **Next**.
Select or add roles to provide permissions for credential.
- 6 Click **Save**.

Note: You can **Edit** or **Remove** the added credentials.

Update and validate AHV cluster credentials

To validate AHV credentials

- 1 On the left, click **Nutanix AHV**, then click the **AHV clusters** tab.
- 2
 - To validate a specific cluster's credentials, locate and select the AHV cluster. Then either click **Validate** from the **Credentials** column or from top bar.
 - To validate the credentials of multiple servers at the same time, locate and select the AHV cluster. Then click **Validate** from the top bar.

Note: NetBackup verifies the current credentials for the selected AHV cluster. If the credentials are not valid, NetBackup indicates **Invalid** under **Credentials**. Use the following steps to update the AHV cluster credentials.

To update AHV cluster credentials

- 1 On the left, click **Nutanix AHV**, then click the **AHV cluster** tab.
- 2 Locate and select the AHV cluster.
- 3 Select **Actions > Edit**.
- 4 Update the credentials as needed.

Note: Adding or updating AHV cluster credentials also automatically starts the discovery of the AHV cluster. 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 9.1 is the minimum version that is supported for a NetBackup media server or client that serves as a backup host.

- 5 Click **Save**.

NetBackup verifies the updated credentials for the selected AHV cluster.

Remove AHV Clusters

Use this procedure to delete AHV clusters.

To remove an AHV cluster

- 1 On the left, click **Nutanix AHV**, then click the **AHV clusters** tab.

The tab lists the names of AHV clusters 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 AHV cluster.
- 3 Select **Actions > Delete**.

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

- 4 If you are sure that you want to delete the AHV cluster, 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: A background task adds the newly discovered VMS that matches the query to the intelligent VM group. This background task runs 30 minutes after the start of the NetBackup Web Management service. After that, the task runs every 30 minutes.

To create an intelligent VM group

- 1 On the left, click **Nutanix AHV**.
- 2 Click the **Intelligent VM groups** tab and then click **Add intelligent VM group**.
- 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 **Clusters** pane, click **Add clusters**.

Note: To create a group, you must have at least one cluster.

- In the **Add clusters** window, select clusters that you want to add.

Note: To add a cluster, you must have view and create permissions on the cluster.

- 5 Perform one of the following:
 - Select the default query: **Include all VMs**.
When the protection plan runs, all VMs that are part of the AHV clusters are added in the intelligent VM group.
 - 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](#).

The following is an example query:

A screenshot of a query builder interface. At the top right, there are buttons for '+ Condition' and '+ Sub-query'. Below these, a single condition is displayed in a light blue box: 'VM display name' (with a dropdown arrow), 'Contains' (with a dropdown arrow), and 'windows' (with a dropdown arrow).

In this example, the query adds to the group any VM that has `windows` 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:

A screenshot of a query builder interface. At the top, there are buttons for 'AND' and 'OR', with 'AND' selected. To the right are buttons for '+ Condition' and '+ Sub-query'. Below, two conditions are listed in separate rows. The first row is 'VM display name' (dropdown), 'Contains' (dropdown), 'windows' (dropdown). The second row is 'Power state' (dropdown), '=', 'ON' (dropdown). Each row has a trash icon on the right.

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

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

A screenshot of a query builder interface. At the top, there are buttons for 'AND' and 'OR', with 'OR' selected. To the right are buttons for '+ Condition' and '+ Sub-query'. Below, two conditions are listed in separate rows. The first row is 'VM display name' (dropdown), 'Contains' (dropdown), 'windows' (dropdown). The second row is 'Power state' (dropdown), '=', 'ON' (dropdown). Each row has a trash icon on the right.

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

- The VMs that have `windows` in their display name (regardless of power state).
- The VMs that have a power state `ON` (regardless of the display name).

7 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.

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.

For the case behavior of each option, see the following topic:

[Query options for creating intelligent VM groups](#)

8 To save the group, click **Add and Manage permissions**.

Note: You can edit, protect, and manage permissions for this group.

- Add a protection plan:
See [“Protect AHV VMs or intelligent VM groups”](#) on page 45.
- Edit or update the intelligent VM group:
See [“Update the intelligent VM group”](#) on page 37.
- Assign permissions to the VM group:
See [“Assign permissions to the intelligent VM group”](#) on page 36.

Query options for creating intelligent VM groups

Table 4-3 Query keywords

Keyword	Description
<code>displayName</code>	The VM's display name. Case-sensitive when the protection plan runs.
<code>powerState</code>	The VM's power state. ON and OFF are case-sensitive.

Table 4-3 Query keywords (*continued*)

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

Table 4-4 Query operators

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

Assign permissions to the intelligent VM group

Things to consider before you assign the permissions to the VM group.

■ View/Update

- All the clusters in the group, you must have VIEW permission.
- Without VIEW permission on any of the cluster, you cant preview VMs of the group in **Virtual machines** tab.
- The cluster for which you dont have permission is displayed with the lock sign.

- The deleted cluster is displayed with **X** sign.
- To add any new cluster in the existing VM group, you must have VIEW permission for the intended cluster.
- To update the VM group, you must have VIEW permission on cluster. However, you can delete a non-existing cluster or a cluster without VIEW permission.
- **Protect**
 - All the clusters in the group must have PROTECT permission.
 - To protect a VM group, you must have PROTECT permission on the all the clusters of the group and also on the VM group.
 - Without PROTECT permission on all the cluster, **Backup Now** is disabled.
 - **Remove protection** is enabled irrespective of permission on the clusters. It is driven only by permission on VM group.

For details on the roles permissions, see [NetBackup Web UI Administrator's Guide](#).

Update the intelligent VM group

You can edit the intelligent VM group.

To edit the intelligent VM group

- 1 On the left, click **Nutanix AHV**.
- 2 Click the **Intelligent VM groups** tab and select the VM group you want to edit.
- 3 In the **Virtual machine** tab, click **Edit**.
In the **Clusters** pane, click **Add clusters**.

Note: You can remove or add the VM groups. To add an Intelligent VM group, See [“Create an intelligent VM group”](#) on page 33.

Remove the intelligent VM group

Use the following procedure to remove an intelligent VM group.

To delete an intelligent VM group

- 1 On the left, click **Nutanix AHV**.
- 2 Locate the group under the **Intelligent VM groups** tab.

- 3 If the group is not protected, click its box and 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**.

Set CHAP for iSCSI

The CHAP settings apply to all AHV clusters that are configured under selected primary server. By default, the configuration is set to one-way CHAP.

Note: For one-way CHAP option, no action is required.

To enable the Mutual CHAP option:

- 1 On the left pane, click **Nutanix AHV**.
- 2 On the top-right, select **AHV settings > CHAP for iSCSI** and select the appropriate Mutual CHAP option.

Note: For Mutual CHAP, NetBackup credential management system auto generates the credentials with prefix `AHV_ISCSI_MUTUAL_AUTO_` for the selected backup or recovery host. The iSCSI Mutual CHAP credential would be seen in the **Credential Management** tab.

Note: By default, auto-generated credentials for the Mutual CHAP option are not visible to the users created by default AHV Administrator role. Security admin / root user must provide credential view permission to the particular user to view those.

This auto-generated credential in the **Credential Management** tab and cannot be edited, it can only be deleted. If it is deleted manually it gets recreated automatically when the next job for which this credential is generated runs.

Add an AHV access host

NetBackup uses a special host that is called a AHV 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 storage container of the virtual machines. The access

host reads the data from the storage container and sends it over the network to the media server.

The AHV access host was formerly called the AHV backup host. 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 AHV access host

- 1 On the left pane, click **Nutanix AHV**.
- 2 On the top- right corner, select **AHV settings > Access hosts**.
NetBackup lists any access hosts that were previously added.
- 3 Click **+ Add**.
- 4 Enter the `Name/FQDN/IP` of the access host and then click **Add**.

Remove an AHV access host

To remove an AHV access host

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

Change resource limits for AHV resource types

Nutanix AHV resource limits control the number of simultaneous backups that can be performed on Nutanix AHV resources. The settings apply to all NetBackup policies for the currently selected primary server.

Resource limits available for Nutanix AHV:

- **Backup Jobs per Host**
- **Backup Jobs per AHV Cluster**
- **Backup Jobs per Storage Container**
- **Snapshot Jobs per AHV Cluster**

Note: For each resource, the default value is 0 (No limit).

To set resource limits for Nutanix AHV resources

- 1** On the left, click **Nutanix AHV**.
- 2** On the top right, click **AHV settings > Resource limits**.

For each resource, the default value is **0** (No limit).

Note: The Snapshot Jobs per AHV Cluster option sets a limit for the number of simultaneous snapshot operations per cluster. It only applies during the snapshot creation phase of a backup. It does not control the number of simultaneous backup jobs. This setting can control the effect that multiple snapshot operations have on the AHV cluster. To override the global snapshot setting for that AHV cluster, add a specific AHV cluster.

- 3** Locate the AHV resource that you want to change and then, click **Edit**.

4 Choose from the following options.

Set a global limit for an AHV 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 AHV 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.

The following example shows a global limit of **2** for all AHV clusters and a limit of **1** for the selected AHV cluster.



Edit cluster limit

Set a limit on number of simultaneous backup jobs that are performed on all AHV clusters. Add a specific AHV cluster to override the global setting for that cluster.

