



# Veeam Backup for Nutanix AHV

---

Version 7

User Guide

January, 2025

© 2025 Veeam Software.

All rights reserved. All trademarks are the property of their respective owners.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form by any means, without written permission from Veeam Software (Veeam). The information contained in this document represents the current view of Veeam on the issue discussed as of the date of publication and is subject to change without notice. Veeam shall not be liable for technical or editorial errors or omissions contained herein. Veeam makes no warranties, express or implied, in this document. Veeam may have patents, patent applications, trademark, copyright, or other intellectual property rights covering the subject matter of this document. All other trademarks mentioned herein are the property of their respective owners. Except as expressly provided in any written license agreement from Veeam, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

#### **NOTE**

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

# Contents

<b>CONTACTING VEEAM SOFTWARE .....</b>	<b>7</b>
<b>ABOUT THIS DOCUMENT .....</b>	<b>8</b>
<b>OVERVIEW .....</b>	<b>9</b>
Solution Architecture .....	10
Prism Central Deployment Scenario .....	11
Standalone Cluster Deployment Scenario .....	13
VM Backup .....	15
Backup Chain .....	17
Backup Methods .....	20
Active Full Backup .....	22
Synthetic Full Backup .....	23
Snapshot Chain .....	24
VM Restore .....	27
Entire VM Restore .....	28
Disk Restore .....	29
File-Level Recovery .....	30
Retention Policies .....	31
<b>PLANNING AND PREPARATION .....</b>	<b>32</b>
System Requirements .....	33
Permissions .....	35
Ports .....	39
Sizing Guidelines .....	49
<b>LICENSING .....</b>	<b>51</b>
<b>DEPLOYMENT .....</b>	<b>52</b>
Installing Nutanix AHV Plug-In Manually .....	53
Installing Plug-In in Unattended Mode .....	55
Upgrading to Veeam Backup for Nutanix AHV 7 .....	58
Uninstalling Nutanix AHV Plug-In .....	59
<b>CONFIGURING BACKUP INFRASTRUCTURE .....</b>	<b>60</b>
Configuring Backup Repositories .....	61
Connecting Nutanix AHV Server .....	62
Adding Nutanix AHV Server .....	63
Editing Nutanix AHV Server Properties .....	72
Rescanning Nutanix AHV Server .....	73
Removing Nutanix AHV Server .....	74
Accessing Nutanix AHV Server Console .....	75
Managing Backup Appliances .....	77

Deploying New Backup Appliance .....	78
Connecting Existing Backup Appliance .....	89
Editing Backup Appliance .....	99
Rescanning Backup Appliance .....	101
Removing Backup Appliance .....	102
Migrating Backup Appliance .....	104
Troubleshooting Backup Appliance .....	105
Managing Workers .....	106
Adding Workers .....	107
Enabling and Disabling Workers .....	113
Editing Workers .....	114
Testing Workers .....	115
Disabling Automatic Worker Updates .....	116
Removing Workers .....	117
<b>ACCESSING BACKUP APPLIANCE .....</b>	<b>118</b>
Changing Password .....	120
<b>CONFIGURING VEEAM BACKUP FOR NUTANIX AHV .....</b>	<b>121</b>
Managing Users .....	122
Adding User Accounts .....	124
Editing User Account Settings .....	125
Changing User Passwords .....	126
Configuring Multi-Factor Authentication .....	127
Enabling SSH on Backup Appliance .....	129
Viewing Network Settings .....	130
Configuring Time Synchronization Settings .....	131
Configuring Email Notification Settings .....	132
Configuring Notifications .....	135
Performing Configuration Backup and Restore .....	136
Backing Up Configuration Settings Manually .....	137
Backing Up Configuration Settings Automatically .....	139
Restoring Configuration Settings .....	141
<b>VIEWING AVAILABLE RESOURCES .....</b>	<b>152</b>
Adding VMs to Job .....	153
<b>PERFORMING BACKUP .....</b>	<b>154</b>
Creating Backup Jobs .....	155
Before You Begin .....	156
Creating Backup Jobs Using Backup Appliance Web Console .....	157
Creating Backup Jobs Using Veeam Backup & Replication Console .....	174
Creating Snapshot Jobs .....	210
Before You Begin .....	211

Creating Snapshot Jobs Using Backup Appliance Web Console .....	212
Creating Snapshot Jobs Using Veeam Backup & Replication Console .....	220
Creating Protection Domain Snapshot Jobs .....	228
Before You Begin .....	229
Creating PD Snapshot Jobs Using Backup Appliance Web Console .....	230
Creating PD Snapshot Jobs Using Veeam Backup & Replication Console .....	238
Analyzing Performance Bottlenecks.....	247
Cloning Jobs.....	249
Starting and Stopping Jobs .....	250
Retrying Jobs .....	251
Editing Job Settings .....	252
Enabling and Disabling Jobs .....	253
Deleting Jobs .....	254
Creating Active Full Backup .....	255
Creating VeeamZIP Backups .....	256
<b>MANAGING BACKUPS AND SNAPSHOTS.....</b>	<b>257</b>
Viewing Backup Properties .....	258
Rescanning Backups.....	259
Verifying Backups .....	260
Exporting Backups .....	261
Copying Backups .....	262
Copying Backups to Tapes.....	263
Deleting Backups.....	264
Deleting Snapshots.....	265
<b>PERFORMING RESTORE.....</b>	<b>266</b>
Performing VM Restore .....	267
Restoring VMs Using Veeam Backup & Replication Console .....	268
Restoring VMs Using Backup Appliance Web Console .....	278
Performing Disk Restore .....	289
Step 1. Launch Virtual Disk Restore Wizard .....	290
Step 2. Select Restore Point .....	291
Step 3. Choose Restore Mode .....	293
Step 4. Specify Target Cluster.....	294
Step 5. Configure Mapping Settings .....	295
Step 6. Specify Reason for Restore .....	296
Step 7. Finish Working with Wizard .....	297
Instant Recovery.....	298
Performing Instant Recovery of Workloads to Nutanix AHV .....	299
Performing Instant Recovery of Workloads to VMware vSphere.....	312
Performing Instant Recovery of Workloads to Hyper-V .....	313

Publishing Disks .....	314
Performing File-Level Restore .....	315
Performing Application Item Restore .....	316
Exporting Disks .....	318
Performing VM Restore to Amazon Web Services .....	319
Performing VM Restore to Microsoft Azure .....	320
Performing VM Restore to Google Cloud .....	321
<b>VIEWING SESSION STATISTICS .....</b>	<b>322</b>
<b>REVIEWING DASHBOARD.....</b>	<b>323</b>
<b>UPDATING BACKUP APPLIANCE .....</b>	<b>325</b>
Configuring Web Proxy .....	326
Checking for Updates.....	327
Installing Updates .....	328
Viewing Update History .....	331
<b>GETTING TECHNICAL SUPPORT.....</b>	<b>332</b>
Viewing Product Details .....	333
Exporting Logs Using Backup Appliance Web Console .....	334
Exporting Logs Using Veeam Backup & Replication Console .....	335
<b>APPENDICES.....</b>	<b>336</b>
Appendix A. Custom Script Samples .....	337
Appendix B. Installing Custom Certificate .....	356
Appendix C. Configuring Bus Type Restore Priority .....	357
Appendix D. Configuring Multiple Networks .....	359

# Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

## Customer Support

Should you have a technical concern, suggestion or question, visit the [Veeam Customer Support Portal](#) to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

## Company Contacts

For the most up-to-date information about company contacts and office locations, visit the [Veeam Contacts Webpage](#).

## Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: [veeam.com/documentation-guides-datasheets.html](http://veeam.com/documentation-guides-datasheets.html)
- Veeam R&D Forums: [forums.veeam.com](http://forums.veeam.com)

# About This Document

This guide is designed for IT professionals who plan to use Veeam Backup for Nutanix AHV. The guide includes system requirements, licensing information and step-by-step deployment instructions. It also provides a comprehensive set of features to ensure easy execution of protection and disaster recovery tasks in Nutanix AHV environments.



# Overview

Veeam Backup for Nutanix AHV is a solution developed for data protection and disaster recovery tasks for the Nutanix AHV environment. With Veeam Backup for Nutanix AHV, you can perform the following operations:

- Create backups of Nutanix AHV VMs and store them in backup repositories.
- Create snapshots of Nutanix AHV VMs and protection domains.
- Create VeeamZIP backups of Nutanix AHV VMs.
- Create several instances (copies) of the same backed-up data in different locations.
- Restore VMs from Nutanix AHV backups and snapshots to the original Nutanix AHV environment.
- Restore VMs from VMware ESXi and Microsoft Hyper-V to the Nutanix AHV environment.
- Restore VMs from oVirt KVM and Proxmox VE backups to the Nutanix AHV environment.
- Restore VMs from Microsoft Azure, Amazon Web Services (AWS) and Google Cloud backups to the Nutanix AHV environment.
- Restore physical machines from backups created by Veeam Agents to the Nutanix AHV environment.
- Restore VMs from Nutanix AHV backups to Microsoft Azure, Amazon Web Services (AWS) and Google Cloud environments.
- Restore VMs from Nutanix AHV backups to VMware vSphere and Microsoft Hyper-V environments.
- Perform Instant Recovery of VMs and physical machines to Nutanix AHV, VMware vSphere and Microsoft Hyper-V environments.
- Restore files and folders of Nutanix AHV VM guest OSes.
- Restore application items (such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Oracle Database and Microsoft SQL Server).
- Restore Nutanix AHV VM disks and attach them to VMs running in Nutanix AHV clusters.
- Export disks of backed-up Nutanix AHV VMs to VMDK, VHD and VHDX formats.
- Mount disks of backed-up Nutanix AHV VMs to any server and access data in the read-only mode.

# Solution Architecture

Starting from version 6.0, Veeam Backup for Nutanix AHV supports 2 deployment scenarios:

- [Prism Central deployment scenario](#) allows you to protect workloads that reside in multiple clusters registered with a Prism Central.

This scenario provides a centralized web console that allows you to manage backup and restore operations performed for workloads in all the registered clusters. Therefore, it reduces time required to install, configure and maintain Veeam Backup for Nutanix AHV,

- [Standalone cluster deployment scenario](#) allows you to protect workloads that reside in a specific cluster.

Even if you add multiple clusters to the backup infrastructure, Veeam Backup for Nutanix AHV will treat each cluster as a dedicated virtual environment. Therefore, backup and restore operations performed for workloads in each cluster will be managed separately.

You can also combine these scenarios to support your own data protection strategy. However, keep in mind that you cannot add to the backup infrastructure both a Prism Central and a standalone Nutanix AHV cluster that is registered with this Prism Central.

# Prism Central Deployment Scenario

In the Prism Central deployment scenario, Veeam Backup for Nutanix AHV architecture comprises the following set of components:

- [Nutanix AHV Prism Central](#)
- [Nutanix AHV clusters](#)
- [Backup server](#)
- [Backup appliance](#)
- [Nutanix AHV Plug-in](#)
- [Backup repositories](#)
- [Workers](#)

## Nutanix AHV Prism Central

The Prism Central is a software appliance that provides a centralized interface for managing multiple clusters in the Nutanix hyper-converged infrastructure (HCI) environment. Veeam Backup for Nutanix AHV uses the Prism Central to access all the registered clusters.

## Nutanix AHV Clusters

A Nutanix AHV cluster is a logical group of Nutanix HCI nodes managed by Nutanix Controller VMs (CVMs). Veeam Backup for Nutanix AHV accesses cluster resources (such as VMs, volume groups, protection domains, storage containers and networks) to perform backup and restore operations.

## Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. The backup server is the configuration, administration and management core of the backup infrastructure. It coordinates backup and restore operations, controls job scheduling and manages resource allocation.

## Backup Appliance

The backup appliance is an architecture component that sits logically between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the backup appliance performs management operations, processes jobs and delivers backup traffic. The backup appliance is a Linux-based VM that resides in one of the Prism Central Nutanix AHV clusters and includes the following components:

- **Backup Appliance Web Console** is an interface that allows you to manage backup and restore operations, and to configure settings of the Nutanix AHV backup appliance.
- **Veeam Updater** is a service that is responsible for installing and scheduling updates for the backup appliance.

# Nutanix AHV Plug-in

Nutanix AHV Plug-in is an architecture component that enables integration between the backup server and the backup appliance. Nutanix AHV Plug-in also allows the backup server to deploy and manage the backup appliance.

## Backup Repositories

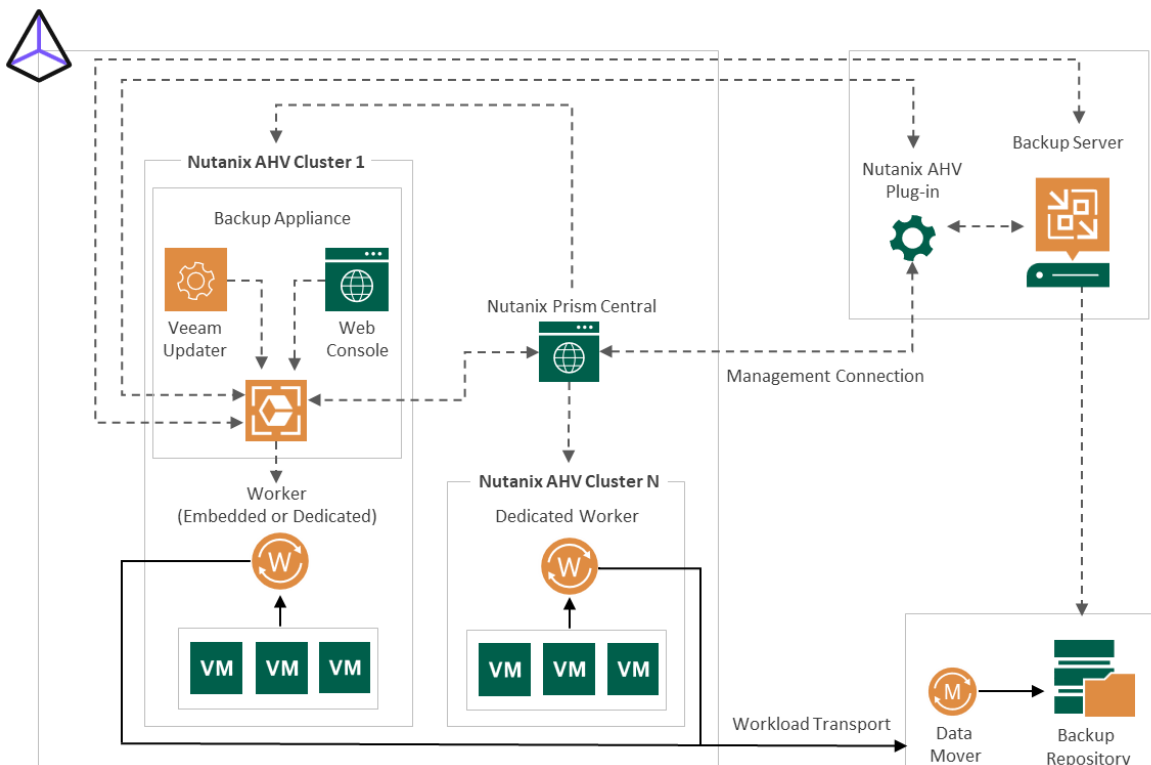
A backup repository is a storage location where Veeam Backup for Nutanix AHV stores backups of protected Nutanix AHV VMs.

To communicate with backup repositories, Veeam Backup for Nutanix AHV uses Veeam Data Mover – the service that is responsible for data processing and transfer. By default, Veeam Data Mover runs on the repositories themselves. If a repository cannot host Veeam Data Mover, it starts on a gateway server – a dedicated component that “bridges” the backup server and workers. For more information, see the Veeam Backup & Replication User Guide, section [Gateway Server](#).

## Workers

A worker is an auxiliary Linux-based VM instance that is responsible for the interaction between the backup appliance and other components of the backup infrastructure. Workers process backup workload and distribute backup traffic when transferring data to and from backup repositories.

The backup appliance comes with a preconfigured embedded worker that can be used in small virtual environments. In large environments, it is recommended to deploy dedicated workers that are distributed among the cluster hosts (nodes) and are automatically launched for the duration of a backup or restore process. For more information on deployment sizing considerations, see [Sizing Guidelines](#).



# Standalone Cluster Deployment Scenario

In the standalone cluster deployment scenario, Veeam Backup for Nutanix AHV architecture comprises the following set of components:

- [Nutanix AHV cluster](#)
- [Backup server](#)
- [Backup appliance](#)
- [Nutanix AHV Plug-in](#)
- [Backup repositories](#)
- [Workers](#)

## Nutanix AHV Cluster

The Nutanix AHV cluster is a logical group of Nutanix HCI nodes managed by Nutanix Controller VMs (CVMs). While performing backup and restore operations, Veeam Backup for Nutanix AHV uses the Nutanix AHV cluster to access Nutanix AHV resources such as VMs, volume groups, storage containers and networks.

## Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. The backup server is the configuration, administration and management core of the backup infrastructure. It coordinates backup and restore operations, controls job scheduling and manages resource allocation.

## Backup Appliance

The backup appliance is an architecture component that sits logically between the backup server and other components of the backup infrastructure. While the backup server administers tasks, the backup appliance performs management operations, processes jobs and delivers backup traffic. The backup appliance is a Linux-based VM that resides in the Nutanix AHV cluster and includes the following components:

- **Backup Appliance Web Console** is an interface that allows you to manage backup and restore operations, and to configure settings of the Nutanix AHV backup appliance.
- **Veeam Updater** is a service that is responsible for installing and scheduling updates for the backup appliance.

## Nutanix AHV Plug-in

Nutanix AHV Plug-in is an architecture component that enables integration between the backup server and the backup appliance. Nutanix AHV Plug-in also allows the backup server to deploy and manage the backup appliance.

## Backup Repositories


A backup repository is a storage location where Veeam Backup for Nutanix AHV stores backups of protected Nutanix AHV VMs.

To communicate with backup repositories, Veeam Backup for Nutanix AHV uses Veeam Data Mover – the service that is responsible for data processing and transfer. By default, Veeam Data Mover runs on the repositories themselves. If a repository cannot host Veeam Data Mover, it starts on a gateway server – a dedicated component that “bridges” the backup server and workers. For more information, see the Veeam Backup & Replication User Guide, section [Gateway Server](#).

## Workers

A worker is an auxiliary Linux-based VM instance that is responsible for the interaction between the backup appliance and other components of the backup infrastructure. Workers process backup workload and distribute backup traffic when transferring data to and from backup repositories.

The backup appliance comes with a preconfigured embedded worker that can be used in small virtual environments. In large environments, it is recommended to deploy dedicated workers that are distributed among the cluster hosts (nodes) and are automatically launched for the duration of a backup or restore process. For more information on deployment sizing considerations, see [Sizing Guidelines](#).



# VM Backup

To produce backups of VMs, Veeam Backup for Nutanix AHV runs backup jobs. A backup job is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, and so on.

While creating [image-level backups](#), Veeam Backup for Nutanix AHV does not install agent software inside VMs to retrieve data. Veeam Backup for Nutanix AHV uses [native Nutanix AHV capabilities](#) instead. During every backup session, Veeam Backup for Nutanix AHV creates a Nutanix AHV live snapshot of each VM added to a backup job. The snapshot is further used to create a VM backup.

## How to Protect VMs

1. Check [system requirements](#) and [account permissions](#).
2. [Add backup repositories](#).
3. [Connect the Nutanix AHV server](#).
4. [Deploy a Nutanix AHV backup appliance](#).
5. [Configure worker settings](#).
6. [Complete the New Backup Job wizard](#).

## How VM Backup Works

Veeam Backup for Nutanix AHV performs VM backup in the following way:

1. Connects to the Nutanix AHV server (Prism Central or Nutanix AHV cluster) over Nutanix REST API and creates a [backup snapshot](#) of the processed VM or a protection domain to which the VM belongs.
2. Launches a worker on the same host where the processed VM resides.  
  
If no worker is deployed on the host, Veeam Backup for Nutanix AHV launches a worker that is deployed on any other host connected to the Nutanix AHV server; if no workers are deployed on hosts, the [embedded worker](#) is used instead.
3. Re-creates VM disks from the snapshot created at step 1, adds them to a temporary volume group and attaches it to the worker.
4. Uses the worker to read data from disks of the volume group, transfers the data to the target repository and stores it in the native Veeam format.

To reduce the amount of data read from snapshots, Veeam Backup for Nutanix AHV uses the changed block tracking (CBT) mechanism: during incremental backup sessions, Veeam Backup for Nutanix AHV compares the new snapshot with the previous one and reads only those data blocks that have changed since the previous backup session. If CBT cannot be used, Veeam Backup for Nutanix AHV reads all data from the snapshot. For more information, see [Changed Block Tracking](#).

Veeam Backup for Nutanix AHV compresses and deduplicates data saved to repositories.

5. Suspends the worker when the backup session completes.

## NOTE

To limit the impact of backup tasks on network performance, Veeam Backup for Nutanix AHV applies [network traffic throttling rules](#) that prevent jobs from utilizing the entire bandwidth available in your environment.



# Backup Chain

When running a backup job, Veeam Backup for Nutanix AHV creates a new backup file in a backup repository during every backup session. A sequence of backup files created during a set of backup sessions makes up a backup chain. Each backup chain contains data for one VM only. If a backup job includes several VMs, Veeam Backup for Nutanix AHV creates one backup chain for each VM processed by the job.

The backup chain includes backup files of the following types:

- VBK – a full backup file stores a copy of the full VM image.
- VIB – incremental backup files store incremental changes of the VM image.
- VBM – backup metadata files store information about the backup job, VMs processed by the backup job, number and structure of backup files, restore points, and so on. Metadata files facilitate import of backups, backup mapping and other operations.

Full and incremental backup files act as restore points for backed-up VMs that let you roll back VM data to the necessary state. To recover a VM to a specific point in time, the chain of backup files created for the VM must contain a full backup file and a set of incremental backup files dependent on the full backup file.

If some file in the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual backup files from the backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backup files in the backup repository. For more information, see [Backup Retention](#).

## Changed Block Tracking

The changed block tracking (CBT) mechanism allows Veeam Backup for Nutanix AHV to increase the speed and efficiency of incremental backups:

- During a full backup session Veeam Backup for Nutanix AHV reads only written data blocks, while unallocated data blocks are filtered out.
- During an incremental backup session, Veeam Backup for Nutanix AHV reads only those data blocks that have changed since the previous backup session.

To detect unallocated and changed data blocks, CBT relies on the Nutanix AHV REST API:

1. During the first (full) backup session, Veeam Backup for Nutanix AHV creates a snapshot of a VM using native Nutanix AHV capabilities. To do that, Veeam Backup for Nutanix AHV sends API requests to access the content of the snapshot and to detect unallocated data blocks.
2. During subsequent sessions, new snapshots are created. Veeam Backup for Nutanix AHV sends API requests to access and to compare the content of the snapshot created during the previous backup session and the snapshot created during the current backup session. This allows Veeam Backup for Nutanix AHV to detect data blocks that have changed since the previous backup session.

## Limitations for Changed Block Tracking

Veeam Backup for Nutanix AHV does not use CBT for backup jobs which include a protection domain with consistency groups that contain two or more entities. If CBT cannot be used, Veeam Backup for Nutanix AHV reads the whole content of processed disks and compares it with backed-up data that already exists in the backup repository. In this case, the completion time of incremental backups may occur to grow.

# Backup Retention

For image-level backups, you can specify retention policy in days or restore points. Veeam Backup for Nutanix AHV retains the number of latest restore points defined in job scheduling settings as described in section [Creating Backup Jobs](#). For backup chains created by jobs without scheduled active or synthetic full backups, Veeam Backup for Nutanix AHV applies forever forward incremental backup retention policy. For backup chains created by jobs that regularly produce active or synthetic full backups, Veeam Backup for Nutanix AHV applies forward incremental backup retention policy.


## NOTE

For backup chains created by jobs that no longer exist, Veeam Backup for Nutanix AHV applies a separate retention mechanism as described in the Veeam Backup & Replication User Guide, section [Background Retention](#).

## Forever Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forever forward incremental backup chain, Veeam Backup for Nutanix AHV performs the following actions at the end of each backup session:


1. Veeam Backup for Nutanix AHV checks the configuration database to detect backup chains that contain redundant restore points:
  - If retention policy is specified in days, Veeam Backup for Nutanix AHV detects backup chains with restore points that are older than the specified time limit.
  - If retention policy is specified in restore points, Veeam Backup for Nutanix AHV detects backup chains where the number of allowed restore points is exceeded.
2. If a redundant restore point exists in a backup chain, Veeam Backup for Nutanix AHV transforms the backup chain in the following way:
  - a. Rebuilds the full backup to include the data of the incremental backup that follows the full backup. To do that, Veeam Backup for Nutanix AHV injects into the full backup data blocks from the earliest incremental backup in the chain. This way, the full backup 'moves' forward in the standard backup chain.



- b. Removes the earliest incremental backup from the chain as redundant – this data has already been injected into the full backup.




- Veeam Backup for Nutanix AHV repeats step 2 for all other redundant restore points found in the backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full backup, Veeam Backup for Nutanix AHV ensures that the backup chain is not broken and that you will be able to recover your data when needed.



## Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forward incremental backup chain, Veeam Backup for Nutanix AHV performs the following actions at the end of each backup session:

- Veeam Backup for Nutanix AHV checks the configuration database to detect forward incremental backup chains where a new full backup has been created (which starts a new backup chain fragment).
- Veeam Backup for Nutanix AHV checks the following:
  - If retention policy is specified in days, Veeam Backup for Nutanix AHV checks whether the period to keep restore points in the new chain fragment has reached the allowed time limit.
  - If retention policy is specified in restore points, Veeam Backup for Nutanix AHV checks whether the number of restore points in the new chain fragment has reached the number of allowed restore points.
- If the new backup chain fragment has reached the limit of allowed restore points, Veeam Backup for Nutanix AHV removes all restore points of the older backup chain fragment.



# Backup Methods

Veeam Backup for Nutanix AHV provides the following methods for creating backup chains:

- **Forever forward incremental**

When the forever forward incremental backup method is used, Veeam Backup for Nutanix AHV creates a backup chain that consists of the first full backup file (VBK) and a set of forward incremental backup files (VIBs) following it. For more information, see [Forever Forward Incremental Backup](#).

This backup method helps you save space on the backup storage because Veeam Backup for Nutanix AHV stores only one full backup file and removes incremental backup files [once the retention period is exceeded](#).

- **Forward incremental**


When the forward incremental backup method is used, Veeam Backup for Nutanix AHV creates a backup chain that consists of multiple full backup files (VBKs) and sets of forward incremental backup files (VIBs) following each full backup file. Full backups created using the synthetic full or active full method split the backup chain into shorter series. This lowers the chances of losing the backup chain completely and makes this backup method the most reliable. For more information, see [Forward Incremental Backup](#).

This backup method requires more storage space than other methods because the backup chains contains multiple full backup files and sometimes Veeam Backup for Nutanix AHV stores more restore points than specified in the retention policy settings due to the specifics of the [forward incremental retention policy](#).

## Forever Forward Incremental Backup

To create a backup chain for a VM protected by a backup job without a full backup schedule, Veeam Backup for Nutanix AHV implements the forever forward incremental backup:

1. During the first (full) backup session, Veeam Backup for Nutanix AHV copies the full VM image and creates a full backup file in the backup repository. The full backup file becomes a starting point in the backup chain.
2. During subsequent backup sessions, Veeam Backup for Nutanix AHV copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backup files in the backup repository. The content of each incremental backup file depends on the content of the full backup file and the preceding incremental backup files in the backup chain.




## Forward Incremental Backup


To create a backup chain for a VM protected by a backup job with scheduled full backups, Veeam Backup for Nutanix AHV implements the forward incremental backup method:

1. During the first (full) backup session, Veeam Backup for Nutanix AHV copies the full VM image and creates a full backup file in the backup repository. The full backup file becomes a starting point in the backup chain.

2. During subsequent backup sessions, Veeam Backup for Nutanix AHV copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backup files in the backup repository. The content of each incremental backup file depends on the content of the full backup file and the preceding incremental backup files in the backup chain.




3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup for Nutanix AHV creates a full backup file and adds it to the backup chain. Incremental restore points produced after this full backup file use it as a new starting point.



# Active Full Backup

In some cases, you may need to regularly create full backups. For example, your corporate backup policy may require that you create full backups on weekends and run incremental backups on work days. To let you conform to these requirements, Veeam Backup for Nutanix AHV allows you to create active full backups (either manually or automatically according to a specific schedule).

To create an active full backup Veeam Backup for Nutanix AHV retrieves VM data from the source cluster where the VM resides, compresses and deduplicates it and writes it to the VBK file in the backup repository. When creating an active full backup, Veeam Backup for Nutanix AHV starts a new backup chain for the VM. All further created incremental backups use the latest active full backup file as a new starting point. The old full backup file from the old backup chain remains on disk until it is automatically deleted according to the retention policy.



Veeam Backup for Nutanix AHV triggers a backup job to create an active full backup even if a regular backup session is not scheduled on this day. The active full backup session starts at the same time when the backup job is scheduled. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule active full backup to be created on Saturday, Veeam Backup for Nutanix AHV will start a backup job session that will produce an active full backup at 12:00 AM on Saturday.


If the backup job is not scheduled to run automatically or is disabled, Veeam Backup for Nutanix AHV will not perform active full backup. If a regular backup session and an active full backup session are scheduled on the same day, Veeam Backup for Nutanix AHV will produce an active full backup only. However, if you run the backup job again on the same day manually, Veeam Backup for Nutanix AHV will perform incremental backup in a regular manner.

# Synthetic Full Backup


In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups that also produce VBK files and contain data of the whole VM. However, while creating synthetic full backups, Veeam Backup for Nutanix AHV connects to the cluster to retrieve only VM data that has changed since recent backup and processes it with the data that is already stored in the backup repository.

To create a synthetic full backup, Veeam Backup for Nutanix AHV performs the following operations:


1. Veeam Backup for Nutanix AHV creates a regular incremental backup and adds it to the backup chain.



2. Veeam Backup for Nutanix AHV creates a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.



3. Veeam Backup for Nutanix AHV deletes the created incremental backup as its data is already incorporated in the synthetic full backup.



When creating a synthetic full backup, Veeam Backup for Nutanix AHV starts a new backup chain for the VM. All further created incremental backups use the latest full backup file as a new starting point. The old full backup file from the old backup chain remains on disk until it is automatically deleted according to the retention policy.

Veeam Backup for Nutanix AHV triggers a backup job to create a synthetic full backup even if a regular backup session is not scheduled on this day. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule synthetic full backup to be created on Saturday, Veeam Backup for Nutanix AHV will start a backup job session that will produce a synthetic full backup at 12:00 AM on Saturday.

If the backup job is not scheduled or is disabled, Veeam Backup for Nutanix AHV will not perform synthetic full backup automatically. If a regular backup session and a synthetic full backup session are scheduled on the same day, Veeam Backup for Nutanix AHV will produce a synthetic full backup only. However, if you run the backup job again on the same day manually, Veeam Backup for Nutanix AHV will perform incremental backup in a regular manner.

# Snapshot Chain


During every backup session, Veeam Backup for Nutanix AHV creates a snapshot for each VM added to the snapshot job or a protection domain snapshot if it is added to a protection domain snapshot job. The snapshot itself is a collection of point-in-time snapshots that Veeam Backup for Nutanix AHV takes using native Nutanix AHV capabilities.

A sequence of Nutanix AHV snapshots created during a set of backup sessions makes up a snapshot chain. Veeam Backup for Nutanix AHV creates the snapshot chain in the following way:

1. During the first backup session, Veeam Backup for Nutanix AHV creates a snapshot that contains all VM data and saves it in the Nutanix AHV cluster where the processed VM resides. This snapshot becomes a starting point in the snapshot chain.
2. During subsequent backup sessions, Veeam Backup for Nutanix AHV creates snapshots that contain only those data blocks that have changed since the previous backup session.

Each Nutanix AHV snapshot in the snapshot chain contains metadata. Metadata stores information about the protected instance and the job that created the snapshot. Veeam Backup for Nutanix AHV uses metadata to identify snapshots created by the Veeam backup service, to detect outdated snapshots, and to load the configuration of source VMs during recovery operations, and so on.

In Veeam Backup for Nutanix AHV, snapshots act as independent restore points for backed-up instances. If you remove any snapshot, it will not break the snapshot chain – Nutanix AHV will include data of the deleted snapshot to a following snapshot. Therefore, you will still be able to roll back instance data to any existing restore point. Note that while removing a protection domain snapshot, Veeam Backup for Nutanix AHV deletes data of all VMs and volume groups included in the domain.



The number of snapshots kept in the snapshot chain is defined by retention policy settings. For more information, see [Snapshot Retention](#).

## Related Topics

- [Creating Snapshot Jobs](#)
- [Creating Protection Domain Snapshot Jobs](#)

## Snapshot Types

In terms of data protection, Veeam Backup for Nutanix AHV allows you to create the following types of snapshots:

- **Backup snapshots**

A backup snapshot is a VM snapshot created by a [backup job](#). Backup snapshots are displayed both in the Nutanix AHV backup appliance web console and in the Veeam Backup & Replication console, and can be used to perform [entire VM restore](#) and [disk restore](#).

Backup snapshots allow Veeam Backup for Nutanix AHV to use the [CBT mechanism](#) while creating backups and to speed up the restore process (in comparison to restore from image-level backups).



- **Snapshots**

A snapshot is a VM snapshot created by a [snapshot job](#) or taken manually in the Prism Element or Prism Central console. Snapshots are displayed both in the Nutanix AHV backup appliance web console and in the Veeam Backup & Replication console. You can use snapshots to [restore VMs to the original Nutanix AHV environment](#).

While taking VM snapshots, Nutanix AHV captures data residing on virtual disks attached to the VMs. To protect data residing on volume groups that are attached to the VMs, volume group (VG) snapshots or protection domain (PD) snapshots are created. VG snapshots capture data of volume groups only, whereas PD snapshots capture data of consistency groups that include VMs and volume groups attached to them.

- **Snapshots on Replica Sites**

[Applies only to the [Prism Central deployment](#)] A snapshot on a replica site is a VM snapshot created and replicated by a Prism Central [protection policy](#). Snapshots on replica sites are displayed neither in the Nutanix AHV backup appliance web console nor in the Veeam Backup & Replication console – these snapshots can be only be found in the Prism Central console.

Snapshots on replica sites allow Veeam Backup for Nutanix AHV to reduce the backup load on the production environment. However, Veeam Backup for Nutanix AHV can use these snapshots only if the following requirements are met for each VM included into the backup scope:

- No volume groups are attached to the VM.
- At least one VM snapshot has been replicated to a remote location since the most recent backup was created.
- The VM disk configuration has not changed since the most recent snapshot was replicated to a remote location.
- Guest processing is disabled for the backup job.

If any of those conditions are not met, Veeam Backup for Nutanix AHV performs backup using [VM snapshots](#) created by backup jobs in the main site.

- **VG snapshots**

A VG snapshot is a volume group snapshot created by a [backup job](#) to produce VM backups. Veeam Backup for Nutanix AHV takes VG snapshots only if the backup scope includes individual virtual machines (not protection domains) with volume groups attached.

VG snapshots are displayed neither in the Nutanix AHV backup appliance web console nor in the Veeam Backup & Replication console. VG snapshots allow Veeam Backup for Nutanix AHV to use the [CBT mechanism](#) while creating backups and to [restore VMs with volume groups](#).

- **PD snapshots**

A PD snapshot is a protection domain snapshot created by a [PD snapshot job](#) or a [backup job](#) to protect data of consistency groups (VMs and volume groups) included into a protection domain. PD snapshots guarantee the consistency of VM and volume group data. Veeam Backup for Nutanix AHV takes PD snapshots only if [application-aware processing is disabled](#) in job settings and the following requirements are met for each protection domain included into the backup scope:

- The protection domain does not contain multiple VMs with the same name.
- VMs and their volume groups belong to the same protection domain.
- VMs and their volume groups are included into one [consistency group](#) of the protection domain.
- [CHAP authentication is disabled](#) for the volume groups.

- [Applies to backup jobs only] At least one consistency group of the protection domain contains more than one VM or volume group.
- [Applies to backup jobs only] Guest processing is disabled for the job.

PD snapshots are displayed both in the Nutanix AHV backup appliance web console and in the Veeam Backup & Replication console. You can use PD snapshots to [restore VMs to the original Nutanix AHV environment](#).

## NOTE

[Recovery points](#) created manually in the Prism Central console cannot be used to protect and recover Nutanix AHV resources with Veeam Backup for Nutanix AHV.

In terms of data consistency, Veeam Backup for Nutanix AHV allows you to create the following types of snapshots:

- **Crash-consistent snapshots**

A crash-consistent snapshot contains the data of virtual disks and volume groups attached to a VM.

- **Application-consistent snapshots**


An application-consistent snapshot contains not only the data of virtual disks and volume groups attached to a VM, but also the data of applications (such as Microsoft Active Directory, Microsoft SQL Server, Microsoft SharePoint, Microsoft Exchange and Oracle) running in the VM guest OS, which allows you to restore the applications without data loss and corruption.

By default, Veeam Backup for Nutanix AHV always tries to create an application-consistent snapshot using [Nutanix Guest Tools](#) when processing a VM. However, if the [requirements for application-consistent snapshots](#) are not met, Veeam Backup for Nutanix AHV creates a crash-consistent snapshot instead.

## Snapshot Retention

For Nutanix AHV native snapshots, you specify retention policy restore points. Veeam Backup for Nutanix AHV retains the number of latest restore points defined in job scheduling settings as described in sections [Creating Snapshot Jobs](#) and [Creating Protection Domain Snapshots Jobs](#).

During every successful backup session, Veeam Backup for Nutanix AHV creates a new restore point. If Veeam Backup for Nutanix AHV detects that the number of restore points in the snapshot chain exceeds the retention limit, it removes the earliest restore point from the chain.



# VM Restore

Veeam Backup for Nutanix AHV offers the following restore options:

- [Entire VM Restore](#) – restores an entire VM from a backup. You can restore one or more VMs at a time, to the original location or to a new location.
- [Disk Restore](#) – restores persistent disks attached to a VM from a snapshot or an image-level backup. You can restore persistent disks to the original location or to a new location.
- [File-level recovery](#) – recovers individual VM files and folders from a backup. You can download the necessary files and folders to a local machine, or restore the files and folders of the source VM to the original location.

You can restore VM data to the most recent state or to any available restore point.

# Entire VM Restore

To restore a VM, Veeam Backup for Nutanix AHV performs the following steps:

1. [Applies only if you perform restore to the original location where the source VM is still present] Connects to the Nutanix AHV server over REST API to power off and remove the source VM.
2. Launches a worker on same host where the processed VM resides.  
  
If no worker is deployed on the host, Veeam Backup for Nutanix AHV launches a worker that is deployed on any other host connected to the Nutanix AHV server; if no workers are deployed on hosts, the [embedded worker](#) is used instead.
3. Connects to the target Nutanix AHV server over REST API and creates a VM in the target location.
4. Creates empty virtual disks in the target location. The number of empty disks equals the number of disks attached to the source VM.
5. Connects to the backup repository and restores backed-up data to the empty disks.  
  
If multiple disks are attached to the source VM, the Nutanix AHV backup appliance restores these disks sequentially, one disk at a time.
6. Attaches the created disks with the restored data to the target VM disk nodes using their original bus.  
  
The maximum number of disk nodes available on Nutanix AHV VMs for each bus type is limited. Veeam Backup for Nutanix AHV can attach to a VM up to 6 SATA, 256 SCSI, 4 IDE and 7 PCI disks. If you restore a VM that has more disks of any of those bus types, Nutanix AHV will attach the disks to remaining nodes of other bus types in the default priority: SATA, SCSI, IDE, PCI. You can modify [the backup appliance configuration](#), to instruct Nutanix AHV to ignore original bus types and to use a specific order of bus types.
7. [Applies only if the VM has volume groups attached] Creates a new volume group with empty disks.
8. [Applies only if the VM has volume groups attached] Connects to the backup repository and restores backed-up data to the empty disks of the volume group.
9. [Applies only if the VM has volume groups attached and you perform restore to the original location where the source VM is still present] Removes the volume group that was attached to the source VM.
10. [Applies only if the VM has volume groups attached] Attaches the created volume group with the restored data to the target VM.

## NOTE

Veeam Backup for Nutanix AHV prioritizes restore tasks higher than other tasks. If multiple VMs are added to the restore session, these VMs are processed in parallel.

To learn how to restore an entire VM, see [Performing VM Restore](#).

# Disk Restore

To restore a VM disk, Veeam Backup for Nutanix AHV performs the following steps:

1. Connects to the Nutanix AHV server over REST API to power off the target VM.
2. Launches a worker on same host where the target VM resides.  
If no worker is deployed on the host, Veeam Backup for Nutanix AHV launches a worker that is deployed on any other host connected to the Nutanix AHV server; if no workers are deployed on hosts, the [embedded worker](#) is used instead.
3. Creates empty virtual disks in the Nutanix AHV infrastructure.
4. Connects to the backup repository and restores backed-up data to the empty disks.
5. [Applies only if you restore the disks to the original VM and if you choose to replace the existing disks] Detaches the original disks from the VM and removes them from the Nutanix AHV infrastructure.
6. Attaches the created disks with the restored data to the target VM.

To learn how to restore a VM disk, see [Performing Disk Restore](#).

# File-Level Recovery

To recover VM files and folders from a backup, Veeam Backup for Nutanix AHV performs the following steps:

1. [Applies only if you perform restore of a VM with an operation system other than Microsoft Windows] Deploys a [helper appliance](#) in the Nutanix AHV cluster.
2. Mounts disks of the VM to either of the following instances:
  - To the backup server or a [mount server](#) – if the VM guest OS is Microsoft Windows.
  - To the helper appliance – if the VM guest OS is a Linux-based operating system.
3. Launches the Veeam Backup Browser.

The Veeam Backup Browser displays the file system tree of the backed-up VM. In the browser, you select the necessary files and folders to restore.
4. Restores the selected files and folders to the original location or to a new location.
5. Detaches the disks from the backup server, mount server or helper appliance.
6. [Applies only if you perform restore of a VM with an operation system other than Microsoft Windows] Removes the helper appliance.

To recover VM files and folders from a backup snapshot, snapshot or PD snapshot, Veeam Backup for Nutanix AHV performs the following steps:

1. Deploys a [helper appliance](#) in the Nutanix AHV cluster.
2. [Applies only if you perform restore from a snapshot or PD snapshot] Creates a temporary VM on the Nutanix AHV cluster.
3. Creates a volume group using disks of the original VM (for a backup snapshot) or of the temporary VM (for a snapshot or PD snapshot).
4. Attaches the volume group to the helper appliance.
5. [Applies only if you perform restore from a snapshot or PD snapshot] Deletes the temporary VM.
6. Launches the Veeam Backup Browser.

The Veeam Backup Browser displays the file system tree of the backed-up VM. In the browser, you select the necessary files and folders to restore.
7. Restores the selected files and folders to the original location or to a new location.
8. Detaches the volume group from the helper appliance.
9. Deletes the volume group and removes the helper appliance.

To learn how to recover individual VM files and folders, see [Performing File-Level Restore](#).

# Retention Policies

Snapshots and image-level backups created by jobs are not kept forever – they are removed according to retention policy settings specified while creating the jobs.

Depending on the data protection scenario, retention policy can be specified:

- **In days** – for image-level backups.

Restore points in the backup chain can be stored only for the allowed period of time. If a restore point is older than the specified time limit, Veeam Backup for Nutanix AHV removes it from the backup chain. Since for retention policy specified in days, the backup chain must contain at least 3 restore points, Veeam Backup for Nutanix AHV may retain restore points for a longer period than configured in the retention policy settings. For more information, see [Backup Retention](#) and [Snapshot Retention](#).

- **In restore points** – for image-level backups and snapshots.

The chain can contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain. For more information, see [Backup Retention](#) and [Snapshot Retention](#).

# Planning and Preparation

Before you start deploying Veeam Backup for Nutanix AHV, check supported virtualization platforms, system requirements, permissions and network ports used for data transmission.



# System Requirements

Before you start deploying Veeam Backup for Nutanix AHV, make sure the Nutanix AHV cluster and the backup infrastructure components meet the following requirements.

Specification	Requirement
Virtualization Platform	<ul style="list-style-type: none"><li>• Veeam Backup for Nutanix AHV is compatible with Nutanix AOS versions 6.5.x - 6.7.x, 6.10.x, 7.0 and Prism Central version pc.2022.6 (or later).</li><li>• An IP address of the cluster and the iSCSI Data Service must be configured in Nutanix AHV cluster settings. For more information, see <a href="#">Nutanix documentation</a>.</li><li>• UEFI boot must be supported in the Nutanix AHV environment. For more information, see <a href="#">Nutanix documentation</a>.</li><li>• Veeam Backup for Nutanix AHV supports Nutanix Cloud Clusters (NC2) used for hybrid multi-cloud deployment.</li></ul>
Veeam Software	<ul style="list-style-type: none"><li>• Veeam Backup &amp; Replication version 12.3.0.310 (or later) with Nutanix AHV Plug-in version 12.7.0.172 (or later) must be deployed on the backup server.</li><li>• Microsoft .NET Core Runtime 8.0 and Microsoft ASP.NET Core Shared Framework 8.0 must be installed on the backup server.</li></ul>
Nutanix AHV Backup Appliance	<p>The Nutanix AHV backup appliance performs management operations and handles backup and restore tasks (if the <a href="#">embedded worker</a> is enabled). If you deploy Veeam Backup for Nutanix AHV using the default configuration, the following compute resources will be allocated to the Nutanix AHV backup appliance:</p> <ul style="list-style-type: none"><li>• CPU: 6 vCPU</li><li>• Memory: 6 GB RAM</li><li>• Disk Space: 100 GB for product installation, internal database files and logs</li></ul> <p>With the default configuration, the Nutanix AHV backup appliance can perform management operations in a cluster with less than 1000 VMs and process less than 100 jobs, and handle up to 4 concurrent backup and restore tasks. While deploying a new Nutanix AHV backup appliance or editing settings of an existing one, you can adjust compute resources allocated to it according to the recommendations described in section <a href="#">Sizing Guidelines</a>.</p>
Workers	<p><a href="#">Dedicated workers</a> process backup workload and distribute backup traffic when transferring data to and from backup repositories. If you deploy a worker using the default configuration, the following compute resources will be allocated to the worker VM:</p> <ul style="list-style-type: none"><li>• CPU: 6 vCPU</li><li>• Memory: 6 GB RAM</li><li>• Disk Space: 100 GB for product installation and logs</li></ul> <p>With the default configuration, the worker can handle up to 4 concurrent backup and restore tasks in parallel. While deploying a new worker or editing settings of an existing one, you can change the maximum number of concurrent tasks. To do that, adjust compute resources allocated to the worker VM according to the recommendations described in section <a href="#">Sizing Guidelines</a>.</p>

## **IMPORTANT**

The Nutanix AHV backup appliance and workers are deployed as backup infrastructure components preconfigured for optimal performance. That is why you must not install any software on VMs running as the Nutanix AHV backup appliance and workers or make any configuration changes to them unless you are requested by Veeam Customer Support.

# Permissions

The accounts used to deploy and administer backup infrastructure components must have the following permissions.

## Backup Server Windows Account

The Windows account used to install Veeam Backup & Replication and Nutanix AHV Plug-in on the backup server must have the following permissions.

Account	Required Permission
Setup Account	The account used to install Veeam Backup & Replication and Nutanix AHV Plug-in must have the Local Administrator permissions on the backup server.
Veeam Backup & Replication User Account	The account used to run Veeam Backup & Replication services must be a <i>LocalSystem</i> account or must have the Local Administrator permissions on the backup server.

## Nutanix AHV Cluster Administrator Account

The Nutanix AHV administrator account that Veeam Backup for Nutanix AHV uses to access the cluster must have privileges of the *Prism Admin* role or higher. For more information on user access control, see [Nutanix documentation](#).

## Performing Guest Processing

To use guest OS processing (application-aware processing, pre-freeze and post-thaw scripts, transaction log processing, guest file indexing and file exclusions), make sure to configure your accounts according to the requirements listed in this section. For more information on guest processing, see [Guest Processing](#).

All user accounts used for guest processing of Windows VMs must have the following permissions:

- *Logon as a batch job* granted.
- *Deny logon as a batch job* not set.

If Veeam Backup & Replication fails to use the **Log on as a batch job** policy, **Interactive Logon** is used.

Other permissions depend on applications that you back up. You can find permissions for backup operations in the following table. For restore operation permissions, see **Permissions** sections in the [Veeam Explorers User Guide](#).

Application	Required Permission
<p><b>Microsoft SQL Server</b></p>	<p>To back up Microsoft SQL Server data, the user whose account you plan to use must be:</p> <ul style="list-style-type: none"> <li>• Local Administrator on the target VM.</li> <li>• System administrator (has the <i>Sysadmin</i> role) on the target Microsoft SQL Server.</li> </ul> <p>If you need to provide minimal permissions, the account must be assigned the following roles and permissions:</p> <ul style="list-style-type: none"> <li>• SQL Server instance-level role: <i>public</i> and <i>dbcreator</i>.</li> <li>• Database-level roles and roles for the model system database: <i>db_backupoperator</i>, <i>db_denydatareader</i>, <i>public</i>; for the master system database – <i>db_backupoperator</i>, <i>db_datareader</i>, <i>public</i>; for the msdb system database – <i>db_backupoperator</i>, <i>db_datareader</i>, <i>public</i>, <i>db_datawriter</i>.</li> <li>• Securables: <i>view any definition</i>, <i>view server state</i>, <i>connect SQL</i>.</li> </ul> <p>If the account does not have enough rights, Veeam Backup &amp; Replication tries to truncate logs using the local SYSTEM account for Microsoft SQL Server 2008 and 2008 R2. For other Microsoft SQL Server versions, Veeam Backup &amp; Replication uses NT AUTHORITY\SYSTEM account.</p>
<p><b>Microsoft Active Directory</b></p>	<p>To back up Microsoft Active Directory data, the account must be a member of the built-in <i>Administrators</i> group.</p>
<p><b>Microsoft Exchange</b></p>	<p>To back up Microsoft Exchange data, the account must have the local Administrator permissions on the machine where Microsoft Exchange is installed.</p>

Application	Required Permission
Oracle	<p>The account specified at the <b>Guest Processing</b> step must be configured in the following way:</p> <ul style="list-style-type: none"> <li>For a Windows-based VM, the account must be a member of both the <i>Local Administrator</i> group and the <i>ORA_DBA</i> group (if OS authentication is used). In addition, if <i>ASM</i> is used, then such an account must be a member of the <i>ORA_ASMADMIN</i> group (for Oracle 12 and higher).</li> <li>For a Linux-based VM, the account must be a Linux user elevated to <i>root</i>. The account must have the <code>home</code> directory created.</li> </ul> <p>To back up Oracle databases, you can specify the following options at the <a href="#">Oracle</a> tab:</p> <ul style="list-style-type: none"> <li>Oracle account with SYSDBA privileges.</li> </ul> <p>You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges.</p> <ul style="list-style-type: none"> <li>Account specified for guest processing. That is, the <b>Use guest credentials</b> option selected.</li> </ul> <p>In this case, the account that was specified at the <b>Guest Processing</b> step must be a member of the <i>ORA_DBA</i> group for a Windows-based VM and <i>OSASM</i>, <i>OSDBA</i> and <i>OINSTALL</i> groups for a Linux-based VM.</p> <p>To perform guest processing for Oracle databases on Linux servers, make sure that the <code>/tmp</code> directory is mounted with the <code>exec</code> option. Otherwise, you will get an error with the permission denial.</p>
Microsoft SharePoint	<p>To back up Microsoft SharePoint server, the account must have the Farm Administrator role.</p> <p>To back up Microsoft SQL databases of the Microsoft SharePoint Server, the account must have the same privileges as that of <a href="#">Veeam Explorer for Microsoft SQL Server</a>.</p>
PostgreSQL	<p>The account specified at the <b>Guest Processing</b> step must be a Linux user elevated to <i>root</i>. The account must have the <code>home</code> directory created.</p> <p>To back up PostgreSQL instances, the account must have the superuser privileges for the PostgreSQL instance. For more information, see <a href="#">PostgreSQL documentation</a>.</p>

Consider the following general requirements when choosing a user account:

- [For guest OS file indexing] For Windows-based workloads, choose an account that has administrator privileges. For Linux-based workloads, choose an account of a root user or user elevated to root.
- [If you plan to use guest processing over network for workloads without listed applications] For Windows-based workloads, choose an account that has administrator privileges. For Linux-based workloads, choose an account of a root user or user elevated to root.
- When using Active Directory accounts, make sure to provide an account in the *DOMAIN|Username* format.
- When using local user accounts, make sure to provide an account in the *Username or HOST|Username* format.

- To process a Domain Controller server, make sure that you are using an account that is a member of the *DOMAIN\Administrators* group.
- To back up a Read-Only Domain controller, a delegated RODC administrator account is sufficient. For more information, see [Microsoft Docs](#).

# Ports

Veeam Backup for Nutanix AHV automatically creates firewall rules for the ports required to allow communication between the Nutanix AHV backup appliance, workers and the backup server.

## IMPORTANT

Some Linux distributions require manual configuration of firewall rules. For more information, see [this Veeam KB article](#).

## Backup Appliance

The following table describes network ports that must be opened to ensure proper communication of the Nutanix AHV backup appliance with other backup infrastructure components.

From	To	Protocol	Port	Notes
Workstation web browser	Nutanix AHV backup appliance	TCP/HTTPS	443	Used to access the Nutanix AHV backup appliance web console.
Nutanix AHV backup appliance	Nutanix REST API	TCP/HTTPS	9440	Used to communicate with Nutanix AHV REST API.
	Backup server	TCP	10006	Used to connect to Veeam Backup & Replication.
	Backup server	TCP	2500 to 3300	Default range of ports used for malware detection metadata transfer.
	Workers	TCP	19000	Used to communicate with workers.
	Nutanix AHV server (Cluster Virtual IP, Cluster iSCSI Data Services IP, Cluster CVM IPs)	TCP/iSCSI	3205, 3260	Used to access disks attached to Nutanix AHV VMs.
	Veeam backup repository (or <a href="#">gateway server</a> )	TCP	2500-3300	Default range of ports used as transmission channels for jobs and restore sessions. For every TCP connection that a job uses, one port from this range is assigned.
	Mail server	SMTP	25	Used to send email notifications. The port number can be changed.

From	To	Protocol	Port	Notes
	Rocky Linux repositories ( <i>mirrors.rockylinux.org,</i> <i>mirrors.fedoraproject.org,</i> <i>rockylinux.map.fastly.net</i> )	TCP/HTTP(S)	80 (443)	Used to get OS security updates, .NET Core updates and PostgreSQL update packages.  The listed mirror URLs are used to get actual URLs that will be used to obtain updates.
	Veeam Update Repository ( <i>repository.veeam.com</i> )  <a href="#">Amazon CloudFront</a> ( <i>cloudfront.net,</i> <i>amazonaws.com</i> )	TCP/HTTPS	443	Used to download Nutanix AHV backup appliance update packages.  <b>Note:</b> Veeam Update Repository uses the Amazon CloudFront service to distribute traffic when downloading product updates.
	Nginx repository ( <i>nginx.org/packages/,</i> <i>nginx.org/packages/keys/</i> )	TCP/HTTPS	443	Used to download Nginx packages required for Nutanix AHV backup appliance web console updates.

## Workers

The following table describes network ports that must be opened to ensure proper communication of workers with other backup infrastructure components.

From	To	Protocol	Port	Notes
Worker	Nutanix REST API	TCP/HTTPS	9440	Used to communicate with Nutanix AHV REST API.
	Backup server	TCP	10006	Used to connect to Veeam Backup & Replication.
	Backup server	TCP	2500 to 3300	Default range of ports used for malware detection metadata transfer.
	Backup appliance	TCP	19001	Used to communicate with the backup appliance.
	Nutanix AHV server (Cluster Virtual IP, Cluster iSCSI Data Services IP, Cluster CVM IPs)	TCP/iSCSI	3205, 3260	Used to access disks attached to Nutanix AHV VMs.



From	To	Protocol	Port	Notes
	Veeam backup repository (or <a href="#">gateway server</a> )	TCP	2500-3300	Default range of ports used as transmission channels for jobs and restore sessions. For every TCP connection that a job uses, one port from this range is assigned.
	Rocky Linux repositories ( <a href="#">mirrors.rockylinux.org</a> , <a href="#">mirrors.fedoraproject.org</a> , <a href="#">rockylinux.map.fastly.net</a> )	TCP/HTTP(S)	80 (443)	Used to get OS security updates, .NET Core updates and PostgreSQL update packages.  <b>Note:</b> The listed mirror URLs are used to get actual URLs that will be used to obtain updates.
	Veeam Update Repository ( <a href="#">repository.veeam.com</a> )  <a href="#">Amazon CloudFront</a> ( <a href="#">cloudfront.net</a> , <a href="#">amazonaws.com</a> )	TCP/HTTPS	443	Used to download Nutanix AHV backup appliance update packages.  <b>Note:</b> Veeam Update Repository uses the Amazon CloudFront service to distribute traffic when downloading product updates.
	Nginx repository ( <a href="#">nginx.org/packages/</a> , <a href="#">nginx.org/packages/keys/</a> )	TCP/HTTPS	443	Used to download Nginx update packages.

## Backup Server

The following table describes network ports that must be opened to ensure proper communication of the backup server with other backup infrastructure components.

From	To	Protocol	Port	Notes
Veeam Backup & Replication console  and Veeam ONE server	Backup server	TCP/HTTPS	8543	Used to communicate with the Platform Service REST API.
FLR helper appliance	Backup server	TCP	2500	Used to connect to the backup server during file-level restore.
Mount Service	Backup server	TCP	9401	Used to connect to the backup server during file-level restore.

From	To	Protocol	Port	Notes
Backup server	FLR helper appliance	TCP	22 2500	Used to connect to the helper appliance during file-level restore. For the full list of ports used for connections to the FLR helper appliance, see the Veeam Backup & Replication User Guide, section <a href="#">Used Ports</a> .
	Backup server	TCP/HTTPS	6172	Used by the AHV Platform Service to enable communication with the Veeam Backup & Replication database.
	Nutanix AHV cluster	TCP/HTTPS	9440	Used by the AHV Platform Service to connect to an Nutanix AHV cluster.
	Nutanix AHV backup appliance	TCP/HTTPS	443	Used by the AHV Platform Service to connect to Nutanix AHV backup appliance.

#### NOTE

For the list of ports used by the backup server to communicate with backup repositories, see the Veeam Backup & Replication User Guide, section [Used Ports](#).

## vPower NFS Service

The vPower NFS Service is a Microsoft Windows service that runs on a Microsoft Windows machine and enables this machine to act as an NFS server. The vPower NFS Service is required to perform such operations as file-level restore and Instant Recovery.

#### NOTE

For the full list of ports required for [Performing File-Level Restore](#), see the Veeam Backup & Replication User Guide, section [Used Ports](#).

From	To	Protocol	Port	Notes
Nutanix AHV cluster	Microsoft Windows server with	TCP UDP	111	Used by the Port Mapper service.

From	To	Protocol	Port	Notes
	the mount server role running vPower NFS Service	TCP UDP	1058+ or 1063+	Used as default mount port. The number of port depends on where the vPower NFS Service is located: <ul style="list-style-type: none"> <li>• 1058+: If the vPower NFS Service is located on the backup server.</li> <li>• 1063+: If the vPower NFS Service is located on a separate Microsoft Windows machine.</li> </ul> If port 1058/1063 is occupied, the succeeding port numbers will be used.
		TCP UDP	2049+	Used as NFS port. If port 2049 is occupied, the succeeding port numbers will be used.

## Guest Processing Components

The following tables describe network ports that must be opened to ensure proper communication of the backup server and backup infrastructure components with the non-persistent runtime components deployed inside the VM guest OS for application-aware processing and indexing.

From	To	Protocol	Port	Notes
Backup server	VM guest OS (Linux)	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.
	Guest interaction proxy	TCP	6190	Used for communication with the guest interaction proxy.
		TCP	6290	Used as a control channel for communication with the guest interaction proxy.
		TCP	445	Port used as a transmission channel.
Guest interaction proxy	VM guest OS (Microsoft Windows)	TCP	445 135	Required to deploy the runtime coordination process on the VM guest OS.
		TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.

From	To	Protocol	Port	Notes
		TCP	49152 to 65535	<p>Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see <a href="#">this Microsoft KB article</a>.</p> <p>Used by the runtime process deployed inside the VM for guest OS interaction.</p> <p><b>Note:</b> If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup &amp; Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the <i>"RPC function call failed"</i> error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see <a href="#">this Microsoft KB article</a>.</p>
VM guest OS	Guest interaction proxy or backup server	TCP	2500 to 3300	Default range of ports used as transmission channels for log shipping.

## Log Shipping Components

The following tables describe network ports that must be opened to ensure proper communication between log shipping components.

- [Log Shipping Server Connections](#)
- [MS SQL Guest OS Connections](#)
- [Oracle Guest OS Connections](#)
- [PostgreSQL Guest OS Connections](#)

### Log Shipping Server Connections

From	To	Protocol	Port	Notes
Backup server		TCP	445 135	Required for deploying Veeam Backup & Replication components.

From	To	Protocol	Port	Notes
	Log shipping server	TCP	6160	Default port used by Veeam Installer Service.
		TCP	6162	Default port used by Veeam Data Mover Service.
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see <a href="#">this Microsoft KB article</a> .  <b>Note:</b> If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the <i>"RPC function call failed"</i> error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see <a href="#">this Microsoft KB article</a> .
Log shipping server	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups.

## MS SQL Guest OS Connections

From	To	Protocol	Port	Notes
Guest interaction proxy	MS SQL VM guest OS	TCP	445 135	Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.

From	To	Protocol	Port	Notes
		TCP	49152 to 65535	Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see <a href="#">this Microsoft KB article</a> .  <b>Note:</b> If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup & Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the <i>"RPC function call failed"</i> error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see <a href="#">this Microsoft KB article</a> .
		TCP	6167	Used by the Veeam Log Shipping Service for preparing the database and taking logs.
MS SQL VM guest OS	Guest interaction proxy	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.
MS SQL VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the MS SQL server has a direct connection to the backup repository.
MS SQL VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.

## Oracle Guest OS Connections

From	To	Protocol	Port	Notes
Guest interaction proxy	Oracle VM guest OS (Microsoft Windows)	TCP	445 135	Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.

From	To	Protocol	Port	Notes
		TCP	49152 to 65535	<p>Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see <a href="#">this Microsoft KB article</a>.</p> <p><b>Note:</b> If you use default Microsoft Windows firewall settings, you do not need to configure dynamic RPC ports. During setup, Veeam Backup &amp; Replication automatically creates a firewall rule for the runtime process. If you use firewall settings other than default ones or application-aware processing fails with the <i>"RPC function call failed"</i> error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see <a href="#">this Microsoft KB article</a>.</p>
		TCP	6167	Used by the Veeam Log Shipping Service for preparing the database and taking logs.
Backup server	Oracle VM guest OS (Linux)	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
Oracle VM guest OS	Guest interaction proxy or backup server	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.
Oracle VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the Oracle server has a direct connection to the backup repository.
Oracle VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.

## PostgreSQL Guest OS Connections

From	To	Protocol	Port	Notes
Backup server	PostgreSQL VM guest OS	TCP	22	Default SSH port used as a control channel.
		TCP	2500 to 3300	Default range of ports used for communication with a guest OS.
PostgreSQL VM guest OS	Backup server	TCP	2500 to 3300	Default range of ports used for communication with a guest interaction proxy.
PostgreSQL VM guest OS	Backup repository	TCP	2500 to 3300	Default range of ports used for communication with a backup repository and transfer log backups. Should be opened if log shipping servers are not used in the infrastructure and the PostgreSQL server has a direct connection to the backup repository.
PostgreSQL VM guest OS	Log shipping server	TCP	2500 to 3300	Default range of ports used for communication with a log shipping server and transfer log backups.



# Sizing Guidelines

This section is intended for professionals who search for a best practice answer to sizing -related issues, and assumes you have already read the whole Veeam Backup for Nutanix AHV User Guide.

Be aware that a best practice is not the only answer available. It will fit in the majority of cases but can also be totally wrong under different circumstances. Make sure you understand the implications of the recommended practices, or request assistance. If in doubt, reach out to Veeam professionals on [Veeam R&D Forums](#).

## Backup Appliance

While allocating compute resources to the Nutanix AHV backup appliance, consider the following recommendations:

Number of VMs in Cluster	Number of Jobs	Resources for Management Tasks	Resources for Embedded Worker
< 1000	< 100	2 vCPU, 2 GB RAM	1 vCPU and 1 GB RAM for each concurrent task
> 1000	>100	2 vCPU, 2 GB RAM	It is not recommended to use the embedded worker for large deployments. <a href="#">Disable the embedded worker</a> and use dedicated workers instead.
> 5000	> 300	3 vCPU, 3 GB RAM	

If you plan to use dedicated workers, it is recommended that you modify the backup appliance settings as follows:

- Set the maximum number of concurrent tasks to 0.
- Allocate resources to the backup appliance that are sufficient for management tasks.

### IMPORTANT

To modify the backup appliance settings, use the Veeam Backup & Replication console as described in section [Editing Backup Appliance](#). Allocating resources to the VM running as the backup appliance in the Nutanix Prism console may cause technical issues.

## Workers

While adding a dedicated worker to the backup infrastructure, consider the following:

- [Applies only to the [Prism Central deployment](#)] It is recommended that workers are deployed in each cluster registered with the Prism Central. If no worker is deployed in the cluster, performance of backup operations will be affected as Veeam Backup for Nutanix AHV will use the embedded worker or a worker deployed in another cluster.
- It is recommended that the number of configured workers does not exceed the number of hosts in the Nutanix AHV cluster.

- Each worker must be provided with sufficient compute resources to handle backup and restore tasks in parallel. The maximum number of concurrent tasks is configured in worker settings – if this number is exceeded, the worker will not start a new task until one of the current tasks finishes.
- It is recommended the total number of concurrent tasks configured for all workers deployed in the cluster does not exceed the [number of physical disks added to the cluster](#). You can change the maximum number of concurrent tasks (the best practice is to allocate 1 vCPU and 1 GB RAM for each additional task) while deploying a new worker or editing settings of an existing one.

### IMPORTANT

To modify the worker settings, use the Veeam Backup for Nutanix AHV web console as described in section [Disabling Automatic Worker Updates](#). Allocating resources to the VM running as a worker in the Nutanix Prism console may cause technical issues.

# Licensing

Veeam Backup for Nutanix AHV is licensed by the number of protected Nutanix AHV VMs. Each Nutanix AHV VM protected with backups consumes one Veeam Universal License instance from the license scope. A Nutanix AHV VM is considered protected if it has a restore point created during the past 31 days. If a Nutanix AHV VM is protected with snapshots only, no license is consumed.

By default, Veeam Backup for Nutanix AHV automatically revokes a license instance from a protected VM if no new restore points have been created during the past 31 days. However, you can manually revoke license instances from protected VMs as described in the Veeam Backup & Replication User Guide, section [Revoking License](#).

## Obtaining New License

You can obtain the following types of licenses for Veeam Backup for Nutanix AHV:

- **Evaluation license** is a free license that can be used for product evaluation. The license is valid for 30 days from the moment of the product download.

To obtain this license, request a trial key on the [Veeam downloads page](#) as described in the Veeam Backup & Replication User Guide, section [Obtaining and Renewing License](#).

- **Subscription license** is a paid license with a limited subscription term. The expiration date of the Subscription license is set to the end of the subscription term. The Subscription license term is normally 1-5 years from the license issue date.

To obtain this license, choose the required subscription term on the [Veeam Backup & Replication Pricing](#) page and contact the Veeam Sales Team.

- **Perpetual license** is a paid license without an expiration date. The Perpetual license typically includes one year period of basic support and maintenance that can be extended.

To obtain this license, [contact a reseller in your region](#).

After you obtain a license, install it on the backup server as described in the Veeam Backup & Replication User Guide, section [Installing License](#).

## Using Existing License

If you already use Veeam Backup & Replication and you have spare Veeam Universal License instances on your backup server, they can be used to protect Nutanix AHV VMs. You can check the number of available license instances in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section [Viewing License Information](#).

If you have a legacy perpetual per-socket license, you must obtain Veeam Universal License instances and merge them with the existing perpetual socket license as described in the Veeam Backup & Replication User Guide, section [Merging Licenses](#).

# Deployment

Starting from version 12.2, the Veeam Backup & Replication solution comes with a plug-in that allows you to add Nutanix AHV servers to the backup infrastructure, and to manage data protection and recovery operations for Nutanix AHV workloads from a single console.

To access the Nutanix AHV Plug-in functionality, you can either deploy a new backup server as described in the [Veeam Backup & Replication User Guide](#) or use a backup server that already exists in your backup infrastructure if it meets the [Veeam Backup for Nutanix AHV system requirements](#).

## Related Topics

- [Installing Nutanix AHV Plug-in Manually](#)
- [Uninstalling Nutanix AHV Plug-in Manually](#)

# Installing Nutanix AHV Plug-In Manually

The plug-in that allows you to protect Nutanix AHV resources comes pre-installed with the default installation package of Veeam Backup & Replication. However, you may require to install a new plug-in version on the backup server manually if some updates or patches become available.

## NOTE

If you use a remote Veeam Backup & Replication console, you do not need to install Nutanix AHV Plug-in on the workstation where the remote Veeam Backup & Replication console is deployed.

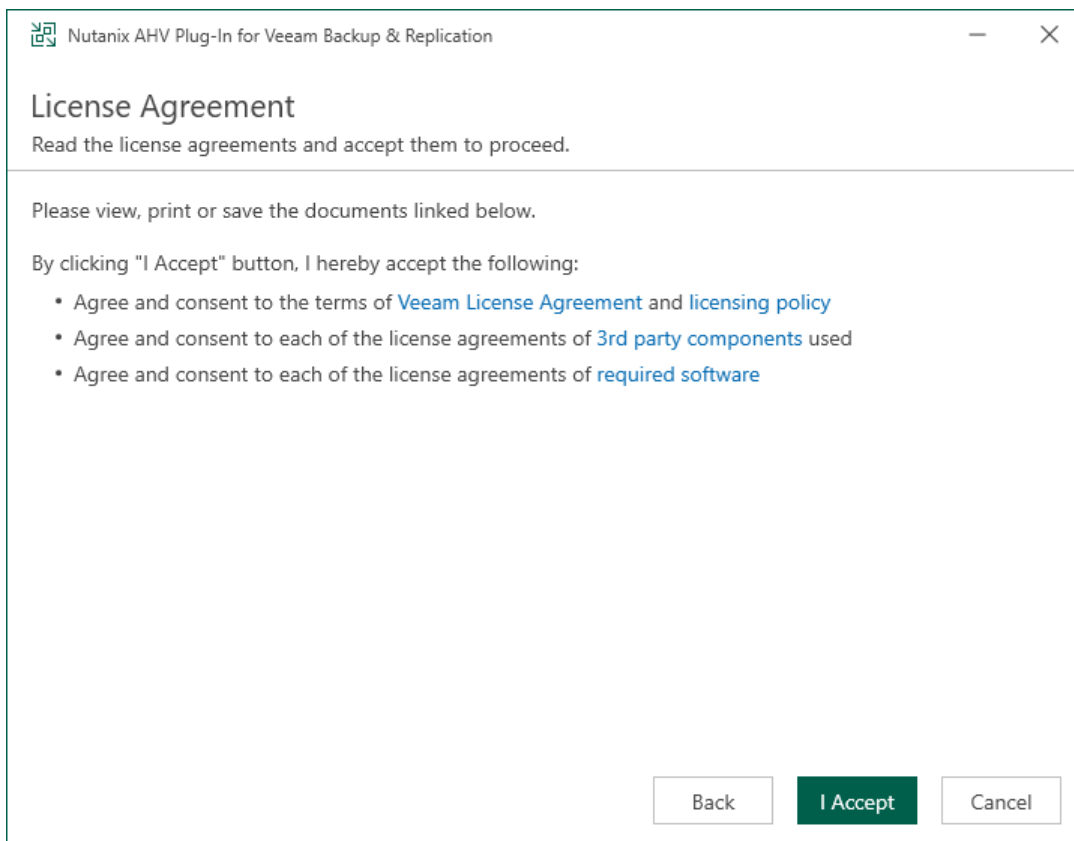
To install Nutanix AHV Plug-in, do the following:

1. Log in to the backup server using an account with the local Administrator permissions.
2. Download the product installation file `NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.zip` from the [Veeam downloads page](#).
3. Open the downloaded archive file and launch the `NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.exe` installation file.

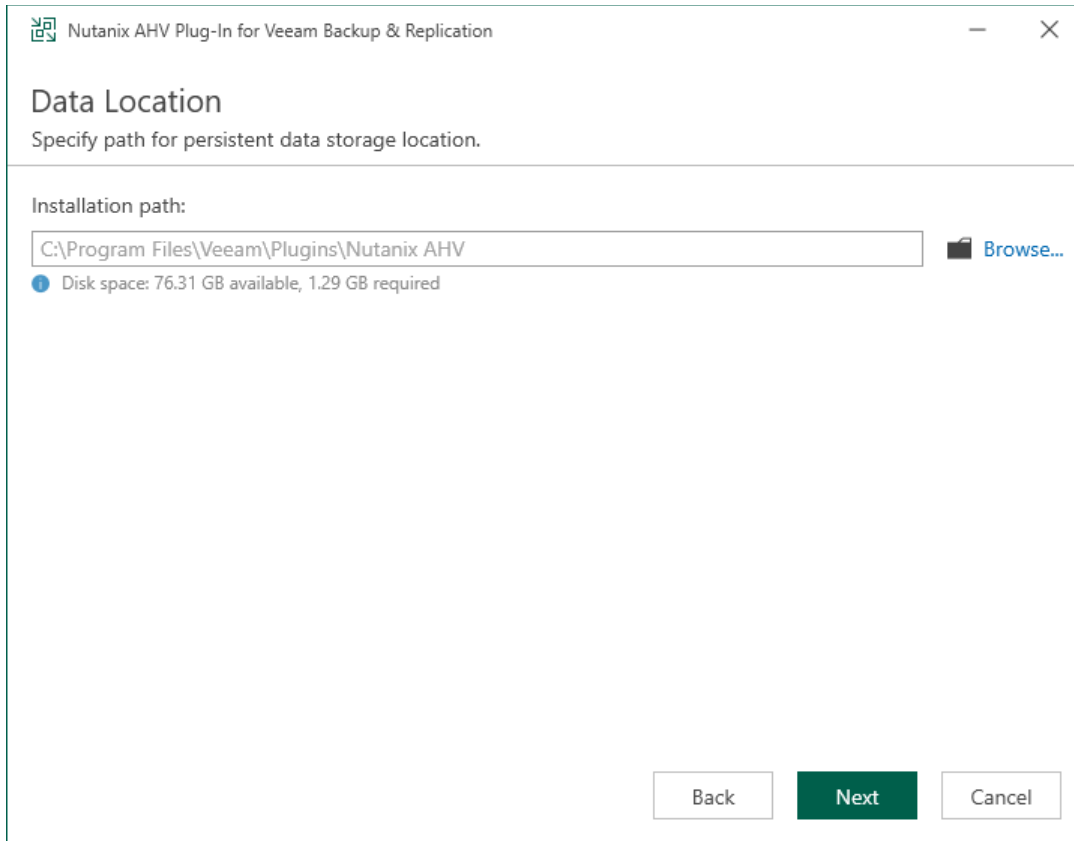
Before proceeding with installation, the installer will check whether you have Microsoft .NET Core Runtime installed on the backup server. In case the required version is missing, the installer will offer to install it automatically. To do that, click **OK**.

4. At the **License Agreement** step of the **Nutanix AHV Plug-In for Veeam Backup & Replication Setup** wizard, read and accept the Veeam license agreement, licensing policy, the 3rd party components and required software license agreement. If you reject the agreements, you will not be able to continue installation.

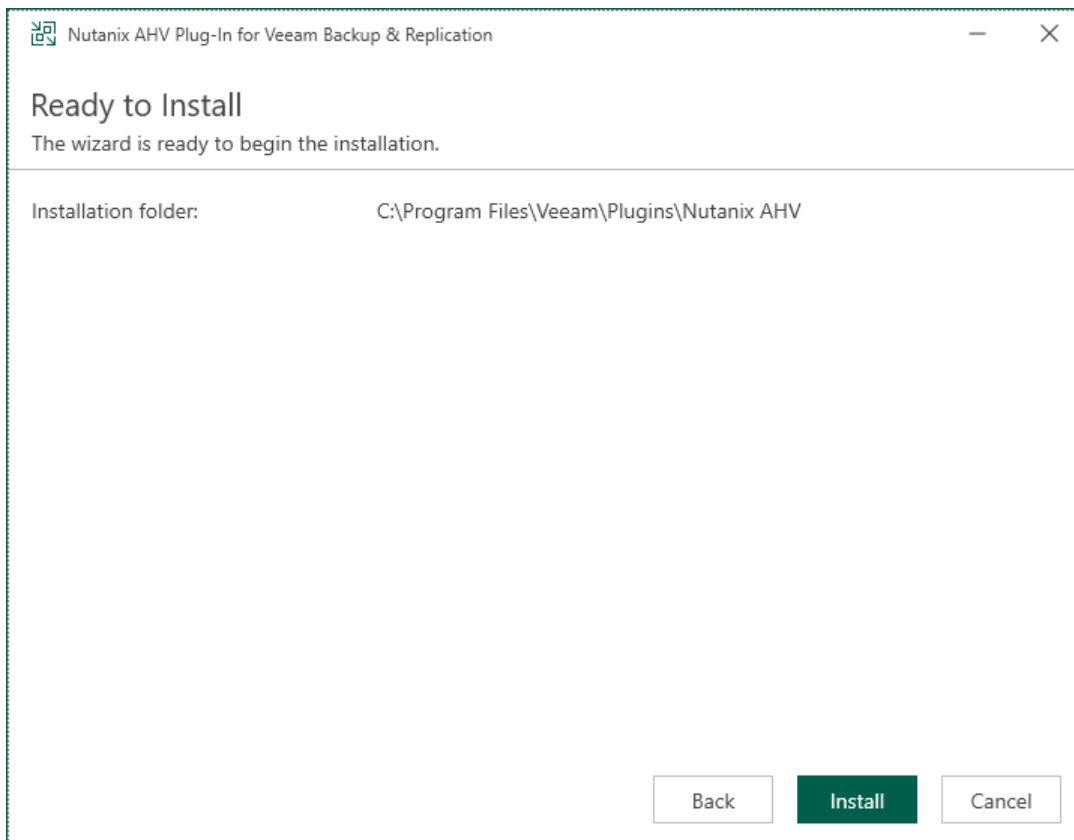
To read the terms of a license agreement, click **View**.



- At the **Installation path** step of the wizard, you can change the installation directory if necessary.



- Click **Install** to begin installation.



# Installing Plug-In in Unattended Mode

You can install Nutanix AHV Plug-in in the unattended mode using the command line interface. The unattended installation mode does not require user interaction – the installation runs automatically in the background, and you do not have to respond to the installation wizard prompts. You can use the unattended installation mode to automate the Nutanix AHV Plug-in installation process in large-scale environments.

To install Nutanix AHV Plug-in in the unattended mode, use either of the following options:

- If Nutanix AHV Plug-in is a part of Veeam Backup & Replication installation package, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Installing Veeam Backup & Replication in Unattended Mode](#).
- If Nutanix AHV Plug-in is delivered as a separate .EXE file, follow the instructions provided in this section.

## Before You Begin

Before you start unattended installation, do the following:

1. Download the `NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.EXE` file as described in section [Installing Nutanix AHV Plug-In Manually](#) (steps 1-3).
2. Check compatibility of the Nutanix AHV Plug-in and Veeam Backup & Replication versions. For more information, see [System Requirements](#).

## Installation Command-Line Syntax

Open the command prompt and run the .EXE file using the following parameters:

```
%path% /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware
```

The following command-line parameters are used to run the setup file:

Parameter	Required	Description
<code>%path%</code>	Yes	Specifies a path to the installation .EXE file on the backup server or in a network shared folder.
<code>/silent</code>	Yes	Sets the user interface level to <i>None</i> , which means no user interaction is needed during installation.
<code>/accepteula</code>	Yes	Confirms that you accept the terms of the Veeam license agreement.
<code>/acceptthirdpartylicenses</code>	Yes	Confirms that you accept the license agreement for 3rd party components that Veeam incorporates.

Parameter	Required	Description
/acceptlicensingpolicy	Yes	Confirms that you accept the Veeam licensing policy.
/acceptrequiredsoftware	Yes	Confirms that you accept the license agreements for each required software that Veeam will install.
/uninstall	No	Uninstalls the plug-in.
/repair	No	Replaces missing files and firewall rules.

## Examples

The following command installs Nutanix AHV Plug-in:

```
NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware
```

The following command repairs Nutanix AHV Plug-in:

```
NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware /repair
```

The following command uninstalls Nutanix AHV Plug-in:

```
NutanixAHVPlugin_12.7.0.172_AHV_7.0.0.241.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware /uninstall
```

Veeam Backup for Nutanix AHV provides the following status codes to report about the installation result:

Code	Description
0	Nutanix AHV Plug-in installation has successfully completed.
1603	Nutanix AHV Plug-in installation has failed.
3010	Nutanix AHV Plug-in installation has successfully completed. The backup server requires rebooting.



## TIP

For detailed logs of Nutanix AHV Plug-in installation, navigate to the `Program Data\Veeam\Setup\Temp\` folder on the backup server and view the following files:

- `VeeamPluginBootstrap.log`
- `NutanixAHVPlugin.log`
- `NutanixAHVPluginUI.log`
- `NutanixAHVPluginProxy.log`

# Upgrading to Veeam Backup for Nutanix AHV 7

You can upgrade Veeam Backup for Nutanix AHV from version 4.0, 4a, 5.0, 5.1, 6 and 6.1 to 7.

Before you start the upgrade process, do the following:

- Download Veeam Backup & Replication version 12.3 from the [Veeam downloads page](#).
- Plan a maintenance period. Typically, the upgrade process takes up to one hour. Make sure there are no jobs currently running or scheduled to run during this period. Wait for the jobs to complete or disable the jobs manually before you start upgrading Veeam Backup for Nutanix AHV.
- Make sure the Nutanix AHV backup appliance is powered on.
- Back up the configuration database of the Nutanix AHV backup appliance. For more information, see the following sections:
  - For Veeam Backup for Nutanix AHV version 4.0 and 4a, see Backup for Nutanix AHV 4.0 User Guide, section [Backing Up Configuration Settings Manually](#).
  - For Veeam Backup for Nutanix AHV version 5.0 and 5.1, see Backup for Nutanix AHV 5 User Guide, section [Performing Configuration Backup Manually](#).
  - For Veeam Backup for Nutanix AHV version 6 and 6.1, see Backup for Nutanix AHV 6 User Guide, section [Performing Configuration Backup Manually](#).

To upgrade Veeam Backup for Nutanix AHV to version 7, do the following:

1. Upgrade your Veeam Backup & Replication server to version 12.3 as described in the Veeam Backup & Replication User Guide, section [Upgrading to Veeam Backup & Replication 12](#). Then, complete the **Components Update** wizard as described in the Veeam Backup & Replication User Guide, section [Server Components Upgrade](#).

Veeam Backup for Nutanix AHV will be upgraded to version 7.

2. [Applies only to the [Prism Central deployment](#)] If Nutanix AHV clusters are registered with a Prism Central, add the Prism Central to the backup infrastructure as described in section [Adding Nutanix AHV Server](#).

While connecting the Prism Central, Veeam Backup for Nutanix AHV will prompt you to select a backup appliance that will be upgraded and automatically configured to manage all the jobs that are currently managed by dedicated backup appliances in the clusters.

3. [Applies only to the [standalone cluster deployment](#)] Upgrade the Nutanix AHV backup appliance to version 7. To do that:
  - a. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
  - b. Navigate to **Backup Proxies > Out of Date**.
  - c. Select the Nutanix AHV backup appliance and click **Upgrade Proxy** on the ribbon.
  - d. In the **Components Update** window, click **Apply**.

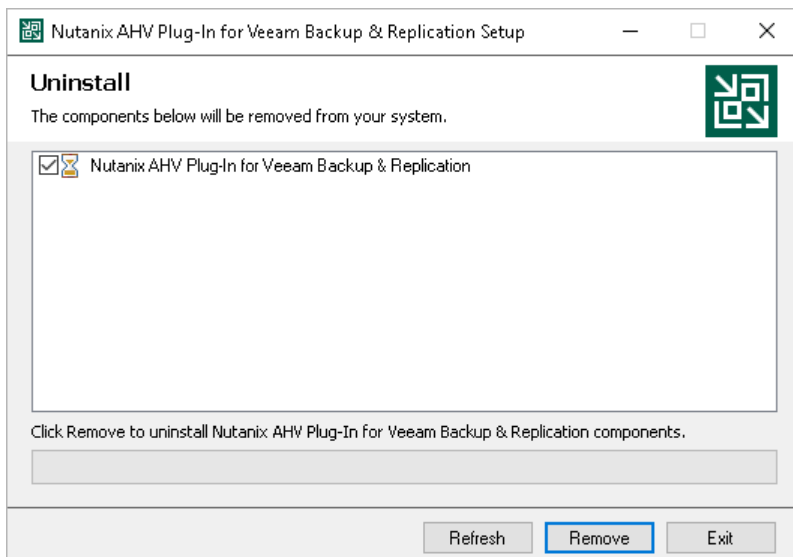
# Uninstalling Nutanix AHV Plug-In

Before you uninstall Nutanix AHV Plug-in, it is recommended to [remove all connected backup appliances](#) from the backup infrastructure. If you keep the backup appliances in the backup infrastructure, the following will happen:

- You will be able to see information on snapshots in the Veeam Backup & Replication console. However, you will not be able to perform any operations with these snapshots.
- You will be able to see information on image-level backups of VMs and perform data recovery operations using these backups. However, you will not be able to perform entire VM restore to Nutanix AHV.
- You will be able to see information on jobs. However, you will only be able to remove these jobs from the Veeam Backup & Replication console.

To uninstall Nutanix AHV Plug-in, do the following:

1. Log in to the backup server using an account with the Local Administrator permissions.
2. Open the **Start** menu and click the **Control Panel** icon.
3. In the **Settings** window, navigate to **System > Apps and Features**.
4. In the program list, select **Nutanix AHV Plug-in for Veeam Backup & Replication**. Then, click **Uninstall**.
5. In the opened window, click **Remove**.



# Configuring Backup Infrastructure

To start working with Veeam Backup for Nutanix AHV, perform a number of steps for its configuration:

1. [Configure backup repositories](#) where Veeam Backup for Nutanix AHV will store backups of Nutanix AHV VMs.
2. [Add to the backup infrastructure the Nutanix AHV cluster or Prism Central](#) that administers Nutanix AHV resources you want to protect.
3. [Deploy a Nutanix AHV backup appliance](#) that will process backup and restore operations.
4. [Deploy workers](#) that will transfer backup traffic.

# Configuring Backup Repositories

A backup repository is a storage location where Veeam Backup for Nutanix AHV keeps backup files. By default, the backup server performs the role of a backup repository. To keep your backups in another storage location, you can configure the following types of repositories:

- **Direct attached storage:** [Microsoft Windows](#) and [Linux](#) virtual and physical machines. [Hardened repositories](#) based on Linux servers are supported.
- **Network attached storage:** [CIFS \(SMB\) shares](#) and [NFS shares](#).
- **Deduplicating storage appliances:** [ExaGrid](#), [Quantum DXi](#), [Dell Data Domain](#), [HPE StoreOnce](#), [Fujitsu ETERNUS](#), [Infinidat InfiniGuard](#).
- **Cloud object storage:** [11:11 Cloud Object Storage](#), [Amazon S3](#), [S3 compatible](#), [Google Cloud](#), [Wasabi Cloud Storage](#), [Veeam Data Cloud Vault](#), [IBM Cloud](#) and [Microsoft Azure Blob](#).

To combine repositories of different types in one repository, you can configure a [scale-out backup repository](#) and add any of supported repositories to its [performance tier](#).

For Linux server, Microsoft Windows server, SMB share, ExaGrid, Quantum DXi, Fujitsu ETERNUS and Infinidat InfiniGuard repositories, you can enable the Fast Clone technology that increases the speed of synthetic backup creation and transformation, reduces disk space requirements and decreases the load on storage devices. With this technology, Veeam Backup for Nutanix AHV references existing data blocks on volumes instead of copying data blocks between files. Data blocks are copied only when files are modified. To learn how to configure a repository to enable this functionality, see the Veeam Backup & Replication User Guide, section [Fast Clone](#).

## IMPORTANT

- Veeam Backup for Nutanix AHV does not support storing backups in [Veeam Cloud Connect](#) and [HPE Cloud Bank Storage](#) repositories. However, you can use them for [storing copies of backups](#) created with Veeam Backup for Nutanix AHV.
- [For scale-out backup repositories] Due to specifics of backup jobs for Nutanix AHV VMs, Veeam Backup for Nutanix AHV always creates a separate backup chain for each VM added to a backup job. Thus, even if you clear the **Use per-VM backup files** check box in the [advanced settings of a scale-out backup repository](#), backups of multiple Nutanix AHV VMs are not stored in a single backup file.

# Connecting Nutanix AHV Server

The Nutanix AHV server (a Prism Central or standalone cluster) allows the backup server to access Nutanix AHV resources such as VMs, storage containers and networks. After you add the Nutanix AHV cluster to the backup infrastructure, you will be able to deploy a Nutanix AHV backup appliance and to manage data protection tasks for Nutanix AHV VMs and protection domains.

# Adding Nutanix AHV Server

To add a Nutanix AHV cluster or Prism Central to the backup infrastructure, do the following:

1. [Launch the New Nutanix AHV Server wizard.](#)
2. [Specify the Nutanix AHV server domain name or IP address.](#)
3. [Enter credentials to access the Nutanix AHV cluster.](#)
4. [Configure helper appliance settings.](#)
5. [Choose a backup appliance.](#)
6. [Apply Nutanix AHV server settings.](#)
7. [Finish working with the wizard.](#)

## Considerations and Limitations

After you add a Prism Central to the backup infrastructure, consider the following:

- If you register a new cluster with the Prism Central, Veeam Backup for Nutanix AHV will automatically add it to the backup infrastructure and you will be able to protect resources in this cluster. For more information, see sections [Performing Backup](#) and [Performing Restore](#).
- If you unregister an existing cluster from the Prism Central, you will not be able to protect resources in this cluster anymore. To protect these resources, you can add the cluster to the backup infrastructure as a standalone cluster. For more information, see [Solution Architecture](#).

