

NetBackup™ for Nutanix AHV Administrator's Guide

Release 11.2.0.1

Last updated: 2026-09-08

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. COHESITY 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.

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to [Cohesity Support](#), to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the [Cohesity Support Portal](#) to create a new case.
- Click the (?) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the [Cohesity Knowledge Base](#).
- Log in to the [Cohesity Support Portal](#) to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1	Overview	8
	Overview of configuring and protecting AHV assets in the NetBackup web UI	8
Chapter 2	RBAC roles for the Nutanix AHV administrator	10
	RBAC roles for the Nutanix AHV administrator	10
	Assign both the Default VMware Administrator and Default AHV Administrator roles to a user	11
	Create a custom role for all Nutanix AHV permissions and additional VMware asset permissions	12
	Create a custom role for all VMware permissions and additional Nutanix AHV asset permissions	13
Chapter 3	Managing AHV clusters	15
	Quick configuration checklist to protect AHV virtual machines	16
	Configure secure communication between the AHV cluster and NetBackup host and Nutanix Prism Central and NetBackup host	20
	Enable the iSCSI initiator service on windows backup host	23
	Install the iSCSI initiator package on Linux backup host	23
	Migrate Java GUI/CLI added clusters into the web UI	24
	Prerequisites to configure Nutanix AHV cluster	25
	About support for Nutanix segmented iSCSI network	25
	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
	Remove AHV Clusters	32
	Add a new Nutanix Prism Central	32
	Add new Prism Central server credentials	34
	Remove Nutanix Prism Central	34
	Create an intelligent VM group	35
	Assign permissions to the intelligent VM group	39

	Update the intelligent VM group	40
	Remove the intelligent VM group	40
	Set CHAP for iSCSI	40
	Add an AHV access host	41
	Remove an AHV access host	42
	Change resource limits for AHV resource types	42
	Change the autodiscovery frequency of AHV assets	46
	Scan for malware	46
	Scanning backup images	46
	Assets by workload type	49
Chapter 4	Managing credentials	51
	Managing AHV cluster credentials	51
	Add new cluster credentials	51
	Update and validate AHV cluster credentials	52
	Managing Nutanix Prism Central credentials	52
	Add new Nutanix Prism Central credentials	53
	Update and validate Nutanix Prism Central credentials	53
	View the credential name that is applied to an asset	54
	Edit or delete a named credential	54
Chapter 5	Instant access	56
	Prerequisites of instant access	56
	Things to consider and limitations before you use the instant access feature	56
	Create an instant access VM	58
	Download files and folders from a VM backup image	60
	Instant access Build Your Own (BYO)	61
	Prerequisites of Instant Access Build Your Own (BYO)	61
	Hardware configuration requirement of Instant Access Build Your Own (BYO)	62
	Frequently asked questions	62
Chapter 6	Protecting AHV virtual machines	66
	Things to know before you protect AHV virtual machines	66
	Protect AHV VMs or intelligent VM groups using Protection plan	67
	Backup AHV VMs or intelligent groups using Policy	68
	Protect AHV VMs within VPC	69
	Protect vTPM enabled AHV VMs	70
	Customize protection settings for an AHV asset	70
	Modify policy for an AHV asset	71

	Schedules and retention	71
	Backup options	71
	Prerequisite to Enable virtual machine quiescing	72
	Remove protection from VMs or intelligent VM groups	73
	View the protection status of VMs or intelligent VM groups	73
Chapter 7	Recovering AHV virtual machines	75
	Things to consider before you recover the AHV virtual machines	76
	About the pre-recovery check	76
	Recover an AHV virtual machine	76
	Recover an AHV VM within VPC	78
	Recover a vTPM enabled AHV VM	79
	About Nutanix AHV agentless files and folders restore	79
	Prerequisites for agentless files and folder recovery	81
	SSH key fingerprint	92
	Recover files and folders with Nutanix AHV agentless restore	93
	Recovery target options	94
	Pre-recovery checks for Nutanix AHV	99
	About Nutanix-AHV agent-based files and folders restore	101
	Prerequisites for agent-based files and folder recovery	101
	Recover files and folders with Nutanix AHV agent based restore	103
	Limitations	104
Chapter 8	Protecting Nutanix Cloud Clusters (NC2)	107
	Protecting Nutanix Cloud Clusters (NC2) on AWS	107
	Protecting Nutanix Cloud Clusters (NC2) on Azure	108
Chapter 9	Troubleshooting AHV operations	109
	Troubleshooting AHV operations: Error during creation of AHV instant access virtual machines	109
	Troubleshooting tips for NetBackup for AHV	111
	Error during AHV credential addition	112
	Error during the AHV virtual machines discovery phase	112
	Errors for the Status for a newly discovered VM	113
	Error run into while backing up AHV virtual machines	114
	Error while restoring AHV virtual machines	122
Chapter 10	API and command line options for AHV	134
	Using APIs and command line options to manage, protect, or recover AHV virtual machines	134
	Additional NetBackup options for AHV configuration	141

Additional information about the rename file	142
--	-----

Overview

This chapter includes the following topics:

- [Overview of configuring and protecting AHV assets in the NetBackup web UI](#)

Overview of configuring and protecting AHV assets in the NetBackup web UI

Table 1-1 Steps to configure and protect AHV assets

Step	Action	Description
Step 1	Sign in to NetBackup web UI as the Default Security Administrator. Then add the AHV user to the Default AHV Administrator role. Or, create a custom RBAC role to fit the permissions that the AHV administrator needs.	Note: To perform the AHV administrator tasks, the Default AHV Administrator role has the minimum required permissions. You can also create a custom role to give different access permissions to your AHV administrators.

Table 1-1 Steps to configure and protect AHV assets (*continued*)

Step	Action	Description
	Configure the following for an AHV cluster: ■ Configure secure communication between the AHV cluster and NetBackup host. ■ (Optional) Configure secure communication between Nutanix Prism Central and NetBackup host. ■ Enable iSCSI on the NetBackup host which you want to use as a backup or restore host. ■ In the Nutanix Prism console, add the backup host to the allowed list. 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.	See “Configure secure communication between the AHV cluster and NetBackup host and Nutanix Prism Central and NetBackup host” on page 20. See “Install the iSCSI initiator package on Linux backup host” on page 23. See “Enable the iSCSI initiator service on windows backup host” on page 23.
Step 3 (Optional)	Configure and Manage Nutanix Prism Central.	See “Add a new Nutanix Prism Central” on page 32.
Step 4	Configure and manage AHV cluster.	See “Prerequisites to configure Nutanix AHV cluster” on page 25.
Step 5	Add and manage credentials.	See “Add new cluster credentials” on page 51.
Step 6	Configure an AHV protection plan.	See the <i>NetBackup™ Web UI Administrator's Guide</i> .
Step 7	Configure an Intelligent VM group.	See “Create an intelligent VM group” on page 35.
Step 8	Protect AHV VMs or intelligent VM groups.	See “Protect AHV VMs or intelligent VM groups using Protection plan” on page 67.
Step 9	Recover a VM.	See “Recover an AHV virtual machine” on page 76.

RBAC roles for the Nutanix AHV administrator

This chapter includes the following topics:

- [RBAC roles for the Nutanix AHV administrator](#)
- [Assign both the Default VMware Administrator and Default AHV Administrator roles to a user](#)
- [Create a custom role for all Nutanix AHV permissions and additional VMware asset permissions](#)
- [Create a custom role for all VMware permissions and additional Nutanix AHV asset permissions](#)

RBAC roles for the Nutanix AHV administrator

NetBackup enables control over which users can access which Nutanix AHV or other workloads using Role Based Access Control (RBAC). Depending on your environment, you may want to configure RBAC for Nutanix AHV administrators in the following ways.

Table 2-1 RBAC roles to use or to create for the Nutanix AHV administrator

Type of access needed	Roles to use or create	
Manage Nutanix AHV assets and configuration and perform restores to Nutanix AHV.	Default AHV Administrator role	

Table 2-1

RBAC roles to use or to create for the Nutanix AHV administrator
(continued)

Type of access needed	Roles to use or create	
Manage Nutanix AHV assets and VMware assets and have cross-hypervisor support.	Default AHV Administrator and Default VMware Administrator roles	With these roles the administrator can also manage the following: Credentials for Nutanix AHV. (Prism Central servers tab or the AHV cluster tab in Workloads > Nutanix AHV .) Credentials for a vCenter, ESX server, etc. (VMware servers tab in Workloads > VMware).
Full access to Nutanix AHV with additional permissions for VMware assets.	Create a custom role.	This role is useful for scenarios when you want to restore from VMware backup images and want to use the AHV workload after recovery.
Full access to VMware with additional permissions for Nutanix AHV assets.	Create a custom role.	This role is useful for scenarios when you use the VMware workload and want to perform cross-hypervisor restores.

Note the following:

- To create an RBAC role, you must have the RBAC Administrator role or the permissions to create roles.
- To create a credential, you must have the RBAC Administrator role or a role that has permissions to create credentials. The **Default Nutanix AHV Administrator** and the **Default VMware Administrator** role can assign a credential to a user, but cannot create a credential in credential management.
- Contact your NetBackup administrator for assistance with creating roles and credentials.

Assign both the Default VMware Administrator and Default AHV Administrator roles to a user

Follow this procedure when you want to grant RBAC access globally to a user for all Nutanix AHV assets and for all VMware assets. This role is useful for scenarios when you use the VMware workload and want to perform cross-hypervisor restores. The benefit with this role is that it is not necessary to manually select the required

permissions for the source and the destination assets. However, it is not possible to restrict the permissions for specific assets with this role. In that case, you must configure a custom role.