+ Add

Cluster	Limits
Global	2
punntx8155-cluster.vxindia.veritas.com	1

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.

Note: After you set resource limits the limits do not take effect until a few jobs run.

Reset the resource limits for all AHV resources

To reset the resource limits for all AHV resources

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

Example - Setting resource limits for a Nutanix cluster with two nodes

Consider the following example:

- The Nutanix cluster has two nodes.
- Each node hosts 40 VMs, so there are 80 VMs in the cluster.
- The **Nutanix-AHV** policy has 20 VMs.

When NetBackup connects to the Nutanix environment for backup, it makes one connection per VM. If no resource limit is set, there are a total of 160 concurrent jobs that run (80 snapshot + 80 backup). See [this article](#).

Nutanix recommends up to 20 connections concurrently per CVM in the cluster, which means that 20 VMs per node are backed up concurrently. In our example, you can enforce a limit of 20 connections with the following settings:

Backup Jobs per Node	20
Backup Jobs per Cluster	40
Backup Jobs per Storage Container	Set any limits based on the characteristics of the storage technology.
Snapshot Jobs per Cluster	10

When a backup begins, the Activity monitor displays the jobs as follows:

- Snapshot jobs: 20
- Active jobs: 10 (snapshot jobs and their backup jobs)
- Queued jobs: 10
- After the active snapshot jobs complete, the queued snapshot jobs become active.

Change the autodiscovery frequency of AHV assets

Automatic discovery of AHV 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 AHV assets

- 1 On the left, click **Workloads > Nutanix AHV**.
- 2 On the right, select **AHV settings > Autodiscovery**.
- 3 Select **Frequency > Edit**.
- 4 Use the up or down arrows to choose how often you want NetBackup to perform autodiscovery of AHV 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 AHV autodiscovery API.

Protecting AHV virtual machines

This chapter includes the following topics:

- [Things to know before you protect AHV virtual machines](#)
- [Protect AHV VMs or intelligent VM groups](#)
- [Edit protection settings for an AHV asset](#)
- [Schedules and Retention](#)
- [Backup options](#)
- [Prerequisite to Enable virtual machine quiescing](#)
- [Remove protection from VMs or intelligent VM groups](#)
- [View the protection status of VMs or intelligent VM groups](#)

Things to know before you protect AHV virtual machines

During protection plan creation there are some validations which are to be taken in consideration:

- If schedule type is Automatic, ensure all the NetBackup versions are as mentioned below.
- Incremental schedules are supported only for backup host version 8.3 or later.
- If you have windows machine as a backup host, make sure the version is 9.1 or later.

- If you want to use Enable virtual machine quiesce option, ensure the backup host is 9.1 or later.

Protect AHV VMs or intelligent VM groups

Use the following procedure to subscribe assets that are AHV 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. In case of protect for Intelligent VM group, ensure that all the clusters forming the group have protect permission.

To protect AHV VMs or VM groups:

- 1 On the left pane, click **Nutanix AHV**.
- 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 If you have the necessary role permissions you can adjust one or more of the following settings:
 - **Schedules and retention**
Change when backups occur and the backup start window.
 - **Backup options**
Adjust the server or host to use for backups.
 - **Advance options**
Enable virtual machine quiesce for the protection plan.
- 5 Click **Protect**.
The results of your choices appear under **Virtual machines** or **Intelligent VM groups**.

Edit protection settings for an AHV asset

If you have the necessary role permissions, you can edit the settings for a protection plan such as schedules.

To edit protection settings for an AHV asset

- 1 On the left, click **Workloads > Nutanix AHV**.
- 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 VM group
On the **Intelligent VM groups** tab, click on the group that you want to edit.
- 3 Click **Customize protection > Continue**.
- 4 If you have the necessary role permissions you can edit one or more of the following settings:
 - **Schedules and retention**
Change the backup start window.
See See [“Schedules and Retention”](#) on page 46.
 - **Backup options**
See See [“Backup options”](#) on page 46.
- 5 Click **Protect**.

Schedules and Retention

- ◆ Start window.
 - Set the window during which a backup can start.

Backup options

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

- 1 Select the server or host as an access 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 Access hosts.

Note: During backup of VM with backup host version prior to 9.1, if the VM with same UUID exists across different cluster. The **Last successful backup** status column for this VM is not updated. However, backup of VM is successful and you can view recovery points and recover.

2 Advanced options

To enable, See “[Prerequisite to Enable virtual machine quiescing](#)” on page 47.

- **Enable virtual machine quiesce**
- **Enable unquiesce snapshots if quiesced snapshots fail.**

By default, I/O on the virtual machine is quiesced before NetBackup creates the snapshot. In most of the 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.

Prerequisite to Enable virtual machine quiescing

- By default, the Nutanix Guest Tools (NGT) feature is disabled for a VM running in a Nutanix cluster. Nutanix recommends installing NGT and in some cases have pre-freeze and post-thaw scripts on the VM when you are planning to take application-consistent snapshots that is enable virtual machine quiescing.

Note: Application consistent backups minimum NetBackup Media Server version required is 9.1.

- To install NGT and add scripts, please see, [here](#).

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 **Protected By** column on 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 **Nutanix AHV**.
- 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, select the VM or the intelligent VM group.
- 3 Click **Remove protection > Yes**.

Under **Virtual machines** or **Intelligent VM group**, the asset is now 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 **Nutanix AHV**.
- 2 On the **Virtual machines** tab or **Intelligent VM groups** tab, select the VM or intelligent VM group.

The **Protection** tab shows the details of the asset subscription plans.

Note: If the asset has been backed up, but status indicates that it has not, see [“Errors for the Status for a newly discovered VM”](#) on page 80.

- 3 If the asset is not protected, click **Add protection** to select a protection plan.
See [“Protect AHV VMs or intelligent VM groups”](#) on page 45.

Recovering AHV virtual machines

This chapter includes the following topics:

- [Things to consider before you recover the AHV virtual machines](#)
- [About the pre-recovery check](#)
- [Recover an AHV virtual machine](#)
- [About Nutanix AHV agentless files and folders restore](#)
- [Prerequisites for agentless files and folder recovery](#)
- [SSH key fingerprint](#)
- [Recover files and folders with Nutanix AHV agentless restore](#)
- [Recovery target options](#)
- [Pre-recovery checks](#)
- [About Nutanix-AHV agent-based files and folders restore](#)
- [Prerequisites for agent based files and folder recovery](#)
- [Recover files and folders with Nutanix AHV agent based restore](#)
- [Limitations](#)

Things to consider before you recover the AHV virtual machines

Ensure that the recovery or backup host, that is added to the AHV access hosts, can communicate with the AHV cluster through port 9440.

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 AHV server and AHV credential validation.
- Availability of the AHV cluster.
- Available space with the storage container.

Recover an AHV virtual machine

You can recover a VM either to an original backup location or to a 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 **Nutanix AHV**.
- 2 Locate and click on the VM.
- 3 Click the **Recovery points** tab. In the calendar view on the left, click the date on which the backup occurred indicated with a green dot.

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

Note: You must have the **Restore Asset** permission on the VM to recover it.

- 4 On the image that you want to recover, select one of the following image recovery options:
 - **Recover**
Recovers from the default copy of the backup image.
 - **Recover from the default copy**

Recovers from the default copy of the backup image. This option is displayed if more than one copy exists.

- **nn copies**

Recovers 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 Review the **Restore to** values in **Recovery target**.

The default values are populated from the backup image of the VM.

- To recover to an alternate location change the default cluster in the **Restore to** option. Then click **Next**.

Note: You must have **View** and **View restore target** permissions on storage container or cluster to list down the expected storage container in target drop-down.

6 Review or change the **Recovery options** values.

For more information about the recovery options, refer to the *NetBackup for AHV Administrator's Guide*.

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.

Note: You must have the Overwrite Asset permission to use this option. See your security administrator if you need this permission.

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.

Create new VM ID instead of existing one

Create a new ID for the VM that is different from the existing value that was set during the backup.

Note: VM ID is the VM UUID.

- 7 Review or change the **Advanced Settings** values.

For more information about the recovery options, refer to the *NetBackup for AHV Administrator's Guide*.

Remove network interfaces	Remove the network interfaces that were set for the VM during the backup.
Retain MAC address	Retain the MAC address that was set for the VM during the backup.

- 8 Click **Next** to run the **Recovery overview**.

This runs the pre-recovery check on values provided in recovery target and recovery options pages. Checks connectivity and existence of AHV cluster and storage containers. Determines whether the storage container has available space and checks other requirements. For more information about the pre-recovery check, refer to the *NetBackup for AHV Administrator's Guide*.

- 9 Click **Start recovery**.

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

For information on the recovery status codes, see the *NetBackup administrator or the 90 Status Codes Reference Guide*.

About Nutanix AHV agentless files and folders restore

NetBackup 9.1 and later support Nutanix AHV agentless files and folders restore. It lets you restore individual files or folders to any target host. The target host can be a virtual machine hosted on AHV or other hypervisors or even a physical machine where the NetBackup client is not installed. This restore uses VxUpdate package of matching target host platform and deploys NetBackup recovery tool on the target host. Agentless files and folders restore performs clean-up of recovery tool and staging location after the restore process is completed. The recovery process uses a NetBackup host as a recovery host that has network connectivity with the target host. This recovery host can be either NetBackup server or a client.