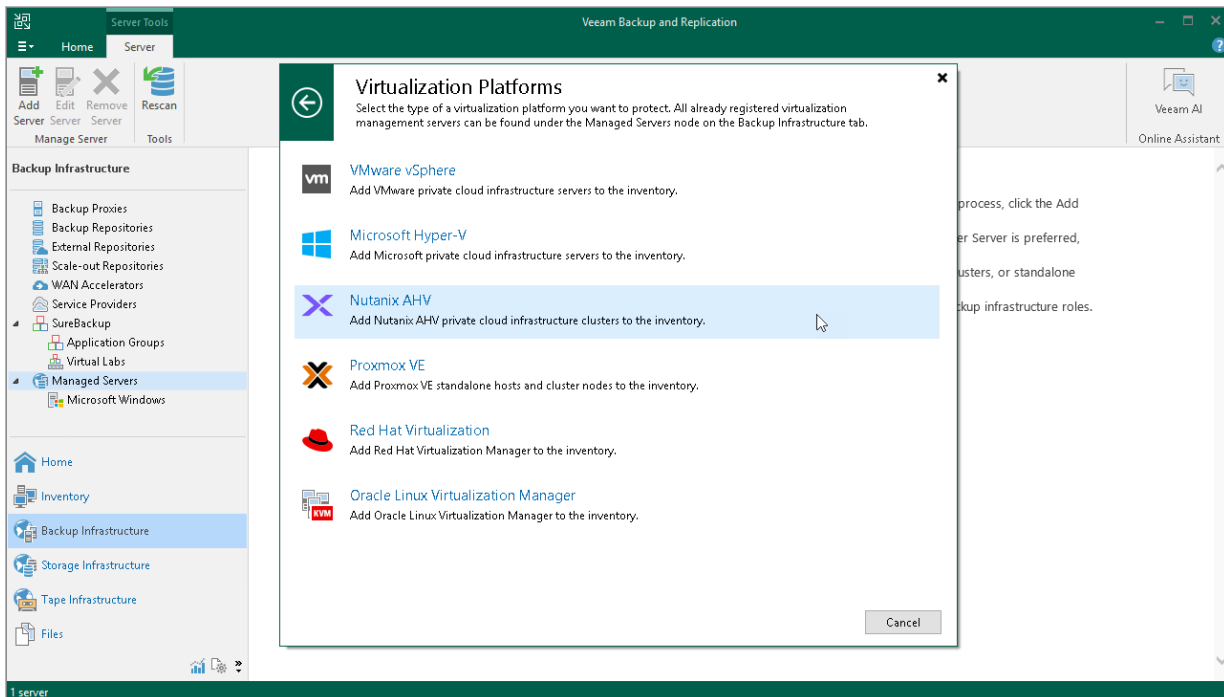
### IMPORTANT

If you unregister the cluster where the backup appliance resides from the Prism Central, you will not be able to protect resources in all clusters managed by the Prism Central. If you still want to unregister this cluster, you must first [back up the appliance configuration](#), then [remove the appliance](#) from the backup infrastructure, [deploy a new appliance](#) in another cluster registered with the Prism Central – and [restore the backed-up configuration](#) to the new appliance.

# Step 1. Launch New Nutanix AHV Server Wizard

To launch the **New Nutanix AHV Server** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.
3. On the ribbon, click **Add Server**.
4. In the **Add Server** window, select **Virtualization Platforms**.
5. In the **Virtualization Platforms** window, select **Nutanix AHV** to launch the **New Nutanix AHV Server** wizard.

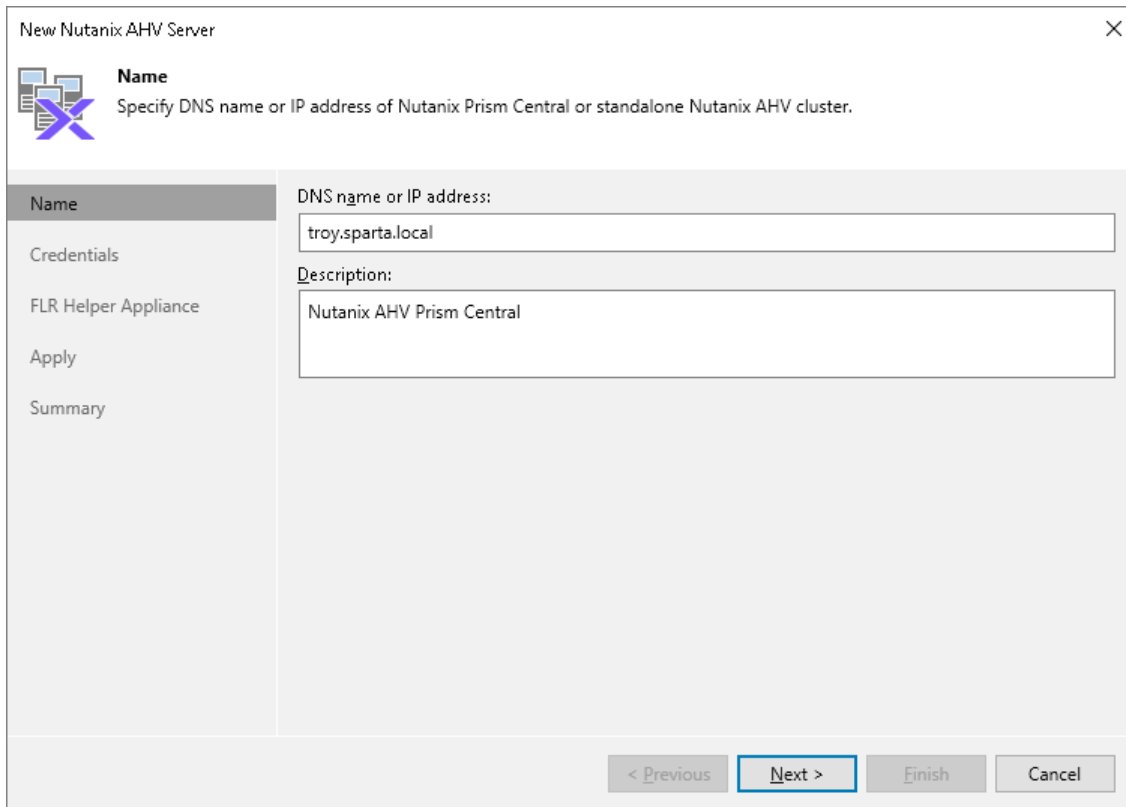




## Step 2. Specify Cluster Domain Name or Address

At the **Name** step of the wizard, do the following:

1. In the **DNS name or IP address** field, enter the FQDN or IP address of the Nutanix AHV standalone cluster or Prism Central .
2. In the **Description** field, provide a description for future reference. The field already contains a default description with information about the user who added the cluster or Prism Central, date and time when it was added.



The screenshot shows a wizard window titled "New Nutanix AHV Server" with a close button (X) in the top right corner. The window is divided into a left sidebar and a main content area. The sidebar contains the following items: "Name" (highlighted), "Credentials", "FLR Helper Appliance", "Apply", and "Summary". The main content area has a header section with a document icon and the text "Name Specify DNS name or IP address of Nutanix Prism Central or standalone Nutanix AHV cluster." Below this, there are two text input fields. The first field is labeled "DNS name or IP address:" and contains the text "troy.sparta.local". The second field is labeled "Description:" and contains the text "Nutanix AHV Prism Central". At the bottom of the window, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

## Step 3. Specify Credentials

At the **Credentials** step of the wizard, do the following specify credentials for an administrator account with the *Prism Admin* role that is used to access the cluster or Prism Central. For more information on Nutanix AHV system administrator roles, see [Nutanix documentation](#).

For credentials to be displayed in the **Credentials** list, they must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary credentials to the Credentials Manager beforehand, you can do this without closing the **New Nutanix AHV Server** wizard. To add an account, do the following:

1. Click **Add**.
2. In the **Credentials** window, specify a user name and password for the account.
3. Click **OK**.

The backup server will connect to the Nutanix AHV cluster or Prism Central and check its TLS certificate. If the certificate is not trusted, the **Certificate Security Alert Window** will display a warning notifying that secure communication cannot be guaranteed. To allow the backup server to connect to the Nutanix AHV cluster or Prism Central using the certificate, click **Continue**.

The screenshot shows the 'New Nutanix AHV Server' wizard at the 'Credentials' step. The main window has a sidebar with 'Name', 'Credentials', 'FLR Helper Appliance', 'Apply', and 'Summary'. The 'Credentials' section is active, displaying a list of credentials and an 'Add...' button. A 'Credentials' dialog box is open, showing fields for Username (admin), Password (masked with dots), and Description (Prism Central administrator credentials). The 'OK' button in the dialog is highlighted with a blue dashed border.

## Step 4. Configure Helper Appliance Settings

To restore [guest OS files](#) or [application items](#), Veeam Backup & Replication deploys on the AHV cluster a helper appliance. The helper appliance is a VM running the Linux operating system with a minimal set of components. The helper appliance occupies 50 MB of storage space and requires 3072 MB RAM.

At the **FLR Helper Appliance** step of the wizard, do the following:

1. [Applies only if you connect a Prism Central] Select the cluster whose helper appliance settings you want to modify and click **Edit**.
2. Click **Choose** next to the **Storage Container** field, and specify a storage container where helper appliance system files will be stored in the **Select Storage Container** window.

For a container to be displayed in the list of the available containers, it must be configured in the Nutanix AHV cluster as described in [Prism Element documentation](#) or in the Nutanix AHV Prism Central as described in [Prism Central documentation](#).

3. Click **Choose** next to the **Network** field, and specify a network to which the helper appliance will be connected.

For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Prism Element documentation](#) or in the Nutanix AHV Prism Central as described in [Prism Central documentation](#).

3. If DHCP is enabled for the selected network, the IP address and DNS settings of the helper appliance can be obtained automatically.

If DHCP is disabled for the selected network, or you want to specify an IP address and configure DNS settings manually, select the **Use the following IP address** option and enter the helper appliance IP address, subnet mask and default gateway. If you select this option, you will be able to perform only one guest OS file restore or application items restore session at a time.

4. To enable FTP access to the helper appliance, select the **Enable FTP server on appliance (advanced)** check box.

New Nutanix AHV Server

### FLR Helper Appliance

Specify a storage container and configure network settings for the FLR helper appliance. Make sure that you choose the same network to which the backup server is connected.

Name

Credentials

**FLR Helper Appliance**

Backup Appliance

Apply

Summary

FLR Appliance Configuration

Storage container:  
default-container-37740068791224 Choose...

Network:  
VM net Choose...

Obtain an IP address automatically  
 Use the following IP address

IP address: . . .

Subnet mask: . . .

Default gateway: . . .

Enable FTP Server on appliance (advanced)

OK Cancel

< Previous Next > Finish Cancel

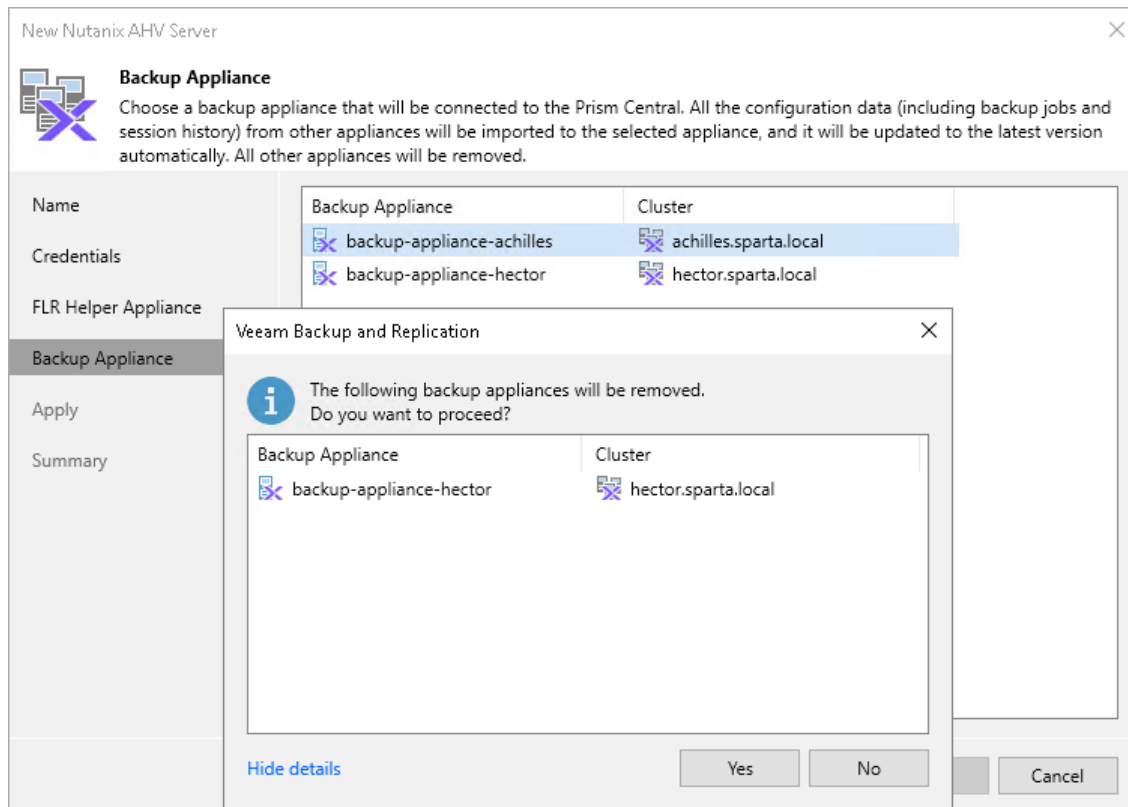
## Step 5. Choose Backup Appliance

[This step applies only if you connect a Prism Central that manages at least one cluster with a backup appliance already added to the backup infrastructure]

At the **Backup Appliance** step of the wizard, select the backup appliance that will be used in the [Prism Central deployment](#). This backup appliance will be automatically configured to manage all the jobs that are currently managed by dedicated backup appliances in the clusters.

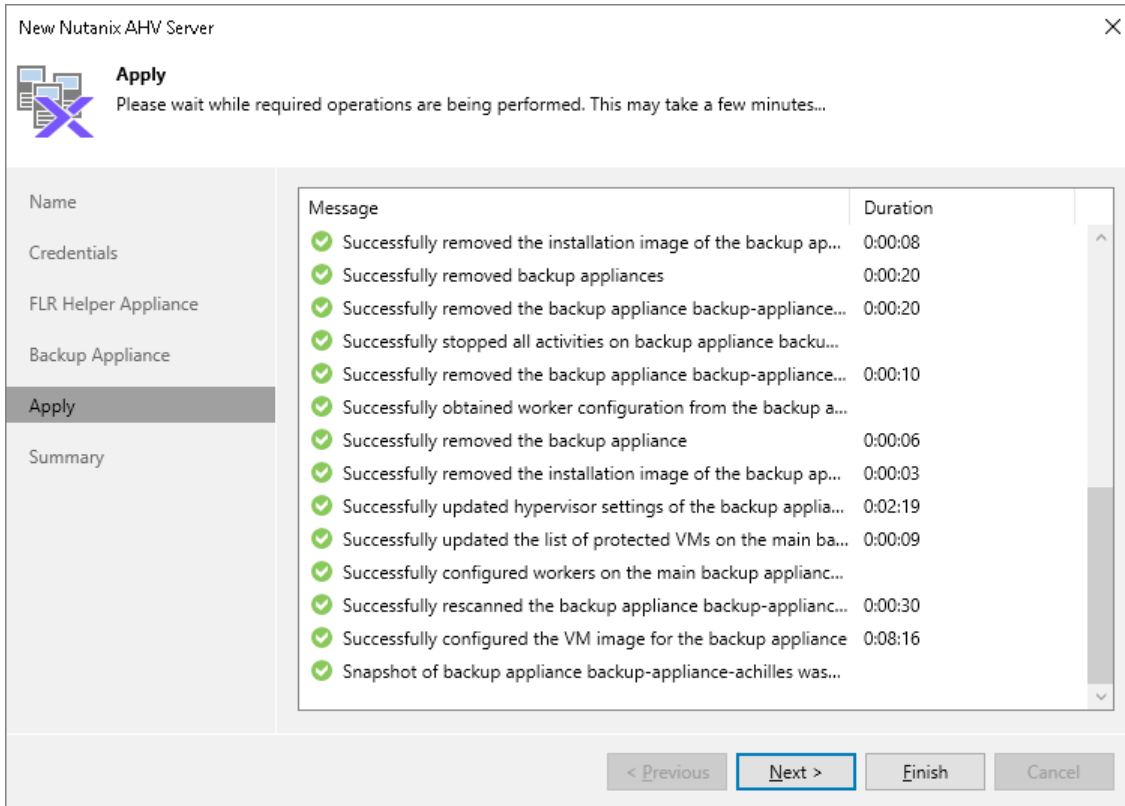
### NOTE

While connecting the Prism Central, Veeam Backup for Nutanix AHV will remove the other backup appliances and will deploy dedicated workers instead of them. The configuration of existing workers will be retained.



## Step 6. Apply Settings

At the **Apply** step of the wizard, wait until the cluster or Prism Central is added to the backup infrastructure and then click **Next**.



New Nutanix AHV Server

**Apply**  
Please wait while required operations are being performed. This may take a few minutes...

Name	Message	Duration
Credentials	✓ Successfully removed the installation image of the backup ap...	0:00:08
FLR Helper Appliance	✓ Successfully removed backup appliances	0:00:20
Backup Appliance	✓ Successfully removed the backup appliance backup-appliance...	0:00:20
<b>Apply</b>	✓ Successfully stopped all activities on backup appliance backu...	0:00:10
Summary	✓ Successfully removed the backup appliance backup-appliance...	0:00:10
	✓ Successfully obtained worker configuration from the backup a...	0:00:06
	✓ Successfully removed the backup appliance	0:00:06
	✓ Successfully removed the installation image of the backup ap...	0:00:03
	✓ Successfully updated hypervisor settings of the backup applia...	0:02:19
	✓ Successfully updated the list of protected VMs on the main ba...	0:00:09
	✓ Successfully configured workers on the main backup applianc...	0:00:30
	✓ Successfully rescanned the backup appliance backup-applianc...	0:00:30
	✓ Successfully configured the VM image for the backup appliance	0:08:16
	✓ Snapshot of backup appliance backup-appliance-achilles was...	

< Previous   **Next >**   Finish   Cancel

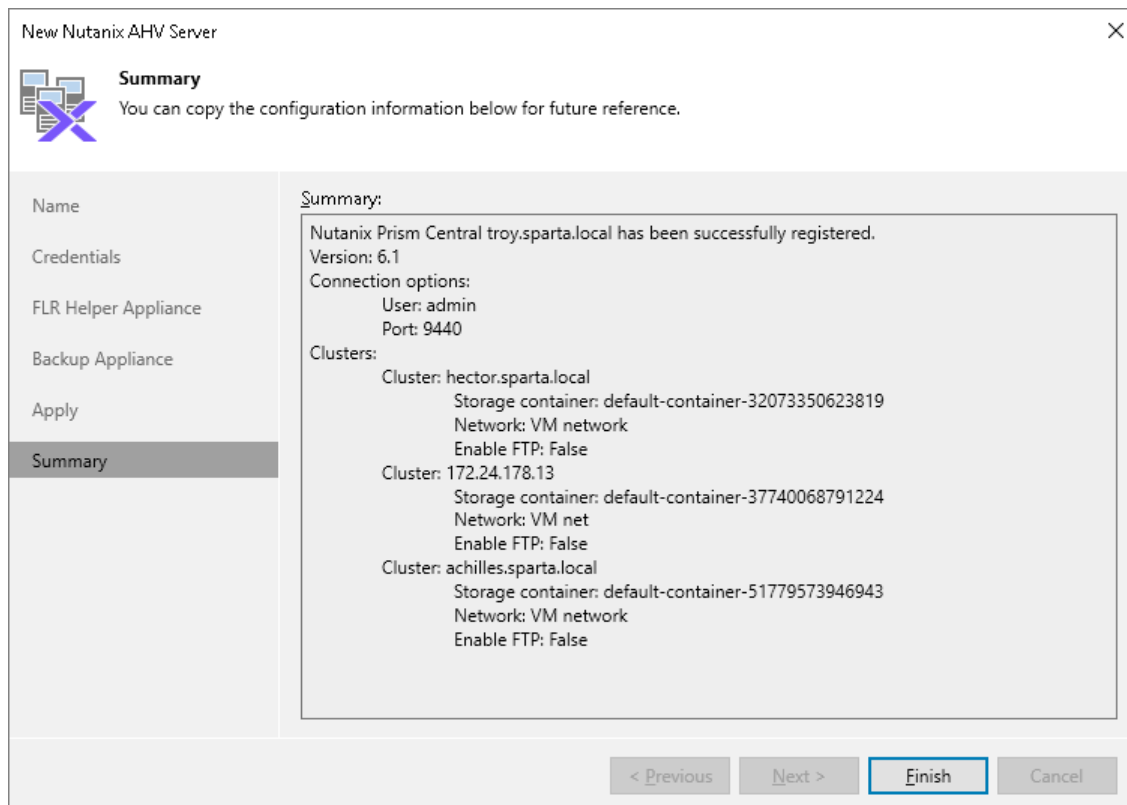
## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, check that the cluster or Prism Central has been successfully added and click **Finish**.

### TIP

You can review details of the cluster or Prism Central registration session in system logs as described in the Veeam Backup & Replication User Guide, section [Viewing History Statistics](#).

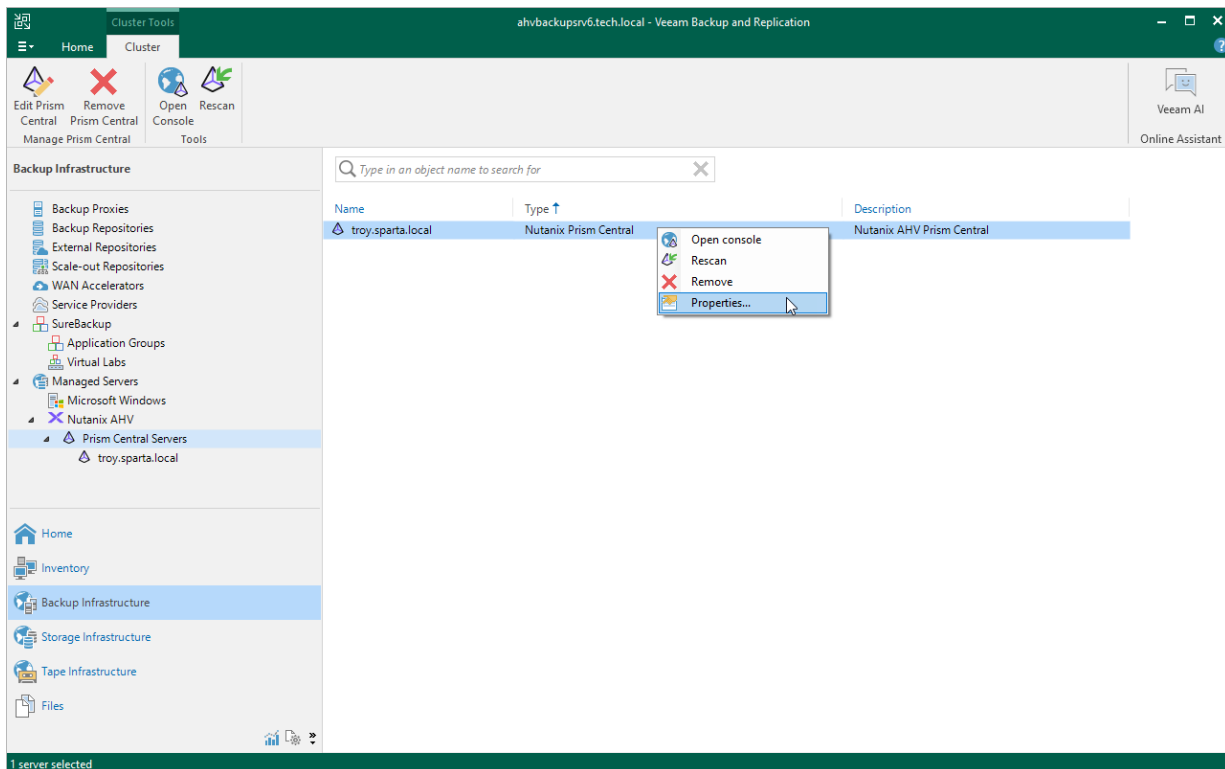
After you complete the wizard, it is required that you configure an Nutanix AHV backup appliance. You can proceed to the **New Nutanix AHV Backup Appliance** wizard immediately, or launch the wizard later as described in section [Managing Backup Appliances](#).



# Editing Nutanix AHV Server Properties

To edit properties of the Prism Central or Nutanix AHV cluster added to the backup infrastructure, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Nutanix AHV**.
3. In the working area, select the Prism Central or Nutanix AHV cluster and click **Edit** on the ribbon, or right-click the Nutanix AHV cluster and select **Properties**.
4. Complete the **Edit Nutanix AHV Cluster** wizard as described in section [Adding Nutanix AHV Server](#).





# Rescanning Nutanix AHV Server

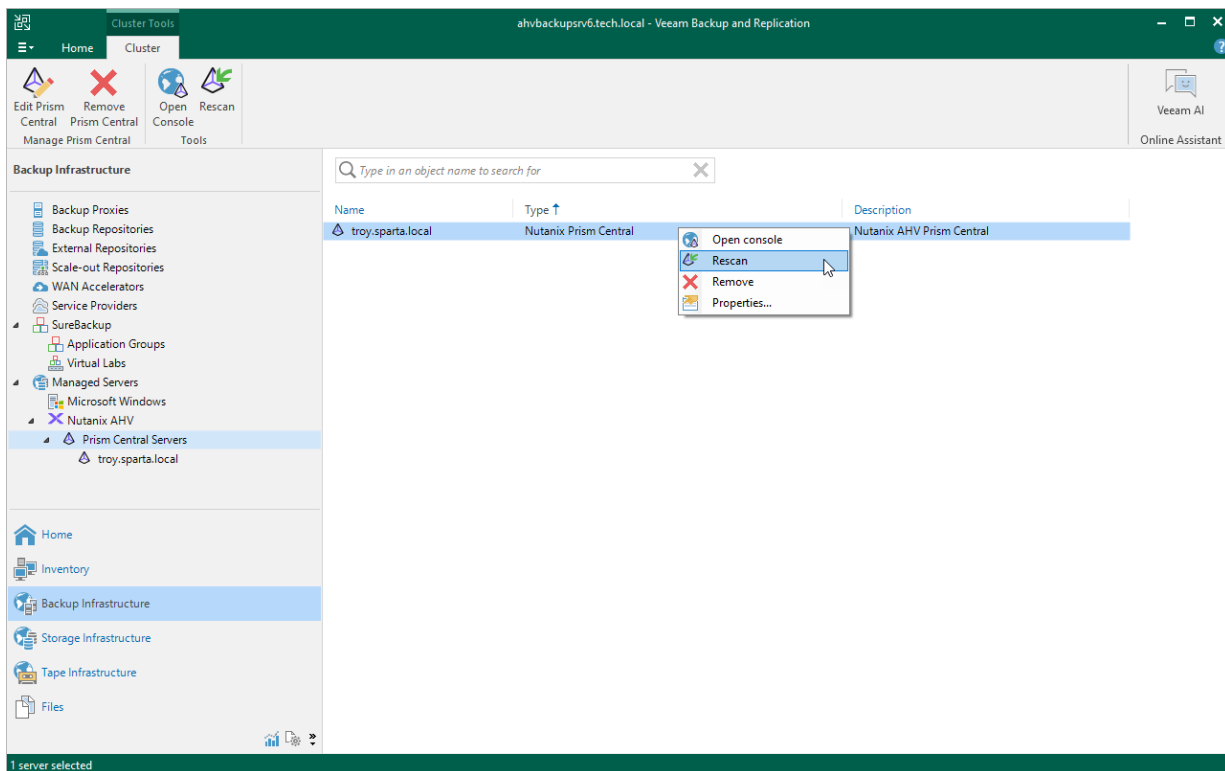
Veeam Backup for Nutanix AHV retrieves information about the Nutanix AHV resources from the Prism Central or Nutanix AHV cluster. However, the data synchronization process may take some time to complete. If you make any changes to the Nutanix AHV environment and want both the Veeam Backup & Replication console and the Nutanix AHV backup appliance web console to display the changes immediately, you can rescan the Prism Central or Nutanix AHV cluster manually.

To rescan the Prism Central or Nutanix AHV cluster, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Nutanix AHV**.
3. In the working area, select the Nutanix AHV cluster and click **Rescan** on the ribbon, or right-click the Nutanix AHV cluster and select **Rescan**.

## TIP

In the **System** window, you can track the progress of the rescan session. You can close the window and check session details later as described in the Veeam Backup & Replication User Guide, section [Viewing History Statistics](#)..



## TIP

Alternatively, in the [Nutanix AHV backup appliance web console](#), click **Configuration** at the top right corner of the Nutanix AHV backup appliance web console, select **Infrastructure** and click **Rescan**.

# Removing Nutanix AHV Server

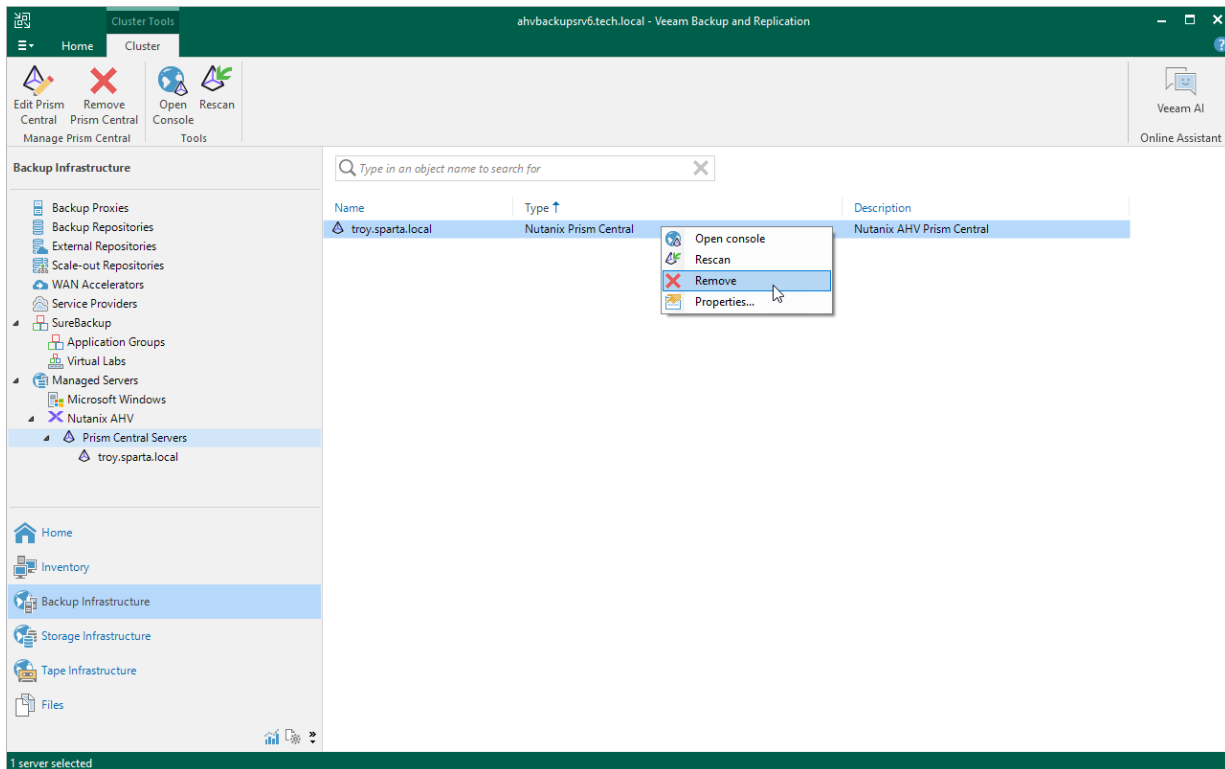
If you do not want to protect resources managed by the connected Prism Central or Nutanix AHV cluster anymore, you can remove it from the backup infrastructure.

## IMPORTANT

Before you remove the Prism Central or Nutanix AHV cluster, you must [remove the Nutanix AHV backup appliance](#) that processes protection jobs for the resources managed by the Prism Central or Nutanix AHV cluster.

To remove the Prism Central or Nutanix AHV cluster from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > Nutanix AHV**.
3. In the working area, select the Prism Central or Nutanix AHV cluster and click **Remove** on the ribbon, or right-click the Prism Central or Nutanix AHV cluster and select **Remove**.

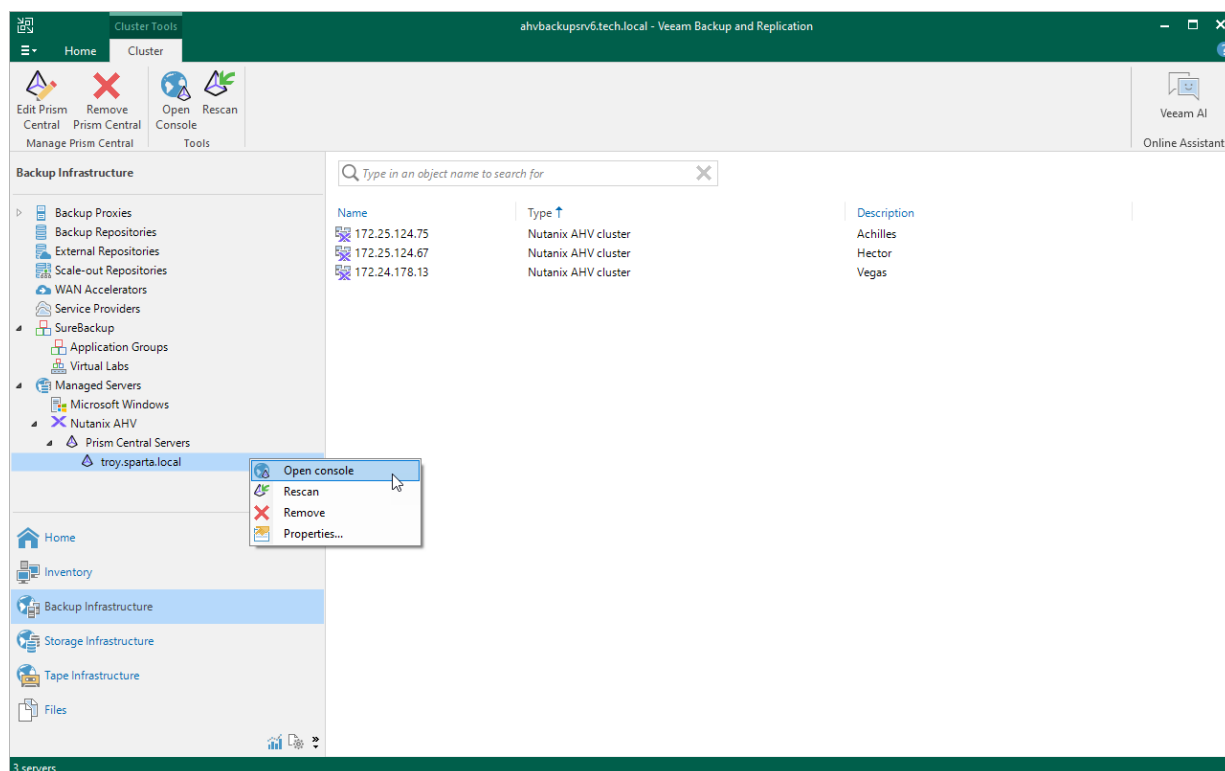


# Accessing Nutanix AHV Server Console

If you want to check the configuration of your Nutanix AHV infrastructure, you can use Veeam Backup for Nutanix AHV to launch the Prism Central console or the Prism Element console.

To access the Prism Central console, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.
3. Select the Prism Central and click **Open Console** on the ribbon, or right-click the Prism Central and select **Open Console**.



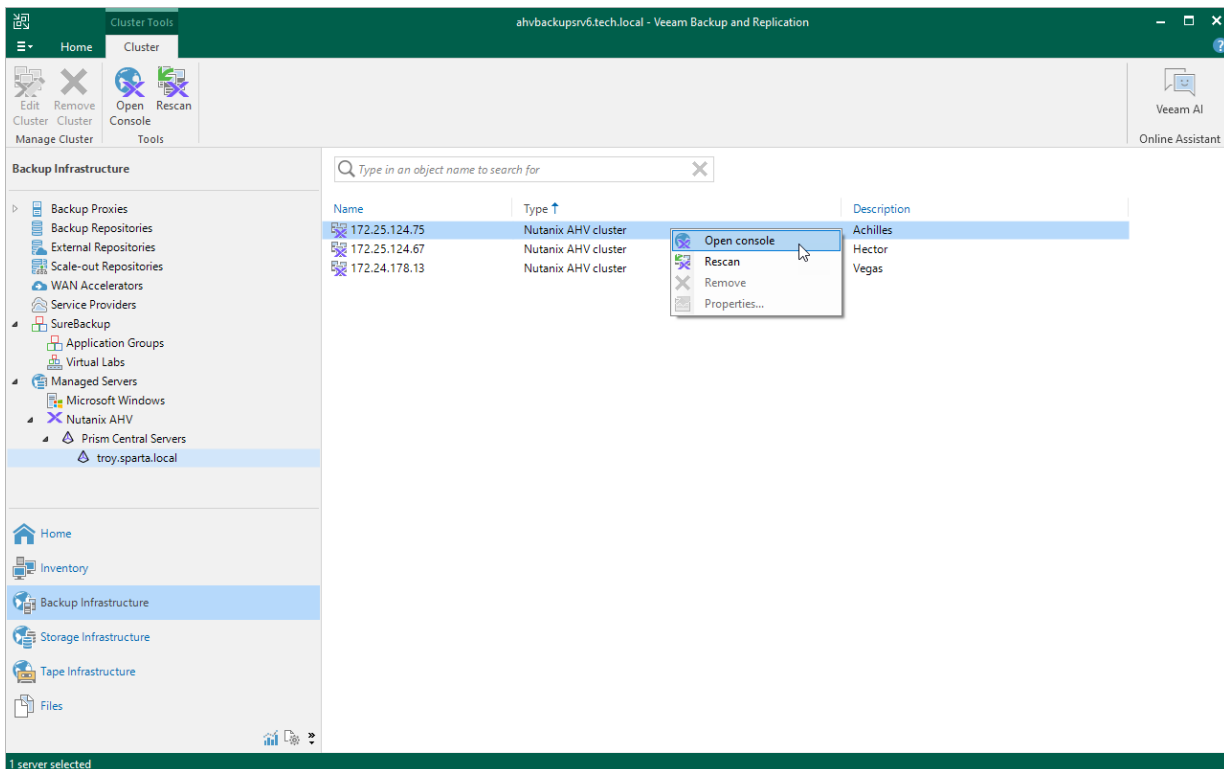
## TIP

You can also open the console using the [web console](#). To do that, switch to the **Configuration** page, navigate to **Infrastructure**, select the Prism Central and click **Open Nutanix Console**.

To access the Prism Element console, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers**.

3. Select a standalone cluster or a cluster registered with the Prism Central and click **Open Console** on the ribbon, or right-click a cluster and select **Open Console**.



## TIP

You can also open the console using the [web console](#). To do that, switch to the **Configuration** page, navigate to **Infrastructure**, select a cluster and click **Open Nutanix Console**.

# Managing Backup Appliances

To be able to back up VMs residing in a standalone cluster or clusters managed by a Prism Central, you must add to the backup infrastructure a Nutanix AHV backup appliance that will process backup jobs and deliver backup traffic to backup repositories.

To add a Nutanix AHV backup appliance, you can either deploy a new Nutanix AHV backup appliance or connect an existing one. Note that you can add only one Nutanix AHV backup appliance for each Prism Central or standalone cluster.

# Deploying New Backup Appliance

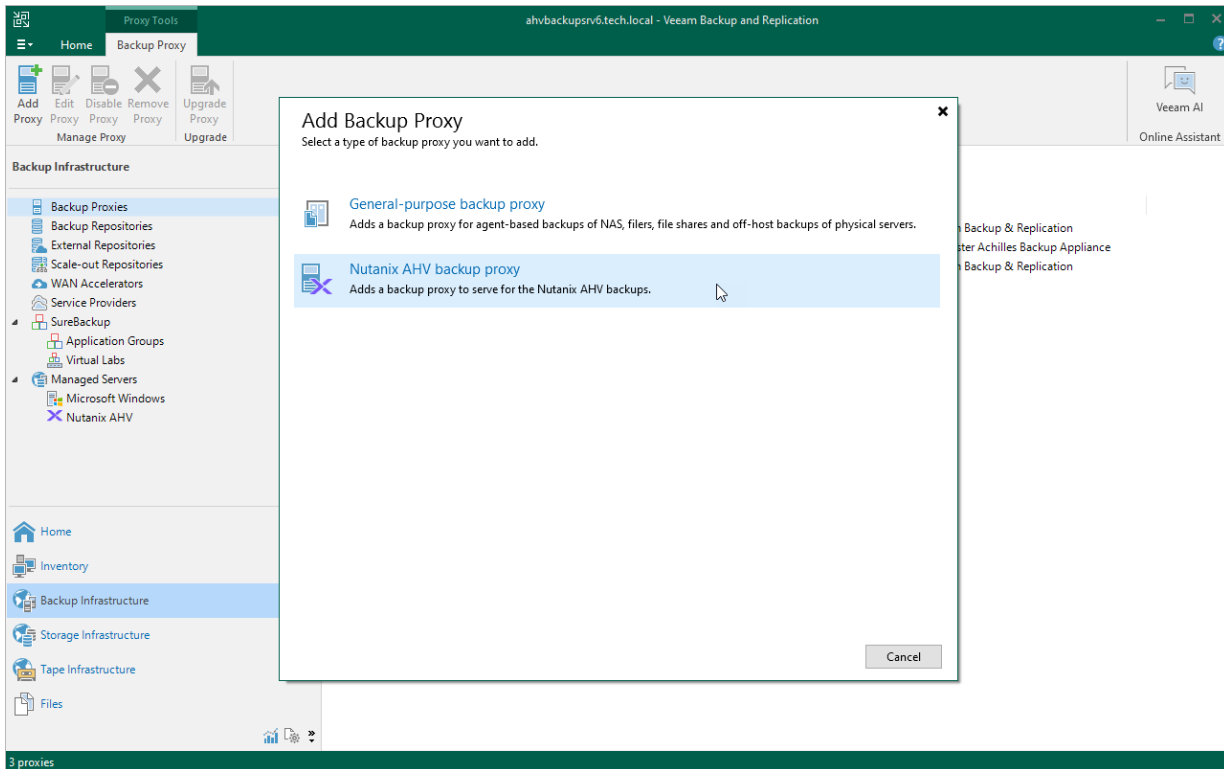
To deploy an Nutanix AHV backup appliance and to add it to the backup infrastructure, do the following:

1. [Launch the New Nutanix AHV Backup Appliance wizard.](#)
2. [Select the backup appliance deployment mode.](#)
3. [Specify backup appliance VM configuration.](#)
4. [Specify credentials for the backup appliance account.](#)
5. [Specify backup appliance network settings.](#)
6. [Grant permissions to the backup appliance.](#)
7. [Apply backup appliance settings.](#)
8. [Finish working with wizard.](#)

# Step 1. Launch New Nutanix Proxy Wizard

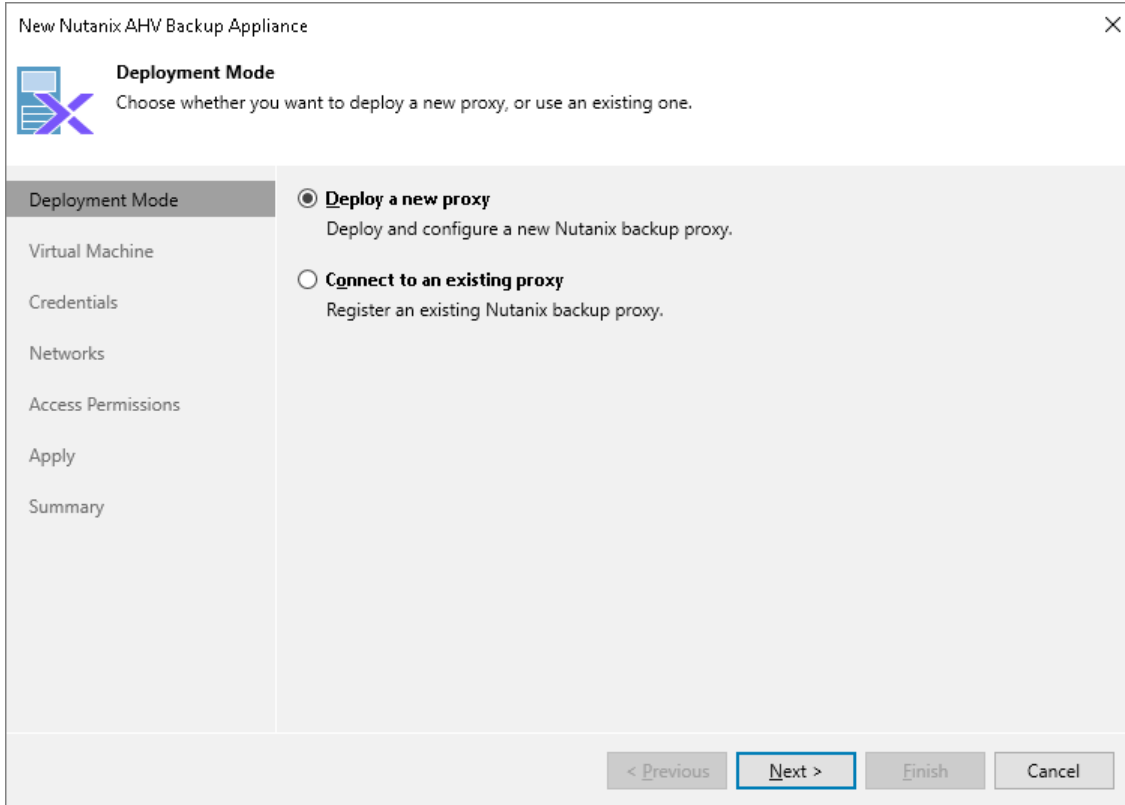
To launch the **New Nutanix AHV Backup Appliance** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. On the ribbon, select **Add Proxy**.
4. Click **Nutanix AHV backup proxy**.



## Step 2. Select Deployment Mode

At the **Deployment Mode** step of the wizard, select the **Deploy a new proxy** option.



The screenshot shows a wizard window titled "New Nutanix AHV Backup Appliance" with a close button (X) in the top right corner. The main heading is "Deployment Mode" with a sub-instruction: "Choose whether you want to deploy a new proxy, or use an existing one." On the left is a vertical navigation pane with the following items: "Deployment Mode" (highlighted), "Virtual Machine", "Credentials", "Networks", "Access Permissions", "Apply", and "Summary". The main content area contains two radio button options:   
1. **Deploy a new proxy** (selected): "Deploy and configure a new Nutanix backup proxy."   
2. **Connect to an existing proxy**: "Register an existing Nutanix backup proxy."   
At the bottom right, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).



## Step 3. Specify VM Configuration

At the **Virtual Machine** step of the wizard, do the following:

1. Click **Choose** next to the **Cluster** field, and specify a Nutanix AHV cluster where the Nutanix AHV backup appliance will be deployed in the **Select Cluster** window.

For a cluster to be displayed in the list of the available clusters, it must be added to the backup infrastructure as described in section [Adding Nutanix AHV Server](#).

### NOTE

In the [Prism Central deployment](#), Veeam Backup for Nutanix AHV requires one backup appliance deployed in any cluster registered with the Prism Central.

2. In the **Name** field, specify a hostname (with domain name excluded) that will be assigned to the Nutanix AHV backup appliance.

The maximum length of the hostname is 63 characters. The hyphen-minus character (-) is supported, but you cannot use it as the first or the last character of the name.

3. [Applies only to the [Standalone Cluster deployment](#)] Click **Choose** next to the **Storage Container** field, and specify a storage container where Nutanix AHV backup appliance system files will be stored in the **Select Storage Container** window.

For a container to be displayed in the list of the available containers, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

### NOTE

In the [Prism Central deployment](#), Veeam Backup for Nutanix AHV automatically selects the *SelfService* storage container that cannot be changed.

4. In the **Proxy description** field, provide a description for future reference. The field already contains a default description with information about the user who added the proxy, date and time when the proxy was added.
5. In the **Max concurrent tasks** field, specify the number of tasks that the Nutanix AHV backup appliance will be able to handle in parallel. If this value is exceeded, the backup appliance will not start processing a new task until one of the currently running tasks finishes.

The default number of concurrent tasks is set to 4. When you change this value, the wizard automatically adjusts the amount of resources that will be allocated to the VM running as the Nutanix AHV backup appliance. If you want to specify the amount of resources manually, click **Advanced**. Note that you must take into account the [Nutanix AHV backup appliance system requirements](#).

### NOTE

When performing data protection and disaster recovery operations, Veeam Backup for Nutanix AHV initiates a new task for each VM that is being processed.

New Nutanix AHV Backup Appliance

**Virtual Machine**  
Configure virtual machine for Nutanix proxy placement.

Deployment Mode

Virtual Machine

Credentials

Networks

Access Permissions

Apply

Summary

Cluster: hector.sparta.local Choose...

Name: backup-appliance-hector

Storage container: default-container-32073350623819 Choose...

Proxy description: Nutanix AHV Cluster Hector Backup Appliance

Max concurrent tasks: 4

Hardware Settings

Number of vCPU: 6

Memory size (GB): 6

[Reset to default settings](#)

OK Cancel

Advanced proxy settings include vCPU and memory sizing settings for proxy VM. Advanced...

< Previous Next > Finish Cancel

## Step 4. Specify Credentials

At the **Credentials** step of the wizard, select credentials for an account that will be created on the Nutanix AHV backup appliance for accessing the web console.

### IMPORTANT

Do not select Active Directory accounts – the Nutanix AHV backup appliance does not support LDAP integration.

For credentials to be displayed in the **Credentials** list, they must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary credentials to the Credentials Manager beforehand, you can do this without closing the **New Nutanix AHV Backup Appliance** wizard. To add credentials, do the following:

1. Click **Add**.
2. In the **Credentials** window, specify a user name and password for the account.

The user name must start with a lowercase Latin letter and must not match Linux system user names (such as *root*, *daemon*). The name can contain only lowercase Latin letters, numeric characters, underscores and dashes. The maximum length of the name is 32 characters.

3. Click **OK**.

The screenshot shows the 'New Nutanix AHV Backup Appliance' wizard at the 'Credentials' step. The main window has a sidebar with options: Deployment Mode, Virtual Machine, Credentials (selected), Networks, Access Permissions, Apply, and Summary. The main area shows a list of credentials with an 'Add...' button. A 'Credentials' dialog box is open, with the following fields:

- Username: administrator
- Password: [masked]
- Description: Backup appliance administrator credentials

The dialog has 'OK' and 'Cancel' buttons. The 'OK' button is highlighted with a red dashed box. The background wizard also has '< Previous', 'Next >', 'Finish', and 'Cancel' buttons at the bottom.

## Step 5. Specify Network Settings

At the **Networks** step of the wizard, do the following:

1. Click **Add** to configure backup appliance network interfaces:
  - a. From the **Network** drop-down list, select a network to which the Nutanix AHV backup appliance network interface will be connected.  
  
For best performance, choose the network that is used for Nutanix Controller VM (CVM). For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).
  - b. In the **Description** field, provide a network interface description for future reference.
  - c. If DHCP is enabled in the selected network, the IP address of the Nutanix AHV backup appliance can be obtained automatically.  
  
If DHCP is disabled in the selected network, or you want to specify an IP address, select the **Use the following IP address** option and enter the worker IP address, subnet mask and default gateway.

To add more network interfaces, repeat the step and specify the network order using the **Up** and **Down** buttons. For more information on multi-network configuration, see section [Appendix D. Configuring Multiple Networks](#).

2. If DHCP is enabled in any network to which the Nutanix AHV backup appliance will be connected, DNS settings of the Nutanix AHV backup appliance can be obtained automatically. To configure DNS settings manually, click **Obtain automatically** and do the following in the **DNS Server Settings** window:
  - a. Select the **Use the following DNS server address** check box.
  - b. Enter the IP addresses of the preferred and alternate DNS servers.
  - c. Click **OK**.

Since the backup appliance are Linux-based VMs, they have the same limitations that apply to machines running the Rocky Linux operating system. That is, DNS settings cannot be configured separately for each network added to the backup appliance.

## IMPORTANT

If you specified the Nutanix AHV server domain name while [adding it to the backup infrastructure](#), ensure that the appliance is able to resolve the name to the cluster IP address.

New Nutanix AHV Backup Appliance

**Networks**  
Specify network settings for the Nutanix backup appliance.

Deployment Mode  
Virtual Machine  
Credentials  
**Networks**  
Access Permissions  
Apply  
Summary

Hostname: backup-applian  
DNS server address  
Networks:  
Order ↑

**Network Settings**

Network: VM network  
Description: Created by AHVBACKUPSRV6\Administrator (Veeam Backup & Replication) at 8/14/2024 4:54 PM.

**Settings**

Obtain an IP address automatically  
 Use the following IP address

IP address: . . .  
Subnet mask: . . .  
Default gateway: . . .

OK Cancel

Add...  
Edit...  
Remove  
↑ Up  
↓ Down

< Previous Next > Finish Cancel

## Step 6. Grant Permissions

At the **Access Permissions** step of the wizard, do the following:

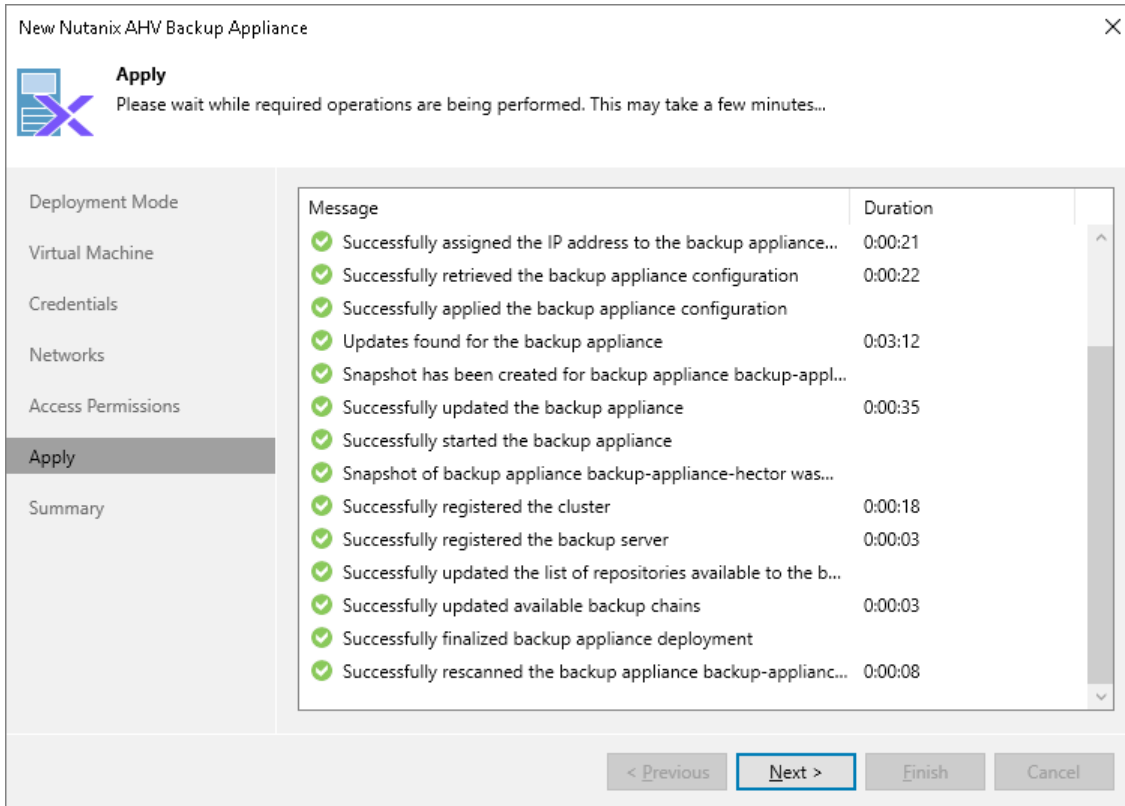
- Select the **Allow access to all backup repositories** option if you want the Nutanix AHV backup appliance to have access to all backup repositories added to the backup infrastructure.
- Select the **Allow access to the following backup repositories** option if you want the Nutanix AHV backup appliance to have access to specific backup repositories only.

If you select the **Allow access to the following backup repositories** option, you must also specify backup repositories to which the Nutanix AHV backup appliance will have access. For a backup repository to be displayed in the **Repository** list, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section [Backup Repository](#).

The screenshot shows the 'New Nutanix AHV Backup Appliance' wizard at the 'Access Permissions' step. The window title is 'New Nutanix AHV Backup Appliance' with a close button (X) in the top right corner. Below the title bar is a Nutanix logo and the text 'Access Permissions' and 'Specify the backup repositories this AHV proxy is allowed to access.' A left-hand navigation pane lists steps: 'Deployment Mode', 'Virtual Machine', 'Credentials', 'Networks', 'Access Permissions' (highlighted), 'Apply', and 'Summary'. The main area contains two radio button options: 'Allow access to all backup repositories' (selected) and 'Allow access to the following backup repositories'. Below these is a list box titled 'Repository' containing one entry: 'Default Backup Repository' with an unchecked checkbox. To the right of the list box are 'Select All' and 'Clear All' buttons. At the bottom of the window are four buttons: '< Previous', 'Apply' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for the Nutanix AHV backup appliance to be added to the backup infrastructure and then click **Next**.



New Nutanix AHV Backup Appliance

**Apply**  
Please wait while required operations are being performed. This may take a few minutes...

Message	Duration
✓ Successfully assigned the IP address to the backup appliance...	0:00:21
✓ Successfully retrieved the backup appliance configuration	0:00:22
✓ Successfully applied the backup appliance configuration	
✓ Updates found for the backup appliance	0:03:12
✓ Snapshot has been created for backup appliance backup-appl...	
✓ Successfully updated the backup appliance	0:00:35
✓ Successfully started the backup appliance	
✓ Snapshot of backup appliance backup-appliance-hector was...	
✓ Successfully registered the cluster	0:00:18
✓ Successfully registered the backup server	0:00:03
✓ Successfully updated the list of repositories available to the b...	
✓ Successfully updated available backup chains	0:00:03
✓ Successfully finalized backup appliance deployment	
✓ Successfully rescanned the backup appliance backup-applianc...	0:00:08

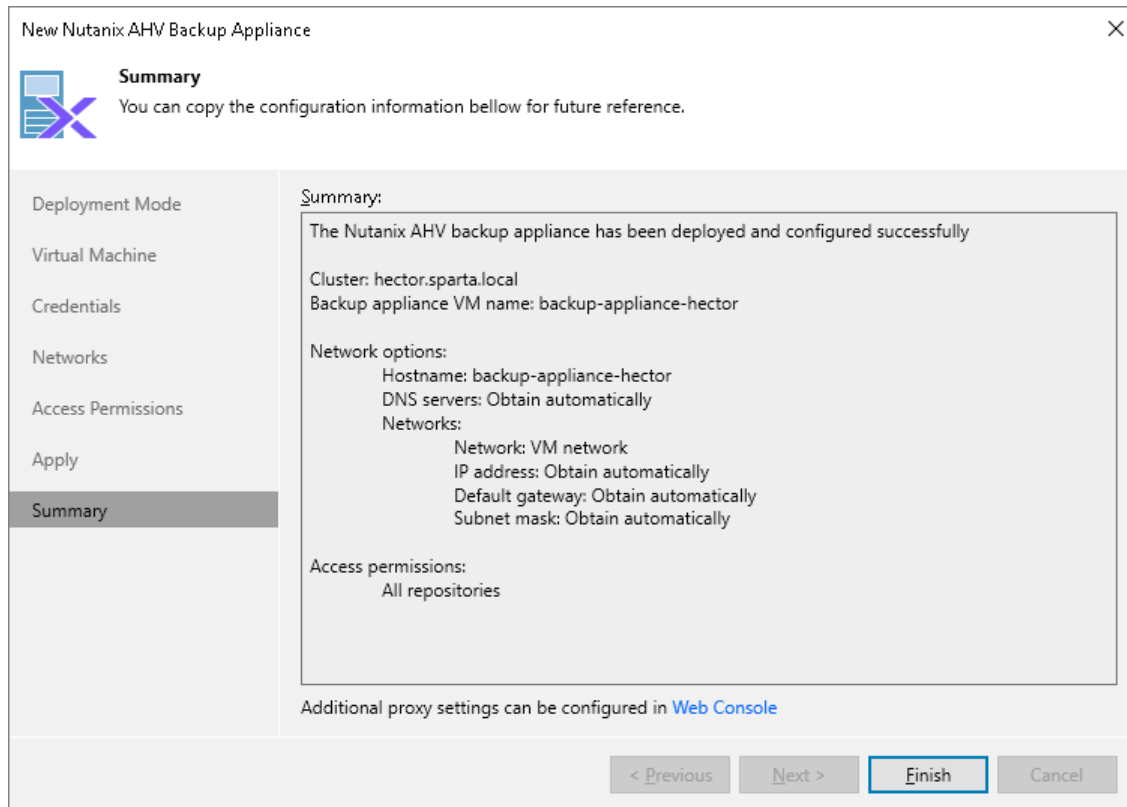
< Previous   **Next >**   Finish   Cancel

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. To open the Nutanix AHV backup appliance web console in your web browser, click **Web Console**.

### TIP

You can review details of the deployment session in system logs as described in the Veeam Backup & Replication User Guide, section [Viewing History Statistics](#).





# Connecting Existing Backup Appliance

If you have a Nutanix AHV backup appliance that has already been deployed but was removed from the backup infrastructure, you can connect it to the backup server. You may also want to connect an existing Nutanix AHV backup appliance in the following situations:

- To upgrade a Nutanix AHV backup appliance from version 4.0, 4a, 5.0, 5.1, 6 and 6.1 to 7.
- To connect a Nutanix AHV backup appliance that was previously connected to another backup server.

To add an existing Nutanix AHV backup appliance to the backup infrastructure, do the following:

1. [Launch the New Nutanix AHV Backup Appliance wizard.](#)
2. [Select the proxy deployment mode.](#)
3. [Specify proxy VM configuration.](#)
4. [Enter credentials for the proxy account.](#)
5. [Check network settings.](#)
6. [Grant permissions to the proxy.](#)
7. [Apply proxy settings.](#)
8. [Finish working with wizard.](#)

After you connect the Nutanix AHV backup appliance, the backup server will retrieve information about all jobs the appliance has ever processed. If the backup server configuration database contains records about Nutanix AHV VM backups and if the backup files are still available in repositories, they will be imported to the Nutanix AHV backup appliance automatically and you will be able to use them to restore [entire VMs](#) and [VM disks](#).

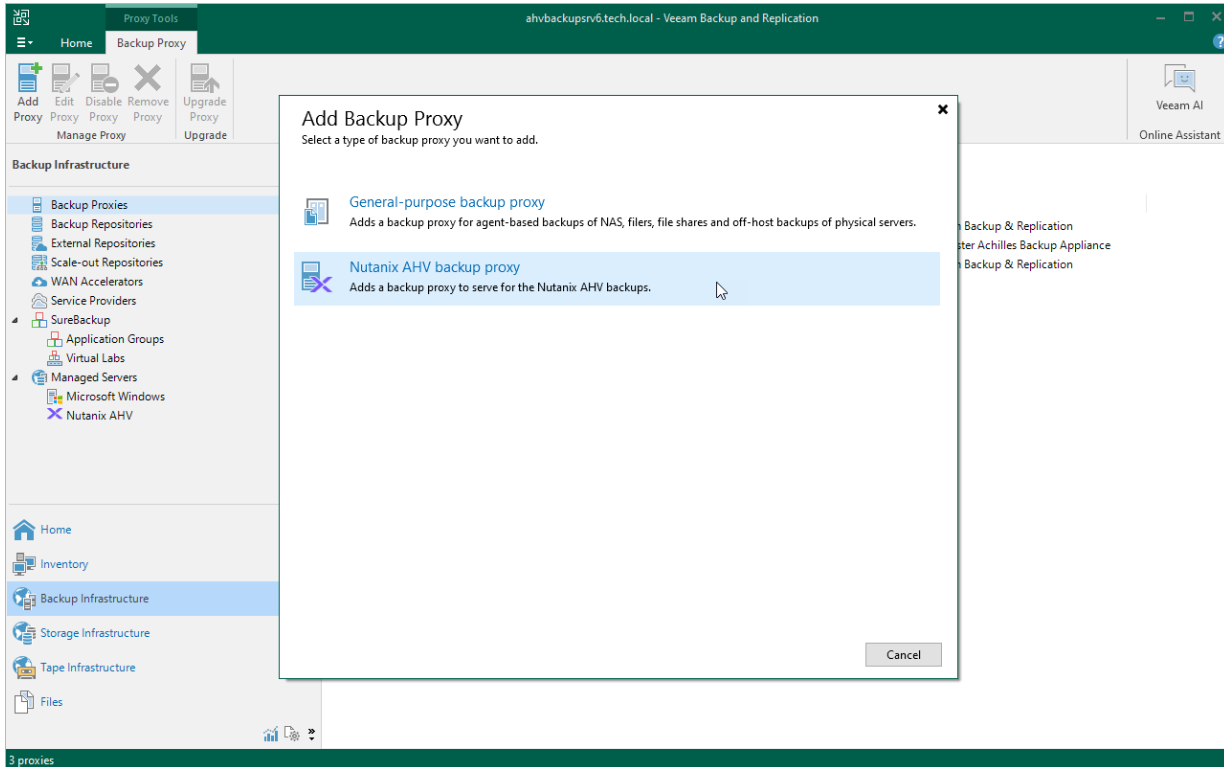
## NOTE

If you connect a Nutanix AHV backup appliance migrated from another cluster to protect VMs migrated from the same cluster, you will need to reconfigure backup jobs. UUIDs of migrated VMs change, therefore, you will need to re-add VMs to a backup job that will start new backup chains for them.

# Step 1. Launch New Nutanix Proxy Wizard

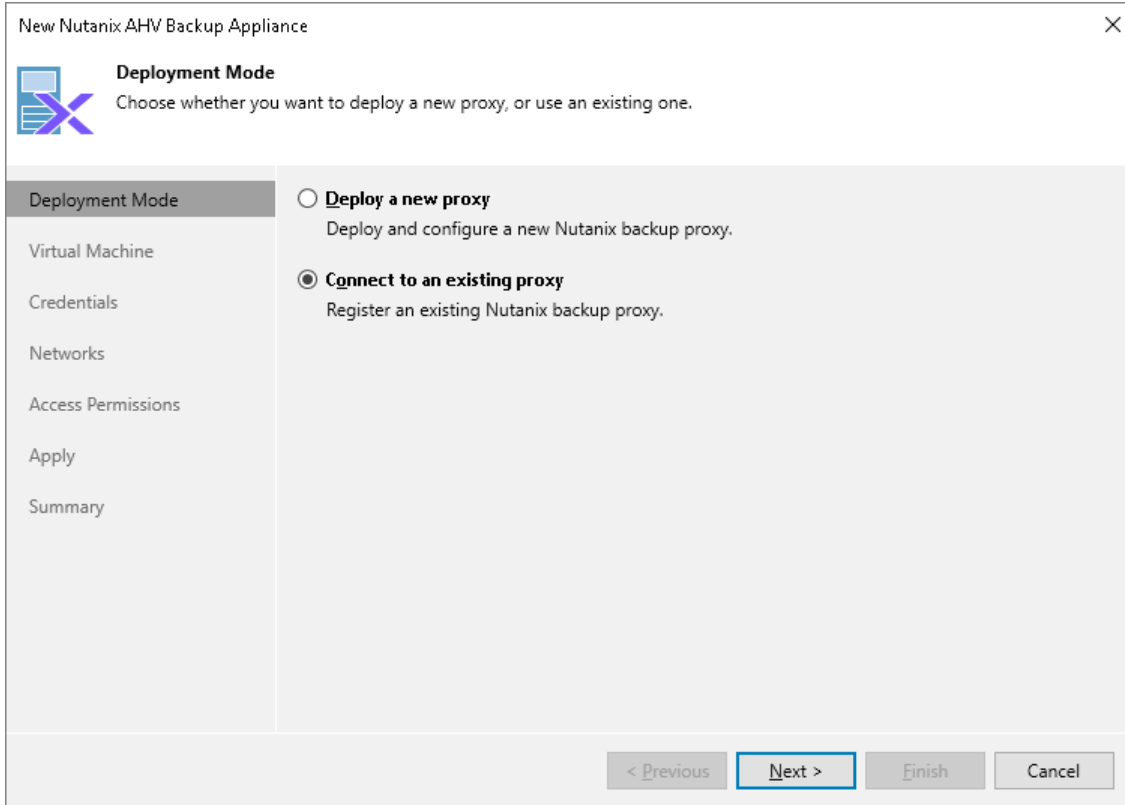
To launch the **New Nutanix AHV Backup Appliance** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. On the ribbon, select **Add Proxy**.
4. Click **Nutanix AHV backup proxy**.



## Step 2. Select Deployment Mode

At the **Deployment Mode** step, select the **Connect to an existing proxy** option.



The screenshot shows a wizard window titled "New Nutanix AHV Backup Appliance" with a close button (X) in the top right corner. The main heading is "Deployment Mode" with a sub-instruction: "Choose whether you want to deploy a new proxy, or use an existing one." On the left is a vertical navigation pane with the following items: "Deployment Mode" (highlighted), "Virtual Machine", "Credentials", "Networks", "Access Permissions", "Apply", and "Summary". The main content area contains two radio button options: "Deploy a new proxy" (unselected) with the description "Deploy and configure a new Nutanix backup proxy.", and "Connect to an existing proxy" (selected) with the description "Register an existing Nutanix backup proxy.". At the bottom right, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

## Step 3. Specify VM Configuration

At the **Virtual Machine** step of the wizard, do the following:

1. Click **Choose** next to the **Cluster** field, and specify a Nutanix AHV cluster where the Nutanix AHV backup appliance is deployed in the **Select Cluster** window.

For a cluster to be displayed in the list of the available clusters, it must be added to the backup infrastructure as described in section [Adding Nutanix AHV Server](#).

2. Click **Choose** next to the **Name** field, and specify the VM running as the Nutanix AHV backup appliance in the **Select Virtual Machine** window.

### NOTE

You cannot change the storage container – it is automatically populated when you select the VM.

3. In the **Proxy description** field, provide a description for future reference. The field already contains a default description with information about the user who added the proxy, date and time when the proxy was added.
4. In the **Max concurrent tasks** field, specify the number of tasks that the Nutanix AHV backup appliance will be able to handle in parallel. If this value is exceeded, the backup appliance will not start processing a new task until one of the currently running tasks finishes. If you change the number of concurrent tasks, the wizard will automatically adjust the amount of resources allocated to the VM running as the Nutanix AHV backup appliance.

When you change number of concurrent tasks, the wizard automatically adjusts the amount of resources that will be allocated to the VM running as the Nutanix AHV backup appliance. If you want to specify the amount of resources manually, click **Advanced**. Note that you must take into account the [Nutanix AHV backup appliance system requirements](#).

## NOTE

When performing data protection and disaster recovery operations, Veeam Backup for Nutanix AHV initiates a new task for each VM that is being processed.

New Nutanix AHV Backup Appliance

**Virtual Machine**  
Select the existing Nutanix proxy virtual machine.

Deployment Mode

Virtual Machine

Credentials

Networks

Access Permissions

Apply

Summary

Cluster:  
achilles.sparta.local Choose...

Name:  
backup-appliance-achilles Choose...

Storage container:  
default-container-51779573946943

Proxy description:  
Nutanix AHV Cluster Achilles Backup Appliance

Max concurrent tasks:  
4

Hardware Settings

Number of vCPUs: 6

Memory size (GB): 6

Reset to default settings

OK Cancel

Advanced proxy settings include vCPU and memory sizing settings for proxy VM. Advanced...

< Previous Next > Finish Cancel

## Step 4. Enter Credentials

At the **Credentials** step of the wizard, specify credentials of a Portal Administrator account that you use to access the Nutanix AHV backup appliance web console.

### NOTE

The following accounts cannot be used to access the Nutanix AHV backup appliance as Portal Administrators:

- Accounts for which [multi-factor authentication \(MFA\)](#) is enabled.
- Active Directory accounts since the Nutanix AHV backup appliance does not support LDAP integration.

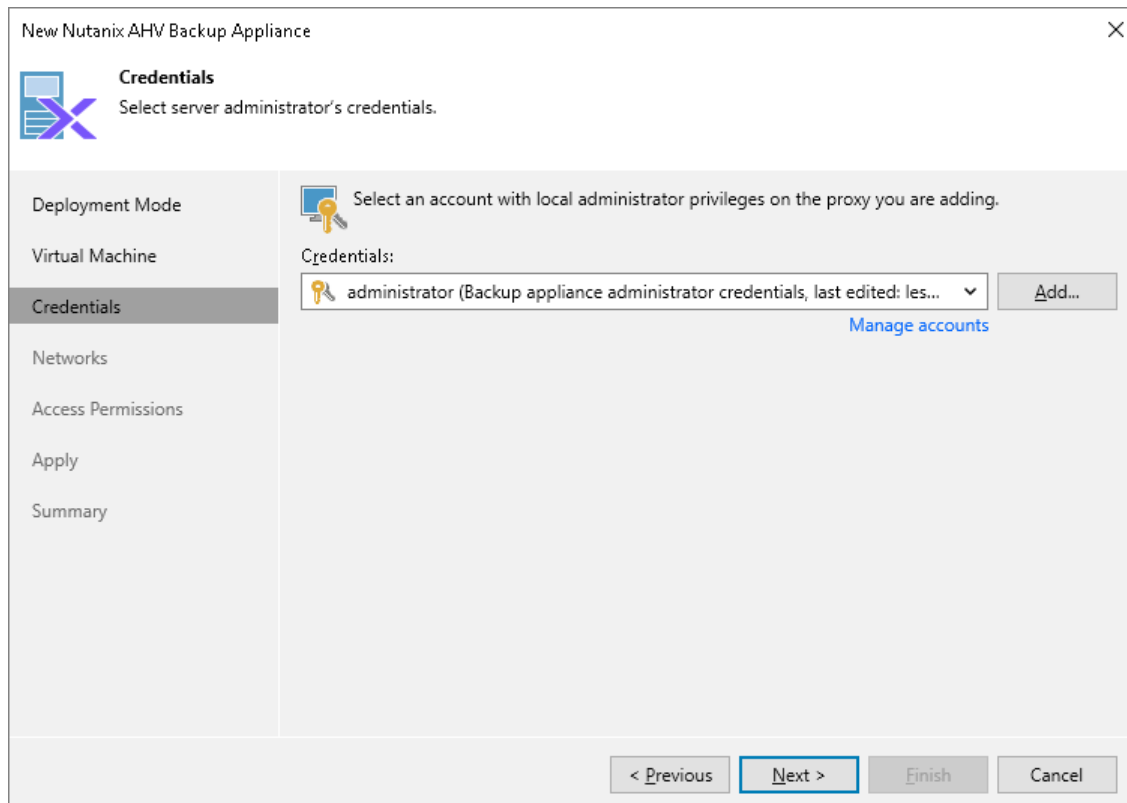
For credentials to be displayed in the **Credentials** list, they must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary credentials to the Credentials Manager beforehand, you can do this without closing the **New Nutanix AHV Backup Appliance** wizard. To add credentials, do the following:

1. Click **Add**.
2. In the **Credentials** window, specify a user name and password for the account.

The user name must start with a lowercase Latin letter and must not match Linux system user names (such as *root*, *daemon*). The name can contain only lowercase Latin letters, numeric characters, underscores and dashes. The maximum length of the name is 32 characters.

3. Click **OK**.

The backup server will connect to the Nutanix AHV backup appliance and check its TLS certificate. If the certificate is not trusted, the **Certificate Security Alert Window** will display a warning notifying that secure communication cannot be guaranteed. To allow the backup server to connect to the Nutanix AHV backup appliance using the certificate, click **Continue**.



The screenshot shows the 'New Nutanix AHV Backup Appliance' wizard window, specifically the 'Credentials' step. The window title is 'New Nutanix AHV Backup Appliance' and it has a close button (X) in the top right corner. The main area is titled 'Credentials' and contains the instruction 'Select server administrator's credentials.' Below this, there is a section for 'Deployment Mode' and 'Virtual Machine', followed by the 'Credentials' section which is currently selected in the left-hand navigation pane. The 'Credentials' section displays a list of accounts with a key icon and the text 'Select an account with local administrator privileges on the proxy you are adding.' The list shows one account: 'administrator (Backup appliance administrator credentials, last edited: les...' with a dropdown arrow. To the right of the list is an 'Add...' button. Below the list is a 'Manage accounts' link. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

## Step 5. Check Network Settings

When you connect an existing Nutanix AHV backup appliance, you can change the DNS server settings and the network to which the Nutanix AHV backup appliance is connected or specify a new IP address for the appliance. To do that, select a configured network interface and click **Edit**, or click **Add** to specify settings for a new network interface. For more information on multi-network configuration, see section [Appendix D. Configuring Multiple Networks](#).

New Nutanix AHV Backup Appliance

**Networks**  
Specify network settings for the Nutanix backup appliance.

Deployment Mode

Virtual Machine

Credentials

**Networks**

Access Permissions

Apply

Summary

Hostname:  
backup-appliance-achilles

DNS server addresses: Obtain automatically

Networks:

Order ↑	Network	IP	Description
1	VM network	DHCP	Created by AHVBACKUP...

Add...

Edit...

Remove

↑ Up

↓ Down

< Previous

Next >

Finish

Cancel

## Step 6. Grant Permissions

At the **Access Permissions** step of the wizard, do the following:

- Select the **Allow access to all backup repositories** option if you want the Nutanix AHV backup appliance to have access to all backup repositories added to the backup infrastructure.
- Select the **Allow access to the following backup repositories** option if you want the Nutanix AHV backup appliance to have access to specific backup repositories only.

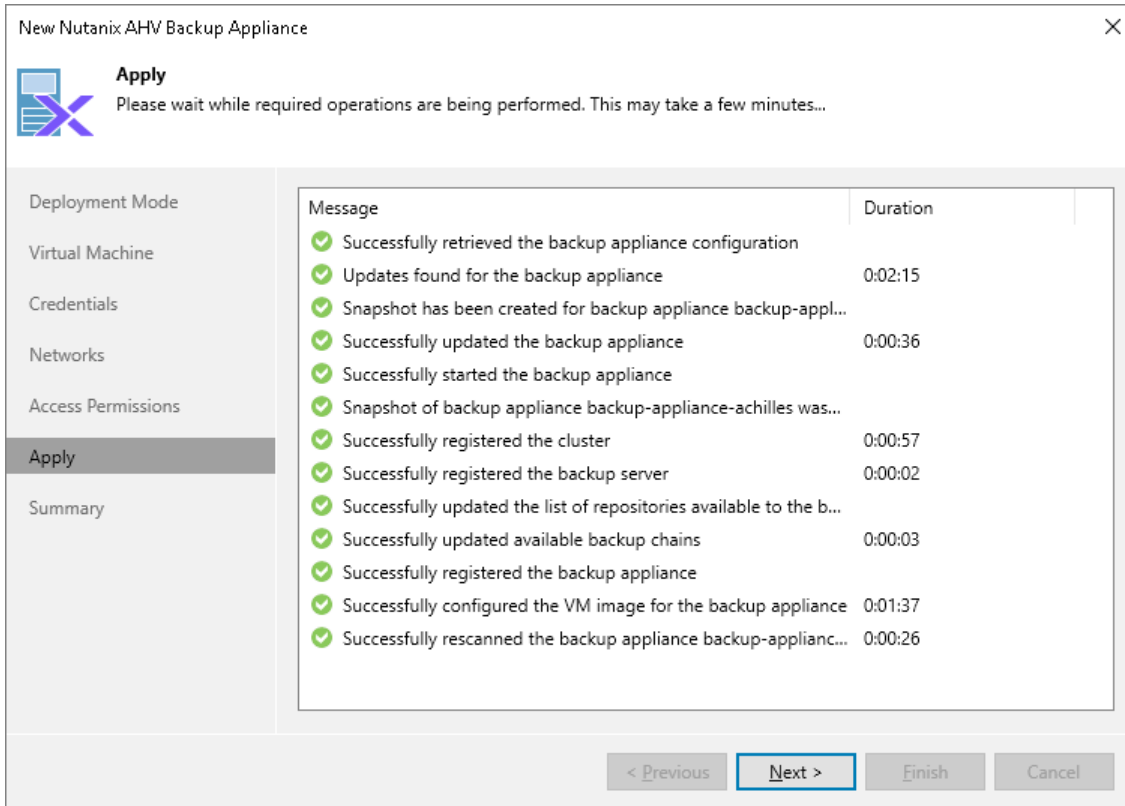
If you select the **Allow access to the following backup repositories** option, you must also specify backup repositories to which the Nutanix AHV backup appliance will have access. For a backup repository to be displayed in the **Repository** list, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section [Backup Repository](#).

The screenshot shows the 'New Nutanix AHV Backup Appliance' wizard window. The title bar reads 'New Nutanix AHV Backup Appliance' with a close button (X) on the right. Below the title bar is a logo and the text 'Access Permissions' followed by the instruction 'Specify the backup repositories this AHV proxy is allowed to access.' A sidebar on the left contains navigation options: 'Deployment Mode', 'Virtual Machine', 'Credentials', 'Networks', 'Access Permissions' (highlighted), 'Apply', and 'Summary'. The main area has two radio button options: 'Allow access to all backup repositories' (selected) and 'Allow access to the following backup repositories'. Below these is a list box titled 'Repository' containing one entry: 'Default Backup Repository' with an unchecked checkbox. To the right of the list box are 'Select All' and 'Clear All' buttons. At the bottom of the window are four buttons: '< Previous', 'Apply' (highlighted with a blue border), 'Finish', and 'Cancel'.



## Step 7. Apply Settings

At the **Apply** step of the wizard, wait for the Nutanix AHV backup appliance to be added to the backup infrastructure and then click **Next**.



New Nutanix AHV Backup Appliance

**Apply**  
Please wait while required operations are being performed. This may take a few minutes...

Message	Duration
✓ Successfully retrieved the backup appliance configuration	
✓ Updates found for the backup appliance	0:02:15
✓ Snapshot has been created for backup appliance backup-appl...	
✓ Successfully updated the backup appliance	0:00:36
✓ Successfully started the backup appliance	
✓ Snapshot of backup appliance backup-appliance-achilles was...	
✓ Successfully registered the cluster	0:00:57
✓ Successfully registered the backup server	0:00:02
✓ Successfully updated the list of repositories available to the b...	
✓ Successfully updated available backup chains	0:00:03
✓ Successfully registered the backup appliance	
✓ Successfully configured the VM image for the backup appliance	0:01:37
✓ Successfully rescanned the backup appliance backup-applianc...	0:00:26

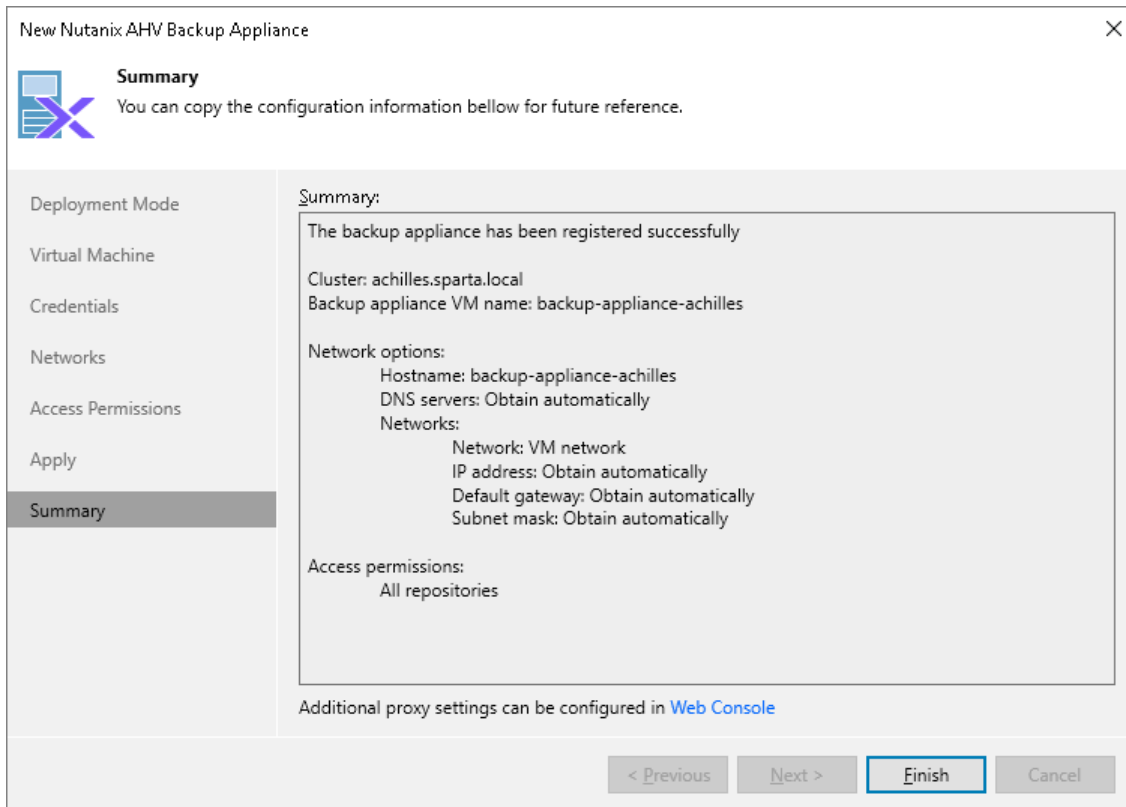
< Previous   **Next >**   Finish   Cancel

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. To open the Nutanix AHV backup appliance web console in your web browser, click **Web Console**.

### TIP

You can review details of the registration session in system logs as described in the Veeam Backup & Replication User Guide, section [Viewing History Statistics](#).



The screenshot shows the 'New Nutanix AHV Backup Appliance' wizard at the 'Summary' step. The window title is 'New Nutanix AHV Backup Appliance' with a close button (X) in the top right corner. On the left, there is a navigation pane with the following items: Deployment Mode, Virtual Machine, Credentials, Networks, Access Permissions, Apply, and Summary (which is highlighted). The main content area is titled 'Summary' and contains the following text: 'You can copy the configuration information below for future reference.' Below this is a large text box with the following summary information: 'Summary: The backup appliance has been registered successfully. Cluster: achilles.sparta.local. Backup appliance VM name: backup-appliance-achilles. Network options: Hostname: backup-appliance-achilles. DNS servers: Obtain automatically. Networks: Network: VM network. IP address: Obtain automatically. Default gateway: Obtain automatically. Subnet mask: Obtain automatically. Access permissions: All repositories.' At the bottom of the main content area, there is a link: 'Additional proxy settings can be configured in [Web Console](#)'. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

# Editing Backup Appliance

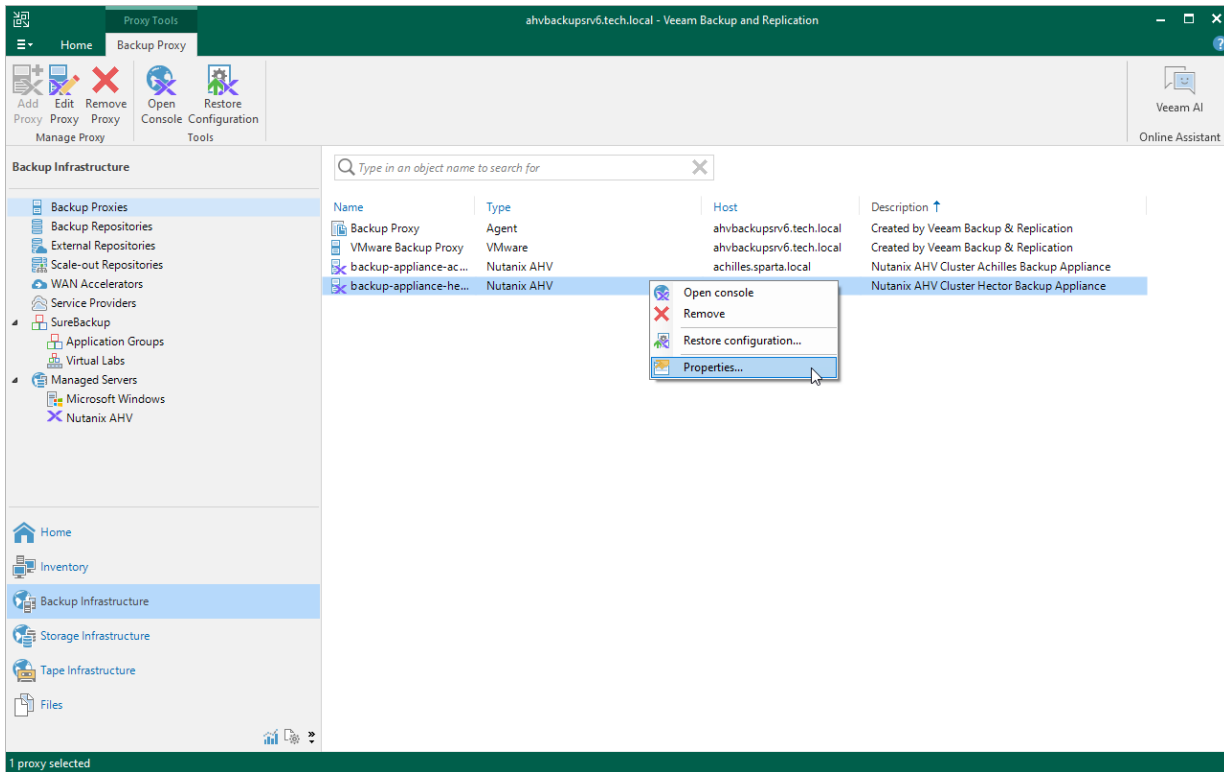
You can edit settings of the Nutanix AHV backup appliance that were specified while adding the appliance to the backup infrastructure.