To assign both the Default VMware Administrator and Default AHV Administrator roles to a user

- 1 On the left, select **Security > RBAC**.
- 2 Select the **Default VMware Administrator** role. Then select the **Users** tab.
- 3 Enter the group name or the username. Then select **Add to list**.
- 4 To return to the roles list, select the **Close** button.
- 5 Select the **Default AHV Administrator** role. Then select the **Users** tab.
- 6 Enter the group name or the username. Then select **Add to list**.
- 7 To view the roles that are assigned to the group or the user, select the **Users** tab.

Create a custom role for all Nutanix AHV permissions and additional VMware asset permissions

Follow this procedure when you want to grant RBAC global access to a user for Nutanix AHV and permissions for specific VMware assets. This role is useful for scenarios when you want to restore from VMware backup images and want to use the AHV workload after recovery.

To create a custom role for all Nutanix AHV permissions and additional VMware asset permissions

- 1 On the left, select **Security > RBAC** and select **Add**.
- 2 Select the **Default AHV Administrator** role. Then select **Next**.
- 3 Provide a **Role name** and a description.
For example, include a description that the role allows users to manage Nutanix AHV and provides specific permissions for VMware.
- 4 Under **Permissions**, select **Edit**.
- 5 Go to **AHV assets**. Note that all permissions are selected already for that workload.
- 6 Go to **VMware assets**.
- 7 Select the following permissions:

- **View**
 - **Manage access**
 - **Restore**
 - **View jobs**
- 8 Select **Assign**.
 - 9 Under **Workloads**, select **Edit**.
 - 10 Choose from the following options:
 - To apply the wanted permissions for the workload to all existing and to future VMware assets, leave the following option enabled: **Apply permissions to all existing and future VMware assets**.
 - To apply the wanted permissions only to specific VMware assets, clear the option: **Apply permissions to all existing and future VMware assets**. Then select the assets to which you want to apply permissions and select **Add**.
 - 11 Select **Assign**.
 - 12 Under **Users**, select **Edit**. Then add the group or users that you want to have this RBAC role.
 - 13 Select **Assign**.
 - 14 When you are done configuring the role, select **Add role**.
 - 15 To view the roles that are assigned to the group or the user, select the **Users** tab.

Create a custom role for all VMware permissions and additional Nutanix AHV asset permissions

Follow this procedure when you want to grant RBAC global access to a user for VMware and permissions for specific Nutanix AHV assets.

To create a custom role for all VMware permissions and additional Nutanix AHV asset permissions

- 1 On the left, select **Security > RBAC** and select **Add**.
- 2 Select the **Default VMware Administrator** role. Then select **Next**.
- 3 Provide a **Role name** and a description.

For example, include a description that the role allows users to manage VMware and provides specific permissions for Nutanix AHV.

- 4 Under **Permissions**, select **Edit**.
- 5 Go to **VMware assets**. Note that all permissions are selected already for that workload.
- 6 Go to **AHV assets > AHV clusters, VMs, and storage containers**.
- 7 Select all permissions except for **Granular restore**.
- 8 Go to **AHV assets > Prism Central**.
- 9 Select all the permissions in that group.
- 10 Do select any permissions for the group **AHV intelligent VM groups**.
- 11 Select **Assign**.
- 12 Under **Workloads**, select **Edit**.
- 13 Choose from the following options:
 - To apply the wanted permissions for the workload to all existing and to future Nutanix AHV assets, leave the following option enabled: **Apply permissions to all existing and future AHV assets**.
 - To apply the wanted permissions only to specific VMware assets, clear the option: **Apply permissions to all existing and future AHV assets**. Then select the assets to which you want to apply permissions and select **Add**.
- 14 Select **Assign**.
- 15 Under **Users**, select **Edit**. Then add the group or users that you want to have this RBAC role.
- 16 Select **Assign**.
- 17 When you are done configuring the role, select **Add role**.
- 18 To view the roles that are assigned to the group or the user, select the **Users** tab.

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 and Nutanix Prism Central 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 the web UI](#)
- [Prerequisites to configure Nutanix AHV cluster](#)
- [About support for Nutanix segmented iSCSI network](#)
- [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](#)
- [Remove AHV Clusters](#)
- [Add a new Nutanix Prism Central](#)
- [Add new Prism Central server credentials](#)
- [Remove Nutanix Prism Central](#)
- [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](#)
- [Scan for malware](#)

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 134.

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

Table 3-1 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 3-1

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 41.
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 and Nutanix Prism Central and NetBackup host” on page 20. - Communication ports See “About the ports that NetBackup uses to communicate with AHV” on page 27.
Manage AHV clusters, Prism Central server, and intelligent VM groups	- Managing AHV clusters See “Add or browse an AHV cluster” on page 28. - Managing Prism Central server See “Add a new Nutanix Prism Central” on page 32. - Managing intelligent VM groups See “Create an intelligent VM group” on page 35. See “Remove the intelligent VM group” on page 40.
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 66. - Protecting virtual machines See “Protect AHV VMs or intelligent VM groups using Protection plan” on page 67.

Table 3-1

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 23. 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 23. 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 3-1

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 42.

Table 3-1

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 - In the NetBackup web UI, click Storage > Disk storage. Then click the Storage servers tab. Add all the supported media servers for load balancing. - Click Storage > Storage units. Select the <i>storage unit name</i>. Under the Media server click Edit. Then 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 and Nutanix Prism Central and NetBackup host

NetBackup can now validate AHV cluster and Prism Central server 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 and AHV Prism Central Server 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.

For Nutanix Prism Central use the `openssl s_client -connect Nutanix Prism Central FQDN:9440 -showcerts < /dev/null`

- 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 the information to 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_path\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_path\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](#).

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

Table 3-2

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. If you have already configured this external CA option, append the Nutanix AHV CA certificates to the existing external certificate trust store. If you have not configured the option, add all the required Nutanix AHV server CA certificates to the trust store and set the option.
ECA_CRL_PATH	Specifies the path to the directory where the certificate revocation lists (CRL) of the external CA are located. If you have already configured this external CA option, append the AHV CRLs to the CRL cache. If you have not configured the option, first add all the required CRLs to the CRL cache. Then set the option.
VIRTUALIZATION_HOSTS_SECURE_CONNECT_ENABLED	This option affects AHV, RHV, and VMware secure communication. Without this option, the secure or insurce communication with workload is decided by each workload and plug-in separately. For Nutanix AHV, secure communication is enabled by default. This option lets you skip the security certificate validation. Disabling this option lets you skip the security certificate validation. Cohesity recommends that you enable secure communication using the <code>ECA_TRUST_STORE_PATH</code> option.

Table 3-2 (continued)

VIRTUALIZATION_CRL_CHECK	Lets you validate the revocation status of the virtualization server certificate against the CRLs. By default, the option is enabled.
--------------------------	---

Enable the iSCSI initiator service on windows backup host

Do one of the following:

- 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.
- 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.

- If you plan to use Nutanix iSCSI segmented network, See [“About support for Nutanix segmented iSCSI network”](#) on page 25. for the backup host network configuration details.

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.

- If you plan to use Nutanix iSCSI segmented network, See [“About support for Nutanix segmented iSCSI network”](#) on page 25. for the backup host network configuration details.

Migrate Java GUI/CLI added clusters into the web UI

Credential management for Java GUI/CLI and web UI are separate.

Warning: The support for Administration Console is discontinued.

- Clusters that are added with the Administration Console or CLIs are not reflected in web UI.
- If there are any existing clusters in Java GUI/ CLI, user must manually add these cluster/clusters and its credentials in web UI.

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

- Once the cluster is added in web UI and if the cluster credential needs to be updated, it must be updated from web UI only.
Consider the following scenario:
 - A cluster exists in both the web UI and the Administration Console.
 - Cluster credential is updated in web UI only.
 - The cluster is deleted from web UI.

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 the web UI, 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 into the web UI and then it is deleted from the Java GUI, the cluster can still be seen in the web UI.
- If a cluster was in the web UI and the Java GUI and its credential was updated in the web UI and then the cluster is deleted from the web UI, backup and

restores may fail as the cluster added in Java UI was not updated. You may have to update the credential from Java UI to get things working.

Prerequisites to configure Nutanix AHV cluster

Prerequisite:

Configure iSCSI data services IP on the Nutanix AHV cluster

- 1 To use **Use segmented iSCSI data service IP** or **Use specified segmented iSCSI data services IP** option, AHV cluster must be configured with segmented iSCSI network interface with volumes(ABS) feature.
- 2 If you plan to choose option **Use iSCSI data services IP** while configuring the cluster, **Nutanix** recommends, 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.

For Linux backup host, jobs fall back to use NFS provided the segmented iSCSI data services IP configuration is correctly done.

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.

About support for Nutanix segmented iSCSI network

NetBackup support separating backup traffic with a use of Nutanix iSCSI segmented network. Separating backup traffic is helpful to reduce load on production resources, dedicating right size resources for better backup-recovery speed and security. By default, backup-recovery traffic flows over Nutanix cluster management network using iSCSI data services IP for initial connection and discovery.

During AHV cluster configuration, choose one of the following options for iSCSI transport:

- Use iSCSI data service IP
- Use segmented iSCSI data service IP
- Use specified segmented iSCSI data service IP

For more details See [“Add or browse an AHV cluster”](#) on page 28.

Backup host network configuration to use Nutanix iSCSI segmented network configuration.