Overview of the files and folders 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 or folders for restore along with the target host credentials. The required credentials are:

- Windows: User must belong to the local administrators group if UAC is disabled. If UAC is enabled, the user must be a domain user, added to the local administrator's group.
 - Linux: User must be a root or sudoer user having all permissions.
2. The primary server sends the requested data to the recovery host.
 3. The recovery host confirms that it has the necessary VxUpdate recovery package to perform restore. If it's not available, the recovery host downloads the required package from the primary server that uses VxUpdate.
 4. The recovery host copies recovery tool of VxUpdate package to target host. Linux recovery and target host use SSH protocol for recovery operation. Windows recovery and target host use WMI, SMB protocol for recovery operation.
 5. The data stream file containing files and folders to be restored is staged at a staging location on a recovery host.
 6. The file created on recovery host staging location is copied to staging location on a target host.
 7. The recovery tool is invoked, and selected files or folders are recovered along with ACLs and metadata details.
 8. NetBackup performs the necessary clean-up even if restore operation is successful or failed. All temporary files stored at the staging location on target host and recovery host are removed. However, in case of failures the evidence is collected from target host to recovery host collection with default configuration.
 9. NetBackup supports the following platforms for target host operating system for agentless file restore:
 - Windows
 - Red Hat Enterprise Linux (RHEL)
 - SUSE Linux (SLES)
 - Ubuntu

For target host operating system version supportability, see [NetBackup client](#) section in [NetBackup Software Compatibility List - 8.1 and later](#).

Prerequisites for agentless files and folder recovery

You can perform files or folder recovery only if the source AHV VM is running on specified operating system such as Red Hat Linux, or SUSE Linux, or Ubuntu, or

windows. Also, the file system must be compatible for creating file system mappings from the full agentless VM backup. For AHV compatibility, see [Support for NetBackup in Virtual Environments](#).

Note: If support for restore of individual files and folders for a non-supported OS is required, protect such VMs with NetBackup standard policy type.

Table 6-1 Prerequisites for files and folder recovery

Step overview	Description and reference
Agent based restore	■ Agent based restore is performed if the target host has NetBackup client or server installed. ■ NetBackup version of such client or server must be 8.1 and later for windows and 8.2 and later for Linux. Note: If you select the Linux version 8.1 or earlier, agentless restore options are displayed. ■ You must specify the NetBackup configured host name in the target host for agent based restore. ■ If the logged on NetBackup user has sufficient permissions, you can browse the list of NetBackup hosts and select one for restore files or folders. If a logged on user does not have sufficient RBAC permissions, target host needs to be manually specified. ■ You must specify the NetBackup configured host name or IP in the target host for agent based restore. If source AHV VM is running on a Linux platform, you can restore files or folders to any supported Linux platform target host. Note: If NetBackup is uninstalled from the target host, you can still initiate the agent based restore, however it fails.

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
Agentless restore	Agentless restore is performed if the target host does not have NetBackup client or server installed. - You need to specify the target host FQDN or IP address. - NetBackup detects if host is a non-NetBackup machine from NetBackup configuration and the agentless restore options are displayed. Note: Both IPv4 and IPv6 IP addresses are supported. In IPv6 the standard CIDR format is not supported.
Target host	- Target host is a host on which you want to restore files or folders from an AHV VM backup. The host name must be in FQDN format or IP address. - You can choose to restore files or folders to any target host, which is deployed on AHV, other hypervisors, or even a physical host. Note: Ensure that the target host is accessible from the recovery host. - Source and target host platforms must be homogeneous. Host files of windows source can be restored to windows target host and Linux source VM files on Linux target host. - The default target host staging directory on target host is the user's home directory. You can provide a custom staging location. Prerequisites: - NetBackup does not create target host staging location, the location must exist with write, and execute permissions. - The target host staging location must have enough space for restore operation. That includes restore file size, NetBackup restore package (~150MB for windows) and (~100 MB for Linux), space for NetBackup operation logs. Note: If staging location path is on a system drive, it must have enough space required for other running processes.

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
Linux target host	■ Agentless target machine must be running on supported OS platforms. For AHV compatibility, see Support for NetBackup Virtual Environment ■ Tar utility should be present on default path on target host and path is added in the system path variable. ■ NetBackup supports host name in ASCII format only. For host name with non-ASCII format you can use the IP address as target host. ■ Maximum number of SSH connections to target host is configurable and default value is 10. ■ SSH port should be open between recovery host and target host. If any firewall is configured, SSH port should be in Exception List in firewall. ■ To restore to network path on target host provide the correct export permissions. For example, <code>rw, sync, no_root_squash.</code>

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
SSH connection requirements	- Agentless restore to Linux target host is performed with the use of SSH service. It must be running on target host. - SSH communication time out on the target host must be greater than 5 minutes. - When you communicate with target host using SSH, NetBackup uses cipher <code>aes256-ctr</code>. - SSH version must be 1.2 or later. - Custom SSH port is supported. Note: Default SSH port is 22. - The following are supported: - Key exchange algorithms: <code>diffie_helman_group_exchange_sha256</code> <code>ecdh_sha2_nistp256</code> <code>cdh_sha2_nistp384</code> <code>ecdh_sha2_nistp521</code> <code>diffie_helman_group14_sha1</code> - Host key <code>ssh-rsa</code> <code>ssh-dss</code> <code>ecdsa-sha2-nistp256</code> <code>ecdsa-sha2-nistp384</code> <code>ecdsa-sha2-nistp521</code> - Hash Method <code>sha256</code> Hex encoded

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
SUDO user restore	■ Sudo user must be already existing on Linux target host. ■ Ensure that the non-root user is already configured in the sudoers file. Example: ■ <sudo-username> ALL = (ALL) ■ <sudo-username> ALL = (ALL) NOPASSWD ■ Linux sudo user must have ownership of the custom staging location along with read, write, and execute permissions. You can use SSH private key instead of password. See “SSH key fingerprint” on page 65.

Table 6-1 Prerequisites for files and folder recovery *(continued)*

Step overview	Description and reference
Windows target host	

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
	■ Agentless target machine must be running on supported OS platforms. For AHV compatibility, see Support for NetBackup Virtual Environment. ■ WMI must be configured and accessible between recovery and target host. For WMI and SMB requirements, see https://www.veritas.com/support/en_US/article.100040135. ■ Accepts host name in ASCII format. For Unicode host name, use the IP address instead of the host name. ■ The following services must be running on your windows hosts: ■ DCOM ■ RPC ■ WMI ■ File and Printer Sharing ■ By default, Admin share is enabled on host. If it is disabled. From GPO, user needs to enable Admin Share on the staging location drive or the drive on which the staging location is located. Note: By default, Administrator users have required permission for WMI and DCOM access. If any issue occurs in DCOM and WMI permissions, refer <i>Microsoft Documentation</i>.. ■ User or Group that is used to assign DCOM and WMI permissions: Out of two ways to assign the DCOM and WMI permissions, use one of the following options: ■ User must be part of Administrators group, you can assign the permissions to the Administrators group. ■ Assign the permissions to the specific user. ■ Supports UAC and non-UAC environments: ■ Built-in administrator and Domain user, added in the local administrator group of target host have required permissions to perform agent less restore. Note: UAC remote restrictions: For local user in Administrator group it is recommended to use agent based restore. But still user can perform

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
	agentless restore by disabling UAC filtering. To disable UAC remote restrictions, see here ■ Staging location requirements: ■ Default location is user's home directory, if custom path is provided, user must have access to it. ■ Must be an absolute path. Note: Soft-links, hard-links, network path, etc. are not supported. ■ It should have enough space for restore operation, includes: ■ The restore file size. ■ NetBackup restore package (~150MB). ■ Space for NetBackup operation logs. Based on verbose level log requirement would differ. Note: If the path is on system drive, it must have enough space required for other running processes. ■ Maximum character limit of the path is 260. However, NetBackup needs around 110 characters for formation of temporary location. Thus, you should pick a path that has less than 150 characters. ■ If staging location and restore location is on same drive, double space of restore size could be needed. ■ Parallel restore jobs with same user are supported. However, if same destination folders are specified the restored data might be in inconsistent state.

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
WMI and SMB requirements	- Agentless restore to Windows target host uses Windows Management Instrumentation (WMI) and Server Message Block (SMB) protocols. - Ensure that WMI and SMB ports are opened in your firewall settings. - Default DCOM port 135 - Default SMB port 445 - Dynamic ports 49152-65535 Note: Your environment can also have a static fixed port. - Encrypt the data transfer over the SMB by enabling SMB encryption. For more details, refer <i>Microsoft Documentation</i>. - Supports the SMB version 3.0. If your host has an older version, you can disable it. Refer, the Microsoft guidelines.