To edit Nutanix AHV backup appliance settings, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the Nutanix AHV backup appliance and click **Edit Proxy** on the ribbon, or right-click the Nutanix AHV backup appliance and select **Properties**.
4. Complete the **Edit Nutanix AHV Backup Appliance** wizard:
  - a. To provide a new description for the Nutanix AHV backup appliance and change the number of tasks that the embedded worker is able to handle in parallel, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 3).
  - b. To change credentials for the account that is used to access the Nutanix AHV backup appliance web console, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 4).
  - c. To change the network to which the Nutanix AHV backup appliance is connected or to specify a new IP address for the appliance, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 5).
  - d. To specify backup repositories the Nutanix AHV backup appliance can access, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 6).
  - e. To save changes made to the proxy settings, click **Finish**.

## IMPORTANT

It is not recommended that you decrease the amount of allocated resources or modify the network settings while the embedded or a dedicated worker is currently transferring data. In this case, Veeam Backup for Nutanix AHV will terminate the related sessions and restart the appliance to update the settings immediately.



# Rescanning Backup Appliance

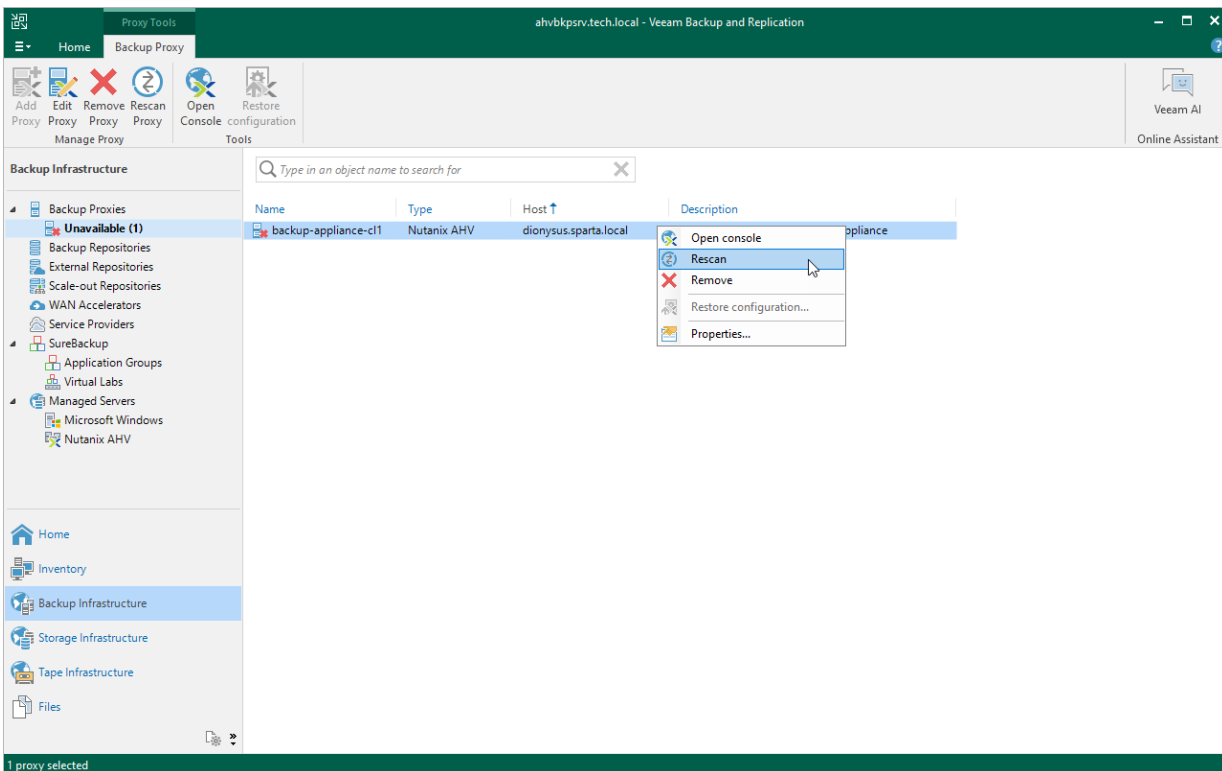
If the Nutanix AHV backup appliance becomes unavailable, you can rescan it to synchronize data with the backup server. The rescan operation will update the appliance configuration and backup job statistics on the backup server.

To rescan the Nutanix AHV backup appliance, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies > Unavailable**.
3. In the working area, select the Nutanix AHV backup appliance and click **Rescan Proxy** on the ribbon, or right-click the Nutanix AHV backup appliance and select **Rescan**.

## TIP

In the **System** window, you can track the progress of the rescan session. You can close the window and check session details later as described in the Veeam Backup & Replication User Guide, section [Viewing History Statistics](#).



# Removing Backup Appliance

You can remove the Nutanix AHV backup appliance from the backup infrastructure if you no longer need it and want to add another appliance to the backup server, or if you want to connect this appliance to another backup server.

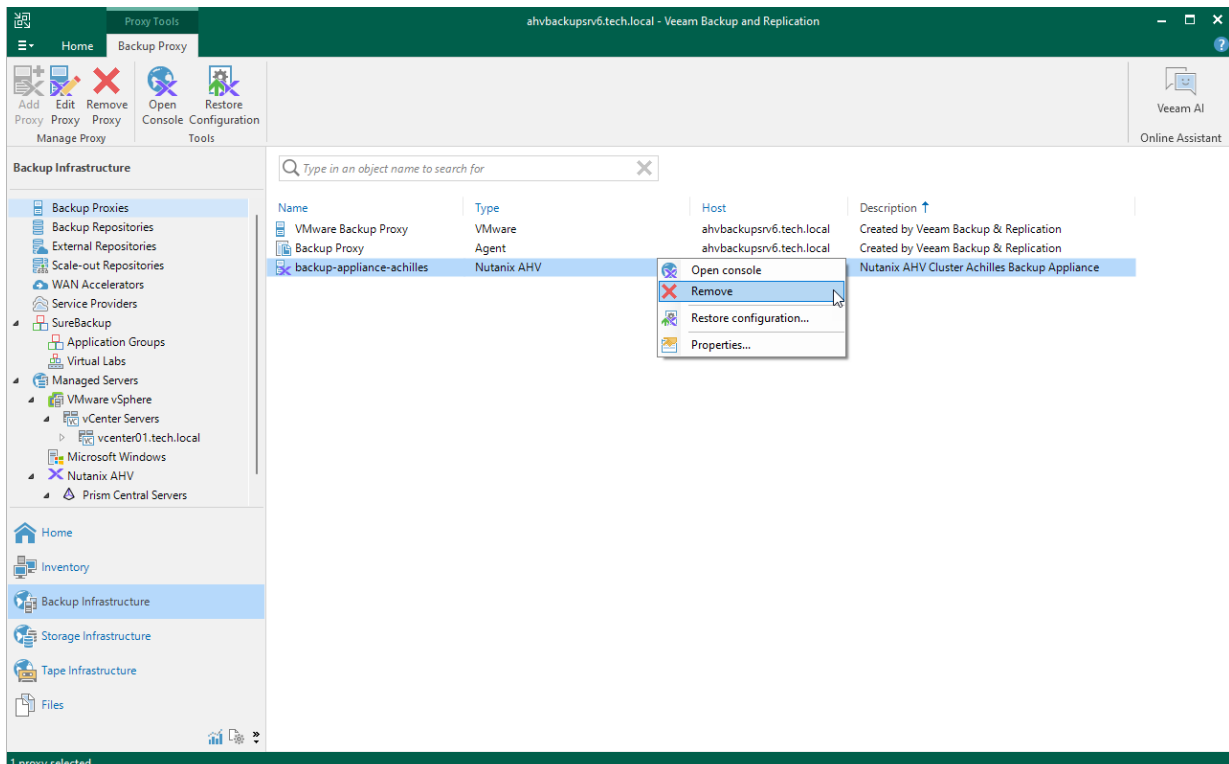
## IMPORTANT

After you remove the Nutanix AHV backup appliance:

- You will not be able to perform [VM backup](#), [entire VM restore](#) and [VM disk restore](#) operations unless you deploy a new Nutanix AHV backup appliance. However, you will still be able to [manage Nutanix AHV VM backups](#) and perform all other restore operation described in section [Performing Restore](#).
- Records about all jobs that have been ever processed by the Nutanix AHV backup appliance will be deleted from the Veeam Backup & Replication configuration database. Backups created by these jobs are displayed under the **Backups > Disk (Orphaned)** node in the **Home** view of the Veeam Backup & Replication console.

To remove the Nutanix AHV backup appliance, do the following:

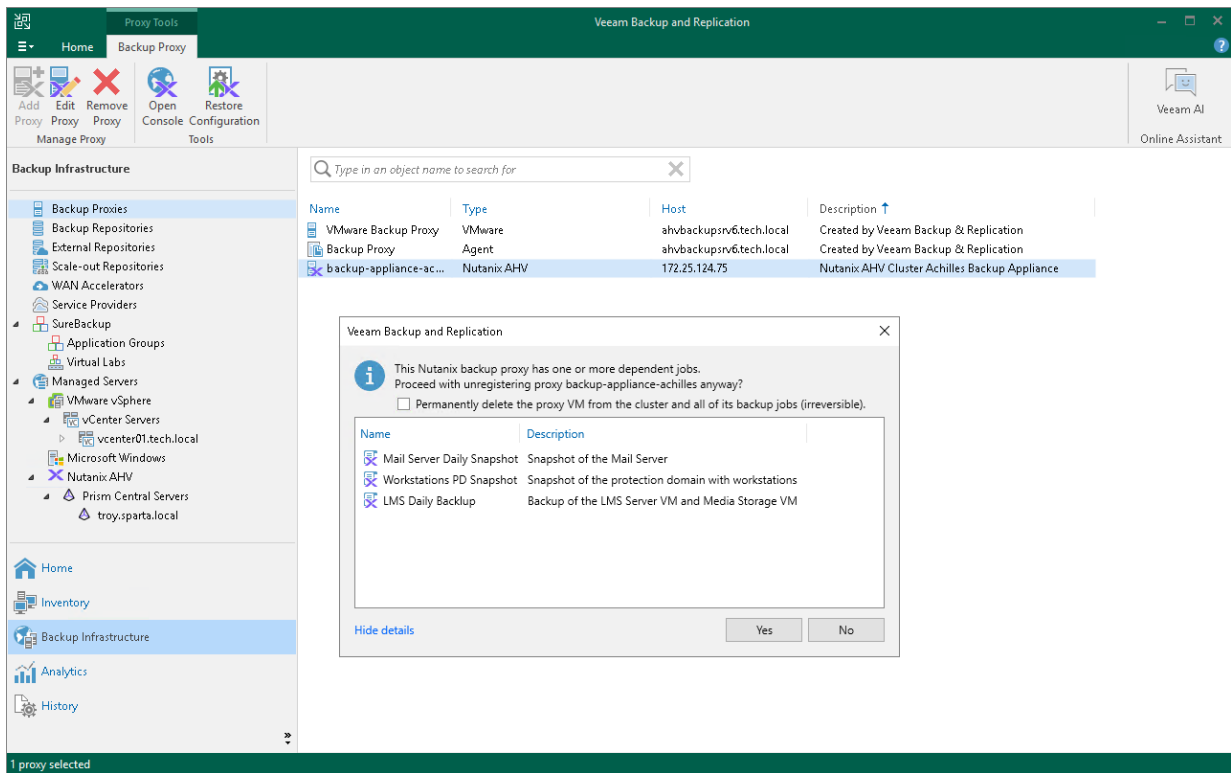
1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the Nutanix AHV backup appliance and click **Remove Proxy** on the ribbon, or right-click the Nutanix AHV backup appliance and select **Remove**.



4. In the **Veeam Backup & Replication** window, choose whether you want to permanently remove the VM running as the Nutanix AHV backup appliance.

## TIP

If you keep the VM, the configuration settings and records about backup jobs ever processed by the appliance will be retained in the appliance database. This can be helpful if you want to [connect the Nutanix AHV backup appliance to another backup server](#).



# Migrating Backup Appliance

You can migrate a Nutanix AHV backup appliance to a new backup infrastructure if you want to keep the specified appliance settings and configured jobs:

1. Add the Prism Central or Nutanix AHV cluster where the Nutanix AHV backup appliance resides to the new backup infrastructure as described in section [Adding Nutanix AHV Server](#).
2. Remove the Nutanix AHV backup appliance from the old backup infrastructure as described section in [Removing Backup Appliance](#).

If you do not remove the Nutanix AHV backup appliance, it will not function properly since the old backup server will try to manage the appliance.

## IMPORTANT

While removing the appliance, do NOT select the **Permanently delete the proxy VM from the cluster and all of its backup jobs (irreversible)** check box. Otherwise, the Nutanix AHV backup appliance VM will be removed from the cluster.

3. Connect the Nutanix AHV backup appliance to the new backup server as described in section [Connecting Existing Backup Appliance](#).
4. Specify new backup repositories in the jobs configured on the appliance as described in section [Editing Job Settings](#).
5. Rescan the Nutanix AHV backup appliance to synchronize data with the backup server as described in section [Rescanning Backup Appliance](#).
6. Import Nutanix AHV backups from new backup repositories as described in section [Rescanning Backups](#).

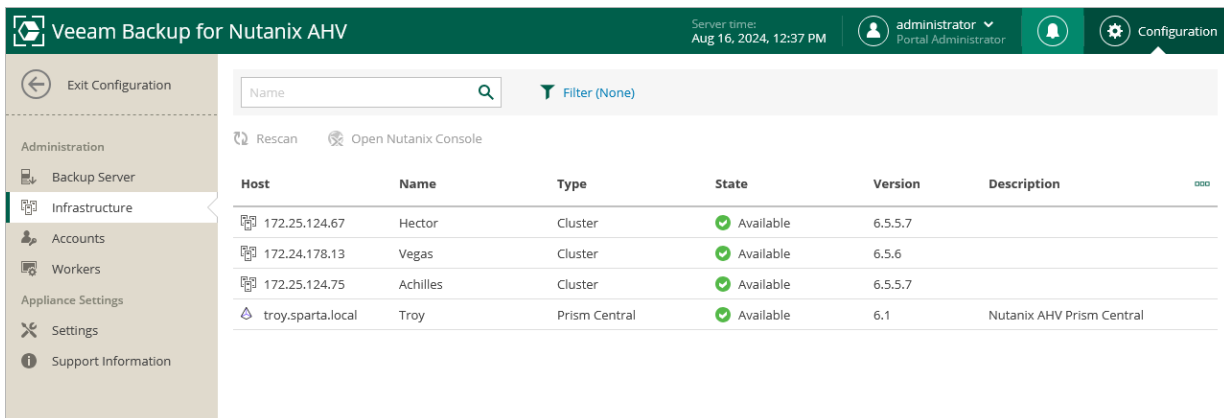


# Troubleshooting Backup Appliance

If the Nutanix AHV backup appliance becomes unavailable in the Veeam Backup & Replication console, you can do the following:

- [Review the backup appliance settings](#) to resolve connection issues.
- [Rescan the backup appliance](#) to resolve synchronization issues.

If the Nutanix AHV backup appliance remains unavailable, [open the web console](#) and switch to the **Configuration** page. Then, in the **Administration** section, check whether the backup appliance is able to connect both to the backup server and to the Nutanix AHV Prism Central or cluster. If you encounter any connection issues, check your network configuration – the appliance must be able to resolve the backup server and cluster domain names to their IP addresses.



The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the title "Veeam Backup for Nutanix AHV", the server time "Aug 16, 2024, 12:37 PM", the user "administrator", and the "Configuration" tab. The left sidebar contains navigation options: Administration, Backup Server, Infrastructure (selected), Accounts, Workers, Appliance Settings, Settings, and Support Information. The main content area features a search bar, a "Filter (None)" button, and two action buttons: "Rescan" and "Open Nutanix Console". Below these is a table listing backup appliances.

Host	Name	Type	State	Version	Description
172.25.124.67	Hector	Cluster	Available	6.5.5.7	
172.24.178.13	Vegas	Cluster	Available	6.5.6	
172.25.124.75	Achilles	Cluster	Available	6.5.5.7	
troy.sparta.local	Troy	Prism Central	Available	6.1	Nutanix AHV Prism Central

# Managing Workers

To perform most data protection and disaster recovery operations (such as creating image-level backups and restoring backed-up data from them), Veeam Backup for Nutanix AHV uses workers. Workers are Linux-based VMs that are responsible for the interaction between the backup appliance and other Veeam Backup for Nutanix AHV components. Workers process backup workload and distribute backup traffic when transferring data to backup repositories.

By default, the worker role is assigned to the backup appliance. However, this is sufficient only [for small deployments](#). [For large deployments](#), it is recommended to deploy dedicated workers as the embedded worker may not have enough bandwidth to process backup traffic. Deploying dedicated workers allows you to increase the maximum number of concurrent backup and restore operations, and to avoid high traffic load on the host running the backup appliance.

Each dedicated worker is launched on a specific host for the duration of a backup or restore operation. While configuring the worker, you can manually select the host or instruct Veeam Backup for Nutanix AHV to choose a host automatically. Manual selection may be helpful if you want to avoid launching workers on specific hosts (for example, production ones), while automatic selection allows Veeam Backup for Nutanix AHV to optimize data transfer and to balance the load on the hosts in the Nutanix AHV cluster. In the latter case, Veeam Backup for Nutanix AHV uses the [AHV affinity functionality](#) to distribute workers among all hosts in the cluster instead of launching multiple workers on one host.

## Worker Lifecycle

When you add a dedicated worker to the backup infrastructure, its configuration is saved to the Veeam Backup for Nutanix AHV configuration database, but no VM is actually deployed in the cluster unless you choose to test the configuration. In the latter case, a VM (worker VM) is deployed and shut down after the test operation completes.

As soon as a backup or restore session starts, Veeam Backup for Nutanix AHV tries to launch the worker and test its configuration. If no worker VM has been previously deployed, Veeam Backup for Nutanix AHV deploys the VM using the worker configuration saved to the configuration database. Veeam Backup for Nutanix AHV checks host affinity settings specified for the worker and chooses a host where the worker VM will run. Then, Veeam Backup for Nutanix AHV powers on the worker VM and installs system updates (if available). When the backup or restore session completes, Veeam Backup for Nutanix AHV shuts down the worker VM so that it can be used for other sessions later.

During the lifecycle, a worker can obtain one of the following statuses:

- **Configured** – the worker configuration is added to the Veeam Backup for Nutanix AHV configuration database.
- **Testing** – the worker configuration is being tested.
- **Updating** – the worker or its configuration is being updated.
- **Working** – the worker is processing a backup or restore operation.
- **Shut Down** – the worker is powered off.
- [Applies only to the embedded worker] **Ready** – the worker is powered on and can be used for data protection operations.

# Adding Workers

To deploy a worker and add it to the backup infrastructure, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Worker wizard.](#)
3. [Specify a worker name.](#)
4. [Specify worker VM configuration.](#)
5. [Specify worker network settings.](#)
6. [Finish working with wizard.](#)

## Before You Begin

Before you add a dedicated worker to the backup infrastructure, consider the following:

- [Applies only to the [Prism Central deployment](#)] It is recommended that workers are deployed in each cluster registered with the Prism Central. If no worker is deployed in the cluster, performance of backup operations will be affected as Veeam Backup for Nutanix AHV will use the embedded worker or a worker deployed in another cluster. Note that you can [enable notifications](#) to get informed when nonoptimal workers process VMs.
- It is recommended that the number of configured workers does not exceed the number of hosts in the Nutanix AHV cluster.
- Each worker must be provided with sufficient compute resources to handle backup and restore tasks in parallel. The maximum number of concurrent tasks is configured in worker settings – if this number is exceeded, the worker will not start a new task until one of the current tasks finishes.
- It is recommended the total number of concurrent tasks configured for all workers deployed in the cluster does not exceed the [number of physical disks added to the cluster](#). You can change the maximum number of concurrent tasks (the best practice is to allocate 1 vCPU and 1 GB RAM for each additional task) while deploying a new worker or editing settings of an existing one.
- If you plan to use dedicated workers, it is recommended that you [modify the backup appliance settings](#) as follows:
  - Set the maximum number of concurrent tasks to 0.
  - Change the amount of allocated resources to according to the recommendations described in section [Sizing Guidelines](#).

Alternatively, you can [disable the embedded worker](#). However, in this case, the backup appliance will be allocated excessive compute resources that will be not used for handling backup traffic.

# Step 1. Launch Add Worker Wizard

To launch the **Add Worker** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Click **Add**.

The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Nutanix AHV', the server time 'Aug 15, 2024, 01:02 PM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Administration (Backup Server, Infrastructure, Accounts, Workers), and Appliance Settings (Settings, Support Information). The main content area is titled 'Workers' and includes a description: 'Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.' Below the description, there is a toggle for 'Obtain updates from rpm repositories' which is 'Enabled'. A search bar contains the text 'Worker'. Below the search bar are action buttons: '+ Add', 'Edit', 'Enable', 'Remove', and 'Test'. A table lists the workers with columns: Worker, Cluster, Affinity, Status, Last Test, Max Concurrency, and Description. The table shows two workers: 'backup-ap...' (Configured) and 'Embedded' (Ready).

Worker	Cluster	Affinity	Status	Last Test	Max Concurrency	Description
backup-ap...	Hector	—	Configured	—	4	Created instead...
Embedded	Achilles	—	Ready	—	4	Nutanix AHV CL...

## Step 2. Specify Worker Name

At the **Worker Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the worker and to provide a description for future reference. The worker name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are only supported: a-z, A-Z, 0-9, -. The maximum length of the description is 1024 characters.

The screenshot shows the 'Add Worker' wizard in Veeam Backup for Nutanix AHV. The top navigation bar includes the product name, server time (Aug 15, 2024, 01:04 PM), user 'administrator', and a 'Configuration' menu. The left sidebar lists 'Worker Info', 'Virtual Machine', 'Network', and 'Summary'. The main area is titled 'Worker Info' and contains the instruction 'Enter a name and description for the worker.' Below this are two input fields: 'Name' with the value 'worker-hector-b' and 'Description' with the value 'Worker residing on host 2, cluster Hector'. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Step 3. Specify Worker VM Settings

At the **Virtual Machine** step of the wizard, do the following:

1. [Applies only to the [Prism Central deployment](#)] Click the link in the **Cluster** field, and specify in the **Choose cluster** window a cluster where the worker will reside.

For a cluster to be displayed in the list of the available cluster, it must be configured in the Nutanix AHV Prism Central as described in [Nutanix documentation](#).

2. Check the **The worker VM will store its files in the container** field to see the storage container that is automatically selected for worker system file.
3. In the **Maximum concurrent tasks** field, specify the number of tasks that the worker will be able to handle in parallel. If this value is exceeded, the worker will not start processing a new task until one of the currently running tasks finishes.

The default number of concurrent tasks is set to 4. When you change this value, the wizard automatically adjusts the amount of resources that will be allocated to the VM running as the worker. If you want to specify the amount of resources manually, click the link below the **Maximum concurrent tasks** field.

### NOTE

When performing data protection and disaster recovery operations, Veeam Backup for Nutanix AHV initiates a new task for each VM that is being processed.

4. To specify a host where the worker will be launched, set the **Host affinity** toggle to *On* and click the link in the **The worker VM will run on the host** field.

If you do not specify host affinity settings, Veeam Backup for Nutanix AHV will automatically define the host to launch the worker.

The screenshot shows the 'Add Worker' wizard in the Veeam Backup for Nutanix AHV interface. The 'Virtual Machine' step is active, showing configuration settings for the worker VM. The settings are as follows:

- Cluster:** Hector
- Storage container:** SelfServiceContainer
- Compute resources:** Maximum concurrent tasks: 4. A link below indicates 'The VM will be allocated 6 vCPU and 6 GB RAM'.
- Affinity settings:** Host affinity is toggled 'On'. The worker VM will run on the host: NTNX-235G5H090054-A.

Navigation buttons at the bottom include 'Previous', 'Next', and 'Cancel'.

# Step 4. Configure Network Settings

At the **Network** step of the wizard, do the following:

1. Click **Add** to configure a worker network interface:
  - a. In the **Description** field, provide a network interface description for future reference.
  - b. From the **Network** drop-down list, select a network to which the worker network interface will be connected.

For best performance, choose the network that is used for Nutanix Controller VM (CVM). For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

- c. If DHCP is enabled in the selected network, the IP address of the worker can be obtained automatically.

If DHCP is disabled in the selected network, or you want to specify an IP address, select the **Use the following IP address** option and enter the worker IP address, subnet mask and default gateway.

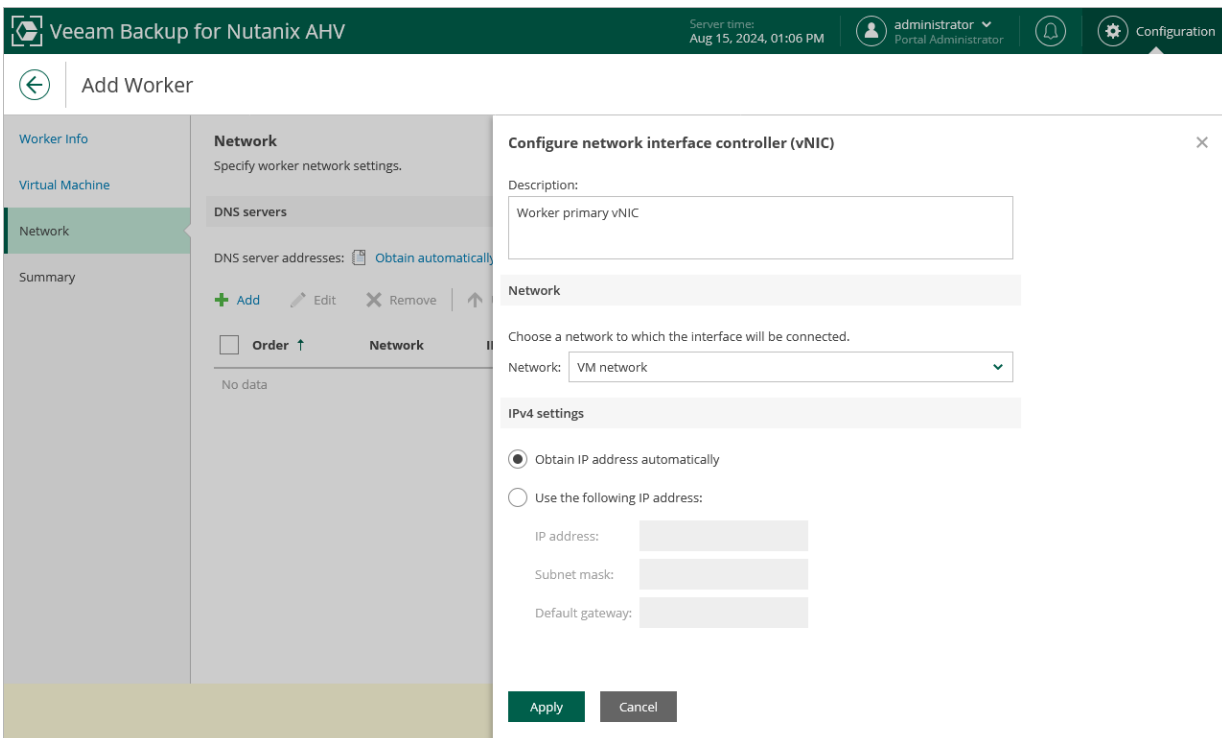
To add more network interfaces, repeat the step and specify the network order using the **Up** and **Down** buttons. For more information on multi-network configuration, see section [Appendix D. Configuring Multiple Networks](#).

2. If DHCP is enabled in the selected network, DNS settings of the worker can be obtained automatically.

If DHCP is disabled in the selected network, or you want to configure DNS settings manually, click the link in the **DNS server addresses** field, select the **Use the following DNS server address** option and enter the IP addresses of the preferred and alternate DNS servers.

## NOTE

Since workers are Linux-based VMs, they have the same limitations that apply to machines running the Rocky Linux operating system. That is, DNS settings cannot be configured separately for each network added to the worker.



# Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

## TIP

If you do not want to test the worker, clear the **Test worker configuration when I click Finish** check box and then click **Finish**.

The screenshot shows the 'Add Worker' wizard in the Veeam Backup for Nutanix AHV interface. The 'Summary' step is active, displaying job summary information and configuration details. The interface includes a top navigation bar with the product name, server time, user information, and configuration settings. A left sidebar lists navigation options: Worker Info, Virtual Machine, Network, and Summary. The main content area is divided into sections for General settings and Network settings, each with a list of key-value pairs. At the bottom, there is a checkbox for 'Test worker configuration when I click Finish' and three buttons: Previous, Finish, and Cancel.

General settings	
Name:	worker-hector-b
Description:	Worker residing on host 2, cluster Hector
Storage container:	SelfServiceContainer
Max concurrent tasks:	4
Resources:	6 vCPU, 6 GB RAM
Host affinity:	NTNX-23SG5H090054-A

Network settings	
DNS server addresses	Obtained automatically
vNIC description:	Worker primary vNIC
Network:	VM network
IP address:	Obtained automatically
Subnet mask:	Obtained automatically
Default gateway:	Obtained automatically

Test worker configuration when I click Finish

Previous Finish Cancel



# Enabling and Disabling Workers

By default, workers are launched when jobs or restore sessions start. However, you can temporarily disable a worker – this may be helpful when you reconfigure a worker and you do not want it to be used for a backup or restore operation. You will still be able to test or enable the disabled worker at any time you need.

To enable or disable a worker, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Select the worker.
4. Click **Enable** or **Disable**.

The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Nutanix AHV', the server time 'Aug 15, 2024, 01:08 PM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Workers' selected. The main content area displays a description of workers and a table of active workers. The 'Obtain updates from rpm repositories' toggle is set to 'Enabled'. The table has columns for Worker, Cluster, Affinity, Status, Last Test, Max Concurrency, and Description. One worker, 'worker-hec...', is selected and highlighted in green, with its status set to 'Testing' and 'Running'.

Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.

Obtain updates from rpm repositories  Enabled ⓘ

Worker  Filter (None) Cluster: All clusters

+ Add Edit Disable Remove Test

<input type="checkbox"/>	Worker ↑	Cluster	Affinity	Status	Last Test	Max Concurr...	Description	⋮
Selected: 1 of 3								
<input type="checkbox"/>	backup-ap...	Hector	—	Configured	—	4	Created instead...	
<input type="checkbox"/>	Embedded	Achilles	—	Ready	—	4	Nutanix AHV CL...	
<input checked="" type="checkbox"/>	worker-hec...	Hector	NTNX-23SG5H0...	Testing	Running	4	Worker residing...	

# Editing Workers

For each worker, you can modify settings specified while adding the worker to the backup infrastructure:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Select the worker and click **Edit**.
4. Complete the **Edit Worker** wizard:
  - a. To provide a new name and description for the worker, follow the instructions provided in section [Adding Workers](#) (step 2).
  - b. To change the cluster where worker resides, to specify a host where the worker is launched, or to modify the number of tasks that the worker is able to handle in parallel, follow the instructions provided in section [Adding Workers](#) (step 3).
  - c. To change the network to which the worker is connected or to specify a new IP address for the worker, follow the instructions provided in section [Adding Workers](#) (step 4).
  - d. To save changes made to the worker settings, click **Finish**.

## IMPORTANT

It is not recommended that you change the worker cluster, decrease the amount of allocated resources, adjust the affinity settings or modify the network settings while the worker is currently transferring data. In this case, Veeam Backup for Nutanix AHV will terminate the related operations, power off the worker and update the settings immediately.

Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.

Obtain updates from rpm repositories  Enabled ⓘ

Worker  Filter (None) Cluster: All clusters

+ Add Edit Disable Remove Test

Worker ↑	Cluster	Affinity	Status	Last Test	Max Concurr...	Description
<input type="checkbox"/> backup-ap...	Hector	—	⚙️ Configured	—	4	Created instead...
<input type="checkbox"/> Embedded	Achilles	—	✅ Ready	—	4	Nutanix AHV CL...
<input checked="" type="checkbox"/> worker-hec...	Hector	NTNX-23SG5H0...	🕒 Testing	🟢 Running	4	Worker residing...

# Testing Workers

Before using a dedicated worker for a backup or restore operation, Veeam Backup for Nutanix AHV automatically tests its configuration – verifies that the worker service can start successfully, checks that the worker can connect to the backup server and to the cluster, and installs available updates.

If you want to ensure that the worker configuration is correct before it is used for a backup or restore operation, you can start a worker configuration test manually:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Select the worker and click **Test**.

Note that you can select and test multiple workers at once.

As soon as Veeam Backup for Nutanix AHV finishes the worker configuration test, the worker will be powered off. To see detailed information on a worker test session, click the worker test status in the **Last Test** column.

Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.

Obtain updates from rpm repositories  Enabled ⓘ

Worker  Filter (None) Cluster: All clusters

+ Add Edit Disable Remove Test

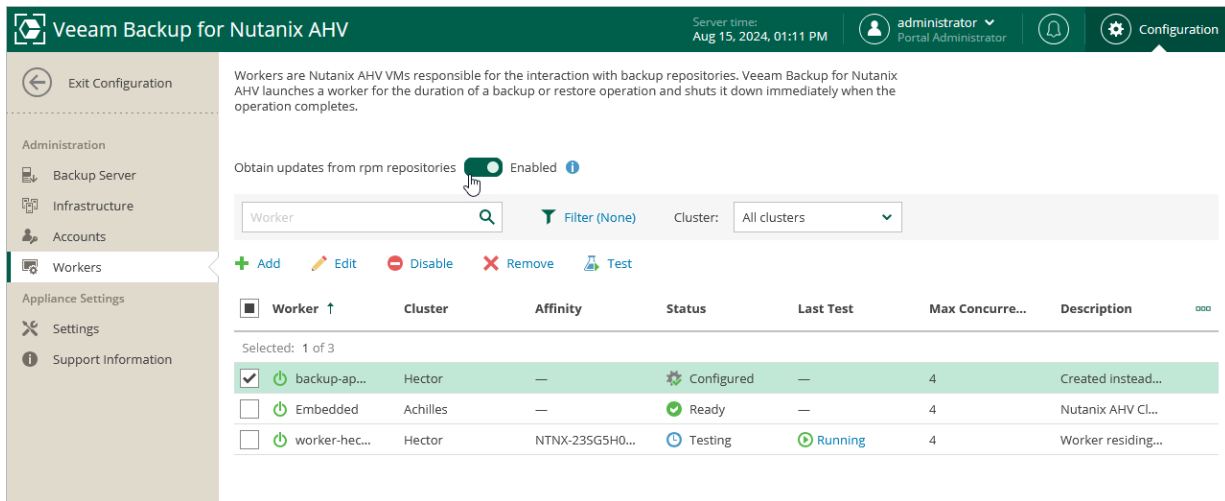
Worker ↑	Cluster	Affinity	Status	Last Test	Max Concurr...	Description
<input checked="" type="checkbox"/> backup-ap...	Hector	—	Configured	—	4	Created instead...
<input type="checkbox"/> Embedded	Achilles	—	Ready	—	4	Nutanix AHV Cl...
<input type="checkbox"/> worker-hec...	Hector	NTNX-23SG5H0...	Testing	Running	4	Worker residing...

# Disabling Automatic Worker Updates

When launching a worker for a backup or restore operation, Veeam Backup for Nutanix AHV automatically downloads updates from Veeam repositories and installs them on the worker. If the worker is not connected to the internet, you can instruct Veeam Backup for Nutanix AHV to [use a web proxy](#) that will provide access to the necessary resources.

If workers do not have access to the internet and no web proxy is used, you can disable automatic updates to avoid connection failures and eliminate session warnings:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Set the **Obtain updates from deb repositories** toggle to *Disabled*.



The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Nutanix AHV', the server time 'Aug 15, 2024, 01:11 PM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Workers' selected. The main content area displays the 'Workers' configuration page. At the top, there is a description: 'Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.' Below this, the 'Obtain updates from rpm repositories' toggle is set to 'Enabled'. A search bar contains the text 'Worker'. Below the search bar are buttons for '+ Add', 'Edit', 'Disable', 'Remove', and 'Test'. A table lists the workers with columns for 'Worker', 'Cluster', 'Affinity', 'Status', 'Last Test', 'Max Concurr...', and 'Description'. The first worker, 'backup-ap...', is selected and has a status of 'Configured'. The second worker, 'Embedded', has a status of 'Ready'. The third worker, 'worker-hec...', has a status of 'Testing' and is currently 'Running'.

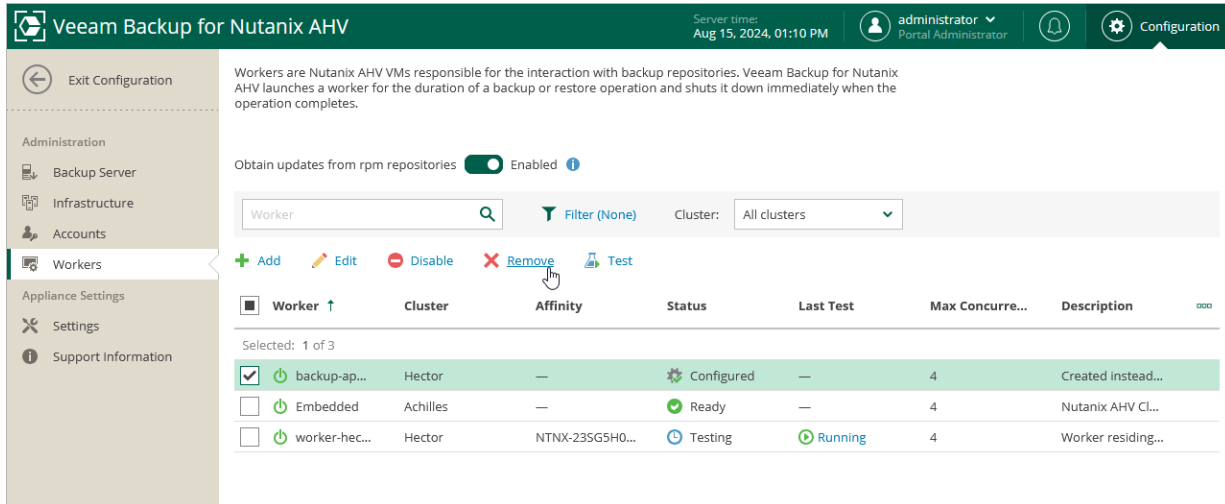
Worker	Cluster	Affinity	Status	Last Test	Max Concurr...	Description
Selected: 1 of 3						
<input checked="" type="checkbox"/> backup-ap...	Hector	—	Configured	—	4	Created instead...
<input type="checkbox"/> Embedded	Achilles	—	Ready	—	4	Nutanix AHV Cl...
<input type="checkbox"/> worker-hec...	Hector	NTNX-235G5H0...	Testing	Running	4	Worker residing...

# Removing Workers

Veeam Backup for Nutanix AHV allows you to permanently remove workers if you no longer need them. Note that you cannot remove a worker while it is transferring data for a backup or restore operation.

To remove a worker, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers**.
3. Select the worker and click **Remove**.



The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Nutanix AHV', the server time 'Aug 15, 2024, 01:10 PM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Workers' selected. The main content area displays a description of workers, a toggle for 'Obtain updates from rpm repositories' (Enabled), a search bar, and a table of workers. The 'Remove' button is highlighted over the 'Remove' link in the table.

Workers are Nutanix AHV VMs responsible for the interaction with backup repositories. Veeam Backup for Nutanix AHV launches a worker for the duration of a backup or restore operation and shuts it down immediately when the operation completes.

Obtain updates from rpm repositories  Enabled ⓘ

Worker  Filter (None) Cluster: All clusters

+ Add Edit Disable Remove Test

Worker ↑	cluster	Affinity	Status	Last Test	Max Concurr...	Description	
Selected: 1 of 3							
<input checked="" type="checkbox"/>	backup-ap...	Hector	—	Configured	—	4	Created instead...
<input type="checkbox"/>	Embedded	Achilles	—	Ready	—	4	Nutanix AHV Cl...
<input type="checkbox"/>	worker-hec...	Hector	NTNX-235G5H0...	Testing	Running	4	Worker residing...

# Accessing Backup Appliance

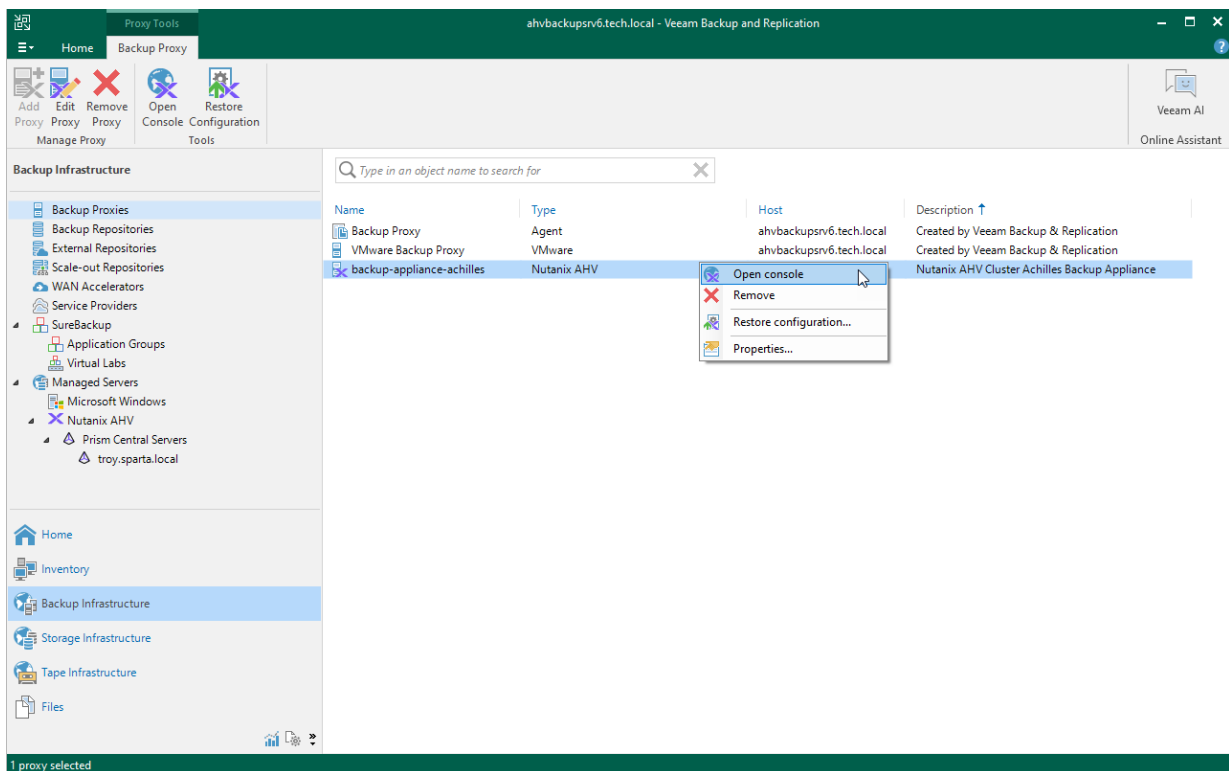
To access Veeam Backup for Nutanix AHV, in a web browser, navigate to the Nutanix AHV backup appliance web address. The address consists of an IPv4 address or DNS hostname of the backup appliance. Note that the website is available over HTTPS only.

## IMPORTANT

The Internet Explorer browser is not supported. To access Nutanix AHV backup appliance, use Microsoft Edge (latest version), Mozilla Firefox (latest version) or Google Chrome (latest version).

To access the Nutanix AHV backup appliance from the Veeam Backup & Replication console, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the Nutanix AHV backup appliance and click **Open Console** on the ribbon, or right-click the Nutanix AHV backup appliance and select **Open Console**.



Alternatively, you can navigate to the IP address or hostname of the Nutanix AHV backup appliance in a web browser.

## NOTE

The web browser may display a warning notifying that the connection is untrusted. To eliminate the warning, you can replace the TLS certificate that is currently used to secure traffic between the browser and the backup appliance with a trusted TLS certificate. To learn how to replace certificates, see [Appendix B. Installing Custom Certificate](#).

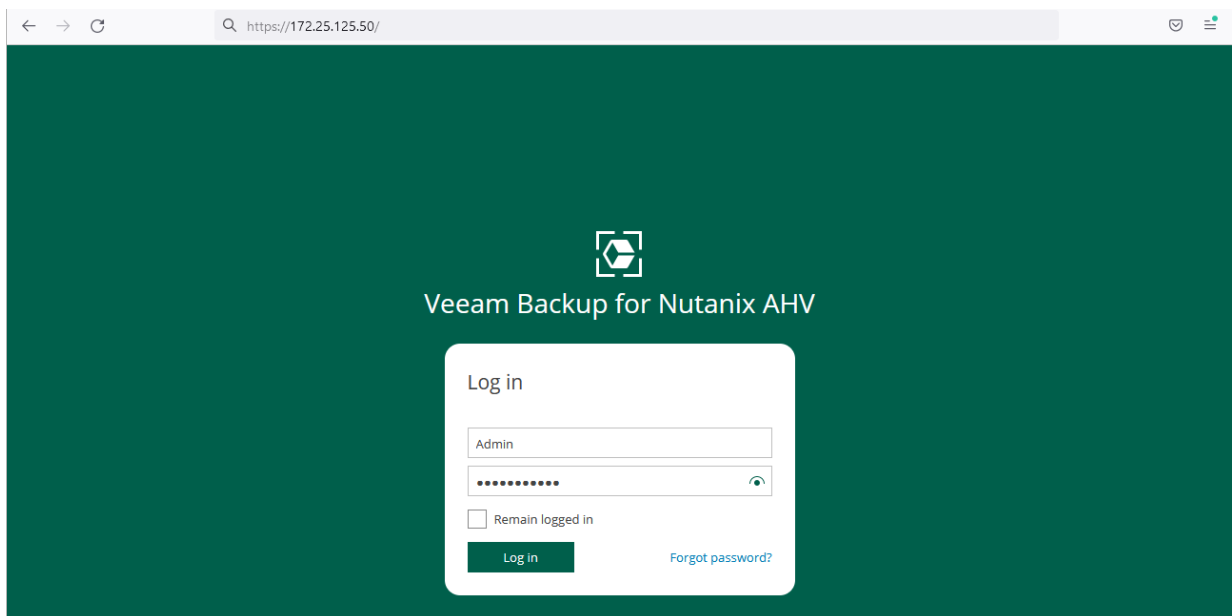
4. On the welcome screen, specify credentials of the user account.

If you are an administrator and you log in for the first time, use credentials of the account that you have specified during the [initial backup appliance configuration](#). In future, you can add other user accounts to grant access to Veeam Backup for Nutanix AHV. For more information, see [Managing Users](#).

#### TIP

If you do not remember the password, you can reset it by running the Nutanix AHV backup appliance in recovery mode. For more information, see [this Veeam KB article](#).

5. [Applies only if [MFA is enabled for your user account](#)] Enter a verification code sent to your trusted device.
6. Select the **Remain logged in** check box to save the specified credentials in a persistent browser cookie so that you do not have to provide credentials every time you access the backup appliance in a new browser session.
7. Click **Log in**.




## Logging Out

To log out, at the top right corner of the web console window, click your user name and then click **Log out**.

# Changing Password

To change the password that you use to access the Veeam Backup for Nutanix AHV web console, do the following:

1. At the top right corner of the web console window, click your user name.
2. Click **Change Password**.
3. In the **Change Password** window, enter the currently used password, enter and confirm a new password, and then click **Change**.





# Configuring Veeam Backup for Nutanix AHV

To configure Veeam Backup for Nutanix AHV, you can perform the following administrative tasks:

- [Manage users and permissions.](#)
- [Enable SSH access to the backup appliance.](#)
- [View network settings.](#)
- [Configure time synchronization settings.](#)
- [Configure settings for update and resource usage notifications.](#)
- [Configure notification settings for automated delivery of backup job results.](#)
- [Back up and restore appliance configuration.](#)

# Managing Users

Veeam Backup for Nutanix AHV controls access to its functionality with the help of user roles. A role defines what operations users can perform and what range of data is available to them in Veeam Backup for Nutanix AHV.

There are 3 user roles that you can assign to users working with Veeam Backup for Nutanix AHV. Actions a user can perform depend on the role.

- **Portal Administrator** – can perform all configuration actions and can also act as a Portal Operator and Restore Operator.
- **Portal Operator** – can create and manage backup jobs, manage the protected data and perform all restore operations.
- **Restore Operator** – can only perform restore operations.

The following table describes the functionality available to users with different roles in the Veeam Backup for Nutanix AHV web console.

Tab	Functionality	Portal Administrator	Portal Operator	Restore Operator
<b>Dashboard</b>	Overview of protected resources	Full	Full	N/A
<b>Jobs</b>	Backup jobs, snapshot jobs, PD snapshot jobs	Full	Full (except for PD backup and snapshot job custom scripts)	N/A
<b>Protected VMs</b>	Restore	Full	Full	Full
	Remove snapshots	Full	Full	N/A
<b>Session Log</b>	Session log	Full	Full	Full
	Stop session execution	Full	Full	Restore sessions only
<b>Configuration</b>				
<b>Backup Server</b>	Backups rescan	Full	N/A	N/A
<b>Nutanix AHV Cluster</b>	Snapshots rescan, auto-protection settings	Full	N/A	N/A

Tab	Functionality	Portal Administrator	Portal Operator	Restore Operator
<b>Accounts</b>	Portal users	Full	N/A	N/A
<b>Workers</b>	Worker instances	Full	N/A	N/A
<b>Settings</b>	Network settings, notifications, time synchronization, mail server configuration	Full	N/A	N/A
<b>Support Information</b>	Updates and logs	Full	N/A	N/A

# Adding User Accounts

To manage access to Veeam Backup for Nutanix AHV, you can create local user accounts.

To add a Veeam Backup for Nutanix AHV user account, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts**.
3. Click **Add**.
4. Complete the **Add Account** wizard:
  - a. At the **Account Info** step of the wizard, specify a name and description for the user account.

The account name must start with a lowercase Latin letter and must not match Linux system user names (such as *root*, *daemon*). The name can contain only lowercase Latin letters, numeric characters, underscores and dashes. The maximum length of the name is 32 characters, the maximum length of the description is 1024 characters.
  - b. At the **General Settings** step of the wizard, select a role and specify a password for the user account. For more information on user roles, see section [Managing Users](#).
  - c. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for Nutanix AHV', server time 'Aug 15, 2024, 11:56 AM', and user information 'administrator Portal Administrator'. The wizard has three steps: 'Account Info', 'General Settings', and 'Summary', with 'Summary' currently selected. The 'Summary' step displays the following account details:

Account	
Name:	p_operator
Description:	Created by administrator (backup-appliance-achilles) at 8/15/2024, 11:54 AM
Role:	Restore Operator
Password:	*****

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

# Editing User Account Settings

For each user account added to the Veeam Backup for Nutanix AHV configuration database, you can modify the settings of the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts**.
3. Select the user account and click **Edit**.
4. Complete the **Edit Account** wizard.
  - a. At the **Account Info** step of the wizard, edit a description of the user account.
  - b. At the **General Settings** step of the wizard, select a new role for the user account.
  - c. At the **Summary** step of the wizard, review summary information and click **Finish**.


The screenshot shows the 'Edit Account' wizard in the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name, the server time (Aug 15, 2024, 11:59 AM), the user (administrator, Portal Administrator), and the Configuration page. The wizard has three steps: Account Info, General Settings, and Summary. The Summary step is currently active. It displays the account details for 'p\_operator', including its description and role. At the bottom, there are three buttons: Previous, Finish, and Cancel.

Account	
Name:	p_operator
Description:	Created by administrator (backup-appliance-achilles) at 8/15/2024, 11:54 AM
Role:	Restore Operator

# Changing User Passwords

As an administrator, you can change the password for a Veeam Backup for Nutanix AHV user account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts**.
3. Select the user account and click **Change Password**.
4. In the **Change Password** window, enter and confirm a new password, and then click **Change**.



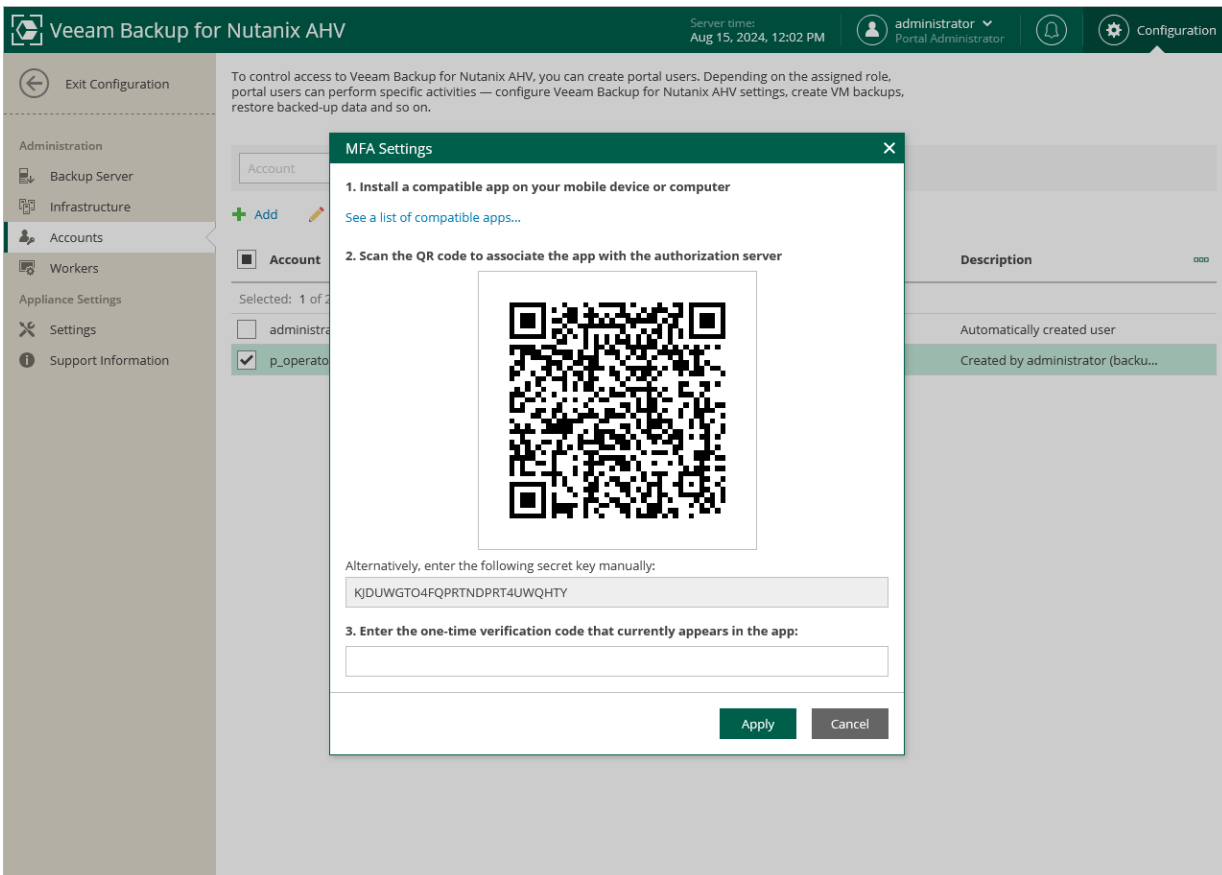
# Configuring Multi-Factor Authentication

Multi-factor authentication (MFA) in Veeam Backup for Nutanix AHV uses the Time-based One-Time Password (TOTP) method that requires users to verify their identity by providing a temporary six-digit code sent by an authentication application to a trusted device.

## Enabling MFA

To enable MFA for a user account, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts**.
3. Select the user account and click **MFA > Enable**.
4. Follow the instructions provided in the **MFA Settings** window:
  - a. Install a TOTP authentication application on a trusted device..
  - b. To associate the authentication application with the authorization server, scan the displayed QR code using the camera of the trusted device. Alternatively, you can open the application on your device and enter a secret key displayed in the **MFA Settings** window.
  - c. Enter a verification code generated by the authentication application.
  - d. Click **Apply**.



## Disabling MFA

To disable MFA for a user account, select the account on the **Accounts** tab and click **MFA> Disable**.

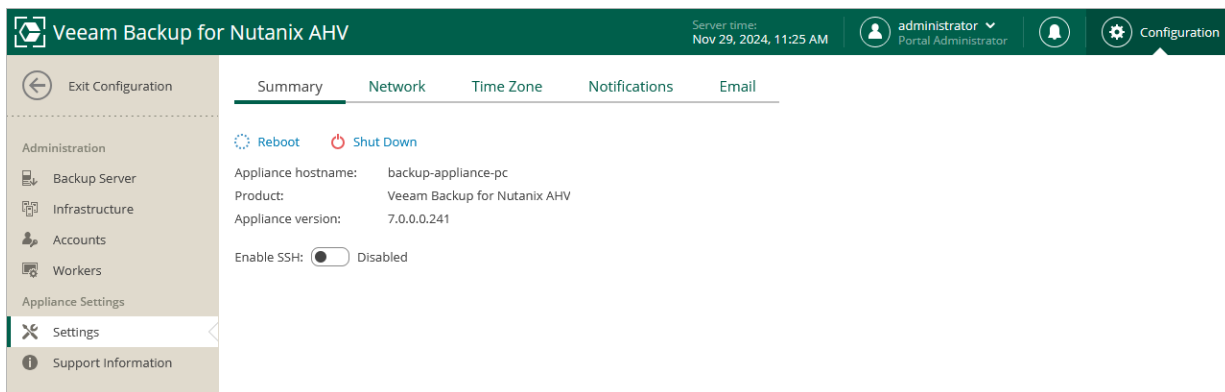


# Enabling SSH on Backup Appliance

For security reasons, SSH is disabled on the Nutanix AHV backup appliance by default. However, you can enable it for the purposes of manual management and troubleshooting:

1. Switch to the **Configuration** page.
2. Navigate to **Settings**.
3. On the **Summary** tab, set the **Enable SSH** toggle to *On*.

When you connect to the Nutanix AHV backup appliance using SSH, enter credentials of the administrator account that you have specified during the [initial backup appliance configuration](#).



# Viewing Network Settings

You can view network settings configured for the appliance while connecting it to the backup server:

1. Switch to the **Configuration** page.
2. Navigate to **Settings > Network**.

To change the network to which the Nutanix AHV backup appliance is connected or to specify a new IP address for the appliance, [edit the appliance properties](#) and follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 5).

The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Nutanix AHV', the server time 'Aug 15, 2024, 12:45 PM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Administration (Backup Server, Infrastructure, Accounts, Workers) and Appliance Settings (Settings, Support Information). The main content area is titled 'Network' and shows the following information:

View network settings of the backup appliance. These settings can be changed in the Veeam Backup & Replication console.

DNS server addresses: 172.24.128.101, 172.24.0.101

Search: Network  Filter (None)

Order ↑	Network	IP Assignment	IP Address	Subnet Mask	Default Gatew...	Description	...
1	VM network	Dynamic	172.25.125.185	255.255.252.0	172.25.124.1	Created by AHV...	


# Configuring Time Synchronization Settings

You can configure what mode Veeam Backup for Nutanix AHV uses for time synchronization while performing data protection and disaster recovery operations:

1. Switch to the **Configuration** page.
2. Navigate to **Settings > Time Zone**.
3. From the **Synchronization mode** list, choose whether you want the time set on the Nutanix AHV backup appliance to be synchronized with the time of specific NTP servers or with the time of the guest OS running on the host where the appliance is deployed.
4. Click **Apply**.

## NOTE

As the backup appliance time zone is automatically synchronized with the backup server time zone, it cannot be changed manually in the Nutanix AHV backup appliance web console.



# Configuring Email Notification Settings

You can specify email notification settings for automated delivery of job results and daily reports. To connect an email service that will be used for sending email notifications:

1. Switch to the **Configuration** page.
2. Navigate to **Settings > Email**.
3. Select the **Enable email notifications** check box.
4. Click the link in the **Email service** field and configure [mail server settings](#).
6. In the **From** field, enter an email address of the notification sender. This email address will be displayed in the **From** field of notifications.
7. In the **To** field, enter an email address of a recipient. Use a semicolon to separate multiple recipient addresses.
8. In the **Subject** field, specify a subject for notifications. You can use the following runtime variables:
  - *%JobName%* – a job name.
  - *%JobResult%* – a job result.
  - *%ObjectCount%* – the number of VMs or PDs in a job.
9. Choose whether you want to receive email notifications in case jobs complete successfully, complete with warnings or complete with errors.
10. Select the **Suppress notifications until the last retry** check box to receive a notification about the final job status. If you do not enable this option, the Nutanix AHV backup appliance will send one notification per every job retry.
11. Click **Apply**.

## TIP

Veeam Backup for Nutanix AHV allows you to send a test message to check whether you have configured the settings correctly. To do that, click **Send Test Email**. A test message will be sent to the specified email address.

## Configuring Mail Server Settings


To configure mail server settings, choose whether you want to employ [Basic \(SMTP\)](#) or [Modern \(OAuth 2.0\)](#) authentication for your mail service.

### Using Basic Authentication

To employ the Basic authentication to connect to your mail server, in the **Email Server Settings** window:

1. From the **Mail service** drop-down list, select *Basic*.
2. In the **Mail server name or address** field, enter a DNS name or an IP address of the SMTP server. All email notifications (including test messages) will be sent by this SMTP server.
3. In the **Port** field, specify a communication port for SMTP traffic. If SSL is enabled, the default SMTP port is 587, otherwise port 25 is used.

4. In the **Timeout** field, specify a connection timeout for responses from the SMTP server.
5. For an SMTP server with SSL/TLS support, select the **Connect using SSL** check box to enable SSL data encryption.
6. If your SMTP server requires authentication, select the **This SMTP server requires authentication** check box and specify credentials that will be used to connect to the SMTP server.
7. Click **Save**.



## Using Modern Authentication

To employ the Modern authentication to connect to your mail server:

1. In **Email Server Settings** window, copy the URL from the **Redirect URL** field.
2. For Veeam Backup for Nutanix AHV to be able to use OAuth 2.0 to access Google Cloud or Microsoft Azure APIs, register a new client application either in the [Google Cloud console](#) or in the [Microsoft Azure portal](#).
3. When registering the application, make sure that the redirect URI specified for the application matches the URL copied from the Veeam Backup for Nutanix AHV web console.
4. Back to the Veeam Backup for Nutanix AHV Web UI, do the following in **Email Server Settings** window:
  - a. Use the **Mail service** drop-down list to choose whether the server that you want to use to send email notifications is a *Google* or *Microsoft* mail server.
  - b. In the **Application client ID** and **Client secret** fields, provide the Client ID and Client secret created for the application as described in [Google Cloud documentation](#) or [Microsoft Docs](#).
  - c. [Applies only if you have selected the **Microsoft** option] In the **Tenant ID** field, provide the ID of an Azure AD tenant in which the application has been registered.

- d. Click **Authorize**. You will be redirected to the authorization page. Sign in using a Google or Microsoft Azure account to validate the configured settings.

The screenshot displays the Veeam Backup for Nutanix AHV configuration interface. The top navigation bar includes the Veeam logo, the title "Veeam Backup for Nutanix AHV", the server time "Aug 27, 2024, 04:45 PM", and the user "administrator Portal Administrator". The left sidebar shows navigation options: Administration (Backup Server, Infrastructure, Accounts, Workers), Appliance Settings (Settings, Support Information), and Exit Configuration. The main content area is titled "Email" and contains the following elements:

- Summary, Network, Time Zone, Notifications, and Email tabs.
- Instruction: "Manage email notifications by providing the mail server details and when you want to get notified."
- Checkbox: "Enable email notifications:" (checked).
- Email service: "Not configured" (with a refresh icon).
- Status: "—"
- From: "vbahv@veeam.com"
- To: "Joe.Smith@veeam.com"
- Subject: "[%JobResult%] %JobName% (%ObjectCount% objects)"
- Link: "Send Test Email" (with an envelope icon).
- Section: "Notify me immediately about job:" with checkboxes for "Success", "Warning", "Failure", and "Suppress notifications until the last retry" (all checked).
- Button: "Apply"

An "Email Server Settings" dialog box is open, containing the following fields and options:


- Mail service: "Google" (dropdown menu)
- Application client ID: (empty text field)
- Client secret: (empty text field)
- Redirect URL: "https://172.25.125.185/smtPOauth" (with a "Copy" button)
- Information message: "To authorize through an application, enter the required information and click Authorize. You will be redirected to the authorization page. After the authorization process completes, you will be redirected back to this dialog."
- Buttons: "Authorize" and "Cancel"

# Configuring Notifications

You can enable notifications for Veeam Backup for Nutanix AHV events that may require your actions:

1. Switch to the **Configuration** page.
2. Navigate to **Settings > Notifications**.
3. Select notifications you want to receive:
  - **Update notifications** remind you about available product updates.
  - **Backup storage notifications** warn you on backup repositories running low on free disk space.
  - **Compute resources notifications** inform you about CPU and RAM consumption on workers breaking configured thresholds for more than 10 minutes. They can also notify you about performance issues caused by worker misconfiguration.
4. Click **Apply**.

You can find backup storage and compute resource notifications in [backup job session logs](#). Update notifications appear at the top right corner of the Veeam Backup for Nutanix AHV window next to the user name.



# Performing Configuration Backup and Restore

You can back up and restore the configuration database that stores data collected from the Nutanix AHV backup appliance for the existing jobs and session records. If the backup appliance goes down for some reason, you can redeploy it and quickly restore its configuration from a configuration backup. You can also use a configuration backup to migrate the configuration of one backup appliance to another backup appliance in the backup infrastructure.

It is recommended that you regularly perform configuration backup for every Nutanix AHV backup appliance present in the backup infrastructure. Periodic configuration backups reduce the risk of data loss and minimize the administrative overhead costs in case any problems with the backup appliances occur.

You can run configuration backup manually on demand, or instruct Veeam Backup for Nutanix AHV to do it automatically on a regular basis. Note that the Nutanix AHV backup appliance configuration database is backed up together with the backup server configuration database. However, the Nutanix AHV backup appliance configuration restore operation does not affect the backup server configuration.



# Backing Up Configuration Settings Manually

While performing configuration backup, Veeam Backup for Nutanix AHV exports data from the configuration database and saves it to a backup file in a backup repository. To back up the configuration database of the backup appliance manually, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. In the **Configuration Backup Settings** window, do the following:
  - a. Select the **Enable configuration backup to the following repository** check box and choose a repository where the configuration backup will be stored. Note that you cannot store configuration backups in scale-out backup repositories and external repositories.  
  
For a backup repository to be displayed in the list of available repositories, it must be added to the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Adding Backup Repositories](#).
  - b. In the **Restore points to keep** field, specify the number of configuration backups you want to keep.
  - c. Select the **Enable backup file encryption** check box.
  - d. From the **Password** drop-down list, select a password.

## IMPORTANT

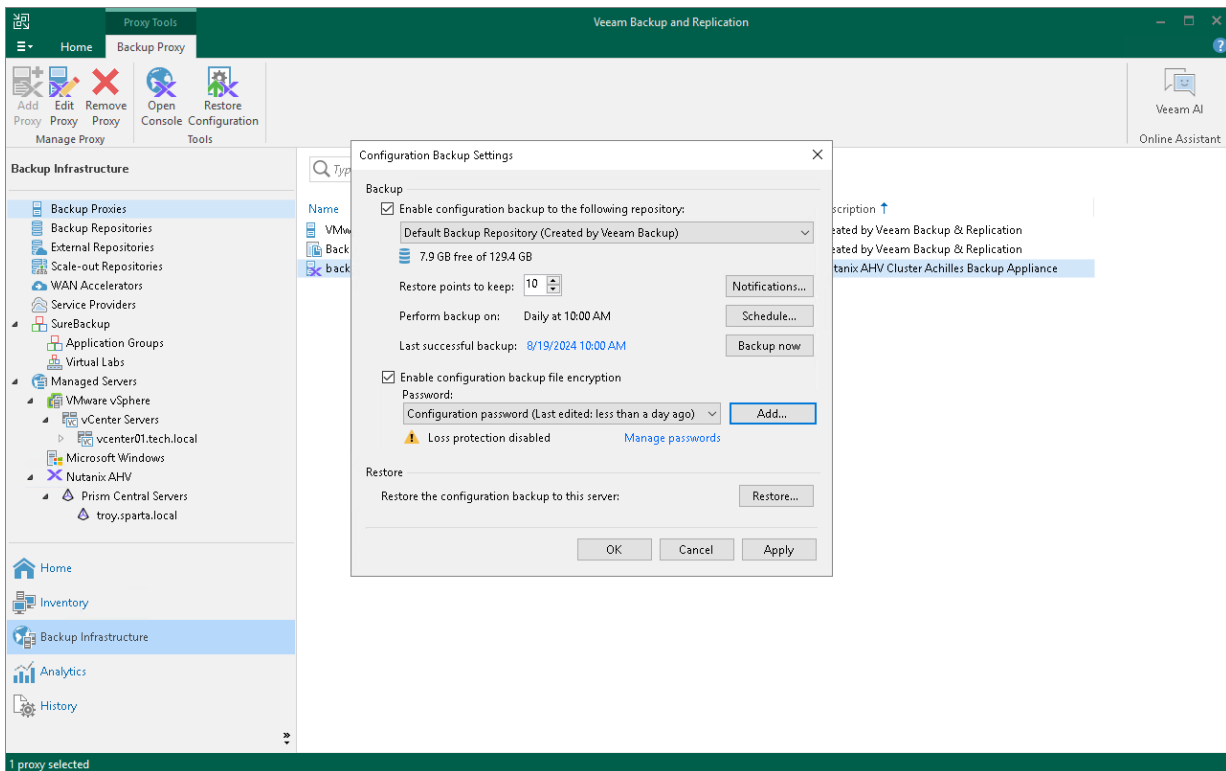
If you do not specify the password, the Veeam Backup for Nutanix AHV configuration database will not be backed up.

For passwords to be displayed in the **Password** list, they must be added to the Password Manager as described in the Veeam Backup & Replication User Guide, section [Password Manager](#). If you have not added the necessary password to the Password Manager beforehand, you can do this without closing **the Configuration Backup Settings** window. To add a password, click **Add** and specify a password and a password hint that will help you remember your password if you forget it.

If you use Veeam Backup Enterprise Manager, you can also enable the Loss protection functionality that can help you decrypt the data in case you have lost or forgotten the password. For more information, see the Veeam Backup Enterprise Manager Guide, section [Managing Encryption Keys](#).

- e. Click **Apply**.
- f. Click **Backup now**.

Once Veeam Backup for Nutanix AHV creates a successful configuration backup, you can use it to [restore configuration data](#).



# Backing Up Configuration Settings Automatically

While performing configuration backup, Veeam Backup for Nutanix AHV exports data from the configuration database and saves it to backup files in a backup repository. To instruct Veeam Backup for Nutanix AHV to back up the configuration database of the backup appliance automatically by schedule, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. In the **Configuration Backup Settings** window, do the following:
  - a. Select the **Enable configuration backup to the following repository** check box and choose a repository where the configuration backup will be stored. Note that you cannot store configuration backups in scale-out backup repositories and external repositories.  
  
For a backup repository to be displayed in the list of available repositories, it must be added to the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Adding Backup Repositories](#).
  - b. In the **Restore points to keep** field, specify the number of configuration backups you want to keep.
  - c. Click **Schedule** and choose whether configuration backups will be created every day or monthly on specific days.
  - d. Select the **Enable backup file encryption** check box.
  - e. From the **Password** drop-down list, select a password.

## IMPORTANT


If you do not specify the password, the Veeam Backup for Nutanix AHV configuration database will not be backed up.

For passwords to be displayed in the **Password** list, they must be added to the Password Manager as described in the Veeam Backup & Replication User Guide, section [Password Manager](#). If you have not added the necessary password to the Password Manager beforehand, you can do this without closing **the Configuration Backup Settings** window. To add a password, click **Add** and specify a password and a password hint that will help you remember your password if you forget it.

If you use Veeam Backup Enterprise Manager, you can also enable the Loss protection functionality that can help you decrypt the data in case you have lost or forgotten the password. For more information, see the Veeam Backup Enterprise Manager Guide, section [Managing Encryption Keys](#).

- f. Click **OK**.

Once Veeam Backup for Nutanix AHV creates a successful configuration backup, you can use it to [restore configuration data](#).



# Restoring Configuration Settings

Veeam Backup for Nutanix AHV offers restore of the configuration database that can be helpful in the following situations:

- The configuration database got corrupted, and you want to recover data from a configuration backup.
- The backup appliance got corrupted, and you want to recover its configuration from a configuration backup.
- The backup appliance went down, and you want to apply its configuration to a new backup appliance.
- You want to roll back the configuration database to a specific point in time.
- You want to apply the backed-up configuration of a Nutanix AHV backup appliance version 4.0 (or later) to a newly deployed Nutanix AHV backup appliance.

When you restore the configuration database of a Nutanix AHV backup appliance, consider the following:

- If the Nutanix AHV backup appliance is still present in the backup infrastructure, you cannot restore its configuration to another Nutanix AHV backup appliance added to same backup infrastructure. This limitation prevents collisions between jobs with the same database ID.
- Network settings of the Nutanix AHV backup appliance remain unchanged. However, you will be able to [change these settings](#) after the configuration restore.
- Configuration settings of dedicated workers will be restored from the configuration backup, and all existing workers will be removed. If any of the settings (such as cluster, worker network settings, host affinity or storage container configuration) is invalid in the current virtual environment, a warning message will be displayed in [configuration restore logs](#). To update worker settings, [modify worker configuration](#) after the configuration restore.
- [Applies only to the [standalone cluster deployment](#)] If you restore the configuration database of a Nutanix AHV backup appliance originally residing in another cluster to protect migrated VMs, you will need to reconfigure backup jobs. UUIDs of migrated VMs change, therefore, you will need to re-add VMs to a backup job that will start new backup chains for them.

You will also need to reconfigure backup jobs and workers, if the original appliance was configured to manage jobs in a Prism Central deployment and the new appliance is connected to a standalone cluster.

- [Applies only to the [Prism Central deployment](#)] If you restore the configuration database of a Nutanix AHV backup appliance originally residing in another cluster registered with the same Prism Central, Veeam Backup for Nutanix AHV will update backup jobs and worker configuration automatically.

## IMPORTANT

Before you start the restore process, stop and disable all jobs that are currently running.

To restore the configuration database, do the following:

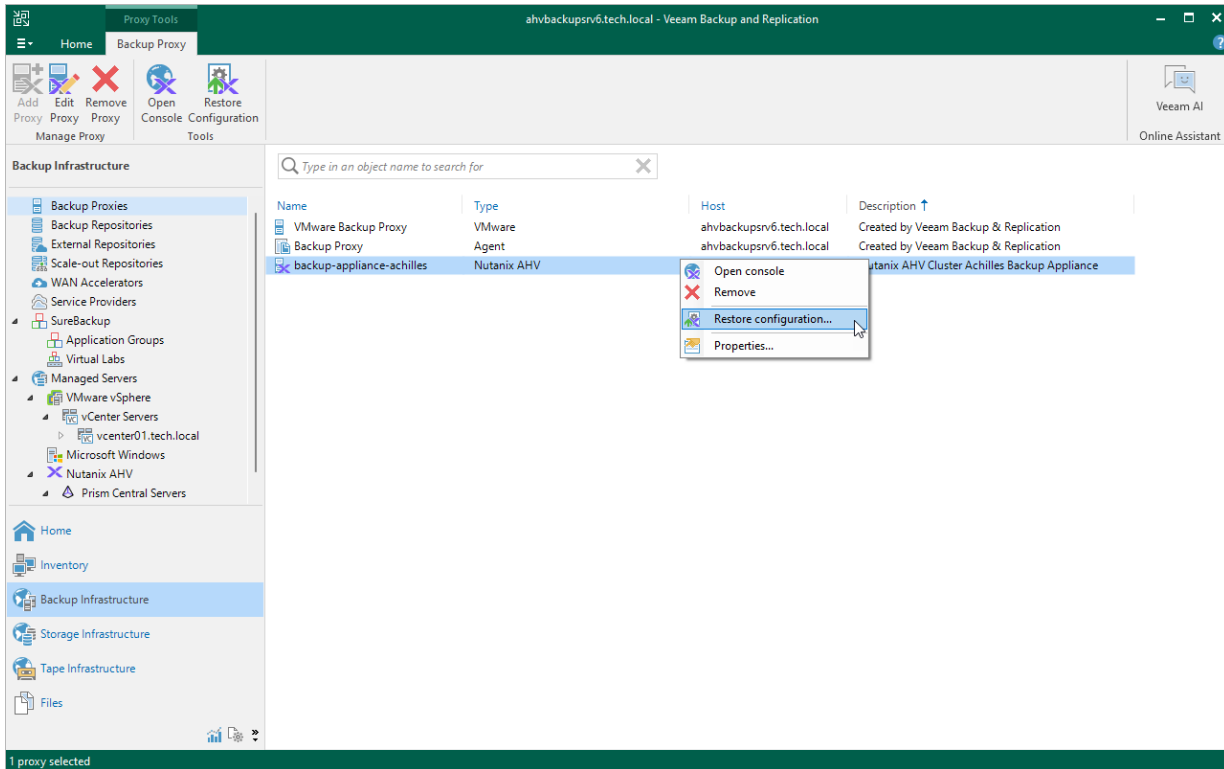
1. [Launch the Configuration Restore wizard.](#)
2. [Choose a backup file.](#)
3. [Review the backup file information.](#)
4. [Provide the encryption password.](#)
5. [Choose restore options.](#)
6. [Specify credentials for the backup appliance account.](#)

7. [Track the restore progress.](#)
8. [Finish working with the wizard.](#)

# Step 1. Launch Configuration Restore Wizard

To launch the **Configuration restore** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the Nutanix AHV backup appliance and click **Restore Configuration** on the ribbon, or right-click the Nutanix AHV backup appliance and select **Restore Configuration**.



## Step 2. Choose Backup File

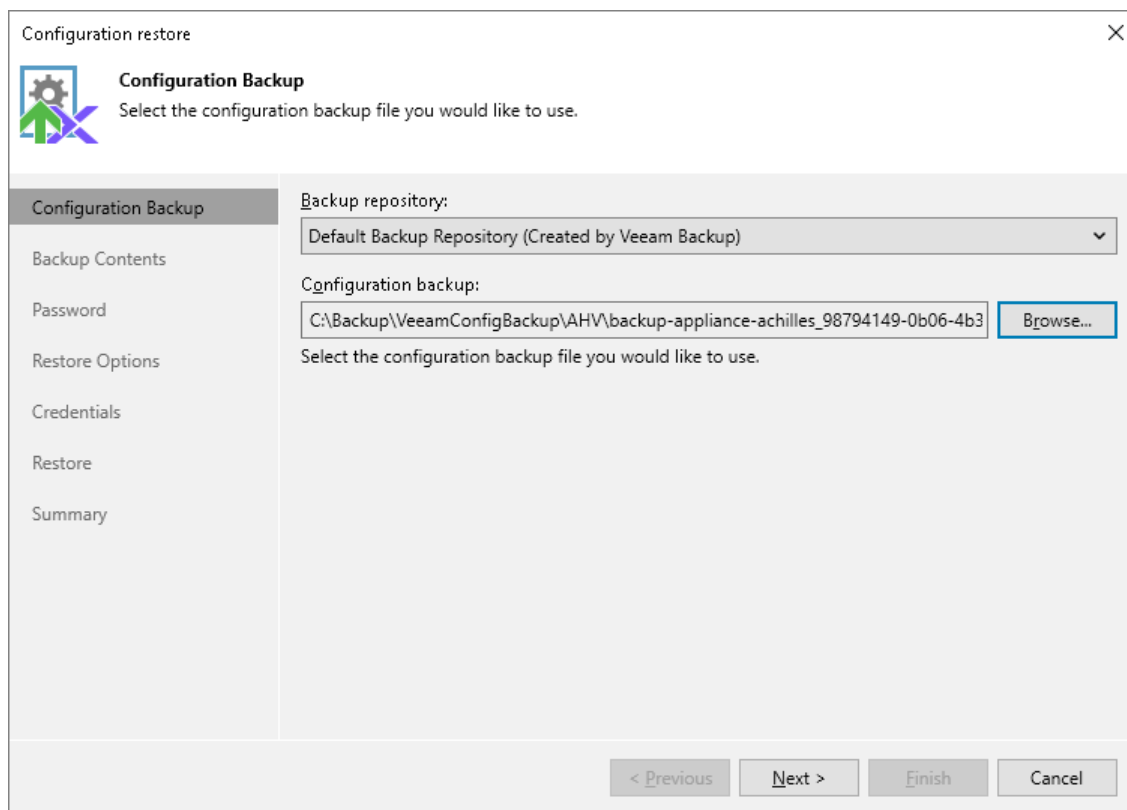
At the **Configuration Backup** step of the wizard, do the following:

1. From the **Backup repository** list, select the backup server or backup repository where the necessary configuration backup file is stored.

For a backup repository to be displayed in the **Backup repository** list, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section [Backup Repository](#). Note that the repository list does not include scale-out backup repositories and external repositories as they cannot store configuration backup files.

2. Click **Browse** and select the necessary file in the **Select file** window.

If the selected configuration backup file is not stored on the backup server, Veeam Backup for Nutanix AHV will copy the file to a temporary folder on the server and automatically delete it from the folder as soon as the restore process completes.



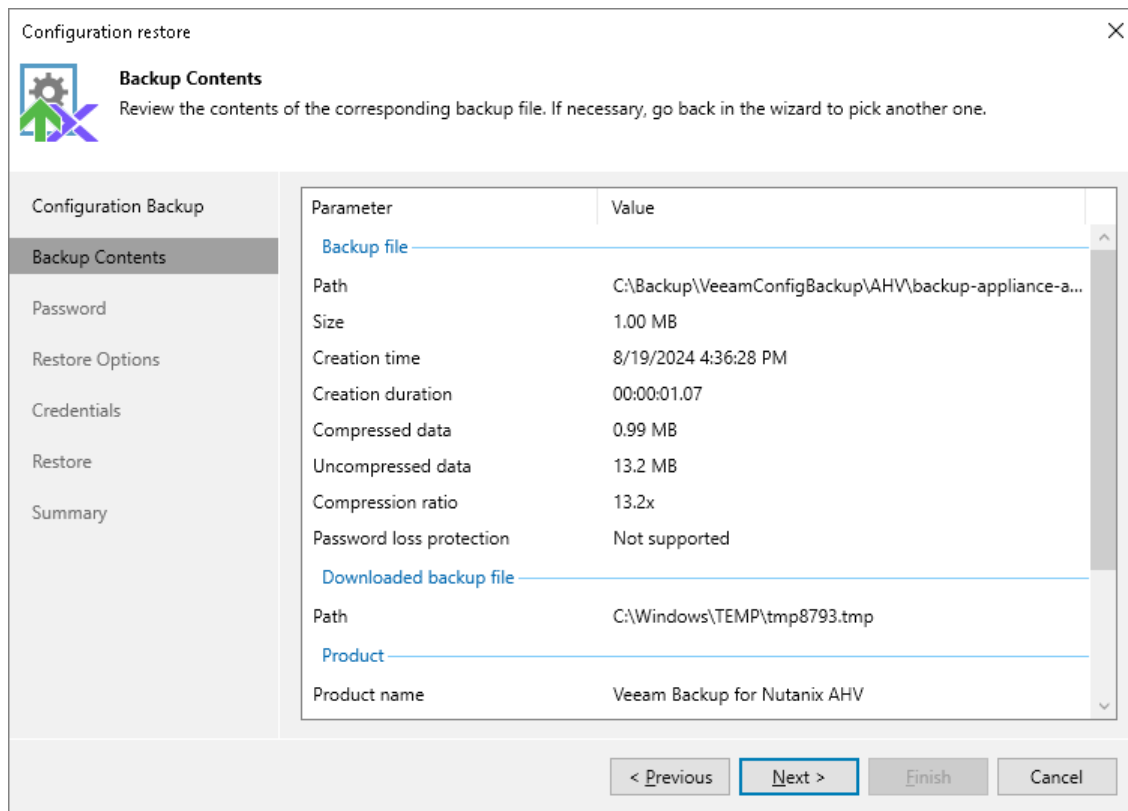


## Step 3. Review Backup Details

Veeam Backup for Nutanix AHV will analyze the content of the selected backup file and display the following information:

- Backup file – the date and time when the backup file was created.
- Downloaded backup file – the temporary location of the configuration backup file on the backup server.
- Product – the version of Veeam Backup for Nutanix AHV that was installed the initial backup appliance.
- Catalogs – configuration data saved in the file (such as the number of configured jobs, users, logged session records and so on).

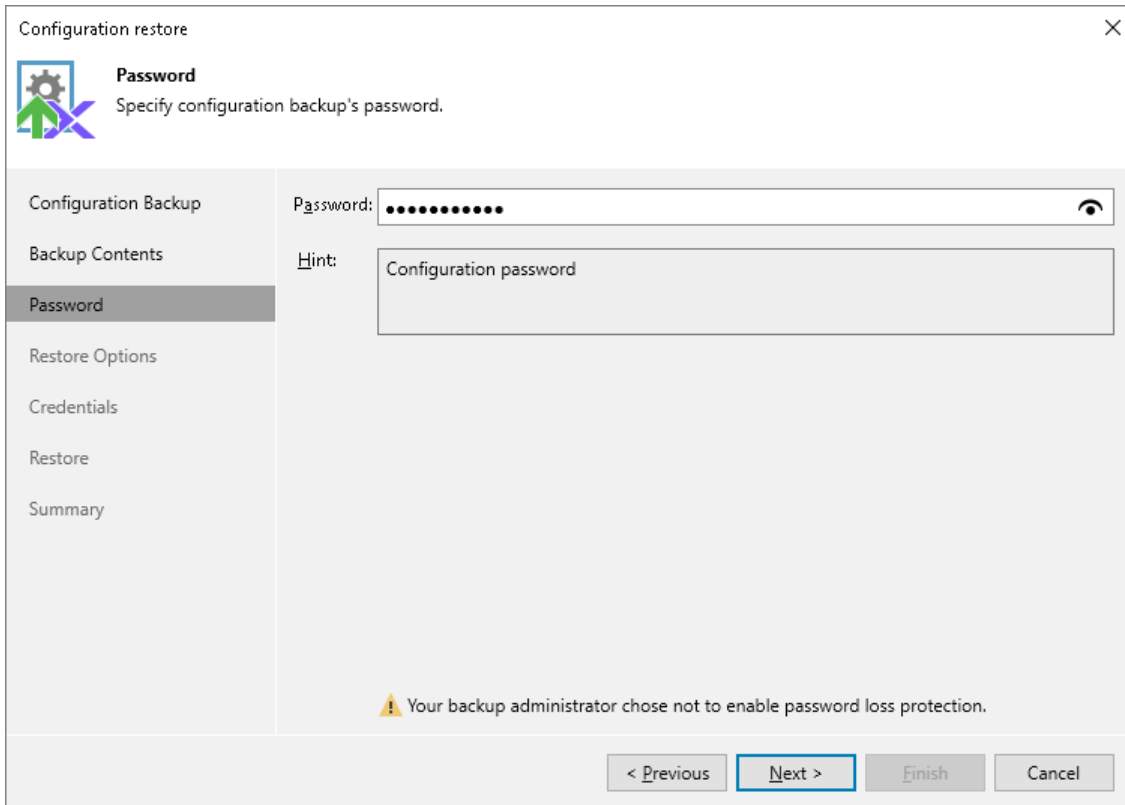
At the **Backup Content** step of the wizard, review the provided information and click **Next** to confirm that you want to use the selected file to restore the configuration data.



## Step 4. Provide Encryption Password

At the **Password** step of the wizard, provide a password that was used to encrypt the file while creating configuration backup.

If you do not remember the password, you can use an alternative way for data encryption. However, this option is available only if password Loss protection was enabled when you created the backup. For more information, see the Veeam Backup & Replication User Guide, section [Decrypting Data Without Password](#).



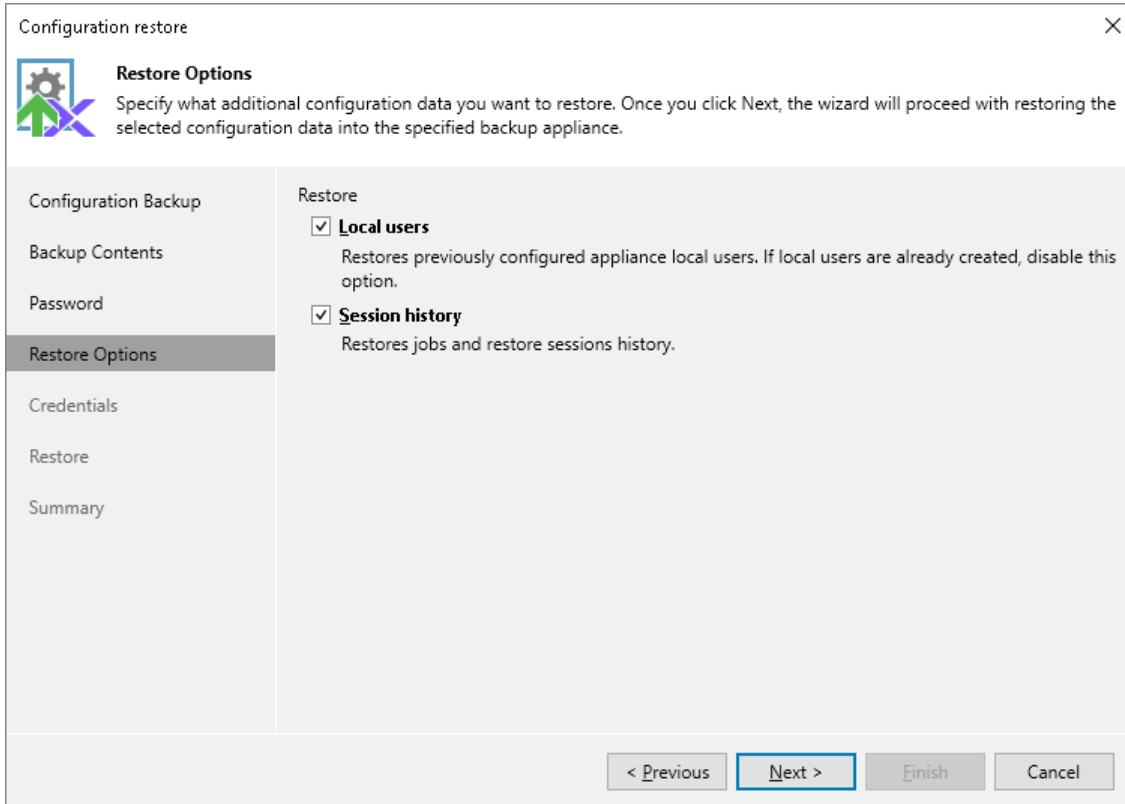
The screenshot shows a window titled "Configuration restore" with a close button (X) in the top right corner. On the left is a navigation pane with the following items: Configuration Backup, Backup Contents, Password (highlighted), Restore Options, Credentials, Restore, and Summary. The main area has a header with a gear icon and the text "Password Specify configuration backup's password." Below this is a "Password:" field containing ten black dots and a toggle icon. Underneath is a "Hint:" field containing the text "Configuration password". At the bottom of the main area is a warning icon and the text "Your backup administrator chose not to enable password loss protection." At the very bottom are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

## Step 5. Choose Restore Options

At the **Restore Options** step of the wizard, you can choose whether you want to restore session logs and user accounts configured on the initial backup appliance (Portal Administrators, Portal Operators and Restore Operators) as well.