- The iSCSI segmented network is on different subnet than the cluster management network, so backup host network must be configured to connect to:

- AHV cluster-management network
- iSCSI-segmented network.

To achieve this, configure the backup host with two VLANs. One corresponding to cluster management network and other corresponding to segmented iSCSI network planned to use for backup-recovery traffic.

- For better performance, use hostname/IP corresponding to segmented network as hostname while installing/configuring NetBackup on the backup host.
- Validate connectivity using the following command on Windows host:
 - Click **Server Manager > Tools > iSCSI initiator**. This opens **iSCSI initiator properties** dialog.
 - Click **Discovery > Discovery portal** and provide the IP address as per configured iSCSI Target Type for the AHV cluster.
 - DEFAULT: iSCSI Data services IP from cluster details page
 - SEGMENTED: segmented iSCSI Data from cluster details page
 - SEGMENTED_SPECIFIC: Virtual IP specified while configuring cluster in NetBackup.
- Validate connectivity using following command on linux host:
 - `iscsiadm -m discovery -t sendtargets -p correct IP as per configured iSCSI targetType`
 - DEFAULT: iSCSI Data services IP from cluster details page
 - SEGMENTED: segmented iSCSI Data from cluster details page
 - SEGMENTED_SPECIFIC: Virtual IP specified while configuring cluster in NetBackup. If there is connectivity problem it shows an error, for example:
`iscsiadm: connect to <IP> timed out`

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 3-3 Ports required by NetBackup to communicate with AHV

Port	Protocol	Destination	Purpose
860, 3260	iSCSI over TCP	*bi-directional	iSCSI provides block-level access to storage devices with the <code>SCSI</code> . iSCSI facilitates data transfers usually over the ethernet.
3205	iSCSI over TCP	*bi-directional	iSNS is able to emulate fibre channel fabric services and manage both iSCSI and fibre channel devices, an iSNS server can be used as a consolidated configuration point for an entire storage network
111	TCP	*bi-directional	Portmapper
2049	TCP	*bi-directional	NFS
9440	TCP	AHV Cluster AHV Prism Central Server	Prism console, REST API

Table 3-3 Ports required by NetBackup to communicate with AHV
(continued)

Port	Protocol	Destination	Purpose
------	----------	-------------	---------

*Ports must be open bi-directional between AHV access host and AHV cluster. Port 9440 is only open inbound to the AHV cluster from the AHV access host.

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 **Workloads > 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 112.

■ **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 the **Use Prism Central for this cluster** check-box to protect Prism Central server-related attributes of virtual machines. For example to capture virtual private cloud network-related attributes, project, category, and owner-related attributes of VM.
See [“Add a new Nutanix Prism Central”](#) on page 32.

Note: Prism Central server must be added in the NetBackup environment, before you select this check-box.

- Select one of the following from **iSCSI transport**.
- **Use iSCSI data service IP**
Uses iSCSI data services IP configured on AHV cluster as an iSCSI target discovery portal and initial connection point.

Note: This option falls back to NFS on a Linux backup host if:
The iSCSI data services IP is not configured on the AHV cluster.
The iSCSI connection is not established on the backup host.

- **Use segmented iSCSI data service IP**

Uses segmented iSCSI data services IP configured on AHV cluster as an iSCSI target discovery portal and initial connection point.

Note: Cluster Validation fails if the configuration is not on an AHV cluster.
Backup/Recovery job fails if the configuration is not on the AHV cluster or the backup host does not have the required network configuration.

- **Use specified segmented iSCSI data service IP**

- In the **Virtual IP address** field, provide the valid IP address.
Provide the virtual IP corresponding to Nutanix segmented iSCSI network interface which you plan to use for backup and recovery iSCSI data traffic.

Uses specified IP address as an iSCSI target discovery portal and initial connection point. NetBackup validates using Nutanix APIs, if the virtual IP address is from any one of the configured segmented iSCSI data services interfaces.

Note: Backup/Recovery job fails if the configuration is not done on the AHV cluster or the backup host do not have the required network configuration.

3 ■ **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 51.

Note: You must associate the credentials of an AHV cluster user who has a Cluster admin role.

4 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](#).

5 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 the selected cluster.

Bulk actions on AHV cluster

You can select one or more AHV clusters 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:**
 - Validates the credentials of the AHV cluster.
 - If you choose **Use segmented iSCSI data service IP**, validates if segmented iSCSI Data service IP address configured on the Nutanix cluster.
 - If you select **Use specified segmented iSCSI data service IP**, validates if the specified virtual IP address is configured as a segmented iSCSI Data service IP address on the Nutanix 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 **Workloads > 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 an AHV cluster to view details.
- 4 Click a virtual machine to view its protection status, recovery points, and restore activity.
- 5 Click **Add protection** to subscribe the selected VMs to a protection plan. You can also select **Backup now**, **Recover**, and **Manage Permission** options.

Note: VMs can be protected using Nutanix-AHV policies in NetBackup web UI 11.2.0.1.

- 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 web UI displays an empty field and the corresponding API shows the -ve value for the free space field for a particular storage container.

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

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

Remove AHV Clusters

Use this procedure to delete AHV clusters.

To remove an AHV cluster

- 1 On the left, click **Workloads > 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**.

Add a new Nutanix Prism Central

You can add and browse Nutanix Prism Central and their credentials.

To add Nutanix Prism Central and respective credentials

1 On the left, click **Workloads > Nutanix AHV** then click the **Prism Central servers** tab.

2 Click **Add** to add an Nutanix Prism Central and enter the following:

- **Prism central server name**
- **REST API port (default: 9440)**
This port must remain open between backup host and AHV cluster.
- **Backup host**

Note: Backup host version must be NetBackup 10.1.1 or later. And operating system must be either Linux (RHEL and SUSE) or Windows.

- **Associate credential**

Do one of the following:

- When adding prism central server credentials for an existing credential, select the category as **AHV Prism Central**. For more information, see *Managing Credentials* in the [NetBackup™ Web UI Administrator's Guide](#).
- See [“Add new Prism Central server credentials”](#) on page 34.

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 Prism Central credentials, click **Add**.

Inline actions on Nutanix Prism Central

You can run the following inline actions on an Nutanix Prism Central:

- **Validate:** Manually validates
- **Edit:** Modify the backup host and Nutanix Prism Central credentials.
- **Delete:** Removes the Nutanix Prism Central.
- **Manage Permissions:** Used to add or manage the permissions on selected Prism Central.

Bulk actions on Nutanix Prism Central

You can select one or more Nutanix Prism Central and run the following bulk actions:

- **Validate**
- **Delete**

Add new Prism Central server credentials

- 1 On left, click **Workloads > Nutanix AHV**, then click the **Nutanix Prism Central** tab.
- 2 Click **Add** to add a new Prism Central server.
- 3 On the **Add AHV Prism Central > Associate credentials** page, click **Add a new credential**.
- 4 Enter the details such as **Credential name**, **Tag**, and **Description**.
- 5 In the **Credentials for the Nutanix Prism Central** part, add the **Username**, **Password**, and **Domain** of the associated Prism Central server.

Note: The associated credential must be of a user who has Prism Central Admin role.

- 6 Click **Next**.
Either select the existing role or add a new role to provide permissions for the credential.
- 7 Click **Save**.

Remove Nutanix Prism Central

Use this procedure to delete one or more Nutanix Prism Centrals.

To remove an Nutanix Prism Central

- 1 On the left, click **Workloads > Nutanix AHV**, then click the **Prism Central servers** tab.
The tab lists the names of Nutanix Prism Centrals that you have access to.
- 2 Locate and select the AHV Prism Central.

- 3 Select one or more prism central servers and click **Actions > Delete**

Note: If you delete a prism central, all virtual machines associated with the deleted prism central server will be backed-up/recovered without project, category, owner and virtual private cloud networking related attributes.

- 4 Unselect the **Disable "Use Prism Central server for this cluster"** for all clusters associated with this Prism Central servers. De-Select if you want to keep it enabled. checkbox, if required and click **Delete**

Note: Once the Nutanix Prism Central is deleted, asset discovery will not be triggered automatically for the clusters of this Prism Central server. Hence, the VM's of these clusters will display the Prism Central server and project on the VM detail page until the next asset discovery is triggered.

Note: If an environment has clusters associated with this prism central with **Use Prism Central server for this cluster** checkbox selected, and the prism central is deleted by unselecting the **Disable "Use Prism Central server for this cluster" for all clusters associated with this Prism Central server** checkbox, then the subsequent backup or restore jobs would fail until the associated prism central server is added.

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 **Workloads > 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:

[+ Condition](#) [+ Sub-query](#)

VM display name	▼	Contains	▼	windows	🗑
-----------------	---	----------	---	---------	---

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:

AND OR

[+ Condition](#) [+ Sub-query](#)

VM display name	▼	Contains	▼	windows	🗑
Power state	▼	=	▼	ON	🗑

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**:

The screenshot shows a query builder interface with two conditions combined using the **OR** operator. The first condition is 'VM display name' contains 'windows'. The second condition is 'Power state' is 'ON'. The interface includes dropdown menus for the field and operator, and a text input for the value. There are also buttons for '+ Condition' and '+ Sub-query'.

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 using Protection plan”](#) on page 67.

- Edit or update the intelligent VM group:
See [“Update the intelligent VM group”](#) on page 40.
- Assign permissions to the VM group:
See [“Assign permissions to the intelligent VM group”](#) on page 39.