Table 6-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
Recovery host	Recovery host is a NetBackup media server/client installed host and use to communicate with provided target host. - NetBackup version of the recovery host must be 9.1 and later must have connectivity to target host. - Linux recovery host must have SSH connectivity to Linux target host and windows recovery host must have WMI and SMB connectivity with windows target host. - Recovery host must be of homogenous platform. Windows recovery host is required to restore files from windows AHV VM to target windows host. Similarly, Linux recovery host is required to restore file from Linux AHV VM to target Linux host. Note: To restore files to Ubuntu target host use either RHEL or SUSE as recovery host. - Recovery host with NetBackup 9.1 server or client is only supported. - Network path as staging location on recovery host works, provided the export permissions are correct. For example, <code>rw, sync, no_root_squash</code>. - Default staging location on recovery host is: - For Linux: <code>{install-path}/openv/var/tmp/staging</code> - For windows: <code>{install-path}\NetBackup\Temp\staging</code> - Default staging location can be changed using <code>bpsetconfig</code>. - Execute <code><NetBackup path>/bin/admincmd/bpsetconfig</code>. - Set <code>AGENTLESS_RHOST_STAGING_PATH = <Path></code>.
Other	Ensure that SUSE target host 'PasswordAuthentication as Yes' in the '/etc/ssh/sshd_config' file. Then restart the 'ssh' service. Note: By default SUSE target hosts have <code>passwordAuthentication</code> value set to No.

SSH key fingerprint

To Obtain the SSH key fingerprint of the Linux target host:

- 1 Use the following command on RHEL or SUSE OS target host to get the SHA256-based RSA key.

```
cat /etc/ssh/ssh_host_rsa_key.pub |awk '{print $2}' |base64 -d |sha256sum  
awk '{print $1}'
```

Note: The output of the commands is the RSA key. Similarly, change the public key path, execute the command to get ecdsa or DSS SSH key fingerprint configured on target host.

- RSA key example:

```
cat /etc/ssh/ssh_host_rsa_key.pub |awk '{print $2}'|base64 -d |  
sha256sum |awk '{print $1}'
```

- Command output:

```
b2352722053ac9f40bc1XXXXXXXXXXXXXXXXXXXXXXXXXXXX419fa241ba9431fd6b9
```

- 2 Copy the RSA fingerprint. You can provide this SSH key fingerprint when you add the target host details. Or you can also verify the displayed SSH key fingerprint after you click, **Fetch SSH Key fingerprint** on the **Recovery Host** page.

To generate SSH private key:

- 1 Execute the following commands on Linux target host:

- `ssh-keygen -t rsa`
- `-t` option supports "ecdsa | rsa | dss"

- 2 You must add/append target host public key in target `vm ~/.ssh/authorized_keys` file.

Recover files and folders with Nutanix AHV agentless restore

To recover files and folders with Nutanix AHV agentless restore

- 1 Ensure that the target host is powered on and has a network connectivity to recovery host to be used in restore process.
- 2 On the left pane, click **Nutanix AHV**.
- 3 Locate and select the AHV VM that contains the files and folders for restore.
This VM would be referred as source VM.
- 4 Click the **Recovery points** tab. In the calendar view, select the date on which the backup occurred.
- 5 The available images are listed in rows with the backup timestamp for each image.
- 6 On the image you want to recover from, click **Recover > Restore files and folders**.
- 7 In the **Select files** pane, specify the files and folders you want to recover then click **Next**. These files or folders are referred further as source files or folders.
- 8 Click **Next**.
- 9 On the **Recovery target** page, do the following:
 - Enter the IP/hostname manually.
 - If required, enter the staging location on the target host.
 - Select the appropriate file restore option.
 - Select the correct recovery host.
 - Add correct credentials for based on OS type.

See ["Recovery target options"](#) on page 67.
- 10 On the **Recovery options** page, select from the followings:
 - **Append string to file names:** Appends the specified string to the destination file names before any file extension. This value only applies to files.
 - **Overwrite existing files:** Overwrites the files or folders if they exist in the destination location with the same name.
 - **Restore directories without crossing mount points**
 - **Create new files for hard links**

- **Rename targets for soft links**

Note: Create new files for hard links and **Rename targets for soft links** options are enabled only to restore everything to a different directory.

11 Click **Next**.

12 On the **Review** page: Review page displays the status of the pre-recovery check. NetBackup performs the pre-recovery validation to confirm if restore job will run successful using your provided inputs.

See “[Pre-recovery checks](#)” on page 71.

- If pre-recovery fails, probable causes of failure are displayed. Click the **Change** button for a specific input that needs to be corrected.
- If pre-recovery is successful, click **Start recovery**.

Recovery target options

Table 6-2 Recovery target options

Step overview	Description and reference
Target host	■ Target host field is pre-populated with the source AHV VM hostname/IP stored during last successful discovery for respective AHV cluster of the VM. Warning: Agent based restore is performed if NetBackup client is installed and configured with provided hostname or IP. ■ If you want to perform restore on another NetBackup client, click search and select the required client from list. Note: Ensure that you select clients with homogenous platforms. ■ If search option is unavailable, manually enter target host. ■ If you want to perform restore on host on which NetBackup client is not installed, enter the host FQDN or IP in target host. Agentless restore options are displayed.

Table 6-2 Recovery target options (*continued*)

Step overview	Description and reference
Agentless restore options	■ Change staging location on target host: If you want to provide a different staging location other than the default staging location, enter the desired path. Staging location path must have only ASCII characters. Note: Default staging location is user's home directory. ■ File restore options: Based on your requirement, select one of the following appropriate files restores options between: ■ Restore everything to original directory ■ Restore everything to different directory Provide different directory path to restore. ■ Flatten existing directory structure Select this option to restore everything to a single directory without creating any subfolders when files are selected from different directories.
Recovery Host	■ Recovery host field is pre-populated with backup host, that was used to during the backup operation for selected AHV VM. Note: Recovery host field is empty if the selected VM and backup host platform are not homogenous. Note: To restore files to Ubuntu target host use either RHEL or SUSE as recovery host. ■ Click on search to select another recovery host. It shows a list of compatible Media servers. If you want to select NetBackup client as a recovery host, click on Media servers > Clients. ■ If search option is not available, manually enter recovery host. Note: Recovery host should be of homogeneous platform as the source VM and NetBackup 9.1 or later server or client must be installed. ■ If you have performed restore on same target host earlier, recovery host is pre-populated with the previously used recovery host based on pre-assigned permissions provided to user performing this restore.

Table 6-2 Recovery target options (*continued*)

Step overview	Description and reference
Linux SSH Connectivity	

Table 6-2 Recovery target options (*continued*)

Step overview	Description and reference
	For selected source Linux VM SSH connectivity the following options are displayed: ■ Target host SSH port Specify SSH port of the target host. Default value is 22. If you have performed restore on same target host earlier, SSH port is pre-populated with previously used value based on pre-assigned permissions by the user performing this restore. ■ Target host SSH key fingerprint To authenticate target host, provide SSH key fingerprint in hexadecimal format. ■ You can either manually enter target host SSH key fingerprint or click on Fetch SSH Key fingerprint. ■ Fetch SSH Key fingerprint: If Fetch SSH Key fingerprint option is not available, you must provide SSH key fingerprint manually. See “SSH key fingerprint” on page 65. ■ If you have performed restore on same target host earlier, SSH key fingerprint is pre-populated with previously used value based on pre-assigned permissions by the user performing this restore. You can overwrite pre-populated value for re-establishing trust. ■ Fetch SSH Key fingerprint ■ Display list of SSH key fingerprint along with NetBackup supported Key types configured on target host. ■ Select one of the listed fingerprints and click OK. NetBackup establishes trust with target host using selected fingerprint. ■ Target host credentials ■ Username Specify target host username. This user must be either root or non-root sudoer. Sudoer user See “Prerequisites for agentless files and folder recovery” on page 54. ■ Provide password Select this option to choose password-based authentication. ■ Password Specify target host password for provided user. ■ Provide SSH private key Select this option to

Table 6-2 Recovery target options (*continued*)

Step overview	Description and reference
	choose SSH private key based authentication. See "SSH key fingerprint" on page 65. ■ SSH private key Specify SSH private key. ■ Key passphrase If SSH private key is created using passphrase, specify key passphrase
Windows WMI Connectivity	■ Username Specify target host username. This user can be domain or local user and must be part of local administrator group. localusername or domain\username is supported format for username. ■ Password Specify target host password for the specified user.

Pre-recovery checks

Table 6-3 Pre-recovery checks

Validation	Description and reference	Input Source
Recovery host space	Checks for the required space on recovery host staging location.	Recovery host
Target host connectivity	Checks if target host is accessible from recovery host.	Target host and Target host port
Target host credential	Checks if provided target host credentials are valid.	Target host credentials
Target host staging location on a local disk	Checks if target host staging location is not a network path.	Target host staging location

Table 6-3 Pre-recovery checks (*continued*)

Validation	Description and reference	Input Source
Target host staging location space	Checks if the required space is available on target host staging location. Note: Required space is total size of selected file with space required for NetBackup restore package and space needed for logs and other files.	Target host staging location
Target host staging location permissions	Checks if provided user is an owner and has RBAC permissions on target host staging location.	Target host staging location
Target host default staging location path	Checks if provided target host staging location path contains valid characters. NetBackup does not support non-ASCII characters in target host staging location path.	Target host staging location
Target host operating system	Checks if target host has a supported OS.	General
VxUpdate package	Checks if required VxUpdate package is available on primary server.	General
Linux target host specific checks		
Target host SSH key fingerprint	Checks if target host SSH key fingerprint is valid to establish trust with target host from the recovery host.	Target host SSH key fingerprint
Tar existence on the target host	Checks if <code>tar</code> is available on target host.	Target host

About Nutanix-AHV agent-based files and folders restore

NetBackup 9.1 and later support Nutanix-AHV agent-based files and folders restore of individual files and folders. The agent-based restore lets you restore individual Nutanix-AHV files to a host that has a NetBackup client. The agent-based target host can be a virtual machine hosted on AHV or other hypervisors or even a physical machine where the NetBackup client is installed.

