# Veeam Plug-in for Nutanix AHV

Version 9

User Guide

February, 2026

**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

# 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
- Veeam R&D Forums: 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 Plug-in for Nutanix AHV is a software component developed for protection and disaster recovery tasks for Nutanix AHV environment. This component comes as part of the Veeam Backup & Replication solution and allows you to perform the following operations:

- Create backups of Nutanix AHV VMs and store them in backup repositories.
- 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, Proxmox VE and Scale Computing HyperCore 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.
- Restore VMs from Nutanix AHV backups to Proxmox VE, oVirt KVM, Scale Computing HyperCore 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, PostgreSQL, 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 Plug-in 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 Plug-in 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 Plug-in 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

Since Veeam Plug-in for Nutanix AHV is integrated with Veeam Backup & Replication, the solution architecture in the Prism Central deployment scenario comprises the following set of components:

- Nutanix AHV Prism Central
- Nutanix AHV clusters
- Backup server
- Veeam Plug-in for Nutanix AHV
- 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 Plug-in 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 Plug-in 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

A backup server is either a Windows-based or Linux-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.

## Veeam Plug-in for Nutanix AHV

Veeam Plug-in for Nutanix AHV is an architecture component that enables integration between the backup server and other components of the backup infrastructure. Veeam Plug-in for Nutanix AHV allows Veeam Backup & Replication to connect to the Nutanix AHV Prism Central, and to perform data protection and disaster recovery tasks with Nutanix AHV resources.

## Backup Repositories

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

To communicate with backup repositories, Veeam Plug-in 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 a Linux-based VM that resides on the Nutanix AHV host and processes backup workloads when transferring data to and from backup repositories. For optimal performance, workers can be deployed in all Nutanix AHV Prism Central clusters and can be distributed among cluster hosts (nodes). For more information on deployment sizing considerations, see Sizing Guidelines.

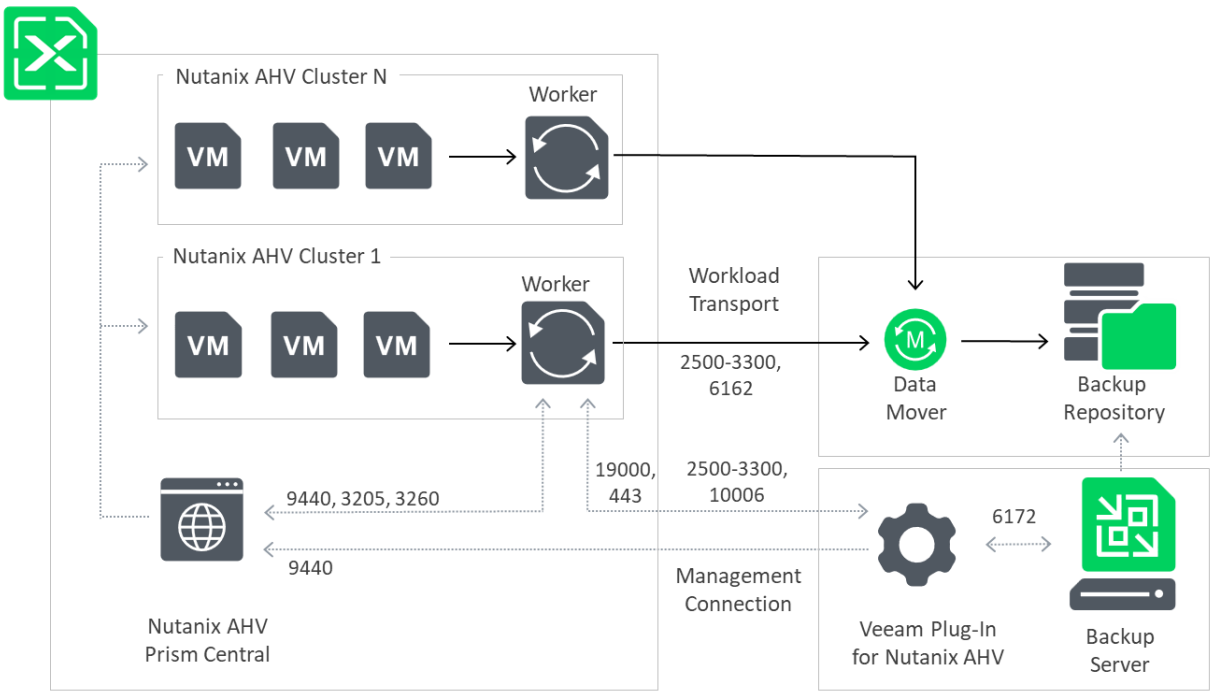

# Standalone Cluster Deployment Scenario

Since Veeam Plug-in for Nutanix AHV is integrated with Veeam Backup & Replication, the solution architecture in the standalone cluster deployment scenario comprises the following set of components:

- Nutanix AHV cluster
- Backup server
- Veeam Plug-in for Nutanix AHV
- 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 Plug-in for Nutanix AHV uses the Nutanix AHV cluster to access Nutanix AHV resources such as VMs, volume groups, storage containers and networks.

## Backup Server

A backup server is either a Windows-based or Linux-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.

## Veeam Plug-in for Nutanix AHV

Veeam Plug-in for Nutanix AHV is an architecture component that enables integration between the backup server and other components of the backup infrastructure. Veeam Plug-in for Nutanix AHV allows Veeam Backup & Replication to connect to the Nutanix AHV cluster and to manage backup and restore tasks to protect Nutanix AHV resources.

## Backup Repositories

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

To communicate with backup repositories, Veeam Plug-in 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 a Linux-based VM that resides in the Nutanix AHV cluster and processes backup workloads when transferring data to and from backup repositories.

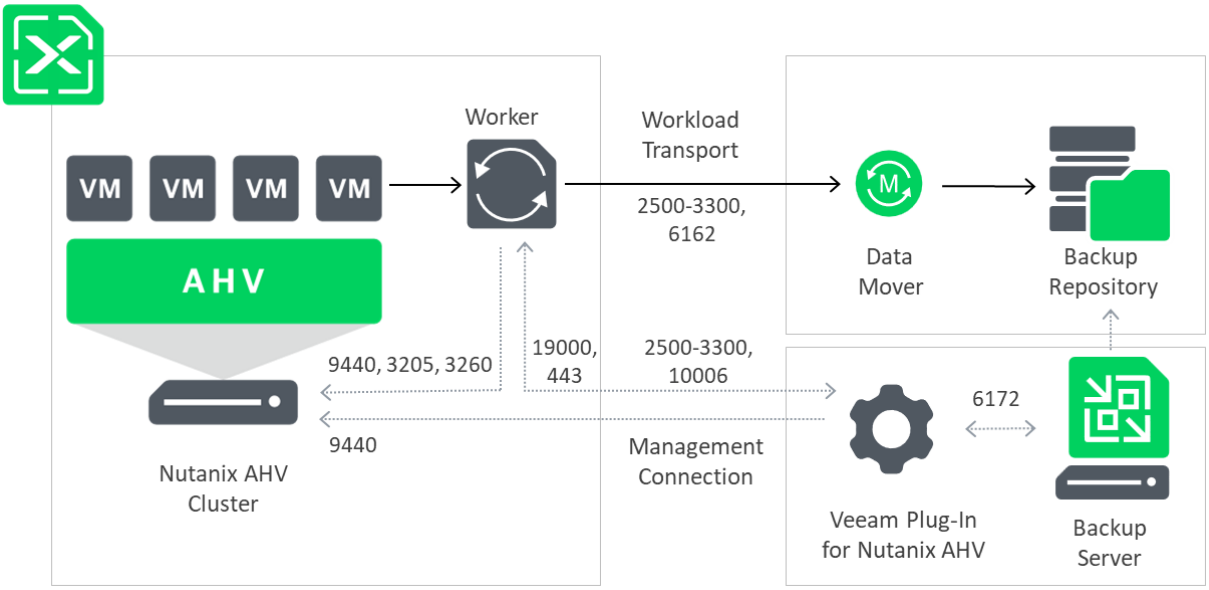

# VM Backup

To produce backups of VMs, Veeam Backup & Replication 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 & Replication does not install agent software inside VMs to retrieve data. Veeam Backup & Replication uses native Nutanix AHV capabilities instead. During every backup session, Veeam Backup & Replication 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. Configure worker settings.
5. Complete the New Backup Job wizard.

## How VM Backup Works

Veeam Backup & Replication 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.
2. Launches a worker on the same host where the processed VM resides.

   If no worker is deployed on the host, Veeam Backup & Replication launches any other Nutanix AHV worker that is added to the backup infrastructure.
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 & Replication uses the changed block tracking (CBT) mechanism: during incremental backup sessions, Veeam Backup & Replication 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 & Replication reads all data from the snapshot. For more information, see Changed Block Tracking.

   Veeam Backup & Replication 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 & Replication 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 & Replication 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 & Replication 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 Plug-in for Nutanix AHV to increase the speed and efficiency of incremental backups:

- During a full backup session Veeam Plug-in for Nutanix AHV reads only written data blocks, while unallocated data blocks are filtered out.
- During an incremental backup session, Veeam Plug-in 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 Plug-in for Nutanix AHV creates a snapshot of a VM using native Nutanix AHV capabilities. To do that, Veeam Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV to detect data blocks that have changed since the previous backup session.

# Limitations for Changed Block Tracking

Veeam Plug-in 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 Plug-in 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

Veeam Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV applies forward incremental backup retention policy.

> **NOTE**
>
> For backup chains created by jobs that no longer exist, Veeam Plug-in 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 Plug-in for Nutanix AHV performs the following actions at the end of each backup session:

1. Veeam Plug-in for Nutanix AHV checks the configuration database to detects backup chains with restore points that are older than the specified time limit.
2. If a redundant restore point exists in a backup chain, Veeam Plug-in for Nutanix AHV transforms the backup chain in the following way:
   a. Rebuilds the full backup to include there data of the incremental backup that follows the full backup. To do that, Veeam Plug-in 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.

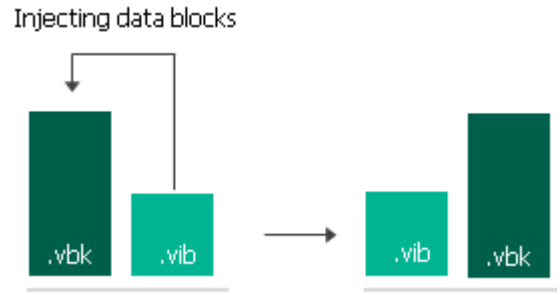


   b. Removes the earliest incremental backup from the chain as redundant — this data has already been injected into the full backup.

3. Veeam Plug-in 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 Plug-in for Nutanix AHV ensures that the backup chain is not broken and that you will be able to recover your data when needed.

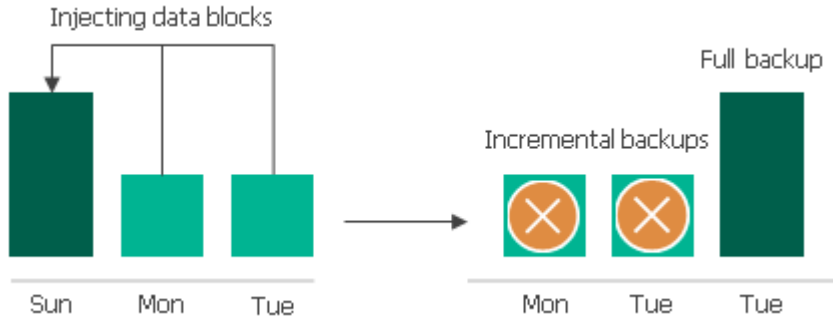


# Forward Incremental Backup Retention Policy

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

1. Veeam Plug-in 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).
2. Veeam Plug-in for Nutanix AHV checks whether the period to keep restore points in the new chain fragment has reached the allowed time limit.
3. If the new backup chain fragment has reached the limit of allowed restore points, Veeam Plug-in for Nutanix AHV removes all restore points of the older backup chain fragment.

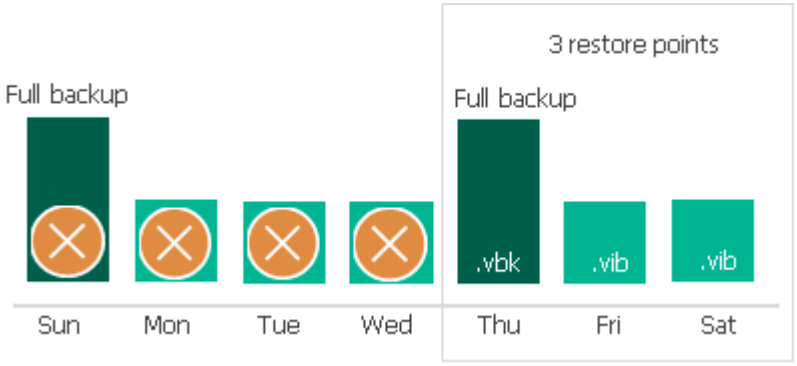

# Backup Methods

Veeam Backup & Replication provides the following methods for creating backup chains:

- **Forever forward incremental**

  When the forever forward incremental backup method is used, Veeam Backup & Replication 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 & Replication 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 & Replication 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 & Replication 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 & Replication implements the forever forward incremental backup:

1. During the first (full) backup session, Veeam Backup & Replication 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 & Replication 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.

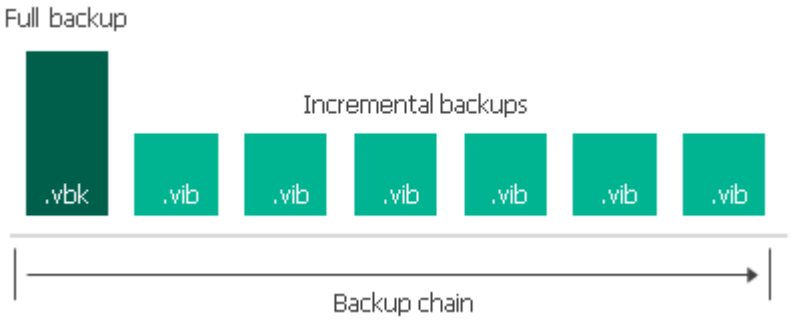


# Forward Incremental Backup

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

1. During the first (full) backup session, Veeam Backup & Replication 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 & Replication 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.

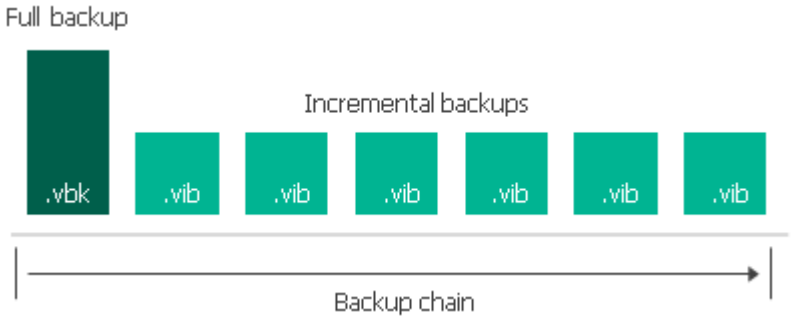


3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup & Replication creates a full backup file and adds it to the backup chain. Incremental restore points produced after this full backup file use it as a new starting point.

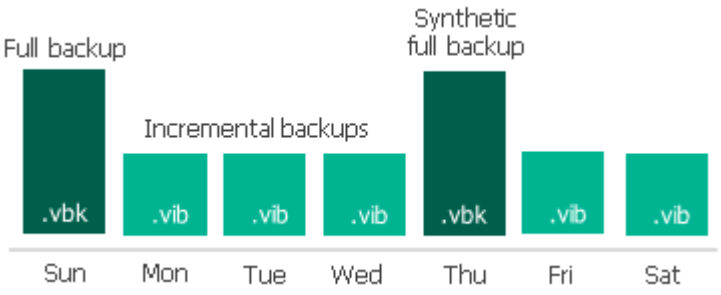

# Active Full Backup

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 & Replication allows you to create active full backups (either manually or automatically according to a specific schedule).

To create an active full backup Veeam Backup & Replication 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 & Replication 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.

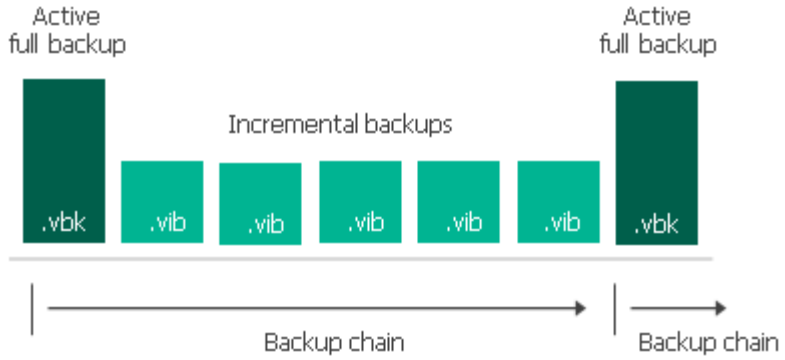


Veeam Backup & Replication triggers a backup job to create an active full backup even if a regular backup session is not scheduled on this day. The active full backup session starts at the same time when the backup job is scheduled. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule active full backup to be created on Saturday, Veeam Backup & Replication 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 & Replication 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 & Replication will produce an active full backup only. However, if you run the backup job again on the same day manually, Veeam Backup & Replication 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 & Replication 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 & Replication performs the following operations:

1. Veeam Backup & Replication creates a regular incremental backup and adds it to the backup chain.

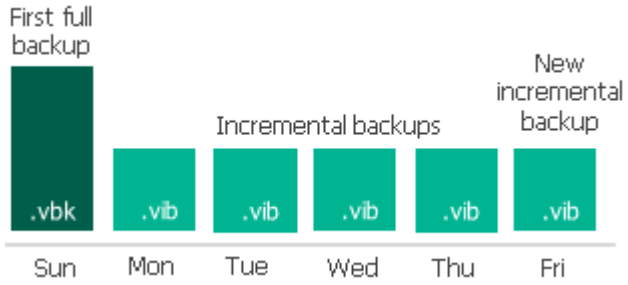


2. Veeam Backup & Replication creates a new synthetic full backup using backup files that are already available in the backup chain, including the newly created incremental backup file.

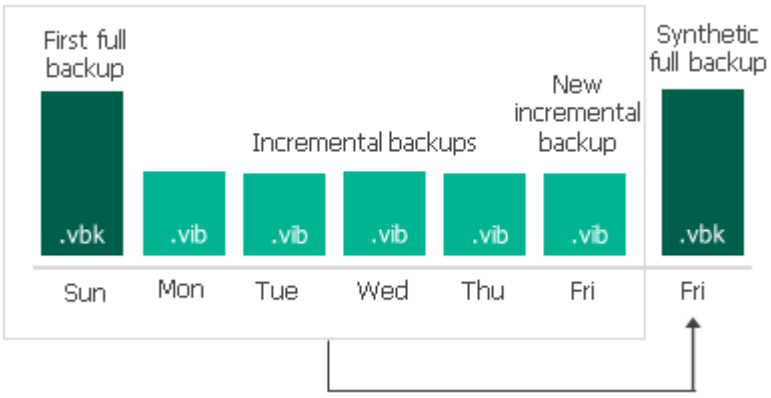


3. Veeam Backup & Replication deletes the created incremental backup as its data is already incorporated in the synthetic full backup.

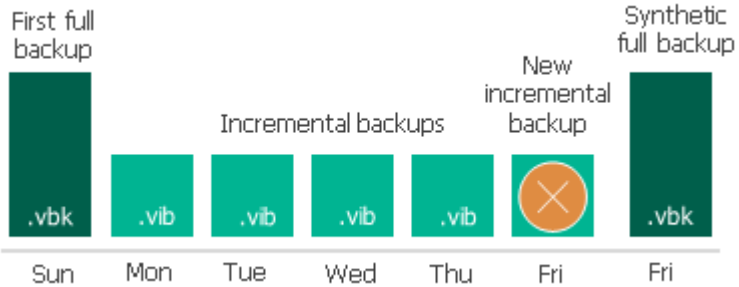


When creating a synthetic full backup, Veeam Backup & Replication starts a new backup chain for the VM. All further created incremental backups use the latest full backup file as a new starting point. The old full backup file from the old backup chain remains on disk until it is automatically deleted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

Veeam Backup & Replication 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 & Replication 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 & Replication 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 & Replication will produce a synthetic full backup only. However, if you run the backup job again on the same day manually, Veeam Backup & Replication will perform incremental backup in a regular manner.

# Snapshot Types

In terms of data protection, Veeam Backup & Replication 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 in the Veeam Backup & Replication console, and can be used to perform entire VM restore and disk restore.

  Backup snapshots allow Veeam Backup & Replication 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 taken manually in the Prism Element or Prism Central console. Snapshots are displayed 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 not displayed 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 & Replication to reduce the backup load on the production environment. However, Veeam Backup & Replication 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 & Replication 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 & Replication takes VG snapshots only if the backup scope includes individual virtual machines with volume groups attached.

  VG snapshots are not displayed in the Veeam Backup & Replication console. VG snapshots allow Veeam Backup & Replication 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 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. Starting from version 8, Veeam Plug-in for Nutanix AHV does not takes PD snapshots — if a PD is included into the backup scope, Veeam Plug-in for Nutanix AHV backs up VMs and their volume groups as if processing individual virtual machines.

> **NOTE**
>
> Recovery points created manually in the Prism Central console cannot be used to protect and recover Nutanix AHV resources with Veeam Backup & Replication.

In terms of data consistency, Veeam Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV creates a crash-consistent snapshot instead.

> **TIP**
>
> As an alternative to application-consistent snapshots, Veeam Plug-in for Nutanix AHV allows you to leverage native Veeam capabilities to create application-consistent backups.

# VM Restore

Veeam Backup & Replication 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 & Replication 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 & Replication launches any other Nutanix AHV worker that is added to the backup infrastructure.
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 & Replication 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 Veeam Plug-in for Nutanix AHV 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 & Replication 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 & Replication 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 & Replication launches any other Nutanix AHV worker that is added to the backup infrastructure.
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 & Replication performs the following steps:

1. Mounts disks of the backed-up VM to the mount host specified for the recovery operation.
2. 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.
3. Restores the selected files and folders to the original location or to a new location.
4. Detaches the disks from the mount host.

To recover VM files and folders from a backup snapshot, snapshot or PD snapshot, Veeam Backup & Replication 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

Image-level backups created by jobs are not kept forever — they are removed according to retention policy specified while creating the jobs as described in section Creating Backup Jobs.

Restore points in the backup chain are stored only for the allowed period of time (in days). If a restore point is older than the specified time limit, Veeam Backup & Replication removes it from the backup chain. To learn how Veeam Backup & Replication applies retention policies to forever forward incremental and forward incremental backup chains, see Backup Retention.

# Planning and Preparation

Before you start using Veeam Backup & Replication, consider the following requirements:

- Hardware and software requirements.
- Permissions that must be assigned to accounts used to deploy and administer backup infrastructure components.
- Network ports that must be open to ensure proper communication of the solution components.
- Sizing Guidelines.

# System Requirements

Before you start using Veeam Plug-in for Nutanix AHV, make sure the Nutanix AHV cluster or Prism Central and the backup infrastructure components meet the following requirements.

| Specification | Requirement |
|---|---|
| Virtualization Platform | • Veeam Plug-in for Nutanix AHV is compatible with:<br>  ○ Nutanix AOS versions 6.8.1.6 and above<br>  ○ Prism Central version pc.2022.6-pc.2024.3.1.10, pc.7.3 and above except 7.3.1.2 and 7.5.0<br>• An IP address of the cluster and the iSCSI Data Service must be configured in Nutanix AHV cluster settings. For more information, see Nutanix documentation.<br>• UEFI boot must be supported in the Nutanix AHV environment. For more information, see Nutanix documentation.<br>• Veeam Plug-in for Nutanix AHV supports Nutanix Cloud Clusters (NC2) used for hybrid multi-cloud deployment. |
| Veeam Software | • Veeam Backup & Replication version 13.0.1.180 (or later) with Nutanix AHV Plug-in version 13.9.0.212 (or later) must be deployed on the backup server. |
| Workers | Workers 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:<br>• CPU: 6 vCPU<br>• Memory: 6 GB RAM<br>• Disk Space: 100 GB for product installation and logs<br><br>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 Sizing Guidelines. |

**IMPORTANT**

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

# Version Compatibility

The following table lists compatible versions of Veeam Backup & Replication and Veeam Backup for Nutanix AHV.

| Product Release | Veeam Plug-in for Nutanix AHV Build | Veeam Backup & Replication Build | Backup Appliance / Worker OS Version |
|---|---|---|---|
| 9 | 13.9.0.212 | 13.0.1.180 | Rocky Linux 9.2 |
| 8 | 13.8.0.582 | 13.0.0.4967 | Rocky Linux 8.10 |
| 7.1 | 12.7.1.12 | 12.3.1.1139 | Rocky Linux 8.10 |
| 7.0 | 12.7.0.172 | 12.3.0.310 | |
| 6.1 | 12.6.1.15 | 12.2.0.334 | |
| 6.0 | 12.6.0.632 | 12.2.0.334 | |
| 5.1 | 12.5.1.8 | 12.1.0.2131<br>12.1.1.56 | Ubuntu 22.04 LTS |
| 5.0 | 12.5.0.465 | 12.0.0.1420<br>12.1.0.2131 | |
| 4a | 12.1.4.5 | 12.0.0.1420 | Ubuntu 20.04 LTS |
| 4.0 | 12.0.4.1020 | 12.0.0.1420 | |

# Considerations and Limitations

When you plan to use Veeam Plug-in for Nutanix AHV, keep in mind the following limitations and considerations.

## Configuration

When configuring Veeam Plug-in for Nutanix AHV, consider the following:

- In the case where a Veeam Backup & Replication certificate is changed it will be necessary to restart the Veeam AHV service in order to facilitate proper internal component communications.
- Veeam ONE 13 supports monitoring, alerting and reporting features for AHV VMs. For the list of supported features, see the What's New document for Veeam ONE 13.
- You can use Veeam Backup Enterprise Manager to file-level restore guest OS files of Nutanix AHV VMs and manage Nutanix AHV VM backup copy jobs. All other operations are not supported.

## Backup Repositories

When configuring repositories for Nutanix AHV VM backups, consider the following:

- Veeam Plug-in 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 Plug-in for Nutanix AHV. You can also use Instant Recovery to restore VMs to Nutanix AHV from those or external repositories.
- If a standalone Veeam Backup & Replication repository extent storing Nutanix AHV VM backups is migrated into a scale-out backup repository, backup jobs will fail. To avoid the issue, update the target repository in backup jobs.
- Only file-level restores are available for backups created by Veeam Plug-in for Nutanix AHV on HPE Cloud Bank Storage

## Workers

When configuring workers, consider the following:

- UEFI boot is required for workers.
- Multiple vNICs may be configured for workers.
- For Prism Central deployments, the "SelfServiceContainer" storage container will always be used for workers.
- By default, installing worker updates is performed on each worker start. However, you can disable automatic updates,for example, if your infrastructure does not have connection to Internet.
- Workers may start on the same host if automatic host affinity is enabled, and the number of workers has exceeded the number of hosts in the cluster.
- [Applies only to the Prism Central deployment] Worker image distribution is optimized so that the image is only populated to a Prism Central-managed cluster when a worker is instantiated on that cluster.
- If you raise the number of concurrent backup tasks on workers, backup jobs may fail due to CVM resource limitations. The CVM on each node of the cluster may need additional resources.

# Backup

When protecting Nutanix AHV resources, consider the following:

- Veeam Plug-in for Nutanix AHV creates forward incremental per-VM backup chains (one backup chain contains data for one VM). When you add several VMs to a backup job, Veeam Plug-in for Nutanix AHV creates individual backup chains on the Veeam backup repository, one for each VM processed by the job. Note that for forward incremental backup chains, you can create active or synthetic full backups. For more information, see the Backup Methods section of the Veeam Backup & Replication User Guide.
- By default, Veeam Plug-in for Nutanix AHV applies the following deduplication and compression settings to backed-up data:
  - Deduplication: *Enabled*
  - Data compression level: *Optimal*
  - Storage optimization: *1MB*

  Due to technical limitations, you cannot change deduplication settings while configuring backup jobs.
- By default, backup encryption is disabled for backed-up data. However, you can enable encryption at the repository level. For more information, see the Access Permissions section of the Veeam Backup & Replication User Guide.
- Since Veeam Backup & Replication does not allow you to assign information about locations to Nutanix AHV clusters, job statistics do not include information on the Nutanix AHV VM data migration between different geographic regions.
- Manual *Move backup* functionality is not supported for Nutanix AHV backups.
- If you add a Prism Central category to a backup job, VMs that are assigned to this category will be processed in size order starting with the largest one.
- If you specify a VM as the source for a backup job, Veeam Plug-in for Nutanix AHV processes volume groups attached to the VM. However, if you back up the VM with the attached volume groups, Veeam Plug-in for Nutanix AHV will create a crash-inconsistent backup.
- Veeam Plug-in for Nutanix AHV does not process volume groups if CHAP authentication is enabled. For more information, see Nutanix documentation.
- Second Health Check of same data corruption returns successful session (disk is skipped from processing)
- Backups cannot be imported from unsupported repository types. This can affect importing from backup copy jobs.
- Backup Copy exclusions cannot be applied to Nutanix AHV jobs and objects.
- For VeeamZIP backups, retention is not supported.
- For VeeamZIP backups, an SMB share that requires authentication cannot be specified as a local or shared folder. However, it can be added to the backup infrastructure and then set as backup repository.
- SureBackup for backups created by Veeam Plug-in for Nutanix AHV is supported in the Backup verification and content scan only verification mode.
- Nutanix CVM cannot be backed up with Veeam Plug-in for Nutanix AHV.
- Veeam Plug-in for Nutanix AHV does not support the Migrate Across Clusters functionality. If a VM is migrated in this way, backup jobs that contain this VM will skip it. If the VM is included into a category, a backup job that protects this category will start a new backup chain for the migrated VM.

# Backup From Replica Cluster

When configuring backup jobs to use backups from replica clusters, consider the following:

- Backup from a replica cluster can be used if a protection policy is configured in Prism Central.
- Backup from replica cluster will not be performed if guest processing (application-aware processing or indexing) is enabled in the backup job.
- If backup cannot be performed from a replica cluster, it will be performed from the original cluster.
- If there is no replicated snapshot on a replica cluster, backup will be performed from the original cluster.
- VMs with volume groups attached cannot be backed up from a replica cluster.
- If VM disks were added or removed after the last replication, VM backup will be performed from the original cluster.
- PD snapshots created using Nutanix AHV Async DR are not supported.

# Guest Processing

When configuring guest processing in backup jobs, consider the following:

- Database log shipping: When upgrading from plug-in version 7.1 the parent backup job must be run to re-initiate log shipping.
- Pre-job and post-job scripts are not supported.
- The .JS, .VBS and .WSF scripts are not supported for pre-freeze and post-thaw (snapshot) operations.
- Persistent guest agent are supported in non-FIPS mode only.
- The file exclusion functionality is not supported.
- Veeam Plug-in for Nutanix AHV cannot use Kerberos authentication while connecting to guest OSes of the processed VMs.
- When restoring a database using Veeam Explorers to the original VM, the VM hostname is used instead of the FQDN name. If Veeam Explorers cannot reach the VM, you can add the FQDN name and the IP address of the VM to the hosts file on the backup server.
- If you import backups created by a job with guest processing enabled, this backup job will truncate transaction logs but it will not store transaction log backups in the repository. To avoid the issue, before running the job, either clone the job and perform active full, or contact contact Veeam Customer Support.
- Image-level, application-aware backups of Veeam Backup for Microsoft 365 servers running on Nutanix AHV clusters are not Veeam Microsoft 365 restore explorers-aware. The behavior described in this article is currently unsupported for AHV backups.

# Restore

When restoring Nutanix AHV resources, consider the following:

- If you restore the VM to Nutanix AOS 7.0 or later, the SCSI Controller and the CPU hot-plug functionality are always enabled (the *scsi_controller_enabled* and *cpu_hotplug_enabled* parameters are automatically set to true).

- If you restore the VM with an affinity policy not to the original cluster, you must manually configure the affinity policy manually before starting the recovered VM. For more information on affinity policies, see Nutanix documentation.
- If you restore the VM from a snapshot of any type, you cannot change the storage container.
- If you restore the VM using entire VM restore, the VM description will not be restored (will be set as empty).
- If you restore the VM from a user snapshot or PD snapshot, you cannot change VM network settings. However, after the VM is restored, you can configure them using the Nutanix Prism console as described in Nutanix documentation. If you choose to restore to different location and choose to disconnect from all networks, the new VM will be created without networks.
- If you restore the VM from a backup stored in the archive tier of the scale-out backup repository, you must first retrieve backup data as described in the Veeam Backup & Replication User Guide, section Retrieving Backup Files. Note that you cannot perform Entire VM restore from backups stored in the archive tier that consists of the Amazon S3 Glacier Instant Retrieval extent. For those backups, you can perform Instant Recovery.
- If you restore the VM from a backup of a VMware, Hyper-V, oVirt KVM, Scale Computing HyperCore or Proxmox VE VM or from a backup created by Veeam Agent, a restored VM may have network connection problems. To resolve the issue, install Nutanix Guest Tools on the restored VM as described in Nutanix documentation.
- You cannot perform Entire VM Restore to Nutanix AHV from a Veeam Cloud Connect repository or an external repository. However, you can use Instant Recovery to restore VMs to Nutanix AHV from Veeam Cloud Connect or external repositories.
- You cannot perform VM restore from a tape to Nutanix AHV. A tape backup needs to be returned to a supported repository to complete the restore operation.
- You cannot perform VM restore from PD snapshots created using Nutanix AHV Async DR.
- If you want to perform file-level restore from volume group disks, you should run FLR from backups of the VM that have the required volume group attached to them.

# Instant Recovery

When restoring Nutanix AHV resources with Instant Recovery, consider the following:

- You can perform instant recovery to VMware and Hyper-V hosts from backups created by Veeam Plug-in for Nutanix AHV. VMware vSphere or Hyper-V hosts must be added to the Veeam Backup & Replication backup infrastructure.
- It is recommended to deploy a dedicated host as a mount server and allocate a minimum of 512 MB of additional RAM for each VM disk that you want to recover at the same time. For example, if you restore a VM with 4 disks, you need an additional 2 GB of RAM on the mount server.
- A Nutanix AHV cluster must be added to the Veeam Backup & Replication backup infrastructure.
- Veeam Plug-in for Nutanix AHV requires 64 MB of RAM for a VM to perform Instant Recovery. For VMs with less than 64 MB of RAM, Veeam Plug-in for Nutanix AHV increases the amount of RAM to 64 MB during the restore process.
- If you perform Instant Recovery using a VM backup stored in the archive tier of the scale-out backup repository, you must first retrieve backup data as described in the Veeam Backup & Replication User Guide, section Retrieving Backup Files. Note that this requirement is not applicable to backups stored in the archive tier that consists of the Amazon S3 Glacier Instant Retrieval extent.

- If you restore a Nutanix AHV VM that has an attached volume group, the disks from the volume group will not be restored.
    - Instant Recovery is not supported:
    - From backups created by Veeam Availability for Nutanix AHV (Veeam Plug-in for Nutanix AHV version 1.0). For those backups, you can perform only entire VM restore.
    - From backups of VMs with the ARM architecture.
    - From file-level backups created by the Kasten platform, Veeam Agent for Linux, Veeam Agent for Microsoft Windows, Veeam Agent for Unix, Veeam Agent for Mac.
- Nutanix VirtIO drivers should be installed initially before the instant recovery process. You cannot add or modify drives in the VM during Instant Recovery launch.

# REST API Limitations

When using REST API for protecting Nutanix AHV resources, consider the following:

- REST API versions earlier than 8 are not supported.
- You cannot obtain information on ordinary backup restore point ID. Use Power Shell instead.
- You cannot use an account with MFA enabled to obtain an authorization token.
- You cannot enable MFA for a user account.

# Permissions

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

## Backup Server Windows Account

The account used to install Veeam Backup & Replication on a Windows-based machine 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 *LocalSystem* account or must have the Local Administrator permissions on the backup server. |

## Nutanix AHV Cluster Administrator Account

The Nutanix AHV administrator account that Veeam Plug-in for Nutanix AHV uses to access the Prism Central or cluster must have privileges of the *Cluster Admin* or *Prism Admin* role. For more information on user access control, see Nutanix documentation.