Query options for creating intelligent VM groups

Table 3-4 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.
<code>vmUuid</code>	The VM's instance UUID. For example: 501b13c3-52de-9a06-cd9a-ecb23aa975d1 Not case-sensitive when the protection plan runs.
<code>StorageContainerName</code>	The name of the storage container. Case-sensitive when the protection plan runs.
<code>Category</code>	Ensure the following: AHV category is applied on VMs in Nutanix Prism Central server. For full search it must be in the CategoryName:Value format.

Table 3-5 Query operators

Operator	Description
<code>Starts with</code>	Matches the value when it occurs at the start of a string. For example: If the value you enter is <code>box</code> , this option matches the string <code>box_car</code> but not <code>flatbox</code> .

Table 3-5 Query operators (*continued*)

Operator	Description
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 **Workloads > 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 35.

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 **Workloads > 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, click **Workloads > 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, click **Workloads > 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 **Workloads > 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 **Workloads > 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.

Scan for malware

NetBackup version 10.5.1 and later provides support for scanning Nutanix assets for malware through the Nutanix AHV workload.

For triggering malware scan, the scan host must be configured. For more information on configuring the scan host, refer to the 'Scan host configurations' chapter in *NetBackup Security and Encryption Guide*.

Scanning backup images

This section describes the procedure for scanning client backup images of a particular policy for malware.

To scan policy of client backup images for malware

- 1 On the left, select **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Scan for malware**.
- 3 In the **Search by** option, select **Backup images**.

Select one of the following scan types:

- **Malware scan** - Select this option to scan images using default malware scan.
- **YARA scan** - Select this option to scan images using YARA rules. Click the **Select threat feeds** option.

On the **Select threat feeds for scanning** dialog box, select the required YARA rules or a zip file of YARA rules that you have uploaded earlier.

For more information, see *NetBackup Security and Encryption Guide*.

All the following steps are applicable for the scan type: Malware scan.

- 4 In the **Scanner host pool** option, search and select the appropriate host pool name from the list of scanner host pools listed in **Select malware scanner host pool**.

Note: Scan host from the selected scan host pool must be able to access the instant access mount created on the storage server which is configured with NFS/SMB share type.

- 5 In the search criteria, review and edit the following:

- **Policy name**
Only supported policy types are listed.
- **Client name**
Displays the clients that have backup images for a supported policy type.
- **Policy type**
Displays all the supported policies which are enabled for malware scanning.

Note: The Nutanix-AHV policy displays Nutanix-AHV images if the backups are taken from a Nutanix-AHV policy and Protection plan backups.

Warning: The **Hypervisor** policy type displays Nutanix AHV and RHV images. NetBackup supports malware scanning only for Nutanix AHV images.

- **Type of backup**
- **Copies**
If the selected copy does not support instant access, then the backup image is skipped for the malware scan.
- **Disk pool**
MSDP (PureDisk), OST (for example, Data Domain) and AdvancedDisk storage type disk pools are listed.
- **Disk type**

MSDP (PureDisk), OST (for example, Data Domain) and AdvancedDisk disk types are listed.

- **Infection status**

The malware infected status of the backup images can be searched based on the following types: infection detected by malware scan, file hash search, not infected, not scanned or all.

- For the **Select the timeframe of backups**, verify the date and the time range or update it.
- On selecting the **Abort malware scan on detecting an infection** option, clean recovery would not be supported for infected images.

6 Click **Search**.

7 Select the search criteria and ensure that the selected scan host is active and available.

8 From the **Select the backups to scan** table select one or more images for scan.

9 Click **Scan for malware**.

10 After the scan is initiated, the **Scan status** is displayed.

The following are the status fields:

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

Hover over the status to view the reason for the failed scan.

Note: Any backup images that fail the validation are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability and for the supported policy types only. For backup images stored on OST targets, malware scanning is supported only if media server with Instant Access capability is also present.

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

Note: You can cancel the malware scan for one or more in progress and pending jobs.

- **Infected - Malware scan aborted**

Assets by workload type

Note: For YARA scanning, only Kubernetes is supported.

To scan the supported assets for malware, perform the following:

- 1** On the left, select the supported workload under **Workloads**.
- 2** Select the resource which has backups completed.
For example, Nutanix AHV.
- 3** Select **Actions > Scan for malware**.
- 4** On the **Malware scan** page, perform the following:
 - Select the date range for the scan by selecting **Start date/time** and **End date/time**.
 - Select **Scanner host pool**.
 - From the **Current infection status** list select one of the following:
 - **Not scanned**
 - **Not infected**
 - **Infection detected by malware scan**
 - **Infection detected by file hash search**
 - **All**
- 5** Click **Scan for malware**.

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

- 6** After the scan starts, you can see the **Scan status** on **Malware detection**, the following fields are visible:
 - **Not scanned**
 - **Not infected**
 - **Infected**
 - **Failed**

Note: Any backup images that fail validation are ignored.

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

Managing credentials

This chapter includes the following topics:

- [Managing AHV cluster credentials](#)
- [Managing Nutanix Prism Central credentials](#)
- [View the credential name that is applied to an asset](#)
- [Edit or delete a named credential](#)

Managing AHV cluster credentials

This section describes the procedures for adding, updating and validating the AHV cluster credentials.

Add new cluster credentials

- 1 On left, click **Workloads > 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 **Workloads > 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 **Workloads > 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.

Managing Nutanix Prism Central credentials

This section describes the procedures for adding, updating and validating the Nutanix Prism Central credentials.

Add new Nutanix Prism Central credentials

- 1 On left, click **Workloads > Nutanix AHV**, then click the **Prism Central servers** tab.
- 2 Click **Add** to add a new prism central server.
- 3 On the **Add AHV Prism Central server > 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 Nutanix Prism Central credentials

To validate Prism Central server credentials

- 1 On the left, click **Workloads > Nutanix AHV**, then click the **Prism Central servers** tab.
- 2
 - To validate specific Prism Central servers credentials, locate and select the Prism Central server. 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 Prism Central server. Then click **Validate** from the top bar.

Note: NetBackup verifies the current credentials for the selected Prism Central server.

If the credentials are not valid, NetBackup indicates **Invalid** under **Credentials**. Use the following steps to update the Prism Central server credentials.

To update Prism Central server credentials

- 1 On the left, click **Workloads > Nutanix AHV**, then click the **Prism Central servers** tab.
- 2 Locate and select the Prism Central server.

- 3 Select **Actions > Edit**.
- 4 Update the credentials as needed.

Note: Adding or updating Prism Central server credentials also automatically starts the discovery of the Prism Central server. When backup host information is provided in the request, it is used to perform validation of credentials as well as for performing the discovery. For discovery, NetBackup 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 Prism Central server.

View the credential name that is applied to an asset

You can view the named credential that is configured for an asset type. If the credentials are not configured for a particular asset, this field is blank.

To view credentials for Nutanix AHV clusters

- 1 On the left, select **Workloads > Nutanix AHV**.
- 2 On the **AHV clusters** tab, locate the **Credential name** column.

Edit or delete a named credential

You can edit the properties for a named credential or delete a named credential NetBackup from the **Credential management**.

Edit a named credential

You can edit a named credential to change the following: credential tag, description, category, authentication details, or permissions. You cannot change the credential name.

Note: Ensure that the credential category that is used for **AHV cluster** is *AHV* and for **Nutanix Prism Central** is *Prism Central*.

To edit a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to edit.
- 3 Select **Edit** and update the credential as needed.
- 4 Review the changes and select **Finish**.

Delete a named credential

You can delete a named credential that you no longer need to use with NetBackup. Be sure to apply another credential to any assets that use the credential you want to delete. Otherwise, backups and restores may fail for those assets.

To delete a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to delete.
- 3 Select **Delete > Delete**.

Instant access

This chapter includes the following topics:

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

Prerequisites of instant access

If you are using instant access, ensure that the WORM instance can access the following port on AHV Cluster Server:

Table 5-1 Port details

Instance	AHV component	Port number
WORM	AHV Cluster Server	9440

Things to consider and limitations before you use the instant access feature

Note the following about the Instant access virtual machines feature:

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

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

- This feature is supported with backup copies that are created from protection plans or policies.
- The download files feature is supported for NetBackup Appliance, NetBackup Virtual Appliance, Flex Appliance, and Build Your Own (BYO) server.

Instant access on Flex WORM storage requires the following services:

- NGINX, NFS, SAMBA, WINBIND (if Active directory is required), SPWS, VPFS
- This feature is limited to 50 concurrent mount points from a Media Server Deduplication Pool (MSDP) media server or from a WORM storage server. If you have a Flex appliance, this feature is limited to 50 concurrent mount points from each node.
- For file or folder download with the **Download** option and the create instant access VM feature, the NetBackup web UI must be able to access the media server with the same name or IP address that the primary server uses to connect to that media server.
- If the media server appliance uses a third-party certificate, you need to create certain configurations on the NetBackup primary server before you use this feature.

For more information, refer to the "Third-party certificates" and "Implementing third-party SSL certificates" sections in the [NetBackup Appliance Security Guide](#).

- A 5-minutes-alive-session threshold is defined in Appliance and BYO web server NGINX. The files and folders that are selected for download must be compressed and downloaded within this threshold.
- To ensure that Instant Access works effectively after the storage server and primary server are upgraded from an earlier NetBackup version, restart the NetBackup Web Service on the upgraded primary server with the following commands:

```
/usr/opensv/netbackup/bin/nbwmc stop
/usr/opensv/netbackup/bin/nbwmc start
```

- If you have to download or restore files or folders from a Windows VM, ensure that the number of Windows registry hives are less than 10000. More information is available about [registry hives](#).

Limitations

- This feature does not support VMs that have a disk in the raw device mapping mode (RDM) or the VMs that have a disk in the Persistent mode.