### IMPORTANT

If you do not select the **Local users** check box, the restore process will start as soon as you click **Next**. In this case, you will not be able to halt the process or edit the restore settings.



The screenshot shows a window titled "Configuration restore" with a close button (X) in the top right corner. Below the title bar is a gear icon and the text "Restore Options". A descriptive paragraph reads: "Specify what additional configuration data you want to restore. Once you click Next, the wizard will proceed with restoring the selected configuration data into the specified backup appliance." On the left side, there is a vertical list of steps: "Configuration Backup", "Backup Contents", "Password", "Restore Options" (highlighted with a dark background), "Credentials", "Restore", and "Summary". The main area is titled "Restore" and contains two checked options:  **Local users** with the description "Restores previously configured appliance local users. If local users are already created, disable this option." and  **Session history** with the description "Restores jobs and restore sessions history." At the bottom right, there are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

## Step 6. Select Credentials

[This step applies only if you have selected the **Local users** option at the **Restore Options** step of the wizard]

At the **Credentials** step of the wizard, specify credentials of a Portal Administrator account used to access the Nutanix AHV backup appliance web console.

### NOTE

The following accounts cannot be used to access the Nutanix AHV backup appliance as Portal Administrators:

- Accounts for which [multi-factor authentication \(MFA\)](#) is enabled.
- Active Directory accounts since the Nutanix AHV backup appliance does not support LDAP integration.

For credentials to be displayed in the **Credentials** list, they must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary credentials to the Credentials Manager beforehand, you can do this without closing the **Configuration restore** wizard. To add credentials, click **Add** and specify a user name and password in the **Credentials** window. Note that the name must start with a lowercase Latin letter and must not match Linux system user names (such as *root*, *daemon*). The name can contain only lowercase Latin letters, numeric characters, underscores and dashes. The maximum length of the name is 32 characters.

If the specified credentials match the credentials stored in the backup configuration database, the user account will be restored. If you do not remember the password, you can specify a new user name and password. In this case, Veeam Backup for Nutanix AHV will restore user accounts from the configuration file and create a new one with credentials that you specify.

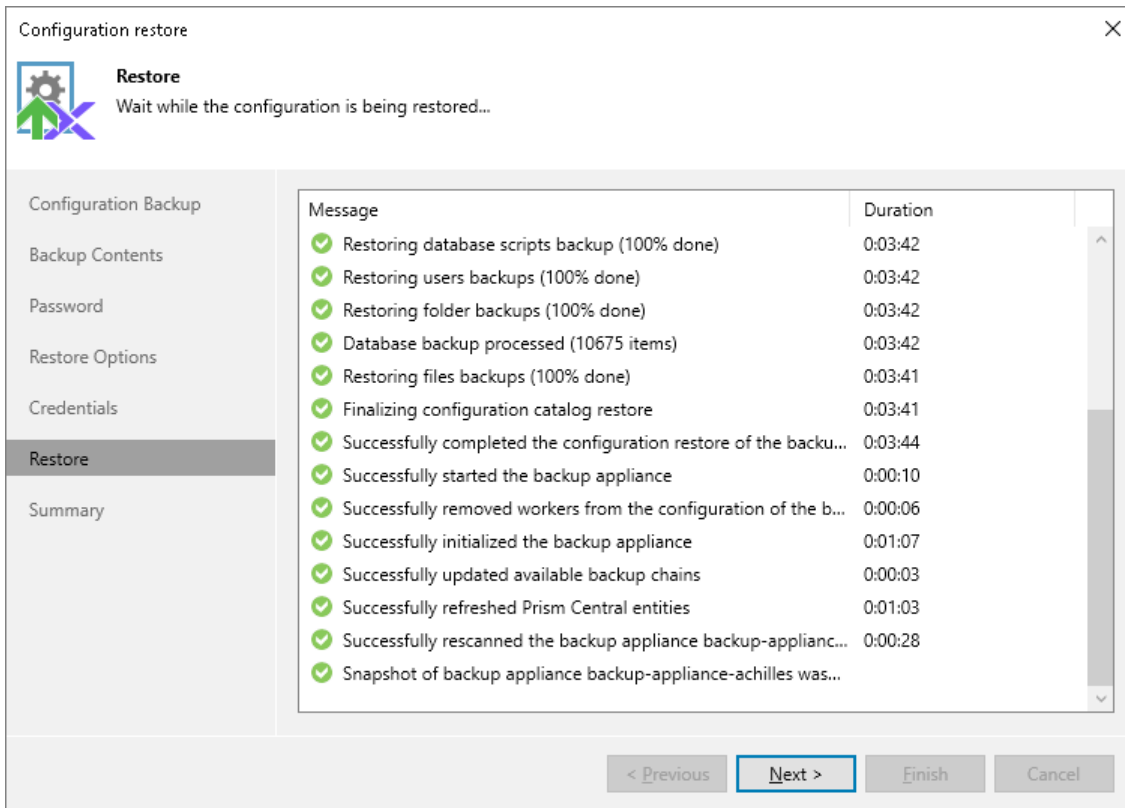
## IMPORTANT

The restore process will start as soon as you click **Restore**. In this case, you will not be able to halt the process or edit the restore settings.

The screenshot shows a 'Configuration restore' dialog box with a 'Credentials' tab selected. The dialog has a title bar with a close button (X) and a navigation pane on the left with the following items: Configuration Backup, Backup Contents, Password, Restore Options, Credentials (selected), Restore, and Summary. The main area contains the following text: 'Credentials Specify server administrator's credentials.' Below this is an instruction: 'Select an account that has administrative privileges on the backup appliance. During restore backup server will verify whether the selected credentials exist on the backup appliance, and create them automatically, if needed.' Underneath is a 'Credentials:' label followed by a dropdown menu showing 'administrator (Backup appliance administrator credentials, last edited: 5 d...)' and an 'Add...' button. A 'Manage accounts' link is positioned below the dropdown. At the bottom of the dialog are four buttons: '< Previous', 'Restore' (highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 7. Track Restore Progress

Veeam Backup for Nutanix AHV will display the results of every step performed while executing the configuration restore. At the **Restore** step of the wizard, wait for the restore process to complete and click **Next**.




The screenshot shows the 'Configuration restore' wizard window. The 'Restore' step is selected in the left-hand navigation pane. The main area displays a list of messages and their durations, indicating the progress of the restore process. The 'Next >' button is highlighted, suggesting the process is complete or nearly complete.

Message	Duration
✓ Restoring database scripts backup (100% done)	0:03:42
✓ Restoring users backups (100% done)	0:03:42
✓ Restoring folder backups (100% done)	0:03:42
✓ Database backup processed (10675 items)	0:03:42
✓ Restoring files backups (100% done)	0:03:41
✓ Finalizing configuration catalog restore	0:03:41
✓ Successfully completed the configuration restore of the backu...	0:03:44
✓ Successfully started the backup appliance	0:00:10
✓ Successfully removed workers from the configuration of the b...	0:00:06
✓ Successfully initialized the backup appliance	0:01:07
✓ Successfully updated available backup chains	0:00:03
✓ Successfully refreshed Prism Central entities	0:01:03
✓ Successfully rescanned the backup appliance backup-applianc...	0:00:28
✓ Snapshot of backup appliance backup-appliance-achilles was...	

## Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, click **Finish** to finalize the process of configuration data restore.



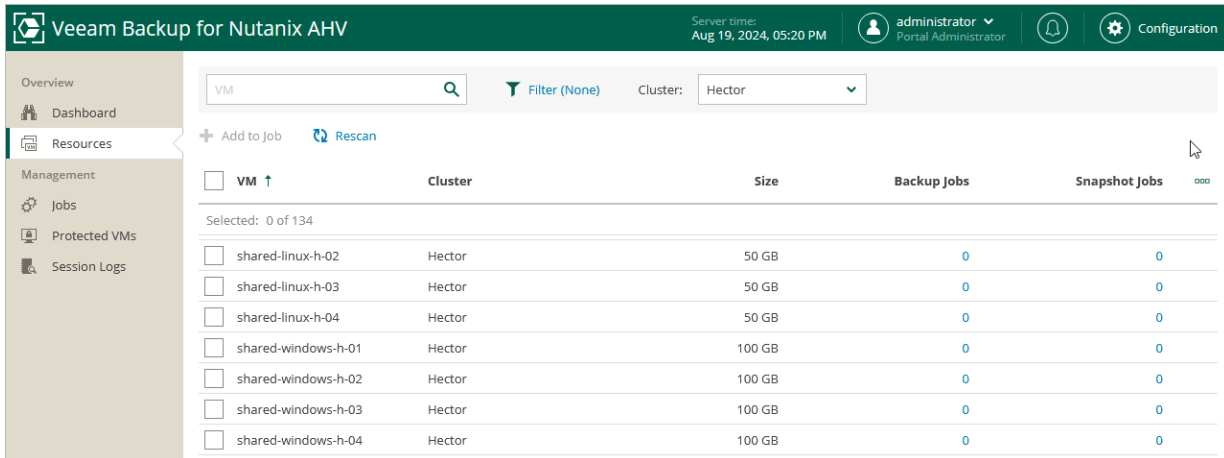
# Viewing Available Resources

After you create backup or snapshot jobs, Veeam Backup for Nutanix AHV rescans the Prism Central or standalone cluster and populates on the **Resources** tab the VM list where you can find the following information:

- **VM** – the name of the VM.
- **Cluster** – the cluster where the VM resides.
- **Size** – the size of the VM.
- **Backup Jobs** – the number of backup jobs that protect the VM.
- **Snapshot Jobs** – the number of snapshot jobs that protect the VM.

On the **Resources** tab, you can perform the following actions:

- Update the information on Nutanix AHV VMs and jobs that protect them. To do that, click **Rescan**.
- View what backup jobs protect a VM. To do that, find the VM in the list and click the link in the **Backup Jobs** column.
- View what snapshot jobs protect a VM. To do that, find the VM in the list and click the link in the **Snapshot Jobs** column.
- Add VMs to existing jobs. For more information, see [Adding VMs to Job](#).



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Nutanix AHV", the server time "Aug 19, 2024, 05:20 PM", and user information "administrator Portal Administrator". The left sidebar contains navigation options: Overview, Dashboard, Resources (selected), Management, Jobs, Protected VMs, and Session Logs. The main content area shows a search bar for "VM", a filter set to "None", and a cluster dropdown set to "Hector". Below this is a table with columns for VM, Cluster, Size, Backup Jobs, and Snapshot Jobs. The table lists seven VMs, all in the "Hector" cluster, with sizes ranging from 50 GB to 100 GB and zero backup or snapshot jobs.

VM	Cluster	Size	Backup Jobs	Snapshot Jobs
<input type="checkbox"/> shared-linux-h-02	Hector	50 GB	0	0
<input type="checkbox"/> shared-linux-h-03	Hector	50 GB	0	0
<input type="checkbox"/> shared-linux-h-04	Hector	50 GB	0	0
<input type="checkbox"/> shared-windows-h-01	Hector	100 GB	0	0
<input type="checkbox"/> shared-windows-h-02	Hector	100 GB	0	0
<input type="checkbox"/> shared-windows-h-03	Hector	100 GB	0	0
<input type="checkbox"/> shared-windows-h-04	Hector	100 GB	0	0



# Adding VMs to Job

If you want to protect additional VMs by configured jobs, you can either [edit the backup job settings](#), or quickly add the VMs to the jobs on the **Resources** tab.


## NOTE

While adding VMs to a job, consider the following:

- If a VM is already added to the exclusion list in the job, the VM will not be protected.
- You cannot add VMs to a protection domain snapshot job. To protect a VM as part of a protection domain, add the VM to the domain as described in [Nutanix documentation](#) and make sure that the protection domain is included into a [protection domain snapshot job](#).

To add a VM to a backup or snapshot job, do the following:

1. Navigate to **Resources**.
2. Select the VM that you want to protect by a job.
3. Click **Add to Job**.
4. In the **Add to Job** window, choose the job that will protect the selected VM and click **Apply**.



## TIP

You can also add a VM to a backup or snapshot job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Inventory** view, navigate to the Nutanix AHV cluster where the VM resides, select the VM and click **Add to Snapshot** or **Add to Backup** on the ribbon.

# Performing Backup

With Veeam Backup for Nutanix AHV, you can protect data in the following ways:

- **Create snapshots of VMs**

A snapshot captures point-in-time image of the processed VM. You can use the snapshot to restore the VM to the original Nutanix AHV environment.

To take VM snapshots, Veeam Backup for Nutanix AHV uses [native Nutanix AHV capabilities](#). VM snapshots are saved locally in the Nutanix AHV cluster where the source VMs reside. To create a VM snapshot, [configure a snapshot job](#).

- **Create snapshots of protection domains**

A protection domain (PD) snapshot captures point-in-time image of consistency groups (VMs and volume groups) that belong to the processed domain. You can use the PD snapshot to restore VMs and volume groups to the original Nutanix AHV environment.

To take PD snapshots, Veeam Backup for Nutanix AHV uses [native Nutanix AHV capabilities](#). PD snapshots are saved locally in the Nutanix AHV cluster where the source PD resources reside. To create a PD snapshot, [configure a protection domain snapshot job](#).

- **Create backups of VMs**

In addition to VM snapshots, you can protect your Nutanix AHV VMs with image-level backups. An image-level backup captures the whole image of the processed VM (including VM configuration, OS data, application data and so on) at a specific point in time.

To create backups, Veeam Backup for Nutanix AHV uses [workers](#) that retrieve VM data from the Nutanix AHV cluster and forward it to a backup repository in the [native Veeam format](#). You can use the backup to restore the VM to the original Nutanix AHV environment or any other supported virtual environment, for example, VMware or Hyper-V. To create a VM backup, [configure a backup job](#) or [perform a VeeamZIP operation](#).

## NOTE

To backup data that resides on Nutanix Files, use the Veeam Backup & Replication file share backup functionality described in the Veeam Backup & Replication User Guide, section [NAS Backup](#).

# Creating Backup Jobs

To produce backups of VMs, Veeam Backup for Nutanix AHV runs backup jobs. A backup job is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, and so on. To create a backup job, you can either use the [Veeam Backup for Nutanix AHV web console](#) or the [Veeam Backup & Replication console](#).

# Before You Begin

Before you create a backup job, consider the following limitations:

- You cannot back up workers and Nutanix Controller VMs.
- You cannot back up Windows 11 VMs due to AOS 6.6 technical limitations related to Virtual Trusted Platform Module (vTPM) support. For more information on vTPM VM configuration, see [Nutanix documentation](#).
- The way Veeam Backup for Nutanix AHV creates image-level backups depends on whether you choose to protect individual virtual machines, protection domains, clusters, categories or the whole Prism Central while creating the job:

If the backup job includes individual virtual machines, as well as the whole Prism Central, category or cluster, Veeam Backup for Nutanix AHV creates a snapshot of each VM and a VG snapshot of each volume group, and then produces image-level backups using the created snapshots. This approach cannot guarantee full consistency of VM and volume group data.

If the backup job includes protection domains, Veeam Backup for Nutanix AHV creates a PD snapshot of all VMs and volume groups added to the same consistency group in each protection domain. This approach guarantees full data consistency. However, Veeam Backup for Nutanix AHV takes PD snapshots only if [application-aware processing is disabled](#) in job settings and the following requirements are met for each protection domain included into the backup job:

- The protection domain does not contain multiple VMs with the same name.
- VMs and their volume groups belong to the same protection domain.
- VMs and their volume groups are included into one [consistency group](#) of the protection domain.
- [CHAP authentication is disabled](#) for the volume groups.

If any of those requirements are not met, Veeam Backup for Nutanix AHV backs up VMs and their volume groups as if processing individual virtual machines.

- When processing a VM, Veeam Backup for Nutanix AHV always tries to produce an application-consistent backup using Nutanix Guest Tools. However, if the [requirements for application-consistent snapshots](#) that are used to create application-consistent backups are not met, Veeam Backup for Nutanix AHV creates a crash-consistent backup instead.

If you want to create an application-consistent backup using the Veeam technology, you can configure [application-aware processing settings](#) for resources included into the backup job. This will also allow you to periodically back up and truncate transaction logs.

- [Applies only to the [Prism Central deployment](#)] If you want to back up VMs using data obtained from snapshots on replica clusters, ensure that you have scheduled Prism Central protection policies to take snapshots more frequent than the backup job runs.
- By default, [backup encryption](#) is disabled for backed-up data. However, you can enable encryption at the repository level as described in the Veeam Backup & Replication User Guide, section [Access Permissions](#).
- [VM guest OS file indexing](#) is not supported for backups created with Veeam Backup for Nutanix AHV.
- Since Veeam Backup & Replication does not allow you to assign [information about locations](#) to Nutanix AHV clusters and backup appliances, job statistics do not include information on the Nutanix AHV VM data migration between different geographic regions.

# Creating Backup Jobs Using Backup Appliance Web Console

To create a backup job using the Nutanix AHV backup appliance web console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Configure general settings.](#)
4. [Selects VMs or protection domains to back up.](#)
5. [Specify a backup repository where backups will be stored.](#)

## NOTE

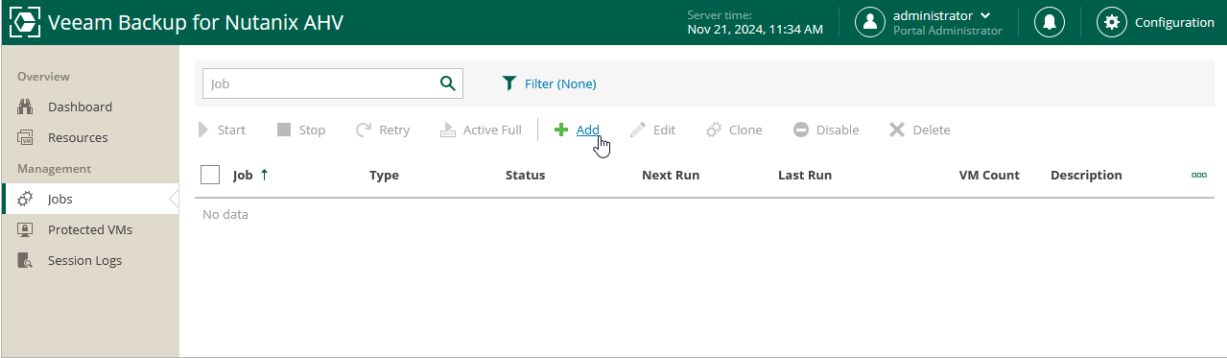
The wizard will also include the [Guest Processing step](#), but guest processing settings cannot be configured in the web console due to technical limitations. You can skip this step, proceed with the wizard – and configure these settings in the Veeam Backup & Replication console later, after you create the backup job.

6. [Create a schedule for the backup job.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch Add Job Wizard

To launch the **Add Job** wizard, do the following:

- 1. Navigate to **Jobs**.
- 2. Click **Add**.



## Step 2. Specify Job Name and Description

At the **Job Info** step of the wizard, select the **Backup job** option. Then, use the **Name** and **Description** fields to specify a name for the new backup job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.

The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV. The interface is in a dark green theme. At the top, the title bar reads 'Veeam Backup for Nutanix AHV' and includes the server time 'Nov 21, 2024, 01:11 PM' and the user 'administrator Portal Administrator'. A navigation pane on the left lists steps: Job Info (selected), Sources, Backup Destination, Guest Processing, Schedule, and Summary. The main area is titled 'Job Info' and contains the instruction 'Enter a name for the job and select a job type.' There are two input fields: 'Name:' with the value 'LMS Daily Backup' and 'Description:' with the value 'Backup of the LMS Server VM and Media Storage VM'. Below these fields are three radio button options: 'Backup job' (selected), 'Snapshot job', and 'PD snapshot job', each with a brief description. At the bottom right, there are 'Next' and 'Cancel' buttons.

## Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify the following backup source settings:

1. [Choose resources to back up.](#)
2. [Choose disks and volume groups to protect.](#)



## Step 3a. Choose Resources

In the **Resources** section of at the **Sources** step of the wizard, specify the backup scope – resources that Veeam Backup for Nutanix AHV will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all VMs in the cluster, only specific VMs or protection domains. In the [Prism Central deployment](#), you can also back up VMs and clusters assigned to a specific category or all VMs managed by a Prism Central.

If you add a protection domain, Veeam Backup for Nutanix AHV will regularly check for new consistency groups (VMs and volume groups) added to the domain and automatically update the job settings to include these groups in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

If you add a category, Veeam Backup for Nutanix AHV will regularly check for new VMs and clusters assigned to the category and automatically update the job settings to include these resources in the backup scope. For a category to be displayed in the list of the available categories, it must be configured in the Nutanix AHV Prism Central as described in [Nutanix documentation](#).

### TIP

As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the resources that you do not want to back up – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

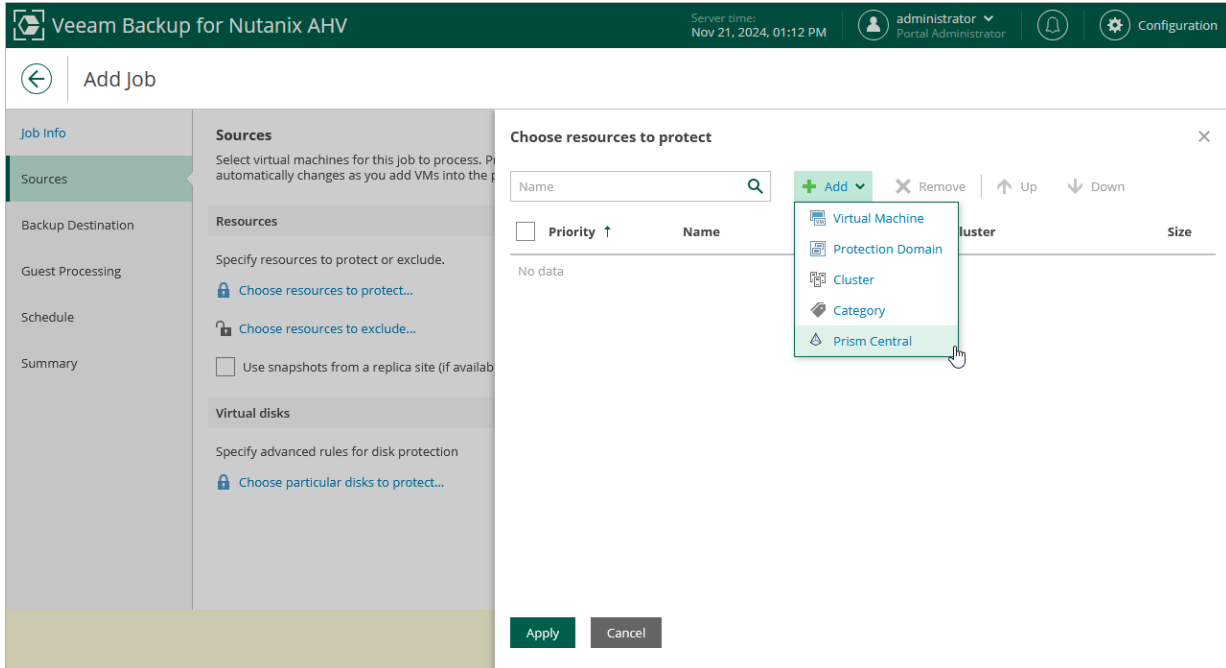
### NOTE

Consider the following:

- If you include a resource into the backup scope for multiple times (for example, an individual VM and a PD that contains this VM), Veeam Backup for Nutanix AHV will process this resource only once.
- If you include a protection domain, category, cluster or Prism Central into the backup scope, VMs in this object are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs – not as a part of the protection domain, category, cluster or Prism Central.

[Applies only to the [Prism Central deployment](#)] To instruct Veeam Backup for Nutanix AHV to obtain [VM data from a replica cluster](#), select the **Use snapshots from a replica site (if available)** check box. If Veeam Backup for Nutanix AHV fails to obtain data from a replica cluster, backup will be still performed using VM data obtained from the main cluster.

By default, jobs process all disks and volume groups attached to VMs included into the backup scope. However, you can protect only specific disks and volume groups of the selected resources. For more information, see [Step 3b. Choose Disks and Volume Groups](#).



## Step 3b. Choose Disks and Volume Groups

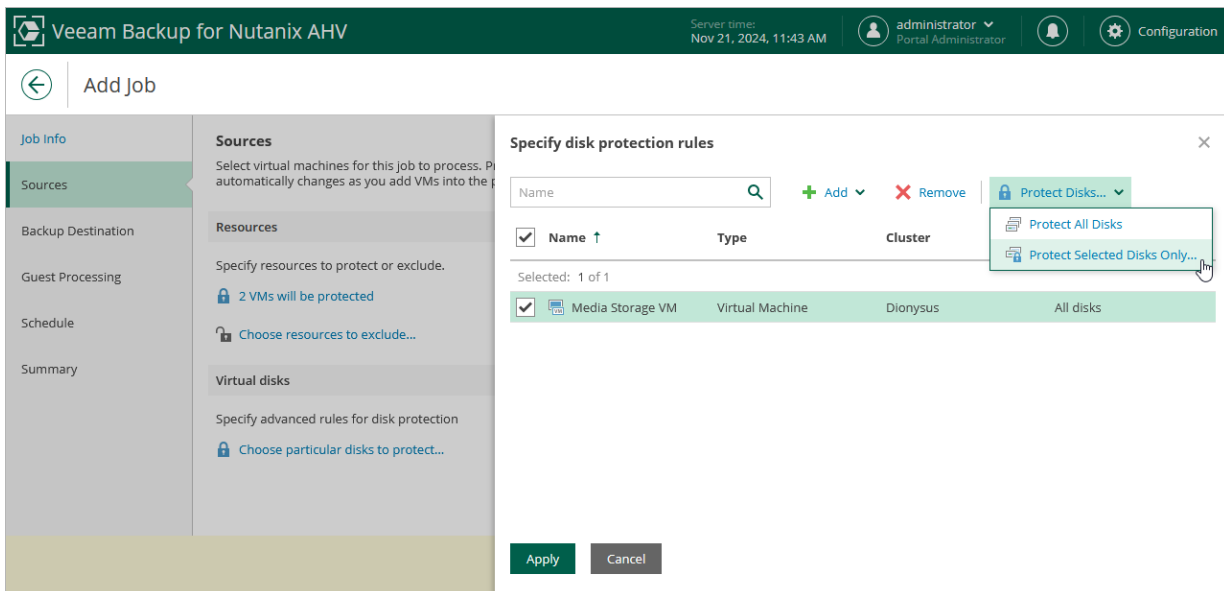
In the **Virtual disks** section of the **Sources** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to back up only specific virtual disks and volume groups related to the specified backup scope:

1. Click **Choose particular disks to protect**.
2. In the **Specify disk protection rules** window, click **Add** and select a resource that you have added to the backup scope at [step 3a](#).
3. Click **Protect Disks > Protect Selected Disks Only**.
4. In the **Protect Selected Disks Only** window, choose a bus type of the disks that you want to back up, select the necessary disks and volume groups, and click **Add**.

Disks and volume groups that you do not select will be excluded from the backup job.

### NOTE

If you configure multiple disk protection rules, specific rules will override general ones. For example, if you add a rule for a protection domain and for a VM included in this domain, Veeam Backup for Nutanix AHV will process the VM disks according to the rule configured for the VM.



# Step 4. Specify Backup Job Settings

At the **Backup Destination** step of the wizard, do the following:

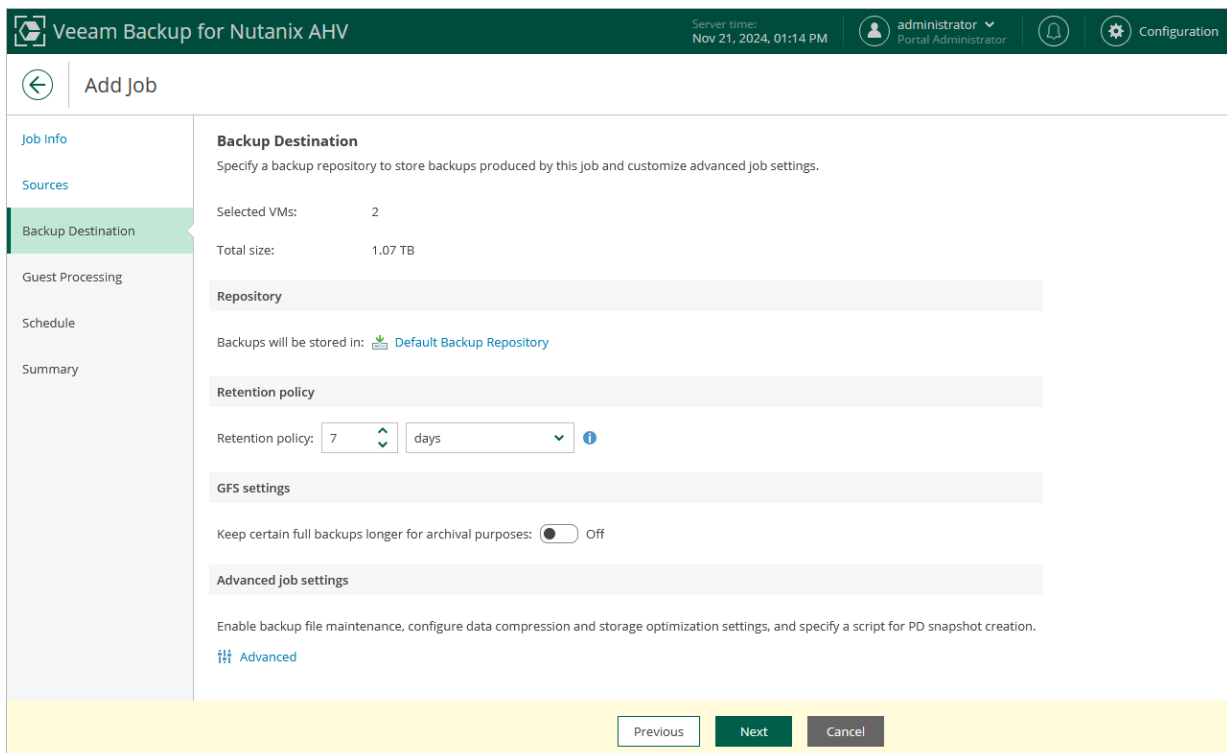
1. From the **Backup repository** drop-down list, select a backup repository where you want to store backups.  
For a backup repository to be displayed in the list of the available repositories, it must be [added to the backup infrastructure](#), and the Nutanix AHV backup appliance must [have access to the repository](#).
2. In the **Retention policy** section, choose a retention policy that Veeam Backup for Nutanix AHV will apply to backups created by the job:
  - Select *days* if you want to keep restore points in a backup chain for the allowed period of time. If a restore point is older than the specified limit, Veeam Backup for Nutanix AHV removes it from the chain.
  - Select *restore points* if you want a backup chain to contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain.

For more information on how Veeam Backup for Nutanix AHV tracks and removes redundant restore points, see [Retention Policies](#).

## NOTE

If the UUID of a VM changes (for example, if the VM was migrated to another cluster), Veeam Backup for Nutanix AHV will be unable to continue the backup chain for this VM. After you re-add the VM to the backup job, Veeam Backup for Nutanix AHV will start a new backup chain for it. However, you will still be able to perform restore operations using backups from the old backup chain.

To help you implement a comprehensive backup strategy, Veeam Backup for Nutanix AHV allows you to [enable long-term retention policy for backups](#) and to [configure backup job advanced settings](#) (for example, enable health check, schedule active and synthetic full backups, plan backup maintenance and upload protection domain custom scripts).



## Configuring GFS Policy Schedules

Grandfather-Father-Son (GFS) policy allows you to leverage full backups for long-term retentions instead of creating a new full backup every time. The mechanism simplifies the backup schedule and optimizes the backup performance.

Veeam Backup for Nutanix AHV re-uses full backups created according to the backup job schedule to achieve the desired retention for GFS policy schedules (weekly, monthly and yearly). Each full backup is marked with a flag of the related GFS policy schedule type: the (W) flag is used to mark full backups created weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for Nutanix AHV uses these flags to control the retention period for the created full backups. Once a flag of a GFS policy schedule is assigned to a full backup, this full backup can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the full backup. If the full backup does not have any other flags assigned, it is removed according to the short-term retention policy settings. For more information on the GFS flag assignment and removal, see the Veeam Backup & Replication User Guide, section [Long-Term Retention Policy \(GFS\)](#).

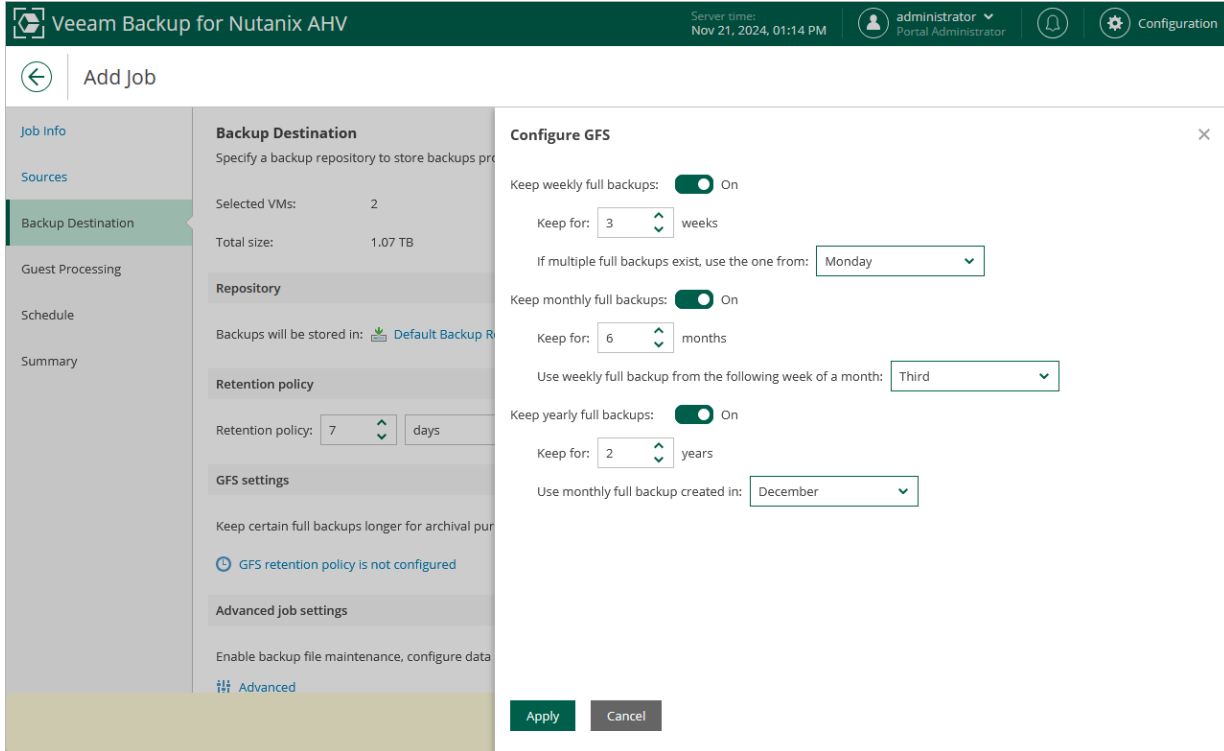
To configure a GFS policy schedule, set the **Keep certain full backups longer for archival purposes** toggle to *On*, and specify the following options in the **Configure GFS** section:

- **Keep weekly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created within a week or on a specific day for a specified number of weeks.
- **Keep monthly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created during a specific week for a specified number of months.
- **Keep yearly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created in a specific month for a number of years.

After you configure the GFS retention policy settings, [schedule active full or synthetic full backups](#). Otherwise, no new full backups will be automatically produced, and Veeam Backup for Nutanix AHV will be unable to leverage them for long-term retentions.

## NOTE

If you choose an object storage repository to store backups produced by the backup job, you cannot enable synthetic full backups. However, if you configure a GFS policy, synthetic backups will be automatically created according to the specified GFS schedule and marked with an appropriate GFS flag.



## Configuring Advanced Settings

To configure backup job advanced settings, do the following:

1. To configure retention settings for backups of VMs that are no longer processed by the backup job, in the **Backup file maintenance** section of the **Advanced job settings** window, set the **Delete backups of VMs that are no longer included in the backup job** toggle to *On*, and specify the number of days during which Veeam Backup for Nutanix AHV will keep backups of VMs excluded from the job.
2. From the **Compression level** drop-down list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. For more information on data compression, see [Compression and Deduplication](#).
3. From the **Storage optimization** drop-down list, select the block size that will be used to process VMs. For more information on the data block sizes and how they affect performance, see [Storage Optimization](#).
4. To instruct Nutanix AHV to freeze applications running on VMs while snapshots are taken, set the **Enable Nutanix Guest Tools quiescence** toggle to *On* and choose how Nutanix AHV will process transaction logs.

Keep in mind that the *Never truncate transaction logs (VSS\_BT\_COPY)* option may significantly increase the amount of storage space consumed by VMs that function as Microsoft Exchange Mail Servers.

## NOTE

This setting will be applied only if guest processing is not configured for the backup job.

- [Applies only if you have included protection domains in the backup job] To instruct Veeam Backup for Nutanix AHV to execute a custom script while running the backup job, set the **Run the following script to create a PD snapshot** toggle to *On*, and click **Browse** to upload a script file. For more information, see [Writing Custom Scripts](#).

## NOTE

When running custom scripts, Veeam Backup for Nutanix AHV uses native Nutanix AHV capabilities to take PD snapshots. If Nutanix AHV fails to create a PD snapshot, the backup job will also fail to complete successfully. For more information on the prerequisites for creating PD snapshots, see [Snapshot Types](#).

The screenshot shows the 'Add Job' configuration window in Veeam Backup for Nutanix AHV. The 'Advanced job settings' tab is active, displaying several configuration sections:

- Backup file maintenance:** 'Delete backups of VMs that are no longer included in the backup job:' is set to *Off*.
- Compression level:** 'Choose the level of data compression to decrease the size of created backup files:' is set to 'Optimal (recommended)'.
- Storage optimization:** 'Specify the data block size to optimize job performance and storage usage:' is set to '1 MB (recommended)'. A yellow tooltip states: 'Changing block size will take effect only after an active full backup is created.'
- Guest quiescence:** 'Enable Nutanix Guest Tools quiescence' is set to *On*. 'Use default Nutanix VSS snapshot type settings' is selected.
- Custom script:** 'Run the following script to create a PD snapshot:' is set to *Off*. A yellow tooltip states: 'Custom scripts help you coordinate application activity before and after a PD snapshot is taken. To learn how to write a script, see the Veeam Backup for Nutanix AHV User Guide.'

At the bottom of the window, there are 'Apply' and 'Cancel' buttons.

## How Health Check Works

When Veeam Backup for Nutanix AHV saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup for Nutanix AHV verifies the availability of data blocks, which are required to restore from the recent point only, and uses the saved values to ensure that the full restore points being verified are consistent.

On the day scheduled for a health check to run, Veeam Backup for Nutanix AHV starts a new health check session. For each restore point in the backup chain, Veeam Backup for Nutanix AHV calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for Nutanix AHV also checks whether data blocks that are required to rebuild the restore point are available.

If Veeam Backup for Nutanix AHV does not detect data inconsistency, the health check session completes successfully. Otherwise, the session completes with an error. Depending on the detected data inconsistency, Veeam Backup for Nutanix AHV performs the following operations:

- If the health check detects corrupted metadata in a full restore point, Veeam Backup for Nutanix AHV marks the backup chain as corrupted in the configuration database. During the next backup job session, Veeam Backup for Nutanix AHV copies the full instance image, creates a full restore point in the backup repository and starts a new backup chain in the backup repository.
- If the health check detects corrupted disk blocks in a restore point, Veeam Backup for Nutanix AHV marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup job session, Veeam Backup for Nutanix AHV copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted, and saves these data blocks to the latest restore point that has been created during the current session.

## Writing Custom Scripts

To coordinate application activity on VMs included into a protection domain before and after a PD snapshot is taken, you can run custom scripts. A custom script performs pre-freeze operations, takes the PD snapshot, and finally performs post-thaw operations.

### IMPORTANT

Veeam does not provide support for cases related to custom scripts.

You can use any programming language to write a script. However, you must specify the used interpreter in the shebang line, for example: `#!/usr/bin/python3`. For more information, see [Appendix A. Custom Script Samples](#).

### NOTE

You cannot use binary files for custom scripts.

While writing custom scripts, use the following arguments:

Argument	Description
<code>--clusterId</code>	[Applies only to <a href="#">Prism Central deployment</a> ] ID of a cluster where the protection domain is configured.
<code>--pdName</code>	Name of a protection domain for which a PD snapshot is created.
<code>--jobName</code>	Name of a job that runs the script.
<code>--logDir</code>	Path to a folder where script logs are stored.
<code>--logLevel</code>	Level of logging.



To specify parameters required for connecting to the Nutanix AHV cluster where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixClusterIp = os.getenv('NUTANIX_CLUSTER_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

To specify parameters required for connecting to the Prism Central where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixPrismCentralIp = os.getenv('NUTANIX_PRISM_CENTRAL_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

You can find script logs in the `custom_script.log` file stored in the `/var/log/nxbackupagent/Backup/<job_name_job_uuid>/<pd_name>/` folder on the backup appliance. The file also contains a JSON string that includes the script execution status, an error description (if any) and an identifier of the PD snapshot if it has been created, for example:

- Script execution succeeded:

```
Result: {"status": "Success", "errorMessage": "", "oob_schedule_id": 7885270}
```

- Script execution failed:

```
Result: {"status": "Failed", "errorMessage": "Unable to get vm_id list for Protection Domain='VeeamBackupProtection': Authentication failed.", "oob_schedule_id": ""}
```

- Script execution completed with a warning:

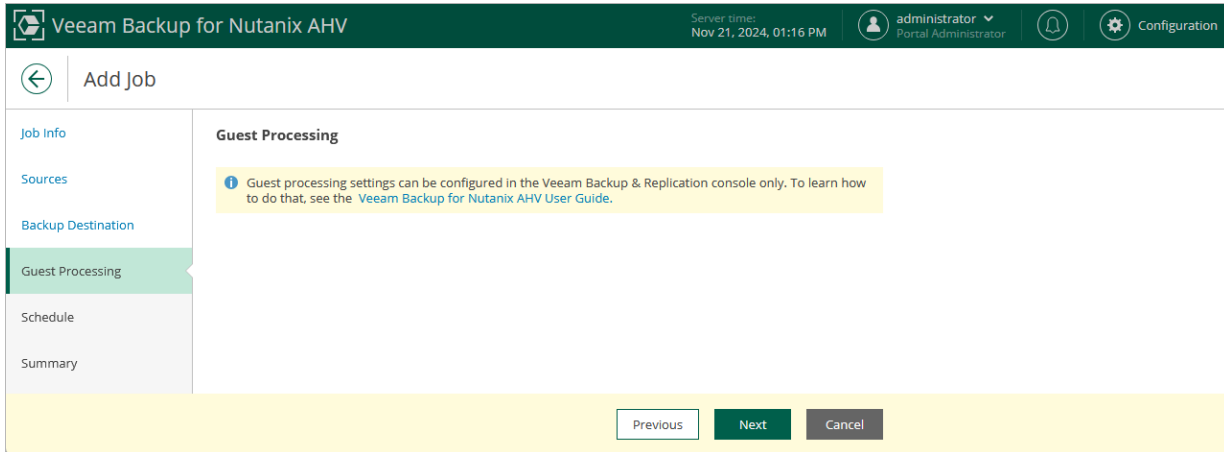
```
Result: {"status": "Warning", "errorMessage": "Pre-freeze took more time than expected", "oob_schedule_id": "44509"}
```

## TIP

You can track the script execution progress on the [Session Logs page](#) in the backup appliance web console. Alternatively, you can [download log files](#) to your local machine.

## Step 5. Skip Guest Processing

If you want to create transactionally consistent backups that will guarantee proper recovery of VM applications, you must enable application-aware processing. However, guest processing settings are not available in the web console. That is why you can skip the **Guest Processing** step, complete the **Add Job** wizard – and then edit the job settings in the Veeam Backup & Replication console to enable application-aware processing as described in section [Configuring Guest Processing Options](#).



## Step 6. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the backup job automatically according to a specific backup schedule. The backup schedule defines how often data of the VMs added to the backup job will be backed up.

1. To start the backup job automatically according to a specific backup schedule, in the **Main job schedule** section, set the **Run job automatically** toggle to *On*, and choose one of the following schedule types:
  - **Daily at this time** – the backup job will create restore points at a specific time on specific days.
  - **Monthly at this time** – the backup job will create restore points once a month on a specific day.
  - **Periodically every** – the backup job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the backup job again if it fails on the first try. To do that, select the **Automatic Retry** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying backup jobs, Veeam Backup for Nutanix AHV processes only those VMs that failed to be backed up during the previous attempt.

2. To [schedule active full backups](#), in the **Full backup schedule** section, set the **Create active full backups periodically** toggle to *On*, and choose whether you want to create active full backups on specific days every week or on specific days of specific months.

Alternatively, you can create active full backups manually when needed. For more information, see [Creating Active Full Backup](#).

### IMPORTANT

Do not schedule synthetic and active full backups to run at the same time. Due to technical limitations, Veeam Backup for Nutanix AHV will be unable to create synthetic full backups according to the specified schedule.

3. To [schedule synthetic full backups](#), in the **Full backup schedule** section, set the **Create synthetic full backups periodically** toggle to *On*, and choose whether you want to create synthetic full backups on specific days every week or on specific days of specific months.
4. To instruct Veeam Backup for Nutanix AHV to periodically [perform a health check](#) for backups created by the backup job, in the **Health check schedule** section, set the **Enable health check** toggle to *On*, and specify a schedule for the health check to run.

## IMPORTANT

- It is recommended that the backup and health check schedules configured for the job do not overlap to avoid data access issues.
- If you have selected an off-premise cloud object storage repository as the target location for backups at [step 4](#), it is recommended that a [helper appliance is configured in the repository settings](#). Otherwise, additional data transfer costs may occur.

The screenshot shows the 'Add Job' configuration page in the Veeam Backup for Nutanix AHV interface. The page is titled 'Add Job' and has a navigation menu on the left with options: Job info, Sources, Backup Destination, Guest Processing, Schedule (selected), and Summary. The main content area is titled 'Schedule' and includes the following settings:

- Schedule**  
Specify how often Veeam Backup for Nutanix AHV should run the job and create full backups. You can also enable health check to verify restore points.
- Main job schedule**  
Run job automatically:  On  
Every day at 10:00 PM, retry 3 times with a 10-minute timeout
- Full backup schedule**  
Create active full backups periodically:  Off  
Create synthetic full backups periodically:  On  
Every Sat
- Health check schedule**  
Enable health check:  On  
Every last Fri of every month at 10:00 PM

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

# Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

## TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.

The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV, specifically the 'Summary' step. The interface includes a top navigation bar with the product name, server time (Nov 21, 2024, 01:18 PM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar lists navigation options: Job info, Sources, Backup Destination, Guest Processing, Schedule, and Summary (which is highlighted). The main content area is titled 'Summary' and contains the following information:

- Summary**: Review the job summary, and click Finish to save and apply the job settings. A 'Copy to clipboard' icon is present.
- General**:
  - Name: LMS Daily Backup
  - Description: Backup of the LMS Server VM and Media Storage VM
- Protected resources**:
  - Virtual machines: 2 VMs
- Backup destination**:
  - Target repository: Default Backup Repository
  - Retention: 7 Days
- Run the job when I click Finish

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish' (highlighted in green), and 'Cancel'.

# Creating Backup Jobs Using Veeam Backup & Replication Console

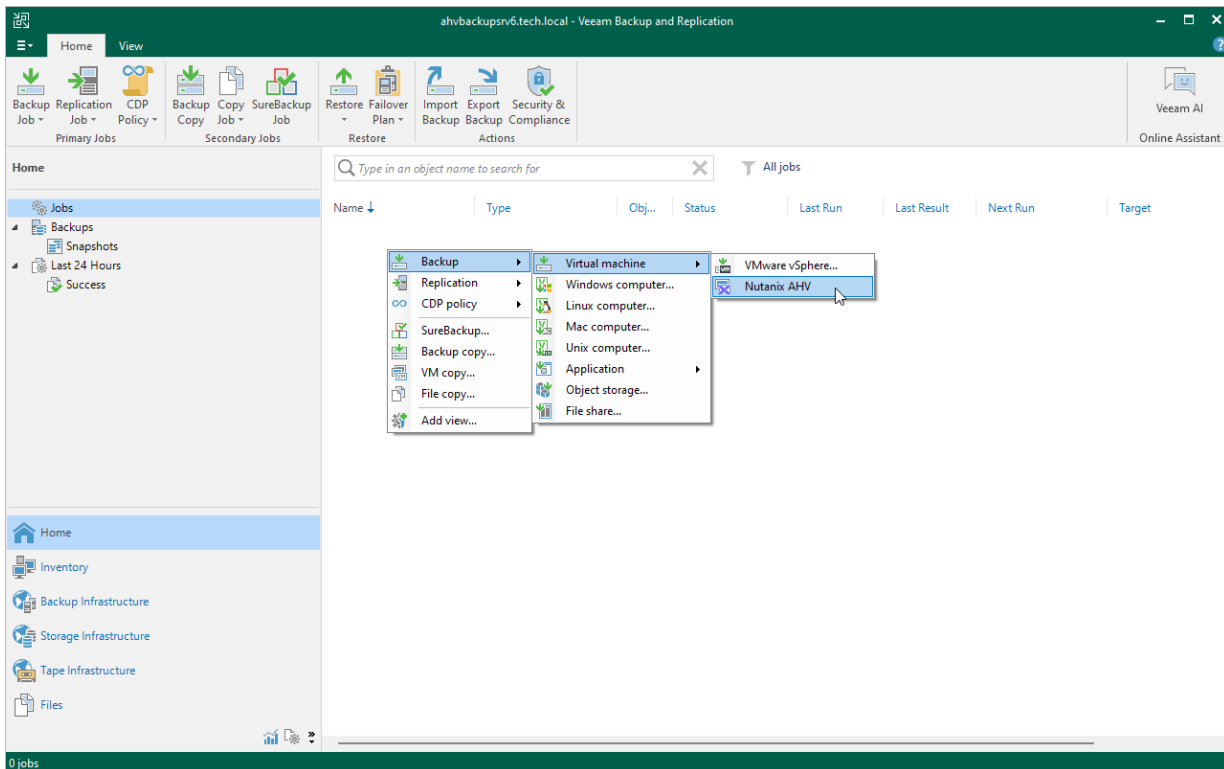
To create a backup job using the Veeam Backup & Replication console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Configure general settings.](#)
4. [Selects resources to back up.](#)
5. [Configure backup target settings.](#)
6. [Enable guest processing.](#)
7. [Create a schedule for the backup job.](#)
8. [Finish working with the wizard.](#)

# Step 1. Launch New Job Wizard

To launch the **New Job** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. On the ribbon, click **Backup Job > Virtual Machine > Nutanix AHV**, or right-click the working area and select **Backup > Virtual machine > Nutanix AHV**.



## Step 2. Specify Job Name and Description

At the **Name** step of the wizard, select the **Backup job** option. Then, use the **Name** and **Description** fields to specify a name for the new backup job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.

New Job

**Name**  
Type in a name and description for this job.

Name: LMS Daily Backup

Description: Backup of the LMS Server VM

Type:

- Backup job  
Configure a regular backup job to create backups of individual VMs.
- Snapshot job  
Configure a snapshot-only job to create snapshots of individual VMs.
- PD snapshot job  
Configure a protection domain snapshot-only job to create a snapshot of specific PD.

< Previous   Next >   Finish   Cancel



## Step 3. Configure Backup Source Settings

At the **Virtual Machines** step of the wizard, specify the following backup source settings:

1. [Choose resources to back up.](#)
2. [Choose disks and volume groups to protect.](#)

## Step 3a. Choose Resources

First, at the **Virtual Machines** step of the wizard, specify the backup scope – resources that Veeam Backup for Nutanix AHV will back up:

1. Click **Add**.
2. In the **Add Objects** window, choose whether you want to back up all VMs in the cluster, only specific VMs or protection domains. In the [Prism Central deployment](#), you can also back up VMs and clusters assigned to a specific category or all VMs managed by a Prism Central.

To view the list of available protection domains, click the **PDs** icon on the toolbar at the top right corner of the window. If you add a protection domain, Veeam Backup for Nutanix AHV will regularly check for new consistency groups (VMs and volume groups) added to the domain and automatically update the job settings to include these groups in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

To view the list of available categories, click the **Categories** icon on the toolbar at the top right corner of the window. If you add a category, Veeam Backup for Nutanix AHV will regularly check for new VMs and clusters assigned to the category and automatically update the job settings to include these resources in the backup scope. For a category to be displayed in the list of the available categories, it must be configured in the Nutanix AHV Prism Central as described in [Nutanix documentation](#).

### TIP

As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Exclusions** and specify the VMs or protection domains that you do not want to back up – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

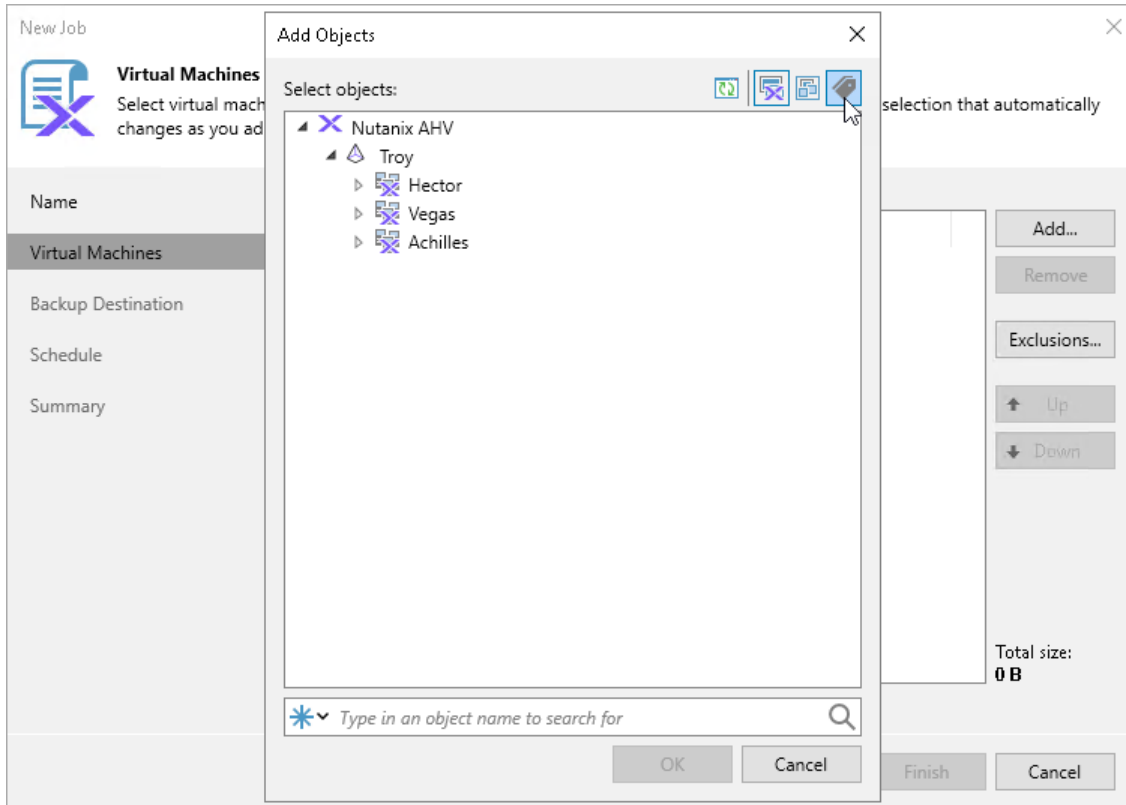
### NOTE

Consider the following:

- If you include a resource into the backup scope for multiple times (for example, an individual VM and a PD that contains this VM), Veeam Backup for Nutanix AHV will process this resource only once.
- If you include a protection domain, category, cluster or Prism Central into the backup scope, VMs in this object are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs – not as a part of the protection domain, category, cluster or Prism Central.

[Applies only to the [Prism Central deployment](#)] To instruct Veeam Backup for Nutanix AHV to [obtain VM data from a replica cluster](#), select the **Backup from Prism Central replica (if available)** check box. Using replica clusters help you reduce impact of backup operations on performance of the production environment. If Veeam Backup for Nutanix AHV fails to obtain data from a replica cluster, backup will be still performed using VM data obtained from the main cluster.

By default, jobs process all disks and volume groups attached to VMs included into the backup scope. However, you can protect only specific disks and volume groups of the selected resources. For more information, see [Step 3b. Choose Disks and Volume Groups](#).



## Related Topics

[Snapshot Types](#)

## Step 3b. Choose Disks and Volume Groups

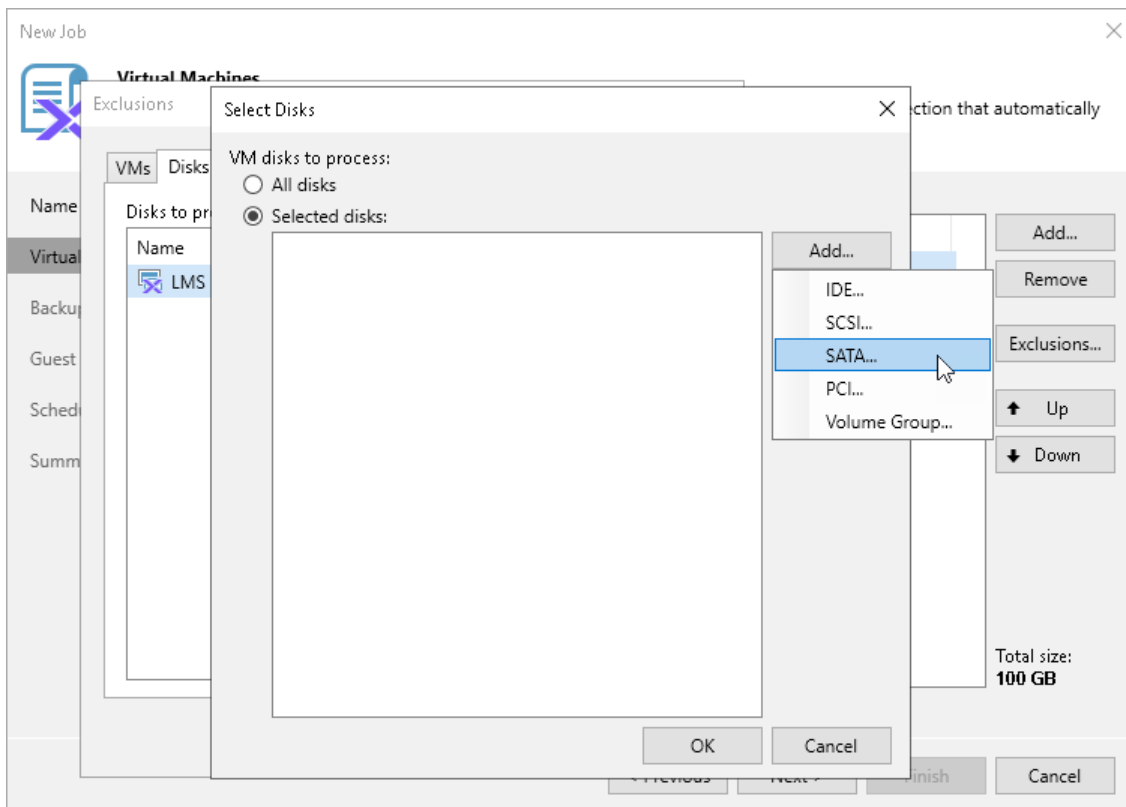
Second, at the **Virtual Machines** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to back up only specific virtual disks and volume groups related to the selected backup scope:

1. Click **Exclusions**.
2. In the **Exclusions** window, switch to the **Disks** tab and click **Add**.
3. In the **Add Objects** window, select a resource that you have added to the backup scope at [step 3a](#), and click **OK**.
4. Back to the **Exclusions** window, select the resource and click **Edit**.
5. In the **Select Disks** window, select the **Selected Disks** option, click **Add** and choose a bus type of the disks that you want to back up. Then, select the necessary disks and volume groups.

Disks and volume groups that you do not select will be excluded from the backup job.

### NOTE

If you configure multiple disk protection rules, specific rules will override general ones. For example, if you add a rule for a protection domain and for a VM included in this domain, Veeam Backup for Nutanix AHV will process the VM disks according to the rule configured for the VM.



## Step 4. Specify Backup Job Settings

At the **Backup Destination** step of the wizard, do the following:

1. From the **Backup repository** drop-down list, select a backup repository where you want to store backups.  
For a backup repository to be displayed in the list of the available repositories, it must be [added to the backup infrastructure](#), and the Nutanix AHV backup appliance must [have access to the repository](#).
2. In the **Retention policy** section, choose a retention policy that Veeam Backup for Nutanix AHV will apply to backups created by the job:
  - Select *days* if you want to keep restore points in a backup chain for the allowed period of time. If a restore point is older than the specified limit, Veeam Backup for Nutanix AHV removes it from the chain.
  - Select *restore points* if you want a backup chain to contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain.

For more information on how Veeam Backup for Nutanix AHV tracks and removes redundant restore points, see [Retention Policies](#).

### NOTE

If the UUID of a VM changes (for example, if the VM migrates to another cluster), Veeam Backup for Nutanix AHV will be unable to continue the backup chain for this VM. After you re-add the VM to the backup job, Veeam Backup for Nutanix AHV will start a new backup chain for it. However, you will still be able to perform restore operations using backups from the old backup chain.

To help you implement a comprehensive backup strategy, Veeam Backup for Nutanix AHV allows you to [enable long-term retention policy for backups](#) and to [configure backup job advanced settings](#) (for example, enable health check, schedule full backups, plan backup maintenance and upload custom scripts).

**New Job** [Close]

**Backup Destination**  
Specify a backup repository to store backups produced by this job and customize advanced job settings.

**Name** Backup repository: Default Backup Repository

**Virtual Machines** 79.7 GB free of 129 GB

**Backup Destination** Retention policy: 7 days [!]

**Schedule**  Keep certain full backups longer for archival purposes [Configure...]

**Summary** 1 weekly, 6 monthly, 2 yearly

Advanced job settings include schedules for active and synthetic full backups, backup file maintenance configuration and specify scripts for PD snapshot creation. [Advanced...]

[< Previous] [Next >] [Finish] [Cancel]

## Configuring GFS Policy Schedules

Grandfather-Father-Son (GFS) policy allows you to leverage full backups for long-term retentions instead of creating a new full backup every time. The mechanism simplifies the backup schedule and optimizes the backup performance.

Veeam Backup for Nutanix AHV re-uses full backups created according to the backup job schedule to achieve the desired retention for GFS policy schedules (weekly, monthly and yearly). Each full backup is marked with a flag of the related GFS policy schedule type: the (W) flag is used to mark full backups created weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for Nutanix AHV uses these flags to control the retention period for the created full backups. Once a flag of a GFS policy schedule is assigned to a full backup, this full backup can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the full backup. If the full backup does not have any other flags assigned, it is removed according to the short-term retention policy settings. For more information on the GFS flag assignment and removal, see the Veeam Backup & Replication User Guide, section [Long-Term Retention Policy \(GFS\)](#).

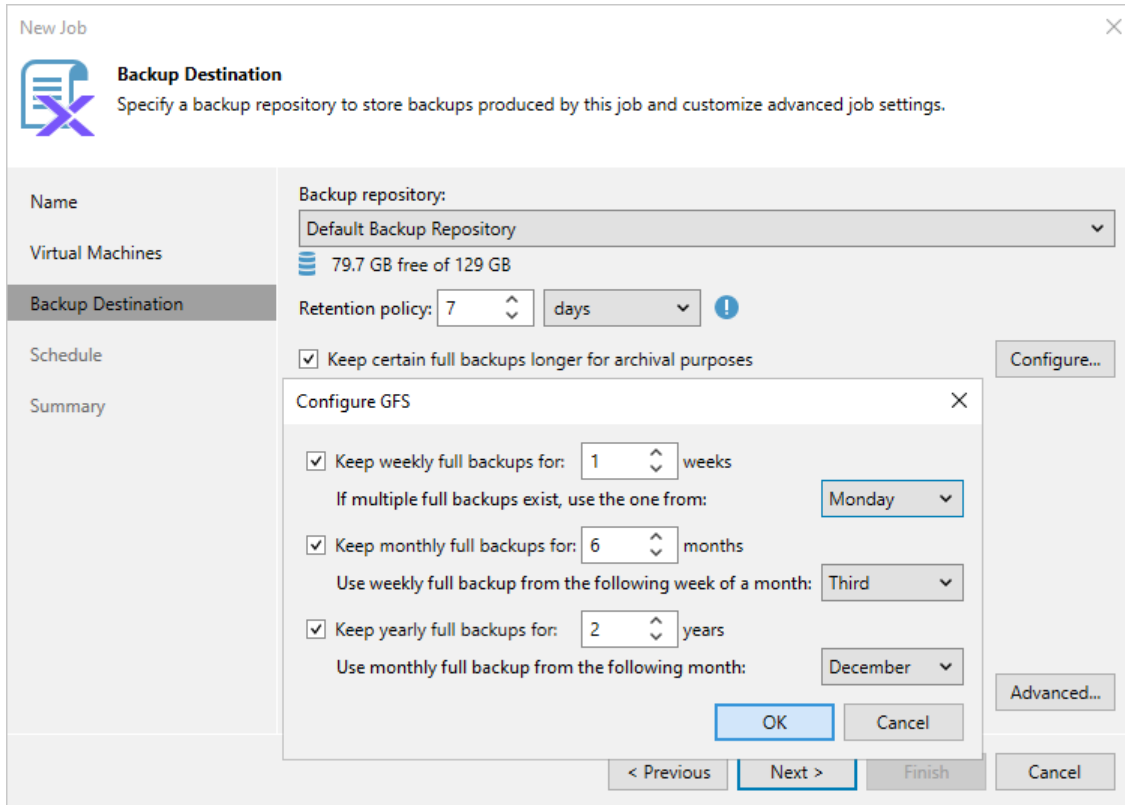
To configure a GFS policy schedule, select the **Keep certain full backups longer for archival purposes** check box and click **Configure**. Then, specify the following options in the **Configure GFS** window:

- **Keep weekly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created within a week or on a specific day for a specified number of weeks.
- **Keep monthly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created during a specific week for a specified number of months.
- **Keep yearly full backups** – Veeam Backup for Nutanix AHV will keep a full backup created in a specific month for a number of years.

After you configure the GFS retention policy settings, [schedule active full or synthetic full backups](#). Otherwise, no new full backups will be automatically produced, and Veeam Backup for Nutanix AHV will be unable to leverage them for long-term retentions.

## NOTE

If you choose an object storage repository to store backups produced by the backup job, you cannot enable synthetic full backups. However, if you configure a GFS policy, synthetic backups will be automatically created according to the specified GFS schedule and marked with an appropriate GFS flag.



The screenshot shows the 'New Job' dialog box with the 'Backup Destination' tab selected. The 'Backup repository' is set to 'Default Backup Repository' with 79.7 GB free of 129 GB. The retention policy is set to 7 days. The 'Keep certain full backups longer for archival purposes' checkbox is checked. The 'Configure GFS' dialog is open, showing the following settings:

- Keep weekly full backups for: 1 weeks
- If multiple full backups exist, use the one from: Monday
- Keep monthly full backups for: 6 months
- Use weekly full backup from the following week of a month: Third
- Keep yearly full backups for: 2 years
- Use monthly full backup from the following month: December

Buttons for 'OK', 'Cancel', 'Advanced...', 'Previous', 'Next >', 'Finish', and 'Cancel' are visible at the bottom of the dialog.

## Configuring Advanced Settings

To configure advanced settings for the backup job, do the following:

1. Click **Advanced**.
2. To [schedule synthetic full backups](#), on the **Backup** tab of the **Advanced settings** window, select the **Create synthetic full backups periodically** check box, click **Configure** and choose whether you want to create synthetic full backups on specific days every week or on specific days of specific months.
3. To [schedule active full backups](#), on the **Backup** tab of the **Advanced settings** window, select the **Create active full backups periodically** check box, click **Configure** and choose whether you want to create active full backups on specific days every week or on specific days of specific months.

Alternatively, you can create active full backups manually when needed. For more information, see [Creating Active Full Backup](#).

## IMPORTANT

Do not schedule synthetic and active full backups to run at the same time. Due to technical limitations, Veeam Backup for Nutanix AHV will be unable to create synthetic full backups according to the specified schedule.

4. To instruct Veeam Backup for Nutanix AHV to periodically [perform a health check](#) for backups created by the backup job, on the **Maintenance** tab of the **Advanced settings** window, select the **Perform backup files health check (detects and auto-heals corruption)** check box, click **Configure** and specify a schedule for the health check to run.

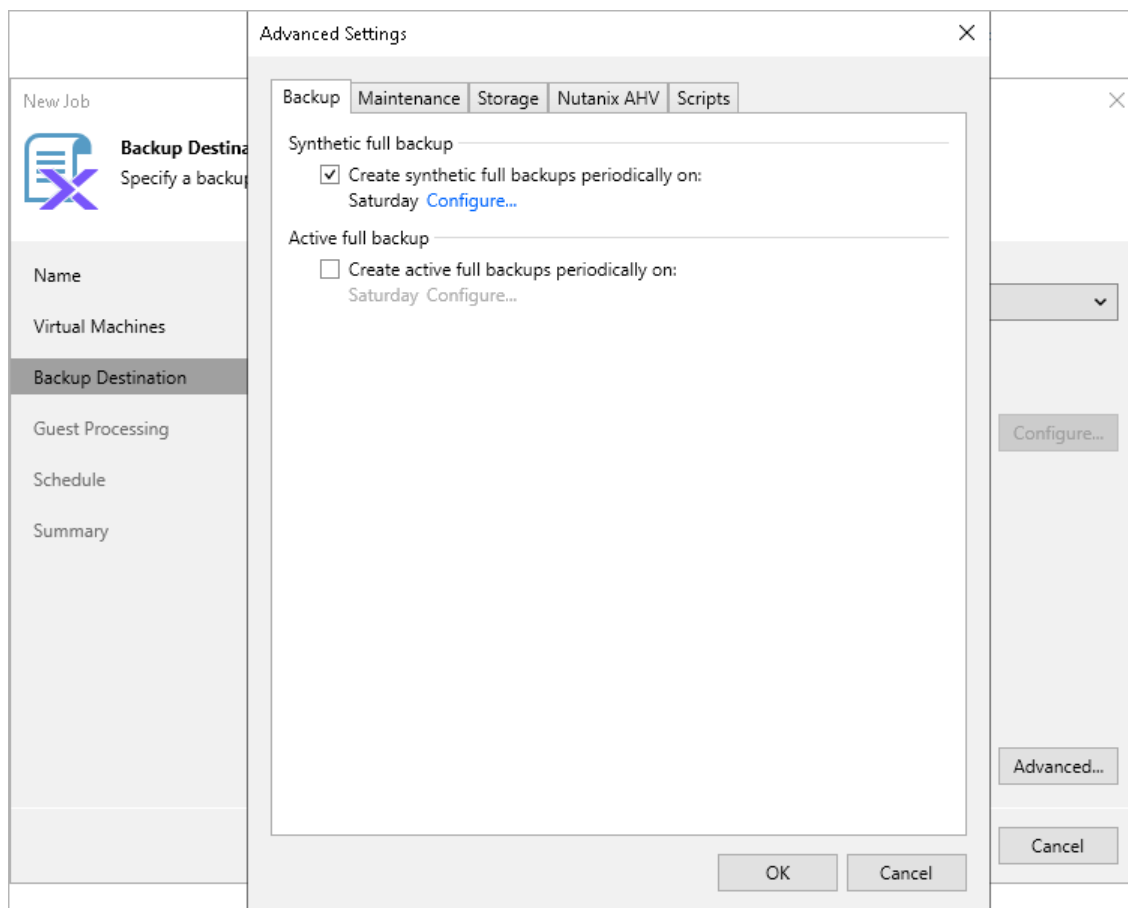
## IMPORTANT

- It is recommended that the backup and health check schedules configured for the job do not overlap to avoid data access issues.
  - If you have selected an off-premise cloud object storage repository as the target location for backups at [step 4](#), it is recommended that a [helper appliance is configured in the repository settings](#). Otherwise, additional data transfer costs may occur.
5. To configure retention settings for backups of VMs that are no longer processed by the backup job, on the **Maintenance** tab of the **Advanced settings** window, select the **Remove deleted items data after** check box, and specify the number of days during which Veeam Backup for Nutanix AHV will keep backups of VMs excluded from the job.
  6. To decrease the size of the backup, on the **Storage** tab of the **Advanced settings** window, from the **Compression level** drop-down list, select a compression level for the backup: *None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*. For more information on data compression, see [Compression and Deduplication](#).
  7. To optimize job performance and storage usage, on the **Storage** tab of the **Advanced settings** window, from the **Storage optimization** drop-down list, select the block size that will be used to process VMs. For more information on the data block sizes and how they affect performance, see [Storage Optimization](#).
  8. To instruct Nutanix AHV to freeze applications running on VMs while snapshots are taken, set the **Enable Nutanix Guest Tools quiescence** toggle to *On* and choose how Nutanix AHV will process transaction logs.  
Keep in mind that the **Never truncate transaction logs** (VSS\_BT\_COPY) option may significantly increase the amount of storage space consumed by VMs that function as Microsoft Exchange Mail Servers.
  9. [Applies only if you have included protection domains in the backup job] To instruct Veeam Backup for Nutanix AHV to execute a custom script while running the backup job, on the **Scripts** tab of the **Advanced settings** window, select the **Run the following script to create a PD snapshot** check box, and click **Browse** to upload a script file. For more information, see [Writing Custom Scripts](#).



## NOTE

When running custom scripts, Veeam Backup for Nutanix AHV uses native Nutanix AHV capabilities to take PD snapshots. If Nutanix AHV fails to create a PD snapshot, the backup job will also fail to complete successfully. For more information on the prerequisites for creating PD snapshots, see [Snapshot Types](#).



## How Health Check Works

When Veeam Backup for Nutanix AHV saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup for Nutanix AHV verifies the availability of data blocks, which are required to restore from the recent point only, and uses the saved values to ensure that the full restore points being verified are consistent.

On the day scheduled for a health check to run, Veeam Backup for Nutanix AHV starts a new health check session. For each restore point in the backup chain, Veeam Backup for Nutanix AHV calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for Nutanix AHV also checks whether data blocks that are required to rebuild the restore point are available.

If Veeam Backup for Nutanix AHV does not detect data inconsistency, the health check session completes successfully. Otherwise, the session completes with an error. Depending on the detected data inconsistency, Veeam Backup for Nutanix AHV performs the following operations:

- If the health check detects corrupted metadata in a full restore point, Veeam Backup for Nutanix AHV marks the backup chain as corrupted in the configuration database. During the next backup job session, Veeam Backup for Nutanix AHV copies the full instance image, creates a full restore point in the backup repository and starts a new backup chain in the backup repository.

- If the health check detects corrupted disk blocks in a restore point, Veeam Backup for Nutanix AHV marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup job session, Veeam Backup for Nutanix AHV copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted, and saves these data blocks to the latest restore point that has been created during the current session.

## Writing Custom Scripts

To coordinate application activity on VMs included into a protection domain before and after a PD snapshot is taken, you can run custom scripts. A custom script performs pre-freeze operations, takes the PD snapshot, and finally performs post-thaw operations.

### IMPORTANT

Veeam does not provide support for cases related to custom scripts.

You can use any programming language to write a script. However, you must specify the used interpreter in the shebang line, for example: `#!/usr/bin/python3`. For more information, see [Appendix A. Custom Script Samples](#).

### NOTE

You cannot use binary files for custom scripts.

While writing custom scripts, use the following arguments:

Argument	Description
<code>--clusterId</code>	[Applies only to <a href="#">Prism Central deployment</a> ] ID of a cluster where the protection domain is configured.
<code>--pdName</code>	Name of a protection domain for which a PD snapshot is created.
<code>--jobName</code>	Name of a job that runs the script.
<code>--logDir</code>	Path to a folder where script logs are stored.
<code>--logLevel</code>	Level of logging.

To specify parameters required for connecting to the Nutanix AHV cluster where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixClusterIp = os.getenv('NUTANIX_CLUSTER_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

To specify parameters required for connecting to the Prism Central where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixPrismCentralIp = os.getenv('NUTANIX_PRISM_CENTRAL_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

You can find script logs in the `custom_script.log` file stored in the `/var/log/nxbackupagent/Backup/<job_name_job_uuid>/<pd_name>/` folder on the backup appliance. The file also contains a JSON string that includes the script execution status, an error description (if any) and an identifier of the PD snapshot if it has been created, for example:

- Script execution succeeded:

```
Result: {"status": "Success", "errorMessage": "", "oob_schedule_id": 7885270}
```

- Script execution failed:

```
Result: {"status": "Failed", "errorMessage": "Unable to get vm_id list for Protection Domain='VeeamBackupProtection': Authentication failed.", "oob_schedule_id": ""}
```

- Script execution completed with a warning:

```
Result: {"status": "Warning", "errorMessage": "Pre-freeze took more time than expected", "oob_schedule_id": "44509"}
```

## TIP

You can track the script execution progress on the [Session Logs page](#) in the backup appliance web console. Alternatively, you can [download log files](#) to your local machine.

## Step 5. Configure Guest Processing Options

At the **Guest Processing** step of the wizard, you can specify the following settings:

- [Enable application-aware processing](#) – to create transactionally consistent backups that will guarantee proper recovery of VM applications, without data loss.

For VMs running Microsoft SQL Server, Oracle Server or PostgreSQL Server applications, you can also instruct Veeam Backup & Replication to periodically back up transaction logs. This will allow you to restore your databases to specific points in time as described in the Veeam Enterprise Manager User Guide, section [Restoring Point-in-Time State](#).

- [Enable guest file system indexing and malware detection](#) – to create a catalog of guest OS files that will allow you to search for specific items during file-level restore. This will also allow you to receive reports about malware files and suspicious system activity detected on VMs included into the backup scope.
- [Choose guest interaction proxies](#) – to select specific Windows servers that Veeam Backup & Replication will use when communicating with guest OSes of Windows VMs included into the backup scope.
- [Manage VM guest OS credentials](#) – to specify credentials that Veeam Backup & Replication will use to access guest OSes of all VMs included into the backup scope.

## Considerations and Limitations

If you enable application-aware processing or guest files system indexing, consider the following:

- Veeam Backup for Nutanix AHV will not be able to [obtain VM data from replica clusters](#).
- Veeam Backup for Nutanix AHV will not be able to [create PD snapshots](#) of protection domains included into the backup scope – it will back up VMs and their volume groups as if processing individual virtual machines.

- Veeam Backup for Nutanix AHV will not be able to [use Kerberos authentication](#) while connecting to guest OSes of the processed VMs.

New Job
✕

**Guest Processing**

Choose guest OS processing options available for running VMs.

Name	<input checked="" type="checkbox"/> <b>Enable application-aware processing</b> Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot. Customize application handling options for individual machines and applications <span style="float: right;">Applications...</span>
Virtual Machines	
Backup Destination	
<b>Guest Processing</b>	<input checked="" type="checkbox"/> <b>Enable guest file system indexing and malware detection</b> Indexing enables global file search functionality, automatic detection of suspicious file system activity and known malware files. Customize advanced guest file system indexing options for individual machines <span style="float: right;">Indexing...</span>
Schedule	Guest interaction proxy: <input type="text" value="Automatic selection"/> <span style="float: right;">Choose...</span>
Summary	Guest OS credentials: <input type="text" value="Select existing credentials or add new"/> <span style="float: right;">Add...</span> <div style="text-align: right; margin-top: 5px;"><a href="#">Manage accounts</a></div> Customize guest OS credentials for individual machines and operating systems <span style="float: right;">Credentials...</span>
	Verify network connectivity and credentials for each machine included in the job <span style="float: right;">Test Now</span>

< Previous
Next >
Finish
Cancel

## Step 5a. Enable Application-Aware Processing

To restore your applications without data loss, you must allow Veeam Backup & Replication to create application-consistent backups. To do that, select the **Enable application-aware processing** check box at the **Guest Processing** step of the wizard.

When creating application-consistent backups, Veeam Backup & Replication takes transactionally consistent VM snapshots while no write operations occur on VM disks. To do that, Veeam Backup & Replication quiesces applications on the processed VMs and creates a consistent view of application data:

- To quiesce VSS-aware applications running on Windows-based VMs (such as MS SQL, MS Exchange, Microsoft Active Directory and Microsoft SharePoint), Veeam Backup & Replication leverages the [Microsoft VSS technology](#).
- To quiesce applications running on Linux-based VMs and non-VSS-aware applications running on Windows-based VMs, Veeam Backup & Replication runs custom scripts before and after the snapshot creation.

## Processing Transaction Logs

If you enable application-aware processing, Veeam Backup & Replication will back up and truncate transaction logs produced by VM applications every time the backup job starts. To change this behavior, you can do either of the following:

- Instruct Veeam Backup & Replication not to process and truncate logs. This will allow third-party backup solutions to perform VM guest-level backup and to maintain consistency of the database state.
- Instruct Veeam Backup & Replication to back up and truncate transaction logs more often. This will allow you to use application-consistent backups to restore your MS SQL, Oracle and PostgreSQL databases to specific points in time.

To configure log processing settings, do the following:

1. Click **Applications**.
2. In the **Application-Aware Processing Options** window, select the necessary resource and click **Edit**. You can configure guest processing settings for multiple resources at a time.

If you want to configure processing settings for a specific VM that is included into a protection domain, cluster, category or Prism Central, you must configure those settings separately. To do that, click **Add**, choose the necessary VM and click **Edit**.

3. To specify how Veeam Backup & Replication will process transaction logs of VSS-aware applications, select the **Process transaction logs with this job** option on the **General** tab of the **Processing Settings** window, switch to the **SQL** tab and follow the instructions provided in section [Microsoft SQL Server Transaction Log Settings](#).

If you do not want Veeam Backup & Replication to process and truncate transaction logs of VSS-aware applications, select the **Perform copy only** option. However, with this option selected, the backup job will produce copy-only backups that cannot be used to restore **MS SQL** databases to specific points in time. For more information on copy-only backups, see [Microsoft Docs](#).

4. To specify how Veeam Backup & Replication will process transaction logs of Oracle database systems, switch to the **Oracle** tab and follow the instructions provided in section [Oracle Archived Redo Log Settings](#).
5. To specify how Veeam Backup & Replication will process transaction logs of PostgreSQL database systems, switch to the **PostgreSQL** tab and follow the instructions provided in section [PostgreSQL WAL Files Settings](#).

6. To specify scripts that Veeam Backup & Replication will use to quiesce non-VSS-aware applications, switch to the **Scripts** tab and follow the instructions provided in section [Pre-Freeze and Post-Thaw Scripts](#).

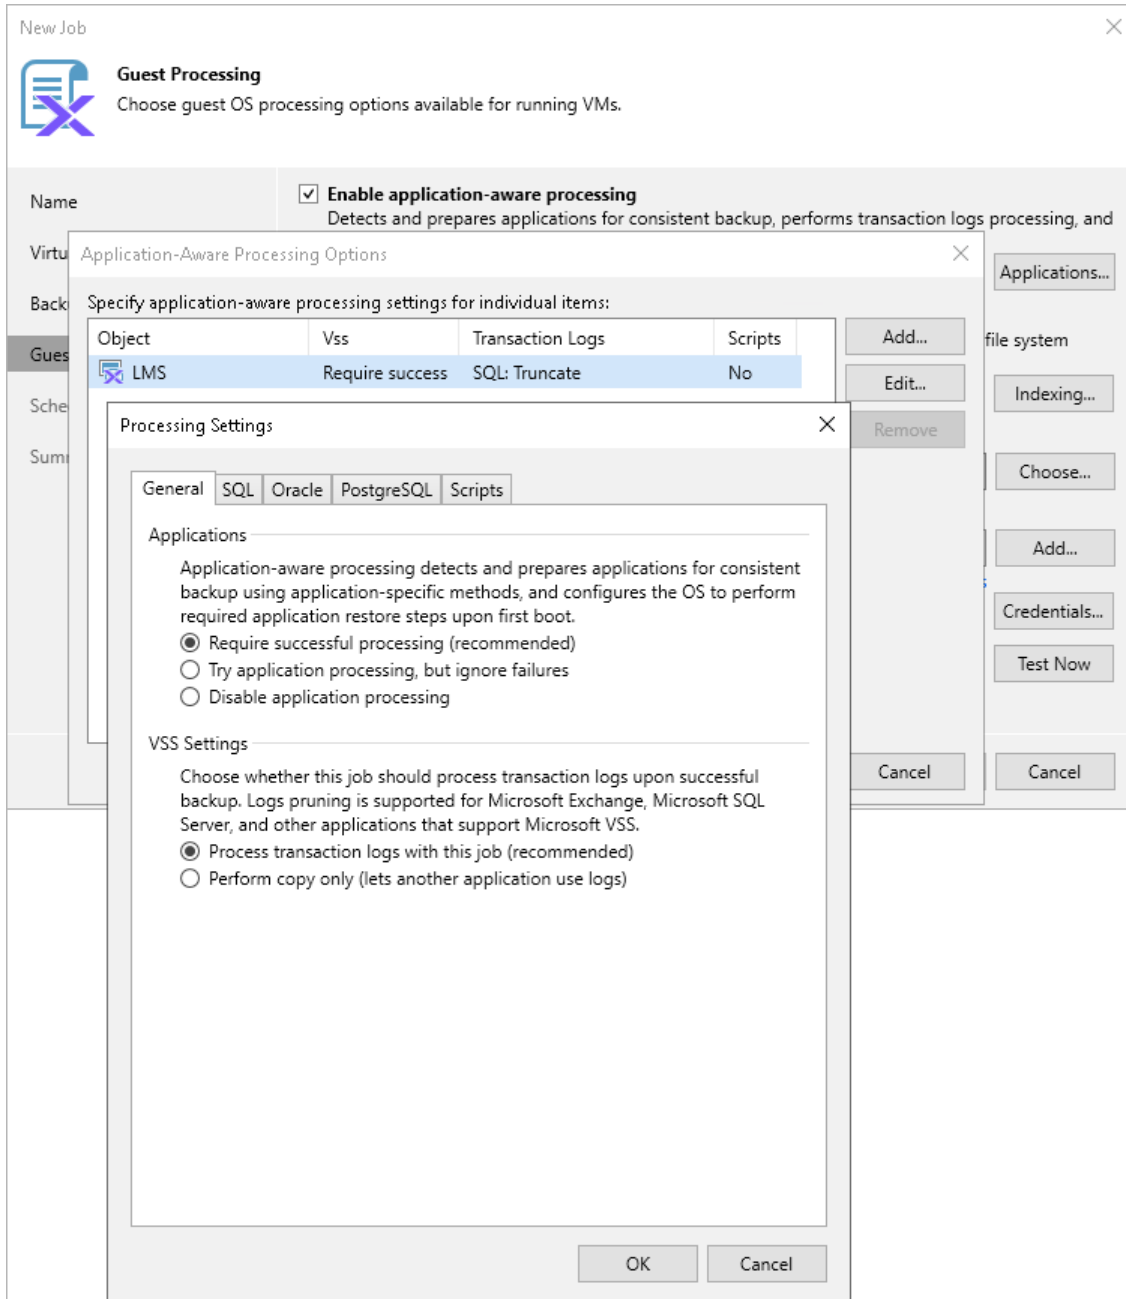
#### TIP

To instruct Veeam Backup & Replication not to perform application-aware processing for the selected resource at all, select the **Disable application processing** option.

## Handling Application-Aware Processing Errors

By default, Veeam Backup & Replication requires application-aware processing to finish without errors for the backup job to complete successfully. In case of an error, Veeam Backup & Replication terminates the backup operation, and the backup job will not process transaction logs until a new image-level backup is created for each of the VMs included into the backup scope.

To change this behavior and instruct Veeam Backup & Replication to proceed with the backup operation, creating a crash-consistent backup instead of an application-consistent backup, switch to the **General** tab of the **Processing Settings** window and select the **Try application processing, but ignore failures** option.



## Microsoft SQL Server Transaction Log Settings

To specify how Veeam Backup & Replication will process transaction logs for VSS-aware applications, switch to the **SQL** tab of the **Processing Settings** window and do the following:

- Select **Truncate logs** to truncate transaction logs of MS SQL and MS Exchange applications after successful backup. The non-persistent runtime components running on the VM guest OS will wait for the backup to complete successfully and then truncate transaction logs. If the job does not manage to back up the VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components or persistent components.



- Select **Do not truncate logs** to preserve transaction logs. When the backup job completes, Veeam Backup & Replication will not truncate transaction logs on the MS SQL or MS Exchange server.

It is recommended that you enable this option for databases that use the Simple recovery model. If you enable this option for databases that use the Full or Bulk-logged recovery model, transaction logs on the VM guest OS may grow large and consume all disk space. In this case, the database administrators must take care of transaction logs themselves.

- Select **Backup logs periodically** to back up transaction logs with Veeam Backup & Replication. Veeam Backup & Replication will periodically copy transaction logs to the backup repository and store them together with the image-level backup of the Microsoft SQL Server VM. During the backup job session, transaction logs on the VM guest OS will be truncated. For more information, see the Veeam Backup & Replication User Guide, section [Microsoft SQL Server Log Backup](#).