Prerequisites for agent based files and folder recovery

- You can perform individual files and folders recovery from source AHV VM backed up image, only if guest operating system and the file system, is compatible for creating file system mappings. Refer Nutanix AHV SCL for guest operating system and file system support for individual files restores, [Support for NetBackup version in virtual environments](#).
- You can perform individual files recovery from source AHV VM backed up. When backup is performed with NetBackup version 9.1 and later for NetBackup primary server, media server and backup host.
- Agent based restore is performed if the target host has NetBackup client or server installed. NetBackup version of such client or target host must be 8.1 or later for Windows and 8.2 or later for Linux.

Note: If you select the Linux version 8.1 or prior, agentless restore options are displayed.

You must specify the NetBackup configured host name or IP in the target host to follow the agent based restore.

- If logged-in NetBackup user has sufficient permission to view NetBackup hosts, then you can browse the list of NetBackup hosts and pick the one where you want to restore files or folders.
- If logged-in user does not have sufficient permissions to view NetBackup hosts, then you need to manually specify the target host. You must specify the NetBackup configured hostname or IP in the target host for the agent based restore

- Following are minimum RBAC permissions required for a user to do agent based files and folders restore. For more details on required RBAC permissions, see [NetBackup Web UI Administrator Guide](#).

Table 6-4 Permissions for all AHV assets

Operation	Description	Additional required operations	Additional optional operations
Granular Restore	Restore individual files or folders from a AHV asset. This permission is required on the Source VM.	Global > NetBackup management > NetBackup backup images > View Global > NetBackup management > NetBackup backup images > View contents Global > NetBackup management > NetBackup hosts > View Assets > Assets > Restore files using client	Assets > Assets > Overwrite files and folders

Recover files and folders with Nutanix AHV agent based restore

- 1 Ensure that the target host is powered on and has a network connectivity to recovery host to be used in restore process.
- 2 On the left, click **Nutanix AHV**.
- 3 Locate and select the AHV VM that contains the files and folders for restore. This VM would be referred further as source VM.
- 4 Click the **Recovery points** tab. In the calendar view, select 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 **Recover > Restore files and folders**.
- 6 On the **Select files** pane, specify the files and folders you want to recover then click **Next**. These files or folders are referred further as source files or folders.
- 7 On the **Recovery target** page, do the following:
 - Select the target host.
 - Target host input must be FQDN or IP address. If you have permission to view hosts, click on the search icon it will show the hosts where NetBackup client is already present select the required host.

Note: Only NetBackup version 8.1 or later are available in the dropdown.

- Select the appropriate file restore option.
- See [“Recovery target options”](#) on page 67.
- 8 On the **Recovery options** page, select one of the following:
 - **Append string to file names:** Append the specified string to the destination file names before any file extension. This value only applies to files.
 - **Overwrite existing files:** Overwrites the files or folders if they exist in the destination location with the same name.
 - **Restore directories without crossing mount points**
To skip over file systems mounted in the selected directories. Clear this checkbox to restore file systems that are mounted in the selected directories
 - **Create new files for hard links**
 - **Rename targets for soft links**

Note: **Create new files for hard links** and **Rename targets for soft links** options are enabled only to restore everything to a different directory.

- 9 Click **Next**.
- 10 On the **Review** page: Review all the previously selected options.
- 11 Click **Start recovery**.

Limitations

- Cross-platform recovery of individual files is not supported. You can restore windows files to only windows guest operating systems and Linux files to only supported Linux guest operating. That implies, the recovery host must be the same platform as the files that you want to restore.
- In a recovery process, NetBackup recreates the links between a hard link and its original file. Only in this case, the link file, and its target file must be restored in the same job.

Note: If each file is restored individually in separate restore jobs, they are restored as separate files and the link is not re-established.

- For dual-boot virtual machines, NetBackup does not support recovery of individual files or folders.
- For client platform and file system support and limitations, see https://www.veritas.com/content/support/en_US/doc/NB_70_80_VE.
- The **Flatten existing directory structure** and **Append string to file names** options are only applicable to files. They are not available for directories.
- 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 such case, the last file 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, do not select **Flatten existing directory structure** when restoring multiple files with the same name.
- If a backup and a restore occur simultaneously on the same VM, one or both jobs can have unexpected results.

Note: If a backup or a restore exits with a non-zero NetBackup Status Code, one possible cause is simultaneous jobs occurring on the same VM.

- If selected restore data contains any hidden files such as `.bashrc`, `.bash_history`, **Append string to file names** restore option is not supported.
- Nutanix agentless restores can only be used to the restore of files and folders.

- Restore job fails, if NetBackup does not have sufficient privileges to the staging directory or if there is insufficient space in the staging directory.

Note: Veritas does not recommend Nutanix AHV agentless restore, if a NetBackup client already exists on the target VM. The NetBackup administrator must use the agent-based restore in such cases.

- On windows target host, restore destination to a mapped drive, is not supported.
- NetBackup does not support the communication with windows target host with a use of `openSSH`. In such cases restore job fails.
- NetBackup does not support non-ASCII characters in target host staging-location path.
- NetBackup only supports NTLM authentication type for windows target host.
- AHV images backed-up prior to 9.1 release, cannot be restored from web UI. To restore these images user must use Java GUI.
- AHV backup images are available on web UI, even if the backup was taken from Java GUI provided the backup host has NetBackup version 9.1 or later.
About the backup images on web UI:
 - If the asset discovery is successful, and after that backup is taken from java GUI, backup images are available on web UI.
 - If the primary sever and backup host is upgraded to 9.1, backup is taken from java GUI, and then if you configure web UI, you must run asset discovery to see the backup images.
 - If the primary sever is upgraded to 9.1 but backup host is version is still prior to 9.1 and backup is taken from java GUI, and then if you configure web UI, you cannot see the backup images even after asset discovery.

Troubleshooting AHV operations

This chapter includes the following topics:

- [Troubleshooting tips for NetBackup for AHV](#)
- [Error during AHV credential addition](#)
- [Error during the AHV virtual machines discovery phase](#)
- [Errors for the Status for a newly discovered VM](#)
- [Error run into while backing up AHV virtual machines](#)
- [Error while restoring AHV virtual machines](#)

Troubleshooting tips for NetBackup for AHV

For more information about AHV 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 AHV-related errors, check **Alerts** on AHV Prism console.
- For backup job failures:
 - Check the **Job details** section for the job in Activity monitor.

- Check the `bpbkar` and `VxMS` logs.
- For AHV-snapshot related errors, check **Alerts** on AHV Prism 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 AHV-related errors, check **Alerts** on AHV Prism console.

Error during AHV credential addition

Table 7-1 Error during AHV credential addition

Error message or cause	Explanation and recommended action
Discovery of virtual machines and credential validations is supported by NetBackup 9.1 or later. The selected server/backup host has NetBackup version 8.3.	Upgrade the server/backup host or select another server/backup host with the required NetBackup version.

Error during the AHV virtual machines discovery phase

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

Table 7-2 Error run into during the AHV virtual machines discovery phase

Error message or cause	Explanation and recommended action
The AHV assets are not discovered after the correct AHV cluster credentials are added. The VM discovery operation fails.	Run discover now and retry the backup. The maximum allowed length of the AHV cluster name is 255 characters, however, if the characters exceed 95, the asset discovery fails. Workaround: ■ Ensure that the AHV cluster name has 95 or fewer characters.

Table 7-2

Error run into during the AHV 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.	Ensure that the query specified in the policy or intelligent VM group is correct. The VMs that need protection are added recently to AHV cluster or the VM configuration has changed and the autodiscovery or discover now was not triggered. ■ The asset discovery does not work if the AHV cluster credentials are added using <code>tpconfig</code>. Workaround: From NetBackup web UI, click Discover for the specified AHV cluster. Ensure that you add the AHV cluster credentials using API or NetBackup web UI.

Errors for the Status for a newly discovered VM

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

Table 7-3 Error for the Status for a newly discovered VM

Error message or cause	Explanation and recommended action
The last successful backup status of VM indicates that it has not been backed up.	In the NetBackup web UI, the last successful backup status for a newly discovered VM does not indicate that it is backed up. In some circumstances, such as Intelligent VM group a new VM is backed up matching the query provided 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. ■ In the NetBackup web UI, the last successful backup status of the VM is not updated indicating that it has not been backed up. Workaround: ■ If you encounter a similar situation, you can still browse the recovery points and recover them. However, it is only after the Discovery is triggered on the cluster and another backup of the VM successfully completes after discovery, then the last successful backup status is updated.