- For Windows restore, the ReFS file system is not supported.
- The Instant Access feature does not support a Windows 10 compact operating system. To verify if your operating system is compressed, run `compact /compactos:query` on the command prompt before backing up your VM. To disable the compression, run `compact /compactos:never` on the command prompt before backing up your VM. You can then use the Instant Access feature for your VM backups.
- The instant access feature does not support hard links. If you create a universal share from an image and the image has hard link files, `vpfsd` shows these hard link files as having 0 bytes size.
- For a Linux VM, mirrored volumes are not supported for instant access.

Create an instant access VM

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

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

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

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

See [“Things to consider and limitations before you use the instant access feature”](#) on page 56.

To create an instant access VM

- 1 On the left, click **Workloads > Nutanix AHV**.
- 2 Locate the VM and click on it.

- 3 Click the **Recovery points** tab, then click the date on which the backup occurred.

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

- 4 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Create instant access virtual machine**.
- 5 Review the **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.

Create new VM ID instead of existing one

Creates 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.

Power on after recovery

Automatically powers on the VM when the recovery is complete.

Enable migration

Automatically moves the storage associated with the instant access VM from the NetBackup storage to a container on the Nutanix cluster.

- 7 Review or change the Advanced options.

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**.

See [“Pre-recovery checks for Nutanix AHV”](#) on page 99.

- 9 Click **Start recovery**.

- 10 Click the **Restore activity** tab to monitor a job's progress. Select a specific job to view its details.
- 11 Click the **Instant access virtual machines** tab to view the details of the new VM.

Download files and folders from a VM backup image

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

Note: More information on using instance access VMs is available: See [“Things to consider and limitations before you use the instant access feature”](#) on page 56.

To download files and folders from a VM backup image

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

The available images are listed in rows with the backup timestamp for each image.
- 4 On the image or the copy of the image that has the option to recover using instant access, click **Recover > Download files and folders**.
- 5 Select the files and click **Add** to add the files in the download list.

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

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

Enter a file name to search for files.

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

- 6 Click **Next**.
- 7 After the download package is created, click **Download**.

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

Instant access Build Your Own (BYO)

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

- Create instant access VMs.
- Migration of data disks to Acropolis Operating System (AOS) cluster.
- Download files and folders.

To use instant access with a BYO VM created with an earlier NetBackup release, you must upgrade to NetBackup 11.0 or later.

Prerequisites of Instant Access Build Your Own (BYO)

Prerequisites (fresh install and upgrade):

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

- For BYO, docker/podman container is used to browse VMDK files. Data related to the container is stored at the following location: `/var/lib/` and requires minimum 20 GB free space.

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

Table 5-2 Hardware configuration requirement

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

Frequently asked questions

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

Table 5-3 Frequently asked questions

Frequently asked question	Answer
How can I enable instant access file browsing (for file download and restore) on BYO after the storage is configured or upgraded without the docker/podman installed?	Perform the steps in the following order: 1 Install the required docker/podman version. 2 Start using the Instant Access feature. For example, you can download files, restore files, and so on.

Table 5-3 Frequently asked questions (*continued*)

Frequently asked question	Answer
How can I enable the Nutanix AHV instant access feature on BYO after storage is configured or upgraded without the nginx service installed?	Perform the steps in the following order: 1 Install the required nginx service version. 2 Ensure that the new BYO nginx configuration entry: <code>/etc/nginx/conf.d/byo.conf</code> is part of the HTTP section of the original: <code>/etc/nginx/nginx.conf</code> file. 3 Run the command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>
How can I resolve the following issue in the <code>vpfs-config.log</code> file that is raised from: Verifying that the MSDP REST API is available via https on port 10087	Perform the steps in the following order: 1 Install the <code>policycoreutils</code> and <code>policycoreutils-python</code> packages through yum tool. 2 Add the following rules that SELinux requires for Nginx to bind on the 10087 port. ■ <code>semanage port -a -t http_port_t -p tcp 10087</code> ■ <code>setsebool -P httpd_can_network_connect 1</code> 3 Run the following command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

Table 5-3 Frequently asked questions (*continued*)

Frequently asked question	Answer
Instant Access for BYO uses a self-signed certificate by default and only supports *.pem external certificate. How do I replace it with a certificate signed by external CA (*.pem certificate), if required?	To configure the external certificate, perform the following steps. If the new certificate is already generated (the certificate must contain long and short host names for the media server), go to step 4. 1 Create the RSA public or private key pair. 2 Create a certificate signing request (CSR). The certificate must contain long and short host names for the media server. 3 The External Certificate Authority creates the certificate. 4 Replace <code><PDDE Storage Path>/spws/var/keys/spws.cert</code> with the certificate and replace <code><PDDE Storage Path>/spws/var/keys/spws.key</code> with the private key. 5 Run the following command to reload the certificate: <pre>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</pre>
How can I disable media automount for the instant access livemount share in gnome? If the automount is enabled, the source folder is mounted from the livemount share in gnome and smaller disks appear. In this scenario, the instant access feature does not work properly. The mounted disk content source is from the <code>.../meta_bdev_dir/...</code> folder under livemount share, while the mount target is in the <code>/run/media/...</code> folder.	Follow the guideline to disable the gnome automount: https://access.redhat.com/solutions/20107

Table 5-3 Frequently asked questions (*continued*)

Frequently asked question	Answer
How can I resolve the following issue in the <code>/var/log/vpfs/vpfs-config.log</code> file? <pre>**** Asking the NetBackup Webservice to trust the MSDP webserver (spws) **** /usr/opensv/netbackup/ bin/nblibcurlcmd failed (1):</pre>	Perform the steps in the following order: 1 Ensure that your NetBackup primary server is up and there is no firewall blocking the connection between the NetBackup primary server and storage server. 2 Run the following command on storage server to verify the connection status: <code>/usr/opensv/netbackup/bin/bpclntcmd -pn</code> 3 After the NetBackup primary server is up and connection between the NetBackup primary server and storage server is allowed, run the following command: <code>/usr/opensv/pdde/vpfs/bin/vpfs_config.sh --configure_byo</code>

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 using Protection plan](#)
- [Backup AHV VMs or intelligent groups using Policy](#)
- [Protect AHV VMs within VPC](#)
- [Protect vTPM enabled AHV VMs](#)
- [Customize protection settings for an AHV asset](#)
- [Modify policy 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:
 - Incremental schedules are supported only for backup host version 8.3 or later.
 - If you have windows computer as a backup host, make sure that the version is 9.1 or later.
 - If you want to use Enable virtual machine quiesce option, ensure that the backup host is 9.1 or later.
 - If Intelligent VM group has category attribute as filter then backup host version should be 10.4 or later.
- To protect Nutanix Prism Central related VM attributes, Nutanix Prism Central configuration is required.

Note: To protect Nutanix Prism Central related VM attributes, ensure that the NetBackup host version is 10.1.1 or later.

- Cohesity recommends to use either of protection plan or policy while protecting a Virtual machine or Intelligent group.
- Cohesity recommends to use `backupId` as recovery point instead of `client` and `filter` in the recovery API.

Note: To protect AHV VMs using policies, ensure that the NetBackup primary server, media server, client / backup hosts are upgraded to NetBackup version 11.2.0.1 or above.

Protect AHV VMs or intelligent VM groups using Protection plan

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, click **Workloads > 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 You can adjust one or more of the following settings:
 - **Schedules and retention**
Change the backup start window.
 - **Backup options**
Select the server or host to use for backups.

Note: If **Automatic** option is selected here and this protection plan is used to protect intelligent VM group having category as a filter, please make sure you have atleast one media server with version 10.4 or later which is associated with storage unit.

- **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**.

Backup AHV VMs or intelligent groups using Policy

The following is the procedure to protect a Nutanix-AHV asset using a policy.

Steps to configure policy to backup an asset

- 1 Log in to the NetBackup web UI.
- 2 Click **Protection** and then click **Policies**.
- 3 Click **Add**. Create policy page is displayed.
- 4 On the **Attribute** tab, perform the following actions:
 - Specify the **Policy name**.
 - Select Nutanix-AHV as the **Policy type**.
 - Configure the other values, as required.

- 5 On the **Schedules** tab, click **Add** and specify the backup schedule parameters.
- 6 Click **Virtual machines** tab and select Intelligent groups or Individual virtual machines options.
- 7 Click **Nutanix-AHV** tab and select the server or host to backup.
- 8 Click **Create**.

Protect AHV VMs within VPC

With the release of NetBackup 10.2, virtual machines hosted on Virtual Private network within the Nutanix Prism Central Server can be protected. NetBackup additionally protects the following attributes of the VMs using the configured Nutanix Prism Central.

- **Project:** A set of users with a common set of requirements or a common structure and function. Projects provide logical groupings of user roles for managing resource usage.
- **Associated categories:** A category is a grouping of entities into a key value pair. Typically, new entities are assigned to a category based on some criteria. Policies can then be tied to those entities that are assigned (grouped by) a specific category value.
- **VPC network attributes:** Primary and secondary IPs assigned to VMs within VPC.
- **Owner of the project:** The user/owner of the project where CALM is co-deployed within the Nutanix Prism Central.

To protect VMs on VPC

- 1 Configure Nutanix Prism Central.

Note: For configured cluster, NetBackup uses Nutanix Prism Central to protect additional attributes for VM, only if **Use Prism Central** check box is selected.

See [“Add a new Nutanix Prism Central”](#) on page 32.

- 2 Add all the Nutanix AHV clusters in NetBackup with **Use Prism Central** check box as selected.