## IMPORTANT

If both Microsoft SQL Server and Oracle Server are installed on one VM, and this VM is processed by a job with log backup enabled for both applications, Veeam Backup & Replication will back up only Oracle transaction logs. Microsoft SQL Server transaction logs will not be processed.

If you choose to back up logs periodically, do the following:

1. In the **Backup logs every <N> minutes** field, specify the frequency for transaction log backup. By default, transaction logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
2. In the **Retain log backups** section, specify retention policy for transaction logs stored in the backup repository.
  - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and transaction log backups.
  - Select **Keep only last <N> days of log backups** to keep transaction logs for a specific number of days. By default, transaction logs are kept for 15 days. If you select this option, you must make sure that retention for transaction logs is not greater than retention for the image-level backups. For more information, see Veeam Backup & Replication User Guide, section [Retention for Transaction Log Backups](#).
3. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport transaction logs:
  - Select **Automatic selection** if you want Veeam Backup & Replication to choose an optimal log shipping server automatically. If the optimal shipping server is busy, Veeam Backup & Replication will direct the data flow to another shipping server so as not to lose data and comply with RPO. The process of transaction logs shipment does not require a dedicated server – Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.
  - To define a log shipping server explicitly, select **Use the specified servers only** and select check boxes next to servers that you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

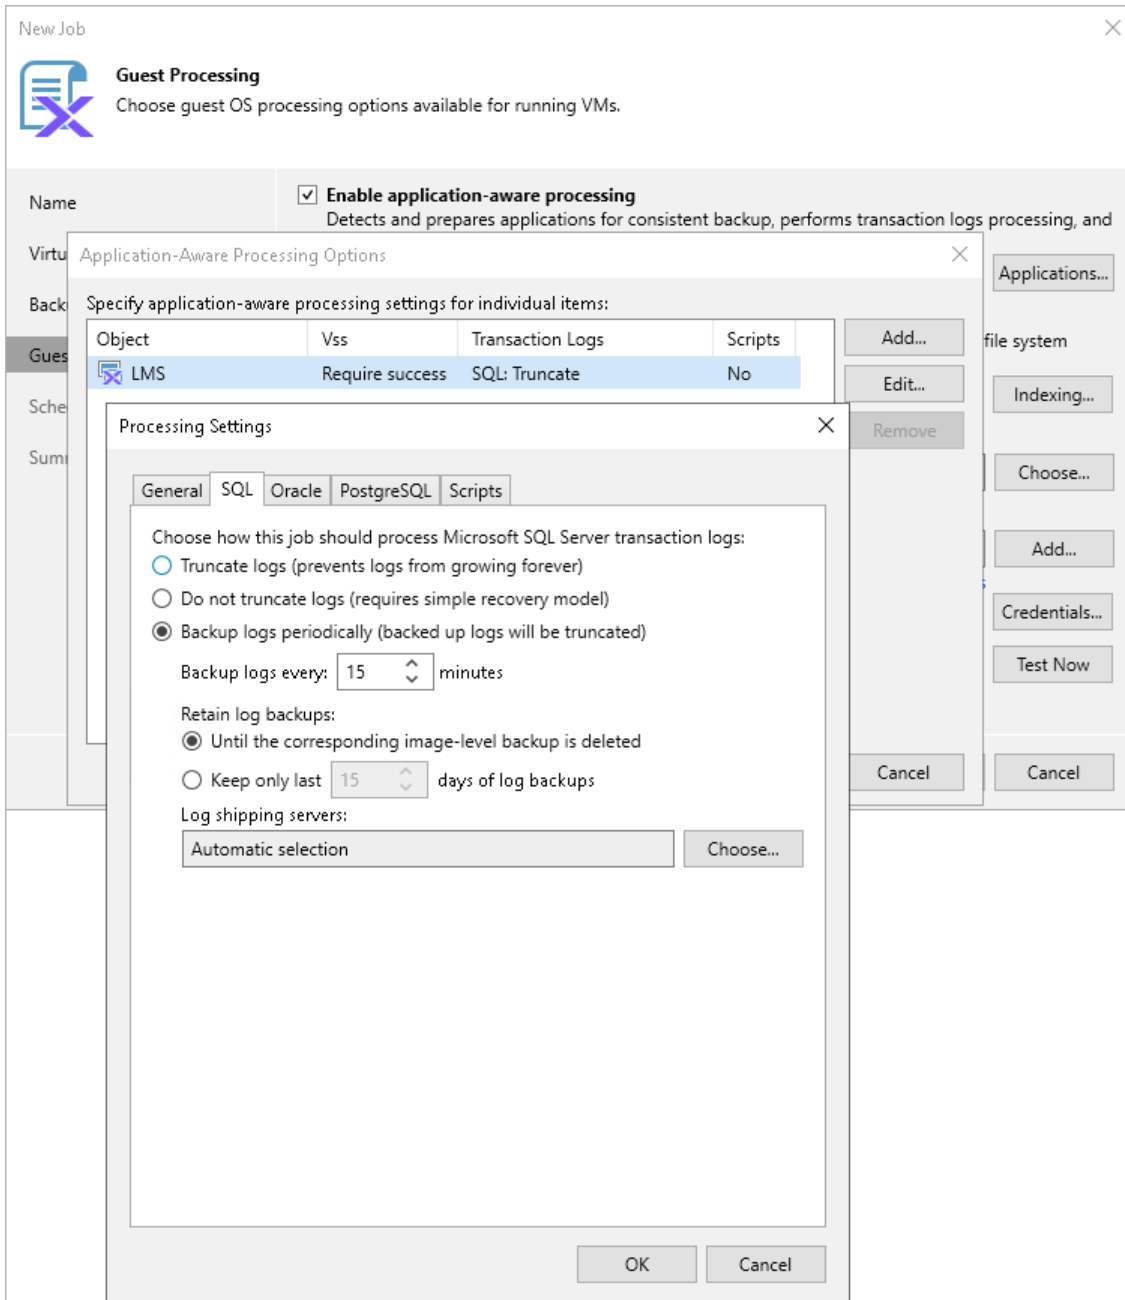
Ensure you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.

For more information on log shipping servers and how they are selected, see Veeam Backup & Replication User Guide, section [Log Shipping Servers](#).

## IMPORTANT

Consider the following:

- Veeam Backup & Replication automatically excludes its configuration database from application-aware processing during backup if the database is hosted without using SQL Server Always On Availability Group. Transaction logs for the configuration database are not backed up.
- If the Veeam Backup & Replication configuration database is hosted using SQL Server Always On Availability Group, you should manually exclude this database from application-aware processing during backup, as described in [this Veeam KB article](#). Otherwise, job processing will fail with the following error: Failed to freeze guest over network, wait timeout.



## Oracle Archived Redo Log Settings

To specify how Veeam Backup & Replication will process archived redo logs of the Oracle database, switch to the **Oracle** tab of the **Processing Settings** window and do the following:

1. In the **Specify Oracle account with SYSDBA privileges** section, specify a user account that Veeam Backup & Replication will use to connect to the Oracle database. The account must have privileges for the Oracle database listed in the section [Permissions for Guest Processing](#).

You can select **Use guest credentials** in the list of user accounts. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the Oracle database.

### NOTE

[For Windows-based machines] Make sure you add the `%ORACLE_HOME%\bin` variable to the environmental path variable on the machine with your Oracle database.

2. In the **Archived logs** section, specify how Veeam Backup & Replication must process archived redo logs on the Oracle VM:
  - Select **Do not delete archived logs** if you want Veeam Backup & Replication to preserve archived logs on the VM guest OS. When the backup job completes, the non-persistent runtime components will not delete archived logs.

It is recommended that you select this option for databases in the NOARCHIVELOG mode. If the database is in the ARCHIVELOG mode, archived logs on the VM guest OS may grow large and consume all disk space. In this case, database administrators must take care of archived logs themselves.

- Select **Delete logs older than <N> hours** or **Delete logs over <N> GB** if you want Veeam Backup & Replication to delete archived logs that are older than <N> hours or larger than <N> GB. The log size threshold refers not to the total size of all logs for all databases but to the log size of each database on the selected Oracle VM.

When the parent backup job (job creating an image-level backup) runs, Veeam Backup & Replication will wait for the backup to complete successfully and then trigger archived log deletion on the Oracle VM over Oracle Call Interface (OCI). If the primary job does not manage to back up the Oracle VM, the logs will remain untouched on the VM guest OS until the next start of the non-persistent runtime components.

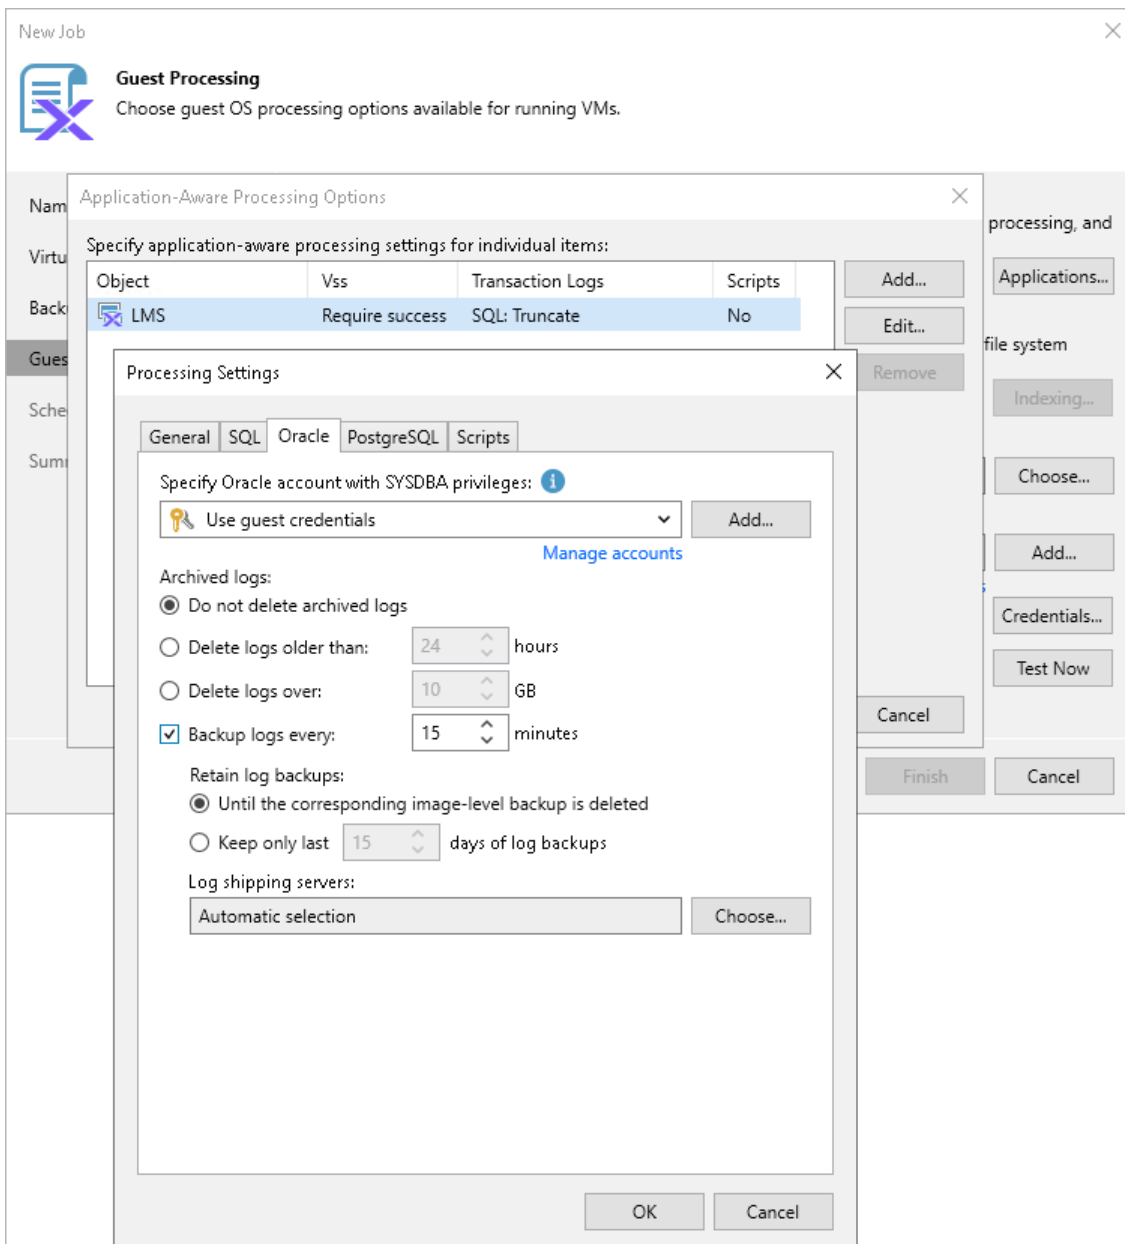
### TIP

Veeam Backup & Replication removes redo logs only after the parent backup job session. To remove redo logs more often, you can schedule the parent job to run more often.

3. To back up Oracle archived logs with Veeam Backup & Replication, select the **Backup log every <N> minutes** check box and specify the frequency for archived log backup. By default, archived logs are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
4. In the **Retain log backups** section, specify the retention policy for archived logs stored in the backup repository:
  - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for image-level backups and archived log backups.
  - Select **Keep only last <n> days** to keep archived logs for a specific number of days. By default, archived logs are kept for 15 days. If you select this option, you must ensure that retention for archived logs is not greater than retention for the image-level backups. For more information, see the Veeam Backup & Replication User Guide, section [Retention for Archived Log Backup](#).

5. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport archived logs:
  - Select **Automatic selection** if you want Veeam Backup & Replication to select an optimal log shipping server automatically. The process of archived logs shipment does not require a dedicated server – Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.
  - Select **Use the specified servers only** to define a log shipping server explicitly. In the **Log Shipping Servers** window, select check boxes next to the servers you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

Ensure you select a server that is not engaged in other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.



## PostgreSQL WAL Files Settings

To specify how Veeam Backup & Replication will process WAL files of the PostgreSQL database, switch to the **PostgreSQL** tab of the **Processing Settings** window and do the following:

1. From the **Specify PostgreSQL account with superuser privileges** drop-down list, select a user account that Veeam Backup & Replication will use to connect to the PostgreSQL instance. The account must have privileges described in section [Permissions](#). If you have not set up credentials beforehand, click the **Manage accounts** link or click **Add** on the right to add credentials.

By default, the *Use guest credentials* option is selected in the list. With this option selected, Veeam Backup & Replication will connect to the PostgreSQL instance under the account. In this case, Veeam Backup & Replication will use the account specified at the **Guest Processing** step of the wizard to access the VM guest OS and connect to the PostgreSQL instance.

Note that if you plan to select the **System user without password file (peer)** authentication method at the step 3 of this procedure, you can add a user account in the Credentials Manager without specifying the password for the account.

2. In the **Specified user is** section, specify how the user will authenticate against the PostgreSQL instance:
  - Select **Database user with password** if the account you specified at the step 2 is a PostgreSQL account, and you entered the password for this account in the **Credentials Manager**.

### NOTE

If you want Veeam Backup & Replication to use the user name map authentication, select **Database user with password** and leave the password field empty. Consider the following:

- **Guest OS credentials** specified at the **Guest Processing** step of the wizard will be used as the System-Username.
- PostgreSQL account specified at the step 2 will be used as the PG-Username.

For more information about the user name maps, see [PostgreSQL documentation](#).

- Select **Database user with password file (.pgpass)** if the password for the account you specified at the step 2 is defined in the `.pgpass` configuration file on the PostgreSQL VM. The password file must be located in the user's home directory. For more information about the password file, see [PostgreSQL documentation](#).
- Select **System user without password file (peer)** if you want Veeam Backup & Replication to use the peer authentication method. In this case, Veeam Backup & Replication will use the account you specified at the step 2 as the OS account and as the PostgreSQL account to connect to PostgreSQL. For more information about the peer authentication method, see [PostgreSQL documentation](#).

### IMPORTANT

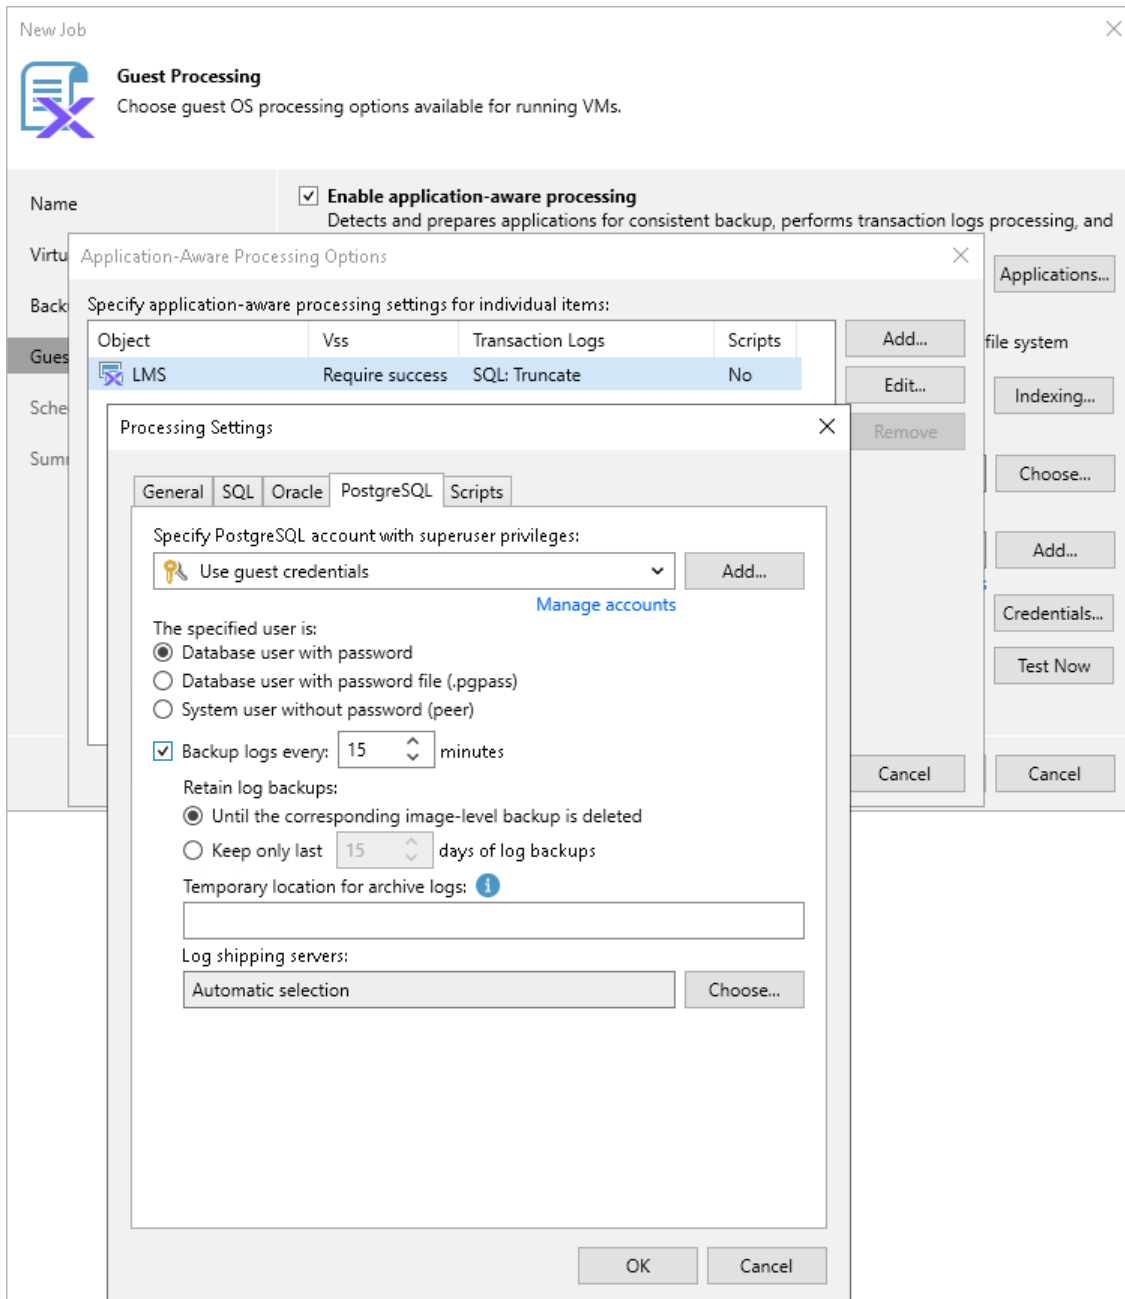
If you have added a new PostgreSQL account and want to use it with the peer authentication method, make sure that you have added this account as a Linux user with [sufficient permissions](#).

3. To back up PostgreSQL WAL files with Veeam Backup & Replication, select the **Backup log every <N> minutes** check box and specify the frequency for WAL files backup. By default, WAL files are backed up every 15 minutes. The maximum log backup interval is 480 minutes.
4. In the **Retain log backups** section, specify the retention policy for WAL files stored in the backup repository:
  - Select **Until the corresponding image-level backup is deleted** to apply the same retention policy for both image-level backups and WAL file backups.

- Select **Keep only last <n> days** to keep WAL files for a specific number of days. By default, WAL files are kept for 15 days. If you select this option, you must make sure that retention for WAL files is not greater than retention for the image-level backups. For more information, see the Veeam Backup & Replication User Guide, section [Retention for PostgreSQL WAL Files](#).
5. In the **Temporary location for archive logs (Staging location for archive logs - for version 12)** section, specify a path to the storage location where you want to keep WAL files.
  6. In the **Log shipping servers** section, click **Choose** to select what log shipping server you want to use to transport WAL files:
    - Select **Automatic selection** if you want Veeam Backup & Replication to select an optimal log shipping server automatically. The process of WAL files shipment does not require a dedicated server – Veeam Backup & Replication can use any Microsoft Windows server added to the backup infrastructure.

- Select **Use the specified servers only** to define a log shipping server explicitly. In the **Log Shipping Servers** window, select check boxes next to the servers you want to use as log shipping servers. The server list includes all Microsoft Windows servers added to the backup infrastructure.

Make sure that you select a server that is not used by other resource-consuming tasks. For example, you may want not to use a server that performs the WAN accelerator role as a log shipping server. For load balance and high availability purposes, it is recommended that you select at least 2 log shipping servers.



## Pre-Freeze and Post-Thaw Scripts

If you plan to back up VMs running applications that do not support VSS, you can specify what scripts Veeam Backup & Replication must use to quiesce the VM. The pre-freeze script quiesces the VM file system and application data to bring the VM to a consistent state before Veeam Backup & Replication triggers a VM snapshot. After the VM snapshot is created, the post-thaw script brings the VM and applications to their initial state.

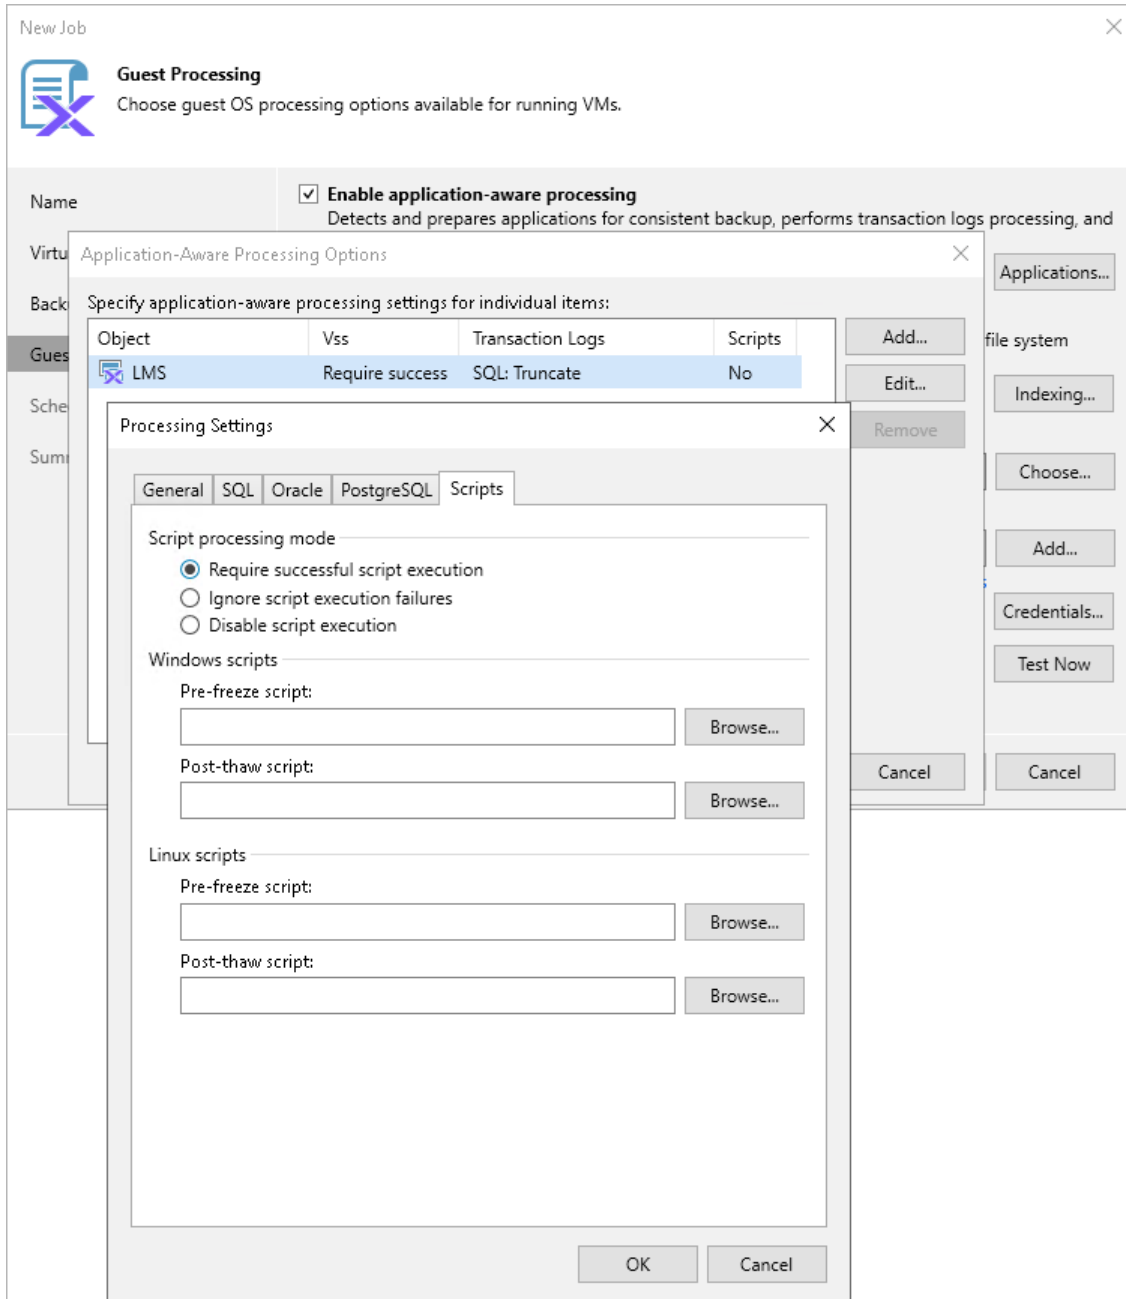
To specify pre-freeze and post-thaw scripts for the job:

1. Switch to the **Scripts** tab.
2. In the **Script processing mode** section, specify the scenario for scripts execution:
  - Select **Require successful script execution** if you want Veeam Backup & Replication to stop the backup process if the script fails.
  - Select **Ignore script execution failures** if you want to continue the backup process, even if script errors occur.
  - Select **Disable script execution** if you do not want to run scripts for the VM.
3. In the **Windows scripts** section, specify paths to pre-freeze and post-thaw scripts for Microsoft Windows VMs. For the list of supported script formats, see the Veeam Backup & Replication User Guide, section [Pre-Freeze and Post-Thaw Scripts](#).



- In the **Linux scripts** section, specify paths to pre-freeze and post-thaw scripts for Linux VMs. For the list of supported script formats, see the Veeam Backup & Replication User Guide, section [Pre-Freeze and Post-Thaw Scripts](#).

If you have added a protection domain, category, cluster or Prism Central with Microsoft Windows and Linux VMs to the job, you can select to execute both Microsoft Windows and Linux scripts for the VM container. When the job starts, Veeam Backup & Replication will automatically determine what OS type is installed on the VM and use the required scripts to quiesce this VM.



## Step 5b. Enable VM Guest OS File Indexing

To be able to recover individual files with 1 click and to search for specific items in Veeam Backup Enterprise Manager during [file-level restore](#), you must enable file indexing to instruct Veeam Backup & Replication to create a catalog of files and folders that belong to VMs included into the backup scope. To do that, select the **Enable guest file system indexing and malware detection** check box at the **Guest Processing** step of the wizard.

### NOTE

If you enable file indexing, Veeam Backup & Replication will scan VM data for suspicious file system activity and malware file presence every time the backup job completes successfully. For more information, see the Veeam Backup & Replication User Guide, section [How Guest Indexing Data Scan Works](#).

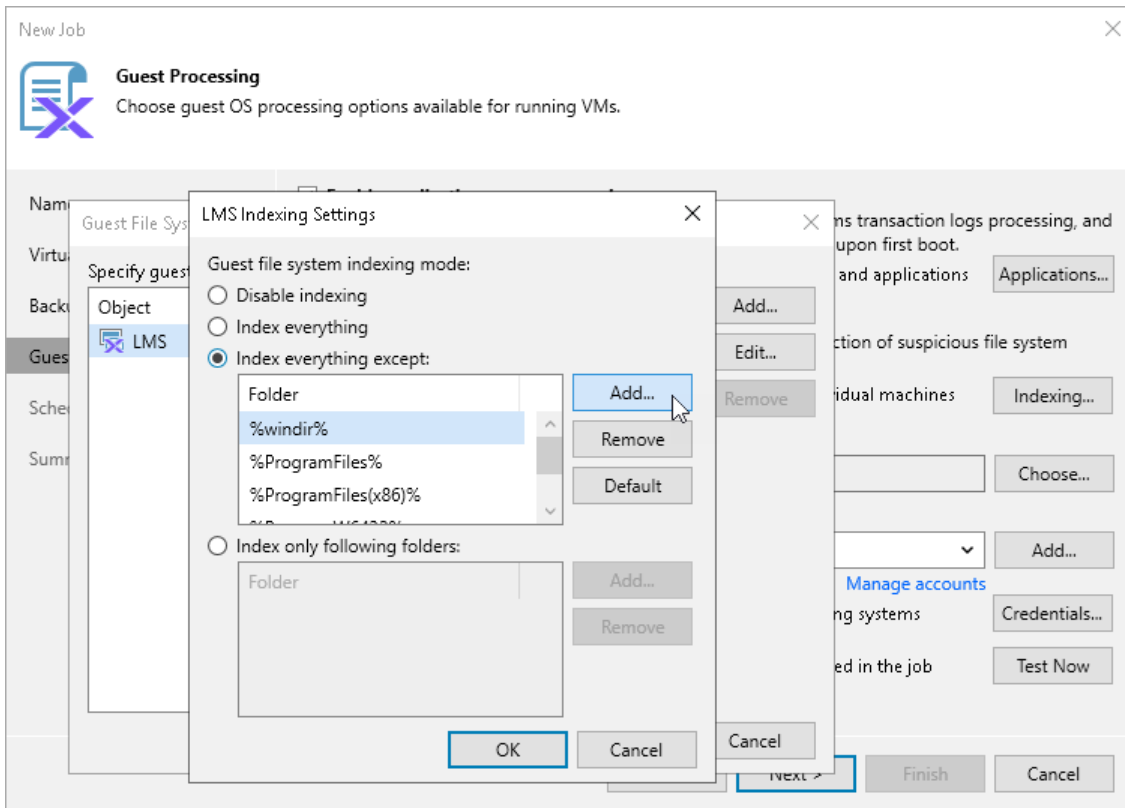
By default, Veeam Backup & Replication will create a catalog of all files and folders for each processed VM – except for system files. To change this behavior and configure indexing settings for specific VMs, do the following:

1. Click **Indexing**.
2. In the **Guest File System Indexing Options** window, select the necessary VMs and click **Edit > Windows indexing** or **Linux indexing**.  
You can configure indexing settings for one or more VMs at a time.
3. In the **Indexing Settings** window, choose whether you want to index files in all guest OS folders, to index files only in specific folders, or not to index any files at all.

If you select the **Index everything except** or **Index only following folders** option, you will be able to modify the list of folders included into the indexing scope – either manually or by using system environment variables (for example, `%windir%`, `%ProgramFiles%` and `%Temp%`).

### IMPORTANT

To allow Veeam Backup & Replication to perform guest OS file indexing for Linux VMs, `openssh`, `gzip` and `tar` utilities must be installed on the processed VMs.



## Step 5c. Choose Guest Interaction Proxy

[This step applies only if you have added Windows-based VMs to the backup scope at the **Virtual Machines** step of the wizard]

To produce transactionally consistent backups and to perform file system indexing, Veeam Backup & Replication communicates with the guest OS of each processed VM to deploy non-persistent runtime components that coordinate guest processing activities such as accessing VM applications and creating a catalog of VM files. Since these activities may significantly increase the load on the backup server in case of a large backup scope, Veeam Backup & Replication distributes the load among all Windows servers added to the backup infrastructure (further referred to as guest interaction proxies).

By default, Veeam Backup & Replication automatically chooses which guest interaction proxy to use for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section [Guest Interaction Proxies](#). You can also manually limit the list of servers that may be used as proxies – to do that, click **Choose**, select the **Prefer the following guest interaction proxy servers** option and then select check boxes next to the necessary Windows servers.

For a Windows server to be displayed in the list of available guest interaction proxy servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section [Adding Microsoft Windows Servers](#).

New Job

### Guest Processing

Choose guest OS processing options available for running VMs.

**Name**

**Virtual Machines**

**Backup Destination**

**Guest Processing**

**Schedule**

**Summary**

**Enable application-aware processing**  
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.  
Customize application handling options for individual machines and applications: **Applications...**

**Enable guest file system indexing and malware detection**  
Indexing enables global file search functionality, automatic detection of suspicious file system activity and known malware files.  
Customize advanced guest file system indexing options for individual machines: **Indexing...**

Guest interaction proxy:  
Automatic selection **Choose...**

**Guest Interaction Proxy**

Guest interaction proxies are used to offload guest processing from backup server. To add proxies, register one or more Windows servers on Backup Infrastructure tab.

Automatic selection  
Most suitable proxy will be selected among all registered Windows servers based on network configuration and current load.

Prefer the following guest interaction proxy servers:  
The job will automatically select most suitable proxy from the following list of selected Windows servers.

Name	Select
<input type="checkbox"/> ahvbackupsrv.tech.local	<input type="checkbox"/>
<input checked="" type="checkbox"/> ahvgpproxy.tech.local	<input checked="" type="checkbox"/>

**Select all** **Clear all**

**OK** **Cancel**

**Add...** **Credentials...** **Test Now** **Cancel**

## Step 5d. Manage VM Guest OS Credentials

If you enable application-aware processing or instruct Veeam Backup & Replication to create a catalog of VM files and folder, you must also specify a user whose credentials will be used to communicate with VM guest OSes. Note the specified account must have the permissions required to perform guest processing. For more information on the required permissions, see [Planning and Preparation](#).

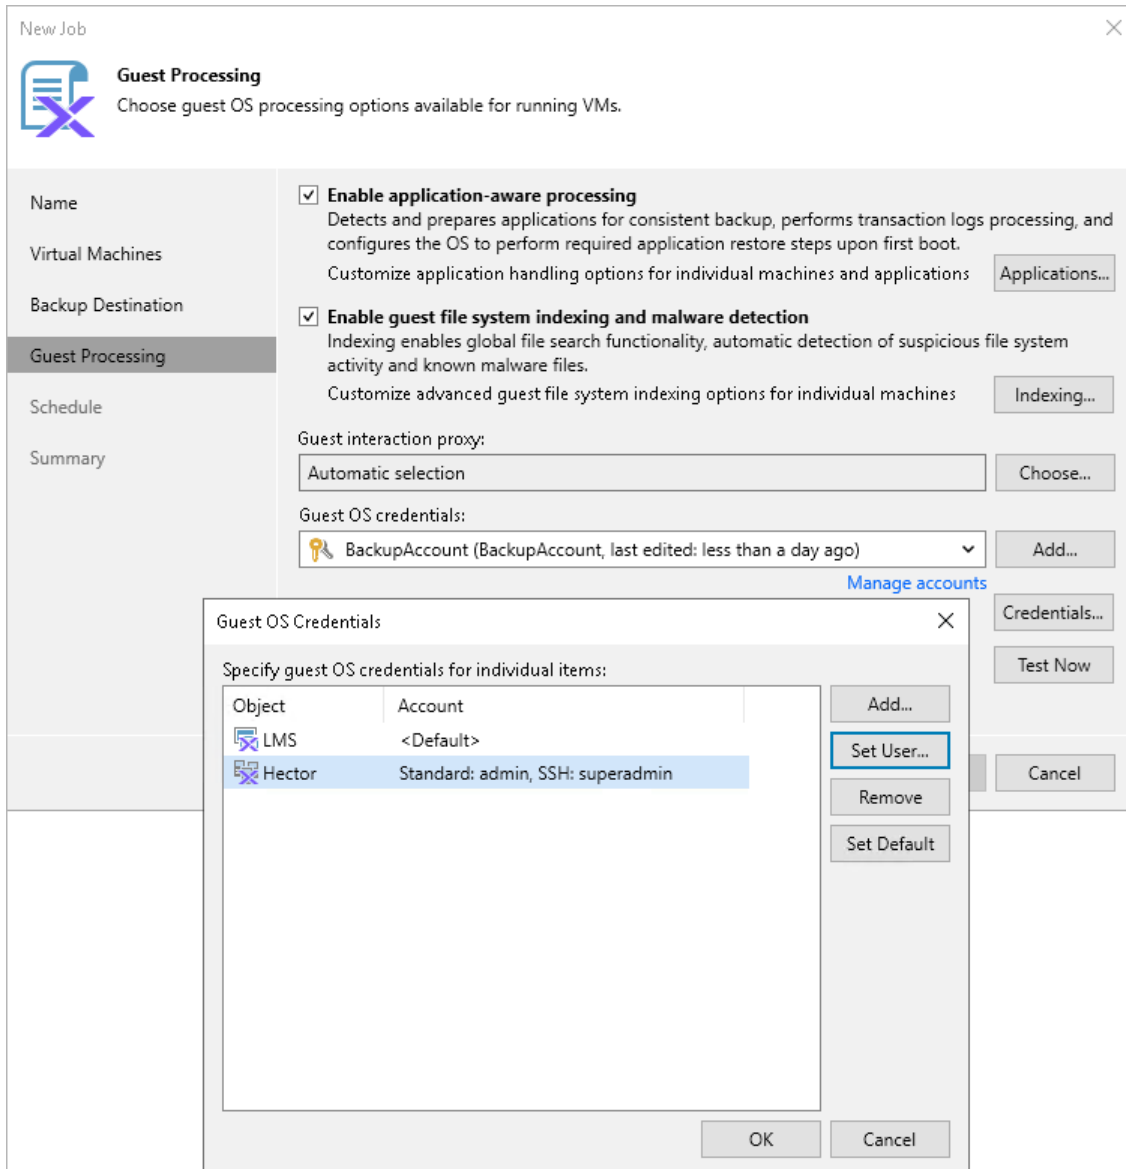
By default, Veeam Backup & Replication uses single credentials to access guest OSes of all VMs included into the backup scope. However, since Windows-based VMs and Linux-based VMs require different types of access credentials, you may need to specify the credentials explicitly for each of the processed VMs. To do that, click **Credentials**, select a VM in the **Guest OS credentials** window, and then click **Set User > Standard credentials** (for a Windows-based VM) or **Set User > SSH credentials** (for a Linux-based VM).

For a user to be displayed in the **Credentials** list, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Credentials Manager](#). If you have not added the necessary user to the Credentials Manager beforehand, you can do it without closing the **New Job** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

### TIP

If the backup scope includes a protection domain, cluster, category or Prism Central, you can specify both Standard and SSH credentials. This will allow Veeam Backup & Replication to access the processed VMs regardless of their guest OSes.

To check whether Veeam Backup & Replication is able to connect to the VM guest OSes using the specified credentials, click **Test Now**.



## Step 6. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the backup job automatically according to a specific backup schedule. The backup schedule defines how often data of the VMs added to the backup job will be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for Nutanix AHV allows you to create schedules of the following types:

- **Daily at this time** – the backup job will create restore points at a specific time on specific days.
- **Monthly at this time** – the backup job will create restore points once a month on a specific day.
- **Periodically every** – the backup job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the backup job again if it fails on the first try. To do that, select the **Retry failed items processing** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying backup jobs, Veeam Backup for Nutanix AHV processes only those VMs that failed to be backed up during the previous attempt.

**New Job** [Close]

**Schedule**  
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

**Name**

**Virtual Machines**

**Backup Destination**

**Schedule**

**Summary**

Run the job automatically

Daily at this time: 10:00 PM [Dropdown] Everyday [Dropdown] [Days...]

Monthly at this time: 10:00 PM [Dropdown] Fourth [Dropdown] Saturday [Dropdown] [Months...]

Periodically every: 1 [Dropdown] Hours [Dropdown]

Automatic retry

Retry failed items processing: 3 [Dropdown] times

Wait before each retry attempt for: 10 [Dropdown] minutes

< Previous [Apply] Finish Cancel

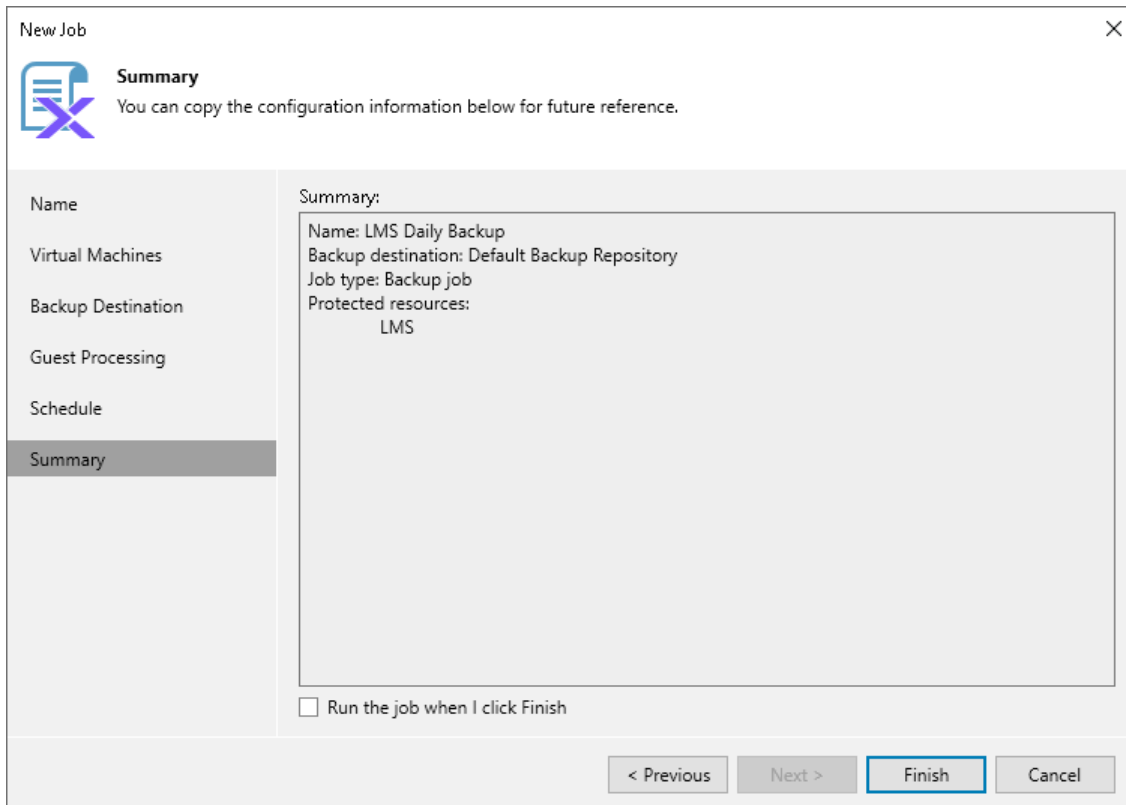


## Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

### TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.



The screenshot shows the 'New Job' wizard in the 'Summary' step. The window title is 'New Job' with a close button (X) in the top right corner. Below the title bar is a blue icon of a document with a checkmark and the word 'Summary'. A message reads: 'You can copy the configuration information below for future reference.' On the left side, there is a vertical navigation pane with the following items: 'Name', 'Virtual Machines', 'Backup Destination', 'Guest Processing', 'Schedule', and 'Summary' (which is highlighted). The main area contains a 'Summary:' box with the following text: 'Name: LMS Daily Backup', 'Backup destination: Default Backup Repository', 'Job type: Backup job', and 'Protected resources: LMS'. Below this box is a checkbox labeled 'Run the job when I click Finish', which is currently unchecked. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

# Creating Snapshot Jobs

To take snapshots of VMs, Veeam Backup for Nutanix AHV runs snapshot jobs. A snapshot job is a collection of settings that define the way data protection operations are performed: what data to protect, when to start the backup process, and so on. To create a snapshot job, you can either use the [Veeam Backup for Nutanix AHV web console](#) or the [Veeam Backup & Replication console](#).

# Before You Begin

Before you create a snapshot job, consider the following limitations:

- You cannot use snapshot jobs to create [PD snapshots](#). If you add a protection domain to a snapshot job, Veeam Backup for Nutanix AHV will create snapshots of individual VMs included into the protection domain.
- You cannot create snapshots of [instantly recovered VMs](#) since the VM disks have not completed migration to the Nutanix AHV cluster yet, which may cause the snapshots to become incomplete in case the [mount server](#) is disconnected.

# Creating Snapshot Jobs Using Backup Appliance Web Console

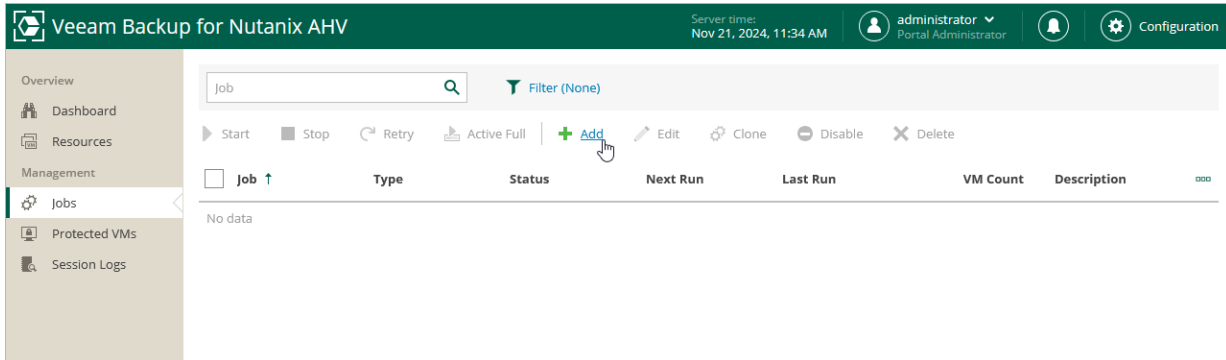
To create a snapshot job using the Nutanix AHV backup appliance web console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Specify a job name.](#)
4. [Selects virtual machines.](#)
5. [Configure snapshot retention settings.](#)
6. [Define a job schedule.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch Add Job Wizard

To launch the **Add Job** wizard, do the following:

1. Navigate to **Jobs**.
2. Click **Add**.



## Step 2. Specify Job Name and Description

At the **Job Info** step of the wizard, select the **Snapshot job** option. Then, use the **Name** and **Description** fields to specify a name for the new snapshot job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.

The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV. The interface is in a dark theme. At the top, the title bar reads 'Veeam Backup for Nutanix AHV' and includes the server time 'Aug 16, 2024, 05:38 PM', the user 'administrator Portal Administrator', and a 'Configuration' button. The main content area is titled 'Add Job' and has a left-hand navigation pane with options: 'Job Info' (selected), 'Sources', 'Backup Destination', 'Schedule', and 'Summary'. The 'Job Info' section contains the instruction 'Enter a name for the job and select a job type.' Below this are two text input fields: 'Name:' with the value 'Mail Server Daily Snapshot' and 'Description:' with the value 'Snapshot of the Mail Server'. At the bottom of the main area are three radio button options: 'Backup job' (unselected), 'Snapshot job' (selected), and 'PD snapshot job' (unselected). Each option has a brief description. At the bottom right of the wizard are two buttons: 'Next' and 'Cancel'.

## Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify the backup scope – choose whether you want to take snapshots of individual VMs, VMs included into protection domains or all VMs residing in the cluster. In the [Prism Central deployment](#), you can also take snapshots of VMs assigned to a specific category or all VMs managed by a Prism Central.

If you add a protection domain, Veeam Backup for Nutanix AHV will regularly check for new VMs added to the domain and automatically update the job settings to include these VMs in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

If you add a category, Veeam Backup for Nutanix AHV will regularly check for new VMs and clusters assigned to the category and automatically update the job settings to include these resources in the backup scope. For a category to be displayed in the list of the available categories, it must be configured in the Nutanix AHV Prism Central as described in [Nutanix documentation](#).

### TIP

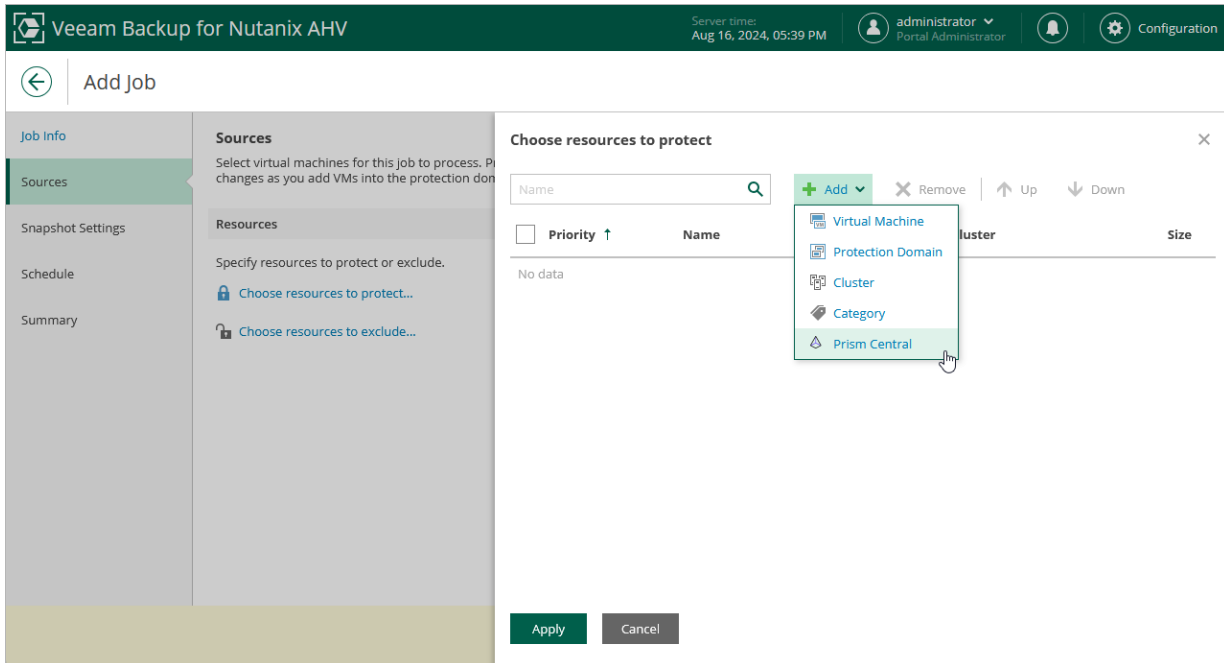
As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the VMs or protection domains that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

## NOTE

If you include a protection domain, category, cluster or Prism Central into the backup scope, VMs in this object are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs – not as a part of the protection domain, category, cluster or Prism Central.

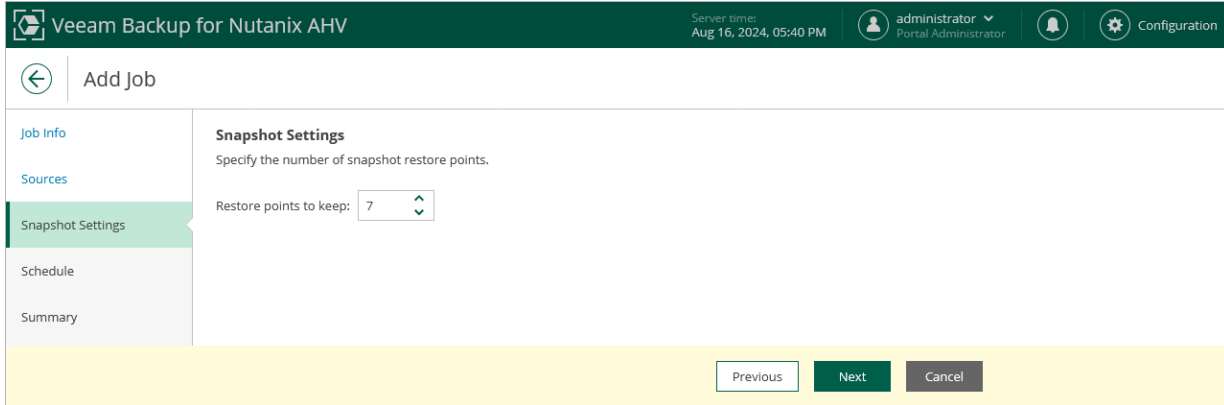




## Step 4. Configure Retention Settings

At the **Snapshot Settings** step of the wizard, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain. For more information, see [Snapshot Retention](#).



The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV. The interface is in a dark green header with the product name 'Veeam Backup for Nutanix AHV' on the left, and server time 'Aug 16, 2024, 05:40 PM', user 'administrator Portal Administrator', and 'Configuration' settings on the right. Below the header, there is a navigation bar with a back arrow and 'Add Job'. A sidebar on the left contains a list of steps: 'Job Info', 'Sources', 'Snapshot Settings' (highlighted in green), 'Schedule', and 'Summary'. The main content area is titled 'Snapshot Settings' and contains the instruction 'Specify the number of snapshot restore points.' followed by a label 'Restore points to keep:' and a numeric input field containing the value '7'. At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Next' (active), and 'Cancel' (disabled).

## Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the snapshot job automatically according to a specific job schedule. The schedule defines how often Veeam Backup for Nutanix AHV will create snapshots of the VMs added to the snapshot job.

To help you implement a comprehensive data protection strategy, Veeam Backup for Nutanix AHV allows you to create schedules of the following types:

- **Daily at this time** – the snapshot job will create restore points at a specific time on specific days.
- **Monthly at this time** – the snapshot job will create restore points once a month on a specific day.
- **Periodically every** – the snapshot job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the snapshot job again if it fails on the first try. To do that, select the **Automatic Retry** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying snapshot jobs, Veeam Backup for Nutanix AHV processes only those VMs whose snapshots were not taken during the previous attempt..

The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV. The 'Schedule' step is active, and the 'Configure main job schedule' dialog is open. The dialog has a close button (X) in the top right corner. Below the title, there is a note: 'Create a schedule to automatically start the job at the specified time. If you do not create a schedule, you will have to start the job manually.' There are three radio button options: 'Daily at this time' (selected), 'Monthly at this time', and 'Periodically every'. The 'Daily at this time' option has a time dropdown set to '10:00 PM' and a frequency dropdown set to 'Everyday'. The 'Monthly at this time' option has a time dropdown set to '10:00 PM', a frequency dropdown set to 'Fourth', a day dropdown set to 'Saturday', and a month list dropdown set to 'January, February, March, April, May, ...'. The 'Periodically every' option has a frequency dropdown set to '1' and a unit dropdown set to 'Hours'. Below these options, there is an 'Automatic retry' toggle switch set to 'On'. Underneath, there are two input fields: 'Retry failed job:' set to '3' times and 'Wait before each retry attempt for:' set to '10' minutes. At the bottom of the dialog, there are 'Apply' and 'Cancel' buttons.

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

### TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.

The screenshot shows the 'Add Job' wizard in the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name, the server time (Aug 16, 2024, 05:41 PM), and user information (administrator, Portal Administrator). The left sidebar contains navigation options: Job Info, Sources, Snapshot Settings, Schedule, and Summary (which is currently selected). The main content area is titled 'Summary' and contains the following information:

- Summary**: Review the job summary, and click Finish to save and apply the job settings. A 'Copy to clipboard' link is available.
- General**:
  - Name: Mail Server Daily Snapshot
  - Description: Snapshot of the Mail Server
- Protected resources**:
  - Virtual machines: 1 VM
- Backup destination**:
  - Retention: 7 Restore Points
- Run the job when I click Finish

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'. The 'Finish' button is highlighted in green.

# Creating Snapshot Jobs Using Veeam Backup & Replication Console

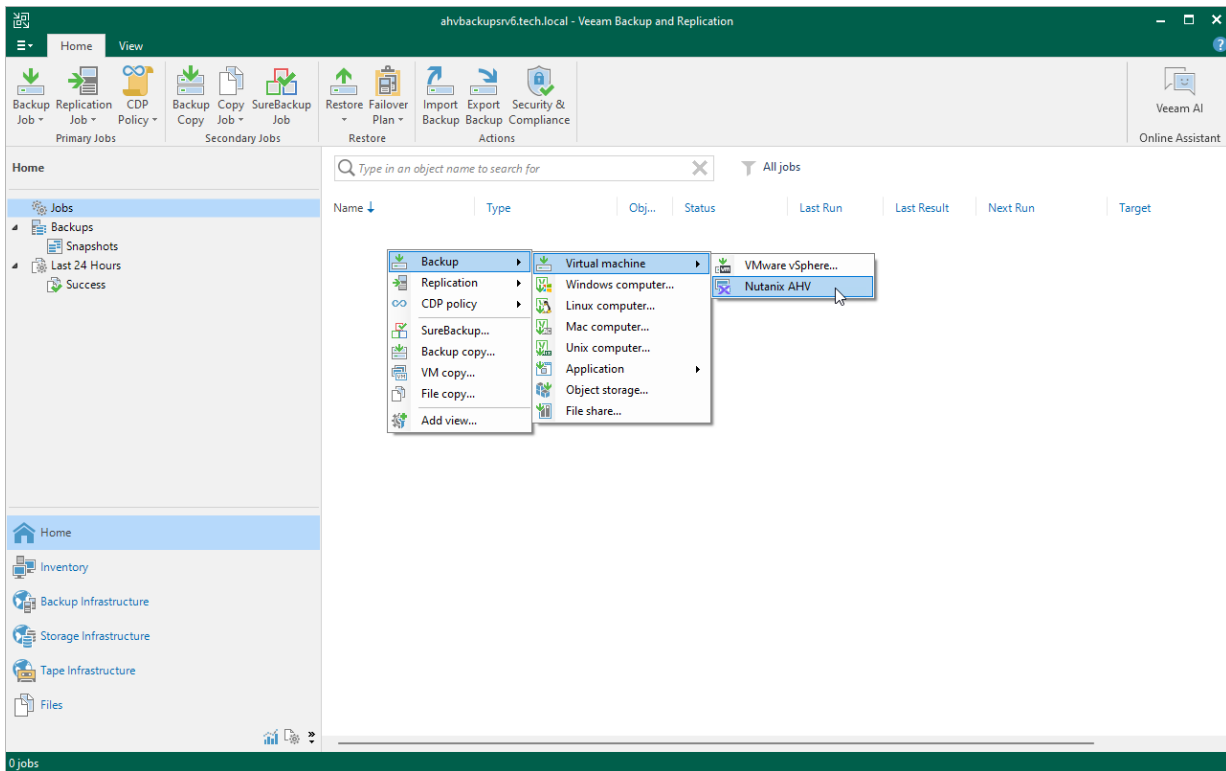
To create a snapshot job using the Veeam Backup & Replication console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Specify a job name.](#)
4. [Selects virtual machines.](#)
5. [Configure snapshot retention settings.](#)
6. [Define a job schedule.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch New Job Wizard

To launch the **New Job** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. On the ribbon, click **Backup Job > Virtual Machine > Nutanix AHV**, or right-click the working area and select **Backup > Virtual machine > Nutanix AHV**.



## Step 2. Specify Job Name and Description

At the **Name** step of the wizard, select the **Snapshot job** option. Then, use the **Name** and **Description** fields to specify a name for the new backup job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.

New Job

**Name**  
Type in a name and description for this job.

Name: Mail Server Daily Snapshot

Description: Snapshot of the Mail Server

Type:

- Backup job  
Configure a regular backup job to create backups of individual VMs.
- Snapshot job  
Configure a snapshot-only job to create snapshots of individual VMs.
- PD snapshot job  
Configure a protection domain snapshot-only job to create a snapshot of specific PD.

< Previous Next > Finish Cancel

## Step 3. Configure Backup Source Settings

At the **Virtual Machines** step of the wizard, specify the backup scope – choose whether you want to take snapshots of individual VMs, VMs included into protection domains or all VMs residing in the cluster. In the [Prism Central deployment](#), you can also take snapshots of VMs assigned to a specific category or all VMs managed by a Prism Central.

To view the list of available protection domains, click the **PDs** icon on the toolbar at the top right corner of the window. If you add a protection domain, Veeam Backup for Nutanix AHV will regularly check for new VMs added to the domain and automatically update the job settings to include these VMs in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

To view the list of available protection domains, click the **Categories** icon on the toolbar at the top right corner of the window. If you add a category, Veeam Backup for Nutanix AHV will regularly check for new VMs and clusters assigned to the category and automatically update the job settings to include these resources in the backup scope. For a category to be displayed in the list of the available categories, it must be configured in the Nutanix AHV Prism Central as described in [Nutanix documentation](#).

### TIP


As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Exclusions** and specify the VMs or protection domains that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

## NOTE

If you include a protection domain, category, cluster or Prism Central into the backup scope, VMs in this object are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs – not as a part of the protection domain, category, cluster or Prism Central.

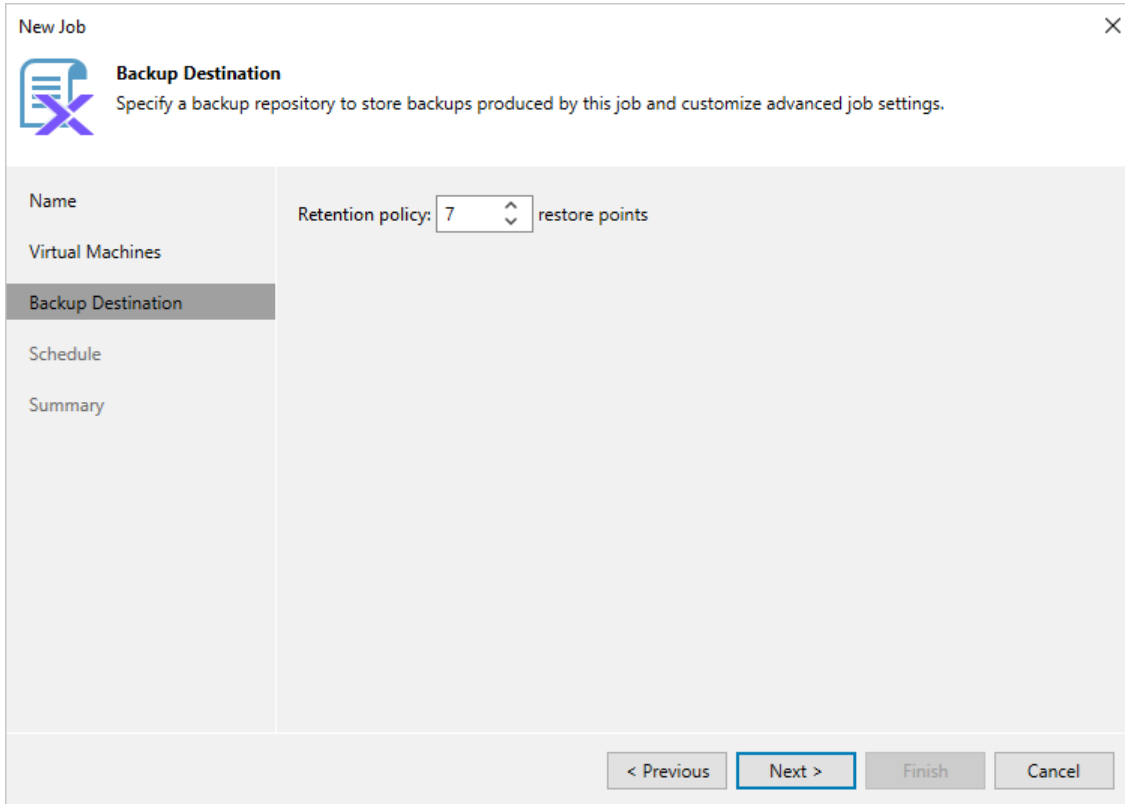




## Step 4. Configure Retention Settings

At the **Backup Destination** step of the wizard, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain. For more information, see [Snapshot Retention](#).



The screenshot shows a 'New Job' wizard window with a close button (X) in the top right corner. The title bar reads 'New Job'. Below the title bar is a blue icon of a document with a checkmark and the text 'Backup Destination' followed by the instruction 'Specify a backup repository to store backups produced by this job and customize advanced job settings.'.

The main area of the wizard is divided into two sections. On the left is a vertical navigation pane with the following items: 'Name', 'Virtual Machines', 'Backup Destination' (which is highlighted with a dark grey background), 'Schedule', and 'Summary'. On the right is the configuration area for the 'Backup Destination' step, which contains the text 'Retention policy: 7' followed by a small up/down arrow icon and the text 'restore points'.

At the bottom of the wizard are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

## Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the snapshot job automatically according to a specific job schedule. The schedule defines how often Veeam Backup for Nutanix AHV will create snapshots of the VMs added to the snapshot job.

To help you implement a comprehensive data protection strategy, Veeam Backup for Nutanix AHV allows you to create schedules of the following types:

- **Daily at this time** – the snapshot job will create restore points at a specific time on specific days.
- **Monthly at this time** – the snapshot job will create restore points once a month on a specific day.
- **Periodically every** – the snapshot job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the snapshot job again if it fails on the first try. To do that, select the **Automatic Retry** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying snapshot jobs, Veeam Backup for Nutanix AHV processes only those VMs whose snapshots were not taken during the previous attempt..

**New Job** [X]

**Schedule**  
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

**Name**  Run the job automatically

**Virtual Machines**  Daily at this time: 10:00 PM [v] Everyday [v] Days...

**Backup Destination**  Monthly at this time: 10:00 PM [v] Fourth [v] Saturday [v] Months...

**Schedule**  Periodically every: 1 [v] Hours [v]

**Summary** **Automatic retry**

Retry failed items processing: 3 [v] times

Wait before each retry attempt for: 10 [v] minutes

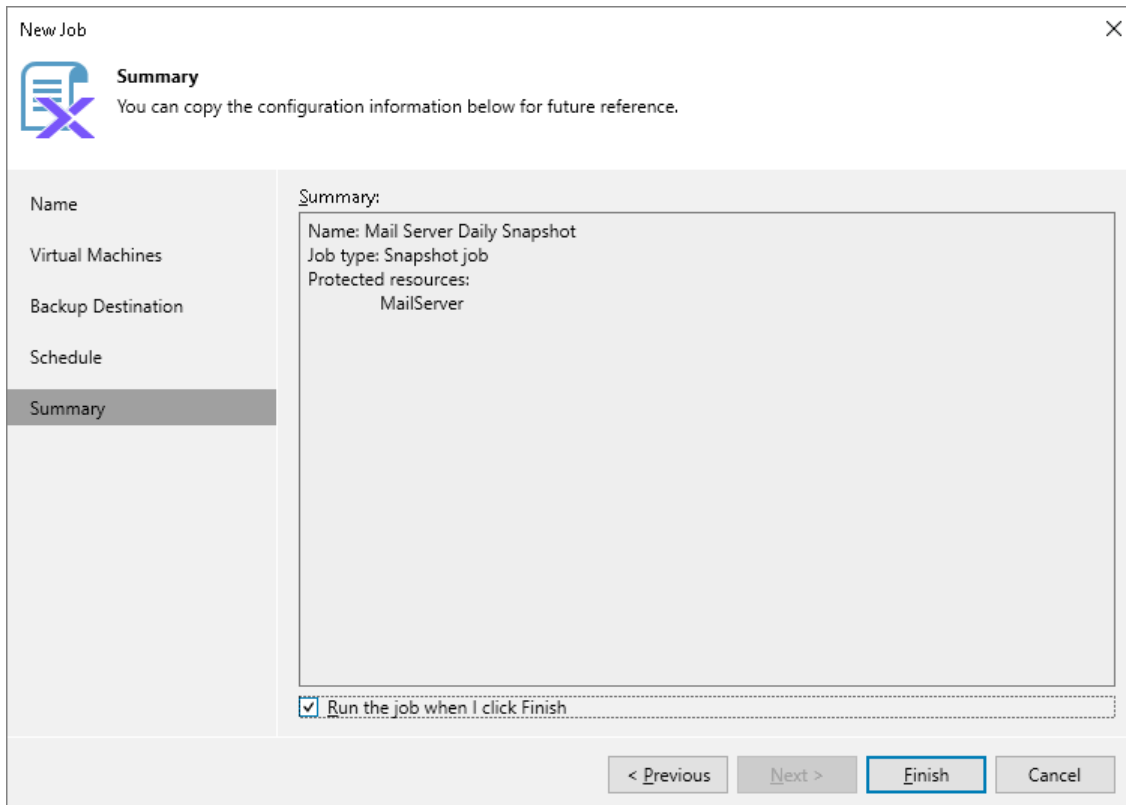
< Previous Apply Finish Cancel

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

### TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.



The screenshot shows the 'New Job' wizard in the 'Summary' step. The window title is 'New Job' with a close button (X) in the top right corner. On the left, there is a navigation pane with the following items: Name, Virtual Machines, Backup Destination, Schedule, and Summary (which is highlighted). The main area contains a 'Summary' section with a text box displaying the following information: 'Name: Mail Server Daily Snapshot', 'Job type: Snapshot job', and 'Protected resources: MailServer'. Below the text box, there is a checked checkbox labeled 'Run the job when I click Finish'. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

# Creating Protection Domain Snapshot Jobs

To take snapshots of protection domains, Veeam Backup for Nutanix AHV runs PD snapshot jobs. A PD snapshot job is a collection of settings that define the way data protection operations are performed: what data to protect, when to start the backup process, and so on. To create a PD snapshot job, you can either use the [Veeam Backup for Nutanix AHV web console](#) or the [Veeam Backup & Replication console](#).

# Before You Begin

Before you create a PD snapshot job, consider the following limitations:

- You cannot create snapshots of a protection domain if it contains multiple VMs with the same name.
- You cannot create snapshots of a protection domain if it contains VMs with attached volume groups that are not included into the protection domain.
- You cannot create snapshots of [instantly recovered VMs](#) since the VM disks have not completed migration to the Nutanix AHV cluster yet, which may cause the snapshots to become incomplete in case the [mount server](#) is disconnected.
- You cannot create snapshots of a protection domain in the *Inactive* state (for example, after having been replicated to a remote site).

# Creating PD Snapshot Jobs Using Backup Appliance Web Console

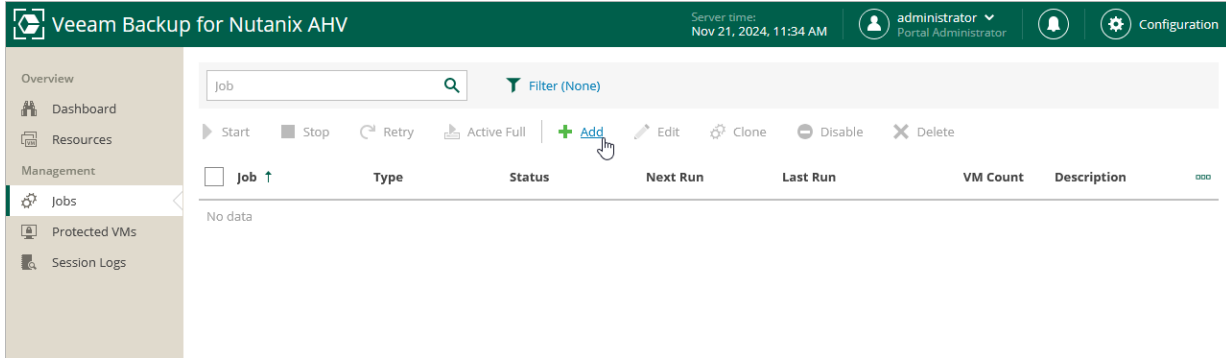
To create a protection domain snapshot job using the Nutanix AHV backup appliance web console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Specify a job name.](#)
4. [Selects protection domains.](#)
5. [Configure snapshot retention settings.](#)
6. [Define a job schedule.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch Add Job Wizard

To launch the **Add Job** wizard, do the following:

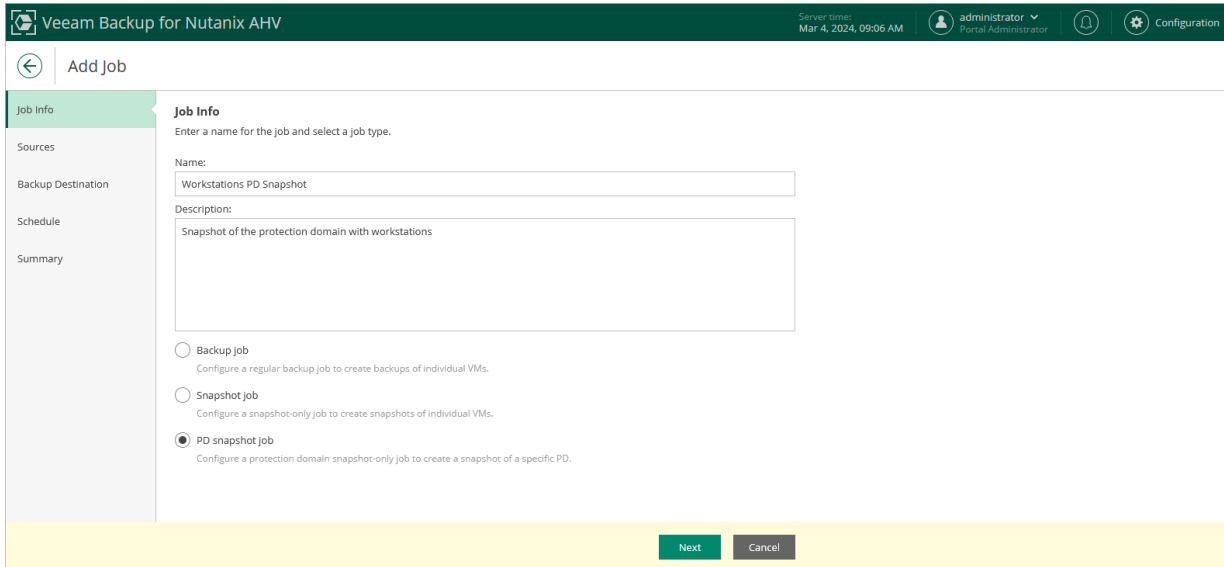
1. Navigate to **Jobs**.
2. Click **Add**.



## Step 2. Specify Job Name and Description

At the **Job Info** step of the wizard, select the **PD snapshot job** option. Then, use the **Name** and **Description** fields to specify a name for the new PD snapshot job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.



The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV. The interface is in a dark theme. At the top, the title bar reads 'Veeam Backup for Nutanix AHV' and includes the server time 'Mar 4, 2024, 09:06 AM', the user 'administrator Portal Administrator', and a 'Configuration' button. The main window has a left sidebar with navigation options: 'Job Info' (selected), 'Sources', 'Backup Destination', 'Schedule', and 'Summary'. The 'Job Info' section is titled 'Job Info' and contains the instruction 'Enter a name for the job and select a job type.' Below this are two text input fields: 'Name' with the value 'Workstations PD Snapshot' and 'Description' with the value 'Snapshot of the protection domain with workstations'. At the bottom of the 'Job Info' section are three radio button options: 'Backup job' (unselected), 'Snapshot job' (unselected), and 'PD snapshot job' (selected). Each option has a brief description. At the bottom of the wizard are 'Next' and 'Cancel' buttons.



## Step 3. Select Protection Domains

At the **Sources** step of the wizard, specify the backup scope – choose whether you want to take snapshots of individual protection domains or all protection domains configured in the cluster or Prism Central.

Veeam Backup for Nutanix AHV will regularly check for new consistency groups (VMs and volume groups) added to the domain and automatically update the job settings to include them in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

### TIP


As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify protection domains that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

### NOTE

If you include a cluster or Prism Central into the backup scope, PDs in this cluster are processed at random. To ensure that the PDs are processed in a specific order, you must add them as standalone PDs – not as a part of the cluster or Prism Central.

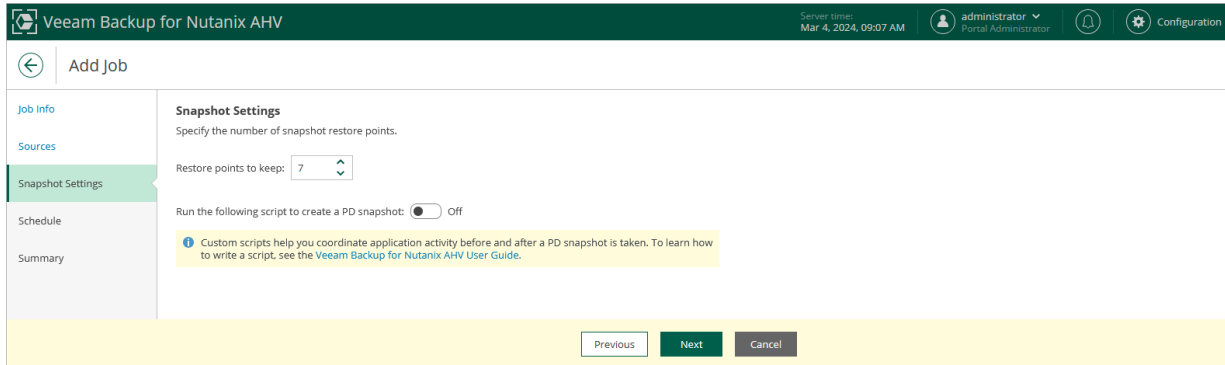


## Step 4. Configure Retention Settings

At the **Snapshot Settings** step of the wizard, specify the number of restore points that you want to keep in a snapshot chain. If the restore point limit is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain. For more information, see [Snapshot Retention](#).

### TIP

To instruct Veeam Backup for Nutanix AHV to execute a custom script while running the PD snapshot job, set the **Run the following script to create a PD snapshot** toggle to *On*, and click **Browse** to upload a script file. For more information, see [Writing Custom Scripts](#).



## Writing Custom Scripts

To coordinate application activity on VMs included into a protection domain before and after a PD snapshot is taken, you can run custom scripts. A custom script performs pre-freeze operations, takes the PD snapshot, and finally performs post-thaw operations.

### IMPORTANT

Veeam does not provide support for cases related to custom scripts.

You can use any programming language to write a script. However, you must specify the used interpreter in the shebang line, for example: `#!/usr/bin/python3`. For more information, see [Appendix A. Custom Script Samples](#).

### NOTE

You cannot use binary files for custom scripts.

While writing custom scripts, use the following arguments:

Argument	Description
--clusterId	[Applies only to <a href="#">Prism Central deployment</a> ] ID of a cluster where the protection domain is configured.
--pdName	Name of a protection domain for which a PD snapshot is created.
--jobName	Name of a job that runs the script.

Argument	Description
--logDir	Path to a folder where script logs are stored.
--logLevel	Level of logging.

To specify parameters required for connecting to the Nutanix AHV cluster where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixClusterIp = os.getenv('NUTANIX_CLUSTER_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

To specify parameters required for connecting to the Prism Central where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixPrismCentralIp = os.getenv('NUTANIX_PRISM_CENTRAL_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

You can find script logs in the `custom_script.log` file stored in the `/var/log/nxbackupagent/Backup/<job_name_job_uuid>/<pd_name>/` folder on the backup appliance. The file also contains a JSON string that includes the script execution status, an error description (if any) and an identifier of the PD snapshot if it has been created, for example:

- Script execution succeeded:

```
Result: {"status": "Success", "errorMessage": "", "oob_schedule_id": 7885270}
```

- Script execution failed:

```
Result: {"status": "Failed", "errorMessage": "Unable to get vm_id list for Protection Domain='VeeamBackupProtection': Authentication failed.", "oob_schedule_id": ""}
```

- Script execution completed with a warning:

```
Result: {"status": "Warning", "errorMessage": "Pre-freeze took more time than expected", "oob_schedule_id": "44509"}
```

## TIP

You can track the script execution progress on the [Session Logs page](#) in the backup appliance web console. Alternatively, you can [download log files](#) to your local machine.

## Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the PD snapshot job automatically according to a specific schedule. The schedule defines how often Veeam Backup for Nutanix AHV will create snapshots of the protection domains added to the PD snapshot job.

To help you implement a comprehensive data protection strategy, Veeam Backup for Nutanix AHV allows you to create schedules of the following types:

- **Daily at this time** – the PD snapshot job will create restore points at a specific time on specific days.
- **Monthly at this time** – the PD snapshot job will create restore points once a month on a specific day.
- **Periodically every** – the PD snapshot job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the job again if it fails on the first try. To do that, select the **Automatic Retry** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying PD snapshot jobs, Veeam Backup for Nutanix AHV processes only those PDs whose snapshots were not taken during the previous attempt.

The screenshot shows the 'Configure main job schedule' dialog box in the Veeam Backup for Nutanix AHV interface. The dialog is titled 'Configure main job schedule' and includes a close button (X) in the top right corner. Below the title is a subtitle: 'Create a schedule to automatically start the job at the specified time. If you do not create a schedule, you will have to start the job manually.'

The dialog contains three radio button options for scheduling:

- Daily at this time:** Selected. Includes a time dropdown set to '10:00 PM' and a frequency dropdown set to 'Everyday'.
- Monthly at this time:** Includes a time dropdown set to '10:00 PM', a frequency dropdown set to 'Fourth', a day dropdown set to 'Saturday', and a month dropdown set to 'January, February, March, April, May, ...'.
- Periodically every:** Includes a frequency dropdown set to '1' and a unit dropdown set to 'Hour'.

Below the scheduling options are the following settings:

- Automatic retry:** A toggle switch is turned 'On'.
- Retry failed job:** A numeric spinner set to '3' with the unit 'times'.
- Wait before each retry attempt for:** A numeric spinner set to '10' with the unit 'minutes'.

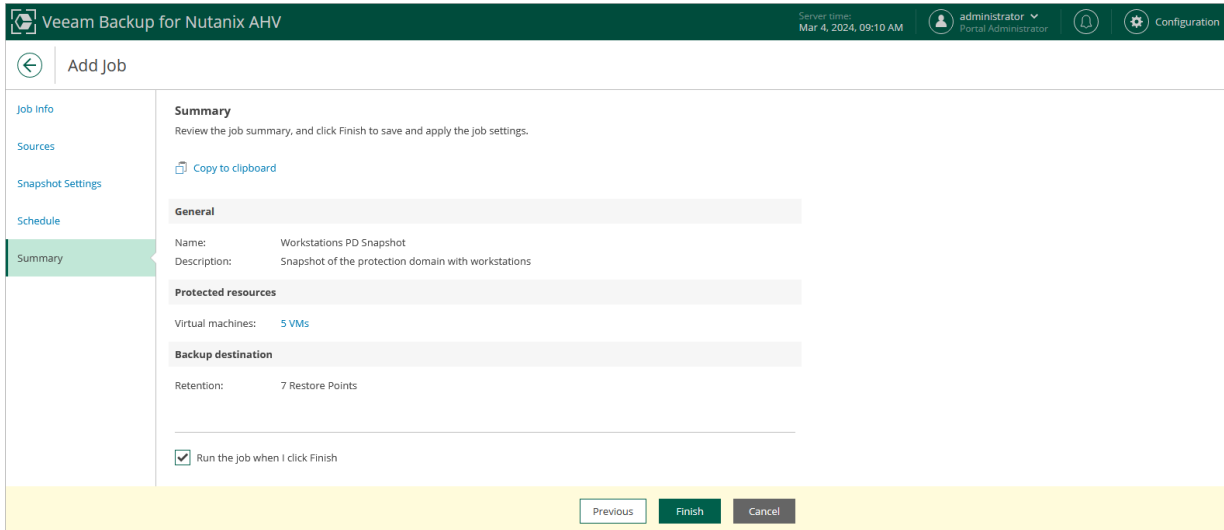
At the bottom of the dialog are two buttons: 'Apply' and 'Cancel'.

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

### TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.



The screenshot shows the 'Add Job' wizard in Veeam Backup for Nutanix AHV, specifically the 'Summary' step. The interface includes a top navigation bar with the product name, server status, user information, and configuration options. A left sidebar lists the wizard steps: Job Info, Sources, Snapshot Settings, Schedule, and Summary (which is currently selected). The main content area displays the job summary, including a 'Copy to clipboard' button, a 'General' section with job name and description, a 'Protected resources' section showing 5 VMs, and a 'Backup destination' section showing 7 Restore Points. At the bottom, there is a checkbox for 'Run the job when I click Finish' which is checked, and three buttons: 'Previous', 'Finish', and 'Cancel'.

Veeam Backup for Nutanix AHV

Server status: Mar 4, 2024, 09:10 AM

administrator Portal Administrator

Configuration

← Add Job

Job Info

Sources

Snapshot Settings

Schedule

Summary

### Summary

Review the job summary, and click Finish to save and apply the job settings.

#### General

Name: Workstations PD Snapshot

Description: Snapshot of the protection domain with workstations

#### Protected resources

Virtual machines: 5 VMs

#### Backup destination

Retention: 7 Restore Points

Run the job when I click Finish

# Creating PD Snapshot Jobs Using Veeam Backup & Replication Console

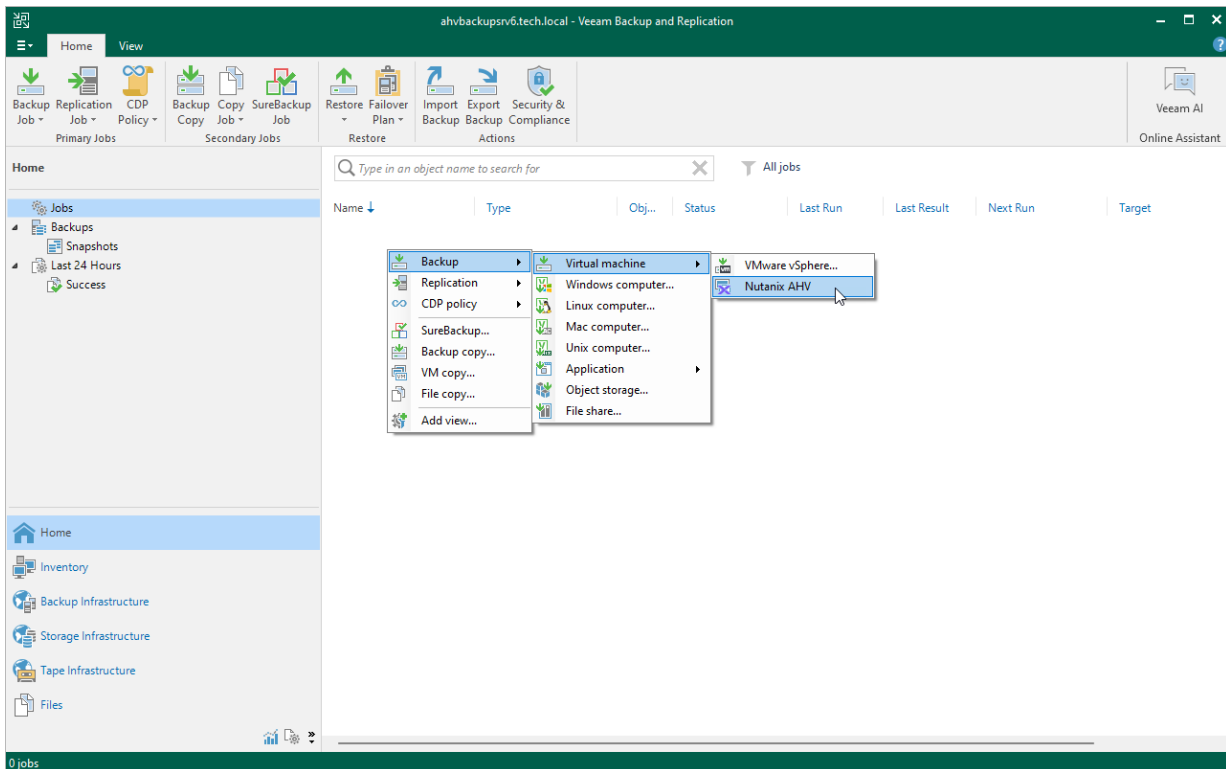
To create a protection domain snapshot job using the Veeam Backup & Replication console, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Add Job wizard.](#)
3. [Specify a job name.](#)
4. [Selects protection domains.](#)
5. [Configure snapshot retention settings.](#)
6. [Define a job schedule.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch New Job Wizard

To launch the **New Job** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. On the ribbon, click **Backup Job > Virtual Machine > Nutanix AHV**, or right-click the working area and select **Backup > Virtual machine > Nutanix AHV**.



## Step 2. Specify Job Name and Description

At the **Name** step of the wizard, select the **PD snapshot job** option. Then, use the **Name** and **Description** fields to specify a name for the new backup job and to provide a description for future reference. The job name must be unique in Veeam Backup for Nutanix AHV.

The maximum length of the name is 40 characters; the following characters are not supported: \ / " ' [ ] : | < > + = ; , ? \* @ & \_ . The maximum length of the description is 1024 characters.

New Job

**Name**  
Type in a name and description for this job.

Name: Workstations PD Snapshot

Description: Snapshot of the protection domain with workstations

Type:

- Backup job  
Configure a regular backup job to create backups of individual VMs.
- Snapshot job  
Configure a snapshot-only job to create snapshots of individual VMs.
- PD snapshot job  
Configure a protection domain snapshot-only job to create a snapshot of specific PD.