Error run into while backing up AHV virtual machines

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

Table 7-4 Error while backing up AHV virtual machines

Error message or cause	Explanation and recommended action
After a NetBackup backup operation, the VM snapshot on the AHV cluster is not deleted.	If a disk attached to the VM is in an inactive state, then the AHV cluster does not delete the VM snapshot after a backup operation is complete. Workaround: - 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.
MSiSCSI service is disabled. Enable the MSiSCSI service on the backup host.	Enable Microsoft iSCSI Initiator Service (MSiSCSI) service on the windows backup host and re-run the job.
Unable to establish a connection. Verify that the iSCSI service is installed and running.	- For Windows: Enable Microsoft iSCSI Initiator Service on backup host. Note: This is displayed as an error only for windows OS. - For Linux, this error is displayed in form of a warning and it falls back to use the NFS for backup/restore. Ensure that the backup host is added in <code>Filesystem Whitelists</code> option in Nutanix UI for backups to work via NFS transport. For Linux to use iSCSI: Install / Enable iSCSI Initiator package on backup host and re-run the job.
Authentication failed. Verify whether the provided initiator CHAP is correct.	Either the provided CHAP key is invalid or the iSCSI initiator name is not unique for each backup/recovery host. Please set iSCSI initiator name unique for each backup/recovery host.
Failed to get an external data service IP address for iSCSI. Re-run the job after setting IP address on the Nutanix cluster: {Nutanix AHV clusterName}.	Set the external data service IP address for iSCSI on the Nutanix AHV cluster. More details See "Configure Nutanix AHV cluster" on page 26. Note: For Linux it will fall back to use the NFS for backup/restore.
Backup hosts that have Linux or Windows OS with existing NetBackup version are not supported for media server load-balancing. Upgrade backup host to the latest version.	This error occurs when the Automatic option is selected for backup hosts in the Nutanix protection plan. Upgrade backup host to latest NetBackup version.

Table 7-4 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
For NetBackup media server load balancing, ensure that the backup hosts have either Red Hat Enterprise Linux, SUSE Linux Enterprise Server or Microsoft Windows operating system.	This error occurs when the Automatic option is selected for backup hosts in the Nutanix protection plan. For Nutanix AHV, supported media server are: ■ Red Hat Enterprise Linux ■ SUSE Linux Enterprise Server ■ Microsoft Windows operating system
Existing version of NetBackup on the media server does not support the incremental backup schedule.	Upgrade NetBackup to the latest version on the backup host.
Unable to set resource limit for specific Nutanix clusters.	If the clusters for which the resource limit is set are deleted from NetBackup environment, in some cases, + Add option is disabled to set the resource limit. Recommended Action Delete the resource limit for deleted clusters and then set the resource limit for rest of the clusters.
Snapshot jobs fail with error code 156 with the following job details: <pre>Critical bpbrm (pid=30139) from client 9c5dcb07-65d2 -4761-b861-9e517edcf5b6_ <Nutanix-cluster> vxindia.veritas.com: FTL - Value 2 that specifies GUID is not supported for the nameuse</pre>	If a protection plan is created using Backup option > Select server or host to use for backups > Automatic and the selected storage unit is configured with media servers with both NetBackup 9.1 or prior versions. And when this protection plan is used to back up the AHV VMs or Intelligent VM group, the snapshot job might fail. Recommended action All the media servers, which are configured in the selected storage unit, must be upgraded to NetBackup 9.1. To avoid job failures as upgrades for other media servers are in progress, in the Protection > Customize Protection > Backup options option, manually select a specific media server or backup host as a server or host to use for backups instead of default Automatic option. It is recommended to use already upgraded media server. Once upgrade of all media servers completes, use Protection > Restore Original Settings to go back to original settings.

Table 7-4 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Error 1 <pre>iscsiadm: Could not login to [iface: default, target: iqn.2010-06.com.nutanix: nbubackup -2d29da9d-f964- 4157-9595-f0319090bb01-tgt0, portal: xx.xx.xx.xx,3260]</pre> <pre>iscsiadm: initiator reported error (24 - iSCSI login failed due to authorization failure)</pre> <pre>iscsiadm: Could not log into all portals</pre> Error 2 <pre>iscsiadm: Could not execute operation on all records: encountered iSCSI database failure</pre> Error 3 <pre>iscsiadm: could not read session targetname: 5</pre> <pre>iscsiadm: could not find session info for session28</pre>	These errors are seen in the successful job details tab of backup/restore jobs. These errors are output of <code>iscsiadm</code> command execution. These errors are intermittent and may occur due to heavy load on the iSCSI network. NetBackup does a retry operation to fix these errors. Once retry operation is successful, backup/restore job is also successful. Recommended action No action is required on NetBackup side. User can still troubleshoot the <code>iscsiadm</code> and ensure correct iSCSI installation/configuration to avoid such errors.