See [“Add or browse an AHV cluster”](#) on page 28.

- 3 For complete information on protecting VMs, see the following section:

See [“Protect AHV VMs or intelligent VM groups using Protection plan”](#) on page 67.

Note: If you select **Automatic** option in a protection plan to Select server or host to use for backups and the storage unit is associated with NetBackup media server version older than 10.2. Then, the backup job might use the older media server as a backup host.

In this case, the backup job is completed without protecting Nutanix Prism central attributes.

Protect vTPM enabled AHV VMs

A Trusted Platform Module (TPM) is used to manage cryptographic keys for security services like encryption and hardware (and software) integrity protection. AHV virtual trusted platform Module (vTPM) is software-based emulation of the TPM 2.0 specification that works as a virtual device.

NetBackup protects vTPM enabled VMs using the Nutanix-AHV policy and Nutanix protection plan.

You must configure the Nutanix Prism Central and Prism cluster credentials.

See [“Add a new Nutanix Prism Central”](#) on page 32.

See [“Managing AHV cluster credentials”](#) on page 51.

Limitations

- Backup of the information stored in the vTPM is not supported by Nutanix.
- vTPM is not supported on Nutanix VMs with machine type other than Q35.

For more information and recommendations, refer to the Nutanix documentation.

Customize protection settings for an AHV asset

You can customize certain settings for a protection plan, including schedules.

To customize 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 You can edit one or more of the following settings:
 - The backup start window.
See “[Schedules and retention](#)” on page 71.
 - **Backup options**
See “[Backup options](#)” on page 71.
- 5 Click **Protect**.

Modify policy for an AHV asset

This section provides the details to edit a policy as per the requirements. The following is the procedure to edit the policy.

Edit policy

- 1 On the left pane, expand **Protection** and then click **Policies**. **Policies** page is displayed.
- 2 Select the required policy and then click **Edit**. **Edit policy** page is displayed.
- 3 Modify the required values and then click **Save**.

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 72.

- **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 require version 9.1 or later for the NetBackup media server version.

- To install NGT and add scripts, 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 **Workloads > 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 **Workloads > Nutanix AHV**.
- 2 On the grid, click **Show or Hide columns**. Click **Protected by policy**.

- 3 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 See [“Errors for the Status for a newly discovered VM”](#) on page 113.

- 4 If the asset is not protected, click **Add protection** to select a protection plan.
See [“Protect AHV VMs or intelligent VM groups using Protection plan”](#) on page 67.

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](#)
- [Recover an AHV VM within VPC](#)
- [Recover a vTPM enabled AHV VM](#)
- [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 for Nutanix AHV](#)
- [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 the backup host can communicate with the AHV cluster and Prism Central server (if installed) 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 **Workloads > 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.

- 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.

Allow overwrite of existing virtual machine	Deletes any VM with the same display name that exists at the destination. That VM must be deleted before the recovery begins. Otherwise, the recovery fails.
Power on after recovery	Automatically powers on the VM when the recovery is complete.
Recovery host	Indicate the host that you want to use to perform the recovery. By default, the recovery host is the one that performed the backup.
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.
Restore VM from snapshot	Lets you restore the VM from snapshot. Note: If snapshot is not available, VM is restored from the backup image.

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

Remove network interfaces

Remove the network interfaces (not valid for alternate restore on a different cluster of the same AHV server)

Note: To restore to an alternate location, you must select this option.

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.

See [“Pre-recovery checks for Nutanix AHV”](#) on page 99.

9 Click **Start recovery**.

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

Recover an AHV VM within VPC

The recovery of VMs on VPC comes with certain limitations listed below:

- For alternate restores, the restore operation is successful if **Remove network interfaces** check box is selected. However, attributes like Project, category, owner information, and VPC-related information are not restored.
- If there was a network configured in VM when backup was triggered and user tries alternate restore with this backup image keeping **Remove network interfaces** check box not selected, then restore operation fails.
- If a VM had Project configured, when the backup was triggered but the project did not exist when the restore was triggered, the restore job fails.
- If a VM had Category configured, when the backup was triggered but the category did not exist when the restore was triggered, the restore job fails.
- If VMs user is not present in the Nutanix Prism Central server or in Project at the time of restore, then restore operation fail.
- Only IP addresses with type ASSIGNED is considered for restore. Learned IP type is ignored and user must manually configure the IP after restore.

- If the VM has NIC with span port enabled, then after restore this is ignored. You must manually add and configure span on the NIC with Nutanix CLI.
- A VM is restored with Project, Category, Owner details, and other VPC-related attributes, if the original location restore has been performed.
- An undefined behavior is observed when trying to backup/restore the VMs of a cluster which have been moved from one Prism Central to another.
- A backup host with version 10.2 or later is required to backup/restore project, category and other Virtual Private Cloud (VPC) related attributes. If a backup host with version lesser than 10.2 is used, backup/restore will complete without capturing VPC related attributes if a VM is not present in VPC environment. If a VM is in VPC environment and if restore is triggered through a backup host with version lesser than 10.2, then the restore might fail.

Recover a vTPM enabled AHV VM

NetBackup recovers vTPM configuration for vTPM enabled AHV VMs.

You must configure the Prism Central and Prism cluster credentials.

See [“Add a new Nutanix Prism Central”](#) on page 32.

See [“Managing AHV cluster credentials”](#) on page 51.

To recover a vTPM enabled VM, refer to the following section:

See [“Recover an AHV virtual machine”](#) on page 76.

Limitations

- Restore of the information stored in the vTPM is not supported by Nutanix.
- vTPM is not supported on Nutanix VMs with machine type other than Q35.
- Instant access VM restore does not support vTPM configuration restore.

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 7-1 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.
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.

Table 7-1 Prerequisites for files and folder recovery (*continued*)

Step overview	Description and reference
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 7-1 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 7-1 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 7-1 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 - There must be a single entry configured for non-root user in sudoers file. - 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 92.

Table 7-1 Prerequisites for files and folder recovery *(continued)*

Step overview	Description and reference
Windows target host	

Table 7-1 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://support.cohesity.com/s/article/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 7-1 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 7-1 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 7-1 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, click **Workloads > 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 94.
- 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 for Nutanix AHV](#)” on page 99.

- 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 7-2 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: An agent-based restore is performed if the NetBackup client is installed and configured with provided hostname or IP. ■ If you want to perform a 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 7-2 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.

Table 7-2 Recovery target options (*continued*)

Step overview	Description and reference
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. ■ In flex scale environment, if all media servers are not listed in media server tab, then user either need view permission on media server or can manually type the media server to proceed. ■ 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 7-2 Recovery target options (*continued*)

Step overview	Description and reference
Linux SSH Connectivity	

Table 7-2 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 92. ■ 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 81. ■ 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 7-2 Recovery target options (*continued*)

Step overview	Description and reference
	choose SSH private key based authentication. See "SSH key fingerprint" on page 92. ■ 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 for Nutanix AHV

Table 7-3 Pre-recovery checks for Nutanix AHV

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 7-3 Pre-recovery checks for Nutanix AHV (*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 a source AHV VM backed up image. The guest operating system and the file system must be compatible for creating file system mappings.
Refer to the *Nutanix AHV SCL for guest operating system and file system support for individual files restores*.
[Support for NetBackup <versions> in virtual environments](#)
- You can perform individual files recovery from a source AHV VM backup. The NetBackup primary server, media server, and backup host must be at NetBackup version 9.1 or later for.
- An agent-based restore is performed if the target host has the NetBackup client or server installed. The client or the target host must be at NetBackup 8.1 or later (Windows) or 8.2 or later (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 or IP in the target host to follow the agent-based restore.

- A user with the necessary RBAC permissions to view NetBackup hosts can browse and select the NetBackup host for a restore of files or folders.
A user without the necessary RBAC permissions must manually specify the NetBackup configured host name or IP for the target host.
- The following are the minimum RBAC permissions required for a user to do agent-based files and folders restore.

Table 7-4 Permissions for all AHV assets

Operation	Description	Additional required operations	Additional optional operations
Granular restore	Restore individual files or folders from an 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

Table 7-5 Permissions for all AHV assets

Operation	Description	Additional required operations	Additional optional operations
Granular restore	Restore individual files or folders from an 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

To 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 **Workloads > 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 94.
- 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 operation of individual files is not supported. You can restore windows files only on 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.

Note: The last restored file is available, when the restore completes.

- If you select the **Flatten existing directory structure** option but do not select **Overwrite existing files**, the restore succeeds, but 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: Cohesity 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 before 9.1 release, cannot be restored from web UI. To restore these images user must use BAR GUI.
- The Linux ext4 file system includes a persistent pre-allocation feature to guarantee disk space for files without padding the allocated space with zeros. When NetBackup restores a pre-allocated file (to any supported ext file system), the file loses its preallocation and is restored as a sparse file. The restored sparse file is only as large as the last byte that was written to the original file. Note also that subsequent writes to the sparse file may be non-contiguous.
- AHV backup images are available on web UI, even if the backup was taken from NetBackup Administration Console 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 NetBackup Administration Console, backup images are available on web UI.
- If the primary sever and backup host are upgraded to 9.1, backup is taken from NetBackup Administration Console, 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 before 9.1 and backup is taken from NetBackup Administration Console. Then, if you configure web UI, you cannot see the backup images even after asset discovery.