< Previous   Next >   Finish   Cancel



## Step 3. Configure Backup Source Settings

At the **Virtual Machines** step of the wizard, specify the backup scope – choose whether you want to take snapshots of individual protection domains or all protection domains configured in the cluster or Prism Central.

Veeam Backup for Nutanix AHV will regularly check for new consistency groups (VMs and volume groups) added to the domain and automatically update the job settings to include them in the backup scope. For a protection domain to be displayed in the list of the available domains, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

### TIP

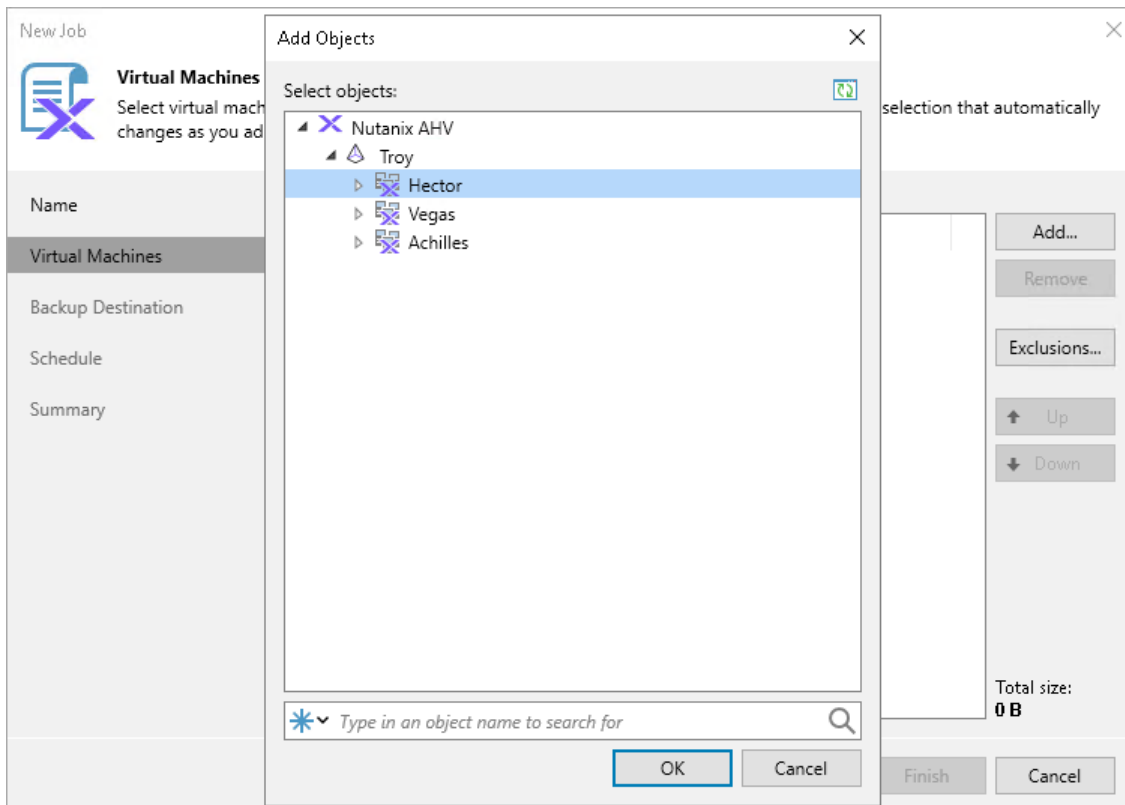
As an alternative to specifying resources explicitly, you can exclude a number of resources from the backup scope. To do that, click **Exclusions** and specify the protection domains that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Nutanix AHV will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup for Nutanix AHV processes resources in the order they are added to the backup scope. However, you can change the order, for example, if you add some mission-critical VMs to the job and want them to be processed first. To change the processing order, select a resource and use the **Up** or **Down** buttons.

### NOTE

If you include a cluster or Prism Central into the backup scope, PDs in this cluster are processed at random. To ensure that the PDs are processed in a specific order, you must add them as standalone PDs – not as a part of the cluster or Prism Central.

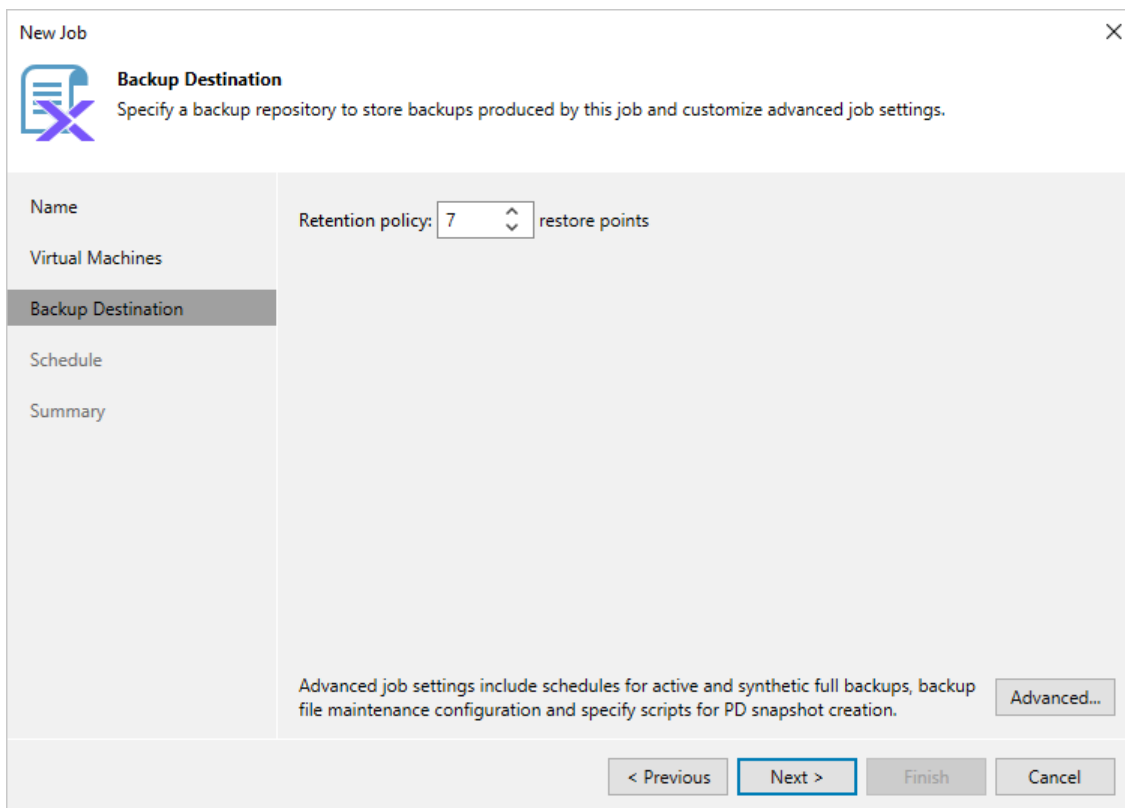


## Step 4. Configure Retention Settings

At the **Backup Destination** step of the wizard, specify the number of restore points that you want to keep in a snapshot chain. If the restore point limit is exceeded, Veeam Backup for Nutanix AHV removes the earliest restore point from the chain. For more information, see [Snapshot Retention](#).

### TIP

To instruct Veeam Backup for Nutanix AHV to execute a custom script while running the PD snapshot job, click **Advanced**, then, in the **Advanced settings** window, select the **Run the following script to create a PD snapshot** check box and click **Browse** to upload a script file. For more information, see [Writing Custom Scripts](#).



The screenshot shows the 'New Job' wizard window with the 'Backup Destination' step selected. The 'Retention policy' is set to 7 restore points. The 'Advanced...' button is visible at the bottom right of the main content area.

Step	Retention policy
Backup Destination	7 restore points

## Writing Custom Scripts

To coordinate application activity on VMs included into a protection domain before and after a PD snapshot is taken, you can run custom scripts. A custom script performs pre-freeze operations, takes the PD snapshot, and finally performs post-thaw operations.

### IMPORTANT

Veeam does not provide support for cases related to custom scripts.

You can use any programming language to write a script. However, you must specify the used interpreter in the shebang line, for example: `#!/usr/bin/python3`. For more information, see [Appendix A. Custom Script Samples](#).

## NOTE

You cannot use binary files for custom scripts.

While writing custom scripts, use the following arguments:

Argument	Description
--clusterId	[Applies only to <a href="#">Prism Central deployment</a> ] ID of a cluster where the protection domain is configured.
--pdName	Name of a protection domain for which a PD snapshot is created.
--jobName	Name of a job that runs the script.
--logDir	Path to a folder where script logs are stored.
--logLevel	Level of logging.

To specify parameters required for connecting to the Nutanix AHV cluster where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixClusterIp = os.getenv('NUTANIX_CLUSTER_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

To specify parameters required for connecting to the Prism Central where VMs included into the protection domain are running, use environment variables such as in the following example:

```
nutanixPrismCentralIp = os.getenv('NUTANIX_PRISM_CENTRAL_ADDRESS')
nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
```

You can find script logs in the `custom_script.log` file stored in the `/var/log/nxbackupagent/Backup/<job_name_job_uuid>/<pd_name>/` folder on the backup appliance. The file also contains a JSON string that includes the script execution status, an error description (if any) and an identifier of the PD snapshot if it has been created, for example:

- Script execution succeeded:

```
Result: {"status": "Success", "errorMessage": "", "oob_schedule_id": 7885270}
```

- Script execution failed:

```
Result: {"status": "Failed", "errorMessage": "Unable to get vm_id list for Protection Domain='VeeamBackupProtection': Authentication failed.", "oob_schedule_id": ""}
```

- Script execution completed with a warning:

```
Result: {"status": "Warning", "errorMessage": "Pre-freeze took more time than expected", "oob_schedule_id": "44509"}
```

#### TIP

You can track the script execution progress on the [Session Logs page](#) in the backup appliance web console. Alternatively, you can [download log files](#) to your local machine.

## Step 5. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Nutanix AHV to start the PD snapshot job automatically according to a specific schedule. The schedule defines how often Veeam Backup for Nutanix AHV will create snapshots of the protection domains added to the PD snapshot job.

To help you implement a comprehensive data protection strategy, Veeam Backup for Nutanix AHV allows you to create schedules of the following types:

- **Daily at this time** – the PD snapshot job will create restore points at a specific time on specific days.
- **Monthly at this time** – the PD snapshot job will create restore points once a month on a specific day.
- **Periodically every** – the PD snapshot job will create restore points repeatedly, with a specific time interval every day.

### TIP

You can instruct Veeam Backup for Nutanix AHV to run the job again if it fails on the first try. To do that, select the **Automatic Retry** check box, and specify the maximum number of attempts to run the job and the time interval between retries. When retrying PD snapshot jobs, Veeam Backup for Nutanix AHV processes only those PDs whose snapshots were not taken during the previous attempt.

**New Job** [Close]

**Schedule**  
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

**Name**  Run the job automatically

**Virtual Machines**  Daily at this time: 10:00 PM [Dropdown] Everyday [Dropdown] [Days...]

**Backup Destination**  Monthly at this time: 10:00 PM [Dropdown] Fourth [Dropdown] Saturday [Dropdown] [Months...]

**Schedule**  Periodically every: 1 [Dropdown] Hours [Dropdown]

**Summary** **Automatic retry**

Retry failed items processing: 3 [Dropdown] times

Wait before each retry attempt for: 10 [Dropdown] minutes

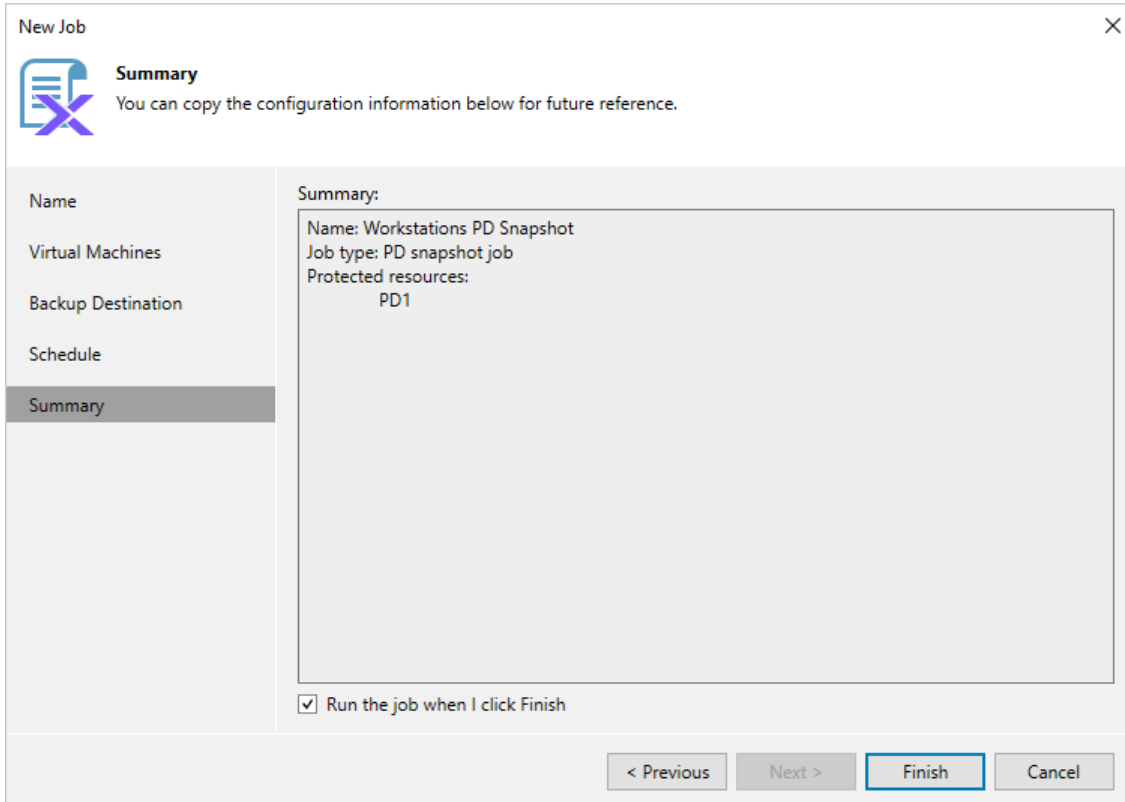
< Previous [Apply] Finish Cancel

## Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. When Veeam Backup for Nutanix AHV starts the job according to the specified schedule, the backup progress will be displayed on the [Session Logs page](#).

### TIP

If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.




The screenshot shows the 'New Job' wizard at the 'Summary' step. The window title is 'New Job' with a close button (X) in the top right corner. On the left, there is a navigation pane with the following items: Name, Virtual Machines, Backup Destination, Schedule, and Summary (which is highlighted). The main area contains a 'Summary' section with a text box displaying the following information: 'Name: Workstations PD Snapshot', 'Job type: PD snapshot job', and 'Protected resources: PD1'. Below the text box is a checked checkbox labeled 'Run the job when I click Finish'. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

# Analyzing Performance Bottlenecks

As any backup application handles a great amount of data, it is important to make sure the data flow is efficient and all resources engaged in the backup process are optimally used. For backup jobs, Veeam provides advanced statistics about the data flow efficiency and lets you identify bottlenecks at the following stages of the data transmission process:

1. Reading VM data blocks from the source.
2. Processing VM data on a worker.
3. Transporting data over the network.
4. Writing data to the target.



While evaluating the data transmission process, Veeam Backup for Nutanix AHV leverages the Veeam Backup & Replication functionality to analyze performance of all the data flow components:

- **Source** – the source disk reader component responsible for retrieving data from the source node.
- **Proxy** – the worker component responsible for processing VM data.
- **Network** – the network queue writer component responsible for getting processed VM data from the worker and sending it over the network to the Target (directly or through the Gateway Server).
- **Target** – the gateway server component responsible for processing VM data, or the target disk writer component responsible for storing data in the backup repository.

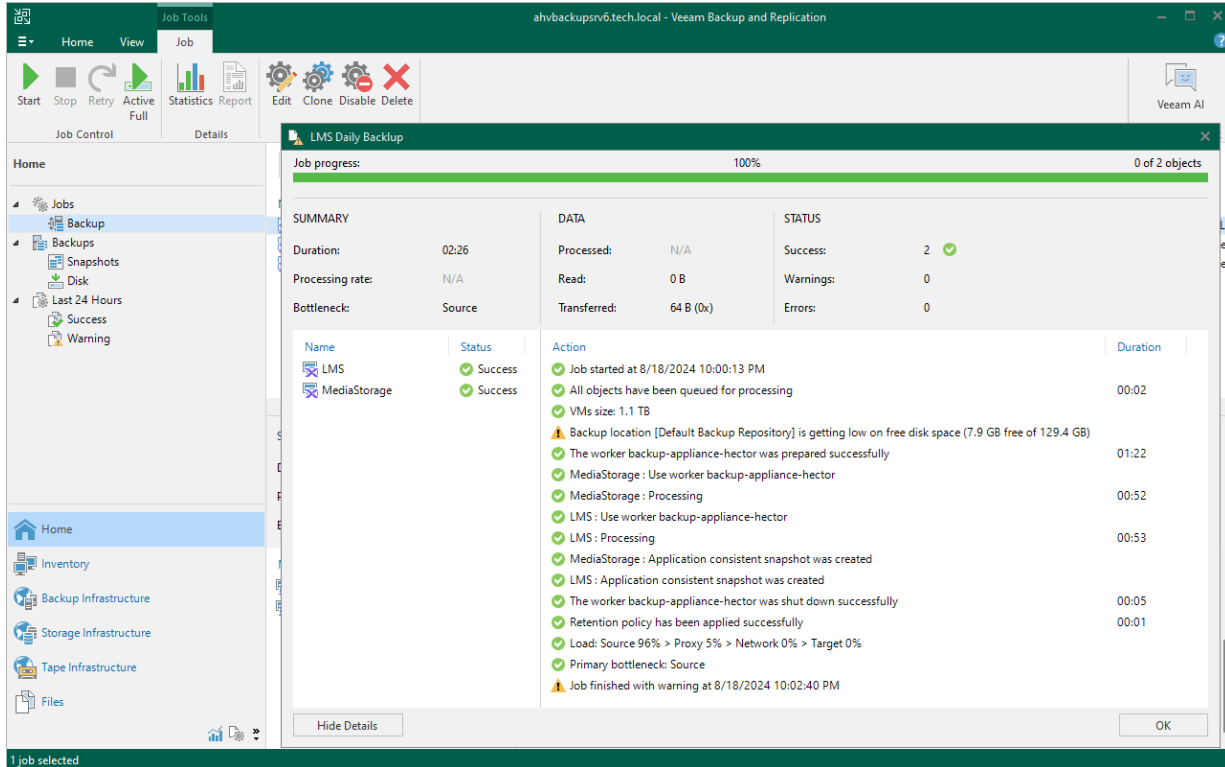
To see the bottleneck statistics for a job or a specific VM processed by the job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, right-click a backup job for which you want to see the bottleneck statistics, and select **Statistics**.

4. In the job session window, click **Show Details**:

- To see the aggregated statistics for the whole job, check the **Load** field in the **Action** column.
- To see the bottleneck statistics for a specific VM, click a VM name and check the **Load** field in the **Action** column.

To learn how to analyze the bottleneck statistics, see Veeam Backup & Replication User Guide, section [Performance Bottlenecks](#).





# Cloning Jobs

You can create a new job by cloning an existing one. Job cloning allows you to create an exact copy of any job with the same job settings.

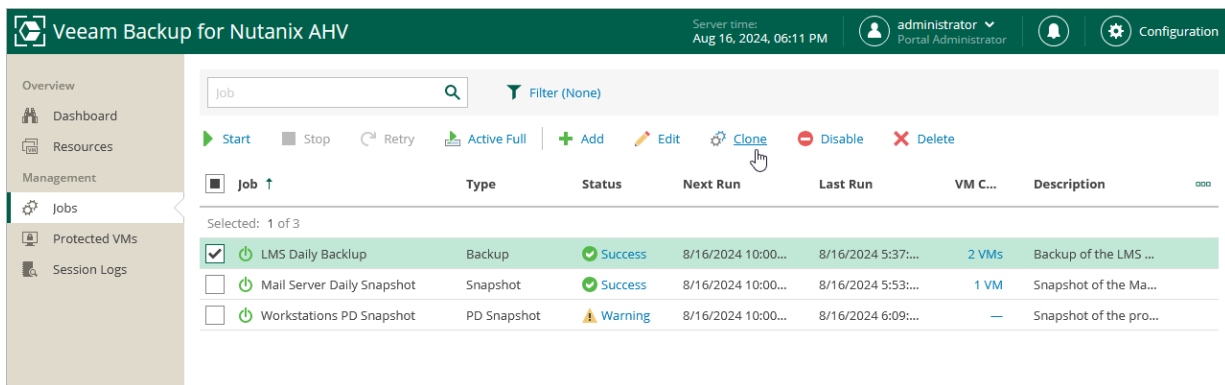
To clone a job, do the following:

1. Navigate to **Jobs**.
2. Select the job.
3. Click **Clone**.

The name of the cloned job is formed by the following rule: *<job\_name\_clone1>*, where *job\_name* is the name of the original job and *clone1* is a suffix added to the original job name. If you clone the same job again, the number in the name will be incremented, for example, *job\_name\_clone2*, *job\_name\_clone3* and so on. To change the name of a cloned job, edit the job as described in section [Editing Job Settings](#).

## NOTE

If the original job is scheduled to run automatically, Veeam Backup for Nutanix AHV disables the cloned job. To enable the cloned job, select it in the job list and click **Enable**.



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Nutanix AHV', the server time 'Aug 16, 2024, 06:11 PM', and the user 'administrator Portal Administrator'. The left sidebar shows navigation options: Overview, Dashboard, Resources, Management, Jobs, Protected VMs, and Session Logs. The main area displays a table of jobs with columns for Job, Type, Status, Next Run, Last Run, VM C..., and Description. A ribbon at the top of the table includes buttons for Start, Stop, Retry, Active Full, Add, Edit, Clone, Disable, and Delete. The 'Clone' button is highlighted with a mouse cursor. The table shows three jobs: 'LMS Daily Backup' (Backup, Success, 8/16/2024 10:00..., 8/16/2024 5:37:..., 2 VMs, Backup of the LMS ...), 'Mail Server Daily Snapshot' (Snapshot, Success, 8/16/2024 10:00..., 8/16/2024 5:53:..., 1 VM, Snapshot of the Ma...), and 'Workstations PD Snapshot' (PD Snapshot, Warning, 8/16/2024 10:00..., 8/16/2024 6:09:..., —, Snapshot of the pro...).

## TIP

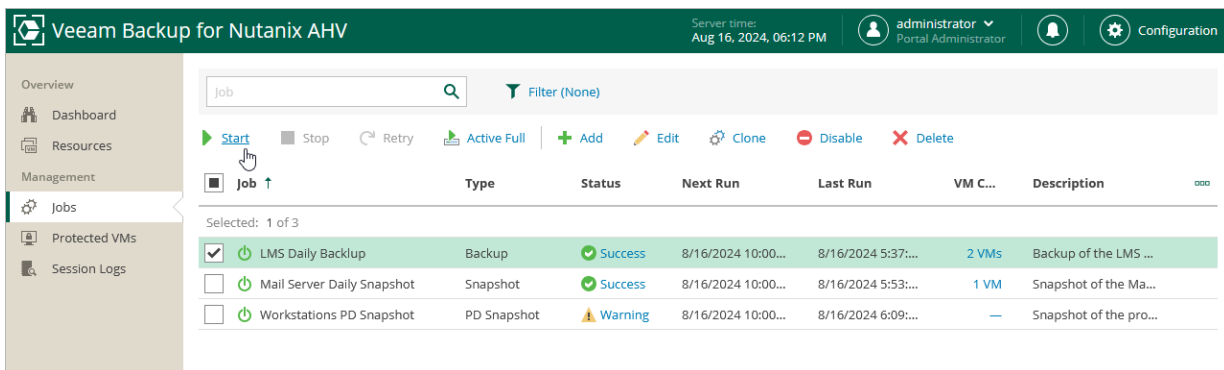
You can also clone a job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Clone** on the ribbon.

# Starting and Stopping Jobs

You can start a job manually, for example, if you want to create an additional restore point and do not want to modify the configured job schedule. You can also stop a job manually if processing of a VM is about to take too long, and you do not want the job to have an impact on the production environment during business hours. When you stop a running job, Veeam Backup for Nutanix AHV creates a new restore point only for those VMs that have already been processed by the time you stop the job.

To start or stop a job, do the following:

1. Navigate to **Jobs**.
2. Select the job.
3. Click **Start** or **Stop**.



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top bar includes the server name, time, user, and configuration options. The left navigation pane shows 'Jobs' selected. The main area displays a table of jobs with the following data:

Job	Type	Status	Next Run	Last Run	VM C...	Description
<input checked="" type="checkbox"/> LMS Daily Backup	Backup	Success	8/16/2024 10:00...	8/16/2024 5:37:...	2 VMs	Backup of the LMS ...
<input type="checkbox"/> Mail Server Daily Snapshot	Snapshot	Success	8/16/2024 10:00...	8/16/2024 5:53:...	1 VM	Snapshot of the Ma...
<input type="checkbox"/> Workstations PD Snapshot	PD Snapshot	Warning	8/16/2024 10:00...	8/16/2024 6:09:...	—	Snapshot of the pro...

## TIP

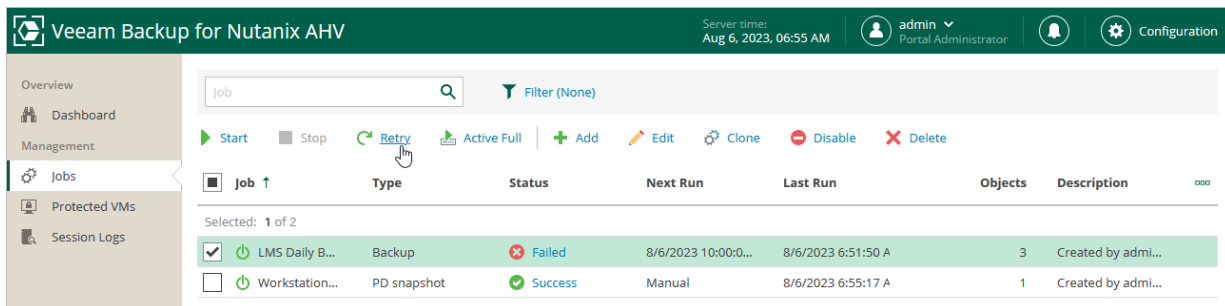
You can also start or stop a job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Start** or **Stop** on the ribbon.

# Retrying Jobs

If a job fails, you can retry the backup operation. When you perform a retry, Veeam Backup for Nutanix AHV restarts the operation only for the failed resources added to the job and does not process VMs that have been processed successfully. As a result, retrying a job takes less time compared to restarting the job for all resources.

To retry a job, do the following:

1. Navigate to **Jobs**.
2. Select the failed job.
3. Click **Retry**.



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top bar displays the server time as Aug 6, 2023, 06:55 AM and the user as admin, Portal Administrator. The left navigation pane includes Overview, Dashboard, Management, Jobs, Protected VMs, and Session Logs. The main area shows a search bar for jobs, a filter set to 'None', and a ribbon with actions: Start, Stop, Retry, Active Full, Add, Edit, Clone, Disable, and Delete. Below the ribbon is a table of jobs with columns for Job, Type, Status, Next Run, Last Run, Objects, and Description. Two jobs are listed: 'LMS Daily B...' (Backup, Failed) and 'Workstation...' (PD snapshot, Success). The 'LMS Daily B...' job is selected, and the 'Retry' button is highlighted.

Job	Type	Status	Next Run	Last Run	Objects	Description
<input checked="" type="checkbox"/> LMS Daily B...	Backup	Failed	8/6/2023 10:00:0...	8/6/2023 6:51:50 A	3	Created by admi...
<input type="checkbox"/> Workstation...	PD snapshot	Success	Manual	8/6/2023 6:55:17 A	1	Created by admi...

## TIP

You can also retry a job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Retry** on the ribbon.


# Editing Job Settings

For each job, you can modify settings configured while creating the job:

1. Navigate to **Jobs**.
2. Select the job.

You can select the job for editing even when it is running.

3. Click **Edit**:
  - To update configuration of a backup job, complete the wizard as described in section [Creating Backup Jobs](#).
  - To update configuration of a snapshot job, complete the wizard as described in section [Creating Snapshot Jobs](#).
  - To update configuration of a protection domain snapshot job, complete the wizard as described in section [Creating Protection Domain Snapshot Jobs](#).



## TIP


You can also start the **Edit Job** wizard using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Edit** on the ribbon.

# Enabling and Disabling Jobs

By default, all created jobs run according to the specified schedules. However, you can temporarily disable a job so that it does not run automatically. You will still be able to enable the disabled job at any time you need.

To enable or disable a backup job, do the following:

1. Navigate to **Jobs**.
2. Select the job.
3. Click **Enable** or **Disable**.



## TIP


You can also enable or disable a job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Enable** or **Disable** on the ribbon.

# Deleting Jobs

You can permanently delete a job from the Veeam Backup for Nutanix AHV configuration database if you no longer need it. When you delete a job, backups created by this job are displayed under the **Backups > Disk (Orphaned)** node in the **Home** view of the Veeam Backup & Replication console. If you want to delete backup files as well, follow the instructions provided in section [Deleting Backups](#).

To delete job, do the following:

1. Navigate to **Jobs**.
2. Select the job.
3. Click **Delete**.



## TIP

You can also delete a job using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Delete** on the ribbon.

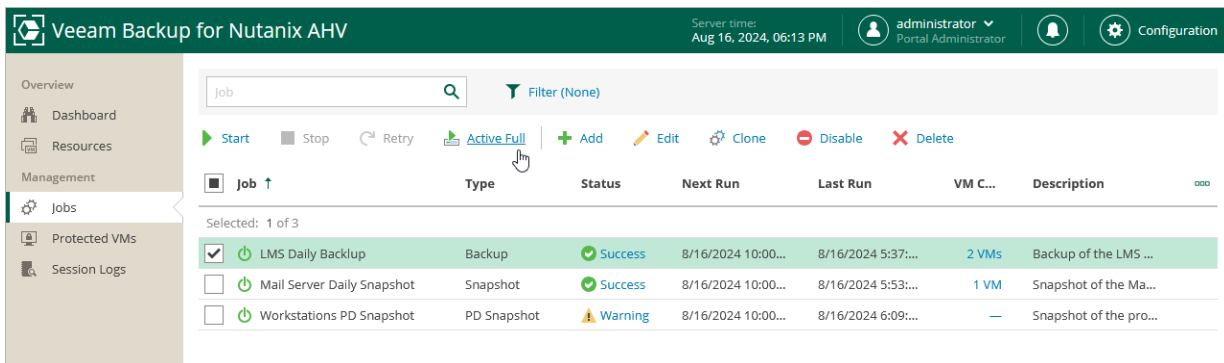
# Creating Active Full Backup

You can manually create an [active full backup](#) for all VMs added to a backup job.

1. Switch to the **Jobs** page.
2. Select the necessary backup job and click **Active Full**.
3. Click **Yes**.

## NOTE

To create active full backup automatically according to a specific schedule, configure backup job settings as described in section [Creating Backup Jobs](#) (step 3).



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Nutanix AHV', the server time 'Aug 16, 2024, 06:13 PM', and the user 'administrator Portal Administrator'. The left sidebar shows navigation options: Overview, Dashboard, Resources, Management, Jobs (selected), Protected VMs, and Session Logs. The main area displays a table of backup jobs with columns for Job, Type, Status, Next Run, Last Run, VM Count, and Description. A ribbon above the table contains actions: Start, Stop, Retry, Active Full (highlighted), Add, Edit, Clone, Disable, and Delete. The table lists three jobs: 'LMS Daily Backup' (Backup, Success, 2 VMs), 'Mail Server Daily Snapshot' (Snapshot, Success, 1 VM), and 'Workstations PD Snapshot' (PD Snapshot, Warning, 0 VMs).

Job	Type	Status	Next Run	Last Run	VM C...	Description
<input checked="" type="checkbox"/> LMS Daily Backup	Backup	Success	8/16/2024 10:00...	8/16/2024 5:37:...	2 VMs	Backup of the LMS ...
<input type="checkbox"/> Mail Server Daily Snapshot	Snapshot	Success	8/16/2024 10:00...	8/16/2024 5:53:...	1 VM	Snapshot of the Ma...
<input type="checkbox"/> Workstations PD Snapshot	PD Snapshot	Warning	8/16/2024 10:00...	8/16/2024 6:09:...	—	Snapshot of the pro...

## TIP

You can also create active full backups using the Veeam Backup & Replication console. To do that, in the inventory pane of the **Home** view, select **Jobs > Backup**, select the job and click **Active Full** on the ribbon.

# Creating VeeamZIP Backups

You can back up one or multiple Nutanix AHV VMs without configuring backup jobs. To do that, you can leverage the VeeamZIP feature — it can be helpful, for example, if you want to create backups for VMs immediately, archive VMs before decommissioning and so on. VeeamZIP produces a full backup that acts as an independent restore point. You can store the backup in a repository added to the backup infrastructure, in a local folder on the backup server or in a network share.

## NOTE

You cannot store VeeamZIP backups in [Veeam Cloud Connect](#) and [HPE Cloud Bank Storage](#) repositories.

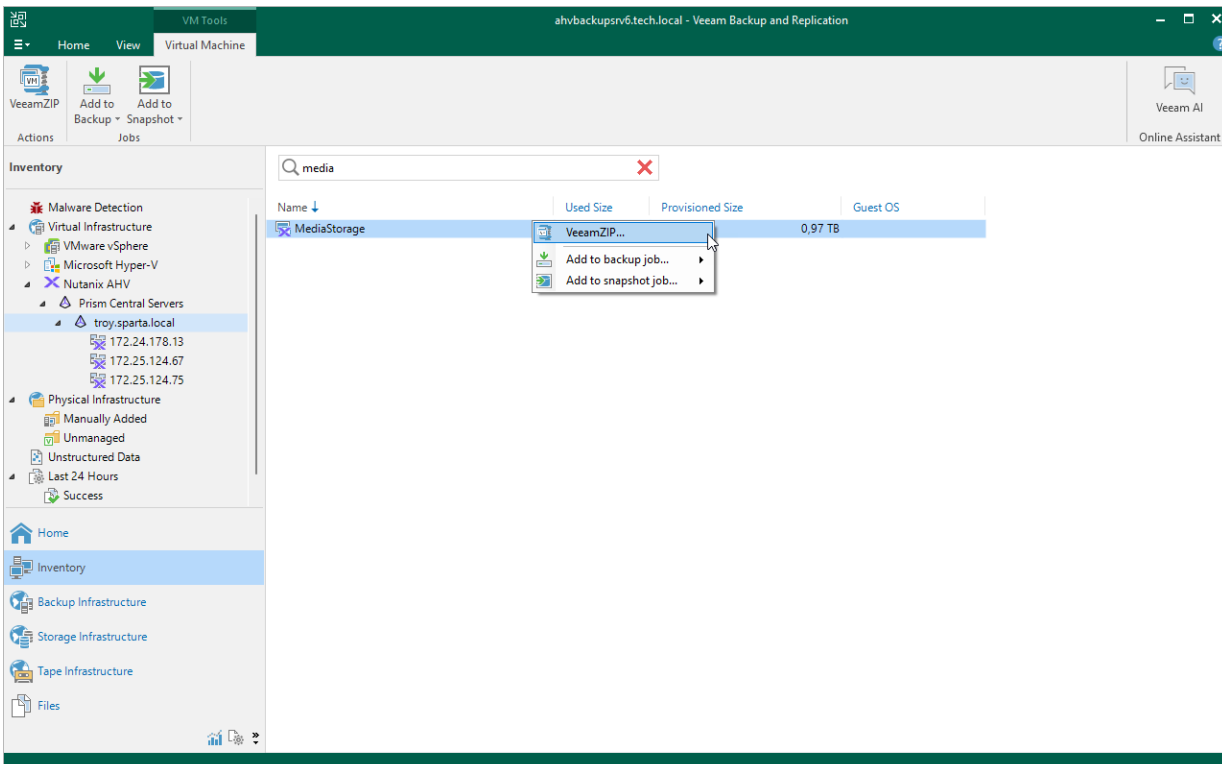
To create a VeeamZIP backup, do the following:

1. In the Veeam Backup & Replication console, open the **Inventory** view.
2. In the inventory pane, select **Nutanix AHV**.
3. In the working area, right-click the VM that you want to back up, and select **VeeamZIP**.
4. Select the destination where the VeeamZIP backup will be stored.

## TIP

You cannot specify an SMB share that requires authentication as a local or shared folder. However, you can [add the SMB share to the backup infrastructure](#) and specify it as backup repository.

The created VeeamZIP backup will be displayed under the **Backups > Disk (Exported)** node in the **Home** view of the Veeam Backup & Replication console.





# Managing Backups and Snapshots

You can perform the following operations with backup files and snapshots:

- [Viewing Backup Properties](#)
- [Rescanning Backups](#)
- [Verifying Backups](#)
- [Exporting Backups](#)
- [Copying Backups](#)
- [Copying Backups to Tapes](#)
- [Deleting Backups](#)
- [Deleting Snapshots](#)

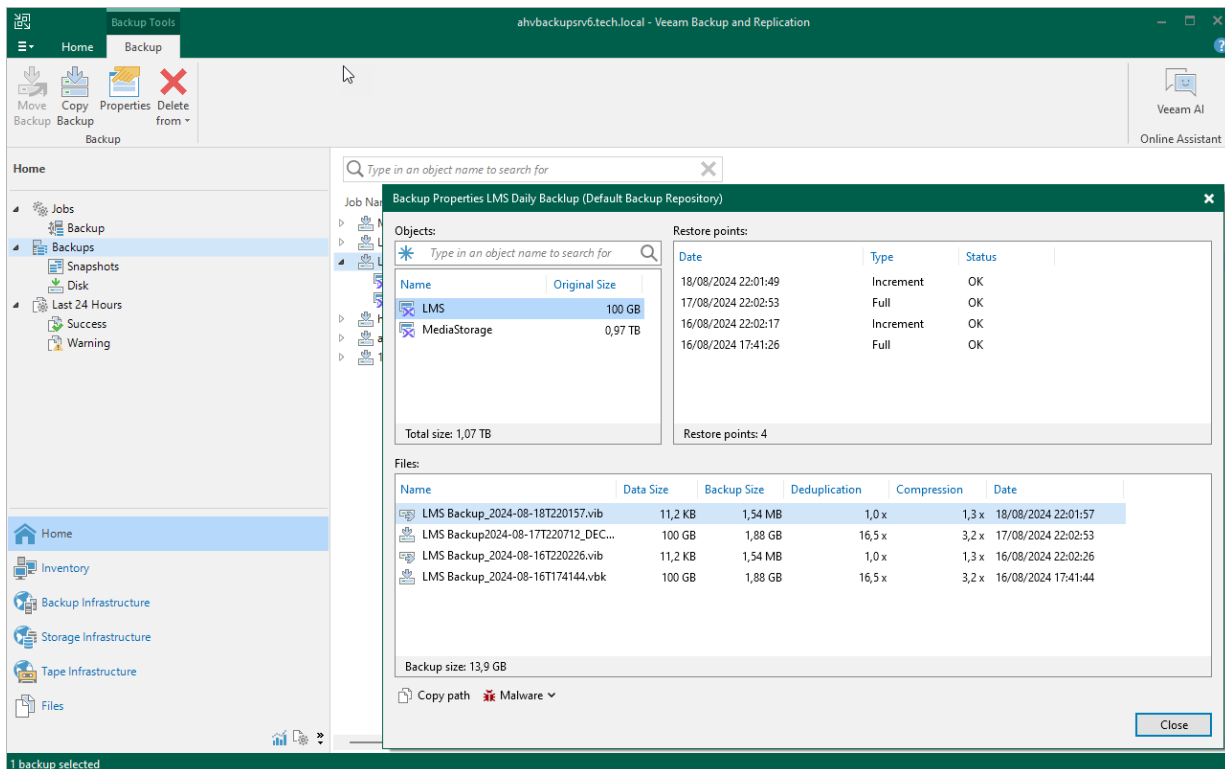
# Viewing Backup Properties

After a backup job successfully creates a backup of a Nutanix AHV VM according to the specified schedule, or after you create an active full backup of a VM manually, the backup is displayed under the **Backups** node in the **Home** view of the Veeam Backup & Replication console. Each backup is represented with a set of properties, such as:

- **Objects** – the names and sizes of backed-up VMs.
- **Restore Points** – the date and time of all restore points created for a VM.
- **Files** – the size of processed VM data, the size and the storage location of backed-up VM data, the ratio of [data deduplication](#) and the ratio of [data compression](#).

To view backup properties, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup and select **Properties**.



# Rescanning Backups

You can import backups created by another Nutanix AHV backup appliance and use them to [restore VMs and VM disks](#).

## NOTE

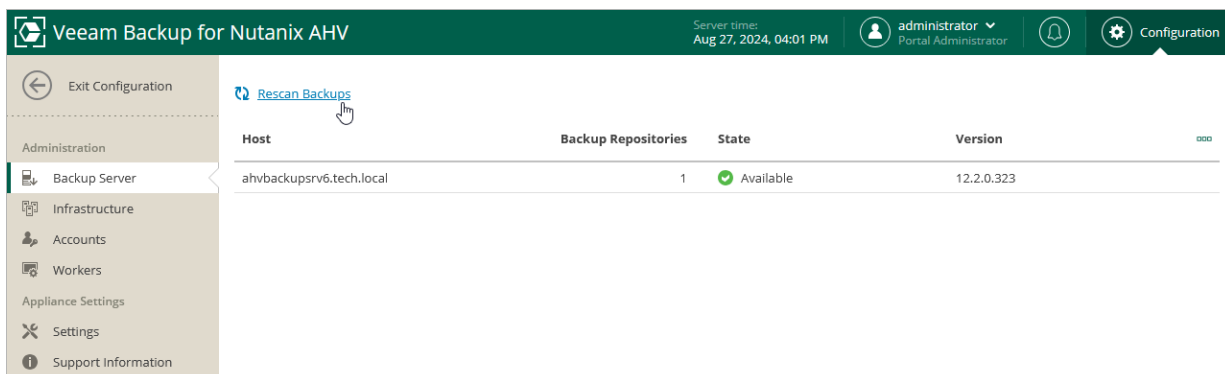
When you import backups, the jobs that produced those backups are not re-created. To re-create the jobs, restore the Nutanix AHV backup appliance configuration as described in section [Restoring Configuration Settings](#).

To import Nutanix AHV VM backups, do the following:

1. Ensure that the [Nutanix AHV backup appliance has access to repositories](#) where the necessary backups are stored.
2. Switch to the **Configuration** page.
3. Navigate to **Backup Server**.
4. Click **Rescan Backups**.

The Nutanix AHV backup appliance will scan the repositories and import all Nutanix AHV VM backups created by Veeam Backup for Nutanix AHV. To track the progress of the import operation, click **Go to Sessions**.

To see the list of imported backups, switch to the **Protected VMs** page.



The screenshot shows the Veeam Backup for Nutanix AHV Configuration page. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Nutanix AHV", the server time "Aug 27, 2024, 04:01 PM", the user "administrator", and the "Configuration" tab. The left sidebar contains navigation options: "Exit Configuration", "Administration", "Backup Server", "Infrastructure", "Accounts", "Workers", "Appliance Settings", "Settings", and "Support Information". The main content area shows a "Rescan Backups" button and a table with the following data:

Host	Backup Repositories	State	Version
ahvbackupsrv6.tech.local	1	Available	12.2.0.323

# Verifying Backups

To perform an integrity check of Nutanix AHV VM backups, Veeam Backup & Replication offers the SureBackup technology that allows you to ensure that the created restore points are not corrupted. For backups of Windows VMs, you can also scan the restore points with antivirus software installed on the backup server, and run YARA rules to detect malware and sensitive data.

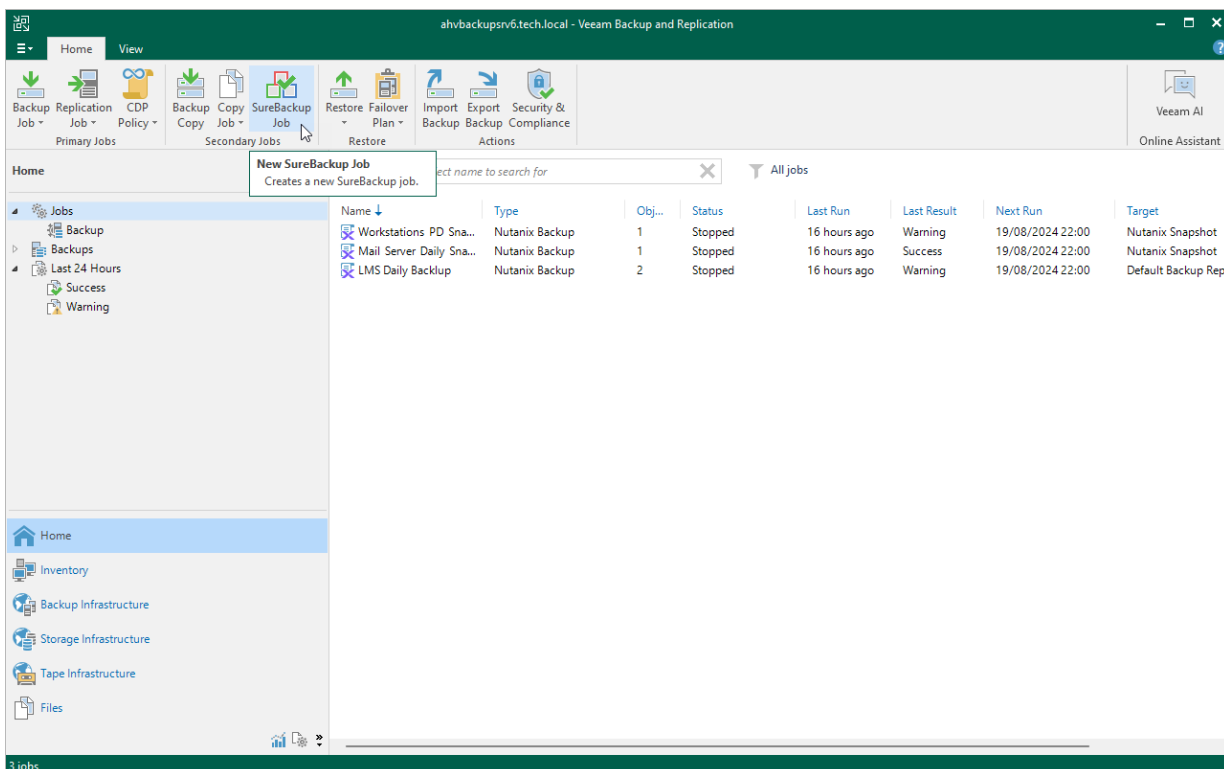
To create a SureBackup job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs** and click **SureBackup Job** on the ribbon.
3. At the **Name** step of the **New SureBackup Job** wizard, select the **Backup verification and content scan only verification** mode, and then complete the wizard as described in the Veeam Backup & Replication User Guide, section [Creating SureBackup Jobs](#).

If any of the verification checks fail for a restore point, Veeam Backup & Replication will mark both this restore point and all subsequent points in the backup chain as *Infected*. To learn how to manage infected restore points, see Veeam Backup & Replication User Guide, section [Managing Malware Status](#).

## TIP

You can scan backups of Windows VMs manually on demand, without creating a SureBackup job. To learn how to do that, see the Veeam Backup & Replication User Guide, section [Scan Backup](#).



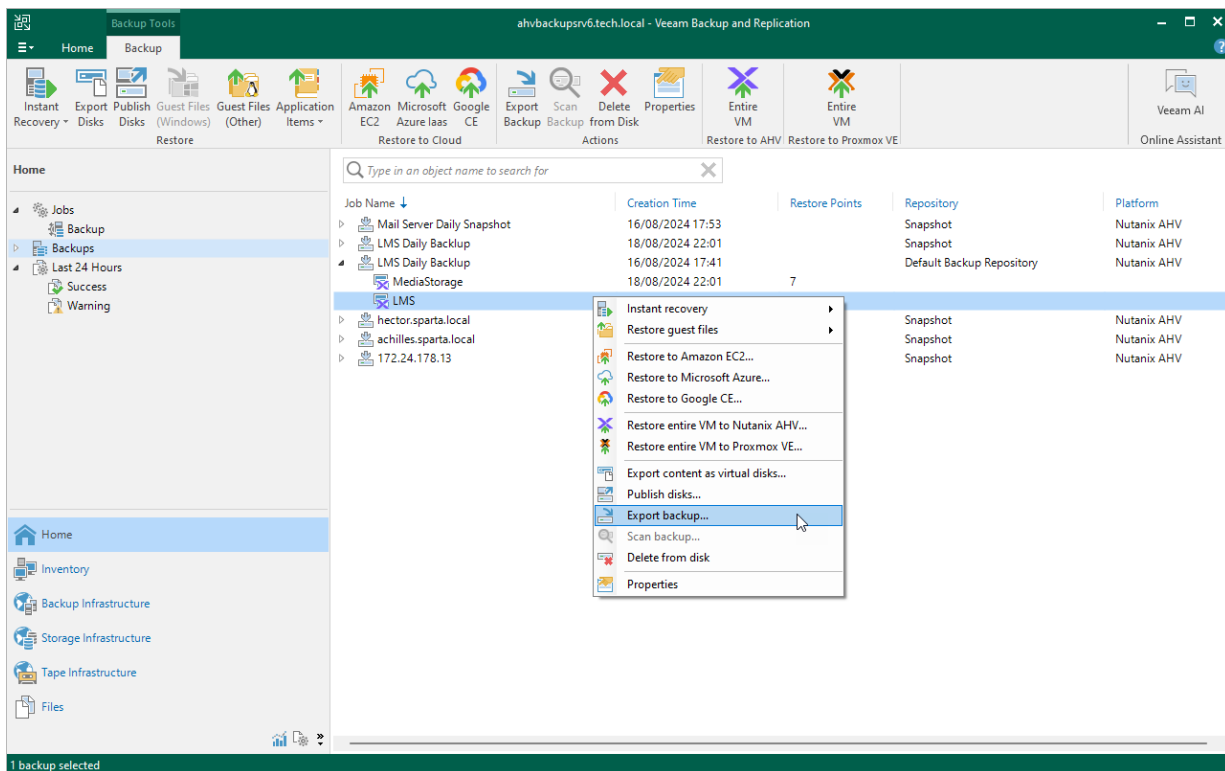
# Exporting Backups

Exporting backups allows you to synthesize a complete and independent full backup file using restore points located in your backup repositories. That is, you can transform any backup chain into a standalone full backup file and save it to the same repository where the selected restore points reside.

To export a backup, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click a VM for which you want to synthesize a full backup file, and select **Export Backup**.
4. Complete the **New Export** wizard as described in the Veeam Backup & Replication User Guide, section [Performing Export](#).

Once the export operation completes, the exported backup will be displayed under the **Backups > Disk (Exported)** node in the **Home** view of the Veeam Backup & Replication console.



# Copying Backups

With backup copy, you can create several instances of a backup and copy them to secondary (target) backup repositories for long-term storage. Target backup repositories can be located in the same site as the source backup repository or can be deployed off-site. Since the backup copy has the same format as the original backup, you can restore VM data directly from the backup copy in case a disaster strikes. For more information on the backup copy functionality, see the Veeam Backup & Replication User Guide, section [Backup Copy](#).

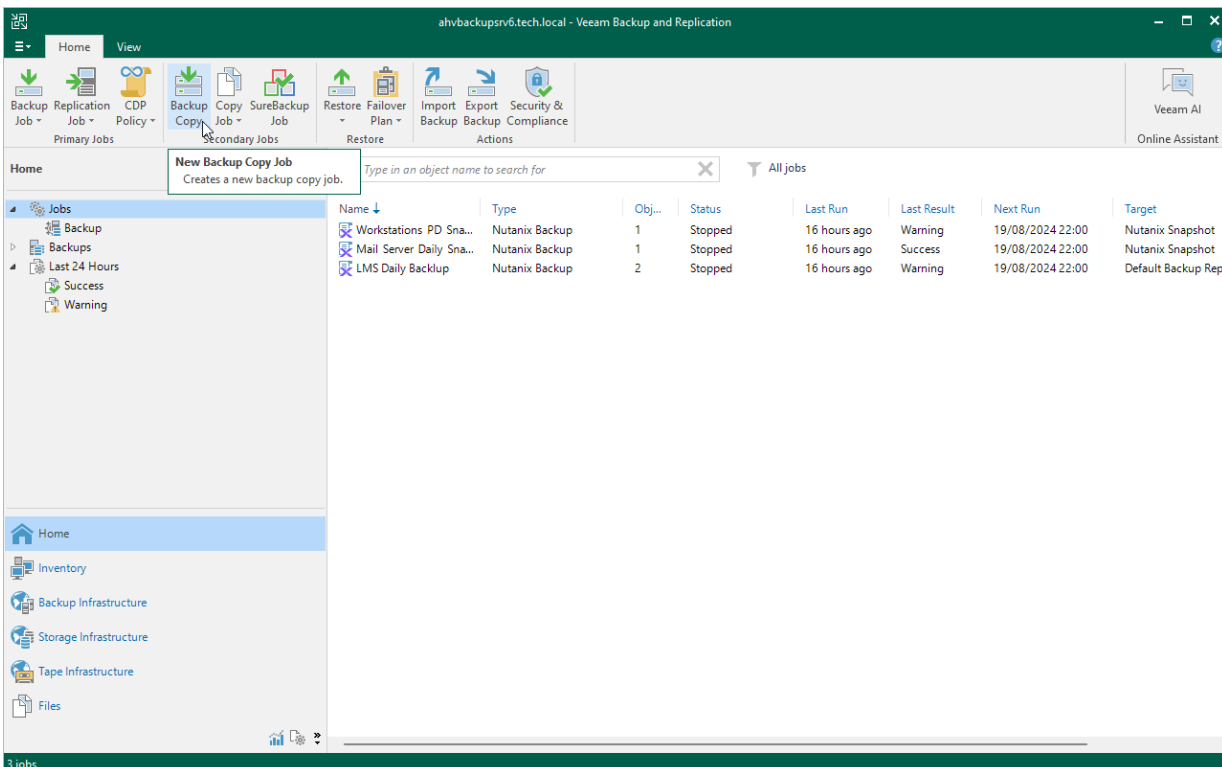
To copy backups to a secondary backup repository, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs** and click **Backup Copy** on the ribbon.
3. Create a backup copy job as described in the Veeam Backup & Replication User Guide, section [Creating Backup Copy Jobs](#).

Note that you can copy backups to [Veeam Cloud Connect repositories](#) if a service provider is added to Veeam Backup & Replication. You can also restore VMs from backup copies stored in those repositories if you use Veeam Cloud Connect version 12.

## TIP

Alternatively, you can create a copy of a backup without configuring a job as described in the Veeam Backup & Replication User Guide, section [Copying Backups](#).



# Copying Backups to Tapes

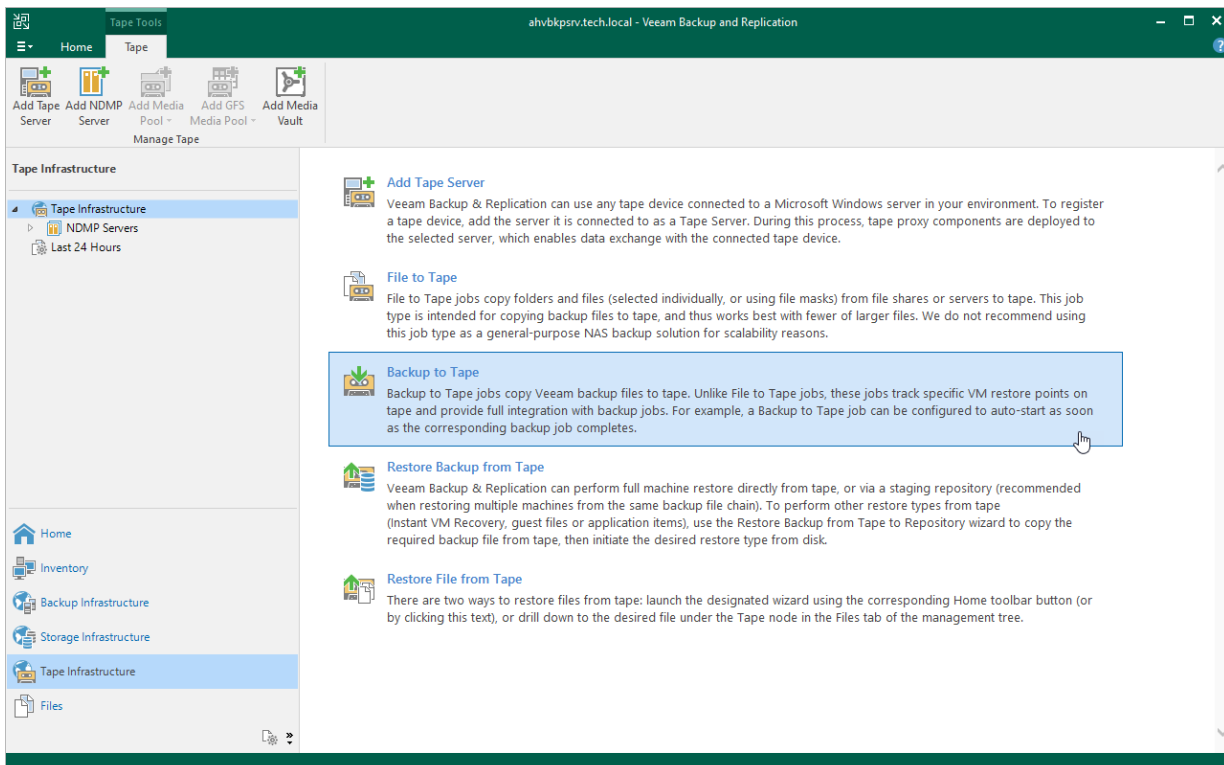
You can create archives of Nutanix AHV VM backups and copy them to tapes for long-term storage. Veeam Backup for Nutanix AHV allows you to manage tape archives the same way you manage backups in backup repositories. However, it usually takes more time to access archived data on tapes than to access backed-up data in repositories. For more information on tapes, see the Veeam Backup & Replication User Guide, section [Tape Devices Support](#).

To archive Nutanix AHV VM backups to tape, do the following:

1. Configure the tape infrastructure:
  - a. Connect tape devices as described in the Veeam Backup & Replication User Guide, section [Tape Devices Deployment](#).
  - b. Perform initial configuration of the tape infrastructure as described in the Veeam Backup & Replication User Guide, section [Getting Started with Tapes](#) (steps 1-3).
2. Create a backup to tape job as described in the Veeam Backup & Replication User Guide, section [Creating Backup to Tape Jobs](#).

## NOTE

You cannot restore Nutanix AHV VMs directly from tapes. To restore a Nutanix AHV VM, you must first restore its backups to a repository as described in the Veeam Backup & Replication User Guide, section [Backup Restore from Tape to Repository](#).



# Deleting Backups

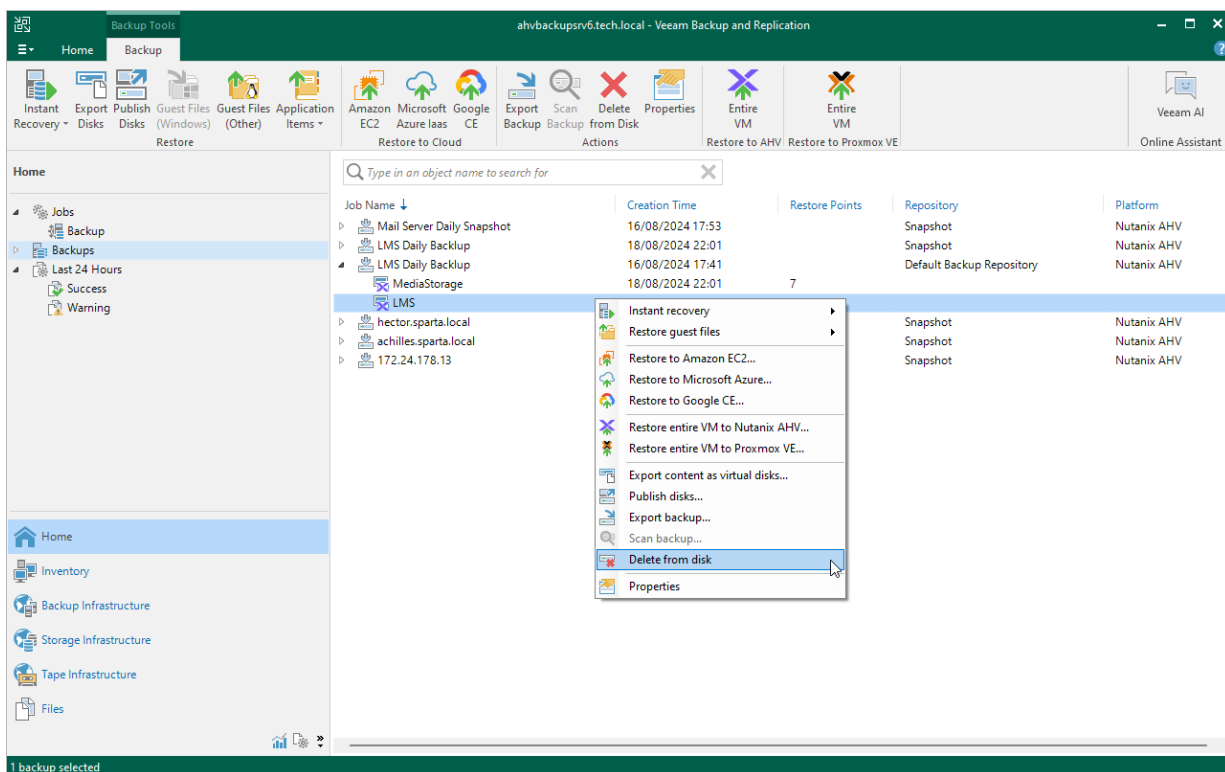
By default, Veeam Backup for Nutanix AHV maintains backups in backup repositories according to retention policy settings saved in the backup metadata. If Veeam Backup for Nutanix AHV detects that the number of restore points in the backup chain exceeds the allowed number, it automatically removes obsolete backups. If necessary, you can delete backups manually.

To delete backup files created for a Nutanix AHV VM by a backup or snapshot job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane of the **Home** view, select **Backups**.
3. In the working area, expand the job that created the backup, right-click the VM name and select **Delete from disk**.

## NOTE

[Applies only to Veeam Backup & Replication version 12.1 and later] If [4-eyes authorization](#) is enabled in Veeam Backup & Replication, deleting backup files will require additional approval from another user with the *Veeam Backup Administrator* role.





# Deleting Snapshots

Veeam Backup for Nutanix AHV applies the [configured retention policy settings](#) to automatically remove snapshots created by jobs. If necessary, you can also remove snapshots manually.


To remove snapshots, do the following:

1. Navigate to **Protected VMs**.
2. Select a VM whose snapshots you want to delete.
3. Click **Remove Snapshots**:
  - Select **All** to remove all VM snapshots including snapshots taken in the Nutanix AHV Prism console.
  - Select **Created by Job** to remove only snapshots created by Veeam Backup for Nutanix AHV jobs.

## NOTE

If the VM is protected with PD snapshots, you can choose whether you want to delete or skip them. If you choose to delete the PD snapshots, Veeam Backup for Nutanix AHV will delete snapshots of the selected VM and of all other VMs included into the protection domain. If you choose to skip the PD snapshots, snapshots of the selected VM included in the protection domain will retain.

4. Click **Remove** to confirm.



# Performing Restore

In various disaster recovery scenarios, Veeam Backup for Nutanix AHV allows you to perform the following operations using backed-up data:

- [Entire VM restore](#) – recover Nutanix AHV VMs to the original location or to a new location.
- [VM disk restore](#) – recover a specific VM disk and attach it to the original VM or to another VM.
- [Instant VM recovery](#) – instantly start a VM directly from a backup.
- [Disk publishing](#) – mount specific disks of a backed-up Nutanix AHV VMs to any server added to the backup infrastructure.
- [File-level restore](#) – recover individual VM guest OS files and folders.
- [Application items restore](#) – restore applications, such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, and Microsoft SQL Server.
- [VM disk export](#) – restore VM disks and convert them to disks of the VMDK, VHD or VHDX format.
- [Restore to AWS](#) – restore Nutanix AHV VMs to Amazon Web Services as EC2 instances.
- [Restore to Microsoft Azure](#) – restore Nutanix AHV VMs to Microsoft Azure as Azure VMs.
- [Restore to Google Cloud](#) – restore Nutanix AHV VMs to Google Cloud as VM instances.

# Performing VM Restore

In case a disaster strikes, you can restore an entire Nutanix AHV VM from a backup or snapshot. Veeam Backup for Nutanix AHV allows you to restore one or more VMs at a time, to the original location or to a new location.

## Supported Workloads

To restore machines to a Nutanix AHV cluster, you can use the following backups and snapshots:

- Snapshots of Nutanix AHV PDs created by Veeam Backup for Nutanix AHV
- Snapshots of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV
- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV (including VMs with volume groups attached and VMs with no disks attached)
- Backups of Microsoft Hyper-V and VMware vSphere VMs created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows and Veeam Agent for Linux
- Backups of VMs created by vCloud Director
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Google Cloud VM instances created by Veeam Backup for Google Cloud
- Backups of oVirt KVM VMs created by Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

VM restore is supported only for snapshots stored in the Nutanix AHV cluster and for backups stored in backup repositories, object storage repositories, and on the performance, capacity and archive tier of a scale-out backup repository (except for backups stored in the archive tier that consists of the Amazon S3 Glacier Instant Retrieval extent; for those backups, you can perform [Instant Recovery](#)).

### NOTE

You cannot restore VMs from backups stored in external repositories, Veeam Cloud Connect repositories, HPE Cloud Bank Storage and on tapes.

## How to Perform VM Restore

To restore a protected VM, you can use either the [Veeam Backup & Replication console](#) or the [Nutanix AHV backup appliance web console](#). However, only the Veeam Backup & Replication console allows you to restore workloads of different types and to choose a location for the restored VMs.

# Restoring VMs Using Veeam Backup & Replication Console

From the Veeam Backup & Replication console, you can restore one or multiple VMs to any Nutanix AHV cluster added to the backup infrastructure.

To restore a protected VM, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Full VM Restore wizard.](#)
3. [Select VMs to restore.](#)
4. [Choose a restore mode.](#)
5. [Specify a target cluster.](#)
6. [Select a storage container where VM virtual disks will be stored.](#)
7. [Specify a new name for the restored VM.](#)
8. [Configure network settings.](#)
9. [Specify a restore reason.](#)
10. [Verify restore settings.](#)

## Before You Begin

Before you perform VM restore, consider the following limitations:

- When restoring a VM from a [snapshot, backup snapshot or PD snapshot](#), Veeam Backup for Nutanix AHV stores virtual disks of the recovered VM in the original storage container.
- When restoring a VM from a [snapshot or PD snapshot](#), Veeam Backup for Nutanix AHV retains the original VM network settings. After the VM is restored, you can change these settings using the Nutanix Prism console as described in [Nutanix documentation](#).
- To restore a VM from a backup stored in the archive tier of a scale-out backup repository, you must first retrieve backup data as described in the Veeam Backup & Replication User Guide, section [Retrieving Backup Files](#).
- A VM restored from a backup created by a solution other than Veeam Backup for Nutanix AHV may become unreachable through the network. To resolve the issue, log in to the VM console using Nutanix AHV Prism Element console and install Nutanix Guest Tools as described in [Nutanix documentation](#).
- When restoring a VM that originally resided on a platform other than Nutanix AHV, Veeam Backup for Nutanix AHV attaches VM disks with the restored data to the target VM disk nodes using their original bus types. Veeam Backup for Nutanix AHV can attach to a VM up to 6 SATA, 256 SCSI, 4 IDE and 7 PCI disks. If the VM has more disks of any of those bus types, Nutanix AHV will attach the disks to remaining nodes of other bus types in the default priority: SATA, SCSI, IDE, PCI. You can [modify the backup appliance configuration](#), to instruct Nutanix AHV to ignore source VM original bus types and to use a specific order of bus types.
- When restoring a VM to a new location, Veeam Backup for Nutanix AHV does not restore the VM affinity policy configuration. Therefore, you must manually configure the affinity policy as described in [Nutanix documentation](#).

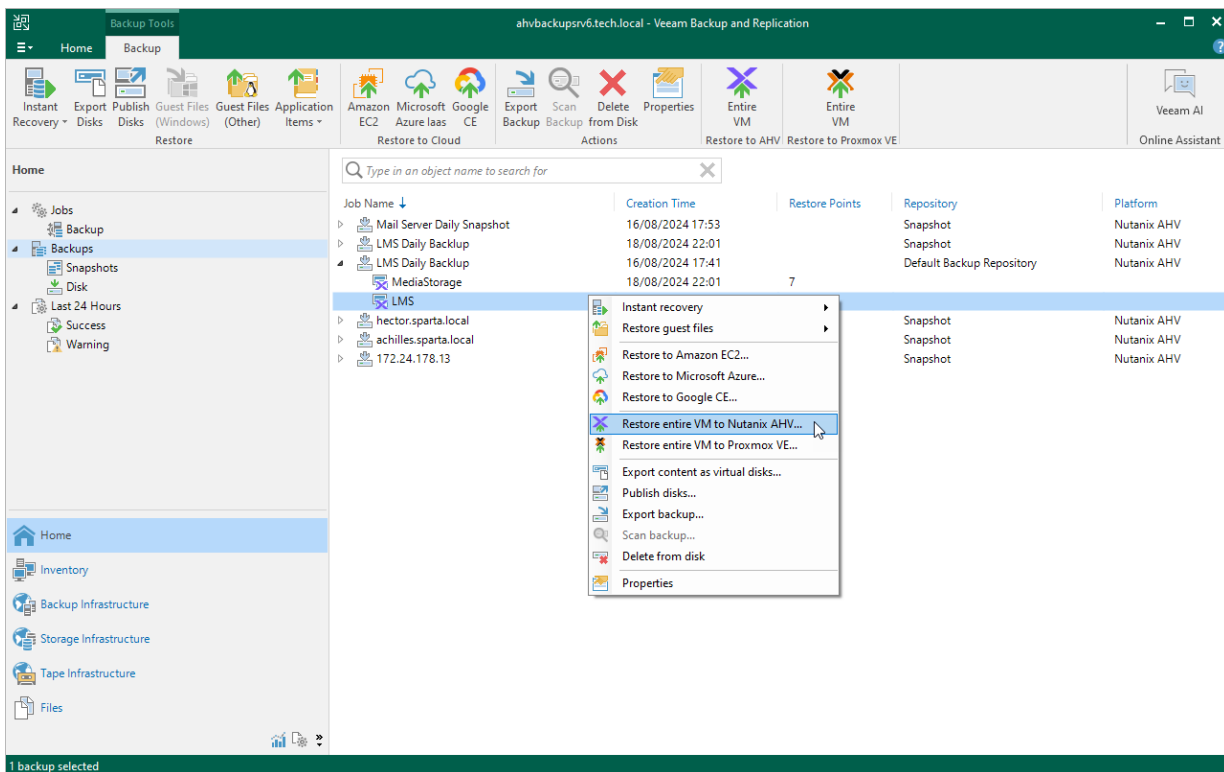
# Step 1. Launch Restore Wizard

To launch the **Full VM Restore to Nutanix AHV** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup or snapshot job, select the VM that you want to restore and click **Entire VM** on the ribbon, or right-click the VM and select **Restore entire VM to Nutanix AHV**.

## TIP

To restore a VM from a snapshot taken in the Nutanix AHV Prism Element console, expand the Nutanix AHV cluster, select the VM that you want to restore and click **Entire VM** on the ribbon, or right-click the VM and select **Restore entire VM to Nutanix AHV**.



## Step 2. Select Restore Point

At the **Virtual Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Backup for Nutanix AHV uses the most recent valid restore point. However, you can restore the VM data to an earlier state.

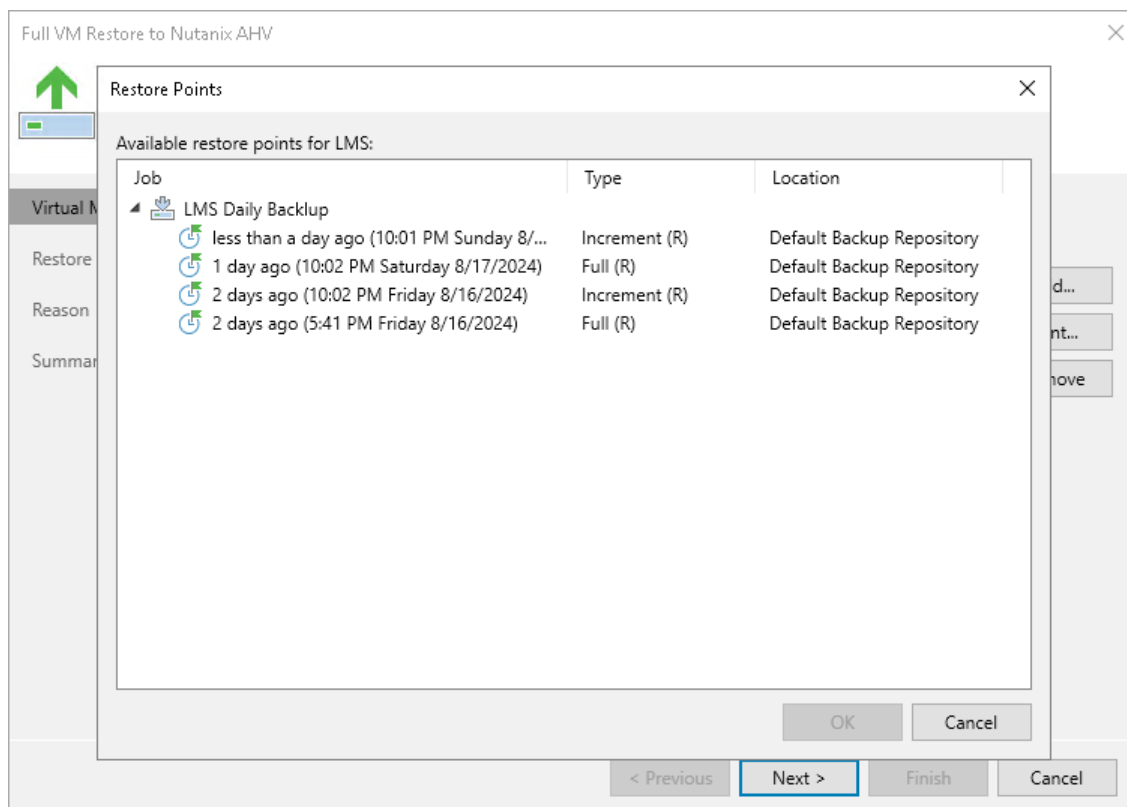
To select a restore point, do the following:

1. Select the VM.
2. Click **Point**.
3. In the **Restore Points** window, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup for Nutanix AHV provides the following information on each available restore point:

- **Job** – the name of the backup job that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the repository where the restore point is stored.

You can use the wizard to restore multiple VMs at a time. To do that, click **Add**, select more VMs to restore and select a restore point for each of them.



## Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected VM to the original or to a custom location.

Full VM Restore to Nutanix AHV

**Restore Mode**  
Specify whether selected VMs should be restored back to the original location, or to a new location or with different settings.

Virtual Machines

**Restore Mode**

Cluster

Storage Container

Name

Network

Reason

Summary

**Restore to the original location**  
Quickly initiate the restore of selected VM to its original location, with the original name and settings. This option minimizes the chance of user input error.

**Restore to a new location, or with different settings**  
Customize the restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the defaults.

< Previous   **Next >**   Finish   Cancel

## Step 4. Specify Target Cluster

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Cluster** step of the wizard, choose the cluster to which the recovered VM will belong. In the Prism Central deployment, you can also choose whether you want the recovered VM to be assigned the same categories as the original VM.

For a cluster to be displayed in the list of the available clusters, it must be added to the backup infrastructure as described in section [Adding Nutanix AHV Server](#).

### NOTE

The **Cluster** step of the **Full VM Restore to Nutanix AHV** wizard is only available when you restore the VM from a backup.

The screenshot shows the 'Full VM Restore to Nutanix AHV' wizard window. The title bar reads 'Full VM Restore to Nutanix AHV'. The main area is titled 'Cluster' and contains the following text: 'By default, original cluster is selected as restore destination for each VM. You can change cluster by selecting desired VM and clicking Cluster. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.' Below this text is a table with the following data:

Name	Cluster
LMS	172.25.124.67

Below the table, there is a checkbox labeled 'Restore VM categories' which is checked. Underneath it, the text reads 'Applies only to VMs managed by Prism Centrals'. At the bottom right of the main area is a button labeled 'Cluster...'. At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.



## Step 5. Select Storage Container

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Storage Container** step of the wizard, choose the storage container where virtual disks of the recovered VM will be stored.

For a container to be displayed in the list of the available containers, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

### NOTE

You cannot choose a storage container when restoring the VM from a snapshot.

The screenshot shows the 'Full VM Restore to Nutanix AHV' wizard window. The 'Storage Container' step is active, indicated by a green arrow icon and a blue bar. The window title is 'Full VM Restore to Nutanix AHV'. Below the title bar, there is a green arrow icon and the text 'Storage Container'. Below that, it says 'By default, original storage container is selected for each VM. You can change them by selecting desired VM, and clicking Container.'

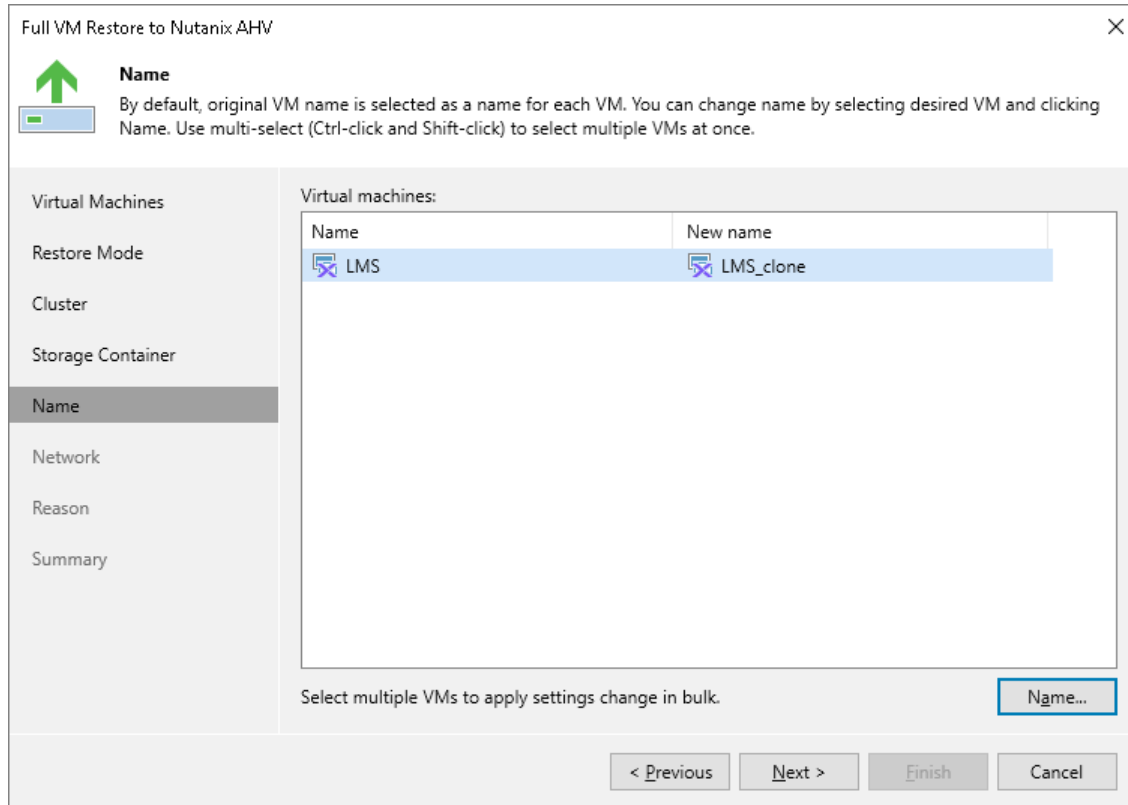
The main area of the wizard is divided into two panes. The left pane is a sidebar with the following items: 'Virtual Machines', 'Restore Mode', 'Cluster', 'Storage Container' (highlighted), 'Name', 'Network', 'Reason', and 'Summary'. The right pane is titled 'Storage container:' and contains a table with the following columns: 'Virtual machine', 'Size', 'Storage container', and 'Cluster'. The table has two rows: one for 'LMS' (selected) and one for 'default-containe...'. The 'LMS' row shows a size of '100 GB' and a storage container of 'NutanixManagem...'. The 'Cluster' column for 'LMS' shows '172.24.178.13'. Below the table, there is a button labeled 'Container...'. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

Virtual machine	Size	Storage container	Cluster
LMS	100 GB	NutanixManagem...	172.24.178.13
default-containe...			

## Step 6. Specify VM Name

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Name** step of the wizard, you can specify a new name for the recovered VM.



# Step 7. Configure Network Settings

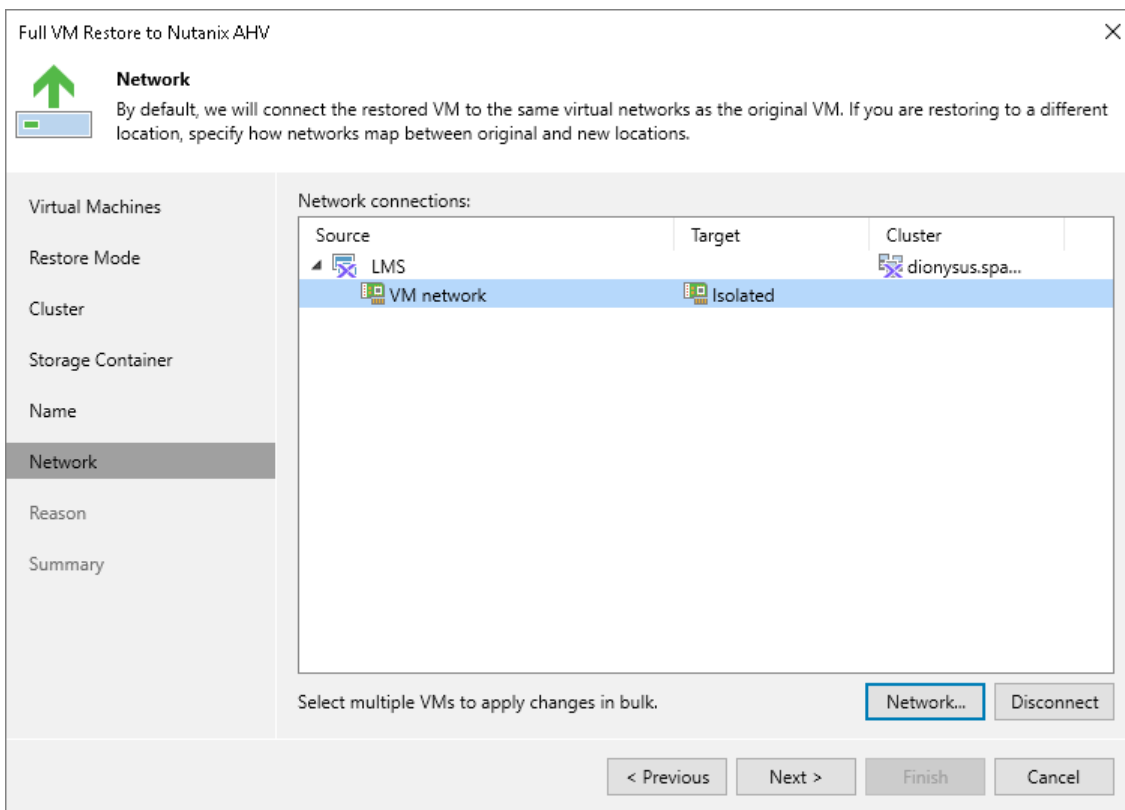
[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, choose a network to which the recovered VM will be connected. If you do not want to connect the VM to any virtual network, select the VM and click **Disconnect**.

For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

## NOTE

You cannot change network settings when restoring the VM from a snapshot. However, you can choose to disconnect the original network.



## Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM. This information will be saved to the session history, and you will be able to reference it later.

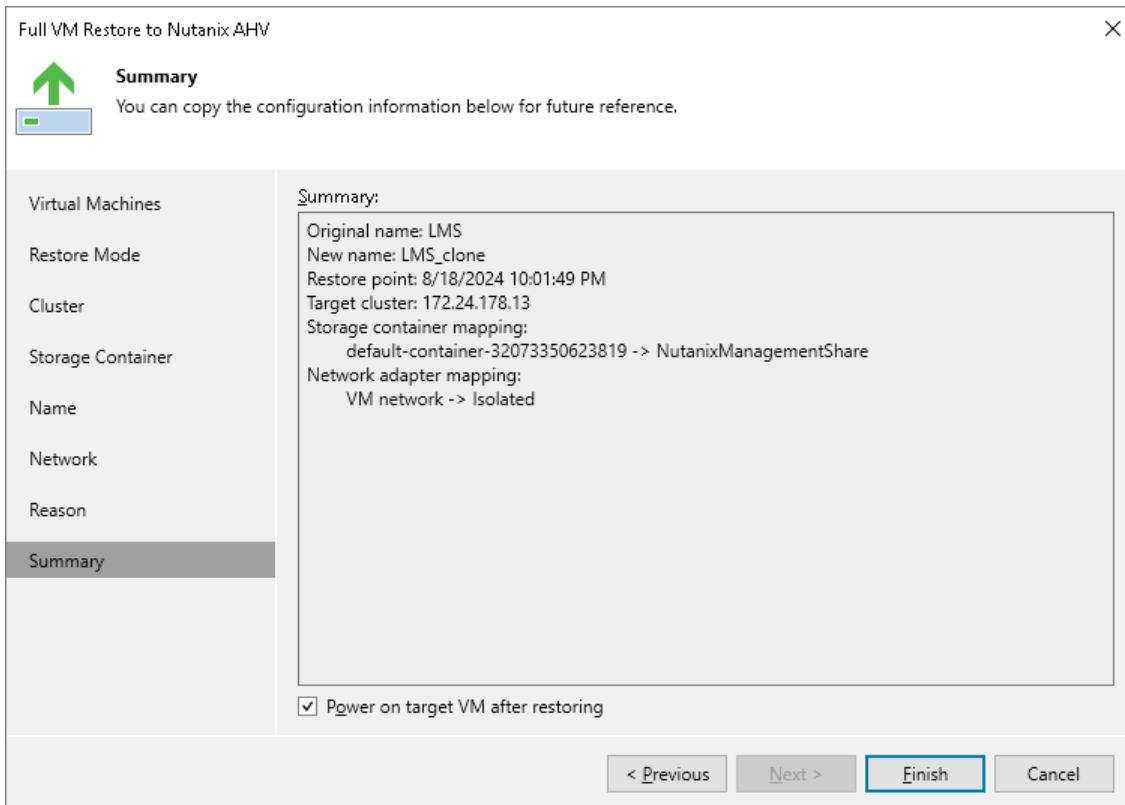
The screenshot shows a wizard window titled "Full VM Restore to Nutanix AHV" with a close button (X) in the top right corner. The window is divided into a left sidebar and a main content area. The sidebar contains a list of steps: "Virtual Machines", "Restore Mode", "Cluster", "Storage Container", "Name", "Network", "Reason" (which is highlighted with a dark grey background), and "Summary". Above the sidebar, there is a green upward-pointing arrow icon and a small progress indicator. The main content area has a heading "Reason" followed by the instruction: "Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference." Below this is a large text input field containing the text "Restoring failed VM". At the bottom of the main content area, there is a checkbox labeled "Do not show me this page again". At the very bottom of the window, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

## Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

### TIP

If you want to start the recovered VM as soon as the restore process completes, select the **Power on target VM after restoring** check box.



The screenshot shows the 'Full VM Restore to Nutanix AHV' wizard at the 'Summary' step. The window title is 'Full VM Restore to Nutanix AHV' with a close button (X) in the top right corner. A green arrow icon is in the top left. The 'Summary' section is active, showing a list of configuration items on the left and their details on the right. The 'Power on target VM after restoring' checkbox is checked. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish' (highlighted with a blue border), and 'Cancel'.

**Full VM Restore to Nutanix AHV** [X]

**Summary**  
You can copy the configuration information below for future reference.

Virtual Machines	Summary:
Restore Mode	Original name: LMS New name: LMS_clone
Cluster	Restore point: 8/18/2024 10:01:49 PM Target cluster: 172.24.178.13
Storage Container	Storage container mapping: default-container-32073350623819 -> NutanixManagementShare
Name	Network adapter mapping: VM network -> Isolated
Network	
Reason	
Summary	

Power on target VM after restoring

< Previous   Next >   **Finish**   Cancel

# Restoring VMs Using Backup Appliance Web Console

In the Nutanix AHV backup appliance web console, you can restore a Nutanix AHV VM from backups and snapshots to the cluster where the original VM belongs.

To restore a protected VM, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Full VM Restore wizard.](#)
3. [Select a restore point.](#)
4. [Choose a restore mode.](#)
5. [Specify a target cluster.](#)
6. [Specify a new name for the restored VM.](#)
7. [Select a container where VM virtual disks will be stored.](#)
8. [Configure network settings.](#)
9. [Specify a restore reason.](#)
10. [Finish working with the wizard.](#)

## Before You Begin

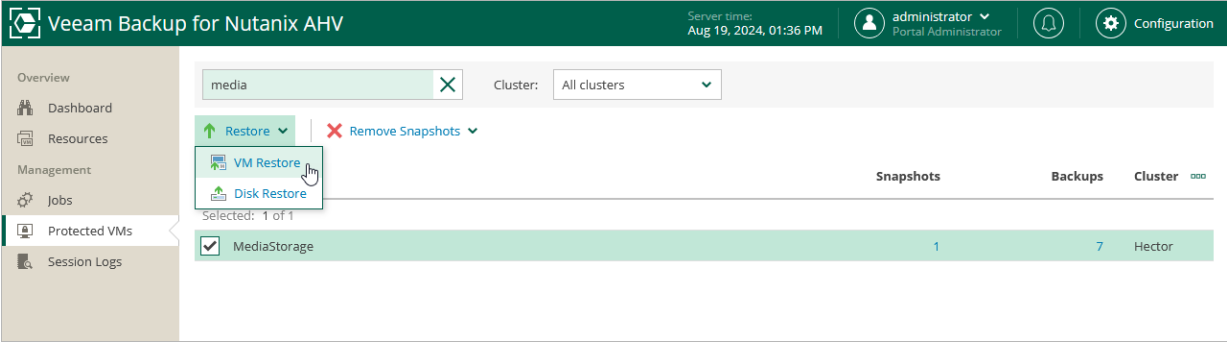
Before you perform Nutanix AHV VM restore, consider the following limitations:

- When restoring the VM from a backup, Veeam Backup for Nutanix AHV stores all virtual disks of the recovered VM in one storage container.
- When restoring a VM from a [snapshot, backup snapshot or PD snapshot](#), Veeam Backup for Nutanix AHV stores virtual disks of the recovered VM in the original storage container.
- When restoring a VM from a [snapshot or PD snapshot](#), Veeam Backup for Nutanix AHV retains the original VM network settings. After the VM is restored, you can change these settings using the Nutanix Prism console as described in [Nutanix documentation](#).
- When restoring a VM to a new location, Veeam Backup for Nutanix AHV does not restore the VM affinity policy configuration. Therefore, you must manually configure the affinity policy as described in [Nutanix documentation](#).