Error while restoring AHV virtual machines

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

Table 7-5 Error run into while restoring AHV 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.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Unable to change the AHV cluster while modifying the recovery destination.	If you cannot see the list of the AHV clusters, you might not have access to the AHV clusters in RBAC. Contact the NetBackup security administrator to resolve this issue.
Pre-recovery check runs successfully when a VM with the same UUID exists in the AHV 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, cannot 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 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 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 AHV 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 AHV VM images, for the recovery host, select the access host in the target domain as a recovery host or select the target primary server.
MSiSCSI service is disabled. Enable the MSiSCSI service on the recovery host.	Enable Microsoft iSCSI Initiator Service (MSiSCSI) service on the windows backup recovery and re-run the job.
Failed to connect to the recovery host.	The recovery host used for agentless restore is not reachable. Recommended action: Ensure that recovery host is reachable from primary server and has a NetBackup media or client installed on it.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
The specified recovery host must be at NetBackup version 9.1 or later to support agentless restores.	Agentless restores of files or folders require recovery host with NetBackup version 9.1 or later. Recommended action: Verify the NetBackup version on recovery host. It should be 9.1 or above. On UNIX NetBackup servers and clients, check the <code>/usr/openv/netbackup/bin/version</code> file. On Windows NetBackup servers, check the <code>install_path\netbackup\version.txt</code> file.
Recovery host staging location does not exist.	staging location path does not exist on recovery host for agentless restore. Recommended action: - Ensure that the default staging location path or the user-configured staging location path for recovery host is valid. NetBackup uses the following on recovery host or as default staging location: - For UNIX: <code>{installpath}/openv/tmp/staging</code>. - For Windows: <code>{installpath}\Netbackup\Temp\staging\</code>. - Ensure that the staging location path that is used exists. For user-configured staging location, check if valid path on recovery host is specified in <code>bp.conf</code> parameter AGENTLESS_RHOST_STAGING_PATH = "path".
Tar image not found at staging location on recovery host.	No tar image was found on recovery host staging location which is required for agentless restore. Recommended action: Contact Veritas Technical support and share <code>bpVMutil</code> log from the recovery host.
Internal error has caused failure of recovery validation.	Internal error was occurred while executing Pre recovery validations for agentless restore. Recommended action: Save the bpVMutil logs on recovery host and contact Veritas Technical support.
Not enough space available on recovery host.	The recovery host may not have enough space to copy the selected files at staging location for agentless restore. Recommended action: Ensure that sufficient free space is available on the recovery host staging location based on the total size of files or folders selected. Or select a different recovery host with sufficient free space for performing agentless restore.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Tar utility is not present on the target host.	Failed to find tar utility on the target host, required for agentless restore. Recommended action: Retry after deploying the tar utility.
Either specified staging location does not exist on target host or the user does not have required permission to access.	Recommended action: Ensure that the target host staging location exists, and the user has sufficient permissions to access the location.
The user does not have required permission on the target host staging location	The user does not have required permission to proceed with restore on the target host. Recommended action: Ensure that the target host staging location exists and the user has minimum Write and Execute permission on the staging location.
The user does not have root/administrator privileges. To restore files and folders, provide user with root/administrator privileges.	The user does not have required permission to proceed with restore on the target host. Recommended action: Provide the credential which is part of local administrator group on the Windows target host. For Linux target host, use the credential which is root or sudo account with ALL permissions.
Admin share of target host is not accessible from the recovery host.	Admin share of remote host is not accessible from the recovery host to perform agentless restore. Recommended action: ■ Ensure that firewall exceptions are set up correctly. ■ Ensure that File and Printer Sharing is enabled. ■ Ensure that GPO/Software Restriction Policy or Antivirus is not blocking access. ■ Ensure that target host is accessible and ensure that the correct credentials are entered and have proper permissions.
For agentless files or folders restore in User Account Control (UAC) environment, please provide credential of domain user which is part of local Administrator group on the Windows target host.	Recommended action: For agentless restore in User Account Control (UAC) environment, provide credential of domain user which is part of local Administrator group on the windows target host.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Agentless restore is not possible.	Received unexpected reason for agentless restore failure. Recommended action: Contact Veritas Technical support and share appropriate logs.
Operating systems do not match. Ensure that the operating system of recovery host matches with the backed-up VM operating system.	Agentless restore is possible only when operating system of recovery host and backed up VM is same. Recommended action: Use alternate recovery host of the same operating system as that of backed up VM.
Failed to retrieve the backup image operating system.	Unable to retrieve operating system of backup image for performing agentless restore. It is an internal error.
Recovery host operating system is not compatible with provided communication mode. Ensure that the operating system of recovery host and provided communication mode are compatible.	Recovery host OS type and communication type provided in the agentless recovery or pre-recovery check request are not compatible. Recommended action: Verify that Recovery host OS type and Communication type must be compatible. If recovery host is: ■ Linux: Communication type must be SSH. ■ Windows: Communication type must be WMI.
Target host SSH private key is invalid.	The 'sshKey' field of the agentless recovery or pre-recovery check request must be valid and non-empty ssh private key of target host. Recommended action: Verify that the <code>sshKey</code> field is specified, if the authentication type is <code>SSH_KEY</code> and that it is not empty.
Target host operating system is not supported for the agentless files or folders restore.	Target host operating system is not supported since agentless restore requires recovery packages to be deployed on target host. Recommended action: Only SUSE Linux Enterprise Server, Microsoft Windows, Red Hat Enterprise Linux (RHEL), and Ubuntu are supported platforms. Refer to the NetBackup Client Compatibility List for the supported platforms for this feature at the following URL:\nhttp://www.netbackup.com/compatibility.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Invalid target host username or password.	The username and password fields in authentication details of the agentless recovery or pre-recovery check request must be specified. Recommended action: Verify that the username and password field in authentication details of the recovery and pre-recovery check request are specified, correct, and are not empty.
Target host staging location path contains non-ASCII characters.	Target host staging location path supports ASCII characters only. Recommended action: Provide custom staging location on target host with ASCII characters only.
Specified path does not exist on the local disk	Target host staging location should not be the network path. Recommended action: Specify a custom staging location on target host which is on its local disk.
WMI connection to the target host is failed.	WMI connection to the target host is failed from recovery host. Recommended action: ■ To connect with WMI and DCOM service, user must have the required permission to connect with the remote WMI service. ■ Firewall exceptions are set up to allow WMI traffic through the firewall. ■ GPO/Software restriction policy or an antivirus is not blocking the access. ■ Ensure that target host is accessible. Validate the given target host credentials. ■ Ensure that the target host trust relationship with domain is intact. When you communicate across domains, two-way trust relationship between those domains must exist.
Unable to find the specified file on the remote server.	Unable to find the specified file on the remote server. Recommended action: Ensure the specified staging location on target host exists or specify another valid staging location.
File exists with same name as the directory.	There is a pre-existing file on target host with same name as that of directory path provided as staging location. Recommended action: Check for a pre-existing file on remote host with the same name and path as that of staging location. If it exists, either rename or remove that file. Or specify an alternate staging location.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Failed to validate administrative privileges for the user.	Target host user doesn't have the administrative privileges for the agentless files and folders restore operation to proceed. Recommended action: Use the credential which is part of local administrator group on the Windows target host. For Linux target host, use the credential which is root or sudo account with ALL permissions.
Failed to connect a network resource using windows API.	Admin share of target host is not accessible from the recovery host for agentless files or folders restore. Recommended action: As a part of agentless files and folders restores operation, SMB Admin share is created from recovery host on target host with the credential provided by user. This error is usually seen when target host for agentless restore has windows OS and admin share of target host is not accessible from the recovery host. Ensure that following things on target host. ■ Firewall Exceptions are set up correctly. ■ File and Printer Sharing is enabled. ■ GPO/Software restriction policy or antivirus is not blocking access. ■ Target host is accessible with valid credentials.
Unable to retrieve user's home directory on the target host. Specify the custom staging location.	User's default staging location, that is home directory can't be retrieved on the target host. A valid custom staging location path must be entered by user. Recommended action: Ensure that user's home directory exists or try with a valid custom staging location.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Failed to establish SSH session with host.	Ensure all of the following criteria are satisfied and then retry. ■ aes256-ctr is the supported cipher used for communication. Ensure that this cipher is supported both in recovery host and target host. ■ Ensure at least one of the following the Hash-based Message Authentication Code (HMAC) protocol is supported on both recovery host and target host: ■ hmac-sha2-256 ■ hmac-sha2-512 ■ Ensure the method used for generating host key is one of the following: ■ ECDSA_SHA2_NISTP256 ■ ECDSA_SHA2_NISTP384 ■ ECDSA_SHA2_NISTP521 ■ SSH_RSA ■ SSH_DSS
Failed to verify SSH key fingerprint of host.	SSH key fingerprint of target host provided is not correct. Recommended action: Verify the SSH key fingerprint of the target host and retry.
Failed to authenticate the host with provided username or password.	Target host authentication is failed with the provided username and password. Recommended action: Verify the username or password of the target host is correct and retry.
Failed to authenticate the host with specified SSH key.	Target host authentication is failed with the provided SSH private key. Recommended action: Verify the SSH private key along with key passphrase if used to generate SSH private key of the target host and retry. Ensure that the corresponding public key is present in the <code>authorized_keys</code> file in <code>/root/.ssh</code> folder at the target host.
Matching SSH Key fingerprint host key method not found on target host.	Unable to find the specified SSH key fingerprint host key method on the target host Recommended action: Ensure that either supported host key method of the specified SSH key fingerprint is available on target host or provide SSH fingerprint of the host key method that is configured on target host.
The restore fails when you restore individual files to a virtual machine that has NetBackup client software.	When you restore individual files to a virtual machine that has a NetBackup client, make sure that a firewall does not interfere with the restore. If a firewall stops the restore, turn off the firewall and retry the restore.

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Mount points not available when restoring files from a Linux virtual machine.	For Linux virtual machines, only the <code>ext2</code>, <code>ext3</code>, <code>ext4</code> and <code>xfs</code> file systems are supported for individual file restore. If a partition is formatted with some other file system, the backup succeeds but NetBackup cannot map the file system addresses of the files. As a result, NetBackup cannot restore individual files from that partition. Only the files that were on <code>ext2</code>, <code>ext3</code>, <code>ext4</code> and <code>xfs</code> partitions can be individually restored. Note: To restore individual files from their original mount points, the "/" (root) partition must be formatted as <code>ext2</code>, <code>ext3</code>, <code>ext4</code> or <code>xfs</code>. If the "/" (root) partition is formatted with a different file system such as ButterFS, the mount points cannot be resolved. In that case, you can restore <code>ext2</code>, <code>ext3</code>, <code>ext4</code> or <code>xfs</code> files from the <code>/dev</code> level (such as <code>/dev/sda1</code>). You cannot restore the files from their original mount point level.
For Linux VMs without persistent device naming, multiple disk controllers such as IDE, SCSI, and SATA may complicate the recovery of individual files.	This issue occurs because of non-persistent device naming, such as <code>/dev/sda</code> and <code>/dev/sdb</code>, may cause unexpected mount point changes after a restart. If the VM has a SCSI disk and SATA disk, the Restore files and folders > Add files and folders navigation interface may show incorrect mount points for the VM's files. For example, the files originally under <code>/vol_a</code> might appear under <code>/vol_b</code> when you browse to restore them. The restore is successful, but the restored files may not be in their original directories. Recommended action: Search for the files on the restored VM and move them to the proper locations. To prevent this issue on Linux VMs with multiple disk controllers, Veritas recommends a persistent device-naming method for mounting the file systems. When persistent naming is in place, device mounting is consistent and this issue does not occur when you restore files from future backups. For persistent device naming, you can mount devices by UUIDs. The following is an example of the <code>/etc/fstab</code> file that contains the devices that are mounted using UUIDs: ■ <code>UUID=93a21fe4-4c55-4e5a-8124-1e2e1460fece /boot ext4 defaults 1 2</code> ■ <code>UUID=55a24fe3-4c55-4e6a-8124-1e2e1460fadf /vola ext3 defaults 0 0</code> To find the device UUIDs, you can use either of the following commands: ■ <code>blkid</code> ■ <code>ls -l /dev/disk/by-uuid/</code>

Table 7-5 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
For Ubuntu VMs without persistent device naming, the Restore files and folders > Add files and folders navigation interface may show incorrect mount points for the VM's files and recovery of individual file may fail.	This issue occurs because of non-persistent device naming and may cause unexpected mount points changes. For the Ubuntu VM, the Restore files and folders > Add files and folders navigation interface may show incorrect mount points for the VM's files. For example, files and folders might appear under <code>/dev/ubuntu-vg/ubuntu-lv</code> when you browse to restore them and recovery of individual files may fail. Recommended action: To prevent this issue on Ubuntu VMs, Veritas recommends a persistent device-naming method for mounting the file systems. When persistent naming is in place, device mounting is consistent and this issue does not occur when you restore files from future backups. For persistent device naming, you can mount devices by UUID. The following is an example of the <code>/etc/fstab</code> file that contains the devices that are mounted using UUIDs: ■ <code>UUID=93a21fe4-4c55-4e5a-8124-1e2e1460fece /boot ext4 defaults 1 2</code> ■ <code>UUID=55a24fe3-4c55-4e6a-8124-1e2e1460fadf /vola ext3 defaults 0 0</code> To find the device UUIDs, you can use either of the following commands: ■ <code>blkid</code> ■ <code>ls -l /dev/disk/by-uuid/</code>
Unable to perform agent-based restore if the elected target host is Linux 8.1.	NetBackup does not support agent-based restore for 8.1 Linux platform. In case of NetBackup 8.1, agent-based restore is supported for windows platform only and not for Linux platform. Recommended action Upgrade Linux target host to 8.2 or later for agent based restores.