## Performing Guest Processing

To allow Veeam Backup & Replication to create application-consistent backups of Windows- and Linux-based VMs, the accounts that will be used to perform guest processing operations (such as transaction log truncation and guest file indexing) must have the permissions listed in this section.

> **NOTE**
>
> The Veeam Backup & Replication console does not provide a possibility to restore application data from application-consistent backups — you can do this using Veeam Explorers only. To see the list of permissions that must be granted to accounts that will be used to perform the restore operations, see the Veeam Explorers User Guide.

## Backup Permissions for Windows-Based VMs

For Windows-based VMs, you must choose an account that has administrator privileges. Note that the *Log on as a batch job* permission must be granted to the account and the *Deny log on as a batch job* policy must not be defined. Other permissions depend on applications that you plan to back up:

| Application | Required Permission |
|---|---|
| **Microsoft SQL Server** | To back up Microsoft SQL Server data, the user whose account you plan to use must have the following permissions:<br>• SQL Server instance-level role: *public* and *dbcreator*.<br>• Database-level roles and roles for the model system database: *db_backupoperator*, *db_denydatareader*, *public*;<br>for the master system database — *db_backupoperator*, *db_datareader*, *public*;<br>for the msdb system database — *db_backupoperator*, *db_datareader*, *public, db_datawriter*.<br>• Securables: *view any definition*, *view server state, connect SQL*.<br>**Tip**: If you do not want to assign the permissions gradually, use an account that has local Administrator permissions on the target VM and system Administrator permissions (with the Sysadmin role) on the target Microsoft SQL Server. |
| **Microsoft Active Directory** | The account used to back up Microsoft Active Directory data or a Domain Controller server must be a member of the built-in *Administrators* group.<br>The account used to back up a Read-Only Domain controller can have permissions of a delegated RODC administrator account. For more information, see Microsoft Docs. |
| **Microsoft Exchange** | The account used to back up Microsoft Exchange data must have the local Administrator permissions on the machine where Microsoft Exchange is installed. |
| **Oracle** | The account used to communicate with VM guest OSes must be a member of both the *Local Administrator* group and the *ORA_DBA* group (if OS authentication is used). In addition, if *ASM* is used, then such an account must be a member of the *ORA_ASMADMIN* group (for Oracle 12 and higher).<br>The account used to back up Oracle databases must have have the following permissions:<br>• Oracle account with SYSDBA privileges.<br>You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges.<br>• Account specified for guest processing. That is, the **Use guest credentials** option selected.<br>In this case, the account that was specified at the **Guest Processing** step must be a member of the *ORA_DBA* group. |

| Application | Required Permission |
| --- | --- |
| **Microsoft SharePoint** | The account used to back up Microsoft SharePoint server data must have the Farm Administrator role.<br><br>The account used to back up Microsoft SQL databases of the Microsoft SharePoint Server must have the same privileges as that of Microsoft SQL Server. |

**TIP**

The account must be specified either in the *DOMAIN|USERNAME* (for Active Directory accounts) or in the *HOST|USERNAME* (for local user accounts) format.

## Backup Permissions for Linux-Based VMs

For Linux-based VMs, you must choose an account of a root user or a user elevated to root. Note that the account must have the `/home` directory created. Other permissions depend on applications that you plan to back up:

| Application | Required Permission |
| --- | --- |
| **Oracle** | The account used to back up Oracle databases must have have the following permissions:<br>• Oracle account with SYSDBA privileges.<br><br>You can use, for example, the SYS Oracle account or any other Oracle account that has been granted SYSDBA privileges.<br>• Account specified for guest processing. That is, the **Use guest credentials** option selected.<br><br>In this case, the account that was specified at the **Guest Processing** step must be a member of the *OSASM*, *OSDBA* and *OINSTALL* groups.<br><br>**Note**: To perform guest processing of Oracle databases running on Linux servers, make sure that the `/tmp` directory is mounted with the `exec` option. Otherwise, you will get a permission denial error. |
| **PostgreSQL** | The account used to back up PostgreSQL instances must have superuser privileges for the PostgreSQL instance. For more information, see PostgreSQL documentation.<br><br>The following permissions must be granted to access the folder used as a temporary location for archive logs:<br>• The user running the PostgreSQL instance must have *read*, *write*, and *execute* (*rwx*) permissions.<br>• The user selected in the backup job settings to access the guest OS must have *read* and *execute* (*rx*) permissions. |

# Ports

Veeam Backup & Replication automatically creates firewall rules for the ports required to allow communication between workers and the backup server.

## 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 Prism Central and standalone clusters | TCP/HTTPS | 9440 | Used to communicate with Nutanix AHV REST API (clusters and Prism Central). |
| | 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. |
| | Nutanix AHV server (cluster virtual IP addresses, cluster iSCSI Data Services IP address, cluster CVM IP addresses) | TCP/iSCSI | 3205, 3260 | Used to access disks attached to Nutanix AHV VMs. |
| | Veeam backup repository (or gateway server) | 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. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | Veeam Update Repository (repository.veeam.com)<br><br>Amazon CloudFront (cloudfront.net, amazonaws.com) | TCP/HTTPS | 443 | Used to download worker update packages.<br><br>**Note**: Veeam Update Repository uses the Amazon CloudFront service to distribute traffic when downloading product updates. |

# 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 |
|---|---|---|---|---|
| Backup server | Backup server | TCP/HTTPS | 6172 | Used by the Platform Service to enable communication with the Veeam Backup & Replication database. |
| | Worker | TCP | 19000 | Used to communicate with workers. |
| | Worker | TCP/HTTPS | 443 | Used by the Platform Service to enable communication with the Veeam Updater service on the worker. |
| | Nutanix AHV cluster | TCP/HTTPS | 9440 | Used by the Platform Service to connect to a Nutanix AHV cluster. |
| | FLR helper appliance | TCP | 22<br>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 Used Ports. |
| 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. |

**NOTE**

For the list of ports used by the backup server to communicate with other backup infrastructure components, 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 the mount server role running vPower NFS Service | TCP<br>UDP | 111 | Used by the Port Mapper service. |
| | | TCP<br>UDP | 1058+ or 1063+ | Used as default mount port. The number of port depends on where the vPower NFS Service is located:<br>• 1058+: If the vPower NFS Service is located on the backup server.<br>• 1063+: If the vPower NFS Service is located on a separate Microsoft Windows machine.<br>If port 1058/1063 is occupied, the succeeding port numbers will be used. |
| | | TCP<br>UDP | 2049+ | Used as NFS port. If port 2049 is occupied, the succeeding port numbers will be used. |

# Guest Processing Components

## Connections with Non-Persistent Runtime 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. |
| | VM guest OS (Microsoft Windows) | TCP | 445<br>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 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>Used by the runtime process deployed inside the VM for guest OS interaction.<br><br>**Note:** 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 *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| | 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 (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. |
| | VM guest OS (Microsoft Windows) | TCP | 445<br>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 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>Used by the runtime process deployed inside the VM for guest OS interaction.<br><br>**Note:** 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 *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| VM guest OS | Guest interaction proxy or backup server | TCP | 2500 to 3300 | Default range of ports used as transmission channels for log shipping.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |

## Connections with Persistent Agent Components

The following table describes network ports that must be opened to ensure proper communication of the guest interaction with the persistent agent components deployed inside the VM guest OS for application-aware processing and indexing.

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server (Windows- | VM guest OS | TCP | 6160<br>11731 | Default port and failover port used by Veeam Installer Service. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| based) or Guest interaction proxy (Windows-based) | (Microsoft Windows) | TCP | 6173<br>2500 | Used by the Veeam Guest Helper for guest OS processing and file-level restore. |

# 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 | Log shipping server | TCP | 445<br>135 | Required for deploying Veeam Backup & Replication components. |
| | | TCP | 6160 | Default port used by Veeam Installer Service. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br><br>**Note:** 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 *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| Log shipping server | Backup repository or gateway server | TCP | 2500 to 3300 | Default range of ports used for communication with a backup repository and transfer log backups.<br><br>By default, the log shipping server connects to the backup repository. However, if the target repository uses a gateway server, the connection will be established with that instead.For more information, see the Veeam Backup & Replication User Guide, section Gateway Servers.<br><br>**Note:** This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |

## MS SQL Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Guest interaction proxy | MS SQL VM guest OS | TCP | 445<br>135 | [Applies to non-persistent runtime components only] Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component. |
| | | TCP | 6160,<br>11737 | [Applies to persistent agent components only] Default port and failover port used by Veeam Installer Service. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 49152 to 65535 | Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br>**Note:** 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 *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| | | TCP | 6167 | Used by the Veeam Log Shipping Service for preparing the database and taking logs. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| MS SQL VM guest OS | Guest interaction proxy | TCP | 2500 to 3300 | Default range of ports used for communication with a guest interaction proxy.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | 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.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## Oracle Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server or Guest interaction proxy | Oracle VM guest OS (Microsoft Windows) | TCP | 135<br>445 | [Applies to non-persistent runtime components only] Required for deploying Veeam Backup & Replication components including Veeam Log Shipper runtime component. |
| | | TCP | 6160, 11737 | [Applies to persistent agent components only] Default port and failover port used by Veeam Installer Service. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 49152 to 65535 | [Applies to non-persistent runtime components only] Dynamic RPC port range for Microsoft Windows 2008 and later. For more information, see this Microsoft KB article.<br>**Note:** 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 *"RPC function call failed"* error, you need to configure dynamic RPC ports. For more information on how to configure RPC dynamic port allocation to work with firewalls, see this Microsoft KB article. |
| | | TCP | 6167 | Used by the Veeam Log Shipping Service for preparing the database and taking logs. |
| | | TCP | 22 | Default SSH port used as a control channel. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | Oracle VM guest OS (Linux) | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | [Applies to persistent agent components only] Default Management Agent port. Required if it is used as a control channel instead of SSH. |
| 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.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | 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.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

## PostgreSQL Guest OS Connections

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| Backup server | PostgreSQL VM guest OS | TCP | 22 | [Applies to non-persistent runtime components only] Default SSH port used as a control channel. |
| | | TCP | 2500 to 3300 | Default range of ports used for communication with a guest OS. |
| | | TCP | 6162 | [Applies to persistent agent components only] Default Management Agent port. Required if it is used as a control channel instead of SSH. |
| PostgreSQL VM guest OS | Backup server | TCP | 2500 to 3300 | Default range of ports used for communication with a guest interaction proxy.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

| From | To | Protocol | Port | Notes |
|---|---|---|---|---|
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | 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.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |
| | | TCP | 6162 | Default port used by Veeam Transport Service (on Linux servers) or Veeam Data Mover Service (on Windows servers). |
| | Log shipping server | TCP | 2500 to 3300 | Default range of ports used for communication with a log shipping server and transfer log backups.<br>**Note**: This range of ports applies to newly installed Veeam Backup & Replication starting from version 10.0, without upgrade from previous versions. If you have upgraded from an earlier version of the product, the range of ports from 2500 to 5000 applies to the already added components. |

# 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 Plug-in 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.

## Workers

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

- It is recommended that workers are deployed in each Nutanix AHV cluster. If no worker is deployed in the cluster, performance of backup operations will be affected as Veeam Backup & Replication will use 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 & Replication 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 Plug-in 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 Plug-in 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 Plug-in 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.

# 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 Veeam Plug-in for Nutanix AHV 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 Plug-in for Nutanix AHV system requirements.

# 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 Veeam Plug-in for Nutanix AHV, do the following:

1. Log in to the backup server using an account with the local Administrator permissions.
2. Download the product installation file `VeeamPluginNutanixAHV_13.9.0.212.zip` from the Veeam downloads page.
3. Open the downloaded archive file and launch the `VeeamPluginNutanixAHV_13.9.0.212.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 **Veeam Plug-In for Nutanix AHV** 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.

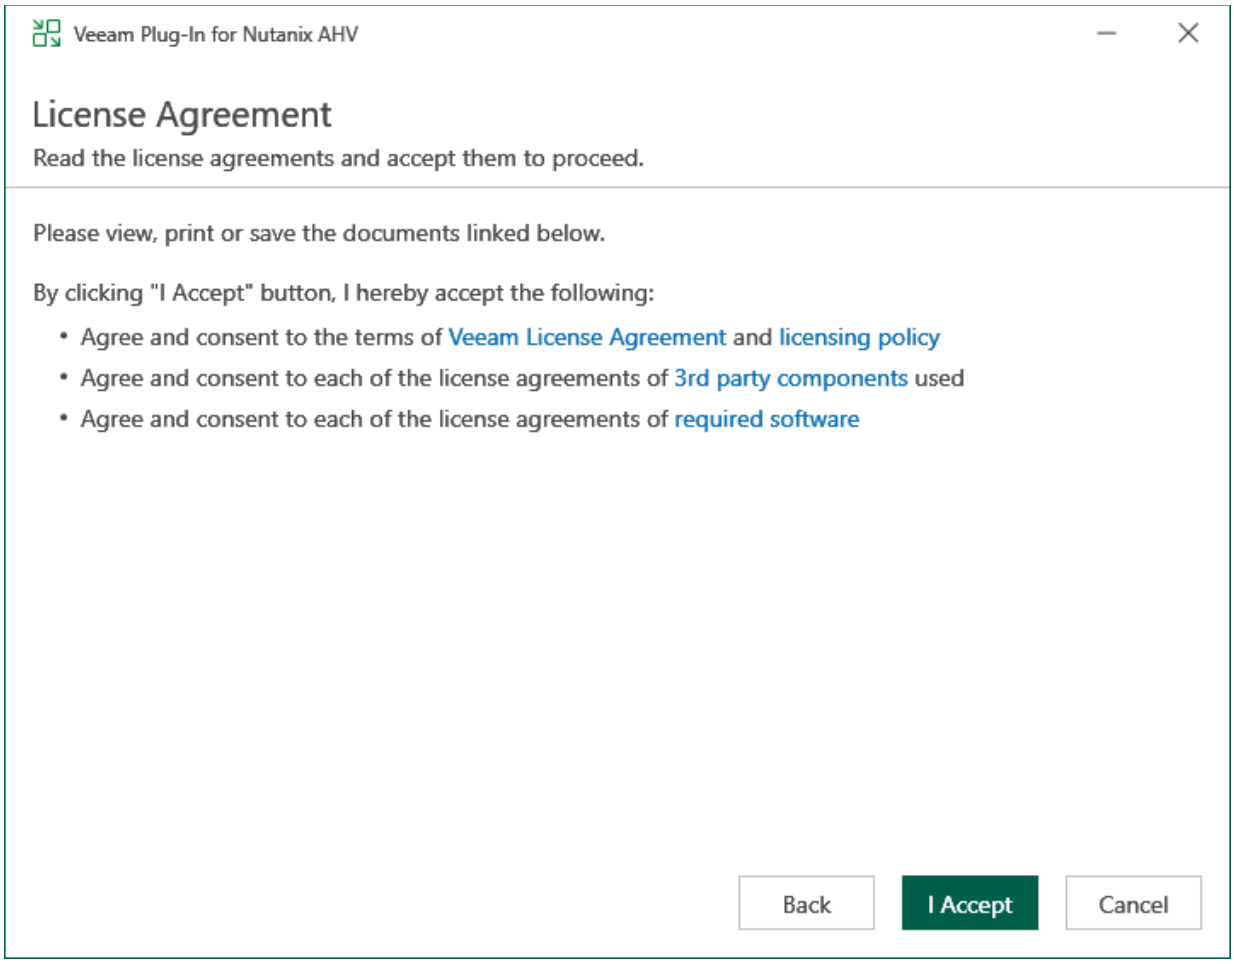

5. At the **Data Location** step of the wizard, you can change the installation directory if necessary.

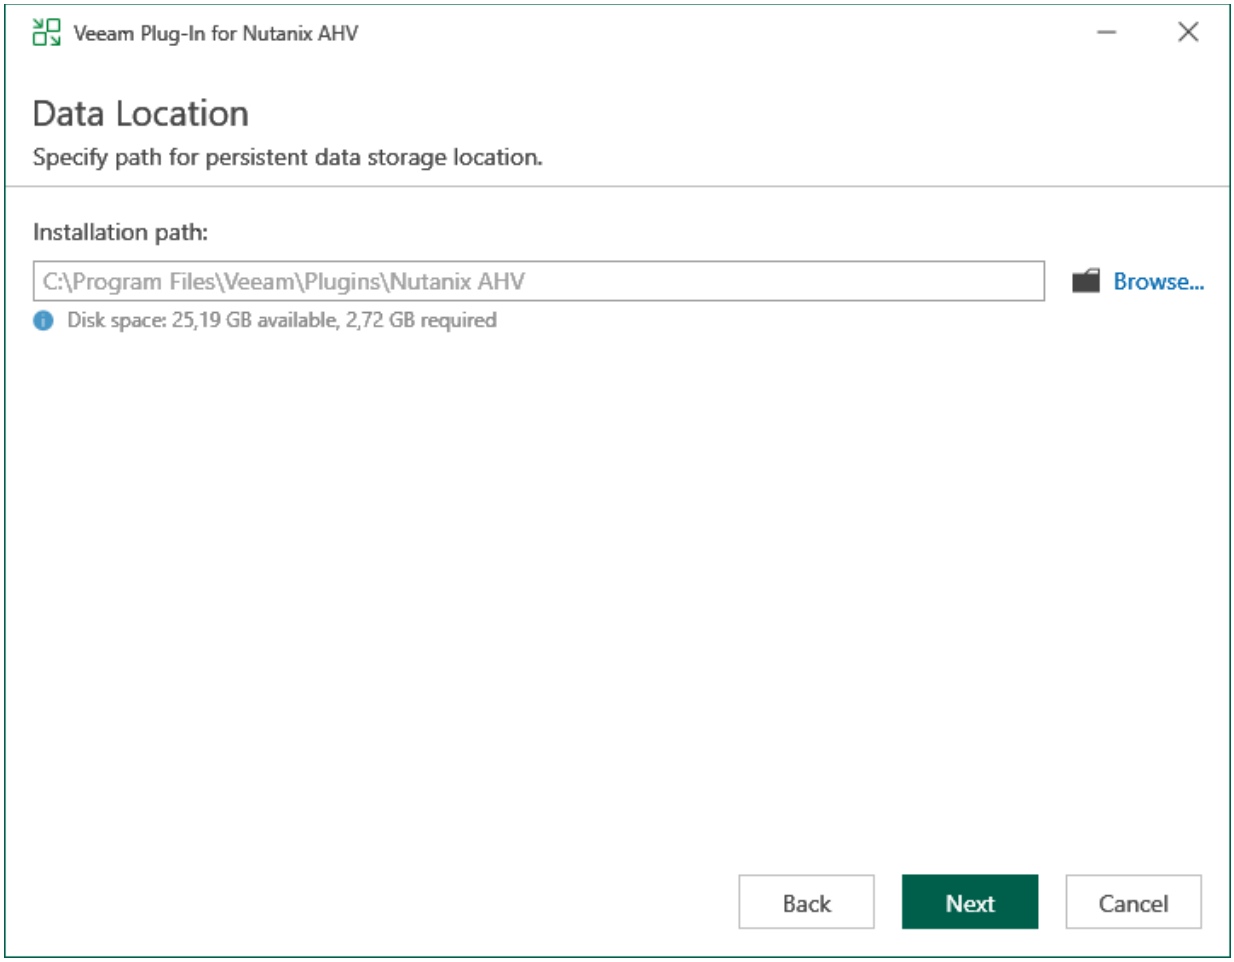

6. Click **Install** to begin installation.

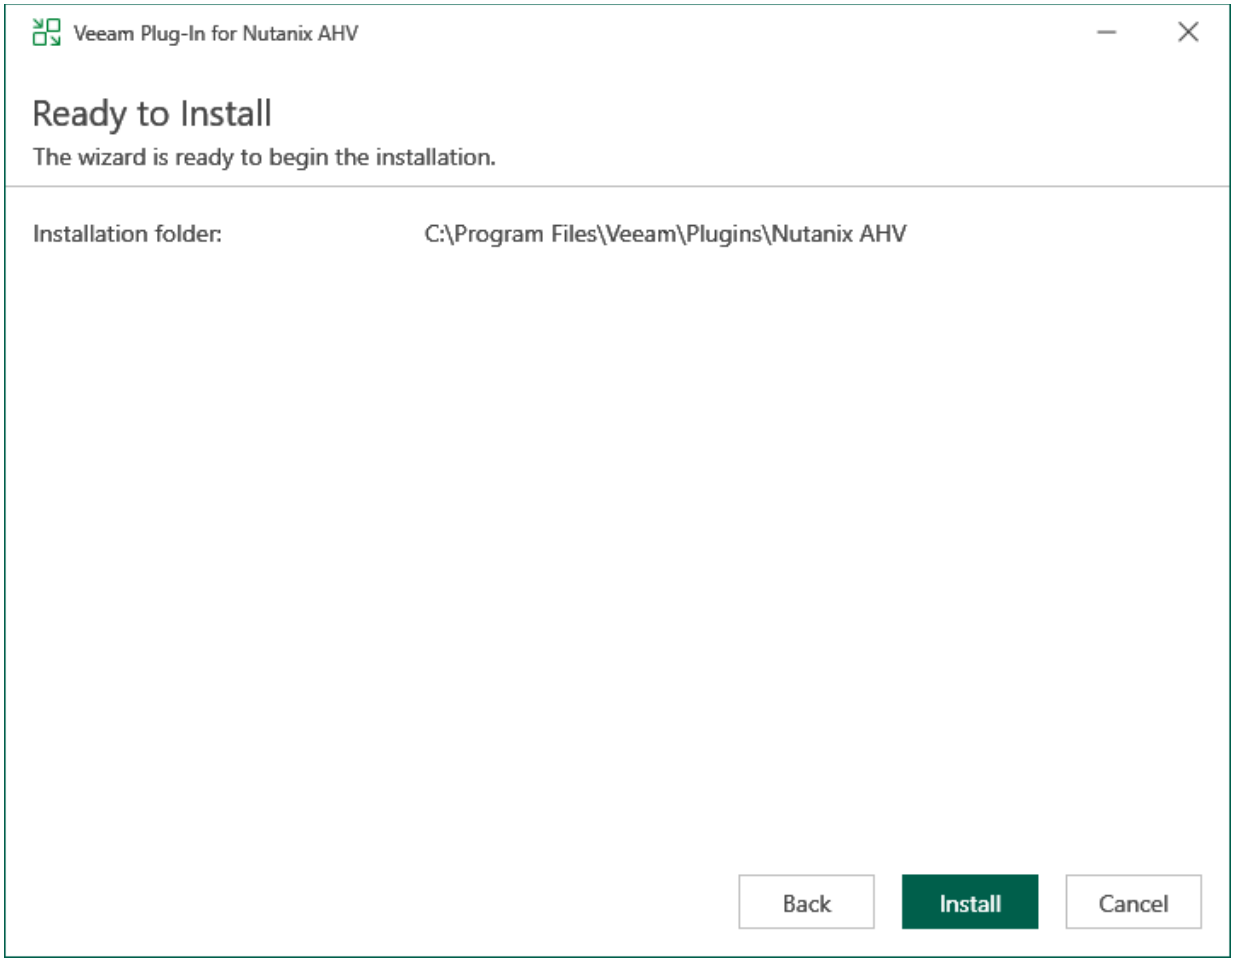

# Installing Plug-In in Unattended Mode

You can install Veeam Plug-in for Nutanix AHV 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 Veeam Plug-in for Nutanix AHV installation process in large-scale environments.

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

- If Veeam Plug-in for Nutanix AHV 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 Veeam Plug-in for Nutanix AHV 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 `VeeamPluginNutanixAHV_13.9.0.212.EXE` file as described in section Installing Nutanix AHV Plug-In Manually (steps 1-3).
2. Check compatibility of the Veeam Plug-in for Nutanix AHV 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 /ac
ceptrequiredsoftware
```

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

| Parameter | Required | Description |
|---|---|---|
| `%path%` | Yes | Specifies a path to the installation .EXE file on the backup server or in a network shared folder. |
| `/silent` | Yes | Sets the user interface level to *None*, which means no user interaction is needed during installation. |
| `/accepteula` | Yes | Confirms that you accept the terms of the Veeam license agreement. |
| `/acceptthirdpartylicenses` | 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 Veeam Plug-in for Nutanix AHV:

```
VeeamPluginNutanixAHV_13.9.0.212.exe /silent /accepteula /acceptthirdpartylicen
ses /acceptlicensingpolicy /acceptrequiredsoftware
```

The following command repairs Veeam Plug-in for Nutanix AHV:

```
VeeamPluginNutanixAHV_13.9.0.212.exe /silent /accepteula /acceptthirdpartylicen
ses /acceptlicensingpolicy /acceptrequiredsoftware /repair
```

The following command uninstalls Veeam Plug-in for Nutanix AHV:

```
VeeamPluginNutanixAHV_13.9.0.212.exe /silent /accepteula /acceptthirdpartylicen
ses /acceptlicensingpolicy /acceptrequiredsoftware /uninstall
```

Veeam Plug-in for Nutanix AHV provides the following status codes to report about the installation result:

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

**TIP**

For detailed logs of Veeam Plug-in for Nutanix AHV 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`

# Uninstalling Nutanix AHV Plug-In

After you remove Veeam Plug-in for Nutanix AHV, 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.

> **TIP**
>
> Before you uninstall Veeam Plug-in for Nutanix AHV, it is recommended to remove all workers from the backup infrastructure.

To uninstall Veeam Plug-in for Nutanix AHV, 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 **Veeam Plug-in for Nutanix AHV**. Then, click **Uninstall**.
5. In the opened window, click **Remove**.

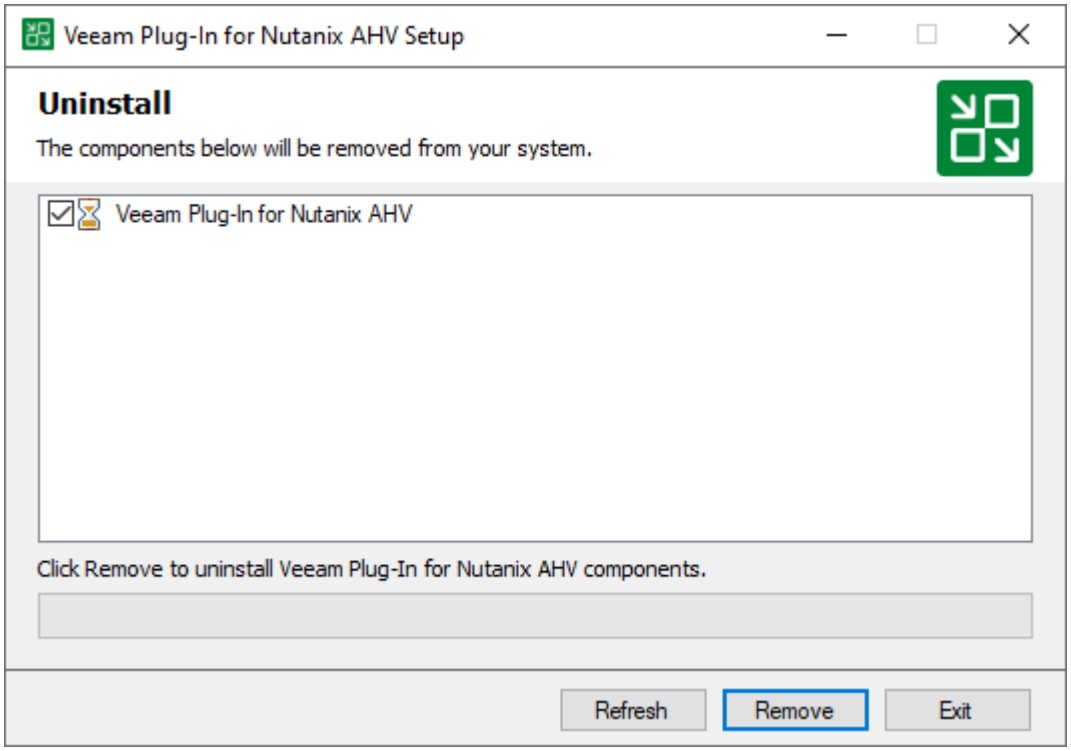

# Upgrading to Veeam Plug-in for Nutanix AHV 9

You can upgrade Veeam Plug-in for Nutanix AHV from version 7.1 or 8 to version 9.

> **IMPORTANT**
>
> To upgrade Veeam Backup for Nutanix AHV version 4.0, 4a, 5.0, 5.1, 6, 6.1 and 7 to version 9, you must first upgrade it to version 7.1 as described in the Veeam Backup for Nutanix AHV 7 User Guide, section Upgrading to Veeam Backup for Nutanix AHV 7.1.

Before you start the upgrade process, do the following:

- [Applies only to Veeam Plug-in for Nutanix AHV version 7.1] Upgrade your Veeam Backup & Replication server to version 12.3.1 or 12.3.2 as described in the Veeam Backup & Replication User Guide, section Upgrading to Veeam Backup & Replication 12.
- Download Veeam Backup & Replication version 13.0.1 from the Veeam downloads page.
- [Applies only to Veeam Plug-in for Nutanix AHV version 7.1] Make sure the Nutanix AHV backup appliance is powered on.

To upgrade Veeam Plug-in for Nutanix AHV to version 9, do the following:

1. Upgrade your Veeam Backup & Replication server to version 13.0.1 as described in the Veeam Backup & Replication User Guide, section Upgrading to Veeam Backup & Replication 13.
2. Complete the **Components Update** wizard as described in the Veeam Backup & Replication User Guide, section Server Components Upgrade.

   After the upgrade process completes, you can switch from the standalone cluster deployment to the Prism Central deployment. To do that, add the Prism Central to the backup infrastructure as described in section Adding Nutanix AHV Server.

> **NOTE**
>
> [Applies only to Veeam Plug-in for Nutanix AHV version 7.1] Since Veeam Plug-in for Nutanix AHV version 9 comes without the backup appliance component, Veeam Backup & Replication will automatically import all the existing configuration settings of your current appliance to the Veeam Backup & Replication configuration database and then remove the backup appliance VM from the Nutanix cluster as soon as the upgrade process completes successfully. If the appliance was used to deliver backup traffic, Veeam Backup & Replication will also deploy a new worker VM and configure all the necessary settings required to perform backup and restore operations — you will be able to adjust these settings later when configuring Veeam Plug-in for Nutanix AHV as described in section Editing Workers.

# Configuring Veeam Plug-in for Nutanix AHV

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

1. Configure backup repositories where Veeam Plug-in 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 workers that will transfer backup traffic.
4. [Optional] Configure email settings and notifications.

# Configuring Backup Repositories

A backup repository is a storage location where Veeam Backup & Replication 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 also 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 & Replication 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 Plug-in 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 Plug-in for Nutanix AHV.
> - [For scale-out backup repositories] Due to specifics of backup jobs for Nutanix AHV VMs, Veeam Plug-in 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 workers and to manage data protection tasks for Nutanix AHV VMs and protection domains.

# Adding Nutanix AHV Server to Backup Infrastructure

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. Apply Nutanix AHV server settings.
5. Finish working with the wizard.

## Considerations and Limitations

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

- If some Prism Central clusters were already added to the backup infrastructure as standalone clusters, they will be automatically configured as clusters registered with the Prism Central. For more information on, see Prism Central Deployment Scenario.
- If you register a new cluster with the Prism Central, Veeam Backup & Replication 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.

# 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.

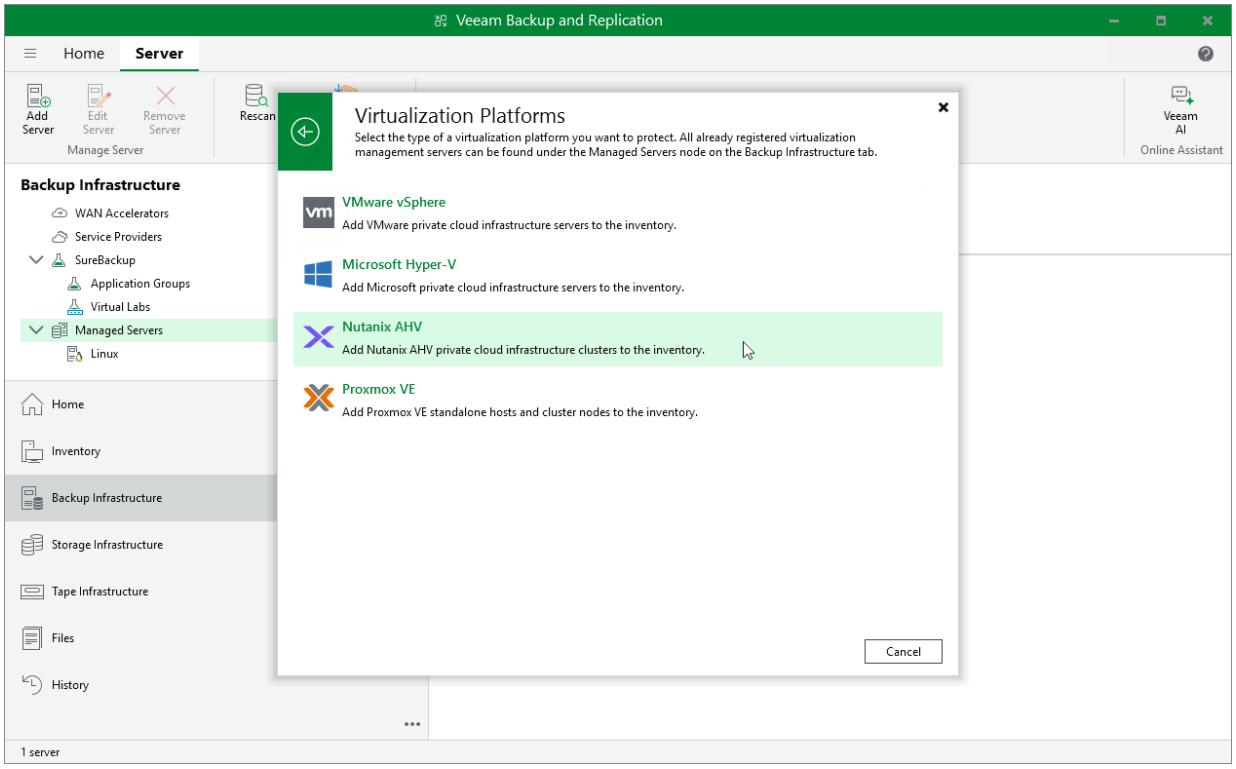

# Step 2. Specify Server Domain Name or Address

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.

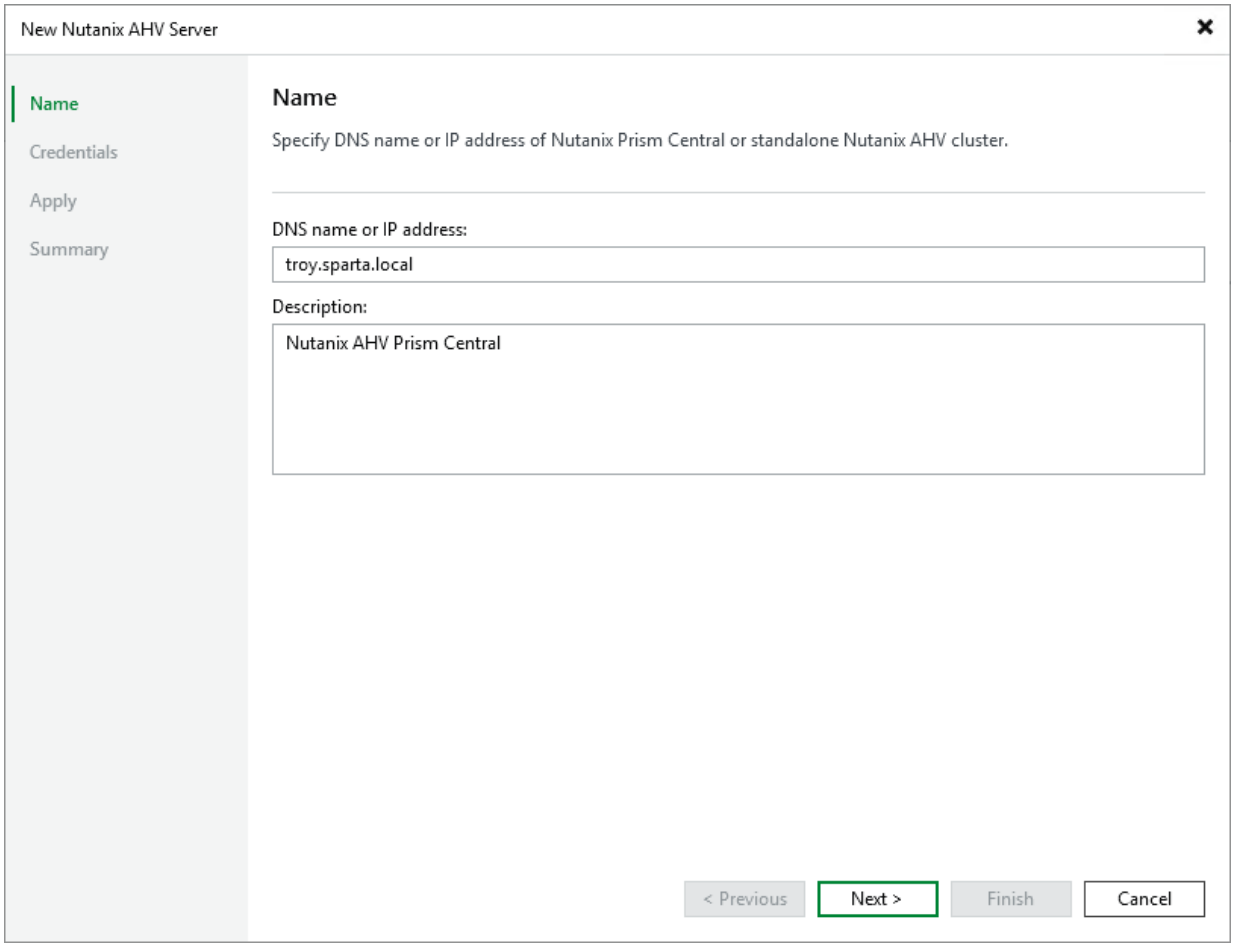

# Step 3. Specify Credentials

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 Sever** 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**.

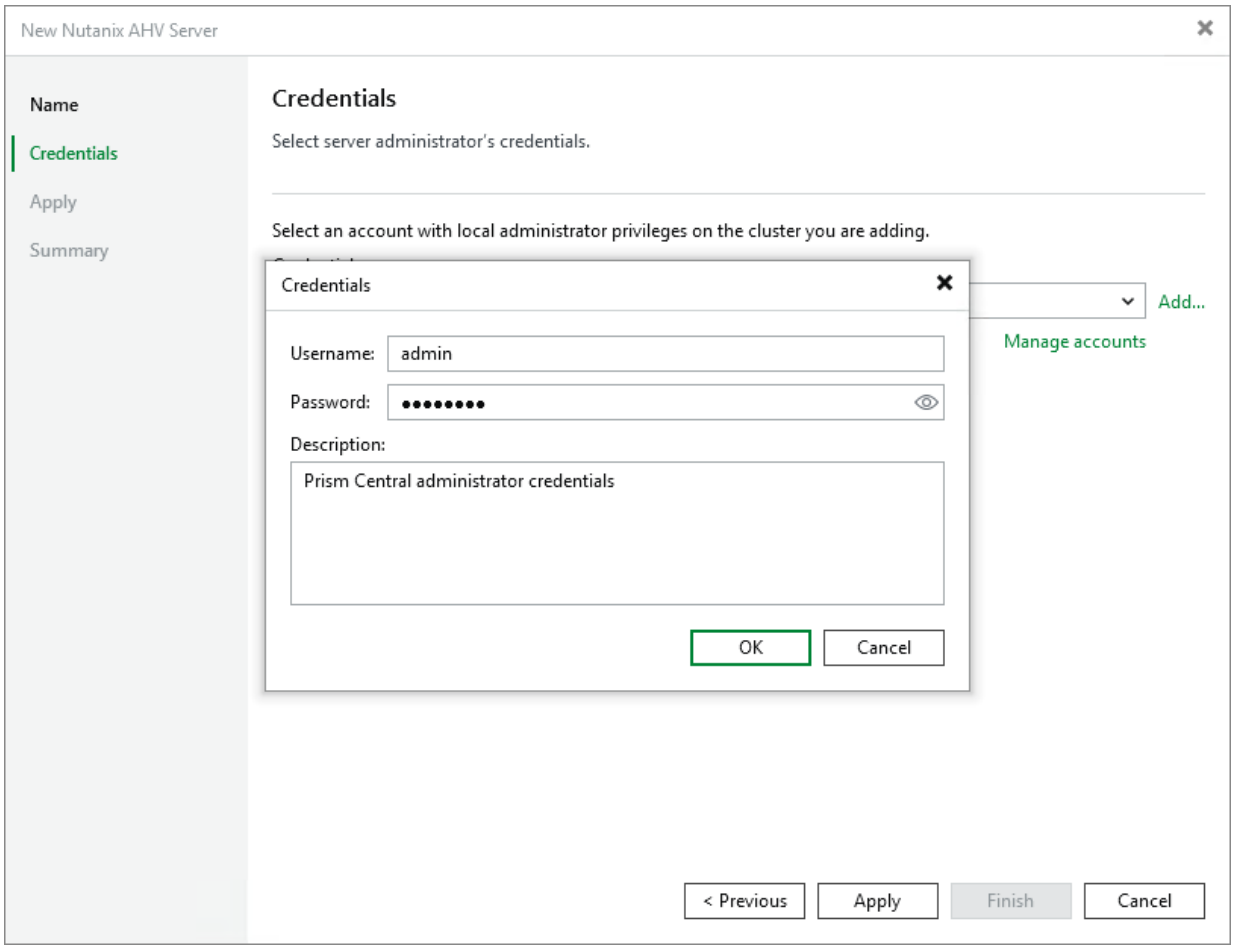

# Step 4. Apply Settings

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

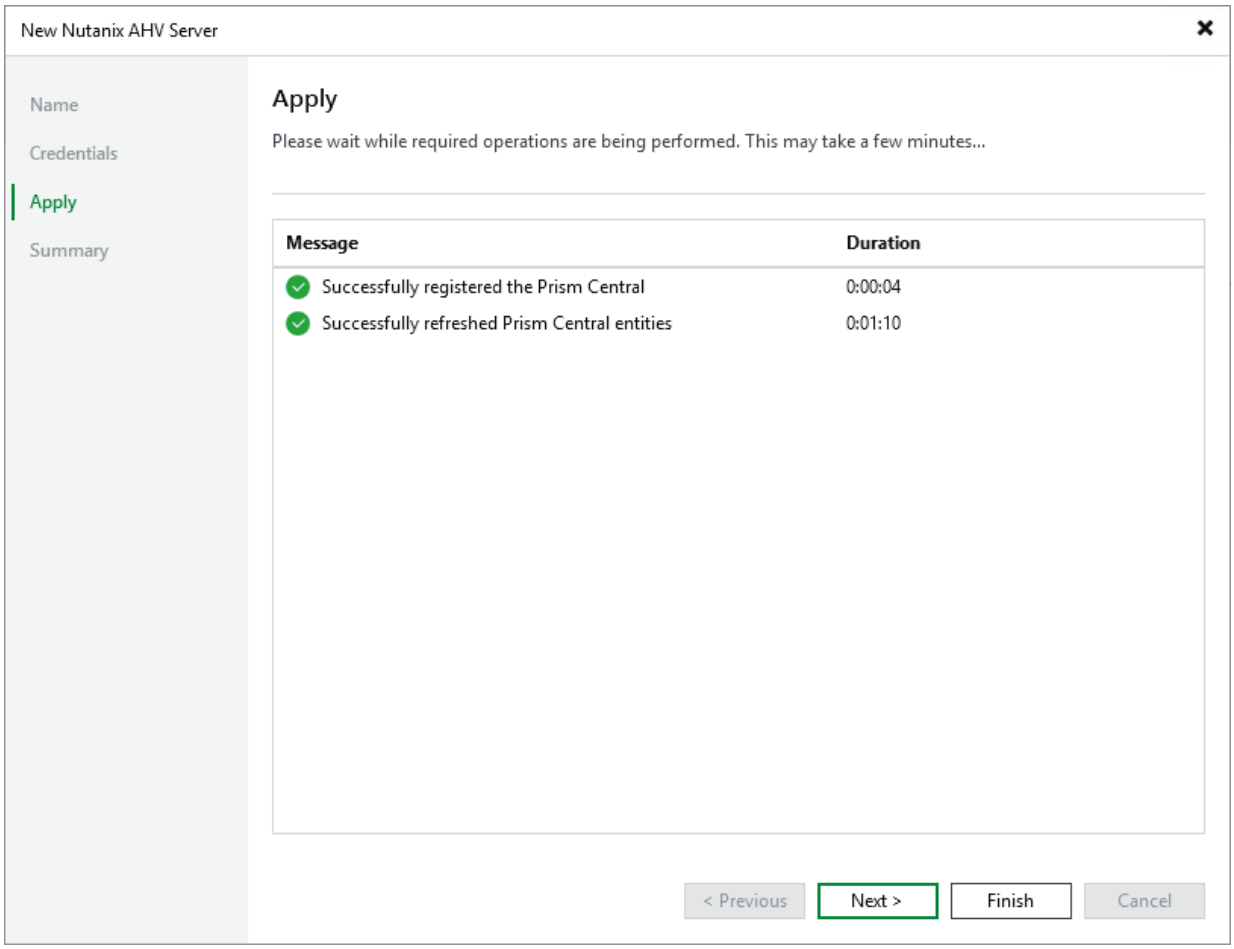


# Step 5 Finish Working with Wizard

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 at least one worker. You can proceed to the **New Nutanix AHV Worker** wizard immediately, or launch the wizard later as described in section Managing Workers.

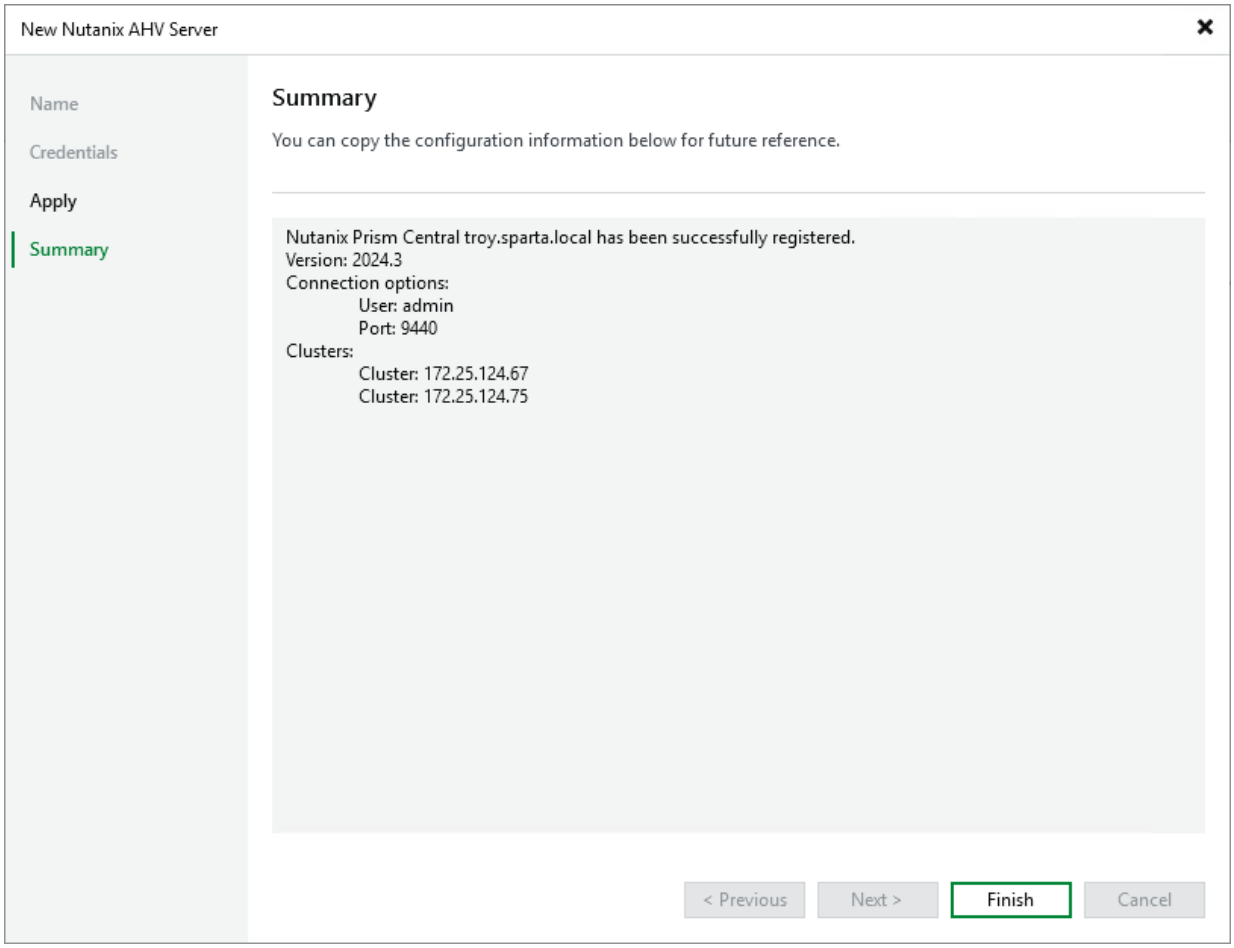

# Editing Nutanix AHV Server Properties

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 to Backup Infrastructure.

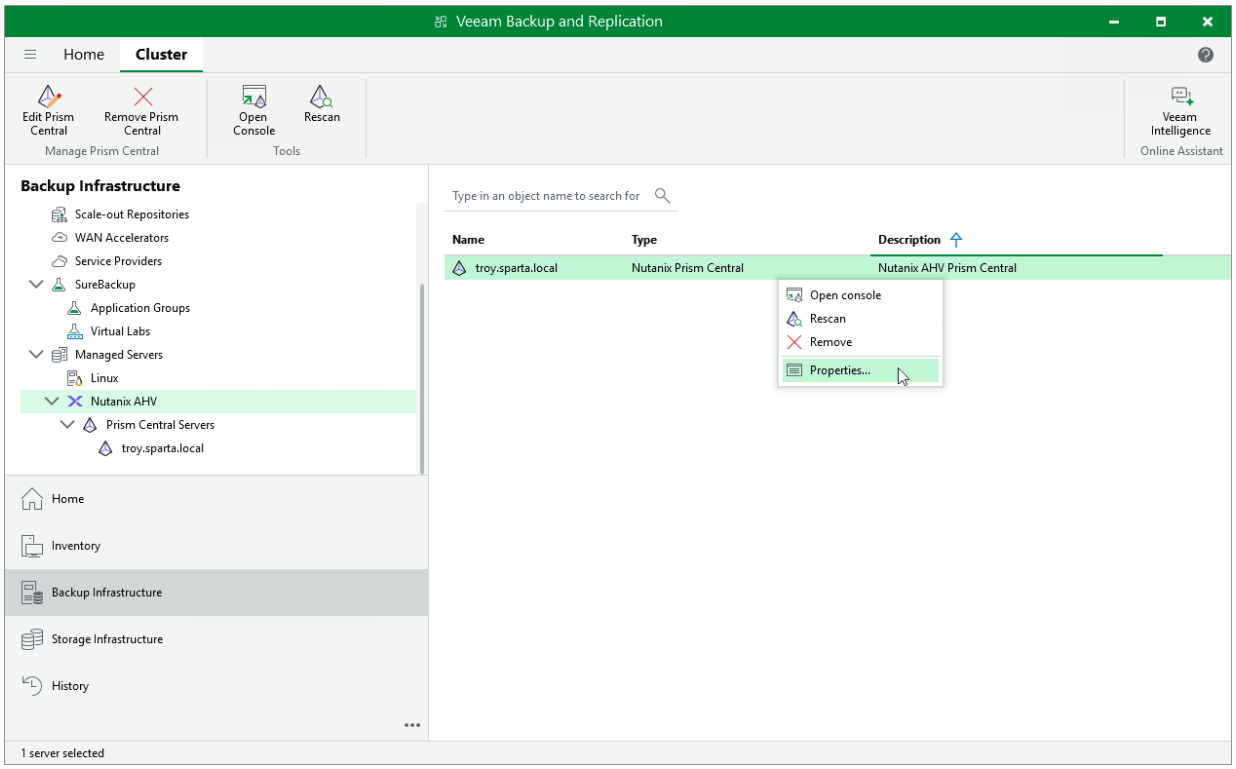

# Rescanning Nutanix AHV Server

Veeam Backup & Replication 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 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 Prism Central or Nutanix AHV cluster and click **Rescan** on the ribbon, or right-click the Prism Central or 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.

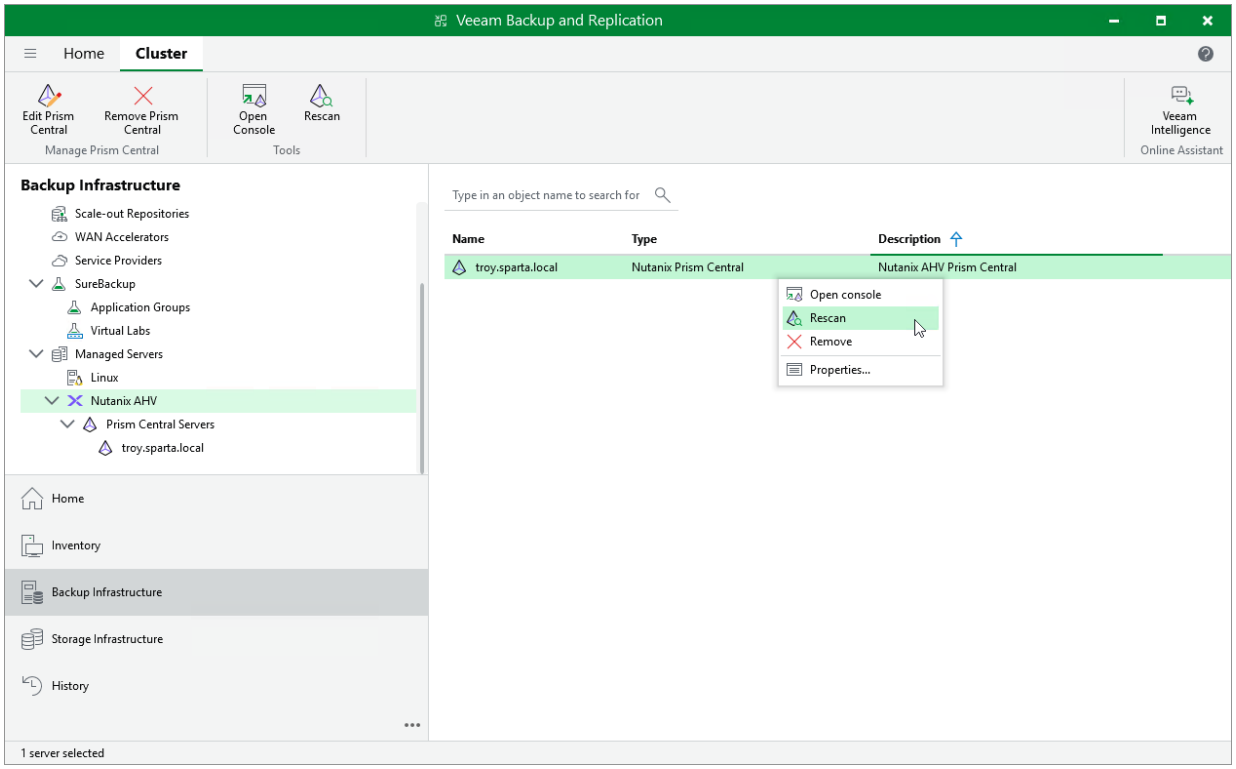

# Removing Nutanix AHV Server

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.

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**.

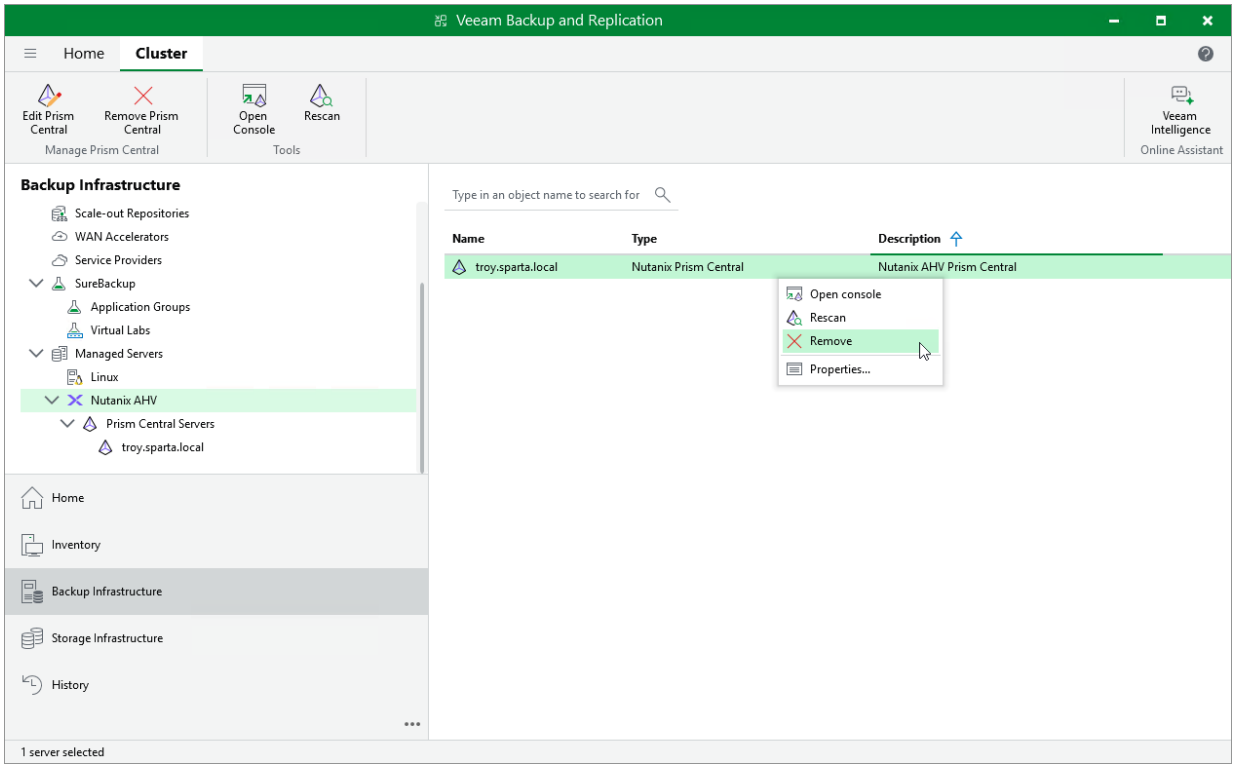

# Accessing Nutanix AHV Server Console

If you want to check the configuration of your Nutanix AHV infrastructure, you can use Veeam Backup & Replication 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 > Nutanix AHV**.
3. Select the Prism Central and click **Open Console** on the ribbon, or right-click the Prism Central and select **Open Console**.

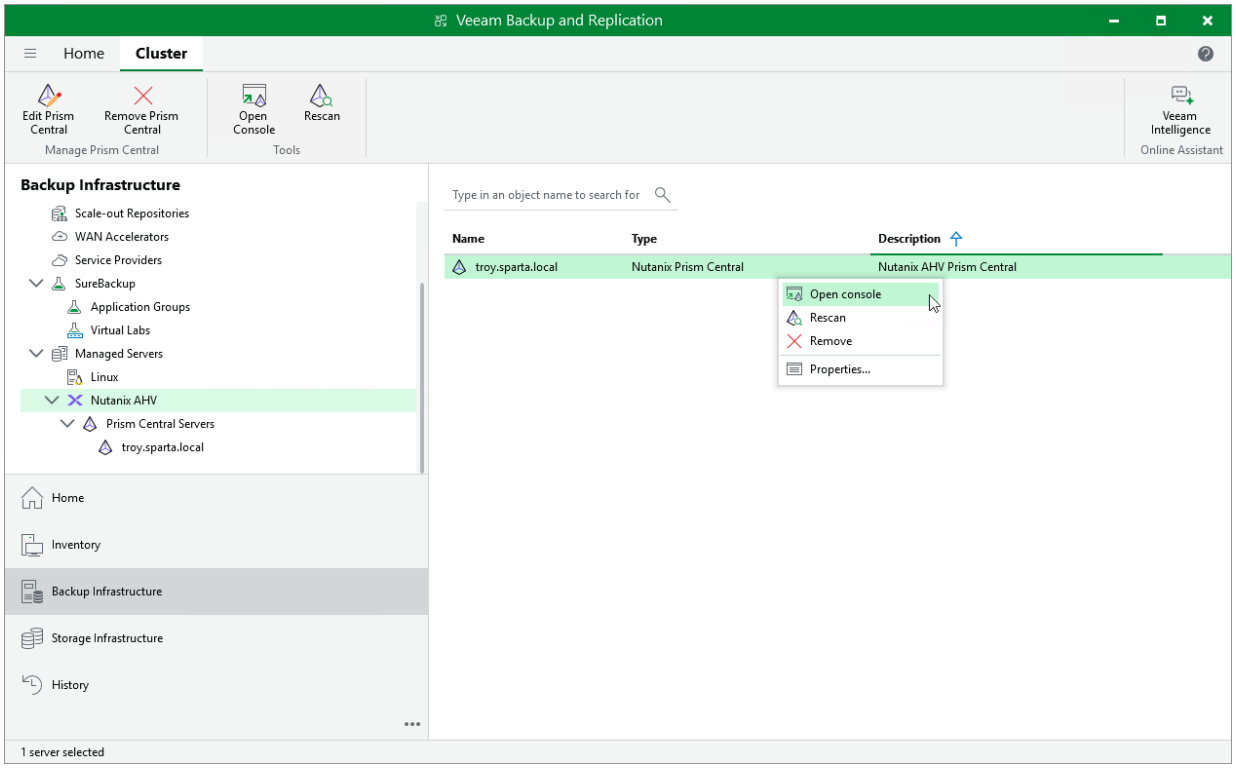


To access the Prism Element console, do the following:

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

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**.

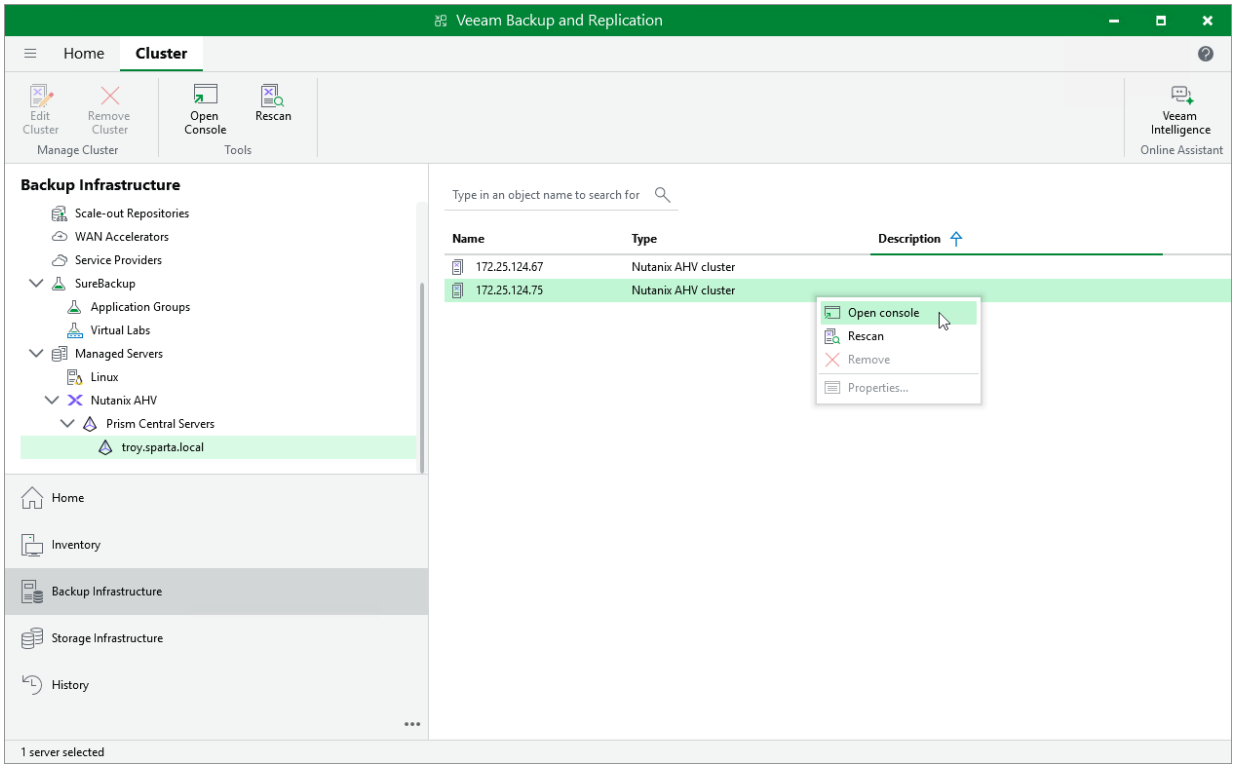

# Managing Workers

To perform most data protection and disaster recovery operations (such as creating image-level backups and restoring backed-up data from them), Veeam Backup & Replication uses workers. Workers are Linux-based VMs that process backup workload and distribute backup traffic when transferring data to backup repositories.

To protect Nutanix AHV resources with Veeam Backup & Replication, you must configure at least one worker. However, it is recommended to deploy multiple workers to increase the maximum number of concurrent backup and restore operations.

Each 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 & Replication 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 & Replication to optimize data transfer and to balance the load on the hosts in the Nutanix AHV cluster. In the latter case, Veeam Backup & Replication 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 worker to the backup infrastructure, its configuration is saved to the Veeam Backup & Replication 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 & Replication tries to launch the worker and test its configuration. If no worker VM has been previously deployed, Veeam Backup & Replication deploys the VM using the worker configuration saved to the configuration database. Veeam Backup & Replication checks host affinity settings specified for the worker and chooses a host where the worker VM will run. Then, Veeam Backup & Replication powers on the worker VM and installs system updates (if available). When the backup or restore session completes, Veeam Backup & Replication 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 & Replication 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.

# 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 worker VM configuration.
4. Specify worker network settings.
5. Finish working with wizard.

## Before You Begin

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

- It is recommended that workers are deployed in each Nutanix AHV cluster. If no worker is deployed in the cluster, performance of backup operations will be affected as Veeam Plug-in for Nutanix AHV will use 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.

# Step 1. Launch Add Worker Wizard

To launch the **New Nutanix AHV Worker** 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 worker**.

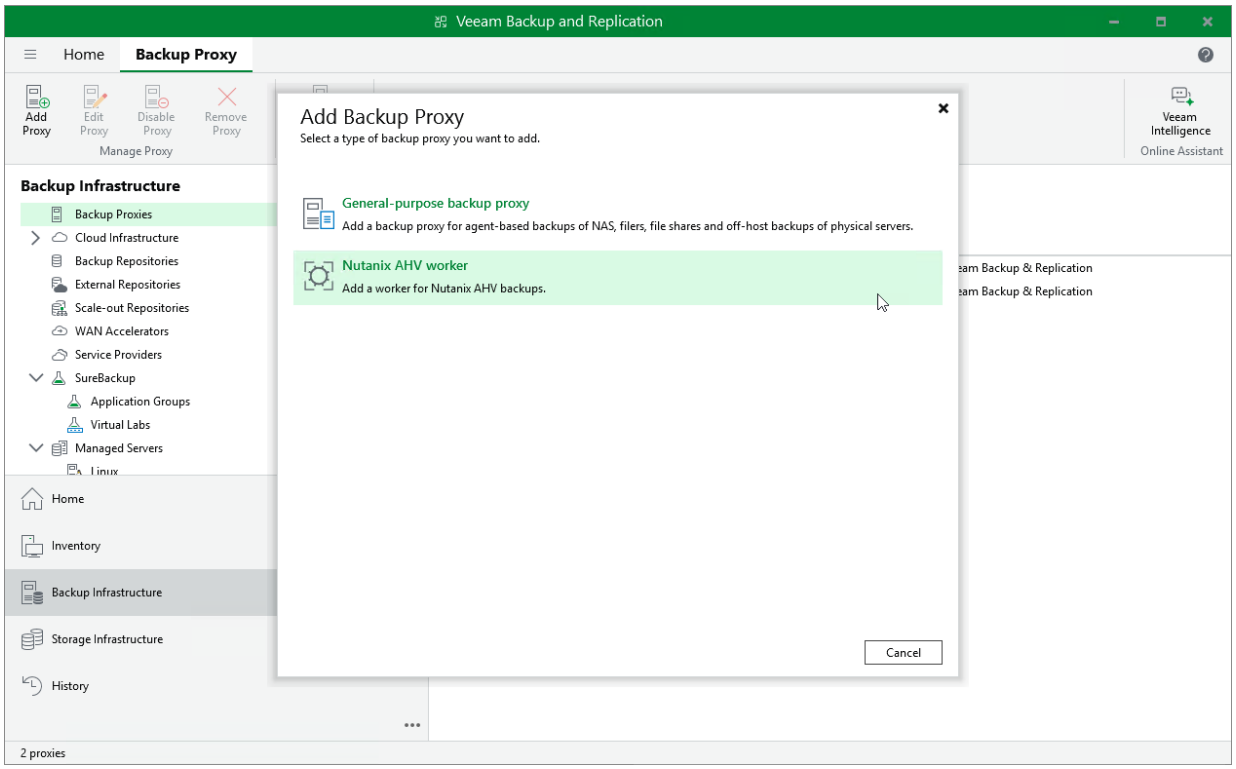

# Step 2. Specify Worker VM Settings

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

1. Click **Choose** next to the **Host** field to specify a cluster where the worker will reside.

   [Applies only to the Prism Central deployment] For a cluster to be displayed in the list of the available clusters, it must be configured in the Nutanix AHV Prism Central as described in Nutanix documentation.

   [Applies only to the standalone cluster deployment] For a cluster to be displayed in the list of the available clusters, it must added to the backup infrastructure as described in section Adding Nutanix AHV Server to Backup Infrastructure.

2. In the **Name** field, specify a name for the worker. The maximum length of the name is 63 characters; the following characters are only supported: a-z, A-Z, 0-9, -.

3. Check the **Storage** field to see the storage container that is automatically selected for worker system file.

4. In the **Description** field, provide a description for future reference. The maximum length of the description is 1024 characters.