# Step 1. Launch Full VM Restore Wizard

To launch the **Full VM Restore** wizard, do the following:

- 1. Navigate to **Protected VMs**.
- 2. Select the VM that you want to restore.
- 3. Click **Restore > VM Restore**.



## Step 2. Select Restore Point

At the **Virtual Machine** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Backup for Nutanix AHV uses the most recent valid restore point. However, you can restore the VM data to an earlier state.

To select a restore point, do the following:

1. Select the VM.
2. Click **Restore Point**.
3. In the **Select restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Nutanix AHV provides the following information on each available restore point:

- **Job Name** – the name of the backup job that created the restore point.
- **Creating Time** – the time and date when the restore point was created.
- **Type** – the type of the restore point:
  - *Backup* – an image-level backup created by a backup job.  
[Applies to the [Prism Central deployment](#)] Only backups are supported for restore to another cluster.
  - *Backup snapshot* – a snapshot created by a backup job.
  - *Snapshot* – a snapshot created by a snapshot job or manually taken in the Nutanix AHV Prism Element console.
  - *PD snapshot* – a snapshot created with a protection domain snapshot.

### TIP

While creating a backup, Veeam Backup for Nutanix AHV takes a VM snapshot that is called backup snapshot. Veeam Backup for Nutanix AHV stores a recent backup snapshot for each backup job. Restore from the backup snapshot is significantly faster than restore from a backup. However, when you restore from the backup snapshot, [limitations are applied](#). For more information on snapshots, see [Snapshot Types](#),



You can use the wizard to restore multiple VMs at a time. To do that, click **Add VM**, select more VMs to restore and select a restore point for each of them.

The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Nutanix AHV', the server time 'Aug 19, 2024, 01:47 PM', and user information 'administrator Portal Administrator'. The main window is titled 'Full VM Restore'. On the left, a sidebar contains 'Virtual Machines', 'Restore Mode', 'Reason', and 'Summary'. The 'Virtual Machines' section is active, showing 'Choose VMs to restore.' with a search box containing 'VM'. Below the search box, there are two checkboxes: 'VM ↑' (checked) and 'MediaStorage' (checked). The 'Selected: 1 of 1' text is visible. The 'Select restore point' dialog is open on the right, featuring a search box for 'Job Name' and a 'Filter (None)' button. Below the search box is a table with the following data:

Job Name	Creation Time ↓	Type
N/A	8/19/2024 1:44:06 PM	Snapshot
LMS Daily Backup	8/18/2024 10:01:49 PM	Backup snapshot
LMS Daily Backup	8/18/2024 10:01:49 PM	Backup
LMS Daily Backup	8/17/2024 10:47:08 PM	Backup
LMS Daily Backup	8/17/2024 10:32:06 PM	Backup
LMS Daily Backup	8/17/2024 10:17:59 PM	Backup
LMS Daily Backup	8/17/2024 10:02:50 PM	Backup
LMS Daily Backup	8/16/2024 10:02:19 PM	Backup
LMS Daily Backup	8/16/2024 5:41:26 PM	Backup

At the bottom of the dialog, there are 'Apply' and 'Cancel' buttons.

## Step 3. Choose Restore Mode

At the **Restore Mode** step, choose whether you want to restore the VM with the original settings or to specify new settings (such as VM network and disk storage settings).

The screenshot shows the Veeam Backup for Nutanix AHV interface. At the top, the header includes the product name, server time (Aug 19, 2024, 01:49 PM), user (administrator), and a Configuration icon. The main title is 'Full VM Restore'. A left-hand navigation pane lists 'Virtual Machines', 'Restore Mode' (highlighted), 'Cluster', 'Name', 'Storage Container', 'Network Settings', 'Reason', and 'Summary'. The 'Restore Mode' section contains the following text and options:

**Restore Mode**  
Specify whether you want to restore VMs to the original location or to a new one, or with different settings.

- Restore to original location  
Quickly restore the selected VMs to their original location, with the same name and settings as the original VMs.
- Restore to new location, or with different settings  
Perform additional configuration steps to restore the selected VMs to a new location or to use settings that differ from the original settings.

At the bottom of the screen, there are three buttons: 'Previous', 'Next', and 'Cancel'.

## Step 4. Specify Target Cluster

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard in the Prism Central deployment]

At the **Cluster** step of the wizard, choose the cluster to which the recovered VM will belong. You can also choose whether you want the recovered VM to be assigned the same categories as the original VM.

### NOTE

The **Cluster** step of the **Full VM Restore** wizard is only available when you restore the VM from a backup.

The screenshot shows the 'Full VM Restore' wizard in the Veeam Backup for Nutanix AHV interface. The 'Cluster' step is active, showing a 'Choose cluster' dialog box. The dialog box contains a search bar and a table with columns for 'Cluster', 'CPU Usage', and 'Memory Usage'. The 'Vegas' cluster is selected. The 'Apply' button is highlighted.

Virtual Machines

Restore Mode

Cluster

Name

Storage Container

Network Settings

Reason

Summary

**Cluster**

For each of the selected VMs, specify a target cluster.

Restore VM categories (Prism Central only)

VM

VM ↑

Selected: 1 of 1

MediaStorage

**Choose cluster**

Cluster


Cluster	CPU Usage	Memory Usage
Hector	33%	55%
Vegas	11%	26%
Achilles	11%	42%

Apply Cancel

# Step 5. Specify VM Name

[This step applies only if you have selected the **Restore to new location**, or with **different settings** option at the **Restore Mode** step of the wizard]

At the **Name** step of the wizard, you can specify a new name for the recovered VM.



## Step 6. Select Storage Container

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Storage Container** step of the wizard, choose the storage container where virtual disks of the recovered VM will be stored.

### NOTE

You cannot choose a storage container when restoring the VM from a snapshot.

For a container to be displayed in the list of the available containers, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).

The screenshot shows the Veeam Backup for Nutanix AHV interface during a 'Full VM Restore' operation. The 'Storage Container' step is active, and a modal window titled 'Select Storage Container' is displayed. The modal shows a table of available storage containers with columns for Name, Free Space, and Total Capacity. The 'Nutanix\_NTNX\_ctr' container is selected.

Name	Free Space	Total Capacity
NutanixManagementShare	10.6 TB	10.69 TB
Veeam_SetupVm_Container_00...	10.6 TB	10.6 TB
Backup_Storage_Container	10.6 TB	10.61 TB
Veeam_SetupVm_Db_Storage_...	10.6 TB	10.6 TB
Nutanix_AT-NTNX_ctr	10.6 TB	10.8 TB
<b>Nutanix_NTNX_ctr</b>	<b>10.6 TB</b>	<b>11.5 TB</b>
objectsm60e6eaf4ea7b4fe3530...	10.6 TB	10.6 TB
objectsd60e6eaf4ea7b4fe3530...	10.6 TB	10.61 TB

# Step 7. Configure Network Settings

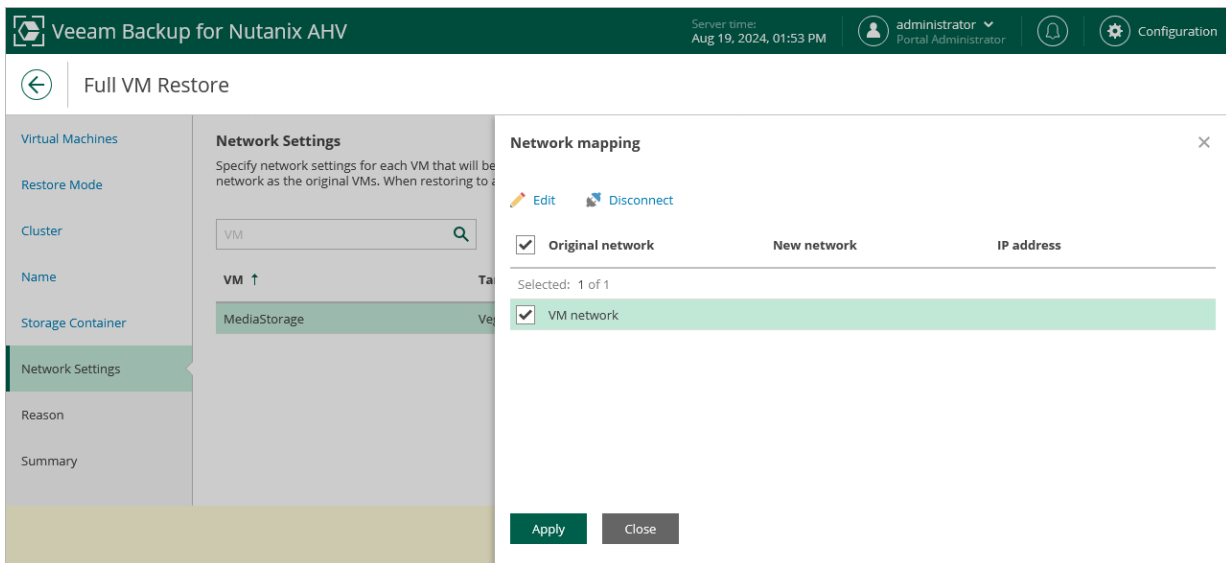
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network Settings** step of the wizard, choose a network to which the recovered VM will be connected. If you do not want to connect the VM to any virtual network, click **Disconnect**.

## NOTE

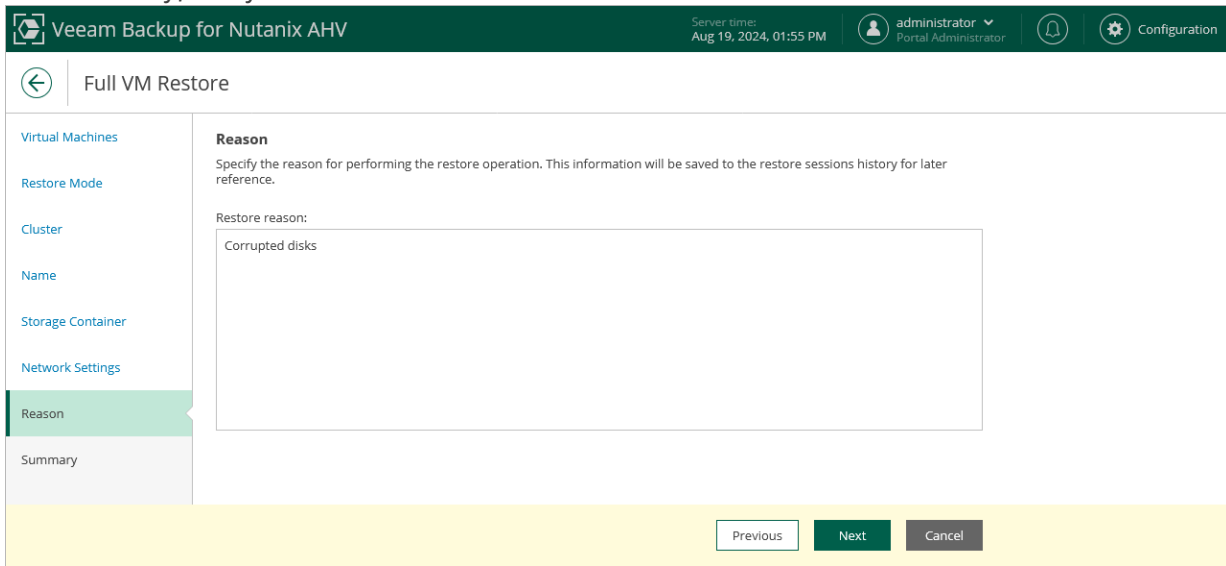
You cannot change network settings when restoring the VM from a PD snapshot or a snapshot created in the Nutanix AHV Prism console.

For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).



# Step 8. Specify Reason for Restore

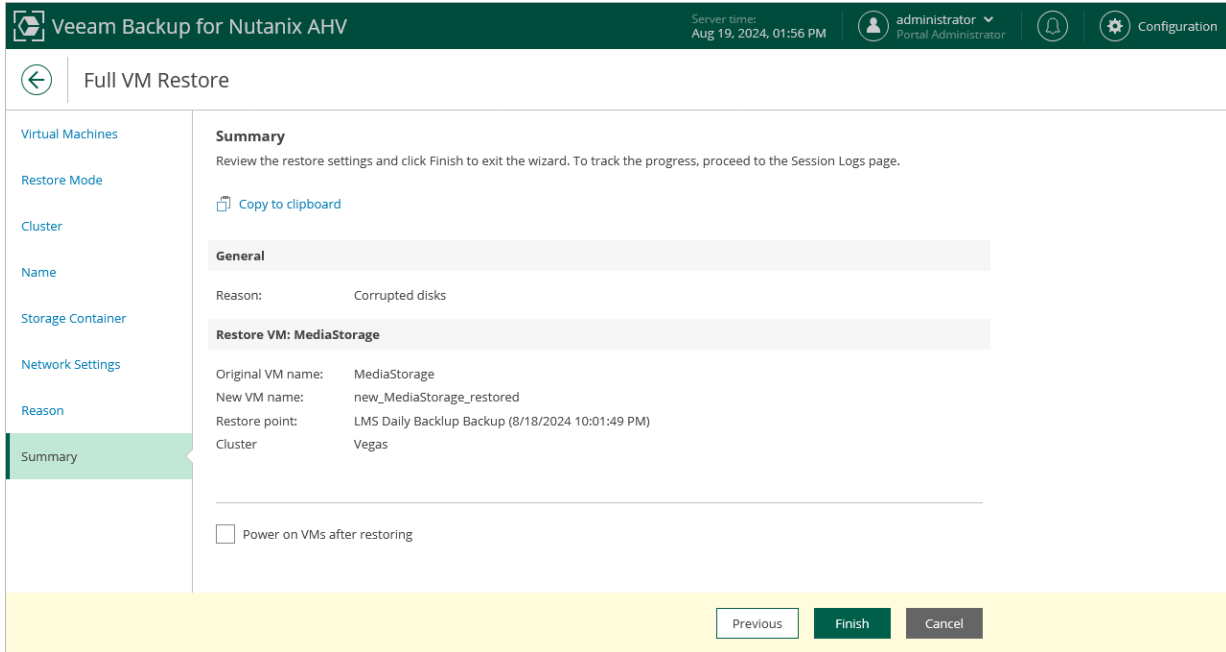
At the **Reason** step of the wizard, specify a reason for restoring the VM. This information will be saved to the session history, and you will be able to reference it later.



# Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

If you want to start the recovered VM as soon as the restore process completes, select the **Power on target VM after restoring** check box.





# Performing Disk Restore

In case a disaster strikes, you can restore disks of a Nutanix AHV VM from a backup or backup snapshot. Veeam Backup for Nutanix AHV allows you to attach the restored disks to the original VM or any other VM in the virtual infrastructure.

## NOTE

You cannot restore disks of volume groups attached to the VM.

## How to Perform Disk Restore


To restore disks attached to a protected VM, do the following:

1. [Launch the Virtual Disk Restore wizard.](#)
2. [Select a restore point.](#)
3. [Choose a restore mode.](#)
4. [Specify a target cluster.](#)
5. [Configure mapping settings.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)

# Step 1. Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** wizard, do the following:

1. Navigate to **Protected VMs**.
2. Select the VM whose virtual disks you want to restore.
3. Click **Restore > Disk Restore**.



## Step 2. Select Restore Point

At the **Virtual Machine** step of the wizard, select a restore point that will be used to restore data. By default, Veeam Backup for Nutanix AHV uses the most recent valid restore point. However, you can restore the data to an earlier state.

### NOTE

While creating a backup, Veeam Backup for Nutanix AHV takes a VM snapshot that is called backup snapshot. Veeam Backup for Nutanix AHV stores a recent backup snapshot for each backup job. Restore from a backup snapshot is significantly faster than restore from a backup.

To select a restore point, do the following:

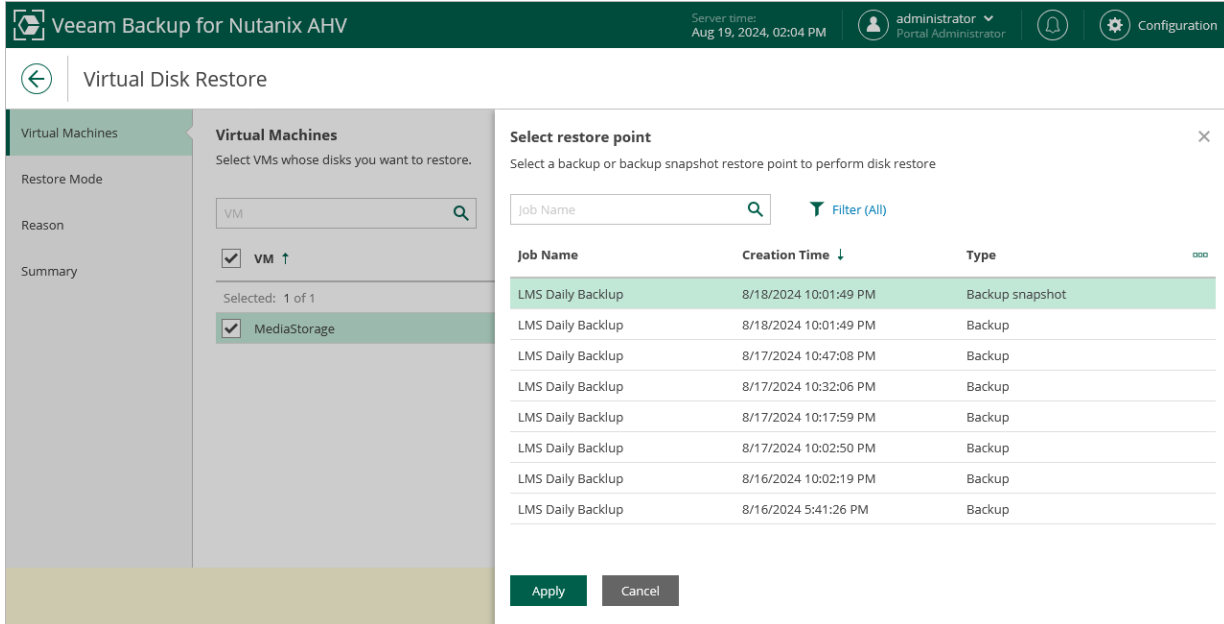
1. Select the VM.
2. Click **Restore Point**.
3. In the **Select restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Nutanix AHV provides the following information on each available restore point:

- **Job Name** – the name of the backup job that created the restore point.
- **Creating Time** – the time and date when the restore point was created.
- **Type** – the type of the restore point:
  - *Backup* – an image-level backup created by a backup job.  
[Applies to the [Prism Central deployment](#)] Only backups are supported for restore to another cluster.
  - *Backup snapshot* – a snapshot created by a backup job.

## TIP

By default, Veeam Backup for Nutanix AHV restores all disks attached to the selected VMs. If you want to exclude specific disks of a VM from restore, do the following. Select a **VM**, click **Exclusions** and select the disks to exclude.



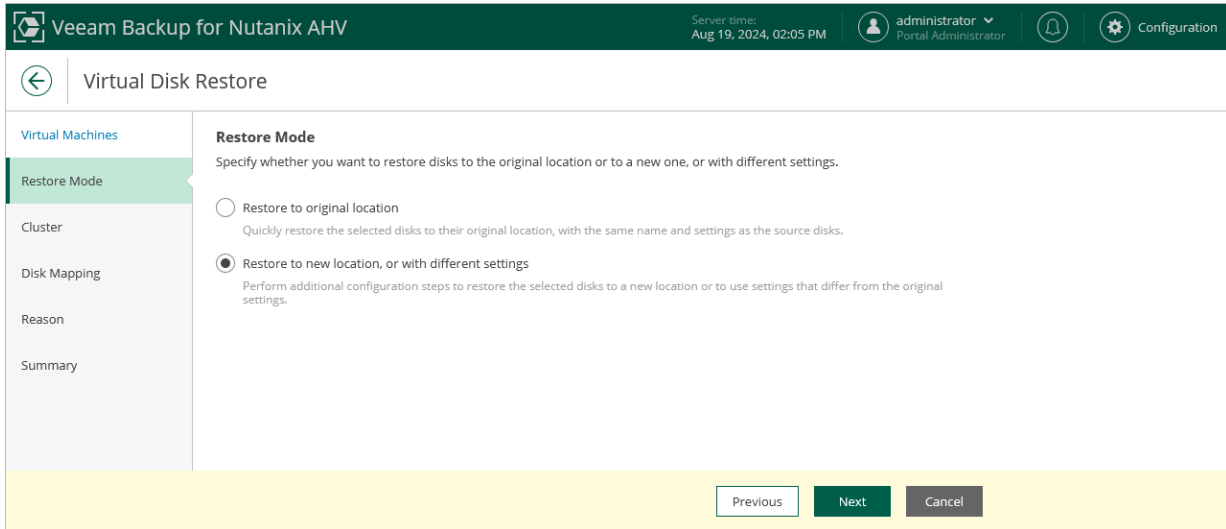
The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for Nutanix AHV', the server time 'Aug 19, 2024, 02:04 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The main window is titled 'Virtual Disk Restore'. On the left, a sidebar shows 'Virtual Machines' selected, with a search bar and a list of selected VMs: 'VM' and 'MediaStorage'. The main area displays a 'Select restore point' dialog with a search bar for 'Job Name' and a 'Filter (All)' button. Below this is a table of backup jobs:

Job Name	Creation Time ↓	Type
LMS Daily Backup	8/18/2024 10:01:49 PM	Backup snapshot
LMS Daily Backup	8/18/2024 10:01:49 PM	Backup
LMS Daily Backup	8/17/2024 10:47:08 PM	Backup
LMS Daily Backup	8/17/2024 10:32:06 PM	Backup
LMS Daily Backup	8/17/2024 10:17:59 PM	Backup
LMS Daily Backup	8/17/2024 10:02:50 PM	Backup
LMS Daily Backup	8/16/2024 10:02:19 PM	Backup
LMS Daily Backup	8/16/2024 5:41:26 PM	Backup

At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

# Step 3. Choose Restore Mode

At the **Restore Mode** step, choose whether you want to restore the disks with the original settings or to specify new settings (such as target VM and disk settings).



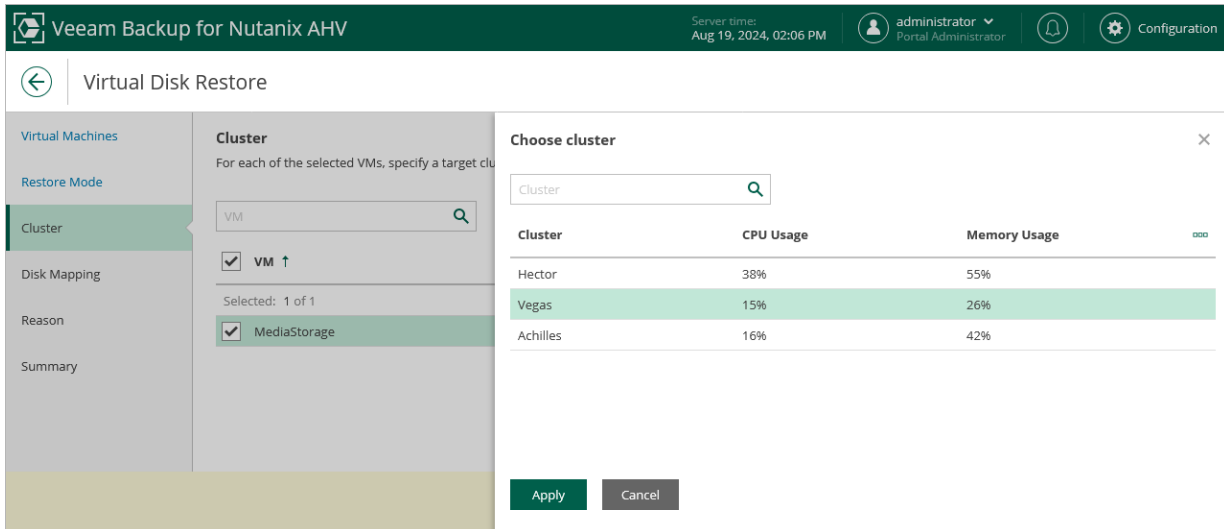
# Step 4. Specify Target Cluster

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard in the Prism Central deployment]

At the **Cluster** step of the wizard, choose the cluster where the disks will be restored.

## NOTE

The **Cluster** step of the **Virtual Disk Restore** wizard is only available when you restore the disks from a VM backup.



The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name, the server time (Aug 19, 2024, 02:06 PM), the user (administrator), and a Configuration icon. The main window is titled 'Virtual Disk Restore' and has a sidebar with navigation options: Virtual Machines, Restore Mode, Cluster (selected), Disk Mapping, Reason, and Summary. The 'Cluster' step is active, showing a search box for VMs and a list of selected items: VM and MediaStorage. A 'Choose cluster' dialog is open, displaying a table of clusters with CPU and Memory usage. The 'Vegas' cluster is highlighted in green.

Cluster	CPU Usage	Memory Usage
Hector	38%	55%
Vegas	15%	26%
Achilles	16%	42%

# Step 5. Configure Mapping Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Disk Mapping** step of the wizard, do the following:

1. Click **Change Disk Mappings**.
2. Choose a target VM to which you want to attach the restored disks.

By default, Veeam Backup for Nutanix AHV attaches the restored disks to the original VM. To attach the disks to another VM, click **Browse**.

## IMPORTANT

During disk restore, Veeam Backup for Nutanix AHV turns off the target VM to reconfigure its settings and attach the restored disk. It is recommended that you stop all activities on the target VM till the restore session completes.

3. Select a virtual disk to restore and click **Edit**.

By default, Veeam Backup for Nutanix AHV attaches the restored disk to the target VM as a new disk. However, if you want the restored disk to replace the existing disk, or if you want to change the disk bus type and to specify a storage container for the restored disk, configure disk settings.

## NOTE

You can select a storage container only if you restore from a backup.

The screenshot shows the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name, server time (Aug 19, 2024, 02:10 PM), user (administrator), and configuration settings. The main window is titled "Virtual Disk Restore" and is divided into several sections:

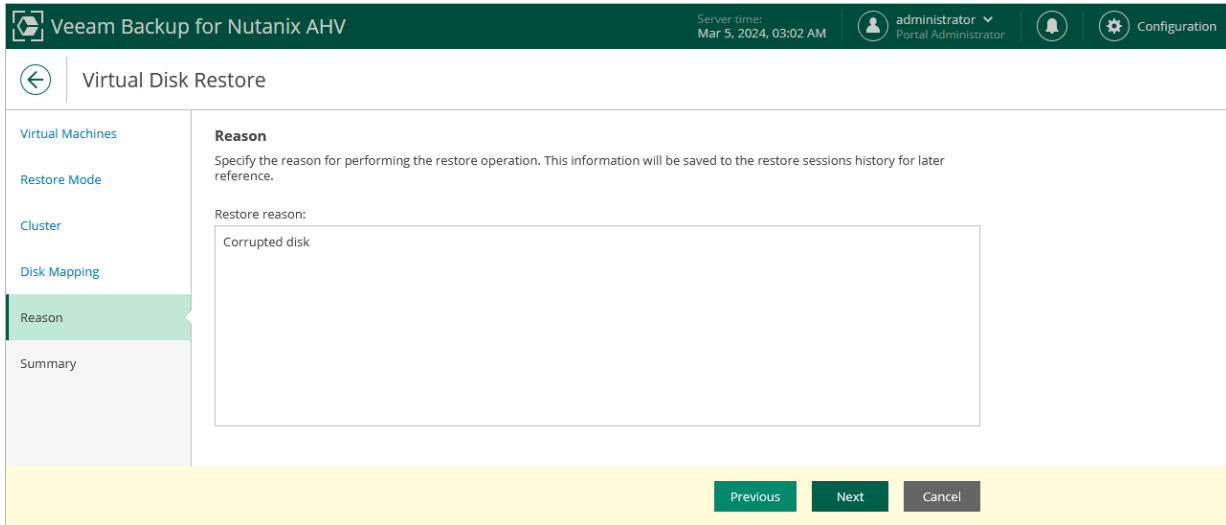
- Virtual Machines:** A sidebar menu with options like "Restore Mode", "Cluster", "Disk Mapping" (selected), "Reason", and "Summary".
- Disk Mapping:** A central panel with a search bar for "VM" and a table showing the mapping of virtual disks. The table has columns for "VM" and "Disk". One entry is visible: "MediaStorage" mapped to "1".
- Target VM:** A modal window showing the configuration for the target VM. It includes a search bar for "Virtual machine" and a dropdown menu showing "phtest". Below this, there is an "Edit" button and a table for "Disk properties".

Original Disk N...	Original Contai...	New Container	Disk Size	New Disk Node
SCSI.0	default-container...	Nutanix_NTNX_ctr	1000 GB	SCSI.0

At the bottom of the "Target VM" modal, there are "Apply" and "Cancel" buttons.

# Step 6. Specify Reason for Restore

At the **Reason** step of the wizard, specify a reason for restoring disks. This information will be saved to the session history, and you will be able to reference it later.





# Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

## TIP

If you want to start the recovered VM as soon as the restore process completes, select the **Power on VM after restore** check box.

The screenshot shows the 'Virtual Disk Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Aug 19, 2024, 02:11 PM), user (administrator), and configuration options. A left sidebar lists navigation options: Virtual Machines, Restore Mode, Cluster, Disk Mapping, Reason, and Summary (which is highlighted). The main content area displays the following information:

- Summary:** Review the restore settings, and click Finish to exit the wizard. A 'Copy to clipboard' button is available.
- General:** Reason: Corrupted disk.
- Restore disks of VM: MediaStorage:**
  - Original VM name: MediaStorage
  - Target VM name: phtest
  - Restore point: LMS Daily Backup Backup (8/18/2024 10:01:49 PM)
  - Cluster: Vegas
  - Disks info: From: SCSI.0, default-container-32073350623819 To: SCSI.0, Nutanix\_NTNX\_ctr

At the bottom, there is a checked checkbox for 'Power on VMs after restore' and three buttons: 'Previous', 'Finish', and 'Cancel'.

# Instant Recovery

With Instant Recovery, you can immediately restore Nutanix AHV VMs as VMware vSphere, Microsoft Hyper-V or Nutanix AHV VMs to your production environment by running them directly from their backups. Instant Recovery helps you improve recovery time objectives and minimize disruption and downtime of production workloads.

# Performing Instant Recovery of Workloads to Nutanix AHV

You can immediately restore virtual or physical machines into a Nutanix AHV cluster by running it directly from a compressed and deduplicated backup file. Before you perform Instant Recovery, check the following prerequisites:

- The Nutanix AHV cluster runs Nutanix AOS 6.0 or later.
- The Nutanix AHV cluster is [added to the backup infrastructure](#).

## Supported Workloads

To recover machines to a Nutanix AHV cluster, you can use the following backups:

- Backups of Nutanix AHV VMs created by Veeam Backup for Nutanix AHV
- Backups of Microsoft Hyper-V and VMware vSphere VMs created by Veeam Backup & Replication
- Backups of virtual and physical machines created by Veeam Agent for Microsoft Windows and Veeam Agent for Linux
- Backups of VMs created by vCloud Director
- Backups of Amazon EC2 instances created by Veeam Backup for AWS
- Backups of Microsoft Azure VMs created by Veeam Backup for Microsoft Azure
- Backups of Google Cloud VMs instances created by Veeam Backup for Google Cloud
- Backups of oVirt KVM VMs created by Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups of Proxmox VE VMs created by Veeam Backup for Proxmox VE

Instant Recovery is not supported:

- From backups of VMs with the ARM CPU architecture
- From file-level backups created by Kasten 10, Veeam Agent for Linux, Veeam Agent for Microsoft Windows, Veeam Agent for Unix, Veeam Agent for Mac

### NOTE

Instant Recovery to a Nutanix AHV cluster is supported only for backups stored in backup repositories, object storage repositories, external repositories, [Veeam Cloud Connect repositories](#), [HPE Cloud Bank Storage](#) and a scale-out backup repository (performance, capacity or archive tier). Instant Recovery from backups stored on tapes is not supported.

## How Instant Recovery Works

When Instant Recovery is performed, Veeam Backup for Nutanix AHV mounts a workload image to a [mount server](#) directly from a compressed and deduplicated backup file. Since there is no need to extract the workload from the backup file and copy it to production storage, you can perform recovery from any restore point in a matter of minutes.

The workload image remains in the read-only state to avoid unexpected modifications. By default, all changes to virtual disks that take place while the recovered workload is running are logged to auxiliary redo log files residing in the Nutanix AHV cluster. These changes are either merged if you choose to migrate the workload to the production environment, or discarded if you choose to revert the recovery operation.

## How to Perform Instant Recovery to AHV

To perform Instant Recovery of a protected workload, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Instant Recovery wizard.](#)
3. [Choose a restore point.](#)
4. [Choose a restore mode.](#)
5. [Select a target cluster.](#)
6. [Select a target storage container.](#)
7. [Specify a name for the restored workload.](#)
8. [Configure network settings.](#)
9. [Specify a restore reason.](#)
10. [Review the configured settings.](#)
11. [Finalize the recovery process.](#)

## Before You Begin

Before you perform Instant Recovery, do the following:

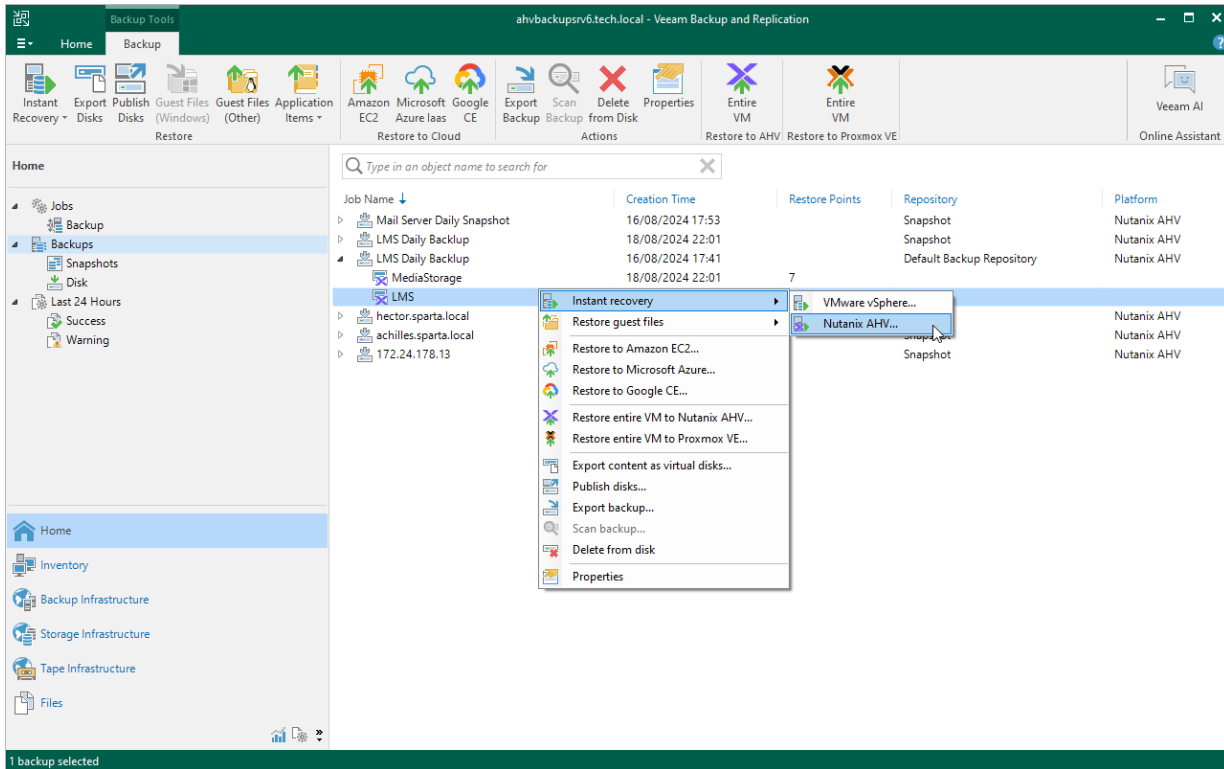
- Power off the original machine if it is still present in the target location.
- Deploy a [dedicated server](#) to mount workload images directly from backups stored in backup repositories and allocate minimum 512 MB of additional RAM for each VM disk that you want to recover. Make sure that the *Server for NFS* role and the *Client for NFS* component are not installed on the server, and that the [Veeam vPower NFS Service](#) is running.
- [Applies only to VMs being restored from backups stored in the archive tier of scale-out backup repositories] Retrieve backup data as described in the Veeam Backup & Replication User Guide, section [Retrieving Backup Files](#). However, this requirement is not applicable to backups stored in the archive tier that consists of the Amazon S3 Glacier Instant Retrieval extent.
- [Applies only to Linux VMs] Make sure that the file systems (also referred to as devices or partitions) listed in the */etc/fstab* file are mounted using UUIDs. Instant Recovery of file systems mounted using device names is not supported as the restored VMs may fail to boot.
- [Applies only to Windows VMs being restored from backups created by solutions other than Veeam Backup for Nutanix AHV] Make sure to install Nutanix VirtIO drivers and Nutanix Guest Tools on the VMs – before the backups are created. You will not be able to add or modify the VM drivers during the recovery operation.

- [Applies only to VMs being restored from backups created by solutions other than Veeam Backup for Nutanix AHV] Veeam Backup for Nutanix AHV attaches VM disks with the restored data to the target VM disk nodes using their original bus types. Veeam Backup for Nutanix AHV can attach to a VM up to 6 SATA, 256 SCSI, 4 IDE and 7 PCI disks. If the VM has more disks of any of those bus types, Nutanix AHV will attach the disks to remaining nodes of other bus types in the default priority: SATA, SCSI, IDE, PCI. You can [modify the backup appliance configuration](#), to instruct Nutanix AHV to ignore source VM original bus types and to use a specific order of bus types.

# Step 1. Launch Instant Recovery Wizard

To launch the **Instant Recovery to Nutanix AHV** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup, select the VM that you want to restore and click **Instant Recovery > Nutanix AHV** on the ribbon, or right-click the VM and select **Instant Recovery > Nutanix AHV**.



## Step 2. Select Restore Point

At the **Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Backup for Nutanix AHV uses the most recent valid restore point. However, you can restore the VM data to an earlier state.

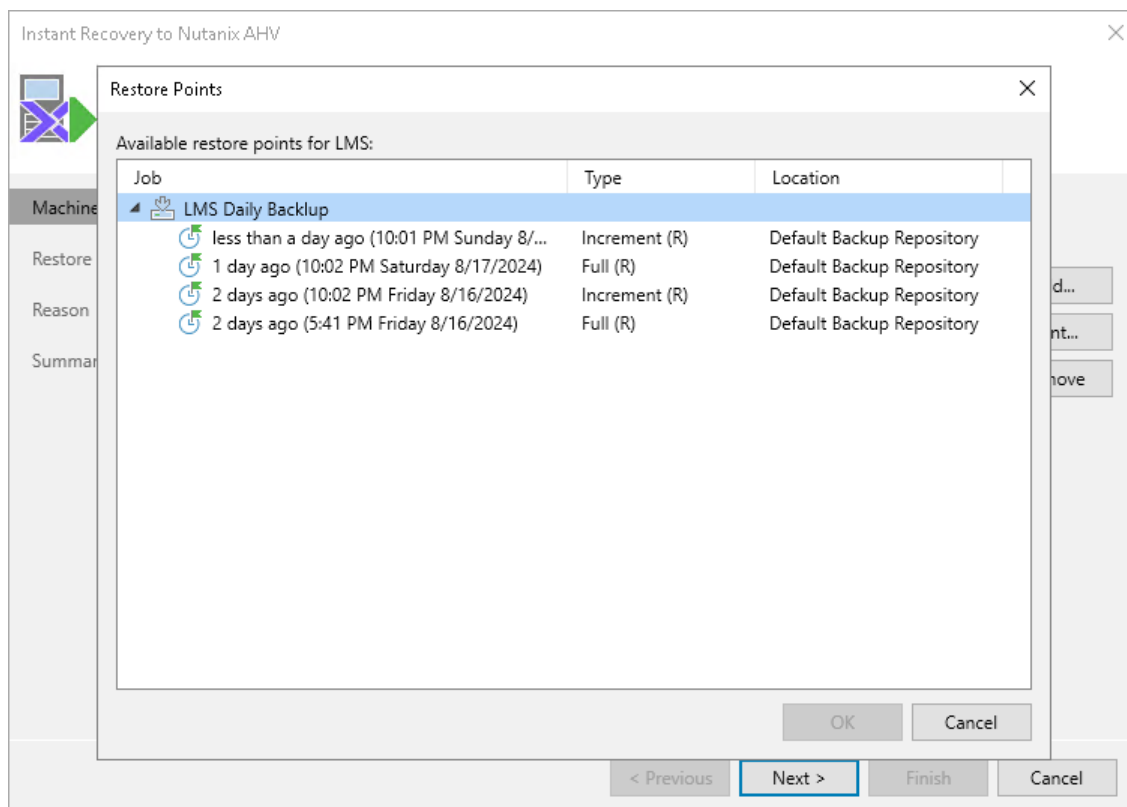
To select a restore point, do the following:

1. Select the VM.
2. Click **Point**.
3. In the **Restore Points** window, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup for Nutanix AHV provides the following information on each available restore point:

- **Job** – the name of the backup job that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the repository where the restore point is stored.

You can use the wizard to restore multiple VMs at a time. To do that, click **Add**, select more VMs to restore and select a restore point for each of them.



## Step 3. Choose Restore Mode

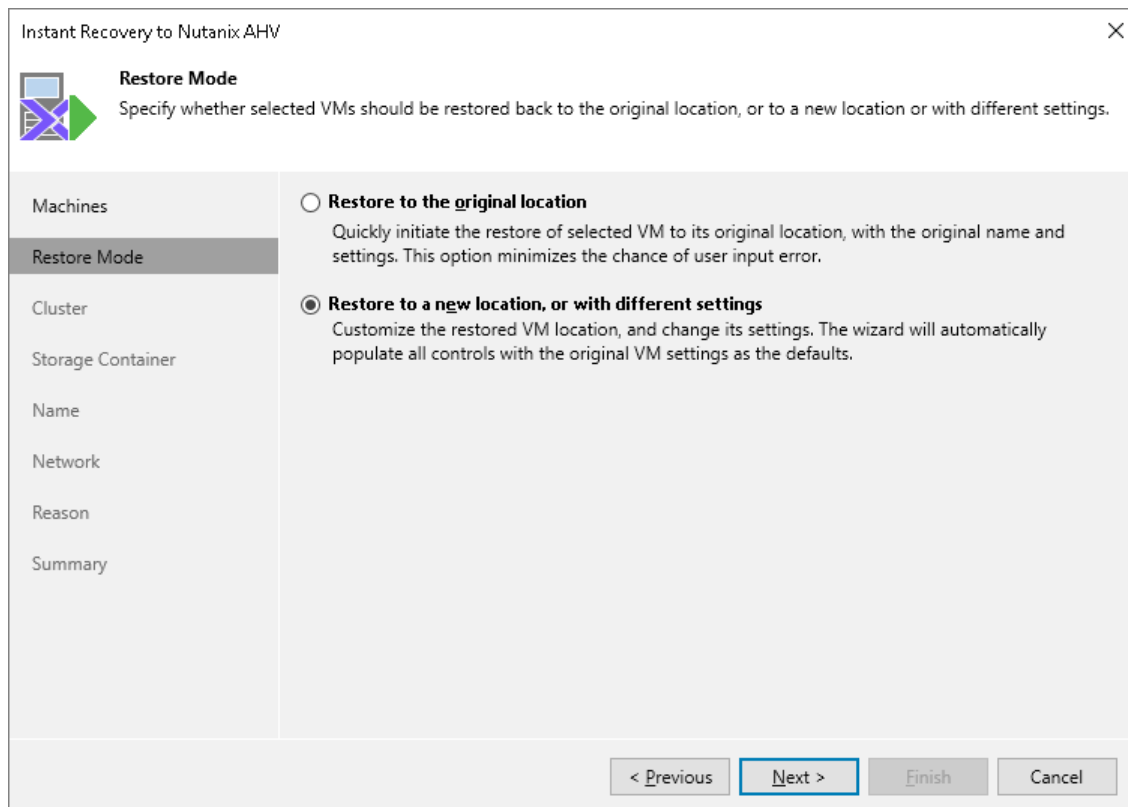
[This step applies only if you restore Nutanix AHV VMs]

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected VM to the original or to a custom location.

To meet minimum requirements for VMs residing on a Nutanix AHV cluster, Veeam Backup for Nutanix AHV allocates 64 MB of RAM to the recovered VM if it originally had less amount of memory.

### IMPORTANT

If you recover a VM with original settings, and the original VM still exists in the virtual infrastructure, the original VM will be removed.





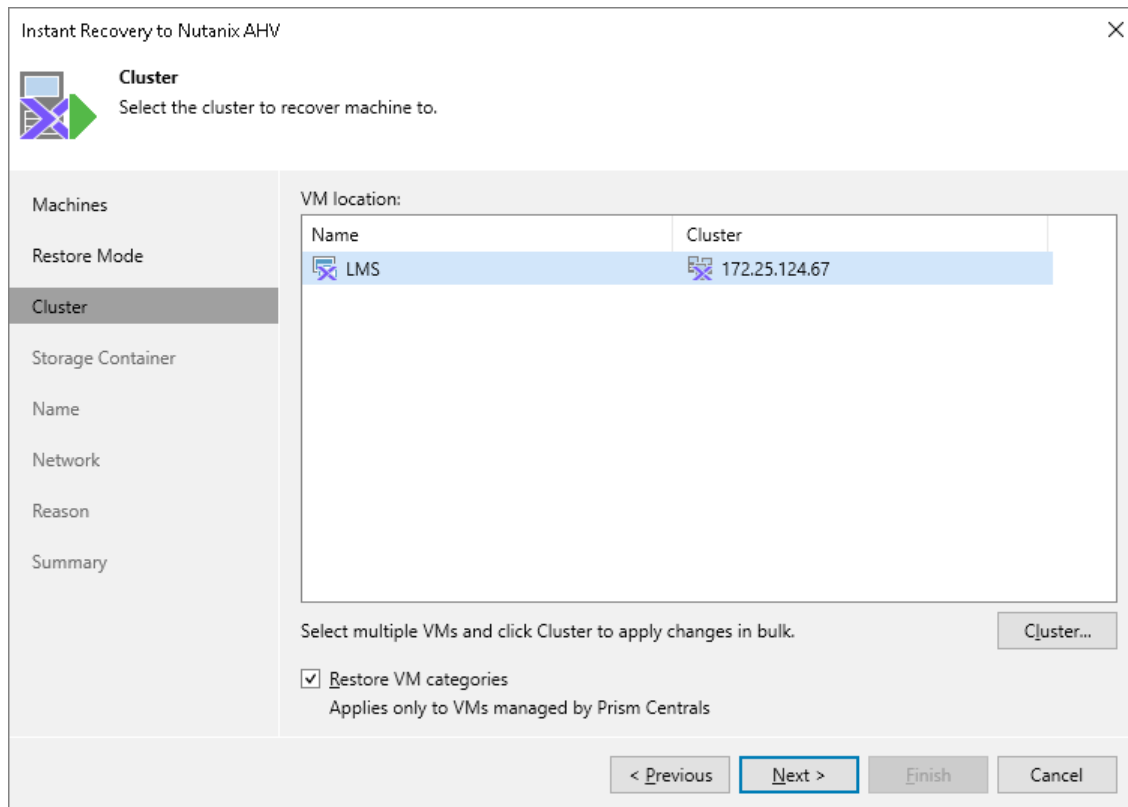
## Step 4. Specify Target Cluster

At the **Cluster** step of the wizard, choose the cluster to which the recovered VM will belong. In the Prism Central deployment, you can also choose whether you want the recovered VM to be assigned the same categories as the original VM.

For a cluster to be displayed in the list of the available clusters, it must be added to the backup infrastructure as described in section [Adding Nutanix AHV Server](#).

### IMPORTANT

If a selected VM has an attached volume group, the disks of the volume group will not be restored.



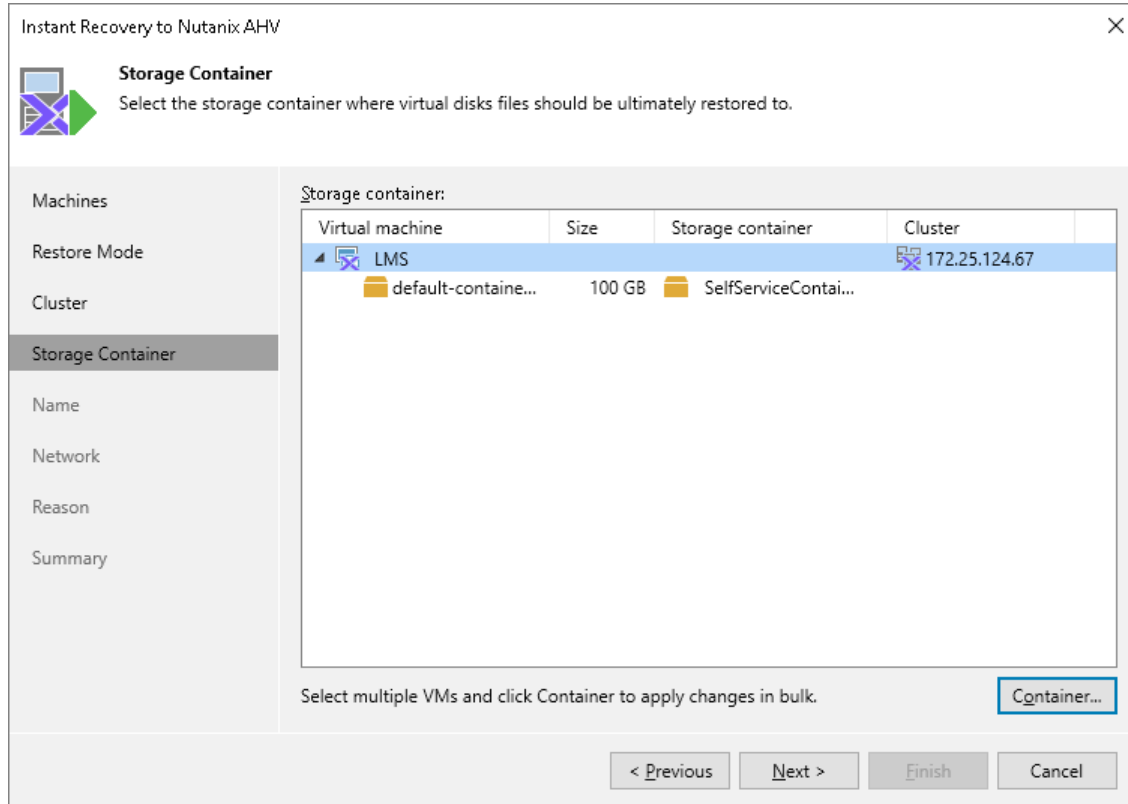
The screenshot shows the 'Instant Recovery to Nutanix AHV' wizard window. The title bar reads 'Instant Recovery to Nutanix AHV' with a close button (X) on the right. The main window has a sidebar on the left with the following items: 'Machines', 'Restore Mode', 'Cluster' (highlighted), 'Storage Container', 'Name', 'Network', 'Reason', and 'Summary'. The main area is titled 'Cluster' and contains the instruction 'Select the cluster to recover machine to.' Below this is a table with the heading 'VM location:'. The table has two columns: 'Name' and 'Cluster'. One row is highlighted in blue, showing 'LMS' in the 'Name' column and '172.25.124.67' in the 'Cluster' column. Below the table, there is a button labeled 'Cluster...'. Underneath the button is a checkbox labeled 'Restore VM categories' with the text 'Applies only to VMs managed by Prism Centrals' below it. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Name	Cluster
LMS	172.25.124.67

## Step 5. Select Storage Container

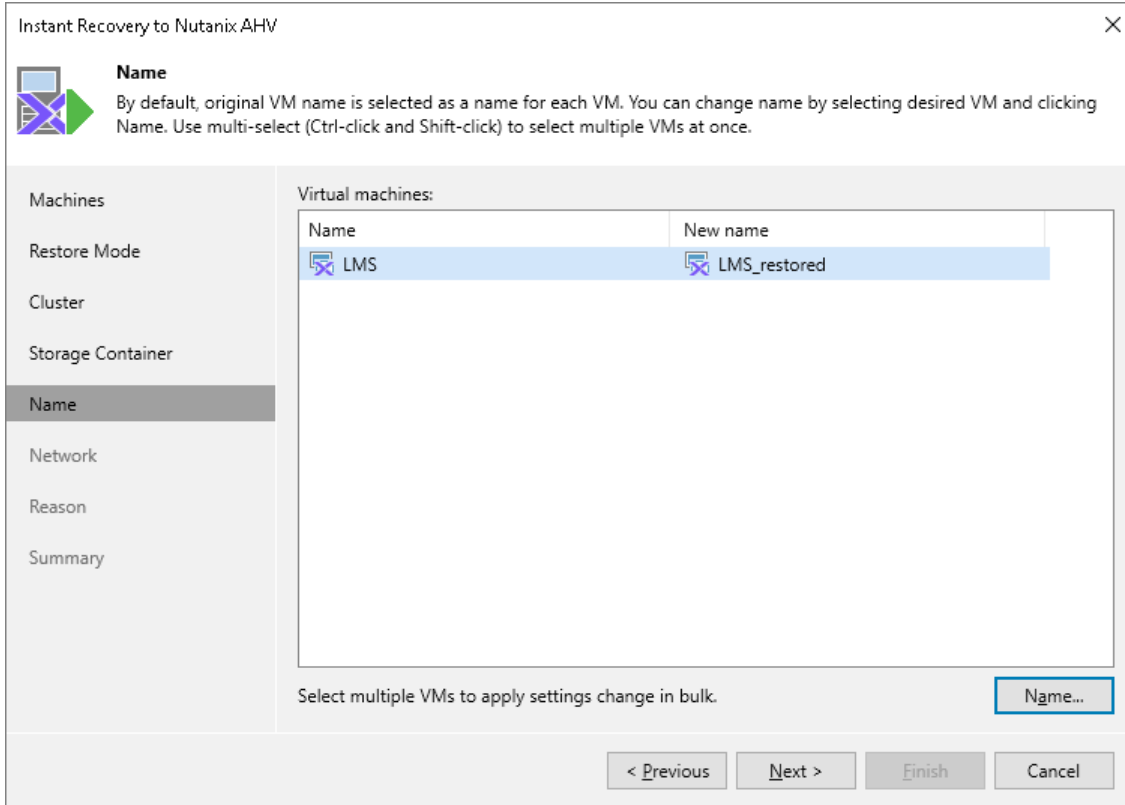
At the **Storage Container** step of the wizard, choose the storage container where virtual disks of the recovered VM will be stored.

For a container to be displayed in the list of the available containers, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).



## Step 6. Specify VM Name

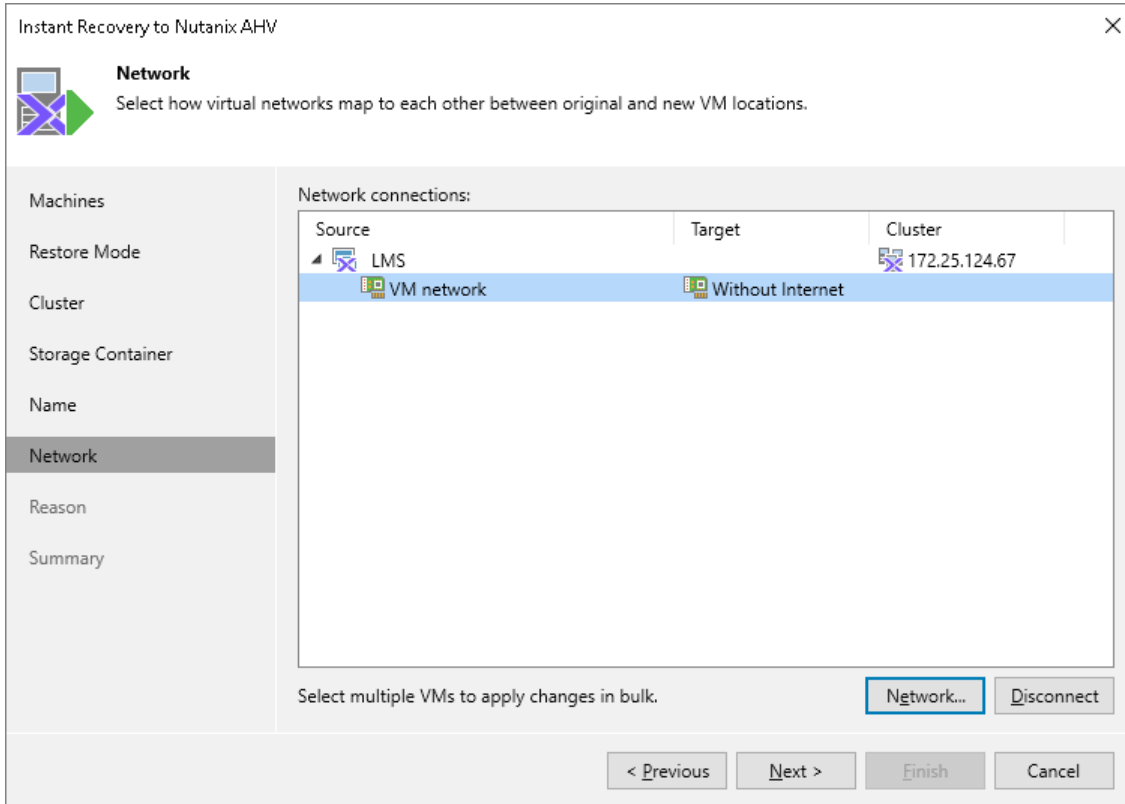
At the **Name** step of the wizard, you can specify a new name for the recovered VM.



# Step 7. Configure Network Settings

At the **Network** step of the wizard, choose a network to which the recovered VM will be connected. If you do not want to connect the VM to any virtual network, select the VM and click **Disconnect**.

For a network to be displayed in the list of the available networks, it must be configured in the Nutanix AHV cluster as described in [Nutanix documentation](#).



## Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM. This information will be saved to the session history, and you will be able to reference it later.

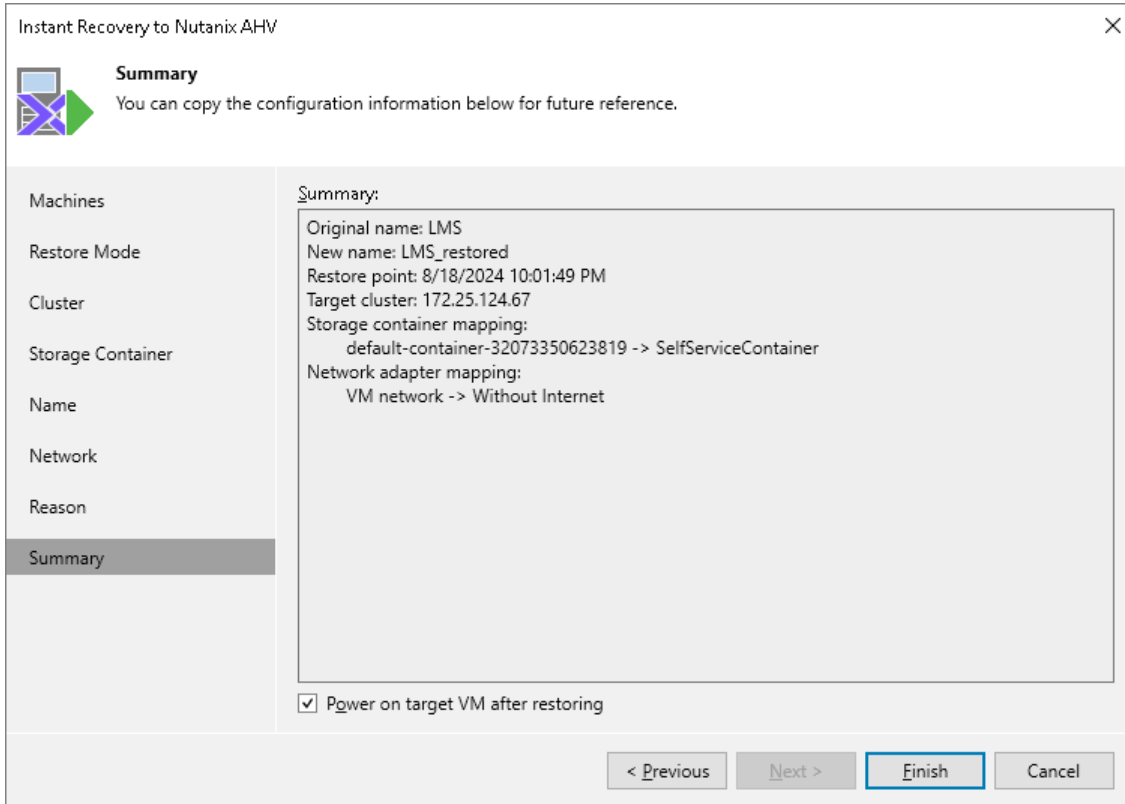
The screenshot shows a wizard window titled "Instant Recovery to Nutanix AHV" with a close button (X) in the top right corner. The window is divided into a left sidebar and a main content area. The sidebar contains a list of steps: Machines, Restore Mode, Cluster, Storage Container, Name, Network, Reason (highlighted), and Summary. The main content area has a heading "Reason" with a sub-heading "Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference." Below this is a text input field labeled "Restore reason:" containing the text "Corrupted data". At the bottom of the main content area is a checkbox labeled "Do not show me this page again". At the bottom of the window are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

## Step 9. Review Configured Settings

At the **Summary** step of the wizard, review summary information and click **Finish**.

### TIP

If you want to start the recovered VM as soon as the restore process completes, select the **Power on target VM after restoring** check box.



The screenshot shows the 'Instant Recovery to Nutanix AHV' wizard at the 'Summary' step. The window title is 'Instant Recovery to Nutanix AHV' with a close button (X) in the top right corner. Below the title bar, there is a 'Summary' section with a small icon of a computer and a green arrow, and the text 'You can copy the configuration information below for future reference.' A left-hand navigation pane lists several steps: 'Machines', 'Restore Mode', 'Cluster', 'Storage Container', 'Name', 'Network', 'Reason', and 'Summary'. The 'Summary' step is currently selected and highlighted. The main content area displays the following configuration details:

**Summary:**  
Original name: LMS  
New name: LMS\_restored  
Restore point: 8/18/2024 10:01:49 PM  
Target cluster: 172.25.124.67  
Storage container mapping:  
default-container-32073350623819 -> SelfServiceContainer  
Network adapter mapping:  
VM network -> Without Internet

At the bottom of the main content area, there is a checked checkbox labeled 'Power on target VM after restoring'. At the very bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

## Step 10. Finalize Instant Recovery

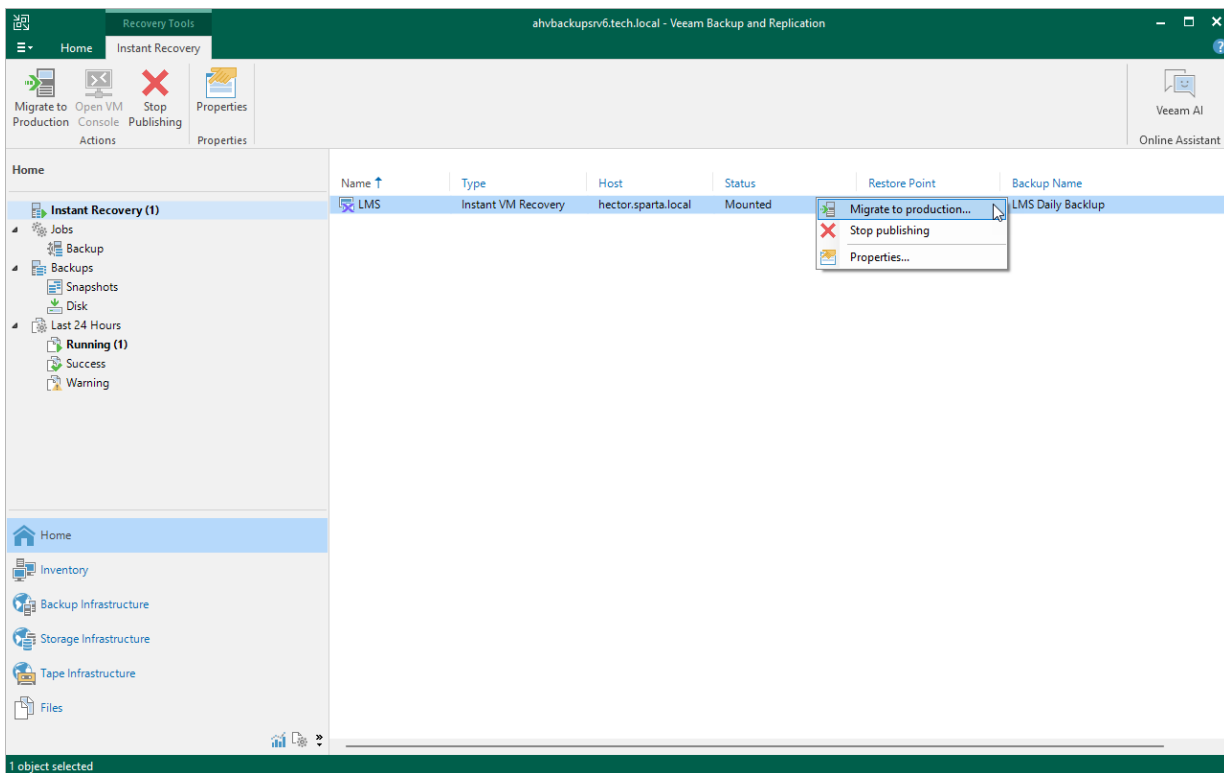
After the VM has been recovered, you can choose whether you want to migrate the VM to the production environment or cancel the recovery operation. When migrating VMs, Veeam Backup for Nutanix AHV transfers VM disk data to the production storage that you have selected as a destination for the recovered VM.

To finalize the instant recovery operation, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Instant Recovery**.
3. In the working area, right-click a VM:
  - To transfer VM disk data to the production storage, select **Migrate to production**.
  - To remove the recovered VM, select **Stop publishing**.

### IMPORTANT

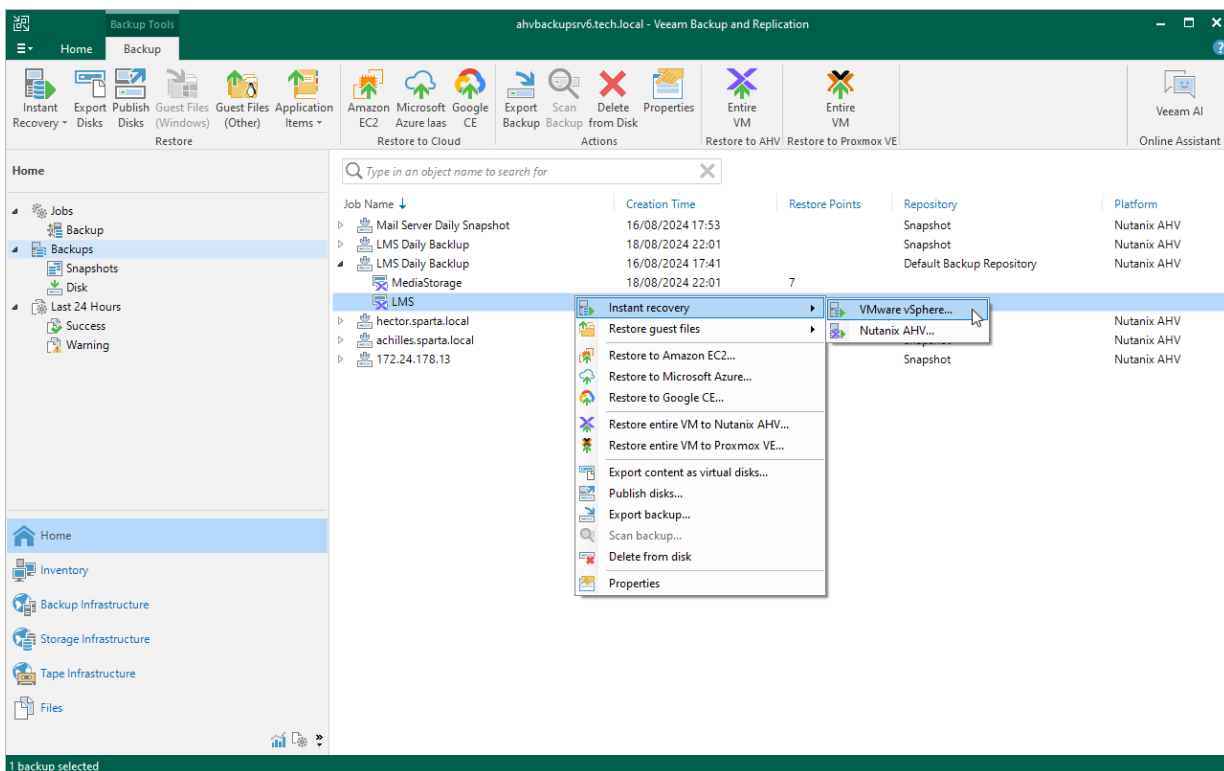
If you stop publishing a VM that was recovered to the same destination where the original VM resided, both the original and recovered VMs will be removed.



# Performing Instant Recovery of Workloads to VMware vSphere

To perform Instant Recovery to VMware vSphere environment, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the VM you want to restore, and select **Instant recovery**.
4. Complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for VMware vSphere, section [Performing Instant VM Recovery of Workloads to VMware vSphere VMs](#).

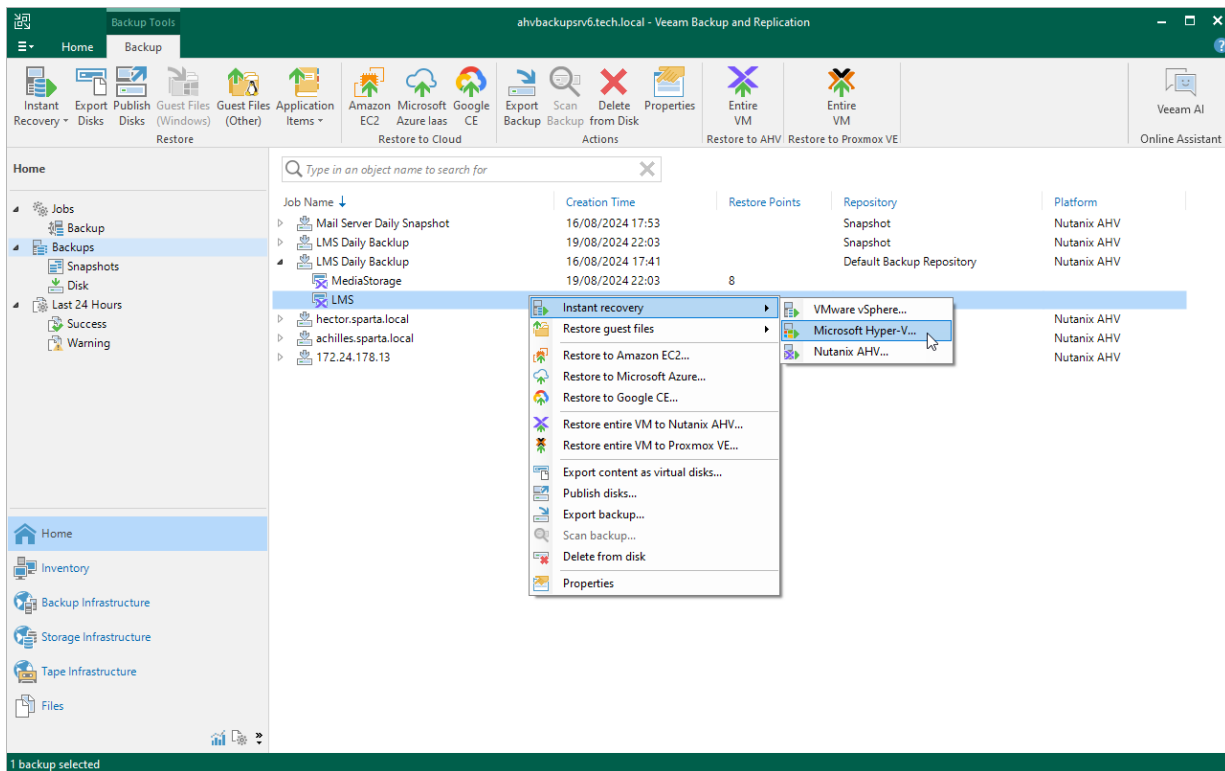




# Performing Instant Recovery of Workloads to Hyper-V

To perform Instant Recovery to Microsoft Hyper-V environment, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the VM you want to restore, and select **Instant recovery**.
4. Complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for Microsoft Hyper-V, section [Performing Instant VM Recovery of Workloads to Hyper-V VMs](#).

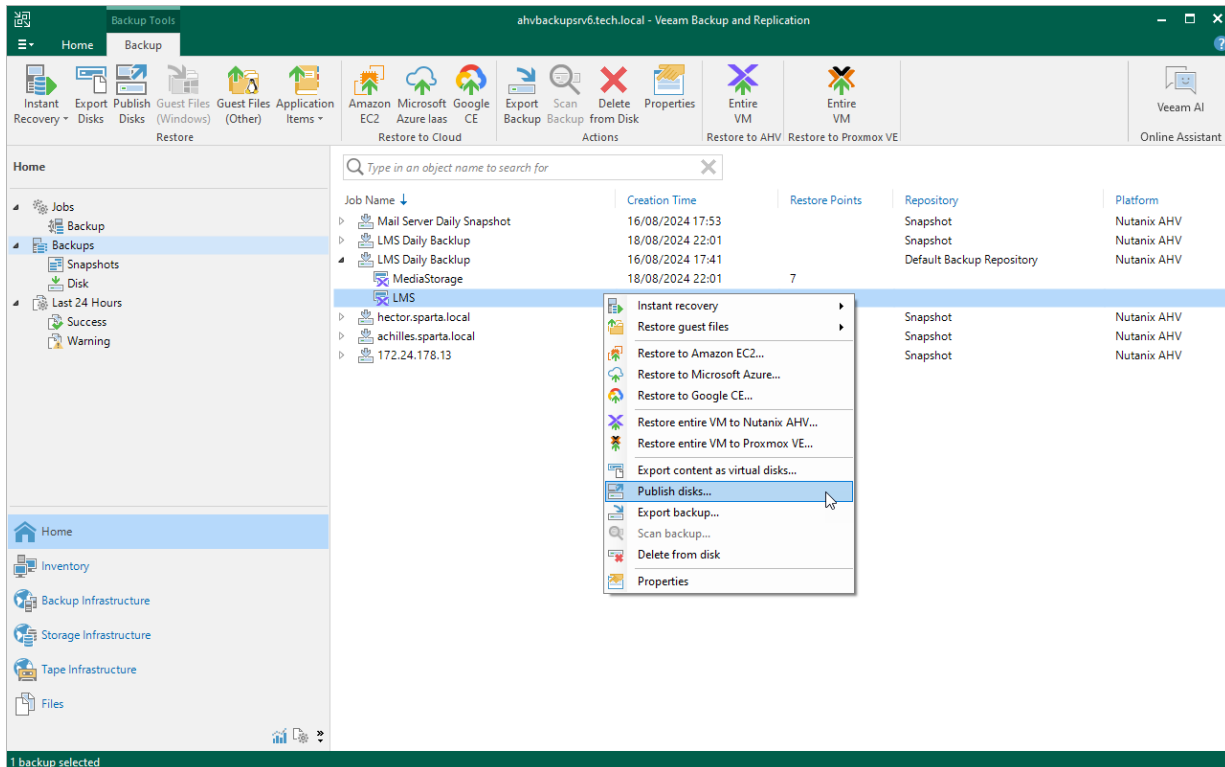


# Publishing Disks

Veeam Backup & Replication allows you to mount specific disks of backed-up Nutanix AHV VMs to any server and to instantly access data in the read-only mode. This can be helpful when you want to copy files and folders as of a point-in-time state to the target server, and perform an antivirus scan of the backed-up data. For more information, see the Veeam Backup & Replication User Guide, section [Disk Publishing \(Data Integration API\)](#).

To publish disks of a Nutanix AHV VM, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to mount and select **Publish disks**.
4. Complete the **Publish Disk** wizard as described in the Veeam Backup & Replication User Guide, section [Publishing Disks](#).



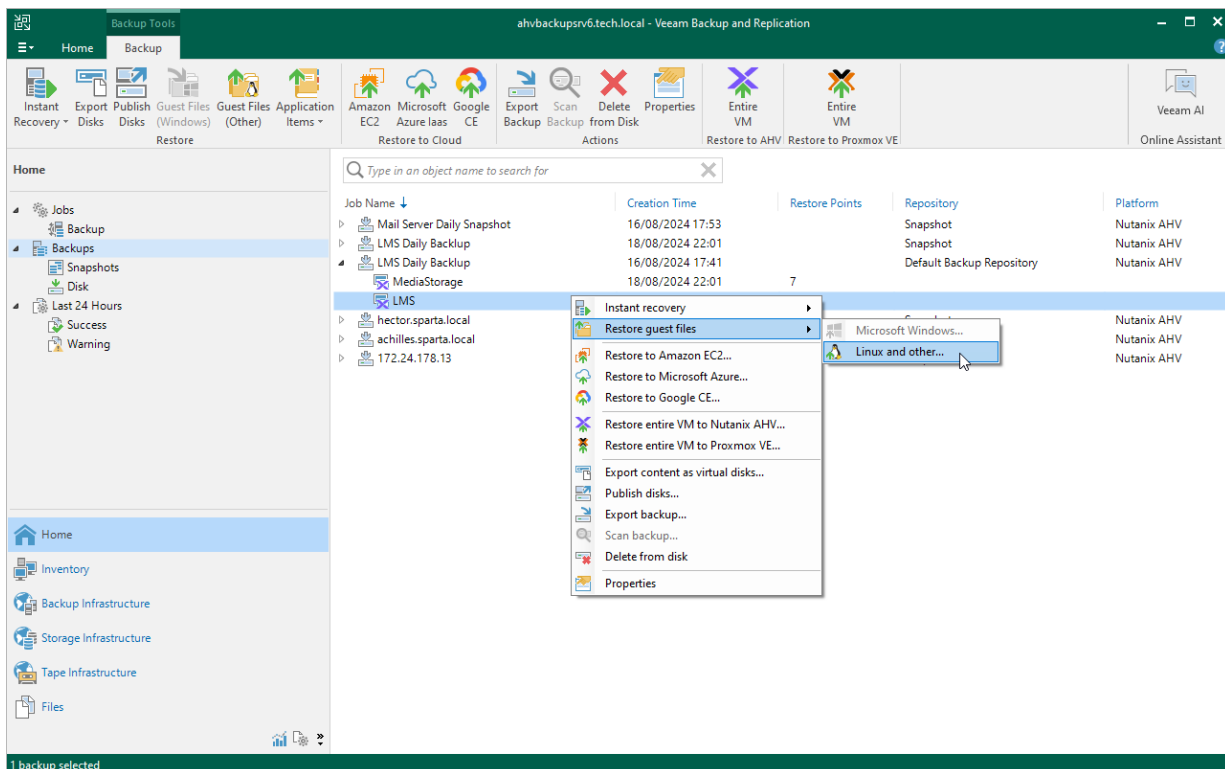
# Performing File-Level Restore

With guest OS file recovery (file-level restore), you can restore individual guest OS files and folders from Nutanix AHV VM snapshots and backups created with Veeam Backup for Nutanix AHV. When restoring files and folders, you do not need to extract the VM image to a staging location or start the VM prior to restore.

## How to Perform File-Level Restore

To restore VM guest OS files and folders, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup or snapshot job, right-click the VM that contains files you want to restore and do the following:
  - If you want to restore files of a Microsoft Windows machine, select **Restore guest files > Microsoft Windows** and complete the **Guest File Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(FAT, NTFS or ReFS\)](#).
  - If you want to restore files of a Linux, Solaris, BSD, Novell Storage Services, Unix or Mac machine, select **Restore guest files > Linux and others** and complete the **Guest File Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(Multi-OS\)](#).



### TIP

Alternatively, you can use Veeam Backup Enterprise Manager to restore guest OS files and folders as described in the Veeam Backup Enterprise Manager Guide, section [Restoring VM Guest OS Files](#).

# Performing Application Item Restore

With application item restore, you can use Nutanix AHV backups or snapshots to restore the following data:

- Microsoft Active Directory objects and containers
- Microsoft Exchange mailboxes, folders and messages
- Microsoft SharePoint sites and lists
- Microsoft SQL Server
- Oracle databases

## NOTE

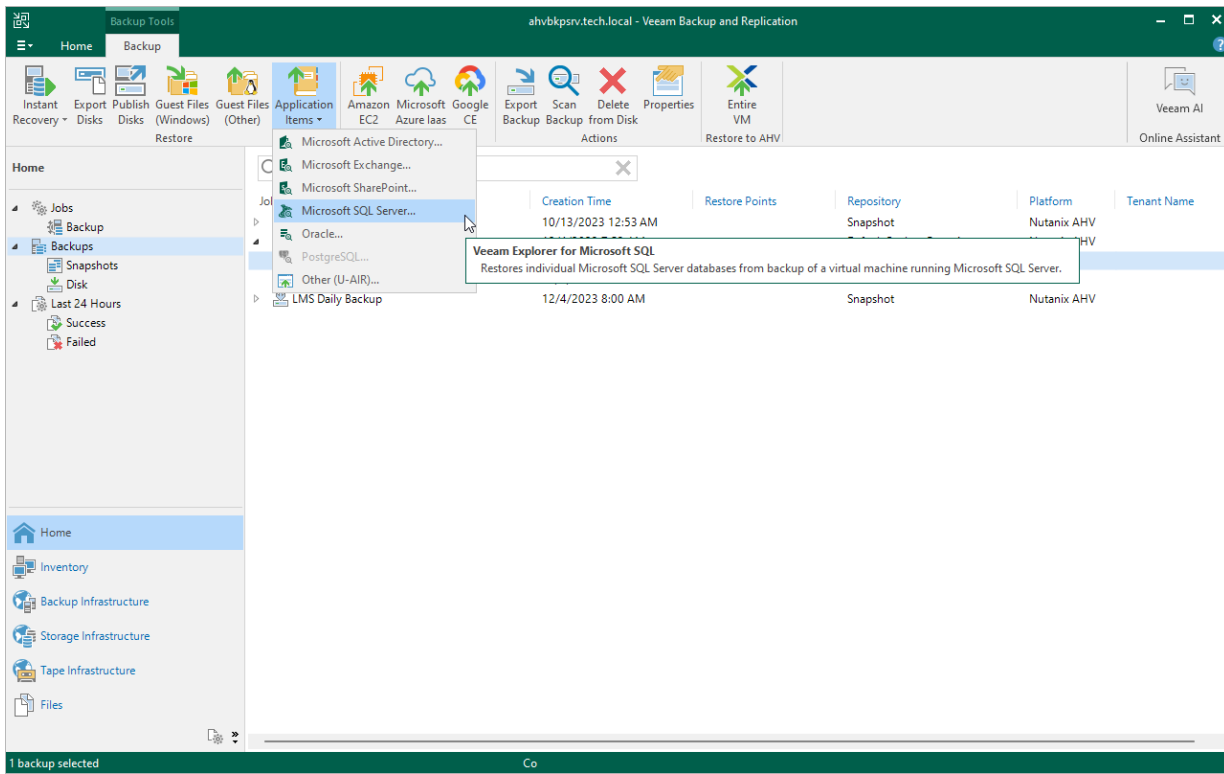
It is recommended that you use [application-consistent backups or snapshots](#) for application item restore. You can check whether a backup or snapshot is application-consistent in [sessions logs of the job](#) that has produced this backup or snapshot.

To restore application items from a Nutanix AHV VM backup or snapshot, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup or snapshot job, select the VM that contains an application you want to restore.
4. Click **Application Items** on the ribbon and then select the application.
5. In the restore wizard, select a restore point that will be used to restore the application, specify a restore reason and click **Browse**.
6. In the Veeam Explorer application, perform the steps described in the [Veeam Explorers User Guide](#).

## TIP

As an alternative to application item restore, you can also [perform file-level restore](#) to recover standalone databases using Veeam Explorers.

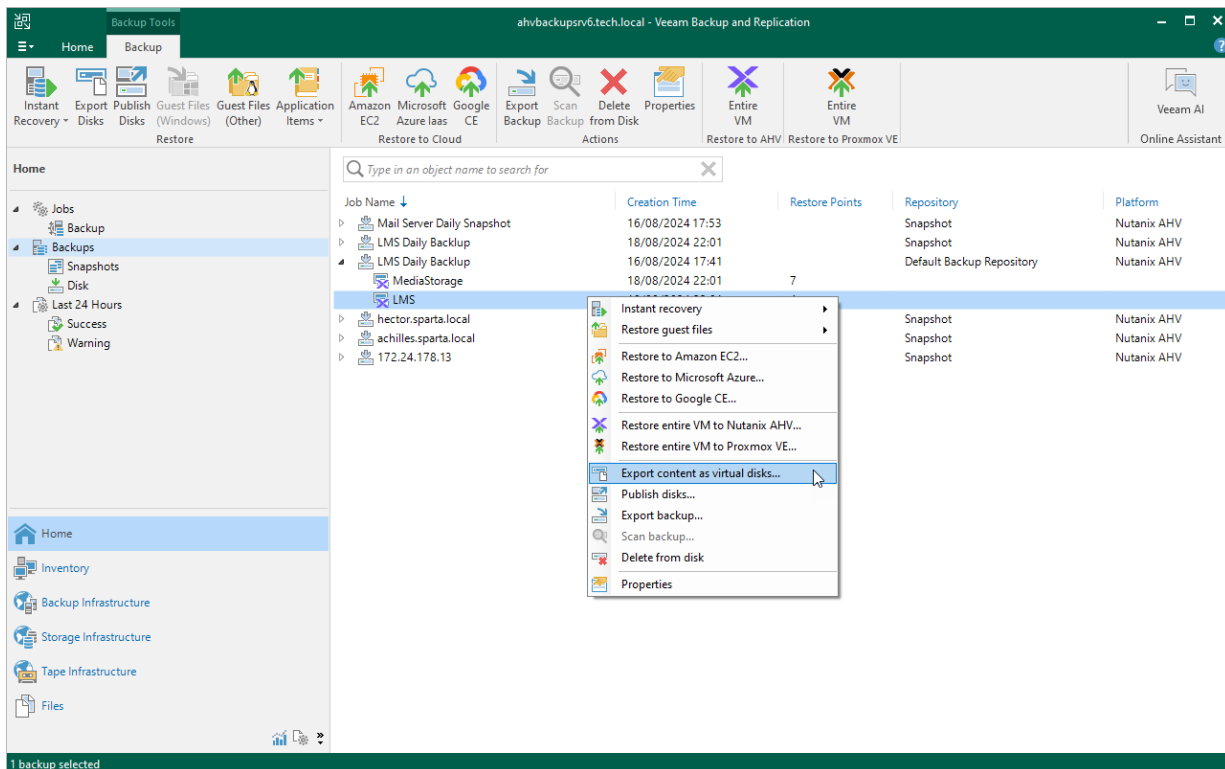


# Exporting Disks

Veeam Backup for Nutanix AHV allows you to export disks, that is, restore disks from Nutanix AHV VM backups and convert them to the VMDK, VHD and VHDX formats. You can save the exported disks to any server added to the backup infrastructure or place the disks on a datastore connected to an ESXi host (for the VMDK disk format only). For more information, see the Veeam Backup & Replication User Guide, section [Disk Export](#).

To export disks of a Nutanix AHV VM, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to export and select **Export content as virtual disks**.
4. Complete the **Export Disk** wizard as described in the Veeam Backup & Replication User Guide, section [Exporting Disks](#).