Protecting Nutanix Cloud Clusters (NC2)

This chapter includes the following topics:

- [Protecting Nutanix Cloud Clusters \(NC2\) on AWS](#)
- [Protecting Nutanix Cloud Clusters \(NC2\) on Azure](#)

Protecting Nutanix Cloud Clusters (NC2) on AWS

Nutanix Cloud Clusters (NC2) is an extension of the Nutanix Cloud Platform, enabling organizations to run Nutanix Cloud Platform software on AWS. It replicates the core Nutanix HCI software used on-premises in a public cloud hyperscaler environment, delivering the same virtualized and software-defined benefits in both private and public clouds.

For more information on Nutanix Cloud Clusters (NC2) on AWS, refer to the Nutanix documentation online.

In the Nutanix Cloud Clusters (NC2) environment, virtual machines can be protected using NetBackup 10.4 and later versions. Once Nutanix clusters and Prism Central are deployed in the Nutanix Cloud Clusters (NC2) environment, they can be configured within the NetBackup Web UI, similar to on-premises Nutanix clusters and Prism Central. After successfully configuring the cluster in NetBackup, the virtual machines will be discovered. Subsequently, these virtual machines can be protected using the protection plan designed for AHV workloads.

Note: NetBackup protects virtual machines in the Nutanix Cloud Clusters (NC2) environment similarly to those in on-prem Nutanix clusters. For more details on protecting Nutanix on-premises virtual machines with NetBackup, refer to the *Managing AHV clusters* chapter in *NetBackup™ for Nutanix AHV Administrator's Guide*.

Protecting Nutanix Cloud Clusters (NC2) on Azure

Nutanix Cloud Clusters (NC2) is an extension of the Nutanix Cloud Platform, enabling organizations to run Nutanix Cloud Platform software on the Microsoft Azure cloud service. It replicates the core Nutanix HCI software used on-premises in a public cloud hyperscaler environment, delivering the same virtualization and software-defined benefits in both private and public clouds.

For more information on Nutanix Cloud Clusters (NC2) on Azure, refer to the Nutanix documentation online.

In the Nutanix Cloud Clusters (NC2) environment, virtual machines can be protected using NetBackup 10.4 and later versions. Once Nutanix clusters and Prism Central are deployed in the Nutanix Cloud Clusters (NC2) environment, they can be configured within the NetBackup Web UI, similar to on-premises Nutanix clusters and Prism Central. After successfully configuring the cluster in NetBackup, the virtual machines will be discovered. Subsequently, these virtual machines can be protected using the protection plan designed for AHV workloads.

Note: NetBackup protects virtual machines in the Nutanix Cloud Clusters (NC2) environment similarly to those in on-prem Nutanix clusters.

For more details on protecting Nutanix on-prem virtual machines with NetBackup, refer to the *Managing AHV clusters* chapter in *NetBackup™ for Nutanix AHV Administrator's Guide*.

Troubleshooting AHV operations

This chapter includes the following topics:

- [Troubleshooting AHV operations: Error during creation of AHV instant access virtual machines](#)
- [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 AHV operations: Error during creation of AHV instant access virtual machines

The following table describes the problems that may occur when you try to create AHV instant access virtual machines.

Table 9-1 Error during creation of AHV instant access virtual machines

Error message or cause	Explanation and recommended action
Failed to create the VM in the Nutanix AHV cluster. Error: Already used VM UUID: <VM UUID>. Return value 114. Failed to create an instant access virtual machine (VM). (4004)	Before creating an Instant Access VM, the system checks if the VM UUID already exists in the Nutanix cluster. If it does, the recovery fails. Workaround: 1 Start the recovery process. 2 Select Create Instant Access Virtual Machine. 3 On the Recovery Options page, select Create new VM ID instead of existing one 4 Continue the recovery and click Start Recovery.
Failed to create an instant access virtual machine (VM). Error details Only the Red Hat Enterprise Linux and SUSE Linux Enterprise Server operating systems are supported.	Instant Access currently supports only RHEL and SLES. If you try to recover unsupported operating system VMs, it fails. Recommended action: Use traditional full restore for media server with unsupported operating systems.
Failed to create an instant access virtual machine (VM). Specification for create VM cannot be prepared, as RetainMacAddress and RemoveNetworkInterface can't be used together.	The API spec becomes invalid if both <code>RetainMacAddress</code> and <code>RemoveNetworkInterface</code> are used at the same time. Recommended action: Do not retain the MAC address if the Network Interface Controllers (NICs) are removed. If the MAC address needs to be retained, NIC cards must remain configured.

Table 9-1 Error during creation of AHV instant access virtual machines
(continued)

Error message or cause	Explanation and recommended action
Failed to create an instant access virtual machine (VM). Error details Create instant access VM capability for Nutanix is supported from NetBackup 11.1 and above.	You cannot create Instant Access VMs with NetBackup 11.0.0.1 and earlier versions. Recommended action: Ensure that the restore is completed with NetBackup version 11.1 or later.
Failed to create the VM in the Nutanix AHV cluster. Error: InUse: MAC address <MAC address> is in use: 18. Return value 114. Failed to create an instant access virtual machine (VM). (4004)	This occurs when the original VM still exists on the Nutanix cluster and the MAC address is already in use, causing a conflict. Recommended action: Ensure the original VM is deleted or powered off before instant access recovery to avoid MAC address duplication.

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 9-2 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 9-3 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 9-3

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 9-4 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 problems that might occur when you back up AHV virtual machines:

Table 9-5 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 the Microsoft iSCSI Initiator Service (MSiSCSI) service on the windows backup host and re-run the job.

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Unable to establish a connection. Verify that the iSCSI service is installed and running.	

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
	- For Windows: Enable the Microsoft iSCSI Initiator Service on the backup host. Note: An error displayed only for windows OS. - For Linux, this error is displayed in the form of a warning and it falls back to use the NFS for backup/restore. If you use segmented iSCSI data services backup/restore fails. Ensure that the backup host is added in an <code>Filesystem Whitelists</code> option in the Nutanix UI for backups to work by NFS transport. For Linux to use iSCSI: Install / Enable iSCSI Initiator package on backup host and re-run the job. - Validate connectivity using the following command on a Linux host: <pre>iscsiadm -m discovery -t sendtargets -p correct IP as per configured iSCSI targetType</pre> Use following IP address based on your iSCSI transport type: - For iSCSI Data services: iSCSI Data services IP from cluster details page - For segmented: Use segmented iSCSI Data from cluster details page SEGMENTED_SPECIFIC. - For specified segmented: Use Virtual IP specified while you configure a cluster in NetBackup. - Validate connectivity using the following command on the Windows host: - Click Server Manager -> Tools -> iSCSI initiator. This opens the iSCSI initiator properties dialog. - Click Discovery > Discovery portal and provide the IP address as per the configured iSCSI Target Type for the AHV cluster. - Default: iSCSI Data services IP from cluster details page - SEGMENTED: Segmented iSCSI Data from cluster details page - SEGMENTED_SPECIFIC: Virtual IP specified while you configure a cluster in NetBackup. - If you encounter the error message using Flex Scale appliances, please note that only the NetBackup primary server does not support iSCSI data path but supports NFS as the data path.

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
	However, if iSCSI data path usage is essential, it is recommended to use a Media Server instead of the Primary Server. The Media Server effectively manage the iSCSI data path, ensuring proper backup and restore functionality. ■ If you encounter this error using Flex appliance as Backup/Recovery host: ■ Edit the configuration to use default option that uses the NFS transport. ■ Use different Backup/Recovery host which has required iSCSI and network configuration. ■ Update the protection plan to use specific Backup host. ■ Use Recovery host with required iSCSI and network configuration.
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 or recovery host. Set the iSCSI initiator name uniquely 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 "Prerequisites to configure Nutanix AHV cluster" on page 25. Note: For Linux it falls back to use the NFS for backup/restore.
NetBackup version is not supported for one or more backup hosts. Use NetBackup version 9.1 or later on all Linux or Windows backup host to use the Automatic backup host option in the Nutanix protection plan.	This error occurs when the Automatic option is selected for backup hosts in the Nutanix protection plan. Upgrade backup host to latest NetBackup version.
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 servers are: ■ Red Hat Enterprise Linux ■ SUSE Linux Enterprise Server ■ Microsoft Windows operating system
The 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.

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Unable to set resource limits for specific Nutanix clusters.	If the clusters for which the resource limit is set are deleted from the NetBackup environment, in some cases, the + Add option is disabled to set the resource limit. Recommended action Delete the resource limit for the deleted clusters and then set the resource limit for the 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> abc.cbush.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 previous 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 the 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 9-5 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 the 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 the retry operation is successful, the backup/restore job is also successful. Recommended action No action is required on the NetBackup side. User can still troubleshoot the <code>iscsiadm</code> and ensure correct iSCSI installation/configuration to avoid such errors.
<pre>iscsid: Ignoring CHAP algorithm request for MD5 due to crypto lib configuration iscsid: Couldn't set CHAP algorithm list</pre>	See, In FIPS enabled environment, NetBackup backup/restore of Nutanix AHV VMs (Virtual Machines) using iSCSI fails

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Error code: 4798 Use Prism Central server for this cluster option is not selected for AHV cluster.	Discovery job during backup may fail with the error for Nutanix Intelligent VM group. Look and fix the following possible causes for Intelligent VM group: ■ Created with category filter as one of the filter queries AND ■ One or more Nutanix clusters are updated after the intelligent VM group is created to use the Use Prism Central server for this cluster option unchecked AND Backup now operation is triggered on such IVMG.
Error message: Unable to find the Prism Central for AHV cluster, server = <i>Server details</i>	Backup job fails with the given error. Look and fix the following possible causes: ■ In an Intelligent VM group consisting of one or more clusters from same Prism Central server and category as one of the filter, when IVM group protection is triggered, Prism Central server is deleted/not accessible. ■ In an Intelligent VM group consisting of two or more clusters from different Prism Central servers and category as one of the filter, when an Intelligent VM group protection is triggered, one or more Prism Central servers are deleted/not accessible.
Error message: Protection plan subscription should fail with error. An invalid API request is encountered error message: backupHost: Backup host with a NetBackup version earlier than 10.4 is not supported for IntelligentVM group Category filter.	If a category filter is used in the Intelligent VM group, subscribing it to a protection plan may fail with the given error. ■ Ensure that the Backup host mentioned in the protection plan must have NetBackup version 10.4 or later.