5. In the **Max 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 worker. If you want to specify the amount of resources manually, click **Advanced proxy settings**.

> **NOTE**
>
> When performing data protection and disaster recovery operations, Veeam Backup & Replication initiates a new task for each VM that is being processed.

6. To specify a host where the worker will be launched, click **Advanced proxy settings**, select the **Host affinity** check box and choose the host.

   If you do not specify host affinity settings, Veeam Backup & Replication will automatically define the host to launch the worker.

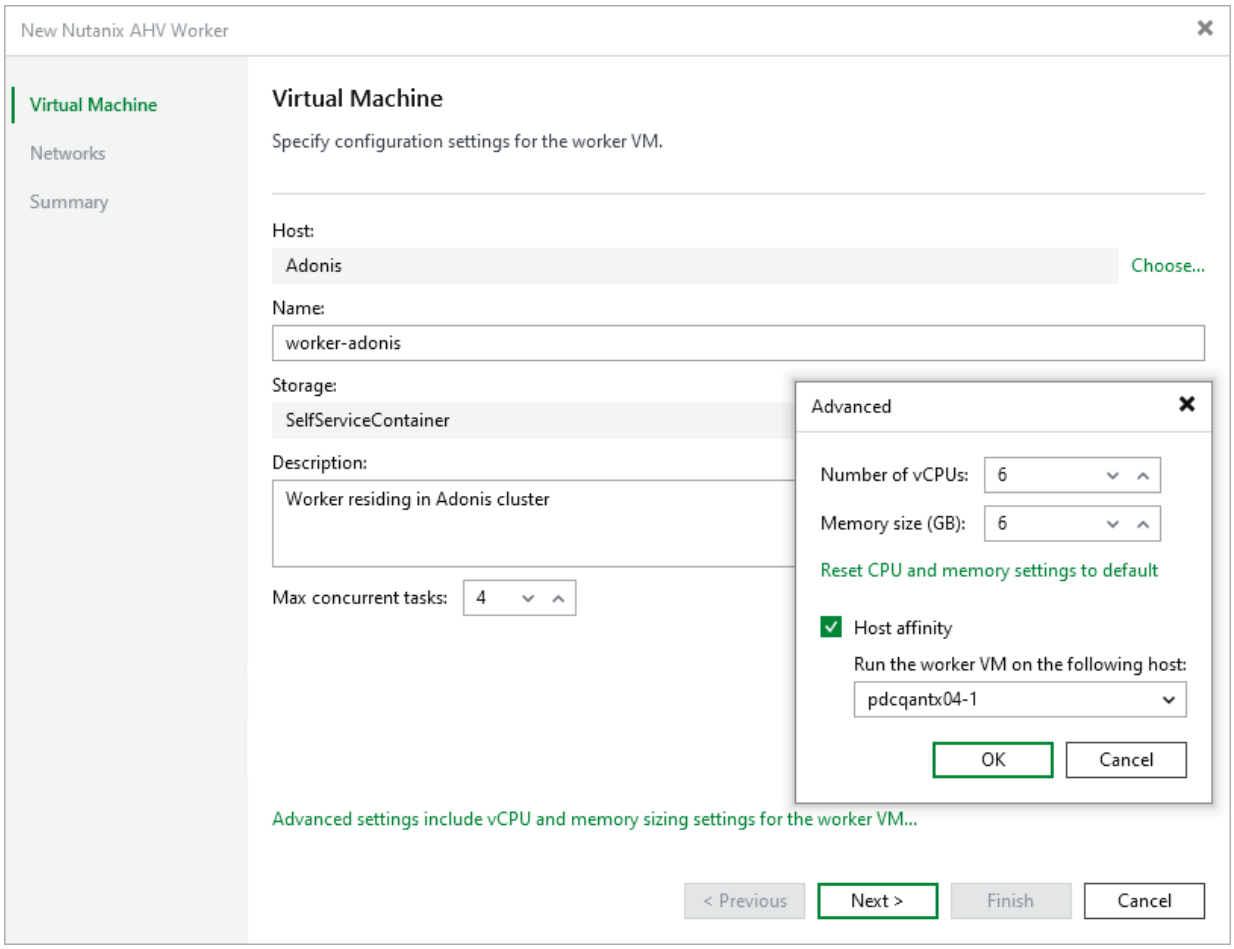

## Step 3. Configure Network Settings

At the **Networks** step of the wizard, choose a network to which the worker VM will be connected:

1. Click **Add**.
2. In the **Network Settings** window, do the following:
   a. Use the **Network** and **Description** fields to select the necessary network and to provide a description for this network connection. 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.

      To optimize worker performance, it is recommended that you select the network to which the Nutanix Controller VMs (CVMs) are connected.

   b. If DHCP is enabled in the selected network, the IP address of the worker VM will be obtained automatically. If DHCP is disabled in the selected network or if you want to specify an IP address manually, select the **Use the following IP address** option and enter the necessary IP address, subnet mask and default gateway.
3. Repeat steps 1–2 to add more network interfaces. For more information on multi-network configuration, see section Appendix. Configuring Multiple Networks.

> **NOTE**
>
> If DHCP is enabled in any of the specified networks, Veeam Backup & Replication will try to obtain DNS settings automatically. If DHCP is disabled or if you want to specify DNS settings manually, click **Obtain automatically**, select the **Use the following DNS server addresses** option in the **DNS Server Settings** window and specify the necessary IP addresses. Keep in mind that DNS settings cannot be configured separately for each network added to the worker.

## Configuring Internet Proxy for Updates

To check for available package updates for the worker, Veeam Backup & Replication automatically connects to Veeam repositories over the internet. If the worker is not connected to the internet, Veeam Backup & Replication uses backup server update settings to obtain the configuration of an internet proxy that provides access to the necessary repositories. However, you can enter specific internet proxy settings that will be used for the current worker. To do that, click **Advanced** and do the following in the **Advanced Settings** window:

1. From the **Internet proxy settings** drop-down list, select *Use custom settings*.
2. In the **Host** field, enter a DNS name or an IPv4 address of the internet proxy.
3. In the **Port** field, enter the port used on the internet proxy for HTTP or HTTPS connections.
4. If your internet proxy requires authentication, select the **Use authentication** check box, and select credentials of the account configured on the proxy to access the internet.

   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 Worker** wizard.

> **TIP**
>
> If the worker does not have access to the internet and no internet proxy is configured for the worker, you can instruct Veeam Backup & Replication not to update it. To do that, clear the **Check for updates online** check box.

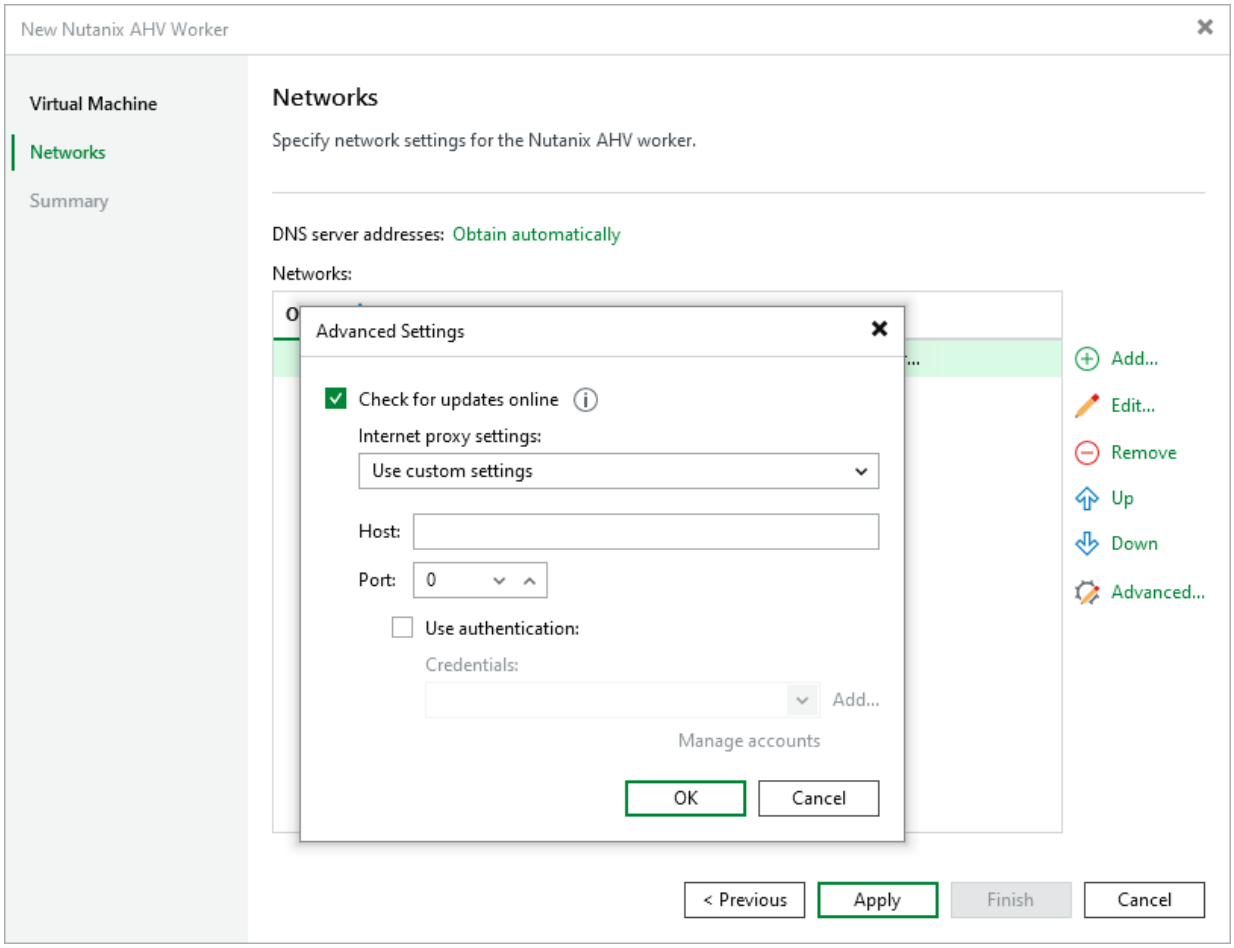

# Step 4. Finish Working with Wizard

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**.

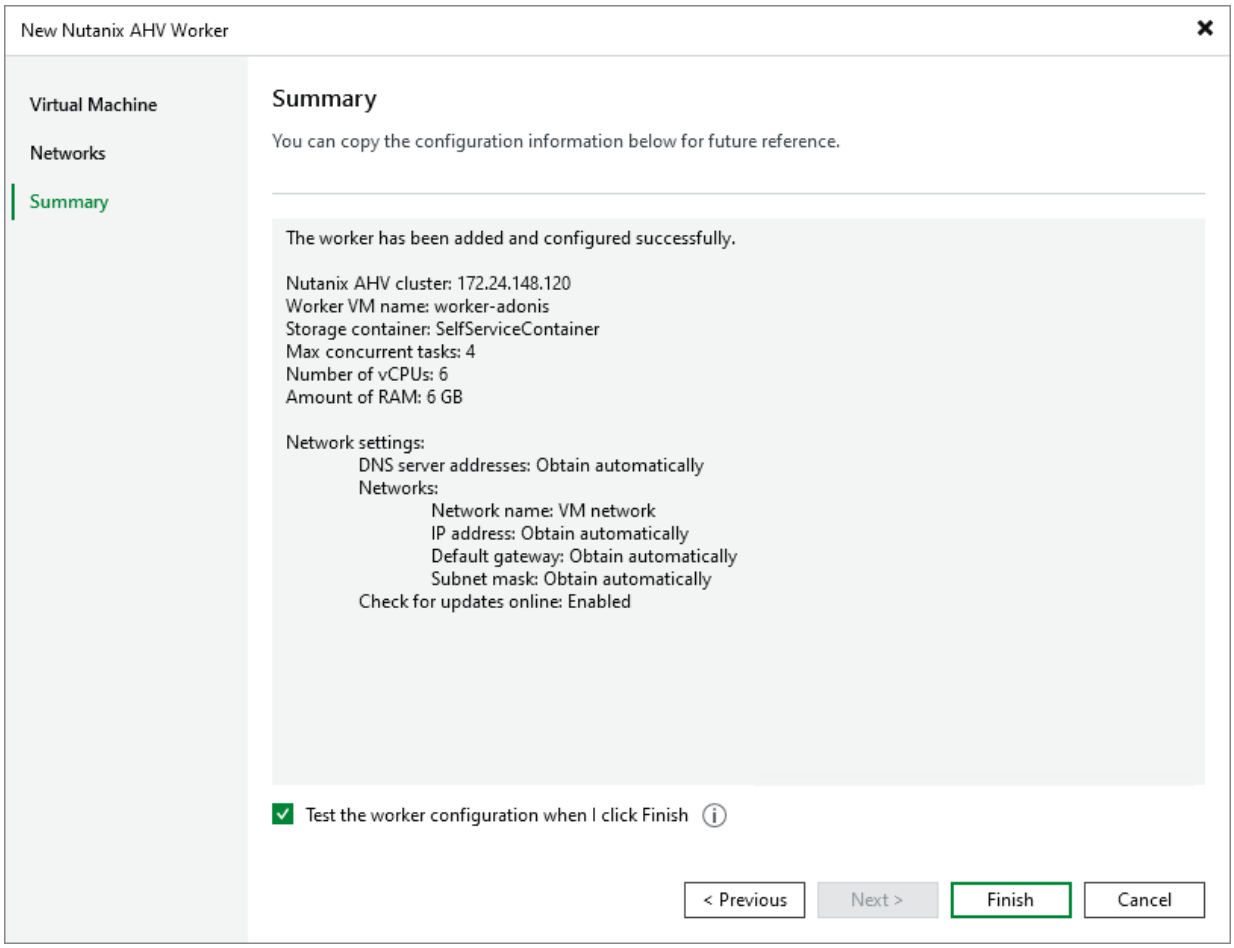

# Enabling and Disabling Workers

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. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Disable Worker** or **Enable Worker** on the ribbon.

   Alternatively, right-click the worker and select **Disable** or **Enable**.

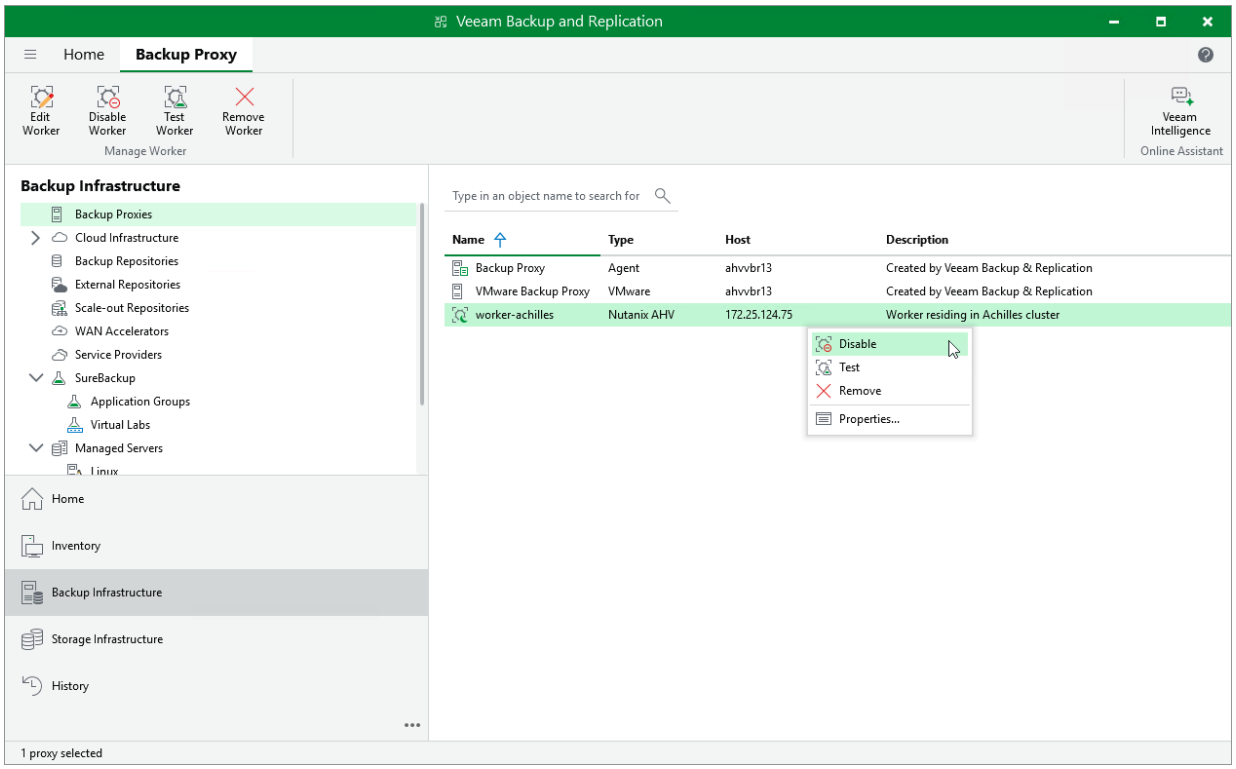

# Editing Workers

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

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Edit Worker** on the ribbon.

   Alternatively, right-click the worker and select **Properties**.
4. Complete the **Edit Nutanix AHV Worker** wizard:

   a. To provide a new name and description for the worker 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 2).

   b. 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 3).

   c. 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 & Replication will terminate the related operations, power off the worker and update the settings immediately.

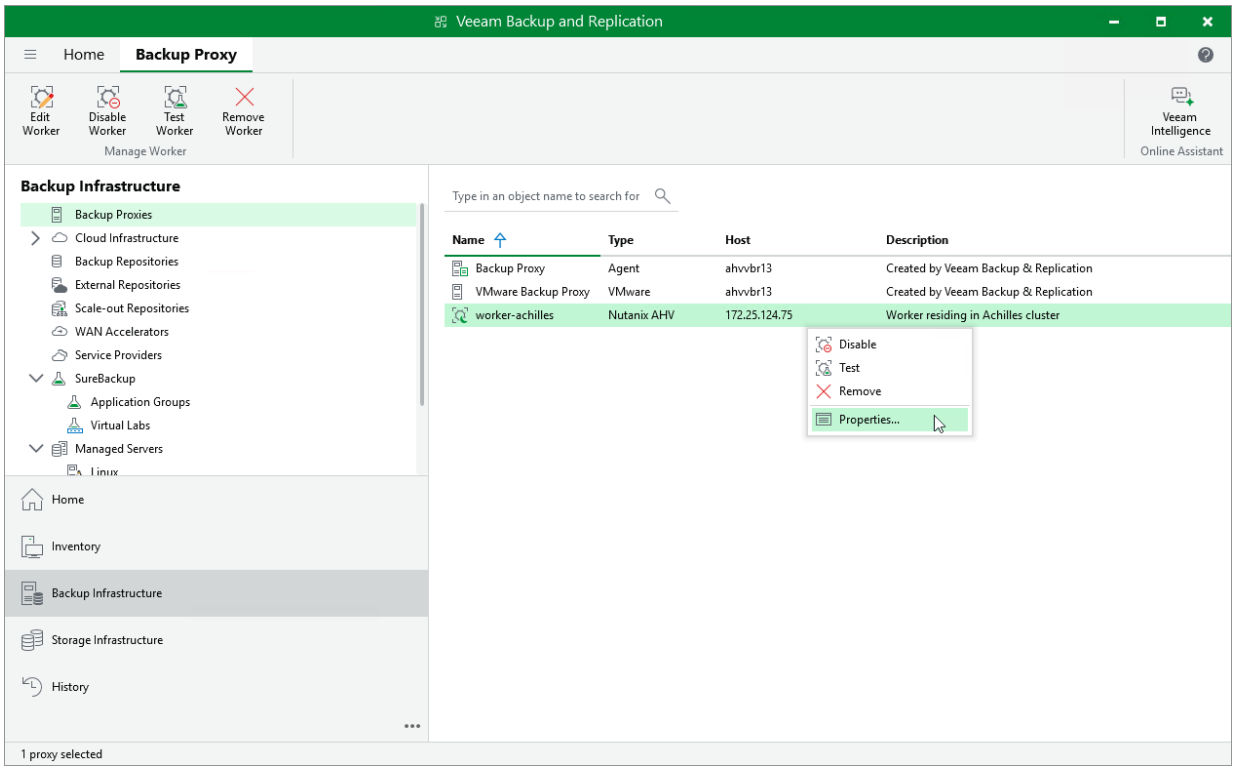

# Testing Workers

Before using a worker for a backup or restore operation, Veeam Backup & Replication 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. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Test Worker** on the ribbon.

   Alternatively, right-click the worker and select **Test**.

> **TIP**
>
> As soon as Veeam Backup & Replication finishes the worker configuration test, the worker will be powered off. You can review details of the test session in system logs as described in the Veeam Backup & Replication User Guide, section Viewing History Statistics.

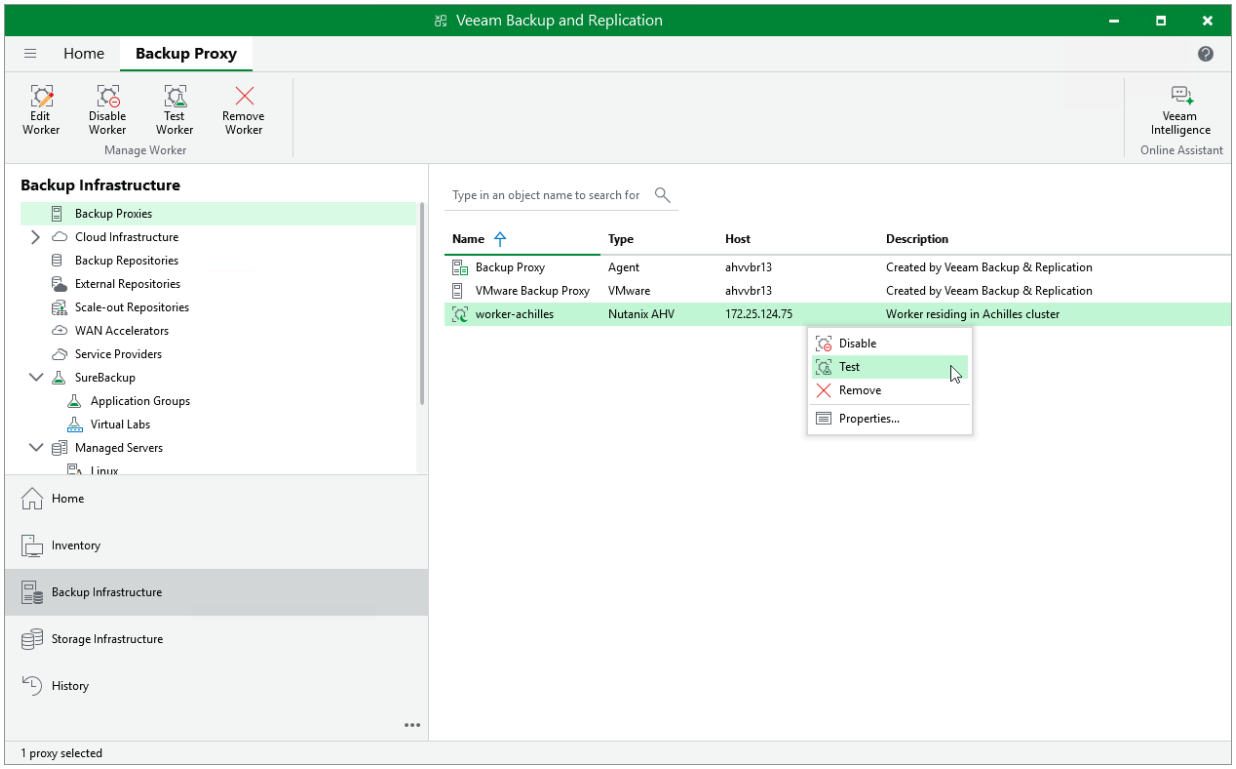

# Disabling Automatic Worker Updates

When launching a worker for a backup or restore operation, Veeam Backup & Replication 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 & Replication to use an internet proxy that will provide access to the necessary repositories.

If a worker does not have access to the internet and no internet proxy is configured for the worker, you can disable automatic updates to avoid connection failures and eliminate session warnings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Edit Worker** on the ribbon.

   Alternatively, right-click the worker and select **Properties**.
4. At the **Networks** step of the **Edit Nutanix AHV Worker** wizard, click **Advanced** and clear the **Check for updates online** check box. Then, click **Finish** to save changes made to the worker settings.

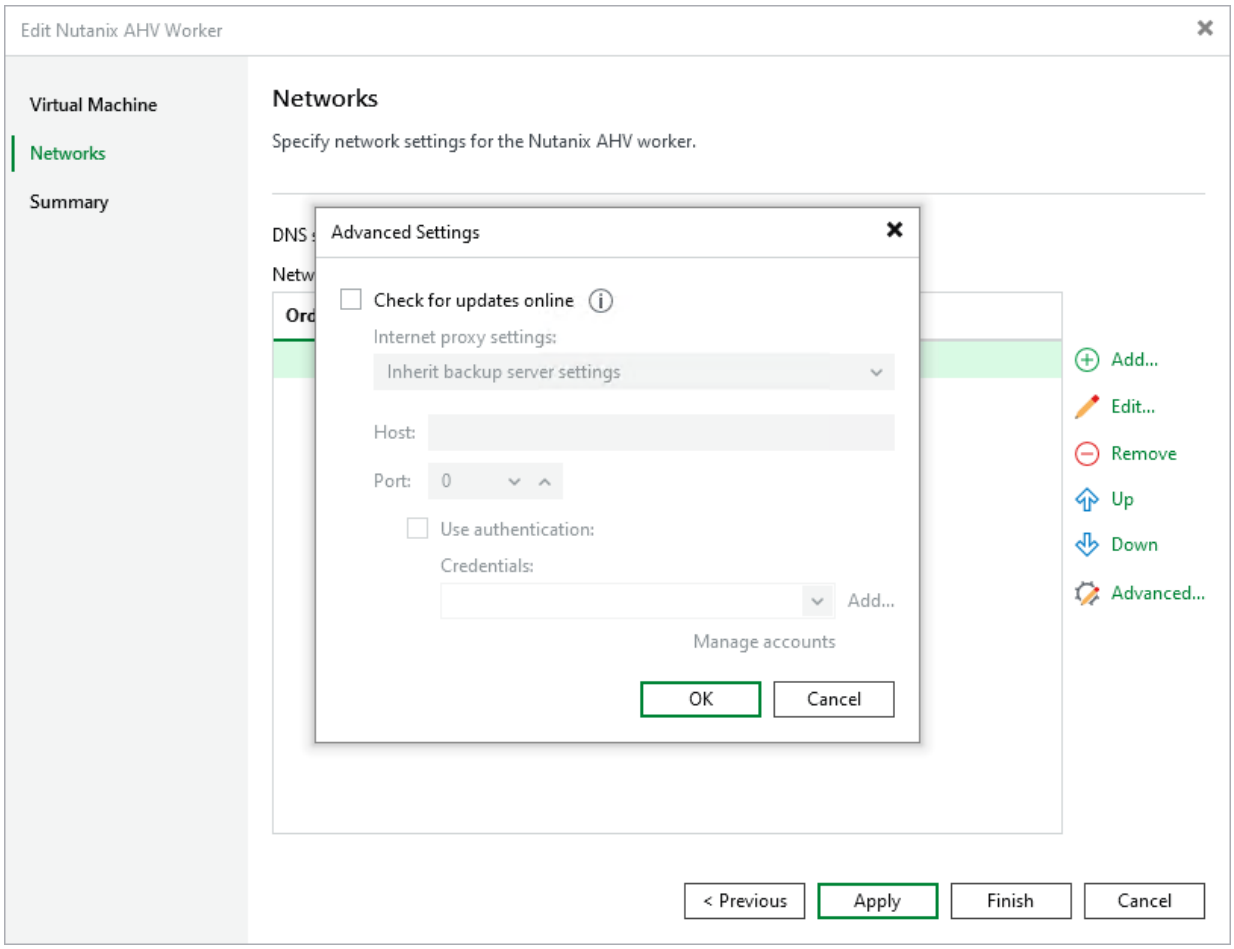

# Removing Workers

Veeam Backup & Replication 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. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the necessary worker and click **Remove Worker** on the ribbon.

   Alternatively, right-click the worker and select **Remove**.
4. In the **Veeam Backup & Replication** window, confirm that you want to permanently delete the worker.

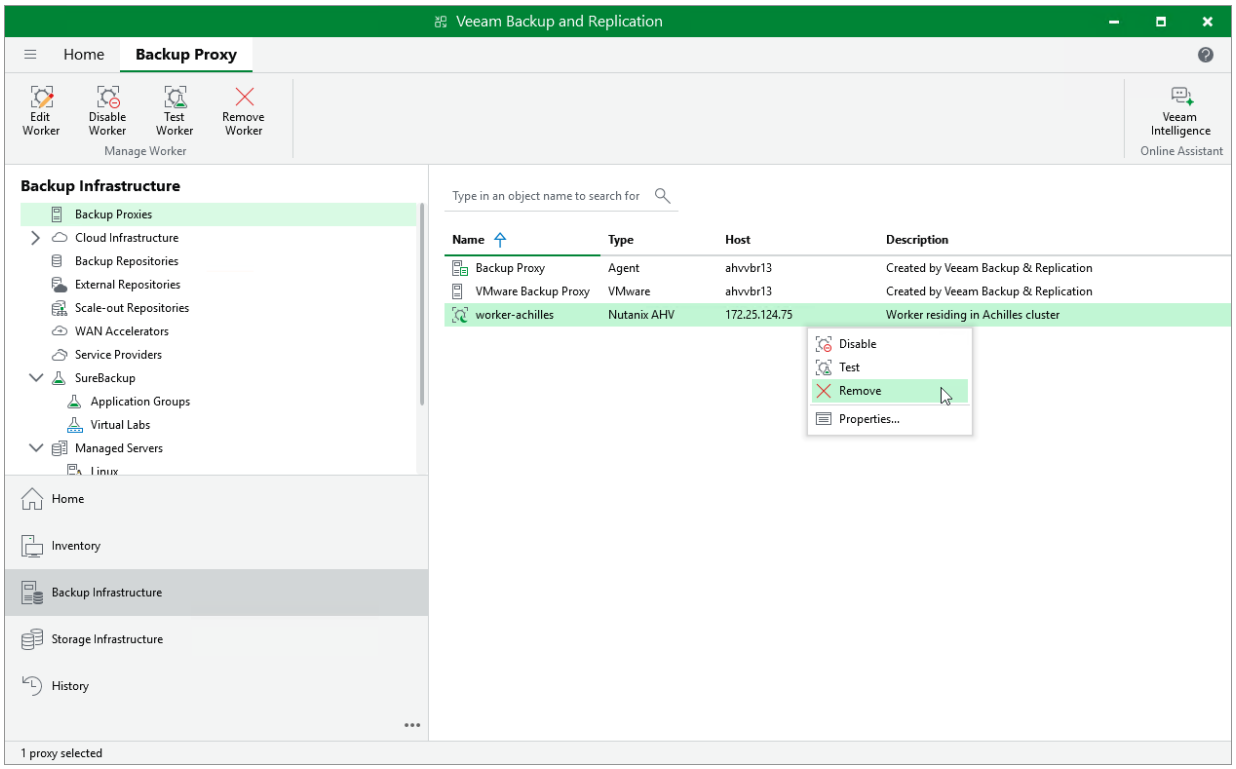

# Configuring General Settings

Veeam Backup & Replication allows you to configure general settings that are applied to all performed operations and deployed architecture components:

- Configure email settings for automated delivery of reports.
- Configure notification settings.

# Configuring Email Notification Settings

You can specify email notification settings for automated delivery of job results. To connect a mail server that will be used for sending email notifications:

1. From the main menu of the Veeam Backup & Replication console, select **Options**.
2. Switch to the **E-mail Settings** tab.
3. Select the **Enable e-mail notifications** check box.
4. 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 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 job retry** check box to receive a notification about the final job status. If you do not select this check box, the Veeam Backup & Replication will send one notification for every job retry.
11. Click **Apply**.

> **TIP**
>
> Veeam Backup & Replication allows you to send a test message to check whether you have configured the settings correctly. To do that, click **Test Message**. 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 SMTP server, Google Gmail or Microsoft 365 authentication for your mail server.

## Using SMTP Server Basic Authentication

To employ the SMTP server basic authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select *SMTP server (basic authentication)*.
2. In the **SMTP server** 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. Click **Advanced** next to the **Mail server** field and configure SMTP server settings:
   a. In the **Port** field, specify a communication port for SMTP traffic. The default SMTP port is 587 (SSL enabled) or 25 (SSL disabled).
   b. In the **Timeout** field, specify a connection timeout for responses from the SMTP server.
   c. For an SMTP server with SSL/TLS support, select the **Connect using SSL** check box to enable SSL data encryption.
   d. 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.

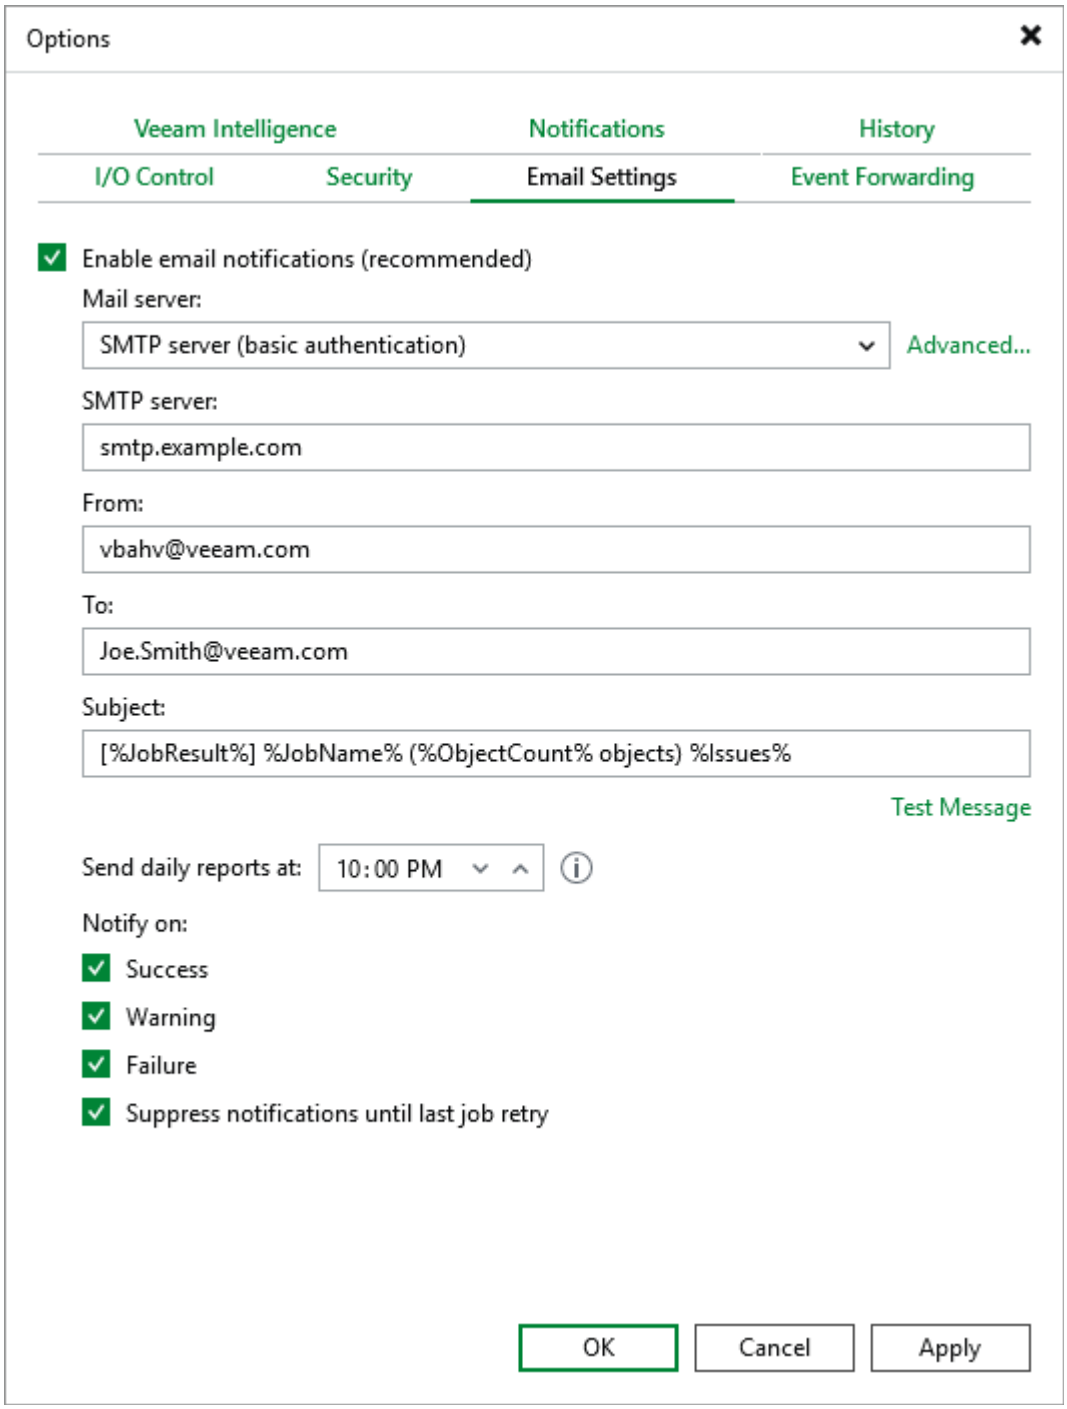

## Using Google Gmail Modern Authentication

To employ the Google Gmail modern authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select *Google Gmail (modern authentication)*.
2. Click **Sign in with Google**. You will be redirected to the authorization page.
3. On the authorization page, specify a Google account to connect to the Veeam Backup & Replication application. Note that you must also select the **Send email on your behalf** check box.

> **TIP**
>
> If you want to use your own web application for email notifications, do the following:
>
> 1. Register a new client application in the Google Cloud console for Veeam Backup & Replication to be able to use OAuth 2.0 to access Google Cloud APIs. When registering the application, it is recommended to use a dedicated service account with granular *SendMail* permissions.
> 2. In the **Options** window, click **Advanced**.
> 3. In the **Advanced** window, select the **Use custom registration settings** check box, and provide the application client ID and client secret created for the application as described in Google Cloud documentation.
> 4. Click **Sign in with Google**. You will be redirected to the authorization page.
> 5. On the authorization page, specify a Google account to connect to the registered application. Note that you must also select the **Send email on your behalf** check box.

If the authentication process completes successful, Veeam Backup & Replication will display a message notifying that the token is valid. If the token gets revoked or if the Google account password changes, click **Re-authorize** to update the configuration settings.

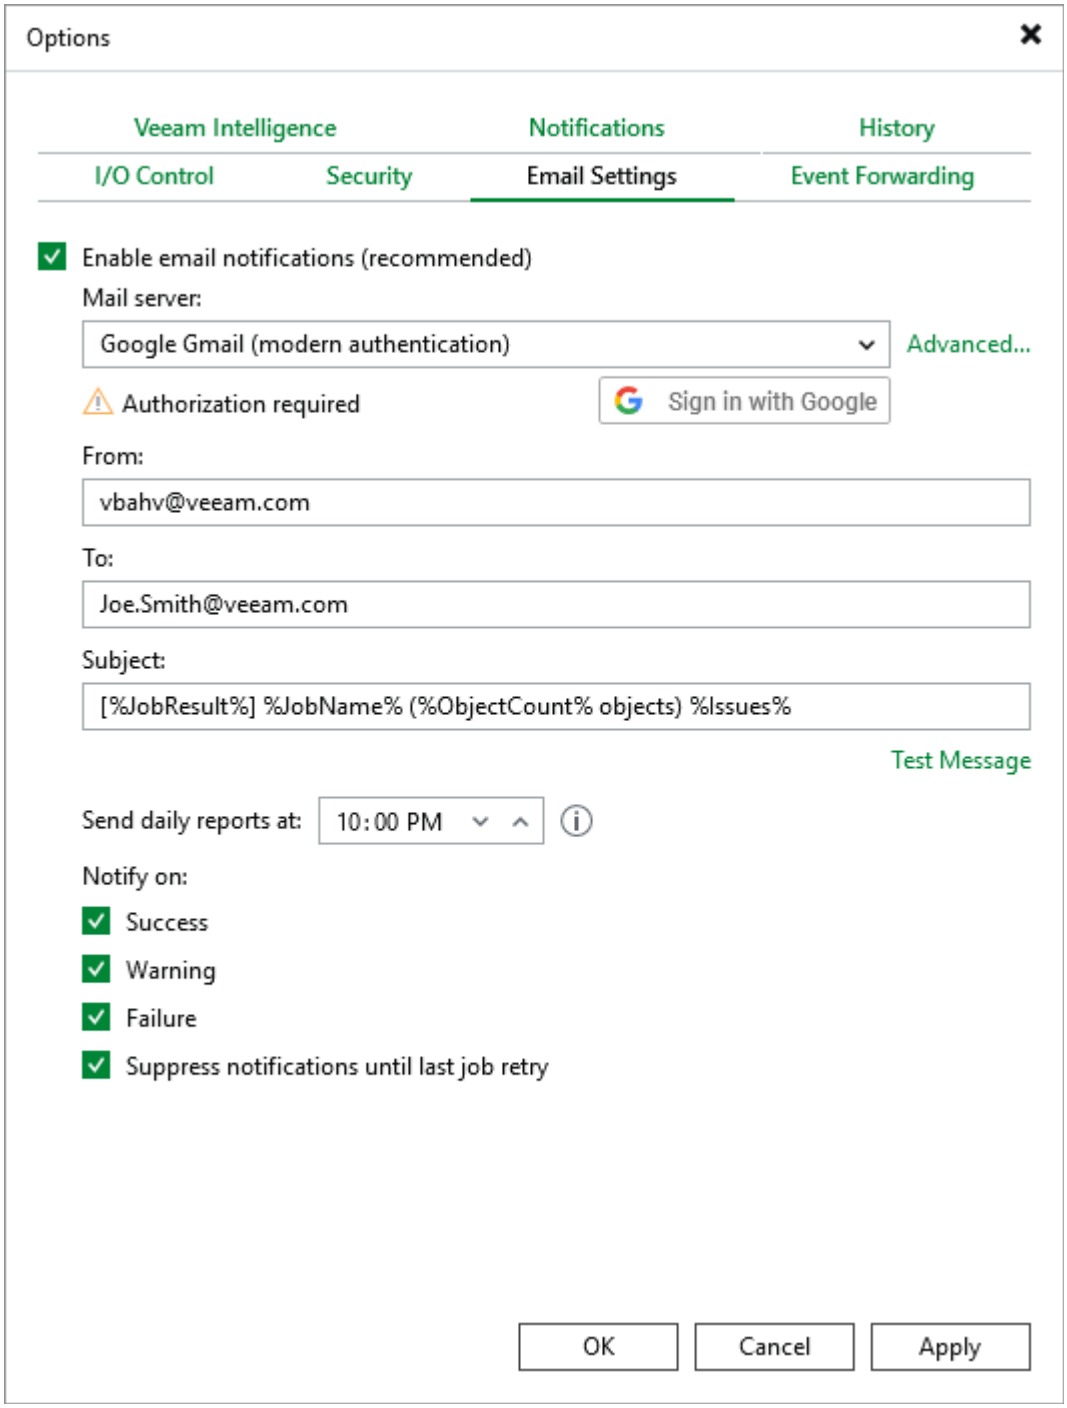