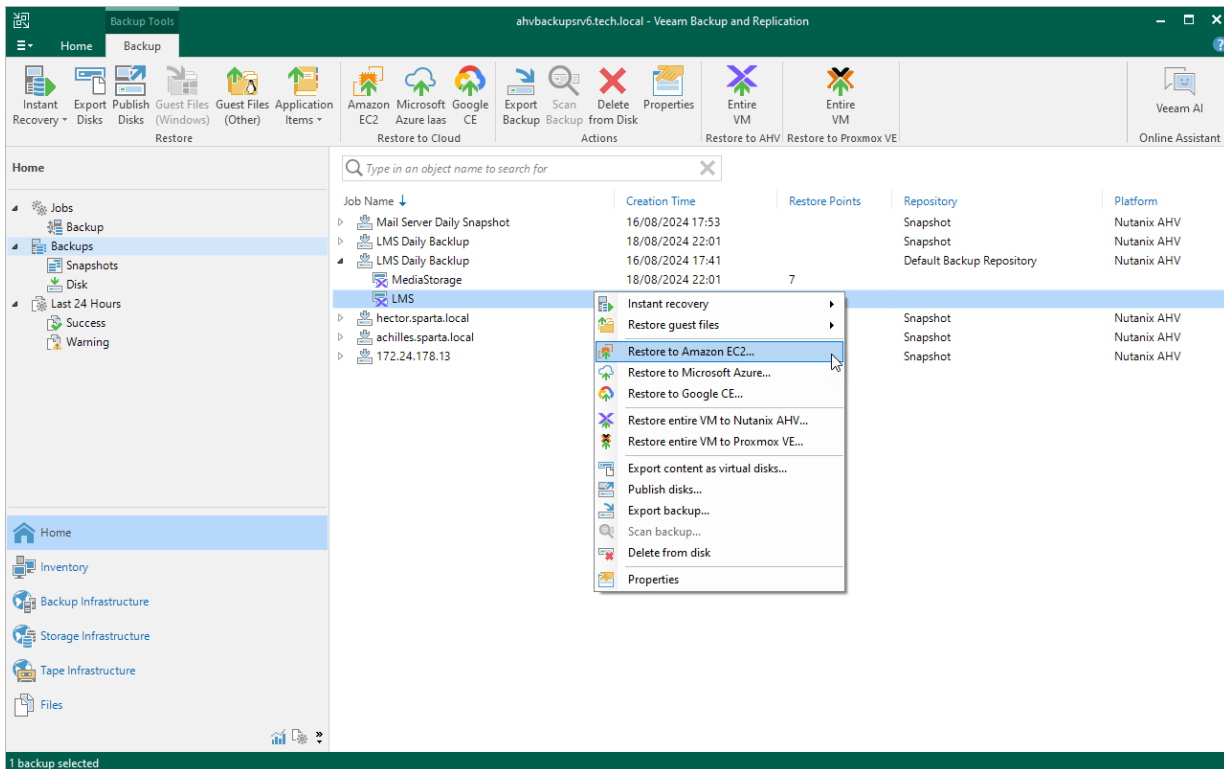


# Performing VM Restore to Amazon Web Services

Veeam Backup for Nutanix AHV allows you to restore Nutanix AHV VMs to Amazon Web Services (AWS) as EC2 instances. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Amazon EC2](#).

To restore a VM to Amazon EC2, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Amazon EC2**.
4. Complete the **Restore to Amazon EC2** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Amazon EC2](#).

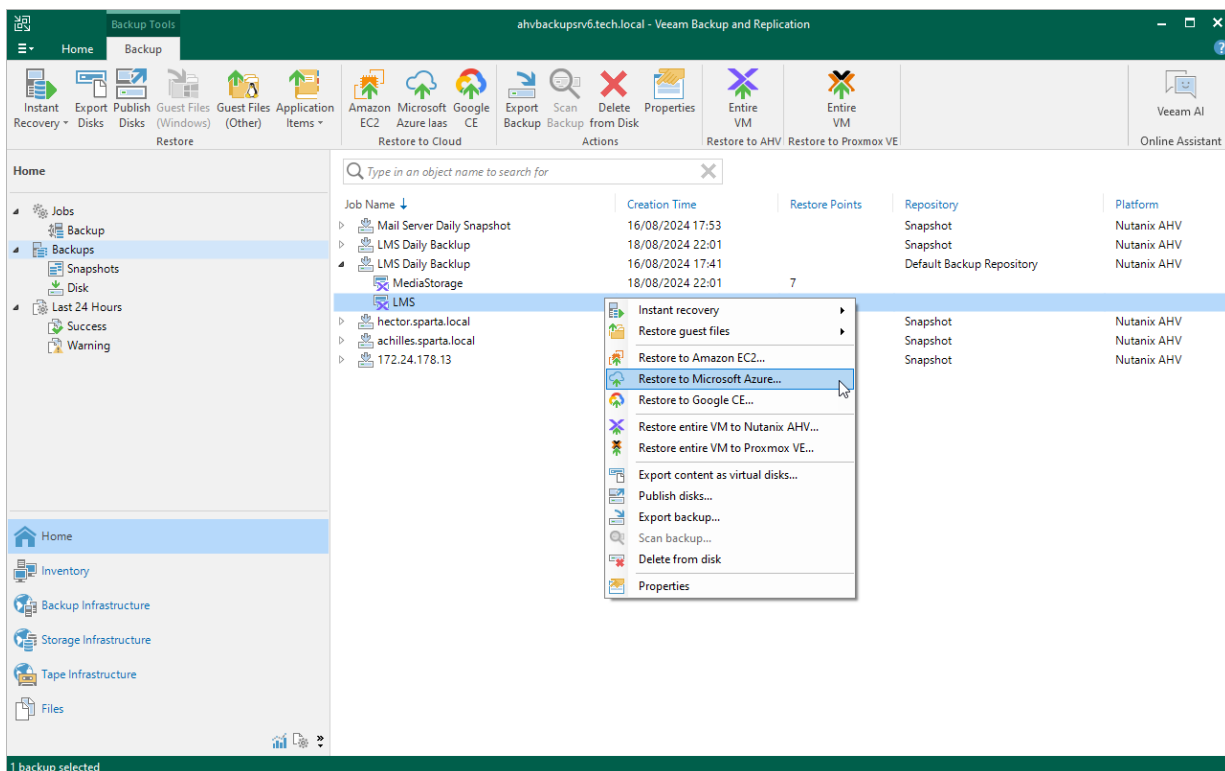


# Performing VM Restore to Microsoft Azure

Veeam Backup for Nutanix AHV allows you to restore Nutanix AHV VMs to Microsoft Azure as Azure VMs. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Microsoft Azure](#).

To restore a VM to Microsoft Azure, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Microsoft Azure**.
4. Complete the **Restore to Microsoft Azure** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Microsoft Azure](#).



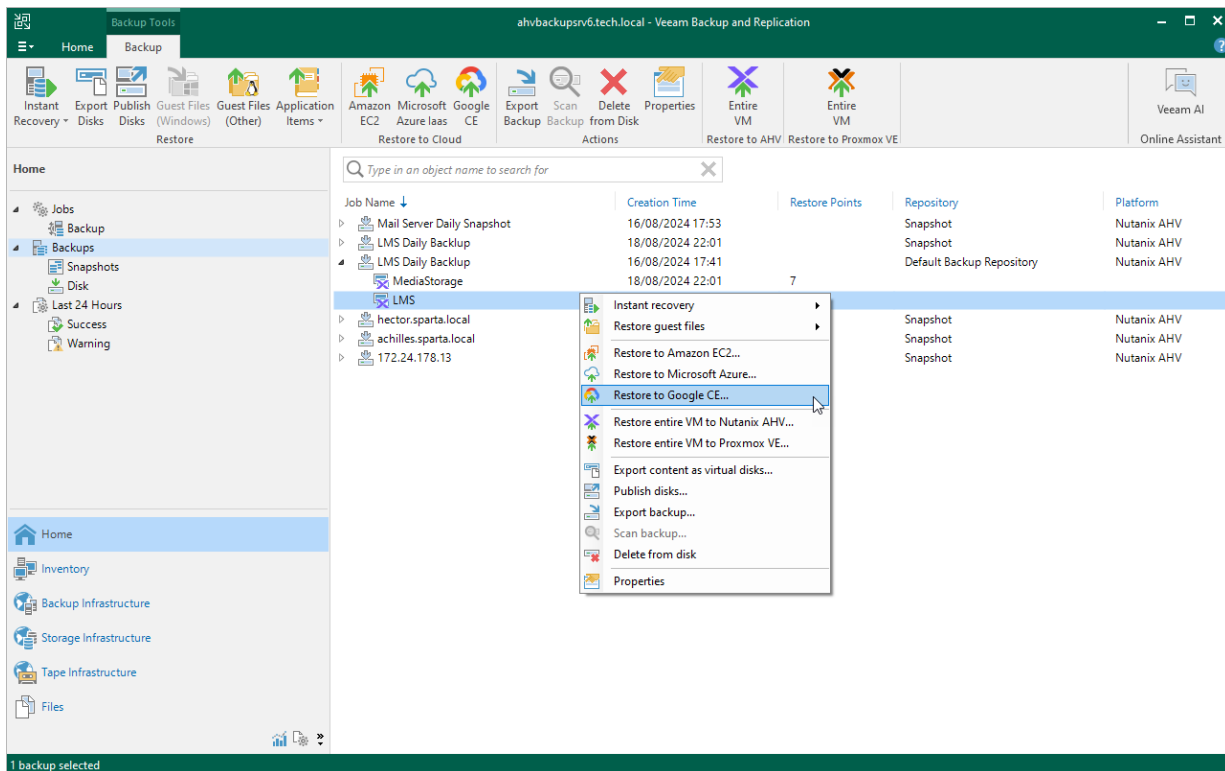


# Performing VM Restore to Google Cloud

Veeam Backup for Nutanix AHV allows you to restore Nutanix AHV VMs to Google Cloud as VM instances. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Google Compute Engine](#).

To restore a VM to Google Cloud, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Google CE**.
4. Complete the **Restore to Google Compute Engine** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Google Compute Engine](#).



# Viewing Session Statistics

For each performed data protection or disaster recovery operation, Veeam Backup for Nutanix AHV starts a new session and stores its records in the configuration database. You can track real-time statistics of all running and completed operations on the **Session Logs** page.

To view the full list of tasks executed during an operation, click the link in the **Status** column.

The screenshot displays the Veeam Backup for Nutanix AHV interface. The top navigation bar includes the Veeam logo, the product name, server time (Aug 19, 2024, 12:15 PM), user information (administrator, Portal Administrator), and a Configuration button. The left sidebar shows navigation options: Overview, Dashboard, Resources, Management, Jobs, Protected VMs, and Session Logs. The main content area shows a list of jobs with a search bar and filters. A modal window titled "Backup job: LMS Daily Backup" is open, displaying session status and logs.

**Backup job: LMS Daily Backup**

**Session Status**

Status	Start Time	End Time	Duration
Warning	8/18/2024 10:00:13 PM	8/18/2024 10:02:40 PM	2 min 26 sec

**Session Logs**

[Copy to Clipboard](#)

Start Time	Status	Description	Execution D.
Selected: 1 of 16			
8/18/2024 10:00:13 PM	Success	Job started at 8/18/2024 10:00:13 PM	—
8/18/2024 10:00:14 PM	Success	All objects have been queued for processing	2 sec
8/18/2024 10:00:14 PM	Success	VMs size: 1.1 TB	—
8/18/2024 10:00:17 PM	Warning	Backup location [Default Backup Repository] is getting low on ...	—
8/18/2024 10:00:17 PM	Success	The worker backup-appliance-hector was prepared successfully	1 min 22 sec
8/18/2024 10:01:39 PM	Success	MediaStorage : Use worker backup-appliance-hector	—
8/18/2024 10:01:39 PM	Success	MediaStorage : Processing	52 sec
8/18/2024 10:01:39 PM	Success	LMS : Use worker backup-appliance-hector	—
8/18/2024 10:01:39 PM	Success	LMS : Processing	53 sec

Close

# Reviewing Dashboard

Veeam Backup for Nutanix AHV comes with a dashboard that provides at-a-glance real-time overview of the protected Nutanix AHV resources and allows you to estimate the overall backup performance. The dashboard includes the following widgets:

- **Sessions in Last 24 Hours** – displays the number of sessions started for system operations as well as for data protection or disaster recovery tasks during the past 24 hours that completed successfully, the number of sessions that completed with warnings, the number of sessions that completed with errors, and the number of sessions that are currently running.

To get more information on the sessions, click either **View Session Logs** or any of the widget rows. In the latter case, the **Session Logs** page will show only those sessions that have the same status as that clicked in the widget.

For more information on the **Session Logs** page, see [Viewing Session Statistics](#).

- **Successful Task Ratio** – displays the number of job runs during a specific time period (the past 24 hours by default).

To specify the time period, click the link next to the **Schedule** icon. To get more information on the created snapshots, backups or archived backups, click any of the widget rows. In the latter case, the **Session Logs** page will show only those sessions during which Veeam Backup for Nutanix AHV created the same items as that clicked in the widget.

For more information on the **Session Logs** page, see [Viewing Session Statistics](#).

- **Protection Status** – displays the number of VMs that are included into snapshot and backup jobs or not included into any job. To get the list of VMs, click a widget row.
- **Backup Repositories** – displays repositories to which the Nutanix AHV backup appliance has access. The widget also shows the amount of storage space that is currently free in each repository. If less than 15% of storage space is free, a repository is **low on space**. If less than 5% of storage space is free, a repository is **out of space**.
- **Top Jobs** – shows top successfully completed jobs for execution time (including retries). For each job, the widget also calculates the growth rate to detect whether it took less or more time for the job to complete in comparison with the previous job run.
- **CPU and RAM Utilization** – is designed to help you avoid possible backup bottlenecks and allocate sufficient resources to the backup appliance.

The **CPU** widget analyzes the amount of backup appliance CPU resources used for processing jobs during the past 24 hours. The **RAM** widget shows memory consumption on the backup appliance during the past 24 hours.

To monitor the availability of compute resources on the backup appliance, you can instruct Veeam Backup for Nutanix AHV to send you notifications if free memory or available CPU drops below a configured threshold. For more information, see [Configuring Notifications](#).

The screenshot displays the Veeam Backup for Nutanix AHV dashboard. The top navigation bar shows the server time as Aug 27, 2024, 04:04 PM, and the user as administrator (Portal Administrator). The left sidebar contains navigation options: Overview, Dashboard, Resources, Management, Jobs, Protected VMs, and Session Logs.

**Sessions in Last 24 Hours**

Category	Count
Errors	0
Warnings	2
Success	2
Running	0

**Successful Task Ratio (Last 24 hours)**

Task Type	Count	Ratio
Backups	1 of 1	100%
Snapshots	1 of 1	100%
PD snapshots	1 of 1	100%

**CPU and RAM Utilization**

**CPU**  
Current load: 1%

**RAM**  
Current usage: 34%

**Protection Status**

- Included in snapshot jobs only: 1 VM
- Included in backup jobs: 2 VMs
- Not included in any job: 528 VMs

**Top Jobs (By duration decrease)**

Job	Duration	Start Time	Percentage
LMS Daily Backup	2 min 28 sec	8/26, 10:00 PM	-42%

**Backup Repositories (1)**

Repository	Status	Free	Capacity
Default Backup Repository	Low on space	1.7 GB	129.45 GB

# Updating Backup Appliance

Veeam Backup for Nutanix AHV allows you to check for available package updates, download and install them right from the Nutanix AHV backup appliance web console.

## NOTE

If the Nutanix AHV backup appliance is not connected to the internet, you can instruct the Veeam Updater service to use a web proxy as described in section [Configuring Web Proxy](#).

It is recommended that you timely install available updates to avoid issues while working with the product. For example, timely installed security updates may help you prevent potential security issues and reduce the risk of compromising sensitive data.

## TIP

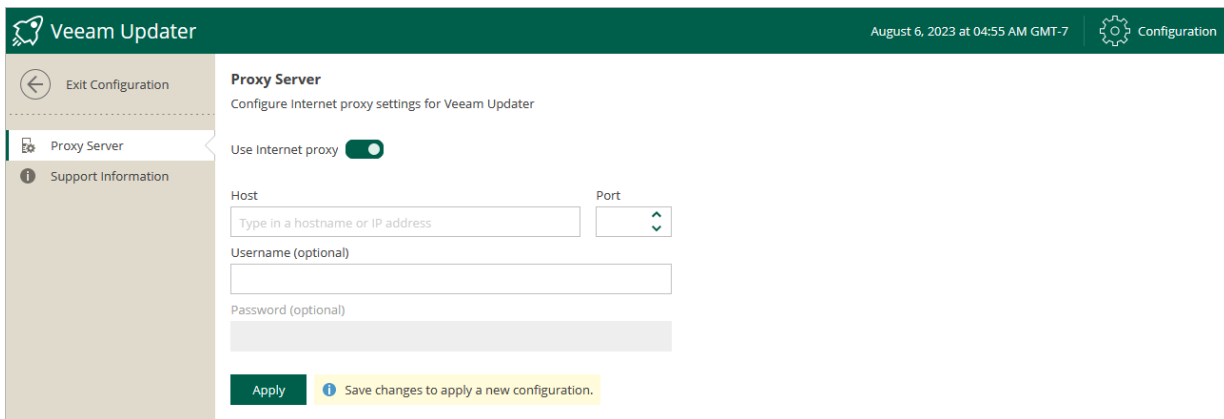
To check for major updates for Veeam Backup for Nutanix AHV, launch the Veeam Backup & Replication console. If a new version of the product is available, the Veeam Backup & Replication console will display the **Components Update** window and prompt you to upgrade the Nutanix AHV backup appliance. For more information, see the Veeam Backup & Replication User Guide, section [Server Components Upgrade](#).

# Configuring Web Proxy

To check for available package updates for the Nutanix AHV backup appliance and workers, the Veeam Updater service running on the backup appliance connects to Veeam repositories over the internet. If the backup appliance is not connected to the internet, you can instruct the Veeam Updater service to use a web proxy that will provide access to the required resources.

To configure connection to the internet through a web proxy, do the following:

1. Open the **Veeam Updater** page:
  - a. Switch to the **Configuration** page.
  - b. Navigate to **Support Information**.
  - c. On the **Updates** tab, click **Check and View Updates**.
2. On the **Veeam Updater** page, do the following:
  - a. At the top right corner, click **Configuration**.
  - b. Navigate to **Proxy Server**.
  - c. Set the **Use Internet proxy** toggle to *On*.
  - d. In the **Host** field, enter the IP address or FQDN of the web proxy.
  - e. In the **Port** field, enter the port used on the web proxy for HTTP or HTTPS connections.
  - f. [Applies only if the web proxy requires authentication] In the **Username** and **Password** fields, enter credentials of the account configured on the web proxy to access the internet.
  - g. Click **Apply**.

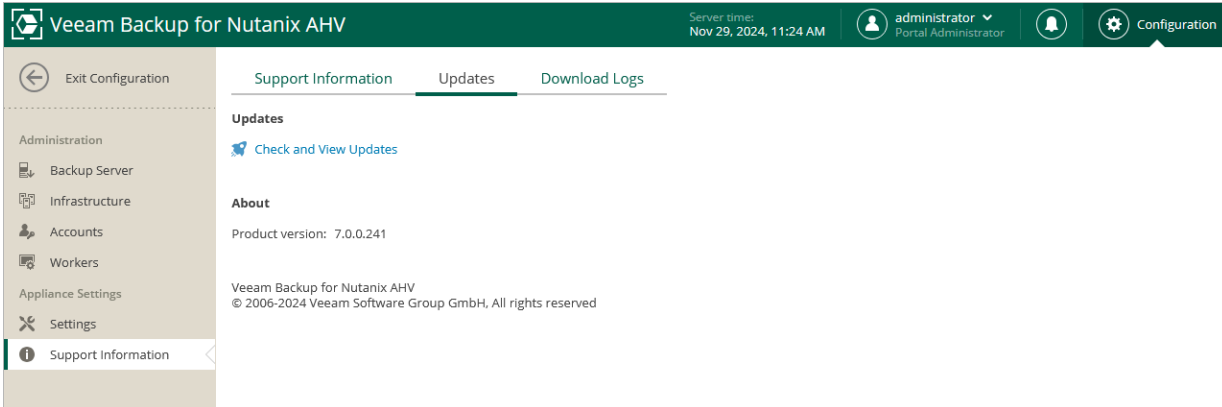


The screenshot shows the Veeam Updater Configuration page for the Proxy Server. The page title is "Proxy Server" and the subtitle is "Configure Internet proxy settings for Veeam Updater". The "Use Internet proxy" toggle is turned on. The "Host" field is labeled "Type in a hostname or IP address" and the "Port" field is a dropdown menu. The "Username (optional)" and "Password (optional)" fields are empty. The "Apply" button is highlighted in green, and a yellow tooltip message says "Save changes to apply a new configuration." The left sidebar shows "Exit Configuration", "Proxy Server", and "Support Information" options. The top right corner shows the date and time "August 6, 2023 at 04:55 AM GMT-7" and a "Configuration" gear icon.

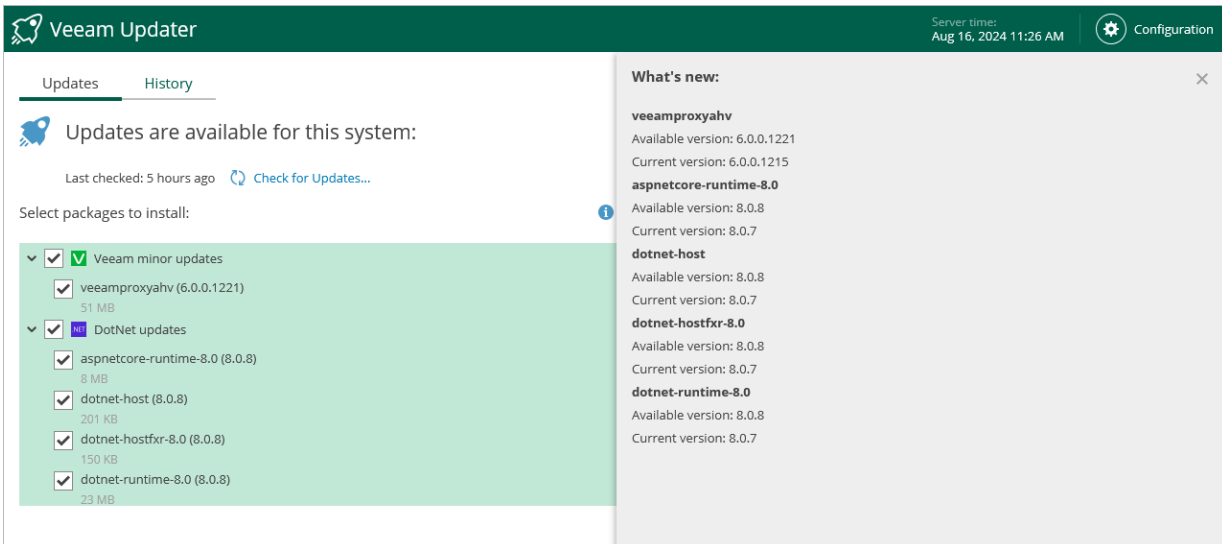
# Checking for Updates

Veeam Backup for Nutanix AHV automatically notifies you about newly released product versions and package updates available for the operating system running on the backup appliance. However, you can check for available updates manually if required:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Updates**.
3. Click **Check and View Updates**.



If new updates are available, Veeam Backup for Nutanix AHV will display them on the **Updates** tab of the **Veeam Backup Updater** page. To view detailed information on an update, select the check box next to the update and click **What's new?**



# Installing Updates

To download and install new product versions and available package updates, you can use either of the following options:

- [Install updates immediately](#)
- [Schedule update installation](#)

You can also [set a reminder to send update notifications](#).

## IMPORTANT

You can update Veeam Backup for Nutanix AHV using the Veeam updater service only. Updating of the backup appliance manually is not supported.

# Installing Updates

## IMPORTANT

Before you install a product update, make sure all jobs are stopped and restore tasks are finished. Otherwise, the update process will interrupt the running activities, which may result in data loss.

To download and install available product and package updates:

1. Open the **Veeam Updater** page:
  - a. Switch to the **Configuration** page.
  - b. Navigate to **Support Information**.
  - c. On the **Updates** tab, click **Check and View Updates**.
2. On the **Veeam Updater** page, do the following:
  - a. In the **Updates are available for this system** section, select check boxes next to the necessary updates.

## IMPORTANT

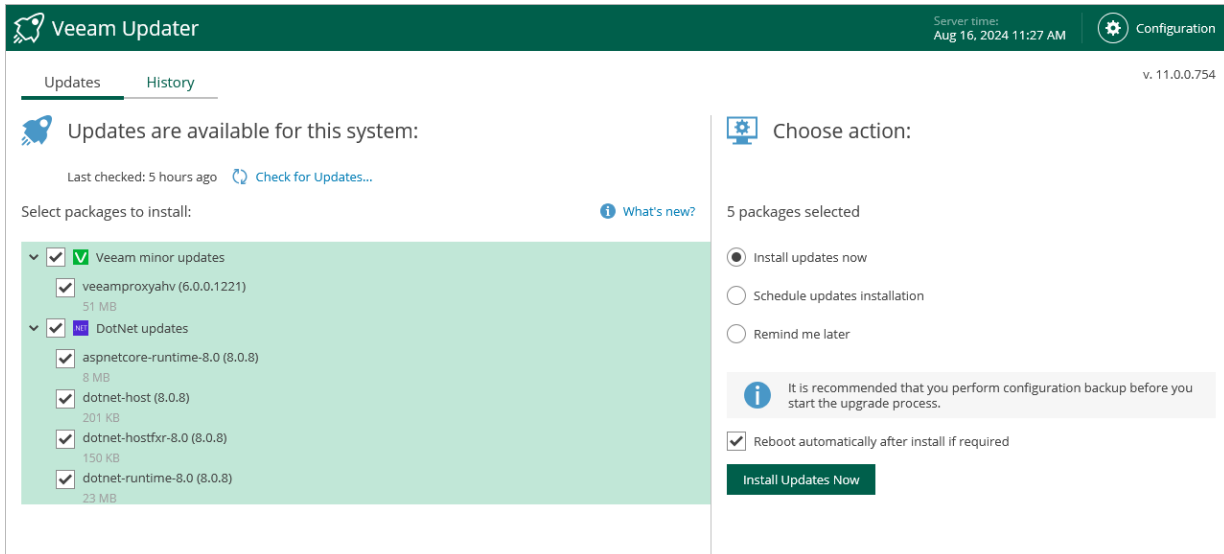
Make sure that you select all available DotNet updates for installation. Since different versions of DotNet components are not always compatible, the backup appliance may become non-operational if only some components are updated.

- b. In the **Choose action** section, select the **Install updates now** option, select the **Reboot automatically after install if required** check box to allow Veeam Backup for Nutanix AHV to reboot the backup proxy if needed, and then click **Install Updates Now**.



## NOTE

The updater may require you to read and accept the Veeam license agreement and the 3rd party components license agreement. If you reject the agreements, you will not be able to continue installation.



The screenshot shows the Veeam Updater interface. At the top, there is a green header with the Veeam logo, the text 'Veeam Updater', the server time 'Aug 16, 2024 11:27 AM', and a 'Configuration' button. Below the header, there are two tabs: 'Updates' (selected) and 'History'. The main content area is divided into two columns. The left column, titled 'Updates are available for this system:', shows a list of updates under the heading 'Select packages to install:'. The updates listed are: 'Veeam minor updates' (51 MB), 'veeamproxiahv (6.0.0.1221)' (51 MB), 'DotNet updates' (8 MB), 'aspnetcore-runtime-8.0 (8.0.8)' (8 MB), 'dotnet-host (8.0.8)' (201 KB), 'dotnet-hostfxr-8.0 (8.0.8)' (150 KB), and 'dotnet-runtime-8.0 (8.0.8)' (23 MB). The right column, titled 'Choose action:', shows '5 packages selected' and three radio button options: 'Install updates now' (selected), 'Schedule updates installation', and 'Remind me later'. There is also a checkbox for 'Reboot automatically after install if required' which is checked. A green 'Install Updates Now' button is at the bottom. A warning message states: 'It is recommended that you perform configuration backup before you start the upgrade process.'

Veeam Backup for Nutanix AHV will download and install the updates; the results of the installation process will be displayed on the [History](#) tab. It may take several minutes for the installation process to complete.

## NOTE

When installing product updates, Veeam Backup for Nutanix AHV restarts all services running on the backup proxy, including the Web UI service. That is why Veeam Backup for Nutanix AHV will log you out when the update process completes.

## Scheduling Update Installation

You can instruct Veeam Backup for Nutanix AHV to automatically download and install available product versions and package updates on a specific date at a specific time:

1. On the **Veeam Updater** page, in the **Updates are available for this system** section, select check boxes next to the necessary updates.
2. In the **Choose action** section, do the following:
  - a. Select the **Schedule updates installation** option and configure the necessary schedule.

## IMPORTANT

When selecting a date and time for the update installation, make sure no jobs are scheduled to run on the selected time. Otherwise, the update process will interrupt the running activities, which may result in data loss.

- b. Select the **Reboot automatically after install if required** check box to allow Veeam Backup for Nutanix AHV to reboot the backup proxy if needed.

### c. Click **Schedule Updates**.

The screenshot shows the Veeam Updater interface. The top bar includes the Veeam logo, the text 'Veeam Updater', the server time 'Aug 16, 2024 11:28 AM', and a 'Configuration' button. Below the top bar, there are tabs for 'Updates' and 'History'. The main content area is divided into two sections. On the left, under 'Updates are available for this system:', it shows 'Last checked: 5 hours ago' and a 'Check for Updates...' button. Below this, there is a list of packages to install, including 'Veeam minor updates' and 'DotNet updates' with their respective sizes. On the right, under 'Choose action:', there are three radio button options: 'Install updates now', 'Schedule updates installation' (which is selected), and 'Remind me later'. The 'Schedule updates installation' option has a dropdown menu set to '11:30 AM' and a date picker set to '16/08/2024'. Below these options, there is a note: 'It is recommended that you perform configuration backup before you start the upgrade process.' and a checkbox for 'Reboot automatically after install if required'. At the bottom right, there is a green 'Schedule Updates' button.

Veeam Backup for Nutanix AHV will automatically download and install the updates on the selected date at the selected time; the results of the installation process will be displayed on the [History tab](#).

## Setting Update Reminder

If you have not decided when to install available product versions and package updates, you can set an update reminder – instruct Veeam Backup for Nutanix AHV to send an update notification later.

To do that, on the **Veeam Updater** page, in the **Choose action** section, do the following:

1. Select the **Remind me later** option and choose when you want to receive the reminder.  
If you select the **Next Week** option, Veeam Backup for Nutanix AHV will send the reminder next Monday.
2. Click **Remind me later**.

This screenshot is similar to the previous one, showing the Veeam Updater interface. In the 'Choose action:' section, the 'Remind me later' radio button is selected. Below it, a dropdown menu is set to 'Next Week'. The 'Remind me later' button is highlighted in green at the bottom right.

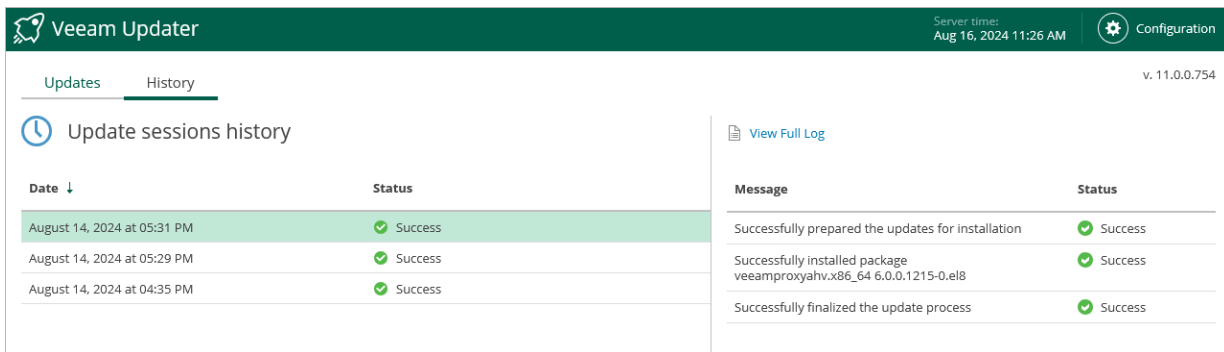
# Viewing Update History

To see the results of the update installation performed on the backup appliance, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Updates**.
3. Click **Check and View Updates**.
4. On the **Veeam Updater** page, switch to the **History** tab.

For each date when an update was installed, the **Veeam Updater** page will display the name of the update and its status (whether the installation process completed successfully, completed with warnings or failed to complete).

To download logs for the installed updates, select the necessary date in the **Date** section, and click **View Full Log**. Veeam Backup for Nutanix AHV will save the logs as a single file to the default download directory on the local machine.



The screenshot shows the Veeam Updater interface. The top navigation bar includes the Veeam Updater logo, server time (Aug 16, 2024 11:26 AM), and a Configuration button. The main content area has two tabs: 'Updates' and 'History', with 'History' selected. Below the tabs, there is a section titled 'Update sessions history' with a clock icon. To the right of this section is a 'View Full Log' link. The 'Update sessions history' section contains a table with two columns: 'Date' and 'Status'. The 'Date' column is sorted in descending order. The table lists three update sessions, all with a 'Success' status. The 'View Full Log' link is positioned above a detailed log table with two columns: 'Message' and 'Status'. This table shows three log entries, all with a 'Success' status.

Date ↓	Status
August 14, 2024 at 05:31 PM	Success
August 14, 2024 at 05:29 PM	Success
August 14, 2024 at 04:35 PM	Success

Message	Status
Successfully prepared the updates for installation	Success
Successfully installed package veeamproxyahv.x86_64 6.0.0.1215-0.el8	Success
Successfully finalized the update process	Success

# Getting Technical Support

If you have any questions or issues with Veeam Backup for Nutanix AHV, you can search for a resolution on [Veeam R&D Forums](#) or submit a support case in the [Veeam Customer Support Portal](#).

When you submit a support case, it is recommended that you provide the Veeam Customer Support Team with the following information:

- [Version information for the product and its components](#)
- Error message or accurate description of the problem you are facing
- Log files

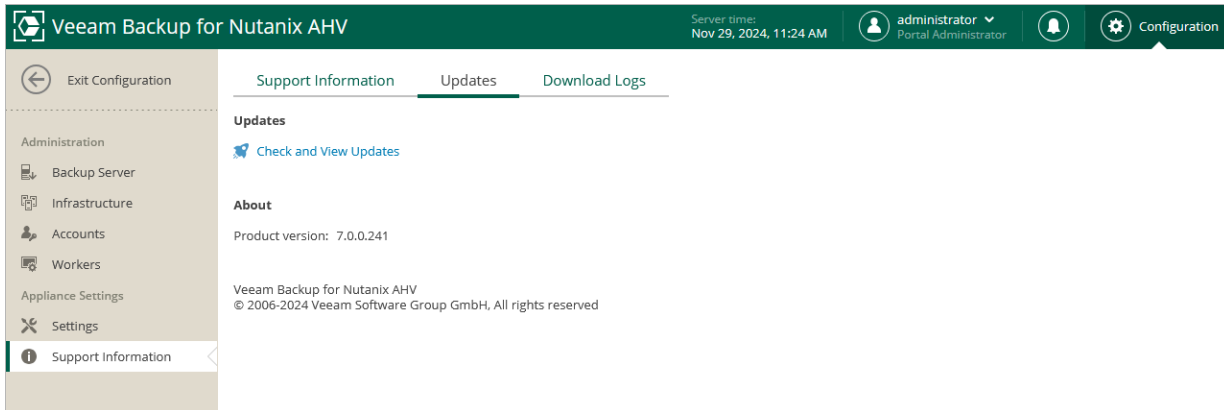
To export logs for Veeam Backup for Nutanix AHV, you must collect logs from both the [Veeam Backup & Replication console](#) and the [Nutanix AHV backup appliance web console](#).

# Viewing Product Details

To view the product details:

1. Switch to the **Configuration** page.
2. Navigate to **Updates**.

The **About** section displays the currently installed version of Veeam Backup for Nutanix AHV.

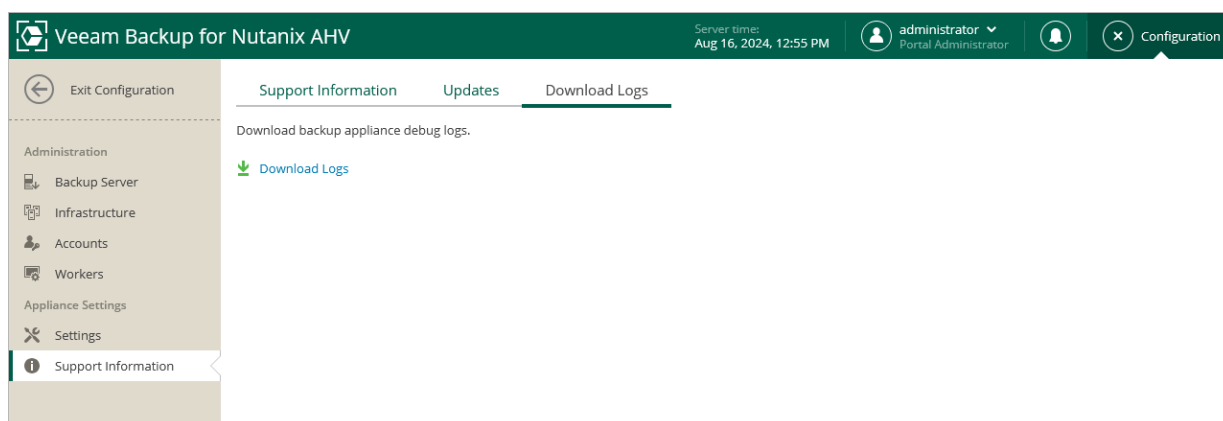


# Exporting Logs Using Backup Appliance Web Console

To download the product logs, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Download Logs**.
3. Click **Download Logs**.
4. In the **Download Logs** window, specify a time interval for which logs must be collected:
  - Select the **Collect logs for the last** option if you want to collect data for a specific number of days in the past.
  - Select the **Collect logs for specified time period** option if you want to collect data for a specific period of time in the past.
4. Click **Download**.

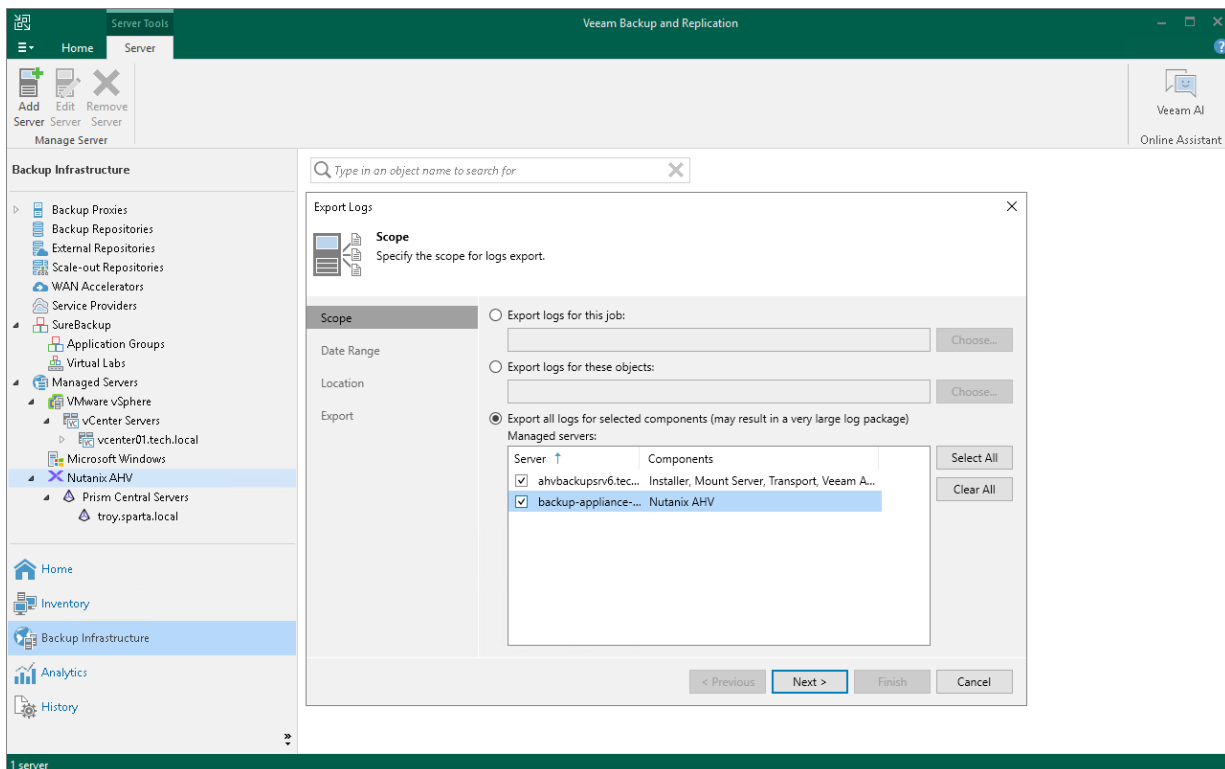
Veeam Backup for Nutanix AHV will collect logs for the specified time interval and save them to the default download folder on the local machine in a single `log.zip` archive.



# Exporting Logs Using Veeam Backup & Replication Console

To collect logs for the Nutanix AHV backup appliance from the Veeam Backup & Replication console, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Help > Support Information**.
2. At the **Scope** step of the **Export Logs** wizard, select the **Export all logs for selected components** option. Then, in the **Managed servers** list, select the backup server and the VM running as the Nutanix AHV backup appliance.
3. Complete the wizard as described in the Veeam Backup & Replication User Guide, section [Exporting Logs](#).



# Appendices

See in this section:

- [Appendix A. Custom Script Samples](#)
- [Appendix B. Installing Custom Certificate](#)
- [Appendix C. Configuring Bus Type Restore Priority](#)
- [Appendix D. Configuring Multiple Networks](#)



# Appendix A. Custom Script Samples

The sample custom scripts are written in Python 3 and contain classes and methods that are used to create a PD snapshot in a [Nutanix AHV cluster](#) and [Prism Central](#). The scripts also include methods that you can use to implement pre-freeze and post-thaw operations:

- `doSomeChecks(pdVms)`
- `freezeVms(self, pdVmsList)`
- `unFreezeVms(self, pdVms)`

## NOTE

The sample scripts do not check whether the provided Nutanix AHV cluster certificate is valid.

# Standalone Cluster Script

```

#!/usr/bin/python3
import requests
import urllib3
urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)
import json
import time
import logging
import argparse
import os
from urllib.parse import urlparse
from urllib.parse import urlunparse
from urllib.parse import urlencode
def quitScriptFailed(message):
    print(""" {"status": "Failed", "errorMessage": "%s"} """ % message)
    exit(0)
# #####
# #####
#
# Initialize
#
# #####
# #####
scriptDir = os.path.dirname(os.path.realpath(__file__))
#
# Process Command Line
#
parser = argparse.ArgumentParser(description='Create vms snapshot for Protection Domain')
parser.add_argument('--pdName', type=str, help='Protection Domain name')
parser.add_argument('--jobName', type=str, help='Job Name for displaying in header')
parser.add_argument('--logFolder', type=str, help='Logging path')
parser.add_argument('--logLevel', type=str, nargs='?', default='DEBUG', help='Log level: debug, info, warning, error, critical. (default=debug)')
args = parser.parse_args()
if args.pdName is None or args.pdName == "":
    quitScriptFailed("Invalid argument: pdName")
if args.jobName is None or args.jobName == "":
    quitScriptFailed("Invalid argument: jobName")
pdName = args.pdName
jobName = args.jobName
logDir = args.logFolder
logLevel = args.logLevel
#
# Init Log
#
from datetime import datetime
logFilePath = "%s/%s" % (logDir, "custom_script.log")
os.makedirs(logDir, exist_ok=True)
LogConsoleDbg = False
logging.addLevelName(logging.FATAL, 'Fatal')
logging.addLevelName(logging.ERROR, 'Error')
logging.addLevelName(logging.WARN, 'Warn ')
logging.addLevelName(logging.INFO, 'Info ')
logging.addLevelName(logging.DEBUG, 'Debug')
logging.basicConfig(
    filename=logFilePath,
    filemode='a',
    format='% (asctime) s.% (msecs) 03d] [% (levelname) s] % (message) s',
    datefmt=' [%Y-%m-%d] [%H:%M:%S]',

```

```

        level=str.upper(logLevel)
    )
# LogConsoleDbg = True
if LogConsoleDbg:
    Formater = logging.getLogger().handlers[0].formatter
    consoleHandler = logging.StreamHandler()
    consoleHandler.setFormatter(Formater)
    logging.getLogger().addHandler(consoleHandler)
logging.info("-----")
logging.info("Start to script execution")
logging.info("    Protection Domain name: %s" % pdName)
logging.info("    Job name:                %s" % jobName)
logging.info("-----")

#
# Init config from env variables
#
try:
    nutanixClusterIp = os.getenv('NUTANIX_CLUSTER_ADDRESS')
    nutanixLogin = os.getenv('NUTANIX_CLUSTER_LOGIN')
    nutanixPass = os.getenv('NUTANIX_CLUSTER_PASSWORD')
except Exception as e:
    # print(e)
    Error = "Failed to get environment variables";
    logging.exception(Error)
    quitScriptFailed(Error)
# #####
#
# Result's
#
# #####
class ScriptResult:
    """
    Describes the attributes and methods required to work with the Nutanix cluster API
    """
    SUCCESS = 0
    WARNING = 1
    FAILED = 2
    def __init__(self, status = SUCCESS, errorMessage = ""):
        self.status = status
        self.errorMessage = errorMessage
        self.oob_schedule_id = ""
    def SetFailed(self, message):
        self.status = ScriptResult.FAILED
        self.errorMessage = message
    def SetWarning(self, message):
        self.status = ScriptResult.WARNING
        self.errorMessage = message
    def makeJsonString(self):
        jResult = dict()
        if self.status == ScriptResult.SUCCESS:
            jResult['status'] = "Success"
        elif self.status == ScriptResult.WARNING:
            jResult['status'] = "Warning"
        elif self.status == ScriptResult.FAILED:
            jResult['status'] = "Failed"

```

```

        jResult['errorMessage'] = self.errorMessage
        jResult['oob_schedule_id'] = self.oob_schedule_id
        return json.dumps(jResult)
def generateResult(result):
    jsonString = result.makeJsonString()
    if result.status == ScriptResult.FAILED or result.status == ScriptResult.WAR
NING:
        logging.error("Custom script execution finished with error. Result: %s"
% jsonString)
    else:
        logging.info("Custom script execution finished. Result: %s" % jsonString
)
    print(jsonString)
    exit(0)
def generateResultFailed(errorMessage, exception = None):
    if isinstance(exception, NutanixException):
        errorMessage += " Error: %s" % str(exception)          # Add error message
of Nutanix API call. E.g.: "Specified protection domain AgentTest-ProtectionDom
ain does not exist"
    elif isinstance(exception, AuthenticationException):
        errorMessage += " Error: Authentication failed."
    elif isinstance(exception, requests.exceptions.ConnectionError):
        errorMessage += " Error: Connect failed."
        # if isinstance(e, NutanixExceptionNotFound):
        generateResult(ScriptResult(ScriptResult.FAILED, errorMessage))
def getFullUrl(url, params):
    paramsStr = urlencode(params)
    urlObj = urlparse(url)
    urlQuery = urlObj.query
    if paramsStr:
        if urlQuery:
            urlQuery = '%s&%s' % (urlQuery, paramsStr)
        else:
            urlQuery = paramsStr
    urlObj = urlObj._replace(query=urlQuery)
    return urlunparse(urlObj)
# #####
# #####
#
# Nutanix
#
# #####
# #####
class AuthenticationException(Exception):
    def __init__(self):
        super(AuthenticationException, self).__init__()
class NutanixException(Exception):
    def __init__(self, message=""):
        super(NutanixException, self).__init__(message)
        self.message = message
class NutanixExceptionNotFound(NutanixException):
    def __init__(self, message=""):
        super(NutanixExceptionNotFound, self).__init__(message)
class NutanixCluster:
    """
    Describes the attributes and methods required to work with the Nutanix clust
er API
    """
    def __init__(self, nutanixClusterIp, nutanixLogin, nutanixPass):
        """Constructor"""

```

```

        self.nutanixClusterIp = nutanixClusterIp
        self.nutanixLogin = nutanixLogin
        self.nutanixPass = nutanixPass
        self.nutanixApi1 = "https://%s:9440/PrismGateway/services/rest/v1" % self.nutanixClusterIp
        self.nutanixApi2 = "https://%s:9440/PrismGateway/services/rest/v2.0" % self.nutanixClusterIp
        self.nutanixApi3 = "https://%s:9440/api/nutanix/v3" % self.nutanixClusterIp

    @staticmethod
    def getErrorMessage(jResponse):
        if "message" in jResponse:
            # for Nutanix API v1 & v2
            return jResponse["message"]
        elif "message_list" in jResponse:
            # for Nutanix API v3
            message = ""
            for jError in jResponse["message_list"]:
                if "message" in jError:
                    if len(message):
                        message += "; "
                    message += jError["message"]
            else:
                raise requests.exceptions.RequestException
            return message
        else:
            raise requests.exceptions.RequestException

    @staticmethod
    def logRequest(level, httpMethod, urlFull, body=''):
        method = httpMethod.upper()
        bodyCanExist = False
        if method == 'POST' or method == 'PUT' or method == 'PATCH':
            bodyCanExist = True
        if bodyCanExist:
            logging.log(level, "HTTP Request (%s %s): %s", method, urlFull, "if body" if body else "none")
            if body:
                logging.log(level, body)
        else:
            logging.log(level, "HTTP Request (%s %s)", method, urlFull)

    @staticmethod
    def logResponse(level, response):
        logging.log(level, "HTTP Response (%s %s): %s", response.status_code, response.reason, "" if response.text else "none")
        if response.text:
            logging.log(level, response.text)

    def executeHttpRequest(self, httpMethod, url, body='', params=''):
        """
        Execute http/https requests
        :param httpMethod: (get, post ...)
        :param url:
        :param body: message body
        :param params: query string parameters in the URL
        :return:
        """
        response = None
        s = requests.Session()
        s.auth = (self.nutanixLogin, self.nutanixPass)
        s.headers.update({'Content-Type': 'application/json; charset=utf-8'})
        urlFull = getFullUrl(url, params)
        reqRespLogged = logging.root.level <= logging.DEBUG
        #

```

```

# Do request
#
self.logRequest(logging.DEBUG, httpMethod, urlFull, body)
if httpMethod == 'get':
    response = s.get(url, data=body, params=params, verify=False, timeou
t=10)
elif httpMethod == 'post':
    response = s.post(url, data=body, params=params, verify=False, timeo
ut=10)
else:
    raise NotImplementedError("HTTP Method '%s' not supported")
self.logResponse(logging.DEBUG, response)
#
# Check to failed request & log
#
isSuccess = response.status_code in (200, 201, 202)
if not isSuccess and not reqRespLogged:
    self.logRequest(logging.ERROR, httpMethod, urlFull, body)
    self.logResponse(logging.ERROR, response)
#
# Process response
#
if isSuccess:
    return response.json()
else:
    isJson = response.headers['Content-Type'].lower().find("application/
json") != -1
    if response.status_code == 401: # Not Found
        raise AuthenticationException()
    if isJson:
        message = self.getErrorMessage(response.json())
        logging.error("Nutanix request failed. Nutanix error: %s" % mess
age)

        if response.status_code == 404: # Not Found
            raise NutanixExceptionNotFound(message)
        else:
            raise NutanixException(message)
    else:
        logging.error("Nutanix request failed.")
        raise requests.exceptions.RequestException
def executeHttpRequestV1(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi1, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def executeHttpRequestV2(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi2, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def executeHttpRequestV3(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi3, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def getPdVms(self, pdName):
    """
    Get vm_id for the Protection Domain
    :param pdName: Protection Domain name
    :return pdVms: returns a list of vms_id for the Protection Domain
    """
    pdVms = []
    try:
        logging.info("Getting vm_id list for Protection Domain='%s'", pdName
)

```

```

        data = self.executeHttpRequestV2('get', "protection_domains/%s" % pd
Name)
        for vms in data["vms"]:
            pdVms.append(vms["vm_id"])
        except Exception as e:
            logging.error("Getting vm_id list from Protection Domain='%s' failed
. Error: %s.", pdName, str(e))
            logging.exception(str(e))
            generateResultFailed("Getting vm_id list for Protection Domain='%s'
failed." % pdName, e)
            logging.info("Getting vm_id list vm_id='%s' success", pdVms)
        return pdVms
    def getVmsInfoEx(self, vmsUuids, mustExist=True):
        """
        Obtaining information about vms, for example ip address
        Similarly, you can get vms data about cpu, memory, etc.
        This method is created as an example.
        :param vmsUuids: list of vm_id
        :return vmsInfo: returns a list of vm ip address
        """
        vmsInfo = []
        logging.info("Getting vms started. UUIDs: %s", vmsUuids)
        for vmUuid in vmsUuids:
            vmInfo = {"uuid": vmUuid, "vm": None}
            try:
                data = self.executeHttpRequestV3('get', "vms/%s" % vmUuid)
                vmInfo["vm"] = data
                vmsInfo.append(vmInfo)
            except Exception as e:
                if isinstance(e, NutanixExceptionNotFound):
                    if not mustExist:
                        vmsInfo.append(vmInfo)
                        continue
                logging.error("Getting vms failed. VM UUID: %s", vmUuid)
                logging.exception(str(e))
                generateResultFailed("Getting vms failed.", e)
            logging.info("Getting vms success. UUIDs: %s", vmsUuids)
        return vmsInfo
    def getVmsInfo(self, vmsUuids, mustExist=True):
        vmsInfo = []
        vmsInfoEx = self.getVmsInfoEx(vmsUuids, mustExist)
        for vmInfo in vmsInfoEx:
            if vmInfo["vm"]:
                # if exist on Nutanix --> add to
list
                vmsInfo.append(vmInfo["vm"])
        return vmsInfo
    def createPdSnapshot(self, pdName):
        """
        Creating vms snapshot for the Protection Domain
        :param pdName:
        :return:
        """
        schedule_id = None
        snapshotInfo = dict()
        body = json.dumps({"app_consistent": "false"})
        try:
            logging.info("Creating vms snapshot for Protection Domain='%s'", pdN
ame)
            data = self.executeHttpRequestV2('post', "protection_domains/%s/oob_
schedules" % pdName, body)

```



```

        schedule_id = data["schedule_id"]
    except Exception as e:
        logging.error("Creating vms snapshot for Protection Domain='%s' failed" % pdName)
        logging.exception(str(e))
        generateResultFailed("Creating vms snapshot for Protection Domain='%s' failed." % pdName, e)
        params = {'oob_schedule_ids': schedule_id}
        state = ''
        isStateAvailable = False
        for i in range(10):
            time.sleep(5) #waiting when information about snapshot be available
            try:
                logging.info("Getting snapshot status oob_schedule_ids='%s' for Protection Domain='%s'", schedule_id, pdName)
                data = self.executeHttpRequestV2('get', "protection_domains/%s/d_r_snapshots/" % pdName, '', params)
                try:
                    state = data["entities"][0]["state"]
                except:
                    state = None
                if state == 'AVAILABLE':
                    isStateAvailable = True
                    break
            except Exception as e:
                logging.error("Getting snapshot status oob_schedule_ids='%s' for Protection Domain='%s' failed", schedule_id, pdName)
                logging.exception(str(e))
                logging.info("Unfreeze vms for Protection Domain='%s'", pdName)
                cluster.unFreezeVms(pdVms)
                generateResultFailed("Getting snapshot status oob_schedule_ids='%s' for Protection Domain='%s' failed." % (schedule_id, pdName), e)
                if not isStateAvailable:
                    logging.error("Unexpected snapshot status='%s' oob_schedule_ids='%s' for Protection Domain='%s'", state, schedule_id, pdName)
                    logging.info("Unfreeze vms for Protection Domain='%s'", pdName)
                    cluster.unFreezeVms(pdVms)
                    generateResultFailed("Snapshot oob_schedule_ids='%s' status error (status='%s') for Protection Domain='%s'." % (state, schedule_id, pdName))
                logging.info("Creating vms snapshot for Protection Domain='%s' success", pdName)
        pdSnapshotId = data["entities"][0]["snapshot_id"]
        pdSnapshotUuid = data["entities"][0]["snapshot_uuid"]
        snapshotInfo['schedule_id'] = schedule_id
        snapshotInfo['pdSnapshotId'] = pdSnapshotId
        snapshotInfo['pdSnapshotUuid'] = pdSnapshotUuid
        return snapshotInfo
    def freezeVms(self, pdVmsList):
        """
        Freezing vms
        :param vmsInfo: vms data needed to freeze
        :return:
        """
        logging.info("Freezing vms for Protection Domain='%s'", pdName)
        # here should be added code for freezing vms"
        logging.info("Freezing vms for Protection Domain='%s' success", pdName)
        pass
    def unFreezeVms(self, pdVms):
        """
        unfreezing vms

```

```

        :param pdVms: vms data needed to unfreeze
        :return:
        """
        logging.info("Unfreezing vms for Protection Domain='%s'", pdName)
        # here should be added code for unfreezing vms"
        logging.info("Unfreezing vms for Protection Domain='%s' success", pdName
    )
        pass
def doSomeChecks(pdVms):
    # here should be added code for PD VMs check
    # for example get vms ip addresses
    logging.info("Do some checks")
    vmsInfo = []
    vmsInfoEx = cluster.getVmsInfoEx(pdVms, False)
    for vmInfo in vmsInfoEx:
        if vmInfo["vm"]: # if exist on Nutanix --> add to list
            vmsInfo.append(vmInfo["vm"])
        else:
            logging.warning("VM with UUID='%s' will not backup. VM in PD does not exist in cluster.", vmInfo["uuid"])
            logging.info("Do some checks success")
            # here should be added error message (in case of failed checks)
# #####
# #####
#
# Main
#
# #####
# #####
if __name__ == "__main__":
    cluster = NutanixCluster(nutanixClusterIp, nutanixLogin, nutanixPass)
    pdVms = cluster.getPdVms(pdName)
    if not pdVms:
        generateResultFailed("Protection Domain '%s' doesn't have vms." % pdName
    )
    doSomeChecks(pdVms)
    cluster.freezeVms(pdVms)
    snapshotInfo = cluster.createPdSnapshot(pdName)
    cluster.unFreezeVms(pdVms)
    logging.info("All steps succeeded")
    result = ScriptResult()
    result.oob_schedule_id = snapshotInfo['schedule_id']
    generateResult(result)

```

# Prism Central Script

```

#!/usr/bin/python3
import requests
import urllib3
urllib3.disable_warnings(urllib3.exceptions.InsecureRequestWarning)
import json
import time
import logging
import argparse
import os
from urllib.parse import urlparse
from urllib.parse import urlunparse
from urllib.parse import urlencode
def quitScriptFailed(message):
    print(""" {"status": "Failed", "errorMessage": "%s"} """ % message)
    exit(0)
# #####
#####
#
# Initialize
#
# #####
#####
scriptDir = os.path.dirname(os.path.realpath(__file__))
#
# Process Command Line
#
parser = argparse.ArgumentParser(description='Create vms snapshot for Protection Domain')
parser.add_argument('--clusterId', type=str, help='Id of a cluster with the corresponding protection domain')
parser.add_argument('--pdName', type=str, help='Protection domain name')
parser.add_argument('--jobName', type=str, help='Job name for displaying in header')
parser.add_argument('--logFolder', type=str, help='Logging path')
parser.add_argument('--logLevel', type=str, nargs='?', default='DEBUG', help='Log level: debug, info, warning, error, critical. (default=warning)')
args = parser.parse_args()
if args.pdName is None or args.pdName == "":
    quitScriptFailed("Invalid argument: pdName")
if args.jobName is None or args.jobName == "":
    quitScriptFailed("Invalid argument: jobName")
clusterId = args.clusterId
pdName = args.pdName
jobName = args.jobName
logDir = args.logFolder
logLevel = args.logLevel
#
# Init Log
#
from datetime import datetime
logFilePath = f"{logDir}/custom_script.log"
os.makedirs(logDir, exist_ok=True)
LogConsoleDbg = False
logging.addLevelName(logging.FATAL, 'Fatal')
logging.addLevelName(logging.ERROR, 'Error')
logging.addLevelName(logging.WARN, 'Warn ')
logging.addLevelName(logging.INFO, 'Info ')
logging.addLevelName(logging.DEBUG, 'Debug')
logging.basicConfig(

```

```

        filename=logFilePath,
        filemode='a',
        format='% (asctime)s.%(msecs)03d] [% (levelname)s] %(message)s',
        datefmt='[%Y-%m-%d] [%H:%M:%S]',
        level=str.upper(logLevel)
    )
# LogConsoleDbg = True
if LogConsoleDbg:
    Formater = logging.getLogger().handlers[0].formatter
    consoleHandler = logging.StreamHandler()
    consoleHandler.setFormatter(Formater)
    logging.getLogger().addHandler(consoleHandler)
logging.info("-----")
logging.info("Start to script execution")
logging.info(f"    Protection Domain name: {pdName}")
logging.info(f"    Job name:                {jobName}")
logging.info("-----")

#
# Init config from env variables
#
try:
    nutanixPrismCentralIp = os.getenv('NUTANIX_PRISM_CENTRAL_ADDRESS')
    nutanixLogin = os.getenv('NUTANIX_PRISM_CENTRAL_LOGIN')
    nutanixPass = os.getenv('NUTANIX_PRISM_CENTRAL_PASSWORD')
except Exception as e:
    # print(e)
    Error = "Failed to get environment variables";
    logging.exception(Error)
    quitScriptFailed(Error)
# #####
# #####
#
# Result's
#
# #####
# #####
class ScriptResult:
    """
    Describes the attributes and methods required to work with the Nutanix cluster API
    """
    SUCCESS = 0
    WARNING = 1
    FAILED = 2
    def __init__(self, status=SUCCESS, errorMessage=""):
        self.status = status
        self.errorMessage = errorMessage
        self.oob_schedule_id = ""
    def SetFailed(self, message):
        self.status = ScriptResult.FAILED
        self.errorMessage = message
    def SetWarning(self, message):
        self.status = ScriptResult.WARNING
        self.errorMessage = message
    def makeJsonString(self):
        jResult = dict()
        if self.status == ScriptResult.SUCCESS:
            jResult['status'] = "Success"

```

```

        elif self.status == ScriptResult.WARNING:
            jResult['status'] = "Warning"
        elif self.status == ScriptResult.FAILED:
            jResult['status'] = "Failed"
            jResult['errorMessage'] = self.errorMessage
            jResult['oob_schedule_id'] = self.oob_schedule_id
            return json.dumps(jResult)
def generateResult(result):
    jsonString = result.makeJsonString()
    if result.status == ScriptResult.FAILED or result.status == ScriptResult.WAR
NING:
        logging.error(f"Custom script execution finished with error. Result: {js
onString}")
    else:
        logging.info(f"Custom script execution finished. Result: {jsonString}")
        print(jsonString)
        exit(0)
def generateResultFailed(errorMessage, exception=None):
    if isinstance(exception, NutanixException):
        errorMessage += f" Error: {str(exception)}" # Add error message of Nutanix
API call. E.g.: "Specified protection domain AgentTest-ProtectionDomain doe
s not exist"
    elif isinstance(exception, AuthenticationException):
        errorMessage += " Error: Authentication failed."
    elif isinstance(exception, requests.exceptions.ConnectionError):
        errorMessage += " Error: Connect failed."
        # if isinstance(e, NutanixExceptionNotFound):
    generateResult(ScriptResult(ScriptResult.FAILED, errorMessage))
def getFullUrl(url, params):
    paramsStr = urlencode(params)
    urlObj = urlparse(url)
    urlQuery = urlObj.query
    if paramsStr:
        if urlQuery:
            urlQuery = f"{urlQuery}&{paramsStr}"
        else:
            urlQuery = paramsStr
    urlObj = urlObj._replace(query=urlQuery)
    return urlunparse(urlObj)
# #####
# #####
#
# Nutanix
#
# #####
# #####
class AuthenticationException(Exception):
    def __init__(self):
        super(AuthenticationException, self).__init__()
class NutanixException(Exception):
    def __init__(self, message=""):
        super(NutanixException, self).__init__(message)
        self.message = message
class NutanixExceptionNotFound(NutanixException):
    def __init__(self, message=""):
        super(NutanixExceptionNotFound, self).__init__(message)
class NutanixCluster:
    """
    Describes the attributes and methods required to work with the Nutanix clust
er API

```

```

"""
def __init__(self, nutanixPrismCentralIp, nutanixLogin, nutanixPass):
    """Constructor"""
    self.nutanixPrismCentralIp = nutanixPrismCentralIp
    self.nutanixLogin = nutanixLogin
    self.nutanixPass = nutanixPass
    self.nutanixApi1 = f"https://{self.nutanixPrismCentralIp}:9440/PrismGate
way/services/rest/v1"
    self.nutanixApi2 = f"https://{self.nutanixPrismCentralIp}:9440/PrismGate
way/services/rest/v2.0"
    self.nutanixApi3 = f"https://{self.nutanixPrismCentralIp}:9440/api/nutan
ix/v3"
    @staticmethod
    def getErrorMessage(jResponse):
        if "message" in jResponse: # for Nutanix API v1 & v2
            return jResponse["message"]
        elif "message_list" in jResponse: # for Nutanix API v3
            message = ""
            for jError in jResponse["message_list"]:
                if "message" in jError:
                    if len(message):
                        message += "; "
                    message += jError["message"]
            else:
                raise requests.exceptions.RequestException
            return message
        else:
            raise requests.exceptions.RequestException
    @staticmethod
    def logRequest(level, httpMethod, urlFull, body=''):
        method = httpMethod.upper()
        bodyCanExist = False
        if method == 'POST' or method == 'PUT' or method == 'PATCH':
            bodyCanExist = True
        if bodyCanExist:
            logging.log(level, "HTTP Request (%s %s): %s", method, urlFull, "" i
f body else "none")
            if body:
                logging.log(level, body)
        else:
            logging.log(level, "HTTP Request (%s %s)", method, urlFull)
    @staticmethod
    def logResponse(level, response):
        logging.log(level, "HTTP Response (%s %s): %s", response.status_code, re
sponse.reason,
            "" if response.text else "none")
        if response.text:
            logging.log(level, response.text)
    def executeHttpRequest(self, httpMethod, url, body='', params=''):
        """
        Execute http/https requests
        :param httpMethod: (get, post ...)
        :param url:
        :param body: message body
        :param params: query string parameters in the URL
        :return:
        """
        response = None
        s = requests.Session()
        s.auth = (self.nutanixLogin, self.nutanixPass)

```

```

s.headers.update({'Content-Type': 'application/json; charset=utf-8'})
urlFull = getFullUrl(url, params)
reqRespLogged = logging.root.level <= logging.DEBUG
#
# Do request
#
self.logRequest(logging.DEBUG, httpMethod, urlFull, body)
if httpMethod == 'get':
    response = s.get(url, data=body, params=params, verify=False, timeou
t=10)
elif httpMethod == 'post':
    response = s.post(url, data=body, params=params, verify=False, timeo
ut=10)
else:
    raise NotImplementedError("HTTP Method '%s' not supported")
self.logResponse(logging.DEBUG, response)
#
# Check to failed request & log
#
isSuccess = response.status_code in (200, 201, 202)
if not isSuccess and not reqRespLogged:
    self.logRequest(logging.ERROR, httpMethod, urlFull, body)
    self.logResponse(logging.ERROR, response)
#
# Process response
#
if isSuccess:
    return response.json()
else:
    isJson = response.headers['Content-Type'].lower().find("application/
json") != -1
    if response.status_code == 401: # Not Found
        raise AuthenticationException()
    if isJson:
        message = self.getErrorMessage(response.json())
        logging.error("Nutanix request failed. Nutanix error: %s" % mess
age)
        if response.status_code == 404: # Not Found
            raise NutanixExceptionNotFound(message)
        else:
            raise NutanixException(message)
    else:
        logging.error("Nutanix request failed.")
        raise requests.exceptions.RequestException
def executeHttpRequestV1(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi1, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def executeHttpRequestV2(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi2, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def executeHttpRequestV3(self, httpMethod, relativeUrl, body='', params=''):
    url = "%s/%s" % (self.nutanixApi3, relativeUrl)
    return self.executeHttpRequest(httpMethod, url, body, params)
def getPdVms(self, pdName):
    """
    Get vm_id for the Protection Domain
    :param pdName: Protection Domain name
    :return pdVms: returns a list of vms_id for the Protection Domain
    """
    pdVms = []

```



```

        params = {"proxyClusterUuid": clusterId}
        try:
            logging.info(f"Getting vm_id list for Protection Domain='{pdName}'")
            data = self.executeHttpRequestV2('get', f"protection_domains/{pdName}
}", params=params)
            for vms in data["vms"]:
                pdVms.append(vms["vm_id"])
        except Exception as e:
            logging.error(f"Getting vm_id list from Protection Domain='{pdName}'
failed. Error: {str(e)}.")
            logging.exception(str(e))
            generateResultFailed(f"Getting vm_id list for Protection Domain='{pd
Name}' failed.", e)
        logging.info(f"Getting vm_id list vm_id='{pdVms}' success")
        return pdVms
def getVmsInfoEx(self, vmsUuids, mustExist=True):
    """
    Obtaining information about vms, for example ip address
    Similarly, you can get vms data about cpu, memory, etc.
    This method is created as an example.
    :param vmsUuids: list of vm_id
    :return vmsInfo: returns a list of vm ip address
    """
    vmsInfo = []
    params = {"proxyClusterUuid": clusterId}
    logging.info(f"Getting vms started. UUIDs: {vmsUuids}")
    for vmUuid in vmsUuids:
        vmInfo = {"uuid": vmUuid, "vm": None}
        try:
            data = self.executeHttpRequestV3('get', f"vms/{vmUuid}", params=
params)
            vmInfo["vm"] = data
            vmsInfo.append(vmInfo)
        except Exception as e:
            if isinstance(e, NutanixExceptionNotFound):
                if not mustExist:
                    vmsInfo.append(vmInfo)
                    continue
            logging.error(f"Getting vms failed. VM UUID: {vmUuid}", )
            logging.exception(str(e))
            generateResultFailed("Getting vms failed.", e)
        logging.info(f"Getting vms success. UUIDs: {vmsUuids}")
    return vmsInfo
def getVmsInfo(self, vmsUuids, mustExist=True):
    vmsInfo = []
    vmsInfoEx = self.getVmsInfoEx(vmsUuids, mustExist)
    for vmInfo in vmsInfoEx:
        if vmInfo["vm"]: # if exist on Nutanix --> add to list
            vmsInfo.append(vmInfo["vm"])
    return vmsInfo
def createPdSnapshot(self, pdName):
    """
    Creating vms snapshot for the Protection Domain
    :param pdName:
    :return:
    """
    schedule_id = None
    snapshotInfo = dict()
    body = json.dumps({"app_consistent": "false"})
    try:

```

```

        logging.info(f"Creating vms snapshot for Protection Domain='{pdName}'")
    """
    data = self.executeHttpRequestV2('post', f"protection_domains/{pdName}/oob_schedules", body, params={"proxyClusterUuid": clusterId})
    schedule_id = data["schedule_id"]
    except Exception as e:
        logging.error(f"Creating vms snapshot for Protection Domain='{pdName}' failed")
        logging.exception(str(e))
        generateResultFailed(f"Creating vms snapshot for Protection Domain='{pdName}' failed.", e)
    params = {
        'oob_schedule_ids': schedule_id,
        "proxyClusterUuid": clusterId
    }
    state = ''
    isStateAvailable = False
    for i in range(10):
        time.sleep(5) # waiting when information about snapshot be available
    try:
        logging.info(f"Getting snapshot status oob_schedule_ids='{schedule_id}' for Protection Domain='{pdName}'")
        data = self.executeHttpRequestV2('get', f"protection_domains/{pdName}/dr_snapshots/", '', params)
        try:
            state = data["entities"][0]["state"]
        except:
            state = None
        if state == 'AVAILABLE':
            isStateAvailable = True
            break
    except Exception as e:
        logging.error(f"Getting snapshot status oob_schedule_ids='{schedule_id}' for Protection Domain='{pdName}' failed")
        logging.exception(str(e))
        logging.info(f"Unfreeze vms for Protection Domain='{pdName}'")
        cluster.unFreezeVms(pdVms)
        generateResultFailed(f"Getting snapshot status oob_schedule_ids='{schedule_id}' for Protection Domain='{pdName}' failed.", e)
    if not isStateAvailable:
        logging.error(f"Unexpected snapshot status='{state}' oob_schedule_ids='{schedule_id}' for Protection Domain='{pdName}'")
        logging.info(f"Unfreeze vms for Protection Domain='{pdName}'")
        cluster.unFreezeVms(pdVms)
        generateResultFailed(f"Snapshot oob_schedule_ids='{state}' status error (status='{schedule_id}') for Protection Domain='{pdName}'.")
        logging.info(f"Creating vms snapshot for Protection Domain='{pdName}' success")
    pdSnapshotId = data["entities"][0]["snapshot_id"]
    pdSnapshotUuid = data["entities"][0]["snapshot_uuid"]
    snapshotInfo['schedule_id'] = schedule_id
    snapshotInfo['pdSnapshotId'] = pdSnapshotId
    snapshotInfo['pdSnapshotUuid'] = pdSnapshotUuid
    return snapshotInfo
def freezeVms(self, pdVmsList):
    """
    Freezing vms
    :param vmsInfo: vms data needed to freeze
    :return:

```

```

    """
    logging.info(f"Freezing vms for Protection Domain='{pdName}''")
    # here should be added code for freezing vms"
    logging.info(f"Freezing vms for Protection Domain='{pdName}' success")
    pass
def unFreezeVms(self, pdVms):
    """
    unfreezing vms
    :param pdVms: vms data needed to unfreeze
    :return:
    """
    logging.info(f"Unfreezing vms for Protection Domain='{pdName}''")
    # here should be added code for unfreezing vms"
    logging.info(f"Unfreezing vms for Protection Domain='{pdName}' success")
    pass
def doSomeChecks(pdVms):
    # here should be added code for PD VMs check
    # for example get vms ip addresses
    logging.info("Do some checks")
    vmsInfo = []
    vmsInfoEx = cluster.getVmsInfoEx(pdVms, False)
    for vmInfo in vmsInfoEx:
        if vmInfo["vm"]: # if exist on Nutanix --> add to list
            vmsInfo.append(vmInfo["vm"])
        else:
            logging.warning(f"VM with UUID='{vmInfo['uuid']}' will not backup. V
M in PD does not exist in cluster.")
            logging.info("Do some checks success")
            # here should be added error message (in case of failed checks)
# #####
# #####
#
# Main
#
# #####
# #####
if __name__ == "__main__":
    cluster = NutanixCluster(nutanixPrismCentralIp, nutanixLogin, nutanixPass)
    pdVms = cluster.getPdVms(pdName)
    if not pdVms:
        generateResultFailed(f"Protection Domain '{pdName}' doesn't have vms.")
    doSomeChecks(pdVms)
    cluster.freezeVms(pdVms)
    snapshotInfo = cluster.createPdSnapshot(pdName)
    cluster.unFreezeVms(pdVms)
    logging.info("All steps succeeded")
    result = ScriptResult()
    result.oob_schedule_id = snapshotInfo['schedule_id']
    generateResult(result)

```

# Appendix B. Installing Custom Certificate

When you upgrade the Nutanix AHV backup appliance to version 4.0 or later, a Veeam certificate is automatically installed on the appliance even if a custom certificate was used before. If you want to replace the Veeam certificate with your custom certificate, the following files are required:

- `certificate.pem` – a file that contains the Nutanix AHV backup appliance certificate in the PEM format. Note that the PFX format is not supported.
- `privatekey.pem` – a file that contains a private key used to generate the certificate.
- `password` – a file that contains a password to decrypt the private key. This file is not required if the private key is not encrypted.

To install the certificate, do the following:

1. [Enable SSH access](#) on the Nutanix AHV backup appliance.
2. Connect to the Nutanix AHV backup appliance using SSH.
3. Upload certificate files to the Nutanix AHV backup appliance.
4. Open the `/opt/VeeamBackupAgent/appsettings.json` file and locate the **BackupApplianceCertificate** configuration section.
5. Update the configuration parameters as in the following example:

```
"BackupApplianceCertificate": {
  "UseCustomCert": true,
  "CustomCertPemFilePath": "<path_to_the_file>/certificate.pem",
  "CustomPrivateKeyPemFilePath": "<path_to_the_file>/privatekey.pem",
  "CustomPasswordFilePath": "<path_to_the_file>/password"
}
```

If the private key is not encrypted, leave the default "" value for the `CustomPasswordFilePath` parameter.

6. Save the `appsettings.json` file.
7. Restart the backup service using the following command:

```
sudo systemctl restart veeamahvbackup
```

## IMPORTANT

If the Nutanix AHV backup appliance cannot find the certificate in the specified folder or any of the certificate files contain incorrect data, the backup service will fail to start, and you will not be able to perform data protection and recovery operations. To troubleshoot the problem, check logs in the `/var/log/veeam_backup/appliance_service/appliance_service.log` file.

8. Update the Nutanix AHV backup appliance configuration using the [Edit Nutanix Proxy wizard](#).

# Appendix C. Configuring Bus Type Restore Priority

When restoring a VM that originally resided on a platform other than Nutanix AHV, Veeam Backup for Nutanix AHV attaches disks with the restored data to the target Nutanix AHV VM taking into account the original disk bus types unless the following limits are exceeded: 6 SATA, 256 SCSI, 4 IDE, 7 PCI disks. Since the maximum number of disk nodes to which disks of a specific bus type can be attached varies depending on the virtualization platform, Veeam Backup for Nutanix AHV may fail to attach some of the VM disks using their original bus types. Those disks will be attached to free nodes of other bus types in the following default priority: SATA, SCSI, IDE, PCI.

You can modify the default priority to define the order in which Veeam Backup for Nutanix AHV will process disks that cannot be attached using their original bus types. You can also instruct Veeam Backup for Nutanix AHV to ignore the original bus types of VM disks. In the latter case, Veeam Backup for Nutanix AHV will attach disks according to the specified bus type priority – this may be useful if some bus type is not configured in the Nutanix AHV environment.

## NOTE

Veeam Backup for Nutanix AHV takes into account the bus type restore priority only when performing the following operations:

- [Restore of an entire VM](#) that originally resided on a platform other than Nutanix AHV.
- [Instant Recovery of any VM](#) (including Nutanix AHV VMs) to Nutanix AHV.

Consider the following example. You want to restore a VMware VM that originally had 30 SATA disks and 2 IDE disks. Depending on the bus type restore priority, Veeam Backup for Nutanix AHV will attach disks to the following nodes of the target VM:

Bus Type Priority	Ignore Original Bus	Target VM Disk Nodes
SATA, SCSI, IDE, PCI (default)	False	<ul style="list-style-type: none"><li>• 6 SATA (originally, 6 SATA)</li><li>• 24 SCSI (originally, 24 SATA)</li><li>• 2 IDE (originally)</li><li>• 0 PCI</li></ul>
SATA, IDE, PCI, SCSI	False	<ul style="list-style-type: none"><li>• 6 SATA (originally, 6 SATA)</li><li>• 4 IDE (originally, 2 IDE and 2 SATA)</li><li>• 7 PCI (originally, 7 SATA)</li><li>• 15 SCSI (originally, 15 SATA)</li></ul>
SCSI, IDE, PCI, SATA	False	<ul style="list-style-type: none"><li>• 24 SCSI (originally, 24 SATA)</li><li>• 2 IDE (originally, 2 IDE)</li><li>• 0 PCI</li><li>• 6 SATA (originally, 6 SATA)</li></ul>
SCSI, IDE, PCI, SATA	True	<ul style="list-style-type: none"><li>• 32 SCSI (originally, 30 SATA and 2 IDE)</li><li>• 0 IDE</li><li>• 0 PCI</li><li>• 0 SATA</li></ul>

To modify the default bus type restore priority, do the following:

1. Close the Veeam Backup & Replication console.
2. Open a plain text editor (for example, Notepad) as Administrator.
3. In the editor, open the `appsettings.json` file located in the `{plug-in location}\Service` folder.  
The default location of Nutanix AHV plug-in is `C:\Program Files\Veeam\Plugins\Nutanix AHV`. However, the location may differ depending on the [specified setup settings](#).
4. Locate the **RestoreDefaults** configuration section.

To instruct Veeam Backup for Nutanix AHV to ignore the original bus types of VM disks, set the following parameter to `true`:

```
"IgnoreOriginalBus": "true",
```

To change the bus type priority, update the following parameter value:

```
"BusesFillingOrder": "SCSI, IDE, PCI, SATA",
```

5. Save the `appsettings.json` file.
6. Restart the Veeam AHV Service.

# Appendix D. Configuring Multiple Networks

Starting from version 6.0, Nutanix AHV allows you to connect the backup appliance and workers to multiple networks. This may be helpful if your corporate policies require that inbound and outbound internet traffic is delivered through a secure network only, or if you want to use a specific network to transfer backed-up data from and to backup repositories.

Since both the backup appliance and workers deployed by Nutanix AHV are Linux-based VMs, they have the same limitations that apply to machines running the Rocky Linux operating system. That is, network routing can only be applied to the networks connected to the network adapters (vNICs) that has been added first while configuring the backup appliance and workers, which mean that these VMs can reach out to endpoints in other networks only through those first vNICs.

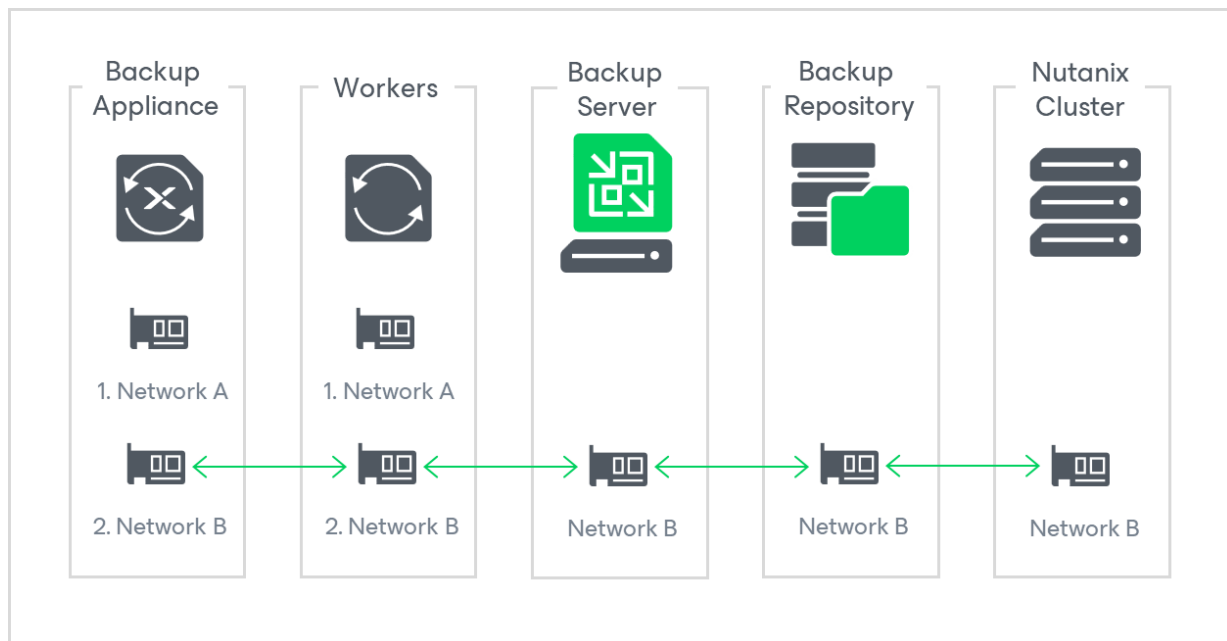
That is why you must consider the following while configuring multiple networks for the backup appliance and workers:

- If you want the backup appliance and workers to obtain updates from online Veeam repositories, you must connect to the first vNIC a network that allows inbound and outbound internet traffic.
- If a backup repository, the backup server, the Nutanix AHV cluster or the Prism Central is not reachable from the network connected to the first vNIC, you must update the backup appliance and worker settings to add one more vNIC and to connect it to the network to which that component is connected.

This section describes examples of valid and invalid network configurations.

## Example 1. Valid Configuration

In this example, the backup appliance, all workers, the backup server, the repository and the Nutanix AHV cluster are connected to Network B, while the backup appliance and the workers are also connected to Network A that allows them to obtain updates from the internet. This configuration is valid since all backup infrastructure components are connected to the same network.

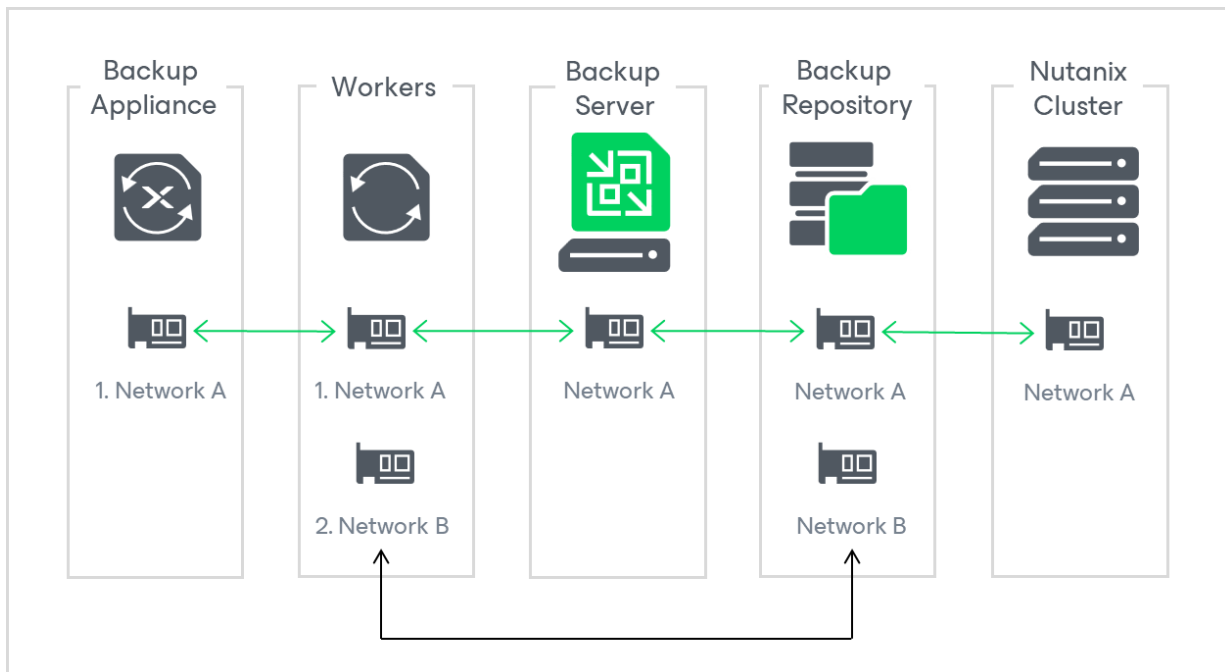


## Example 2. Valid Configuration

In this example, the backup appliance, all workers, the backup server, the repository and the Nutanix AHV cluster are connected to Network A, while the workers and the backup repository are also connected to Network B that is [configured as a preferred network](#) to deliver traffic to the backup repository. This configuration is valid since all backup infrastructure components are connected to the same network.

### NOTE

The backup appliance and workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.



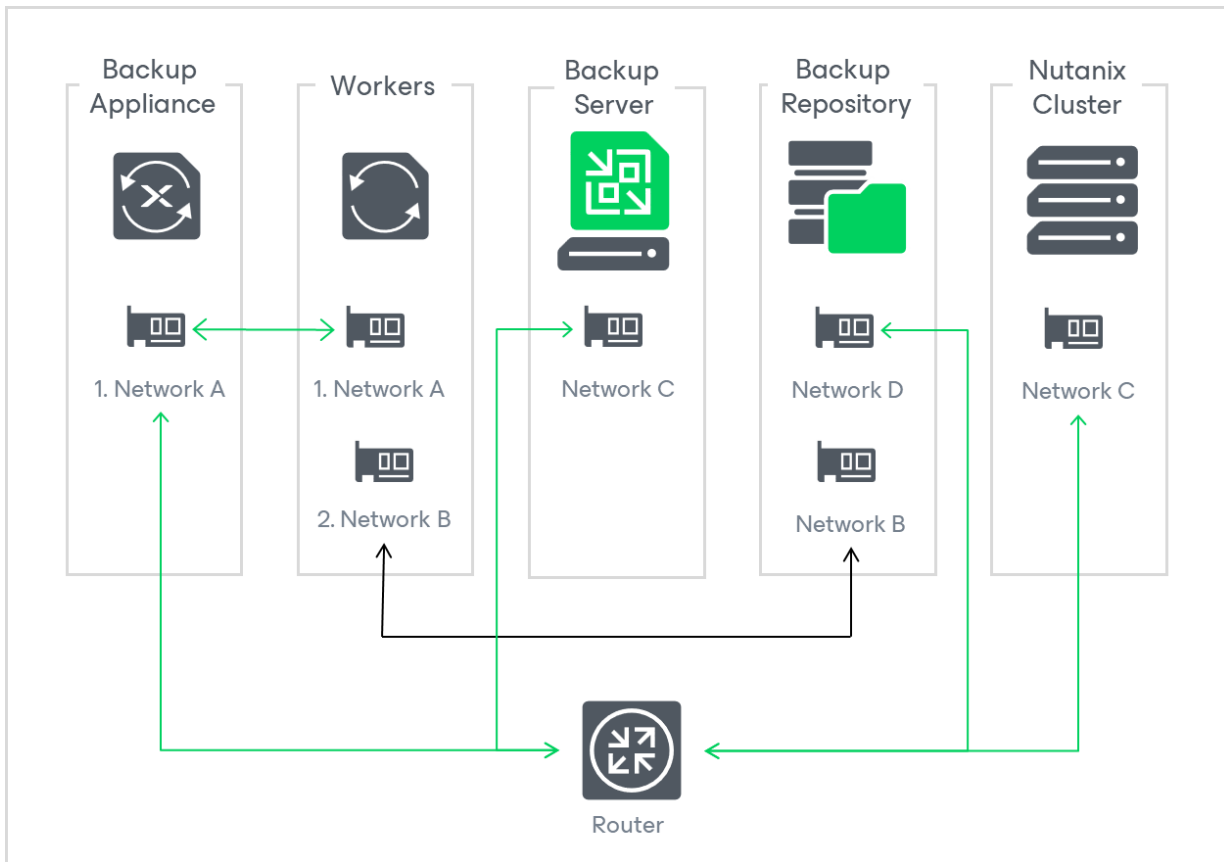
## Example 3. Valid Configuration

In this example, the backup appliance and all workers are connected to Network A using their first vNICs, while the workers are also connected to Network B that is [configured as a preferred network](#) to deliver traffic to the backup repository. Also, you have a router configured to forward traffic between networks A, C and D. This configuration is valid since the backup appliance and the workers can use Network A to communicate with other backup infrastructure components through the router.



## NOTE

The backup appliance and workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.



## Example 4. Invalid Configuration

In this example, the backup appliance and all workers are connected both to Network A using their first vNICs and to Network B using their second vNICs, while the backup server, the backup repository and the Nutanix cluster are connected to Network C. Also, you have a router configured to forward traffic between networks B and C. This configuration is invalid since the workers and the backup appliance cannot use Network B to communicate with other backup infrastructure components through the router.

To make the configuration valid, do either of the following:

- Change your network configuration to connect Network A to the router.

- Add more vNICs to the backup appliance and the workers. Then, connect these vNICs to Network C.