Table 9-5 Error while backing up AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Backup job fails with error code 800. <pre>Error nbjm(pid=113200) NetBackup status: 800, EMM status :Use NetBackup media server version 10.4 or later to protect Nutanix Intelligent VM groups with category filters.</pre> . <pre>Error nbpem(pid=113293) backup of client MEDIA_SERVER exited with status 800 (resource request failed).</pre>	Description 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 earlier than NetBackup 10.4 versions. And when this protection plan is used to backup the Intelligent VM group with category attribute as filter, the backup job will fail. Recommended action: At least one of the media server, which is configured in the selected storage unit, must be upgraded to NetBackup v10.4 or later.
Backup job fails with following error messages: Error 1 <pre>Begin Application Resolver:Resolver Discovery</pre> Error 2 <pre>Error nbpem(pid=98395) Invalid URI.</pre> Error 3 <pre>Error nbpem (pid=98395) backup of client falcna12c3.abcus.com exited with status 4232 Invalid Discovery Query URI).</pre>	Description When the Intelligent VM group is subscribed with a protection plan that have a backup host version 10.3 or older, and the Intelligent VM group is modified with the category filter. Then, once the backup job runs, it fails with an error message. As the category filter is not known to the older versions of backup hosts. Recommended action: Upgrade the backup host to 10.4 or later by customizing the protection plan. More details See “Customize protection settings for an AHV asset” on page 70.

Error while restoring AHV virtual machines

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

Table 9-6 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.
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)	The pre-recovery check compares the VM display name instead of UUID to find out if the 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 the 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 backup.	During the recovery of imported AHV VM images, 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 the Microsoft iSCSI Initiator Service (MSiSCSI) service on the Windows backup recovery and re-run the job.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Failed to connect to the recovery host.	The recovery which is host used for agentless restore is not reachable. Recommended action: Ensure that the recovery host is reachable from the primary server and has the NetBackup media server or the client software installed on it.
The specified recovery host must be at NetBackup version 9.1 or later to support agentless restores.	Agentless restores of files or folders require a recovery host with NetBackup version 9.1 or later. Recommended action: Verify the NetBackup version on the recovery host. It should be 9.1 or later. On UNIX NetBackup servers and clients, verify the <code>/usr/openv/netbackup/bin/version</code> file. On Windows NetBackup servers, verify 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 the recovery host is valid. NetBackup uses the following on the recovery host as the 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, verify if a valid path on the 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 the recovery host staging location. This image is required for agentless restore. Recommended action: Contact Cohesity Technical Support and share <code>bpVMutil</code> log from the recovery host.
Internal error has caused failure of recovery validation.	An internal error occurred while running the pre-recovery validations for agentless restore. Recommended action: Save the bpVMutil logs on recovery host and contact Cohesity Technical support.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Not enough space available on recovery host.	The recovery host may not have enough space to copy the selected files at the 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 to perform the agentless restore.
Tar utility is not present on the target host.	Failed to find the tar utility on the target host. This utility is 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 the minimum write and run permission on the staging location.
The user does not have root/administrator privileges. To restore files and folders, provide user with root or 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 the 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 the target host is not accessible from the recovery host.	The admin share of the remote host is not accessible from the recovery host to perform an agentless restore. Recommended action: ■ Ensure that firewall exceptions are set up correctly. ■ Ensure that File and Printer Sharing is enabled. ■ Ensure that the GPO/Software Restriction Policy or Antivirus does not block access. ■ Ensure that the target host is accessible and ensure that the correct credentials are entered and have proper permissions.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
For agentless files or folders restore in User Account Control (UAC) environment, provide credential of domain user which is a part of local Administrator group on the Windows target host.	Recommended action: For an agentless restore in the User Account Control (UAC) environment, provide the credential of the domain user. This user is the part of the local Administrator group on the Windows target host.
Agentless restore is not possible.	Received an unexpected reason for agentless restore failure. Recommended action: Contact Cohesity 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 the same. Recommended action: The alternate recovery host must have the same operating system as the VM that was backed up.
Failed to retrieve the backup image operating system.	Unable to retrieve operating system of backup image for performing an 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.	The recovery host OS type and the communication type that was provided in the agentless recovery or pre-recovery check request are not compatible. Recommended action: Verify that Recovery host OS type and the communication type is compatible for the recovery host: ■ Linux: The communication type must be <code>SSH</code>. ■ Windows: The communication type must be <code>WMI</code>.
Target host SSH private key is invalid.	The <code>sshKey</code> field of the agentless recovery or pre-recovery check request must be valid and a non-empty ssh private key of the 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.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
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 the 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 Software Compatibility List for the supported platforms for this feature.
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 the 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 a custom staging location on the 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 the 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 the WMI and the DCOM services the 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 does not block the access. ■ Ensure that the 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, a two-way trust relationship between those domains must exist.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Unable to find the specified file on the remote server.	Unable to find the specified file on the remote server. Recommended action: Ensure that the specified staging location on the target host exists or specify another valid staging location.
File exists with same name as the directory.	On the target host a pre-existing file with same name exists as that in the directory path of the staging location. Recommended action: Check for a pre-existing file on a remote host that has the same name and path as the staging location. If it exists, either rename or remove that file. Or specify an alternate staging location.
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 the 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.	The admin share of the 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 usually occurs when the target host for agentless restore has the Windows OS and the admin share of the target host is not accessible from the recovery host. Ensure that the following requirements are met on the target host. ■ Firewall Exceptions are set up correctly. ■ File and Printer Sharing is enabled. ■ GPO/Software restriction policy or antivirus does not block the access. ■ Target host is accessible with valid credentials.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
Unable to retrieve user's home directory on the target host. Specify the custom staging location.	The user's default staging location that is on the home directory can't be retrieved on the target host. Enter a valid custom staging location path. Recommended action: Ensure that the user's home directory exists or try with a valid custom staging location.
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 that the method used that is used to generate the 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.	The SSH key fingerprint of the target host that was 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 that the SSH private key and the key passphrase are valid that were used to generate the SSH private key of the target host. Then 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.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
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 a supported host key method of the specified SSH key fingerprint is available on the target host. Or provide the SSH fingerprint of the host key method that is configured on the 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.
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 <code>ButterFS</code>, 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 /dev level (such as <code>/dev/sda1</code>). You cannot restore the files from their original mount point level.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
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 9-6 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, Cohesity 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>
Virtual machine creation failed, cannot proceed with restore. bpVMutil pid=3144	If a backup host with NetBackup version 10.1.1 or earlier is used to restore VMs in Virtual Private Cloud (VPC) environment, the restore job fails. Recommended action Use a backup host version NetBackup 10.1.1 or later to restore VMs which are in a VPC environment.
Restore from snapshot job completes with partial successful status.	Restore from snapshot job completes with partial successful status if the AHV cluster does not have the correct configuration as per the iSCSI transport option. Workaround Verify and fix the following error based on iSCSI transport setting: ■ For default: The iSCSI data services IP is configured. ■ For segmented: The segmented IP address is not configured. ■ For segmented_specified: The segmented iSCSI interface is not configured or the specified IP address does not match with virtual IP of any of the segmented iSCSI interfaces that are configured.

Table 9-6 Error run into while restoring AHV virtual machines (*continued*)

Error message or cause	Explanation and recommended action
A backup host with a NetBackup version earlier than 11.2.0.1 is not supported for the Nutanix-AHV policy.	Cohesity recommends upgrading NetBackup to latest version.
NetBackup status: 213, EMM status: NetBackup media server version is too low for the operation. No storage units available for use(213).	Check the storage unit that is used in the Nutanix-AHV policy. The media server on which the storage unit was created should be at NetBackup version 11.0 or later. For more details check the logs at path: <code>/usr/openv/logs/nbwebsevice</code>

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 135.
- See [the section called “Set iSCSI CHAP settings APIs”](#) on page 135.
- See [the section called “Create an AHV VM backup policy”](#) on page 136.
- See [the section called “Pre-Recovery check for AHV VM at the original location”](#) on page 137.
- See [the section called “Pre-Recovery check for AHV VM at a different location”](#) on page 138.
- See [the section called “Restore the AHV VM at the original location”](#) on page 138.
- See [the section called “Restore the AHV VM to an alternate location”](#) on page 140.

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 10-1 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 10-2 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 10-3 Create an AHV VM backup policy

API or command line options	Important variables and options
POST /netbackup/config/policies/	- policyType is Hypervisor. - policyType is Nutanix-AHV using web UI. - backuphost is a hostname of a NetBackup client that performs backups on behalf of the virtual machines. - 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 10-4 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 10-5 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 10-6 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 10-6 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 142.

Restore the AHV VM to an alternate location

Table 10-7 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 10-7 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 142.

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 10-8

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

`change vmname to newVMname`

Recover the virtual machine to a different AHV container

`change /<original_container1>/<disk_uuid1>
to /<alternate_container1>`

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.