## Using Microsoft 365 Modern Authentication

To employ the Microsoft 365 modern authentication to connect to your mail server, do the following in the **Options** window:

1. From the **Mail server** drop-down list, select *Microsoft 365 (modern authentication)*.
2. Click **Authorize now**. You will be redirected to the authorization page.

3. On the authorization page, specify an Exchange Online account to connect to the Veeam Backup & Replication application. Note that you must also select the **Consent on behalf of your organization** check box.

   To sign in with Exchange Online credentials, you may need to turn off the Internet Explorer Enhanced Security Configuration option in Server Manager as described in Microsoft documentation.

**TIP**

If you want to use your own web application for email notifications, do the following:

1. Register a new client application in the Microsoft Azure portal for Veeam Backup & Replication to be able to use OAuth 2.0 to access Microsoft Azure APIs. When registering the application, it is recommended to use a dedicated service account with granular *SendMail* permissions.
2. In the **Options** window, click **Advanced**.
3. In the **Advanced** window, select the **Use custom registration settings** check box, and provide the application client ID and tenant ID created for the application as described in Microsoft documentation.
4. Click **Authorize now**. You will be redirected to the authorization page.
5. On the authorization page, specify a Exchange Online account to connect to the registered application. Note that you must also select the **Send email on your behalf** check box.

If the authentication process completes successful, Veeam Backup & Replication will display a message notifying that the token is valid. If the token gets revoked or if the Microsoft account password changes, click **Re-authorize** to update the configuration settings.

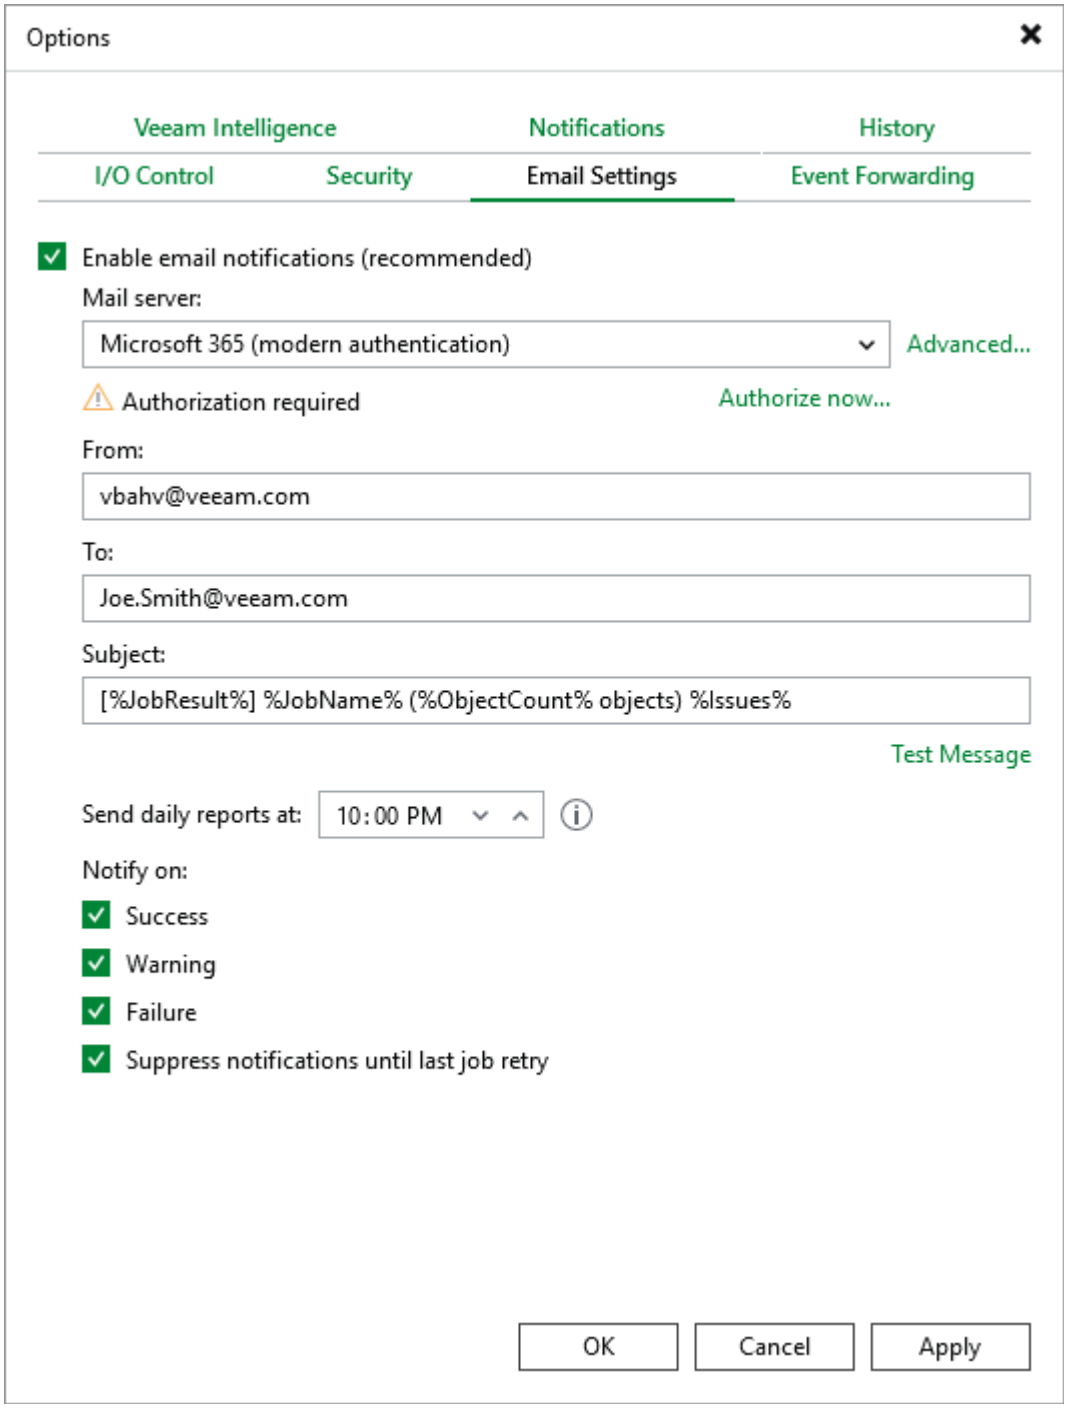

# Configuring Notifications

You can enable notifications for Veeam Backup & Replication events that may require your actions:

1. From the main menu of the Veeam Backup & Replication console, select **Options**.
2. Switch to the **Notifications** tab.
3. In the **Backup storage** section, choose whether you want to receive notifications when backup repositories used as target locations for VM backups start running out of free space. While processing VMs included into backup jobs, Veeam Backup & Replication analyzes the amount of storage space left in target repositories and displays warnings in job session details if a specific threshold is breached.
4. In the **Production datastores** section, choose whether you want to receive notifications when Nutanix AHV storage disks used as target locations for VM snapshots start running out of free space. While processing VMs included into backup jobs, Veeam Backup & Replication analyzes the amount of space left on target storage disks and displays warnings in job session details if a specific threshold is breached.

> **TIP**
>
> If Veeam Backup & Replication detects a target storage disk that is about to run out of free space while processing a VM, it will either skip the VM from processing or create a snapshot of the VM anyway, which may result in storage disruptions in the production environment. To avoid the latter, you can instruct Veeam Backup & Replication to skip VMs from processing if a specific threshold is breached.

5. In the **Support expiration** section, choose whether you want to receive notifications when the Production Support and Maintenance agreement included into your Subscription license is about to expire. When Veeam Backup & Replication detects that there are less than 14 days left before the support expiration date, it sends an email notification to the recipient specified in the general email settings.

   For more information on how to track the support expiration date, see the Veeam Backup & Replication User Guide, section Viewing License Information.

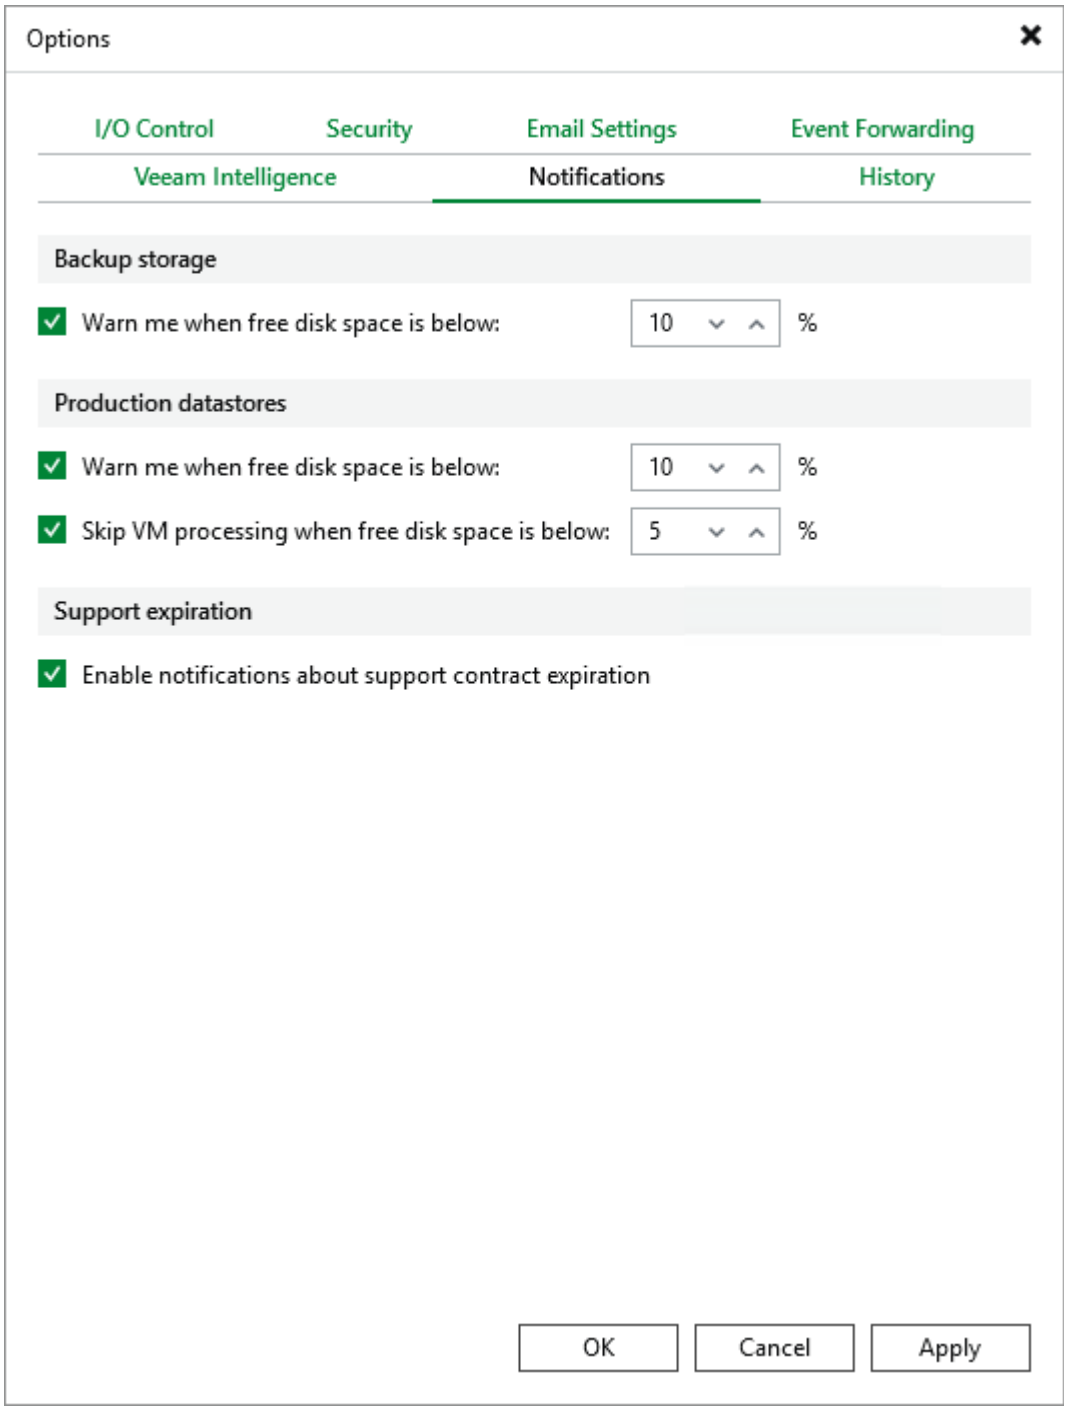

# Performing Backup

To produce VM backups, Veeam Backup & Replication 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.

One backup job can be used to process multiple VMs, but you can back up each VM with one backup job at a time. If a VM is added to more than one backup job, it will be processed only by the backup job that started earlier.

> **NOTE**
>
> To back up data that resides on Nutanix Files, use the Veeam Backup & Replication file share backup functionality described in the Veeam Backup & Replication User Guide, section Unstructured Data.

# Creating Backup Jobs

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. Select 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.

# Before You Begin

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

- You cannot back up workers and Nutanix Controller VMs and Prism Central VMs.
- You cannot use Veeam Plug-in for Nutanix AHV to back up iSCSI disks mounted to VMs — they are skipped from processing automatically. However, you can use Veeam Agent for Microsoft Windows or Veeam Agent for Linux to protect those disks.
- If the backup job includes individual virtual machines, as well as the whole Prism Central, category, protection domain or cluster, Veeam Backup & Replication 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.
- [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.
- Since Veeam Backup & Replication does not allow you to assign information about locations to Nutanix AHV clusters and workers, job statistics do not include information on the Nutanix AHV VM data migration between different geographic regions.

# 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**.

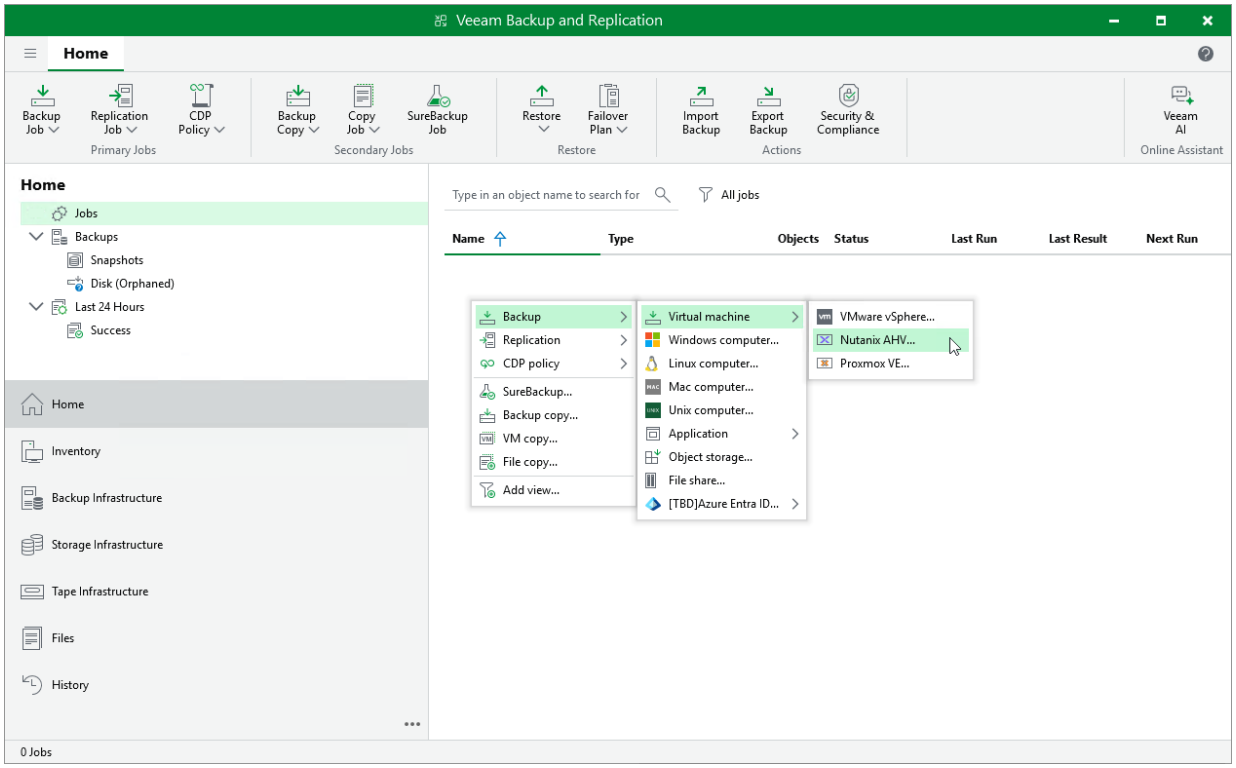

# Step 2. Specify Job Name and Description

At the **Name** step of the wizard, 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 & Replication.

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

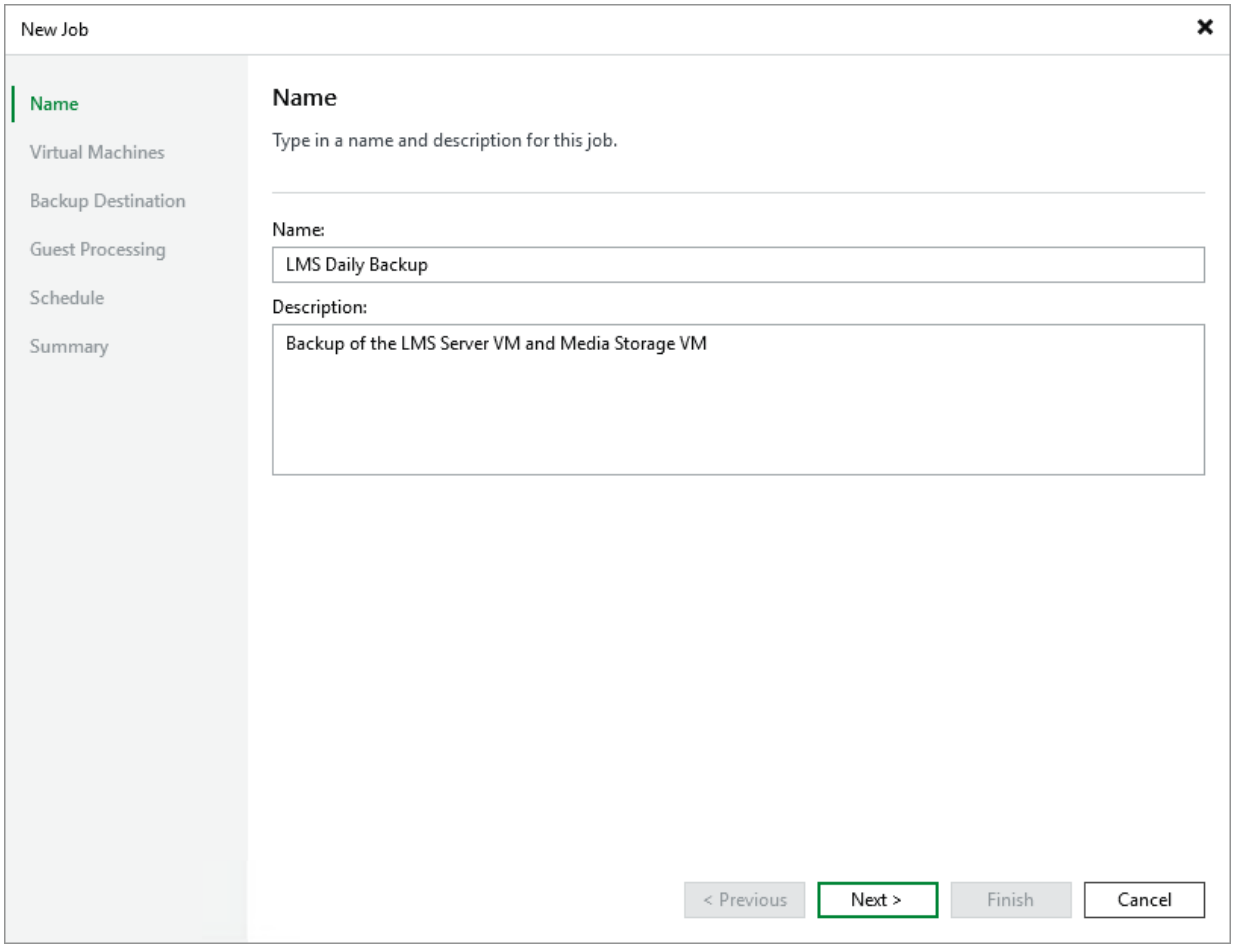

# Step 3. Configure Backup Source Settings

At the **Virtual Machines** step of the wizard, specify the backup scope — resources that Veeam Backup & Replication will back up.

## Step 3a. Choose Virtual Machines

Specify VMs that will be included into the backup scope:

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 & Replication 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 & Replication 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, protection domains, cluster or categories 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 & Replication will still not process the resource because the list of excluded resources has a higher priority.

While running the job, Veeam Backup & Replication 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 & Replication 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 & Replication 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 & Replication fails to obtain data from a replica cluster, backup will be still performed using VM data obtained from the main cluster.

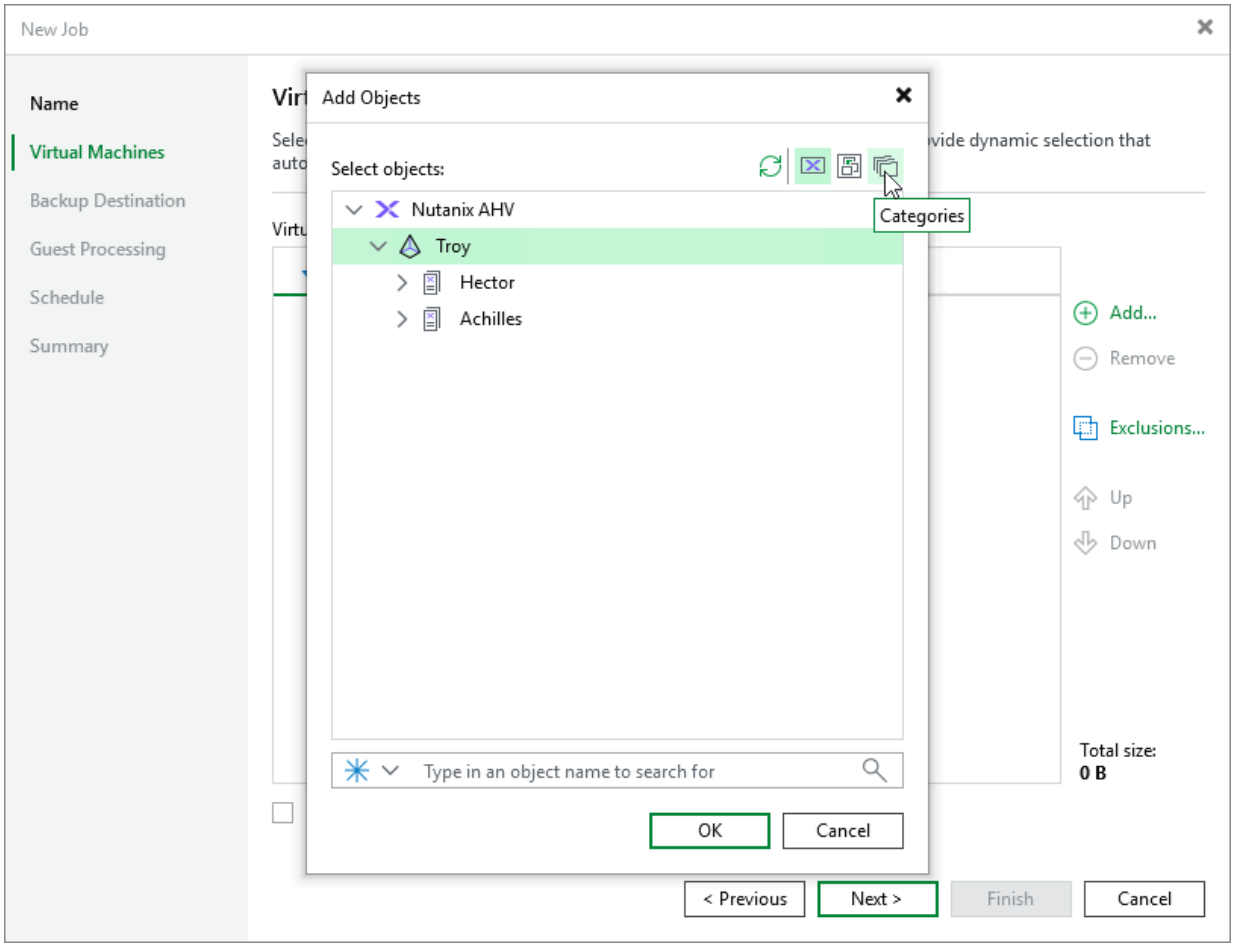


# Step 3b. Choose Disks and Volume Groups

By default, jobs process all disks and volume groups attached to VMs included into the backup scope. However, you can instruct Veeam Backup & Replication 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 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 & Replication will process the VM disks according to the rule configured for the VM.

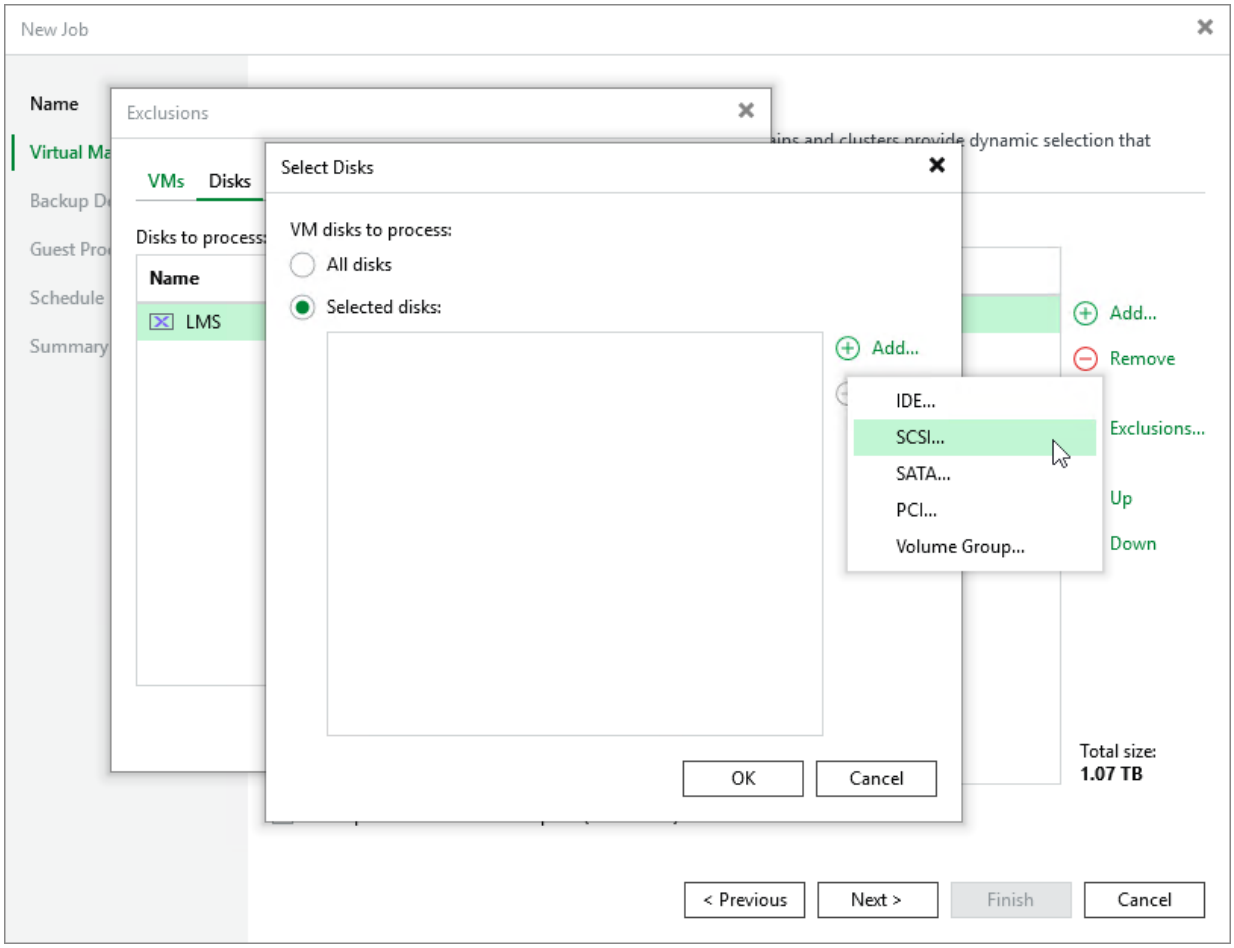


# Related Topics

Snapshot Types

# Step 4. Configure Backup Destination 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.

2. In the **Retention policy** section, choose how long Veeam Backup & Replication will keep restore points in a backup chain. If a restore point is older than the specified limit, Veeam Backup & Replication will remove it from the chain. For more information on how Veeam Backup & Replication tracks and removes redundant restore points, see Retention Policies.

   Keep in mind that since every backup chain must contain at least 3 restore points, Veeam Backup & Replication may ignore the configured retention policy settings and retain restore points for longer periods of time. For more information, see Backup Retention.

> **NOTE**
>
> If the UUID of a VM changes (for example, if the VM migrates to another cluster), Veeam Backup & Replication will be unable to continue the backup chain for this VM. After you re-add the VM to the backup job, Veeam Backup & Replication 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 & Replication 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 customize email notifications).

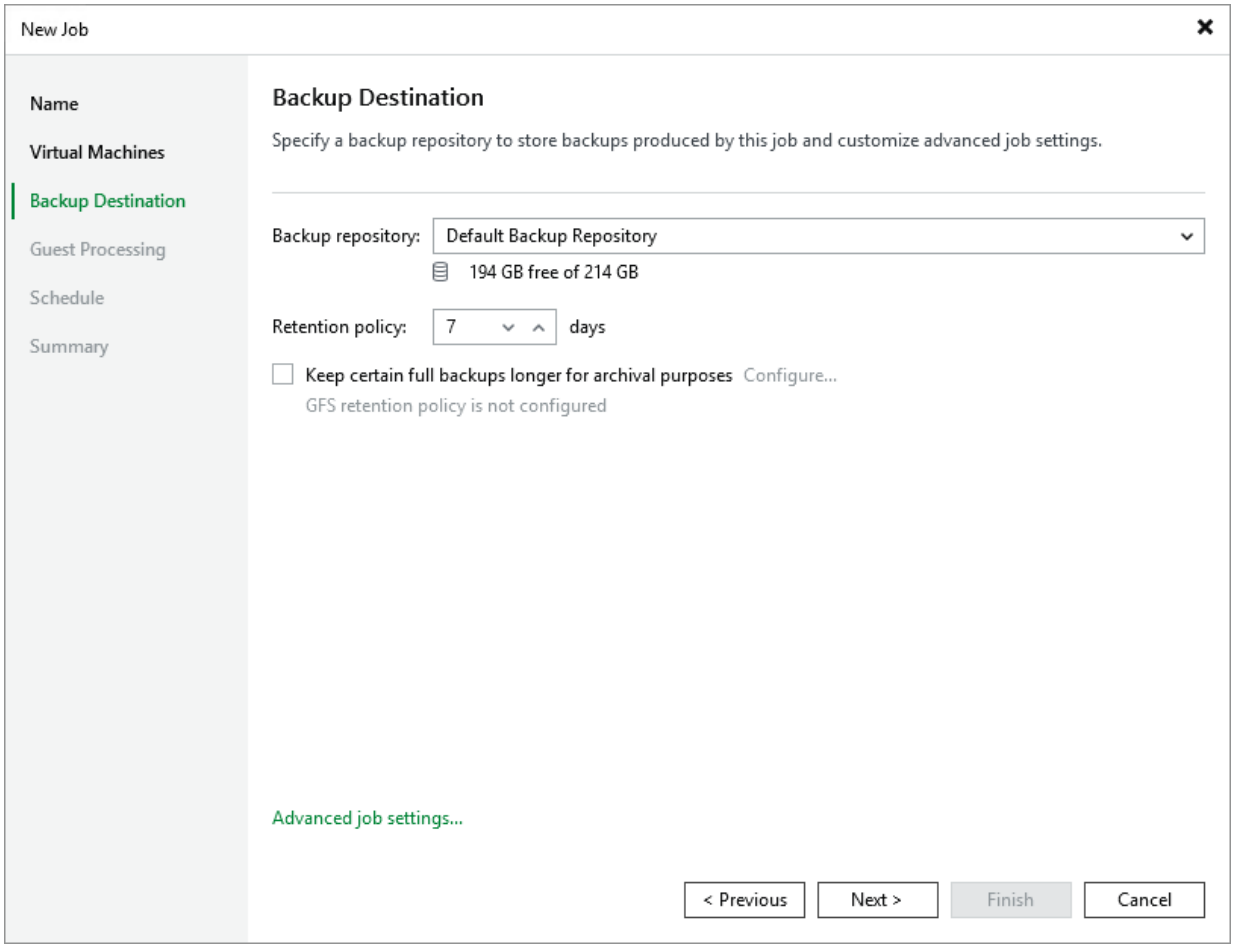


# Configuring GFS Policy Schedules

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 & Replication 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 & Replication 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 & Replication 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 & Replication will keep a full backup created during a specific week for a specified number of months.
- **Keep yearly full backups** — Veeam Backup & Replication 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 & Replication 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.

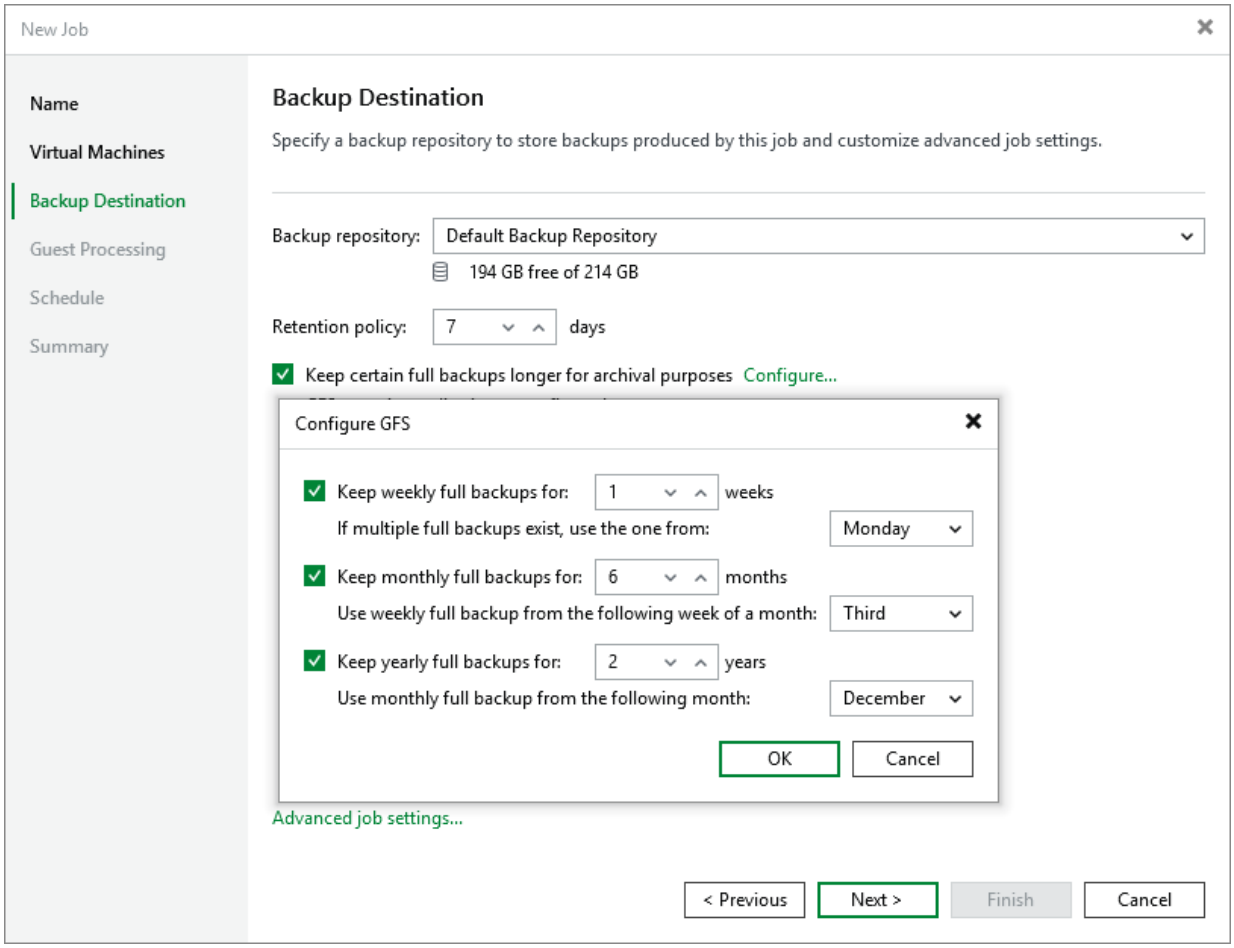


# Configuring Advanced Settings

In the **Advanced settings** window, you can schedule full backups, configure backup job maintenance settings, enable Nutanix Guest Tools quiescence, specify backup file storage settings and customize email notifications.

# Backup Settings

To instruct Veeam Backup & Replication to create full backups according to a specific schedule, switch to the **Backup** tab and do the following:

1. To schedule synthetic full backups, select the **Create synthetic full backups periodically** check box, click **Configure** and choose whether you want to create these backups on specific days on a weekly or monthly basis.
2. To schedule active full backups, select the **Create active full backups periodically** check box, click **Configure** and choose whether you want to create these backups on specific days on a weekly or monthly basis.

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

**IMPORTANT**

- Synthetic full backups cannot be scheduled if an object storage repository is selected as the target location for backups.