API and command line options for AHV

This chapter includes the following topics:

- [Using APIs and command line options to manage, protect, or recover AHV virtual machines](#)
- [Additional NetBackup options for AHV configuration](#)
- [Additional information about the rename file](#)

Using APIs and command line options to manage, protect, or recover AHV virtual machines

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

Following sections are part of this topic:

- See [the section called “Add an AHV cluster”](#) on page 95.
- See [the section called “Set iSCSI CHAP settings APIs”](#) on page 95.
- See [the section called “Create an AHV VM backup policy”](#) on page 96.
- See [the section called “Pre-Recovery check for AHV VM at the original location”](#) on page 97.
- See [the section called “Pre-Recovery check for AHV VM at a different location”](#) on page 98.
- See [the section called “Restore the AHV VM at the original location”](#) on page 98.
- See [the section called “Restore the AHV VM to an alternate location”](#) on page 100.

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 an AHV cluster

Table 8-1 Add an AHV cluster

API or command line options	Important variables and options
POST /netbackup/asset-service/queries GET /netbackup/asset-service/queries/{aqcId}	■ <code>clusterName</code> is the name of the AHV cluster. ■ <code>backuphost</code> is hostname of a NetBackup client. ■ <code>credentialName</code> are credentials associated with AHV cluster. Note: The credential must exist with <code>credentialName</code> mentioned.
tpconfig command	■ <code>virtual_machine</code> is the name of the AHV cluster. ■ <code>vm_type</code> is 9. The number 9 stands for AHV cluster.

Set iSCSI CHAP settings APIs

Table 8-2 Set iSCSI CHAP settings APIs

API or command line options	Important variables and options
GET /netbackup/config/iscsi-settings/{workloadType}	■ <code>workloadType</code> specify the supported workload. ■ Obtains the global iSCSI settings for the specified workload type.
POST /netbackup/config/iscsi-settings/{workloadType}	■ Changes the global iSCSI setting for the specified workload type. ■ <code>authType</code> is the authentication type. For example: ■ <code>ONEWAY_CHAP</code> ■ <code>MUTUAL_CHAP_AUTOMATIC</code> ■ <code>passwordRenewalIntervalDays</code> is applicable only for Mutual CHAP Automatic option. Note: Valid value is 1 - 365 days.

Create an AHV VM backup policy

Table 8-3 Create an AHV VM backup policy

API or command line options	Important variables and options
POST /netbackup/config/policies/	- policyType is Hypervisor. - backuphost is a hostname of a NetBackup client that performs backups on behalf of the virtual machines. - Add Add useVirtualMachine = 6 for Nutanix AHV. - snapshotMethodArgs can have the following values to back up a VM using VM UUID: - In backupSelections > selections, use the filter option as Nutanix-ahv:/?filter=uuid Equal <uuid_filter>" to filter AHV VMs of a specific UUID. Apart from UUID, you can use the other filter criteria mentioned for Intelligent VM groups.
admincmd command	- In bpplclients -add <discoveryhost> Hypervisor Hypervisor, the hypervisor discovery host is a allowedlisted Windows or Linux host. - In bpplinfo, the policy type (-pt) is Hypervisor. - In bpplinclude, use the filter option as Nutanix-ahv:/?filter=uuid Equal <uuid_filter>" to filter AHV VMs of a specific UUID. - In bpplinfo - Value of use_virtual_machine is 6 for AHV VMs. - Value of snapshot_method is Hypervisor_snap.

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*.

Pre-Recovery check for AHV VM at the original location

Table 8-4 Pre-Recovery check for AHV VM at the original location

API or command line options	Important variables and options
<code>POST /netbackup/recovery/workloads /nutanix-ahv/scenarios/full-vm /pre-recovery-check</code>	■ <code>client</code> is identifier that was used at the time of backup. It can either be the <code>displayName</code> or the <code>UUID</code>. ■ <code>ahvCluster</code> is the name of the alternate AHV cluster. ■ <code>recoveryHost</code> is server that is to be used as the VM recovery host to perform this pre-recovery check. ■ <code>vmDisks</code> represents one or more virtual machine disks. ■ <code>source</code> is the source path of the virtual machine disk. This must be of the <code>/storage_container/disk_uuid</code> format. ■ <code>destination</code> is the destination path of the virtual machine disk. This should be of the format <code>/storage_container</code>. ■ Set the following values: <code>powerOnAfterRecovery</code> <code>overwriteExistingVm</code> <code>removeNetworkInterfaces</code> <code>retainVmGuid</code> <code>retainNicMacAddress</code>

Pre-Recovery check for AHV VM at a different location

Table 8-5 Pre-Recovery check for AHV VM at a different location

API or command line options	Important variables and options
<pre>POST /netbackup/recovery/workloads /nutanix-ahv/scenarios/full-vm /pre-recovery-check</pre>	■ <code>client</code> is identifier that was used at the time of backup. It can either be the <code>displayName</code> or the <code>UUID</code>. ■ <code>ahvCluster</code> is the name of the alternate AHV cluster. ■ <code>recoveryHost</code> is server that is to be used as the VM recovery host to perform this pre-recovery check. ■ <code>vmDisks</code> represents one or more virtual machine disks. ■ <code>source</code> is the source path of the virtual machine disk. This must be of the <code>/storage_container/disk_uuid</code> format. ■ <code>destination</code> is the destination path of the virtual machine disk. This should be of the format <code>/storage_container</code>. ■ Set the following values: <pre>powerOnAfterRecovery overwriteExistingVm removeNetworkInterfaces retainVmGuid retainNicMacAddress</pre>

Restore the AHV VM at the original location

Table 8-6 Restore the AHV VM at the original location

API or command line options	Important variables and options
<pre>POST /netbackup/recovery/workloads/ahv/ scenarios/full-vm/recover</pre>	■ <code>client</code> is identifier that was used at the time of backup. It can either be the <code>display name</code> or the <code>UUID</code>. ■ <code>recoveryHost</code> is server that is to be used as the VM recovery host to perform this recovery. ■ Set the following values: <pre>powerOnAfterRecovery overwriteExistingVm removeNetworkInterfaces retainVmGuid retainNicMacAddress</pre>

Table 8-6

Restore the AHV VM at the original location *(continued)*

API or command line options	Important variables and options
bprestore command	■ <code>vmproxy</code> specifies the name or the FQDN of the backup host. ■ <code>vmserver</code> is the name of the AHV cluster. ■ <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> Note: For a Windows NetBackup host, you must add an empty line at the end of the rename file entries. See “Additional information about the rename file” on page 102.

Restore the AHV VM to an alternate location

Table 8-7 Restore the AHV VM to an alternate location

API or command line options	Important variables and options
POST /netbackup/recovery/workloads/ahv /scenarios/full-vm/recover	■ <code>client</code> is identifier that was used at the time of backup. It can either be the <code>displayName</code>) or the <code>UUID</code>. ■ <code>ahvCluster</code> is the name of the alternate AHV cluster. ■ <code>recoveryHost</code> is server that is to be used as the VM recovery host to perform this recovery. ■ <code>vmDisks</code> represents one or more virtual machine disks. ■ <code>source</code> is the source path of the virtual machine disk. This should be of the format <code>/storage_container/disk_uuid</code>. ■ <code>destination</code> is the destination path of the virtual machine disk. This should be of the format <code>/storage_container</code>. ■ Set the following values: powerOnAfterRecovery overwriteExistingVm removeNetworkInterfaces retainVmGuid retainNicMacAddress

Table 8-7 Restore the AHV VM to an alternate location (*continued*)

API or command line options	Important variables and options
bprestore command	■ <code>vmproxy</code> specifies the name or the FQDN of the backup host. ■ <code>vmserver</code> is the name of the AHV cluser. ■ Use the following values to modify the VM configuration: ■ <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> Note: For a Windows NetBackup host, you must add an empty line at the end of the rename file entries. See “Additional information about the rename file” on page 102.

Additional NetBackup options for AHV configuration

Use the following NetBackup command options for additional AHV configuration:

`NUTANIX_AUTODISCOVERY_INTERVAL` option for NetBackup servers. This option controls how often NetBackup scans the AHV clusters to discover virtual machines to display in the NetBackup web UI.

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

1. NetBackup primary server

2. Access host, client, or proxy server
3. Media server

Table 8-8

Usage	Important variables and options
POST /netbackup/asset-service/queries	■ <code>clusterName</code> is the name of the AHV cluster. ■ <code>backuphost</code> is hostname of a NetBackup client. ■ <code>credentialName</code> are credentials associated with AHV cluster.
GET /netbackup/asset-service/queries/{agcId}	
tpconfig command	■ <code>virtual_machine</code> is the name of the AHV cluster. ■ <code>vm_type</code> is 9. The number 9 stands for AHV cluster.

Additional information about the rename file

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

Create or modify the rename file in the `/usr/openv/tmp` directory for the following scenarios:

- Recover the VM to an alternate container
- Recover the VM to the same or an alternate container with a modified VM name

If the rename file is not available, then you must create it and save it as `rename.txt` on the NetBackup primary server.

To set the alternate location or modify the configuration, add the following lines in the rename file in the given format:

Scenario	Line to add in the rename file
Change Virtual Machine Name	<code>change vmname to newVMname</code>
Recover the virtual machine to a different AHV container	<code>change /<original_container1>/<disk_uuid1> to /<alternate_container1></code>

Sample rename file

The following `rename.txt` lets you change the VM name.

```
change vmname to newVMname
```

After making the required changes in the rename file, you can run the `bprestore` command.