- Do not schedule synthetic and active full backups to run at the same time. Due to technical limitations, Veeam Backup & Replication will be unable to create synthetic full backups according to the specified schedule.

# Maintenance Settings

To specify how Veeam Backup & Replication will maintain backups created by the backup job, switch to the **Maintenance** tab and do the following:

1. To instruct Veeam Backup & Replication to periodically perform a health check for the backups, 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.
2. To configure retention settings for backups of VMs that are no longer processed by the backup job, select the **Remove deleted items data after** check box, and specify the number of days during which Veeam Backup & Replication will keep these backups.

**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 you configure a helper appliance in the repository settings. Otherwise, additional data transfer costs may occur.

# Storage Settings

To specify storage settings for backup files created by the backup job, switch to the **Storage** tab and do the following:

1. To decrease the size of the files, select a compression level from the **Compression level** drop-down list (*None*, *Dedupe-friendly*, *Optimal*, *High* or *Extreme*). For more information on data compression, see the Veeam Backup & Replication User Guide, section Compression and Deduplication.

2. To optimize job performance and storage usage, select a block size from the **Storage optimization** drop-down list. Veeam Backup & Replication will use this size to "split" VM images into separate data blocks when processing VMs — the more data blocks there are, the more time is required to process the VM images. For more information on how data block sizes affect performance, see the Veeam Backup & Replication User Guide, section Storage Optimization.

# Guest Quiescence Settings

To instruct Nutanix AHV to freeze applications running on VMs while snapshots are taken, switch to the **Nutanix AHV** tab, select the **Enable Nutanix Guest Tools quiescence** check box and choose whether you want to truncate transaction logs. Keep in mind that if you select the **Never truncate transaction logs (VSS_BT_COPY)** option, it may significantly increase the amount of storage space consumed by VMs running as Microsoft Exchange Mail Servers.

> **NOTE**
>
> Veeam Backup & Replication prioritizes guest processing settings over guest quiescence settings. If you enable application-aware processing for a backup job, Veeam Backup & Replication will ignore the configured guest quiescence settings.

# Notification Settings

To instruct Veeam Backup & Replication to send email notifications on the backup job results, switch to the **Notifications** tab, select the **Send email notifications** check box and specify an email address of a recipient; use a semicolon to separate multiple recipient addresses. For Veeam Backup & Replication to be able to send email notifications, you must configure a mail server as described in section Configuring Email Notification Settings.

> **NOTE**
>
> Email notifications on the backup job results will be also sent to recipients configured in the global notification settings.

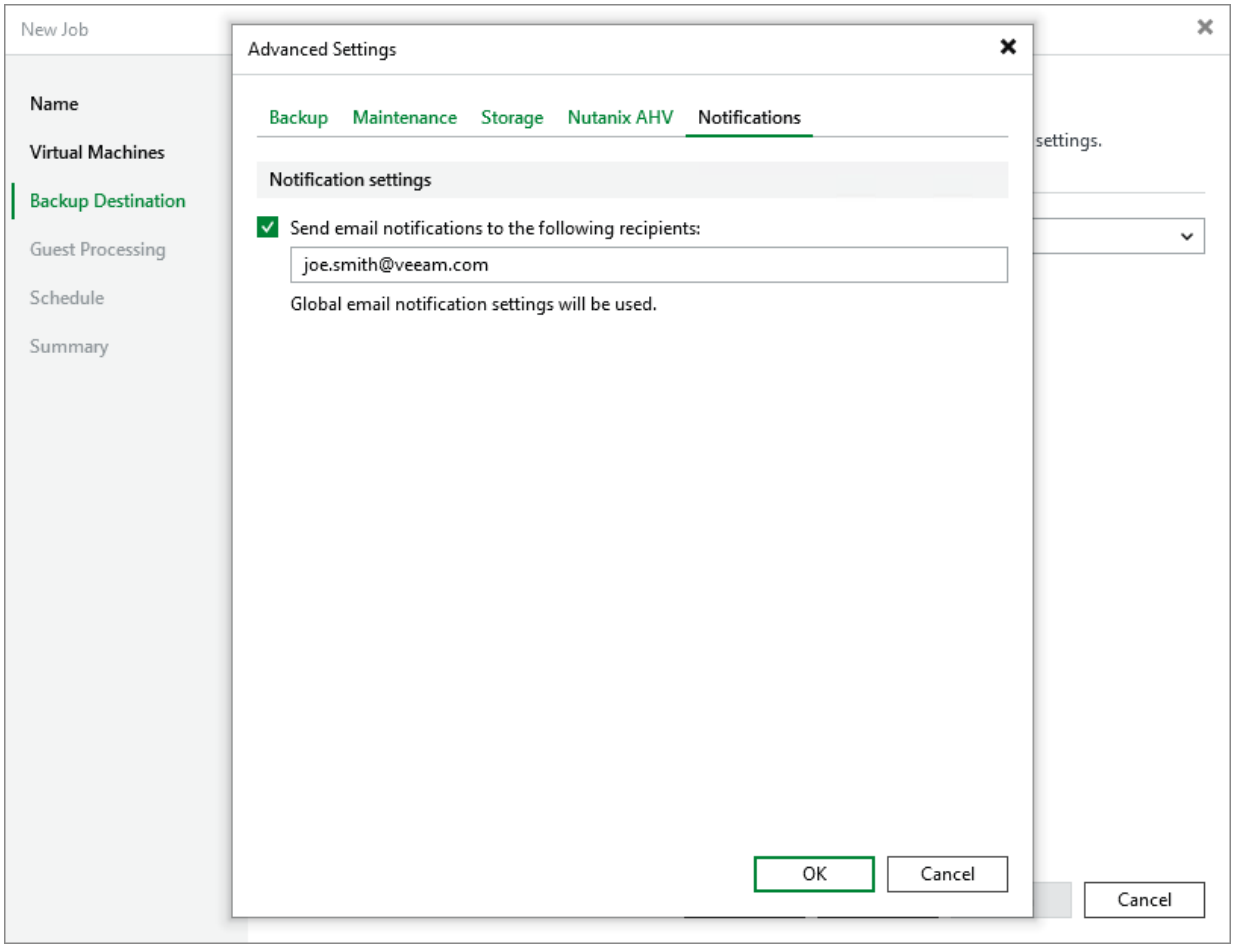


## How Health Check Works

When Veeam Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV starts a new health check session. For each restore point in the backup chain, Veeam Plug-in 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 Plug-in for Nutanix AHV also checks whether data blocks that are required to rebuild the restore point are available.

If Veeam Plug-in 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 Plug-in for Nutanix AHV performs the following operations:

- If the health check detects corrupted metadata in a full restore point, Veeam Plug-in for Nutanix AHV marks the backup chain as corrupted in the configuration database. During the next backup job session, Veeam Plug-in 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 Plug-in 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 Plug-in 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.

# Step 5. Specify 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 servers that Veeam Backup & Replication will use when communicating with guest OSes of 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 Plug-in for Nutanix AHV will not be able to obtain VM data from replica clusters.

- Veeam Plug-in 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 Plug-in for Nutanix AHV will not be able to use Kerberos authentication while connecting to guest OSes of the processed VMs.

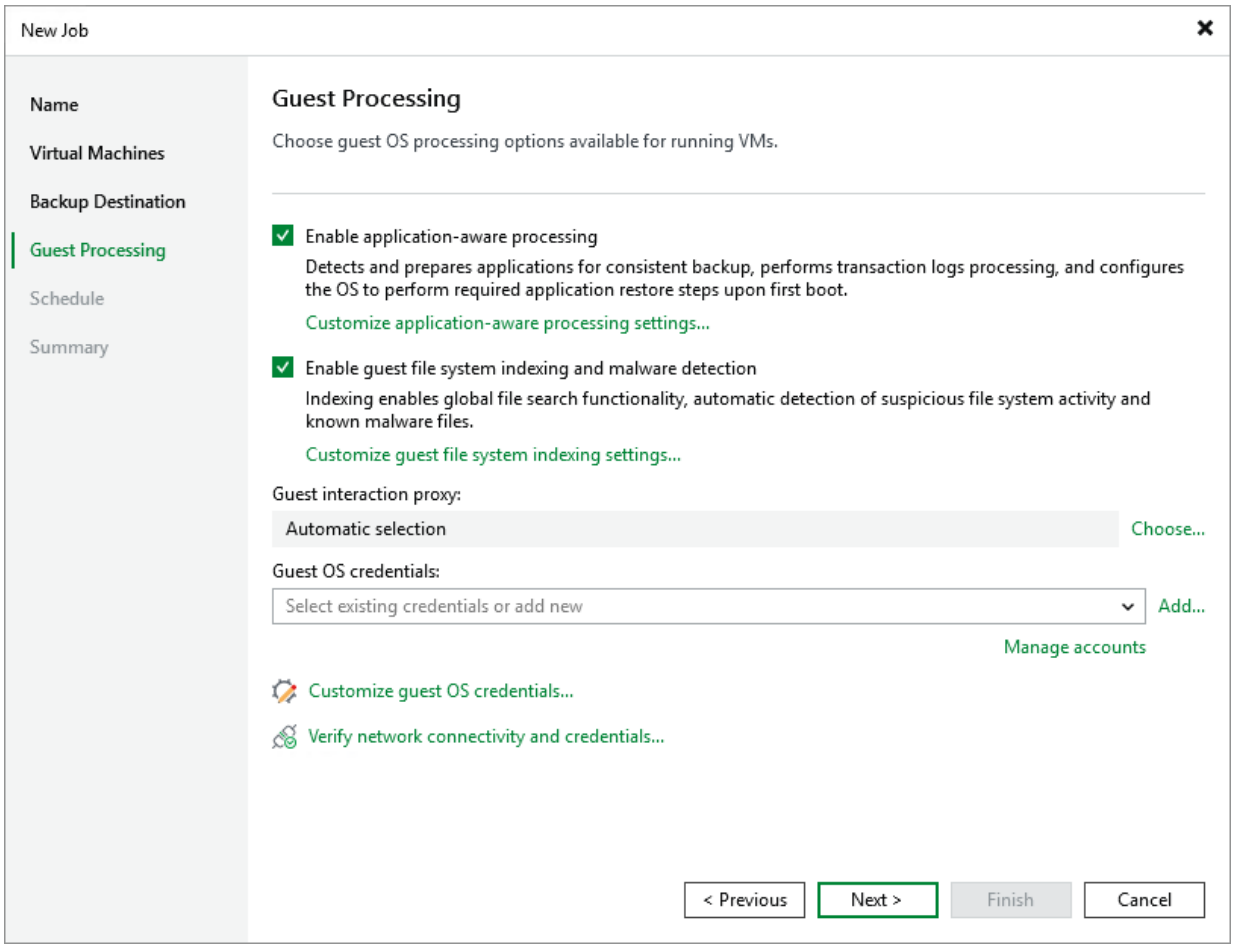


# Related Topics

- Requirements and limitations for PostgreSQL WAL files backup
- Requirements and limitations for Oracle archived redo logs backup

# 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.

For more information on supported applications that can be protected with application-consistent backups, see the Veeam Backup & Replication User Guide, section Supported Platforms and Applications.

# 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, complete the following steps:

1. Click **Customize application-aware processing settings**.
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. In the **Processing Settings** window, do the following:
   - 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, switch to the **SQL** tab and follow the instructions provided in section Specifying 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.
   - To specify how Veeam Backup & Replication will process transaction logs of Oracle Server applications, switch to the **Oracle** tab and follow the instructions provided in section Specifying Oracle Archived Redo Log Settings.

- To specify how Veeam Backup & Replication will process transaction logs of PostgreSQL Server applications, switch to the **PostgreSQL** tab and follow the instructions provided in section Specifying PostgreSQL WAL Files Settings.
- 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 Specifying 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**.

# Enabling Persistent Agent Components

For Veeam Backup & Replication to be able to create transactionally-consistent backups of a processed VM, this VM must run a number of runtime components that enable access the VM guest OS. By default, these components are installed temporarily and then removed automatically as soon as the backup operation completes; this approach eliminates error-prone manual steps but requires multiples ports to be open on the VM. To change this behavior, you can instruct Veeam Backup & Replication to use persistent agent components that are installed permanently and then run in the background while no backup operations are performed; this approach increases the security of guest processing, requires a very limited number of ports to be open on the processed VM — but you will have to take some additional configuration steps to install the agent deployment kit on the VM.

> **IMPORTANT**
>
> If you want to enable persistent agent components, consider the following:
>
> - Persistent agents can be used only for application-aware processing of Windows-based VMs.
> - To use persistent agents, Veeam Backup & Replication requires a Windows Deployment Kit deployed on the protected VM beforehand. If Veeam Backup & Replication fails to connect to the VM using persistent agent components, no backup will be created for this VM. For security reasons, Veeam Backup & Replication will not try to install non-persistent runtime components for application-aware processing.

For more information on components required for application-aware processing , see the Veeam Backup & Replication User Guide, section Non-Persistent Runtime Components and Persistent Agent Components.

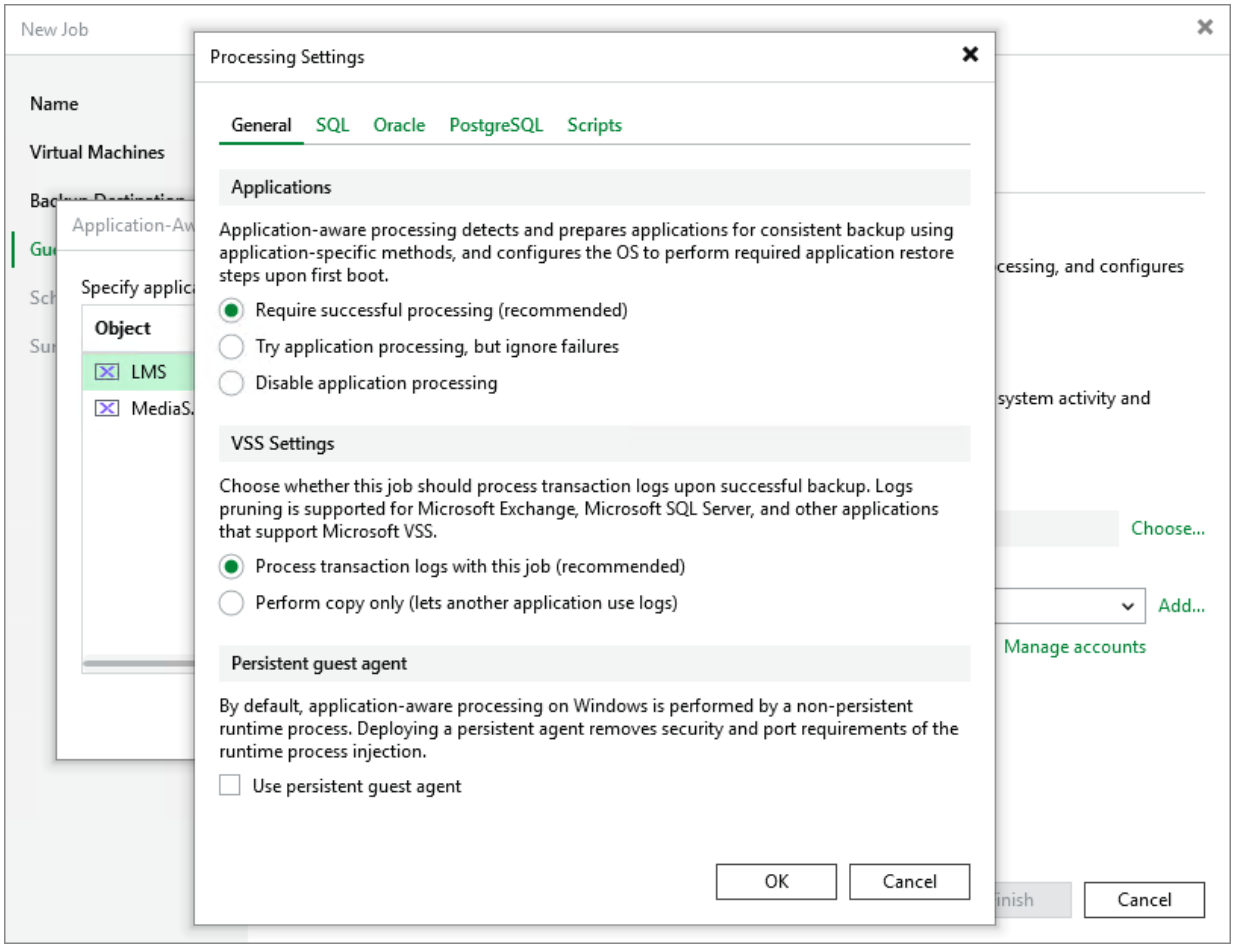


## Specifying Microsoft SQL Server Transaction Log Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the Microsoft SQL Server application and truncates transaction logs after each successfully completed backup session — this will allows you to restore Microsoft SQL Server databases using specific backups. To protect mission-critical Microsoft SQL Server databases, you can instruct Veeam Backup & Replication to create secondary restore points with transaction logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

**NOTES**

- Veeam Backup & Replication stores image-level backups and transaction log backups in the same repository.
- If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary transaction log backups will be created.

To back up Microsoft SQL Server transaction logs periodically, do the following:

1. Switch to the **SQL** tab and select the **Backup logs periodically** option.
2. In the **Backup logs every** field, specify how frequently you want transaction logs to be backed up. The maximum field value is 480 minutes.

3. In the **Retain log backups section**, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove transaction log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4.
   - Select the **Keep only last <N> days of log backups** option if you want to retain transaction logs for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that image-level backups must always be kept for a longer period than the related transaction log backups.

   For more information on how Veeam Backup & Replication retains transaction logs, see the Veeam Backup & Replication User Guide, section Microsoft SQL Server Log Backup.
4. In the **Log shipping servers** section, choose whether you want to use a specific Windows server to transfer transaction log backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary Windows servers.

   For a Windows server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section Adding Microsoft Windows Servers. Keep in mind that the list will also include Linux servers added to the backup infrastructure; however, Linux servers cannot be used as log shipping servers due to technical limitations in the current version.

**TIPS**

- It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
- It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators or backup servers).

You can also choose not to truncate logs at all. However, keep in mind that this option requires databases to use the simple recovery model. Otherwise, transaction logs may grow large and increase the storage space consumption significantly. For more information on recovery models used by Microsoft SQL databases, see Microsoft Docs.

# Considerations and Limitations

When you configure transaction log settings, consider the following:

- If a processed VM runs Microsoft SQL Server along with Oracle Server and transaction log backup is enabled for both applications, the Microsoft SQL Server transaction logs will not be backed up — Veeam Backup & Replication will create transaction log backups for the Oracle Server application only.
- If a processed VM runs Microsoft SQL Server that hosts the Veeam Backup & Replication configuration database, its transaction logs will not be backed up:
  - If the Microsoft SQL Server has the SQL Server Always On availability groups feature disabled, the configuration database will be excluded from application-aware processing automatically.

- If the Microsoft SQL Server has the SQL Server Always On availability groups feature enabled, you will have to exclude the configuration database from application-aware processing manually, as described in this Veeam KB article.

For more information on the SQL Server Always On availability group feature, see Microsoft Docs.

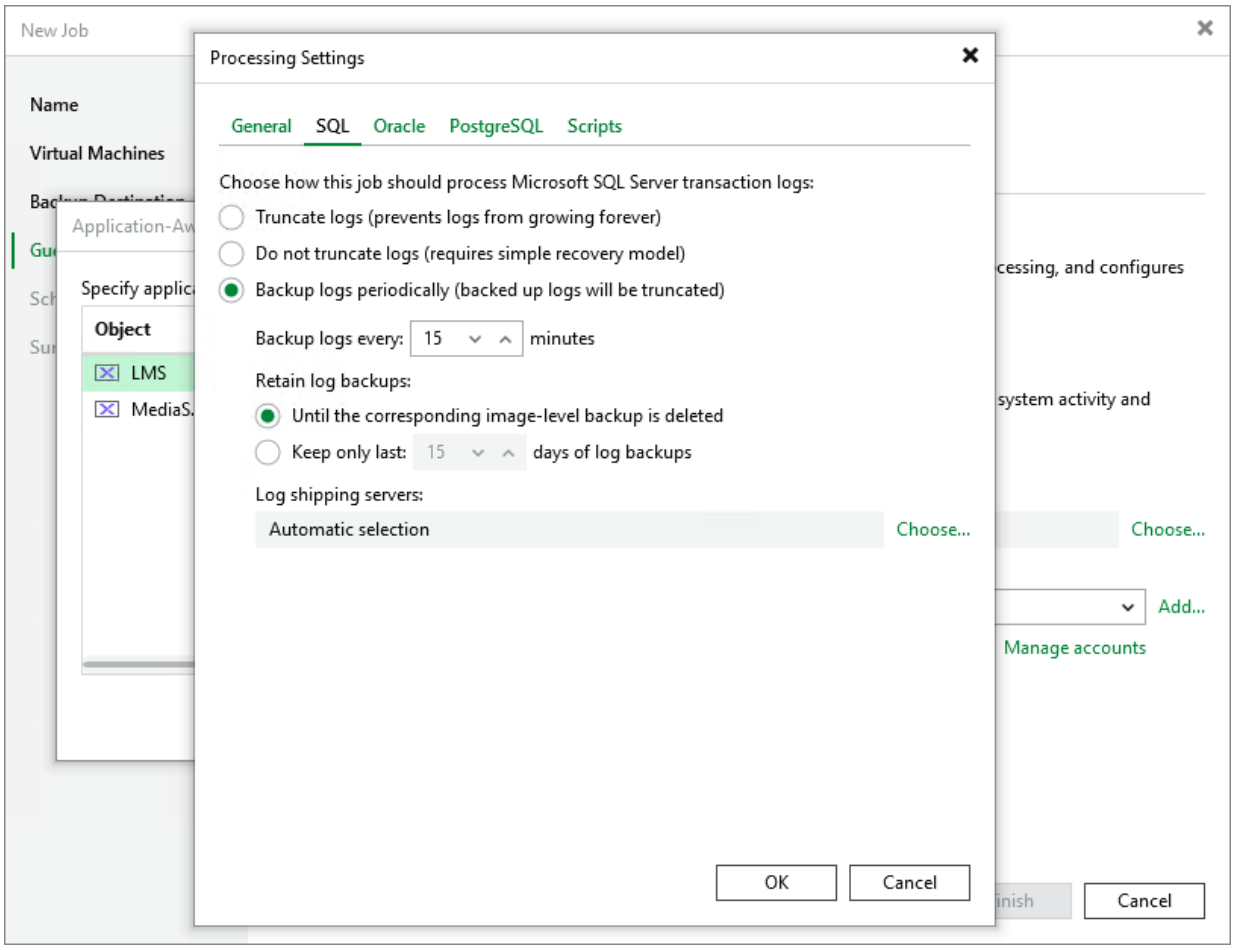


## Specifying Oracle Archived Redo Log Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the Oracle application and does not truncate archived redo logs after each successfully completed backup session — this allows you to restore Oracle databases using specific backups. To protect mission-critical Oracle databases, you can instruct Veeam Backup & Replication to create secondary restore points with archived redo logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

> **NOTES**
>
> - Veeam Backup & Replication stores image-level backups and archived redo log backups in the same repository.
> - If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary archived redo log backups will be created.

To back up Oracle archived redo logs periodically, do the following:

1. Switch to the **Oracle** tab.
2. In the **Backup logs every** field, specify how frequently you want archived redo logs to be backed up. The maximum field value is 480 minutes.

3. In the **Retain log backups** section, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove archived redo log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4.
   - Select the **Keep only last <N> days of log backups** if you want to retain archived redo log backups for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that archived redo logs backups must always be retained for a longer period than image-level backups.

   For more information on how Veeam Backup & Replication retains archived redo logs, see the Veeam Backup & Replication User Guide, section Retention for Archived Log Backup.
4. In the **Log shipping servers** section, decide whether you want to use a specific server to transfer archived redo logs backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary servers.

   For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers. Keep in mind that the list will also include Linux servers added to the backup infrastructure; however, Linux servers cannot be used as log shipping servers for processing Windows-based VMs due to technical limitations in the current version.

**TIPS**

- It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
- It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators or backup servers).

You can choose to keep the default **Do not delete archived logs** option, but in this case archived redo logs may grow large and increase the storage space consumption significantly. That is why it is recommended that you choose to remove archived redo logs that are older than a specific time limit or whose size exceeds a specific storage threshold. Keep in mind that the selected option will apply to logs of each processed Oracle database individually — and only after the backup job completes successfully.

# Configuring Access to Oracle Data

To access databases of the processed Oracle applications, Veeam Backup & Replication uses accounts with SYSDBA privileges — by default, these are the accounts you specify for accessing the VM guest OSes. To change this behavior, you can choose another account from the **Specify Oracle account with SYSDBA privileges** drop-down list.

For an account to be displayed in the list of available accounts, 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 account 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.

# Considerations and Limitations

When you configure transaction log settings, consider the following:

- For Oracle databases running in the NOARCHIVELOG mode, Veeam Backup & Replication is not able to create restore points with archived redo logs — only image-level backups will be created. For more information on how to choose between database modes, see Oracle documentation.
- For Veeam Backup & Replication to be able to access Oracle command-line tools when performing application-aware processing for Windows-based VMs, the *%ORACLE_HOME%\bin* directory must be added to the `PATH` system variable. For more information on how to set Oracle environment variables, see Oracle documentation.

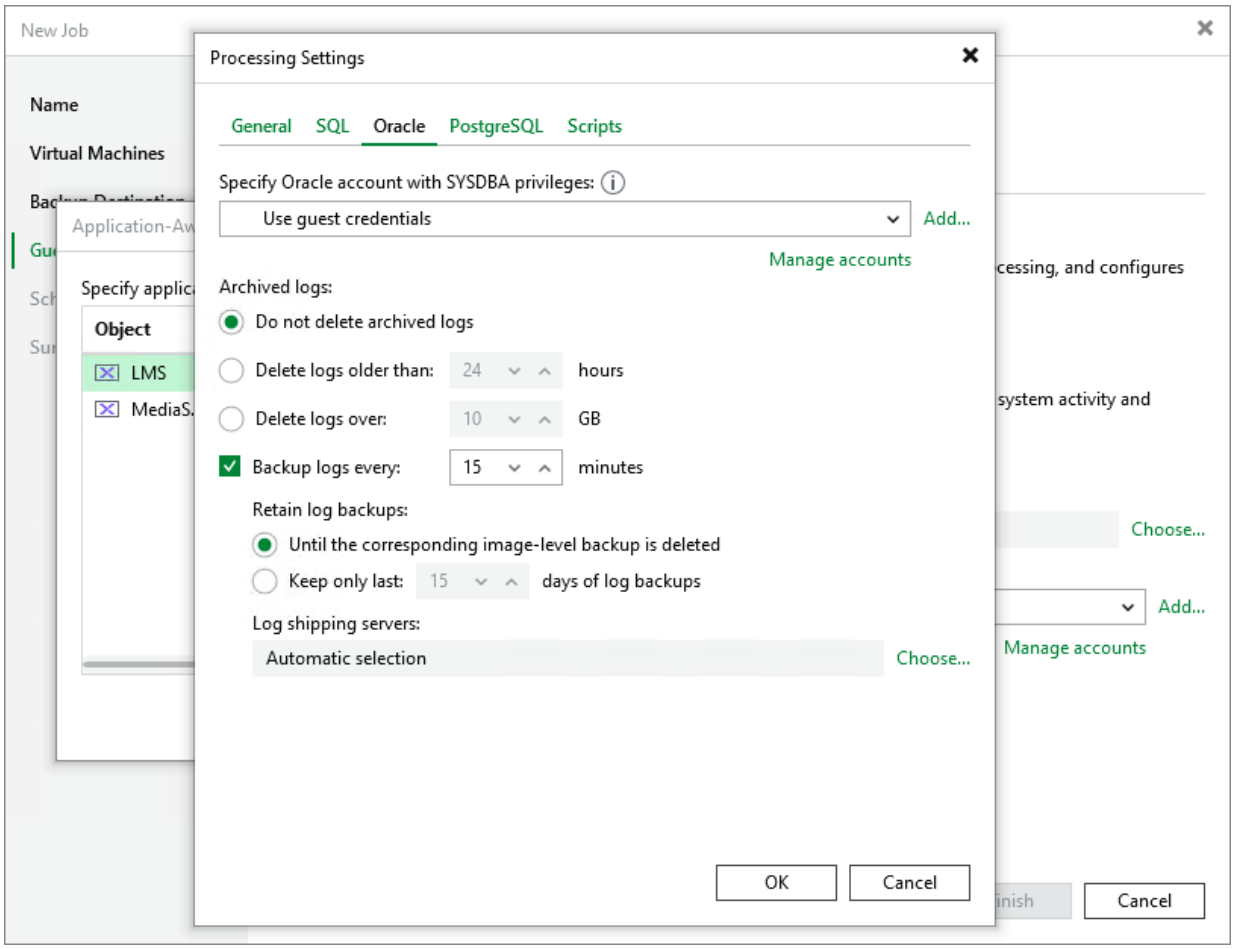


## Specifying PostgreSQL WAL Files Settings

By default, Veeam Backup & Replication creates application-consistent image-level backups of VMs running the PostgreSQL Server application and does not truncate write ahead logs (WAL) after each successfully completed backup session — this allows you to restore PostgreSQL Server databases using specific backups. To protect mission-critical PostgreSQL Server databases, you can instruct Veeam Backup & Replication to create secondary restore points with WAL logs in addition to primary image-level backups — this will allow you to restore your databases to specific points in time.

NOTES

- Veeam Backup & Replication stores image-level backups and WAL log backups in the same repository.
- If Veeam Backup & Replication fails to produce a primary image-level backup, no secondary WAL log backups will be created.

To back up PostgreSQL WAL logs periodically, do the following:

1. Switch to the **PostgreSQL** tab.
2. In the **Backup logs every** field, specify how frequently you want WAL logs to be backed up. The maximum field value is 480 minutes.
3. In the **Retain log backups** section, choose either of the following options:
   - Select the **Until the corresponding image-level backup is deleted** option if you want to remove WAL log backups and the related image-level backups at the same time, according to the retention policy settings specified at step 4 of the wizard.
   - Select the **Keep only last <N> days of log backups** if you want to retain WAL log backups for a specific time period, regardless of the retention policy settings specified for image-level backups. Note that WAL log backups must always be retained for a longer period than image-level backups.

   For more information on how Veeam Backup & Replication retains WAL logs, see the Veeam Backup & Replication User Guide, section Retention for PostgreSQL WAL Files.
4. In the **Temporary location for archive logs** section, specify the path to a folder on the PostgreSQL machine where Veeam Backup & Replication will temporarily store archive logs until they are backed up.

   Keep in mind that you must create the folder beforehand manually. Also, make sure that there is enough free space in this folder for the log files and required permissions are granted to the user account.
5. In the **Log shipping servers** section, decide whether you want to use a specific server to transfer WAL log backups or let Veeam Backup & Replication choose it automatically to reduce the load on the backup server.

   By default, Veeam Backup & Replication automatically chooses a log shipping server for each of the processed VMs based on network settings and rules listed in the Veeam Backup & Replication User Guide, section Log Shipping Servers. You can also manually limit the list of machines that may be used as log shipping servers — to do that, click **Choose**, select the **Use the specified servers only** option and then select check boxes next to the necessary servers.

   For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers.

TIPS

- It is recommended that you choose at least 2 log shipping servers for load balancing and high availability purposes.
- It is recommended that you do not choose servers that are engaged in permanent tasks consuming resources (such as WAN accelerators and backup servers).

# Configuring Access to PostgreSQL Data

To access databases of the processed PostgreSQL Server instances, Veeam Backup & Replication uses accounts with superuser privileges — by default, these are the accounts you specify for accessing the VM guest OSes. To change this behavior, you can choose another account from the **Specify PostgreSQL account with superuser privileges** drop-down list. For an account to be displayed in the list of available accounts, 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 account 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.

After you choose an account, you must also explicitly specify whether this account is a PostgreSQL database account (whose password is stored either in the Credentials Manager or in a configuration file) or a system account. In the latter case, make sure that the account has all the permissions required to access the PostgreSQL Server instance; for more information on the required permissions, see Planning and Preparation.

> **IMPORTANT**
>
> - If the password that is stored in the Credential manager is empty, Veeam Backup & Replication will be able to use this account by leveraging username map authentication. In this case, the `SYSTEM-USERNAME` variable value will be set to the name of the account used to access the VM guest OS, while the `PG-USERNAME` variable value will be set to the name of the account that you have selected from the **Specify PostgreSQL account with superuser privileges** drop-down list. For more information on username map authentication, see PostgreSQL documentation.
> - If the password is stored in a .PGPASS configuration file, make sure that it is located in the */home* directory of the selected account. For more information on .PGPASS files, see PostgreSQL documentation.

Depending on the scope of resources that you have specified at step 5a of the wizard, Veeam Backup & Replication will use the selected account in the following way:

- If the scope includes an individual VM, the account will be used to access the PostgreSQL Server instance running on this specific VM.
- If the scope includes multiple individual VMs, the account will be used to access the PostgreSQL Server instance running on each of these VMs.

- If the scope includes a VM container (such as protection domain, cluster, category or Prism Central), the account will be used to access every PostgreSQL Server instance running on VMs in this container.

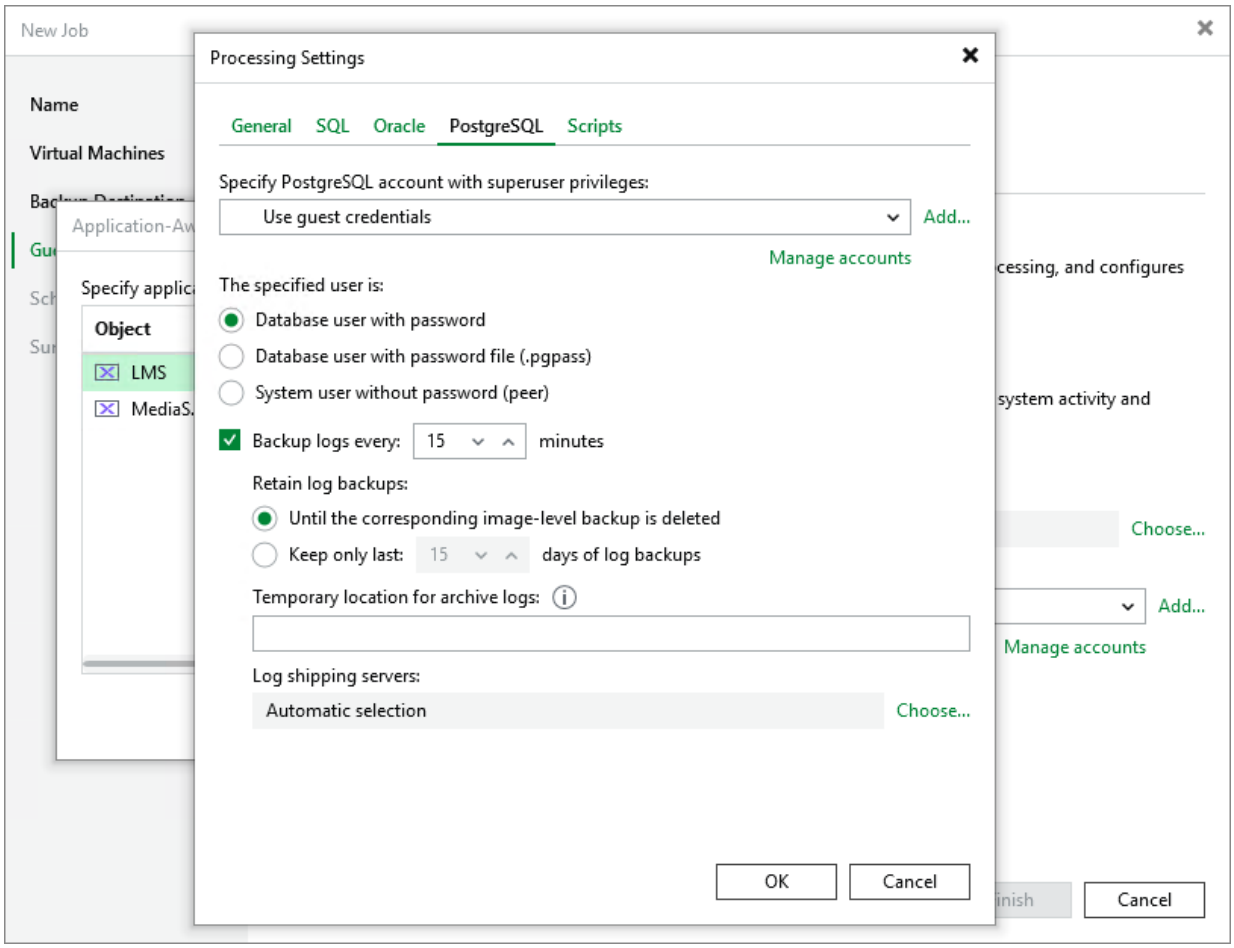


## Specifying Pre-Freeze and Post-Thaw Scripts

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, choose a scenario for script execution:
   - Select the **Require successful script execution** option if you want Veeam Backup & Replication to stop the backup process if the script fails.
   - Select the **Ignore script execution failures** option if you want to continue the backup process, even if script errors occur.
   - Select the **Disable script execution** option 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.

4. 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.

> **TIP**
>
> 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.

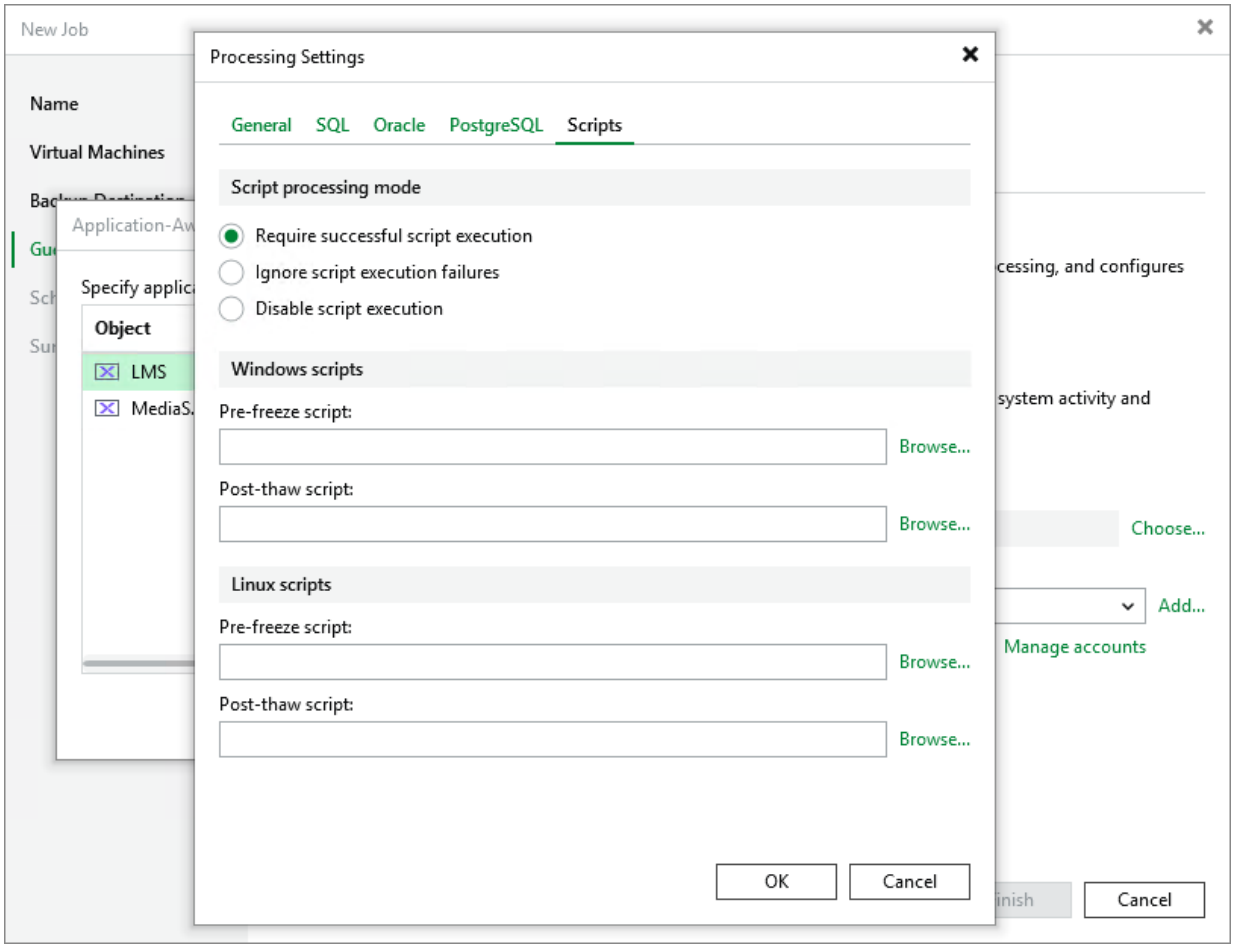

# Step 5b. Enable VM Guest OS File Indexing

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 a specific VM, do the following:

1. Click **Customize guest file system indexing settings**.
2. In the **Guest File System Indexing Options** window, select the necessary VM 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.

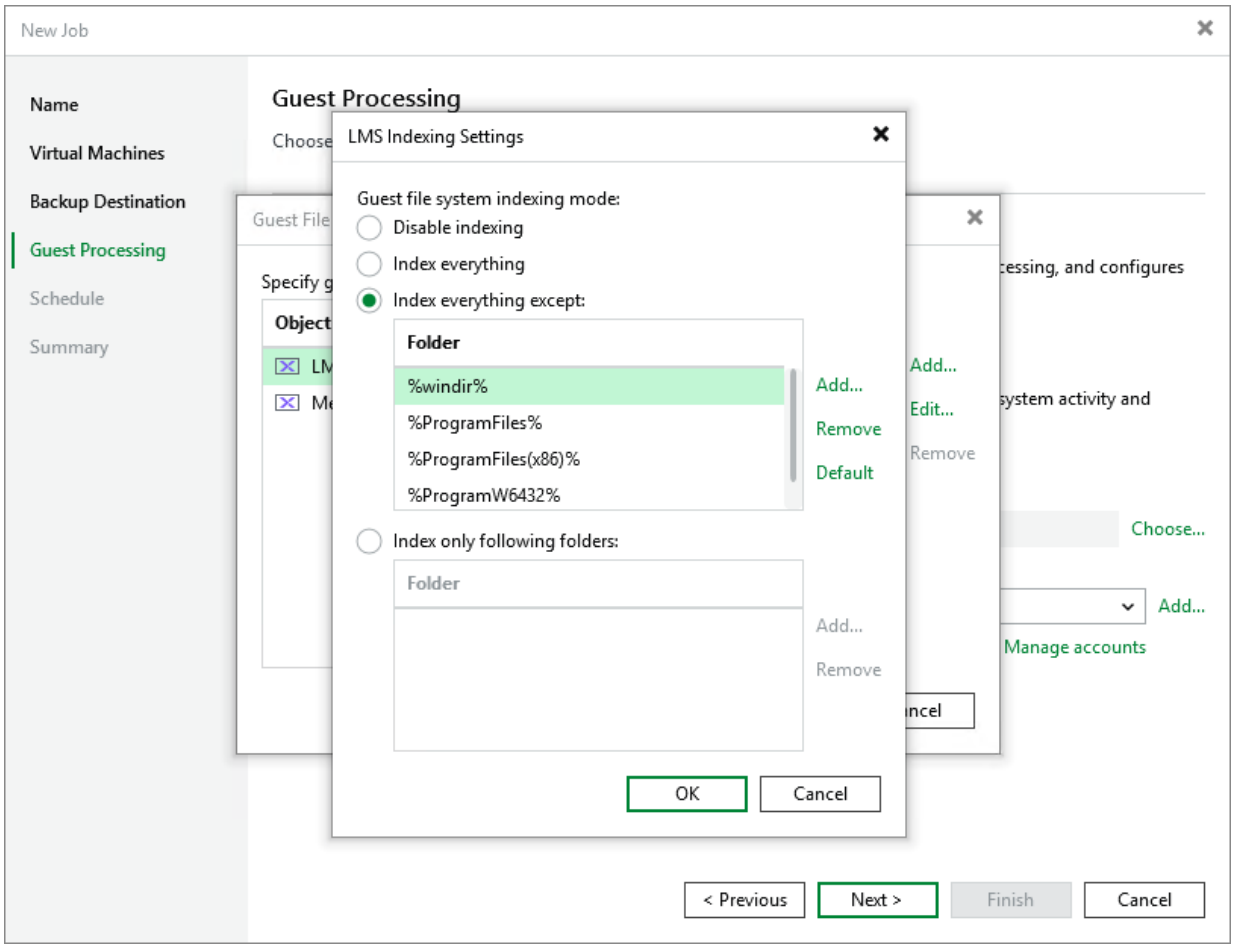
New Job
Name
Virtual Machines
Backup Destination
Guest Processing
Schedule
Summary
Guest Processing
LMS Indexing Settings
Guest file system indexing mode:
Disable indexing
Index everything
Index everything except:
Folder
%windir%
%ProgramFiles%
%ProgramFiles(x86)%
%ProgramW6432%
Add...
Remove
Default
Index only following folders:
Folder
Add...
Remove
OK
Cancel
Add...
Edit...
Remove
Choose...
Add...
Manage accounts
< Previous
Next >
Finish
Cancel

# Step 5c. Choose Guest Interaction Proxy

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 Microsoft Windows and Linux 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 servers.

For a server to be displayed in the list of available log shipping servers, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, sections Adding Microsoft Windows Servers and Adding Linux Servers.

> **IMPORTANT**
>
> Due to technical limitations, Linux-based proxies cannot access Windows guest OSes in the current version. That is why if you have added Windows-based VMs to the backup scope at step 3 of the wizard, you must also add at least one Microsoft Windows server to the backup infrastructure.

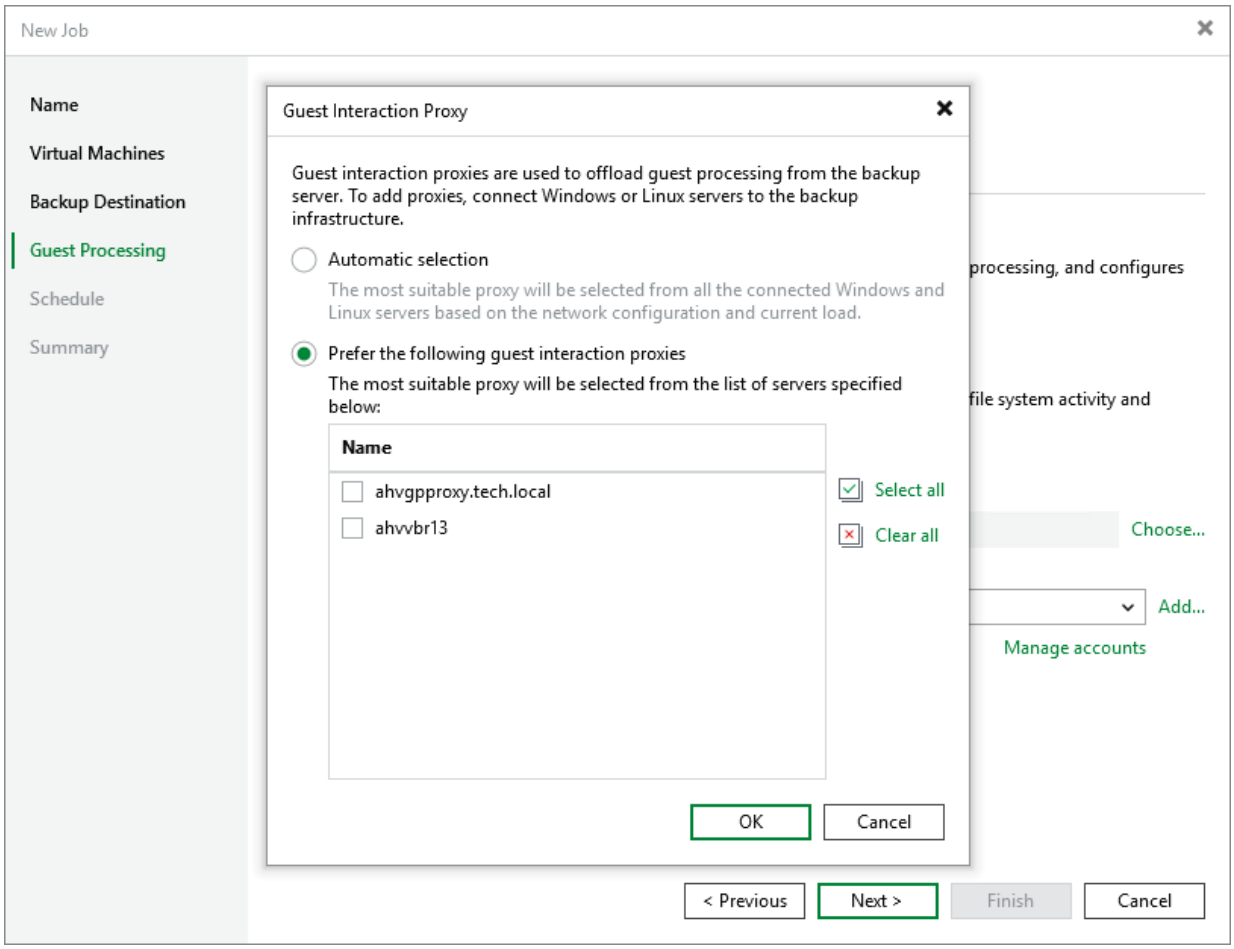

# Step 5d. Manage VM Guest OS Credentials

If you enable application-aware processing or instruct Veeam Backup & Replication to create a catalog of VM files and folders, you must also specify a user whose credentials will be used to communicate with VM guest OSes. Note the specified user 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 **Verify Network Connectivity**.

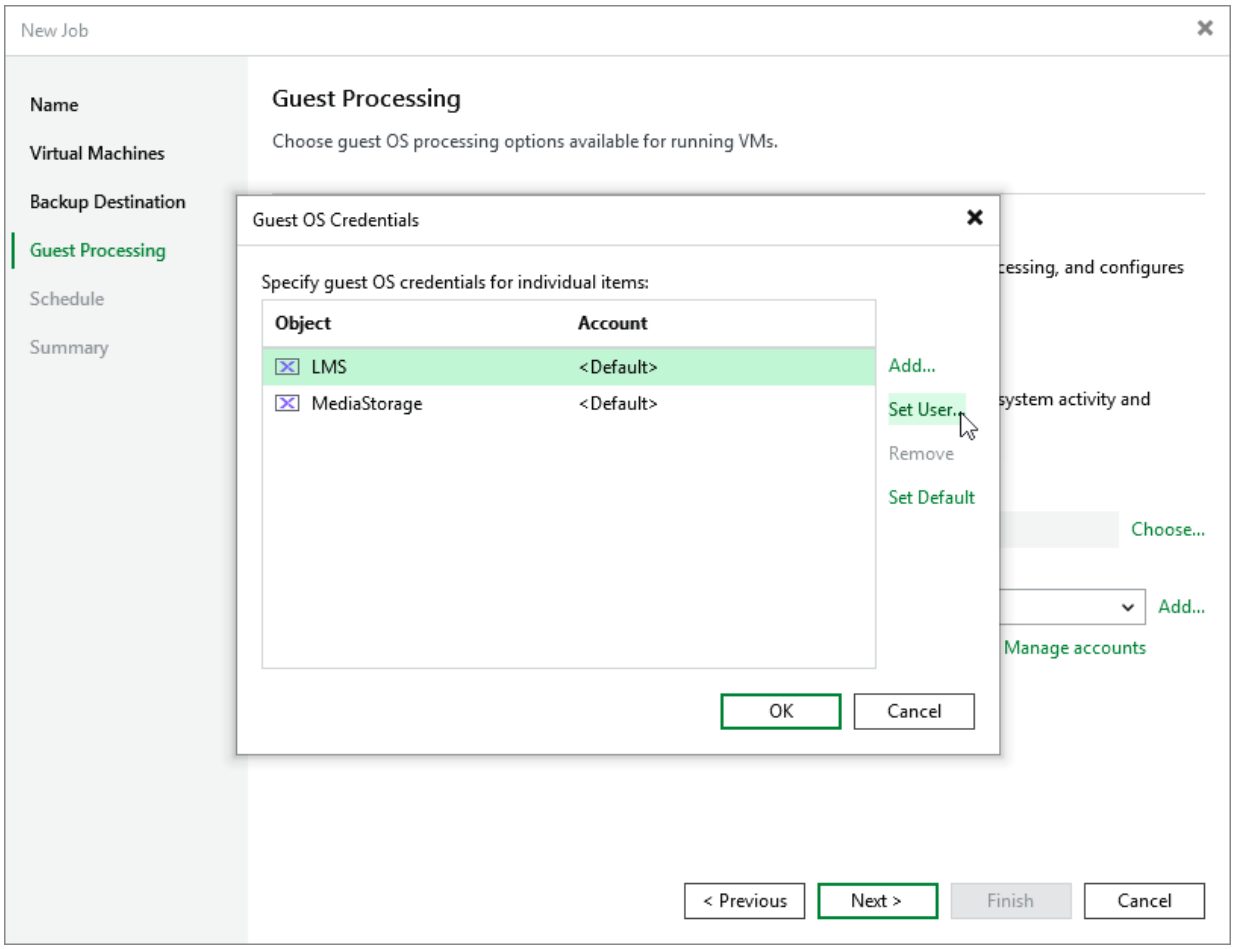

# Step 6. Specify Job Scheduling Options

At the **Schedule** step of the wizard, you can instruct Veeam Backup & Replication 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 & Replication 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 & Replication 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 & Replication processes only those VMs that failed to be backed up during the previous attempt.

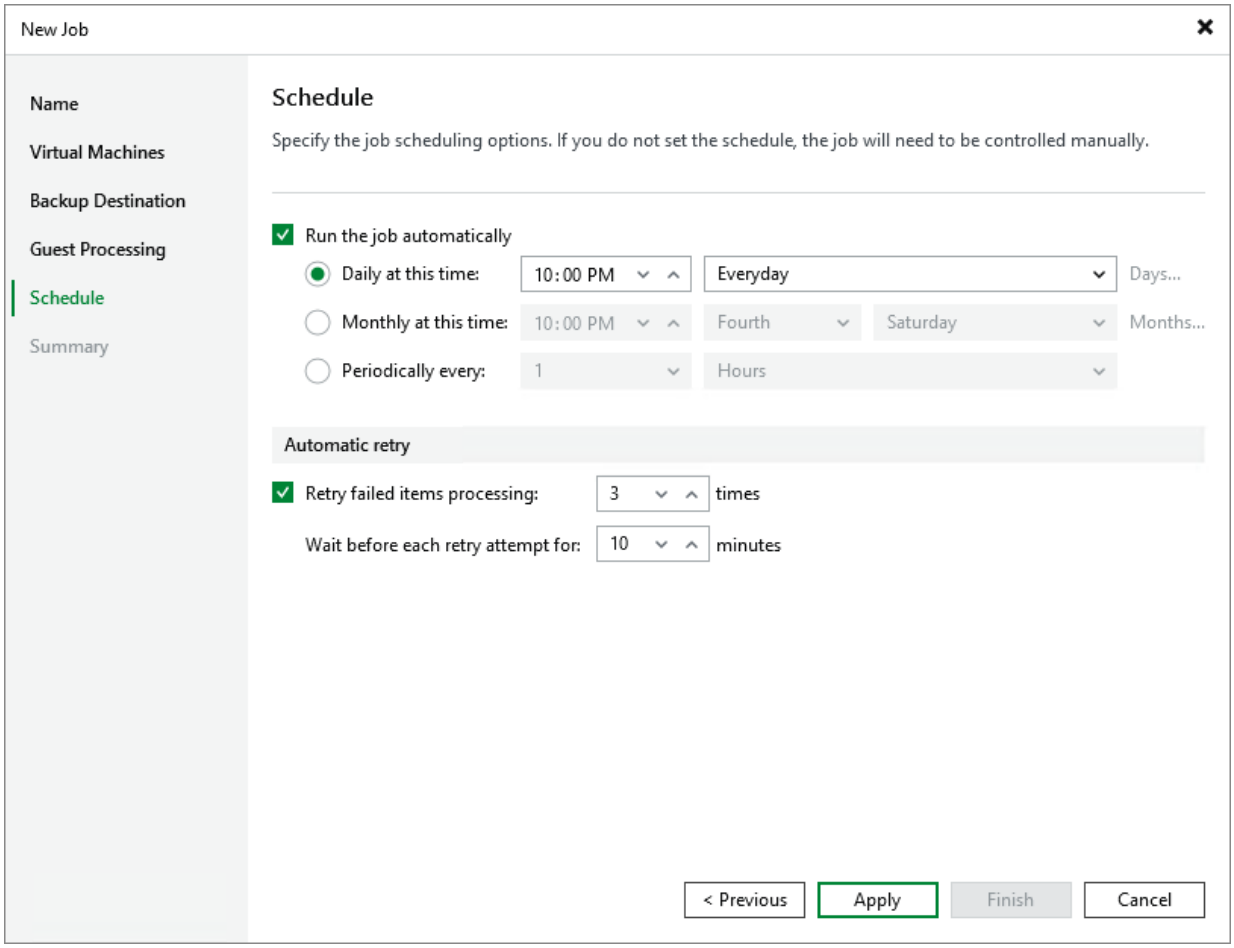

# Step 7. Finish Working with Wizard

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

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

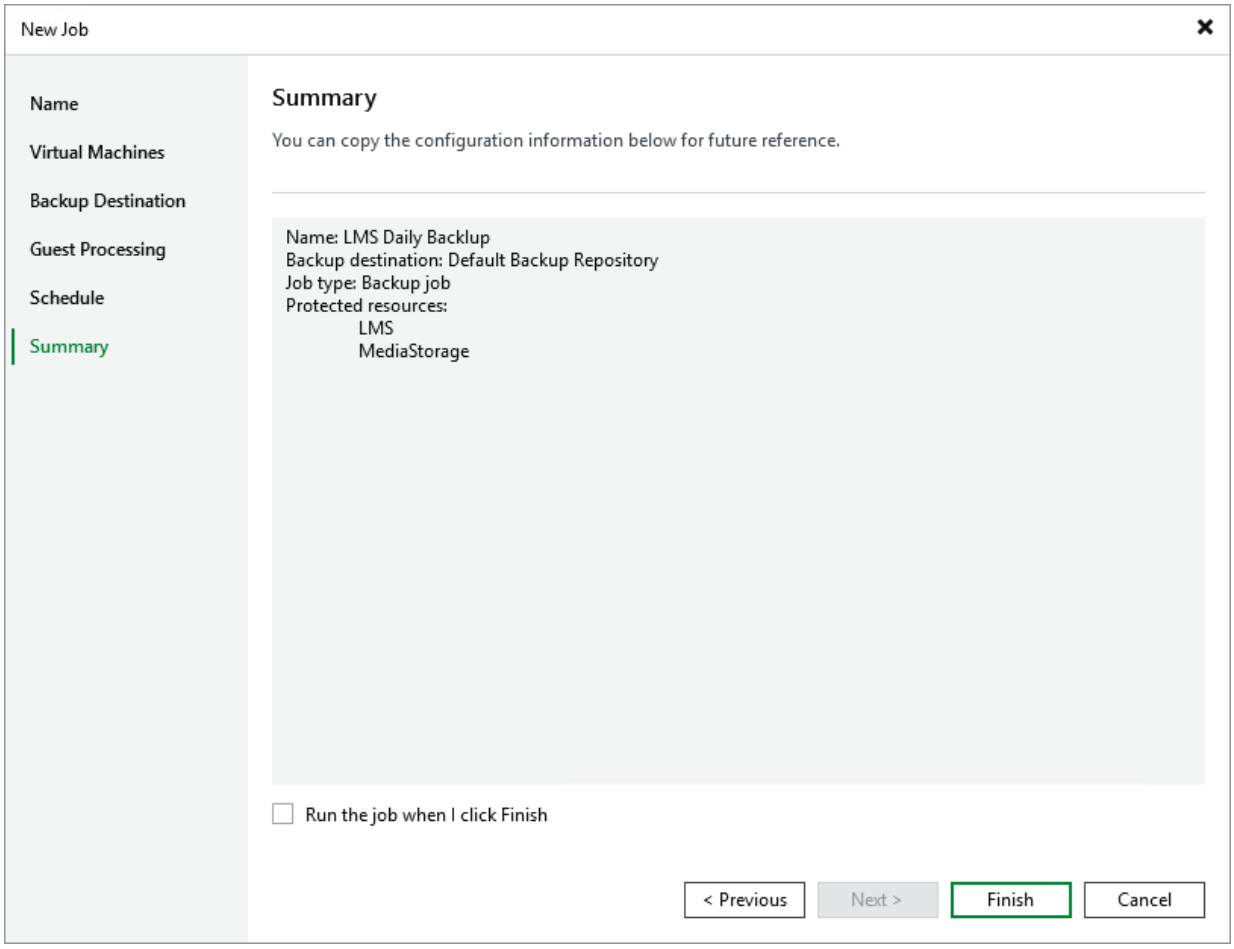

# Analyzing Performance Bottlenecks

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.

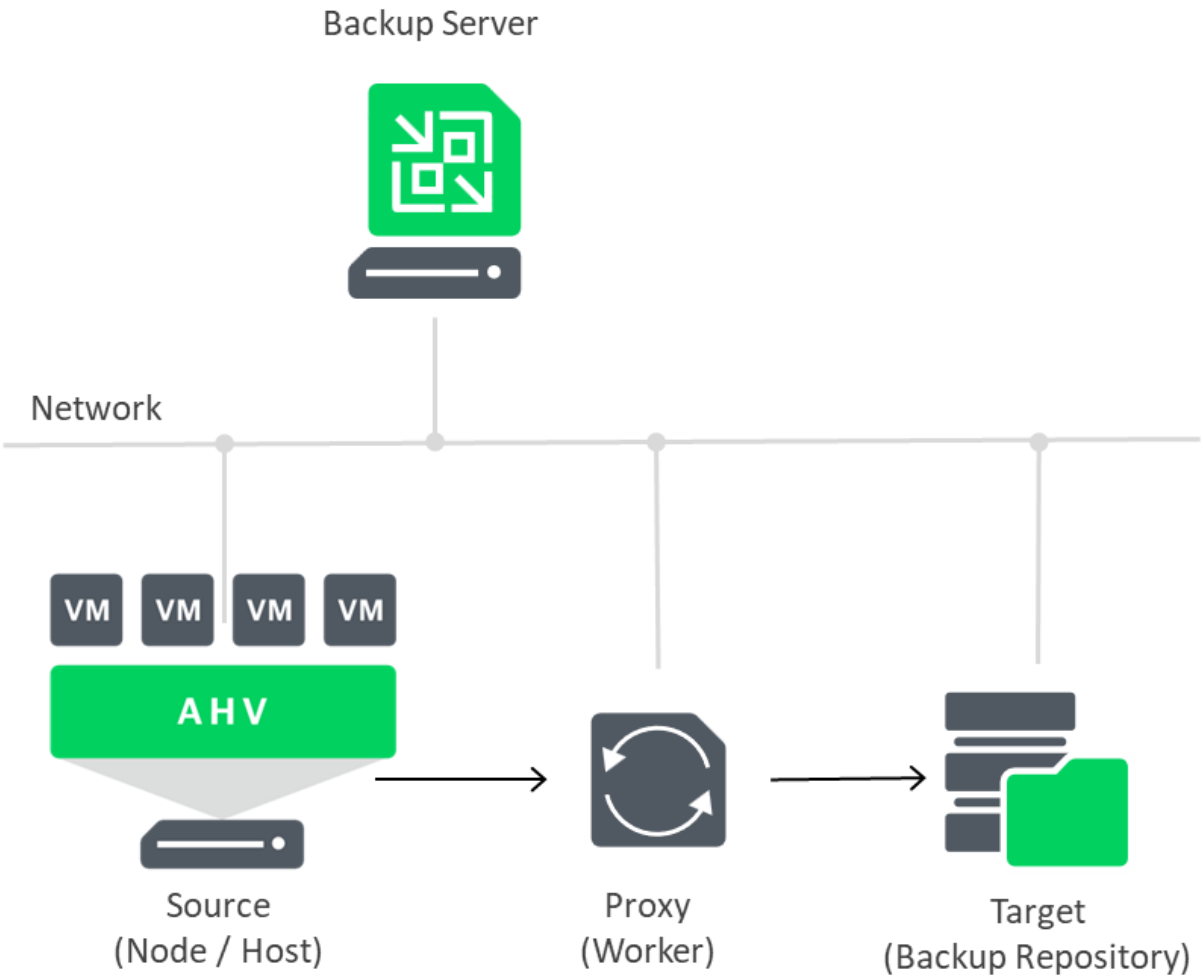


While evaluating the data transmission process, Veeam Backup & Replication analyzes performance of all the data flow components:

- **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, check the bottleneck statistics:
   - 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.

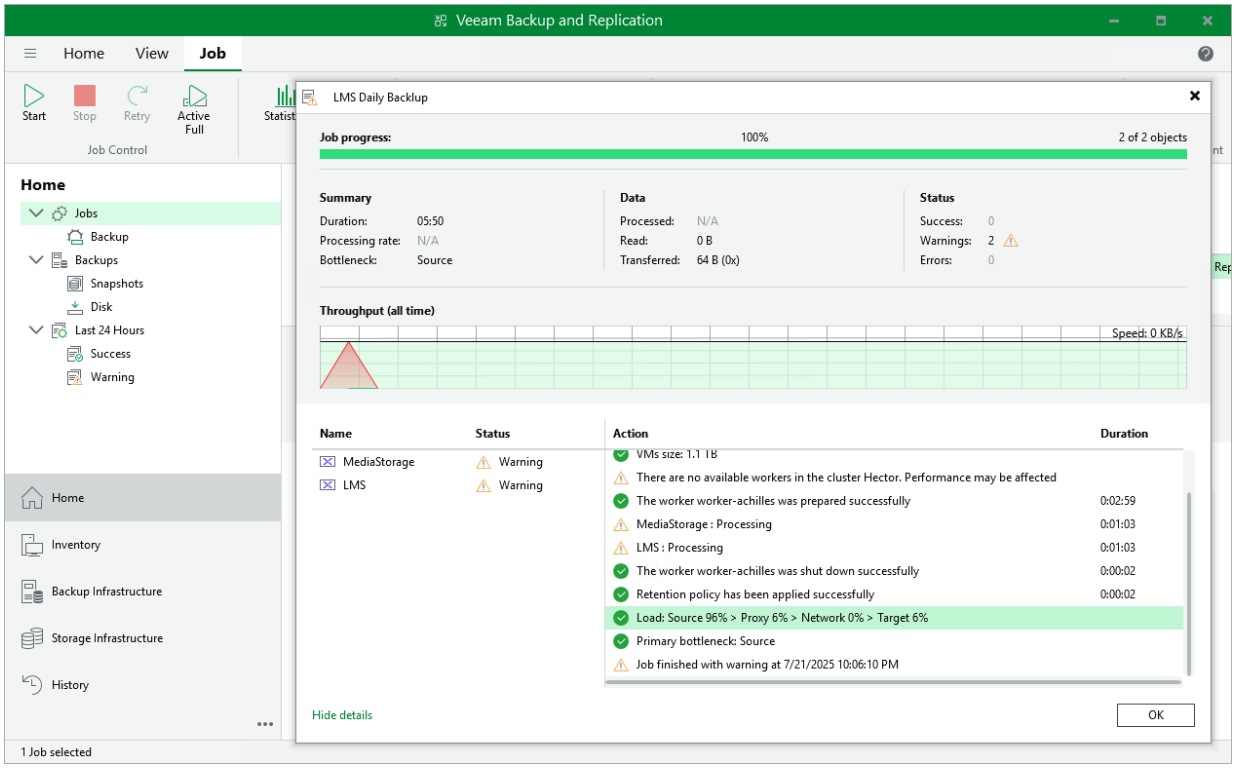

# Adding VMs to Job

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 from the **Inventory** view.

> **NOTE**
>
> If a VM is already added to the exclusion list in the job, the VM will not be protected.

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

1. In the Veeam Backup & Replication console, open the **Inventory** view.
2. In the inventory pane, select **Nutanix AHV** and expand the Prism Central or cluster where the VM resides.
3. It the working area, select the VM that you want to back up, click **Add to backup job** on the ribbon and choose the necessary job.

   Alternatively, right-click the VM, select **Add to backup job** and choose the necessary job.

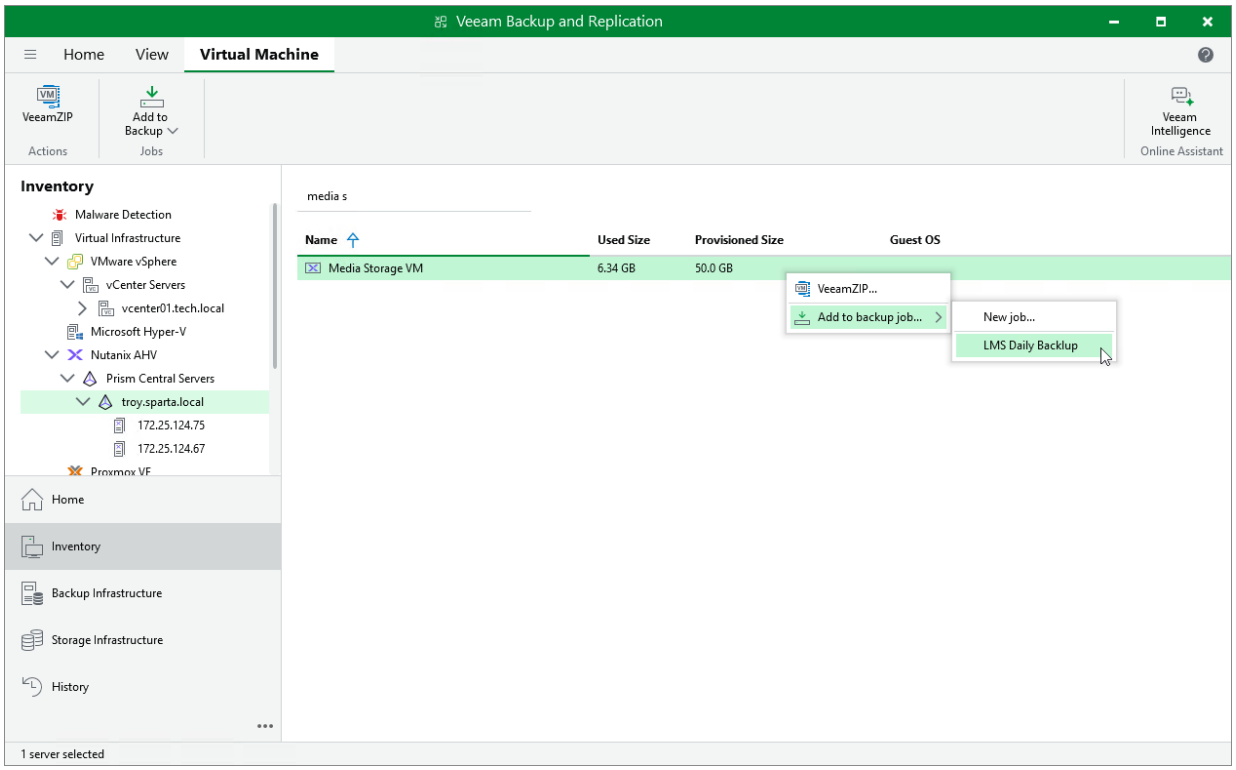

# Cloning Jobs

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. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Clone** on the ribbon.

   Alternatively, right-click the job and select **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 Plug-in for Nutanix AHV disables the cloned job. To enable the cloned job, select it in the job list and click **Enable**.

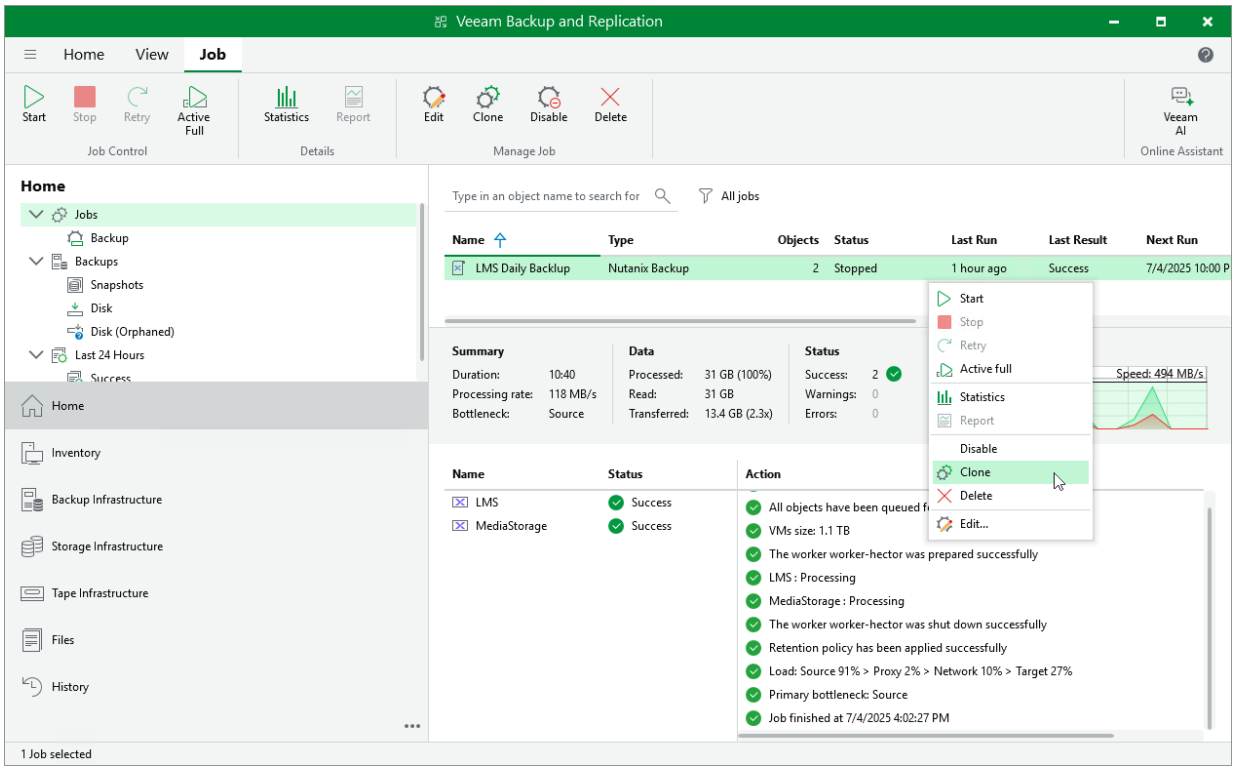

# Starting and Stopping Jobs

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 & Replication 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. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Start** or **Stop** on the ribbon.

   Alternatively, right-click the job and select **Start** or **Stop**.

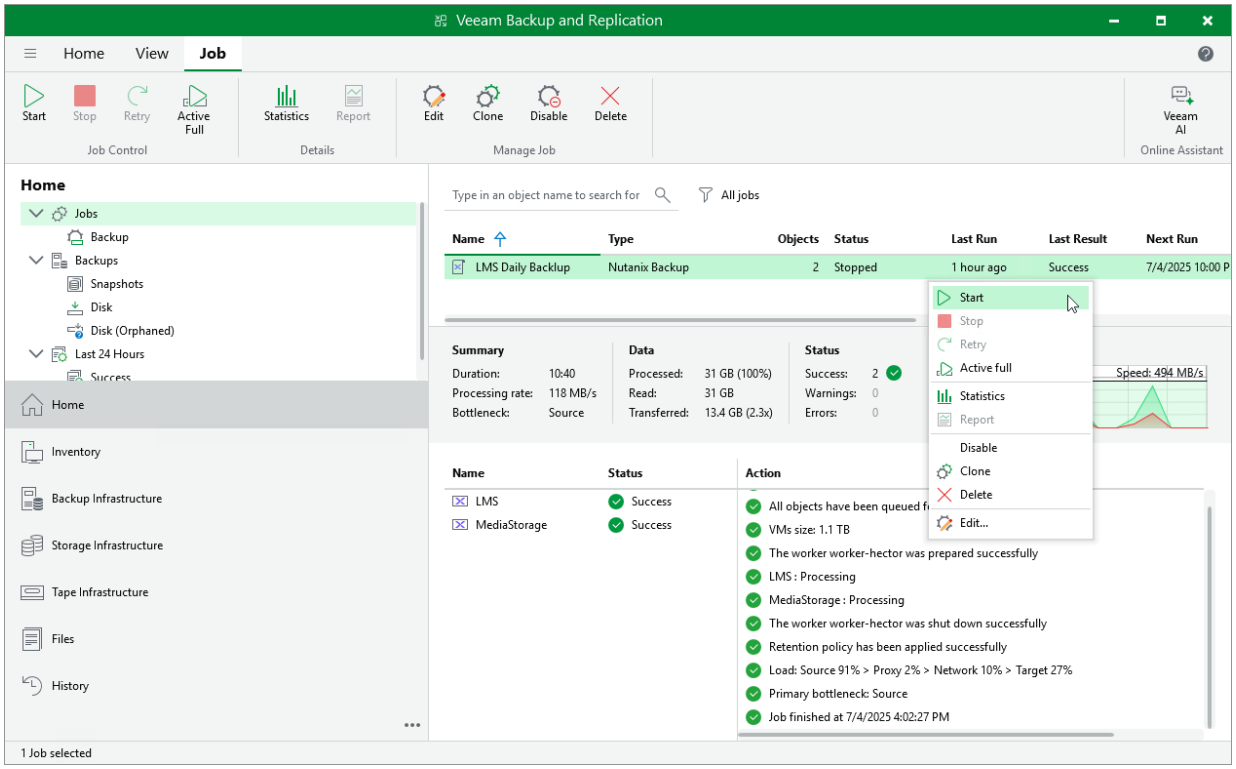

# Retrying Jobs

If a job fails, you can retry the backup operation. When you perform a retry, Veeam Backup & Replication 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. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Retry** on the ribbon.

   Alternatively, right-click the job and select **Retry**.

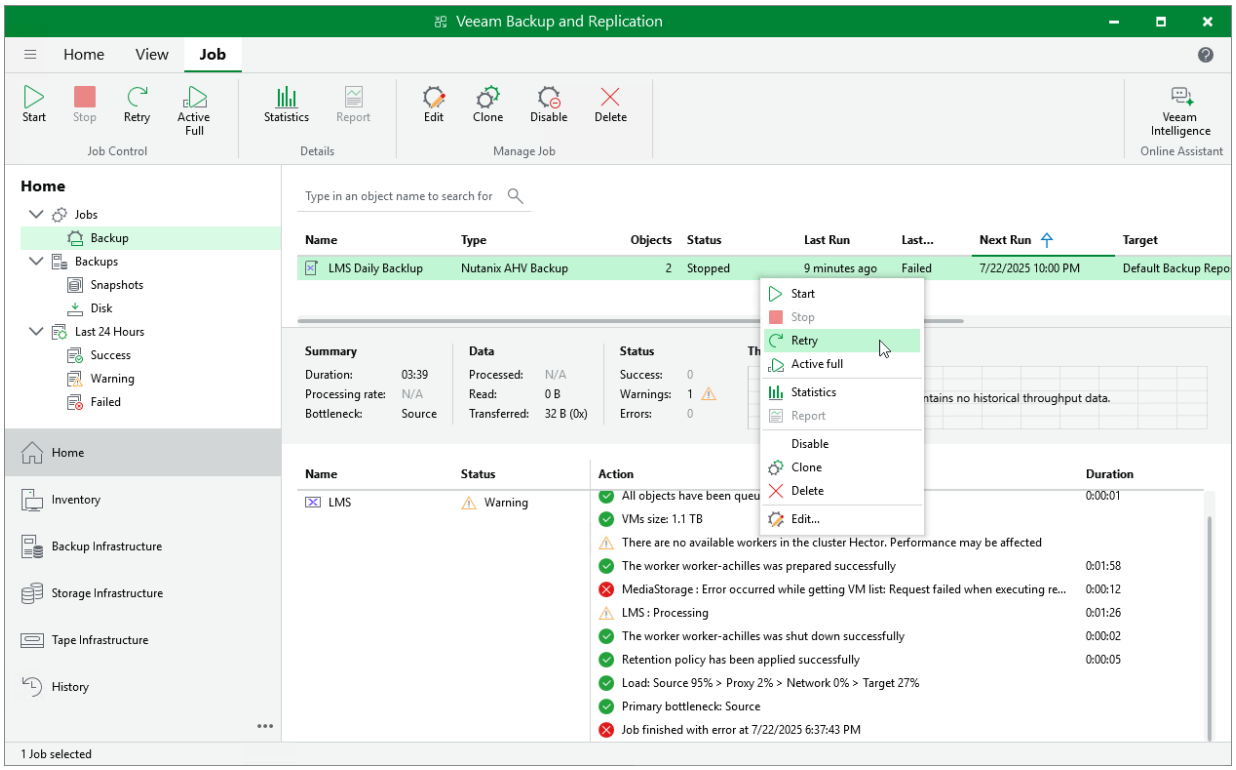

# Editing Job Settings

For each job, you can modify settings configured while creating the job:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the necessary job and click **Edit** on the ribbon.

   Alternatively, right-click the job and select **Edit**.
4. Complete the **Edit Job** wizard:

   a. To provide a new name and description for the job, follow the instructions provided in section Creating Backup Jobs (step 2).

   b. To edit the backup scope, follow the instructions provided in section Creating Backup Jobs (step 3).

   c. To change the backup repository where backups are stored, to configure backup job retention settings, to schedule active and synthetic full backups, to configure health checks and email notifications, follow the instructions provided in section Creating Backup Jobs (step 4).

   d. To modify settings for application-aware processing of VMs included into the backup scope, follow the instructions provided in section Creating Backup Jobs (step 5).

   e. To modify the job schedule and configure automatic retry settings, follow the instructions provided in section Creating Backup Jobs (step 6).

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

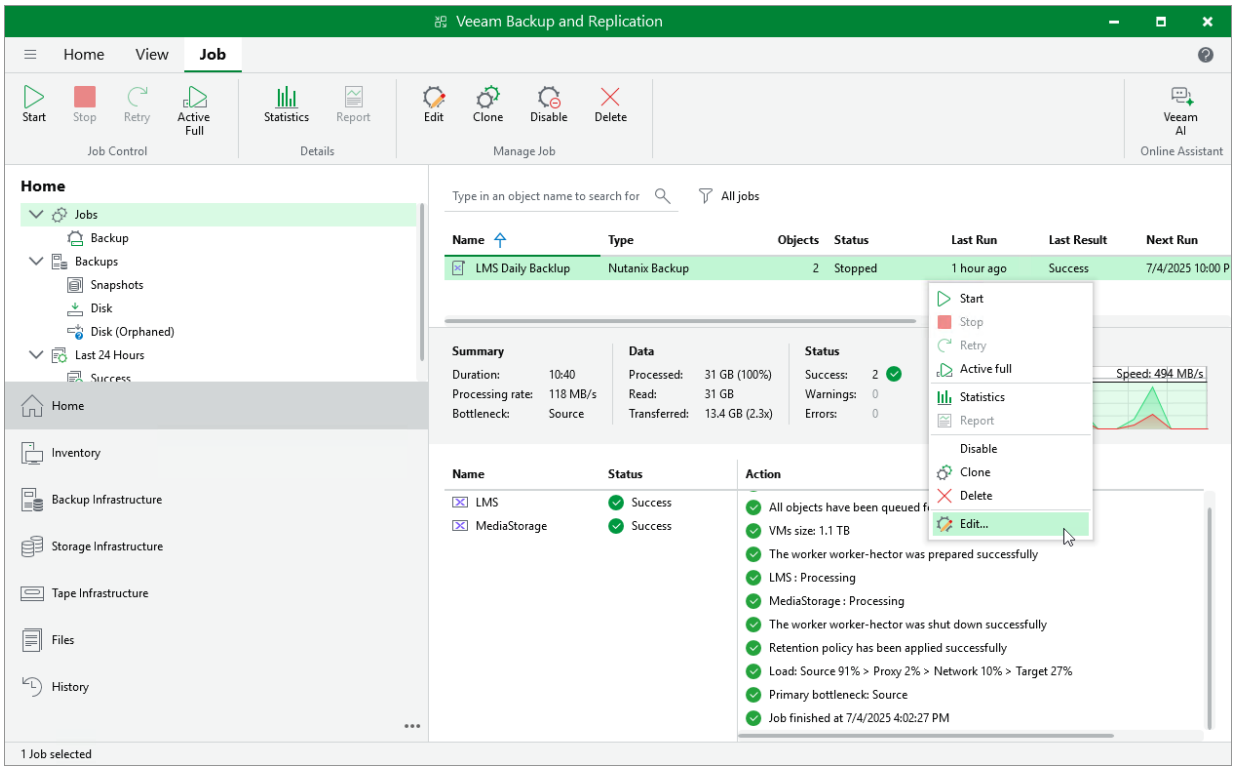

# Enabling and Disabling Jobs

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. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Disable** on the ribbon.

   Alternatively, right-click the job and select **Disable**.

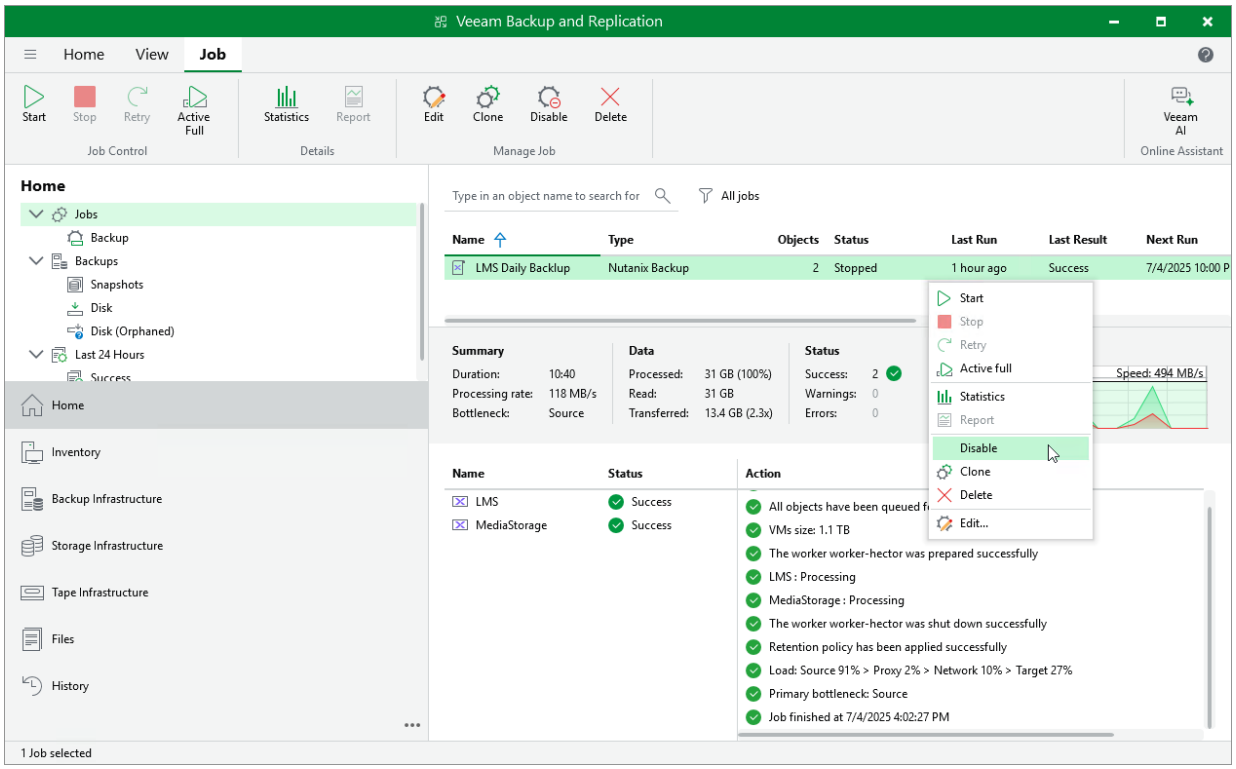

# Deleting Jobs

You can permanently delete a job from the Veeam Backup & Replication 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 a backup job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Delete** on the ribbon.

   Alternatively, right-click the job and select **Delete**.

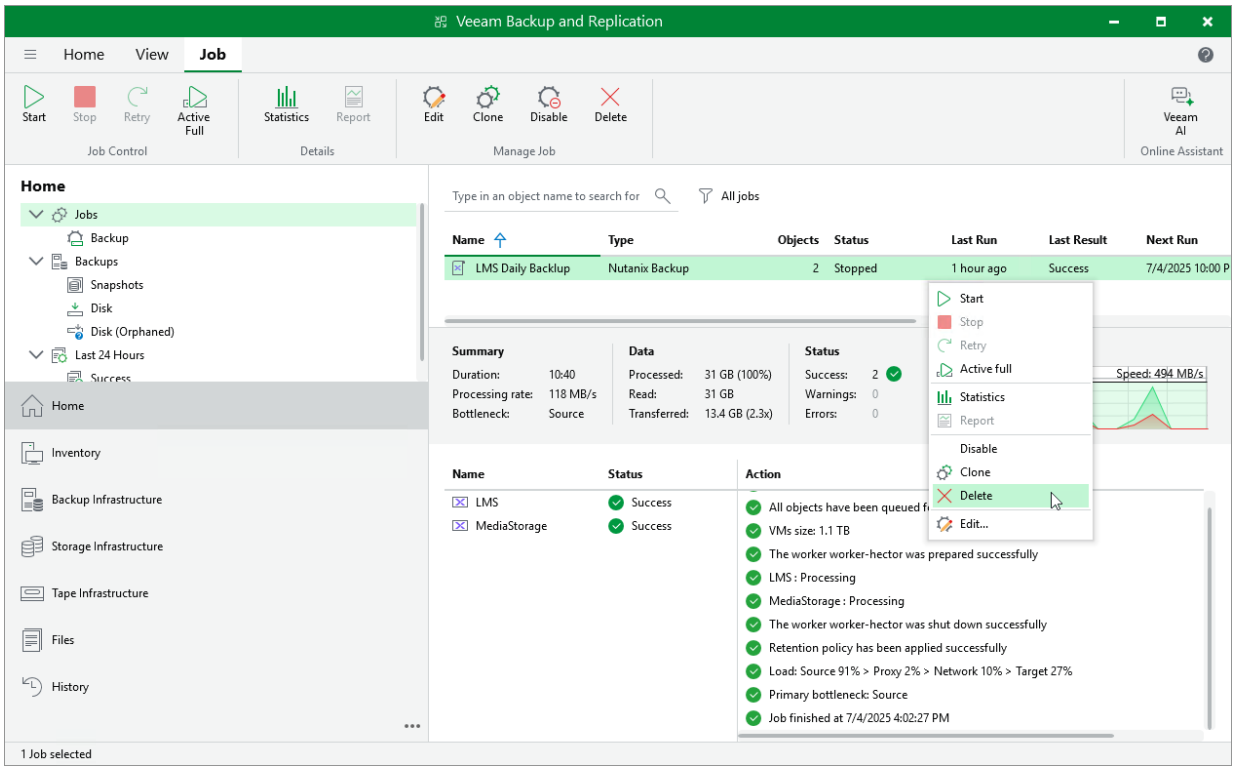

# Creating Active Full Backup

You can manually create an active full backup for all VMs added to a backup job.

1. Open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Active Full** on the ribbon.

   Alternatively, right-click the job and select **Active full**.

> **TIP**
>
> To create active full backup automatically according to a specific schedule, configure backup job settings as described in section Creating Backup Jobs (step 4).

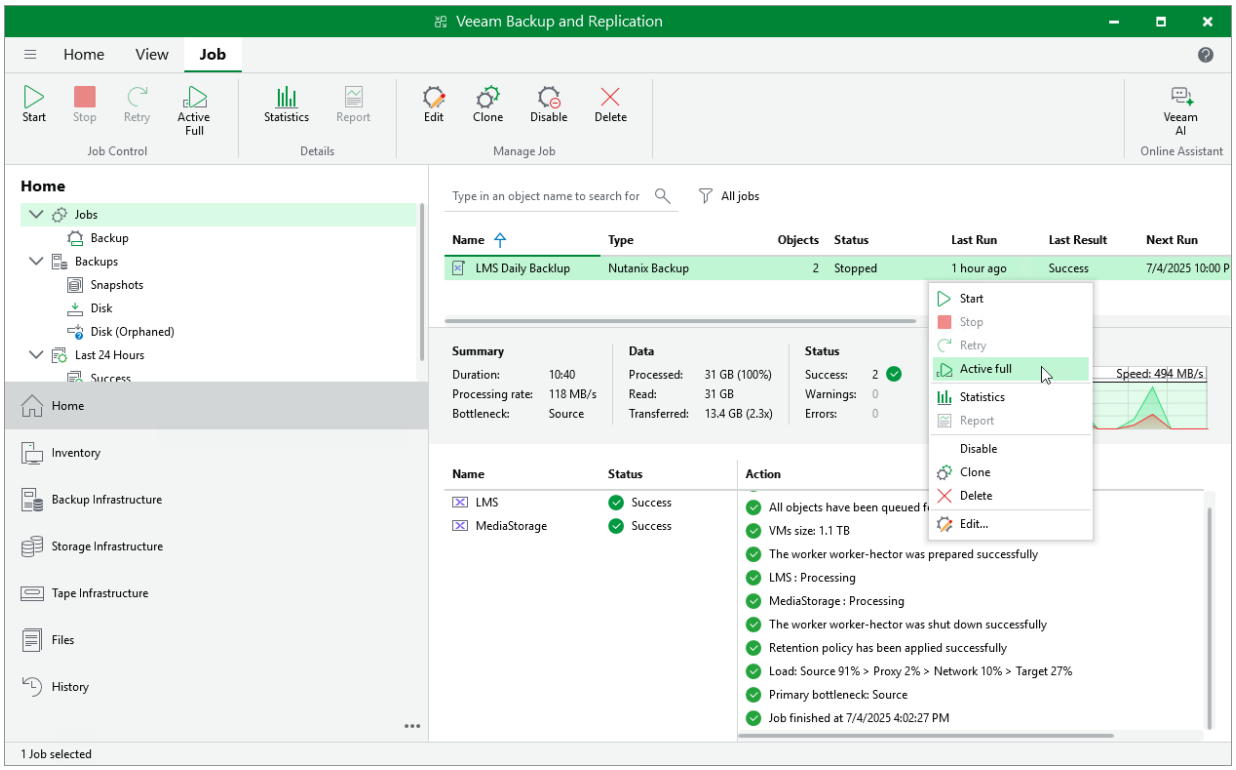

# Creating VeeamZIP Backups

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** and expand the Prism Central or cluster where the VM resides.
3. It the working area, select the VM that you want to back up and click **VeeamZIP** on the ribbon.

   Alternatively, right-click the VM 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.

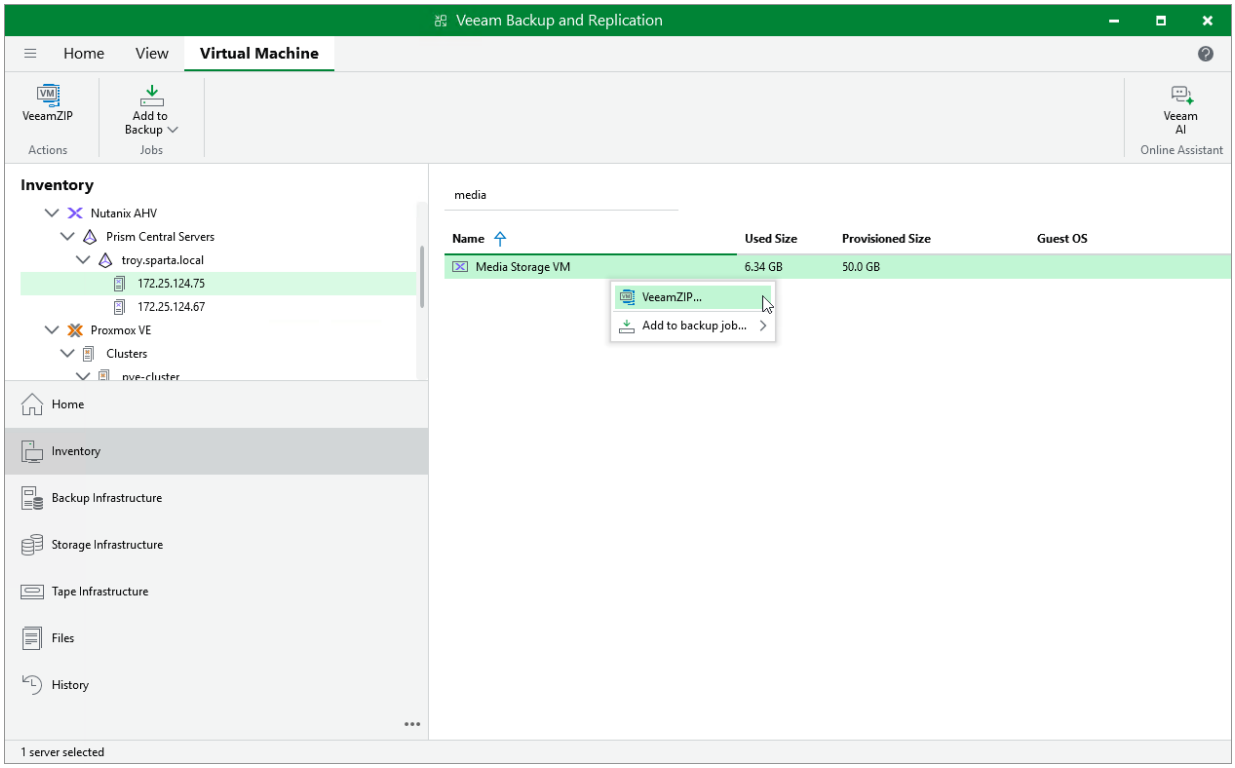

# Managing Backups

You can perform the following operations with backup files:

- Viewing Backup Properties
- Verifying Backups
- Exporting Backups
- Copying Backups
- Copying Backups to Tapes
- Deleting Backups

# 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 and the collection of restore points created for this backup is represented with a set of properties, such as:

- **Object Name** – the name of a protected VM.
- **Original Size** – the total amount of disk space allocated to the VM.
- **File Name** – the name of a restore point.
- **Data Size** – the amount of processed VM data.
- **Backup Size** – the amount of backed-up VM data.
- **Data Reduction** – the ratio between the data size and backup size.
- **Date** – the date and time when the restore point was created.
- **Immutable Until** – the date when the immutability period for the restore point will expire.
- **Status** – the result of the most recent malware scan performed for the restore point.

> **TIP**
>
> If a restore point is marked as *Infected* but you know that this point is clean, you can change its status manually. To do that, select the restore point and click **Malware** > **Mark as clean**. To learn how to manage infected restore points, see Veeam Backup & Replication User Guide, section Managing Malware Status.

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 job name and select **Properties**.

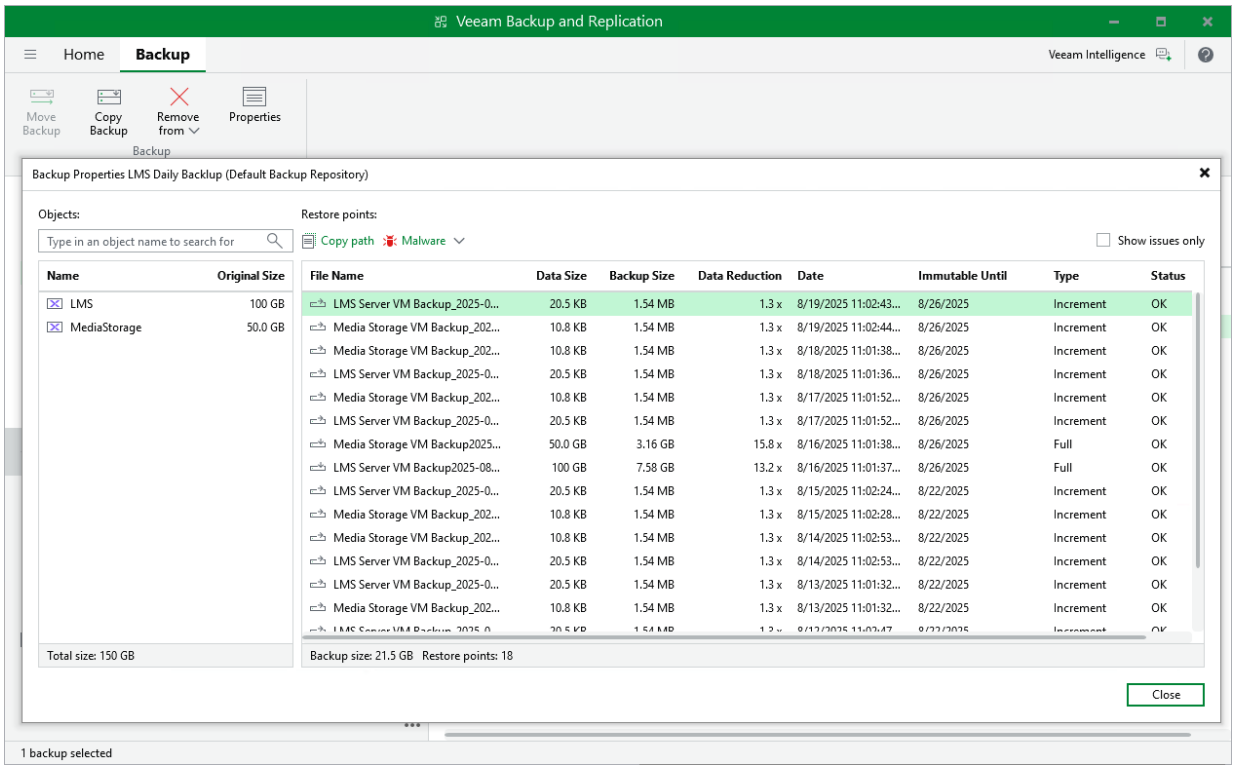

# Verifying Backups

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. 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 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.

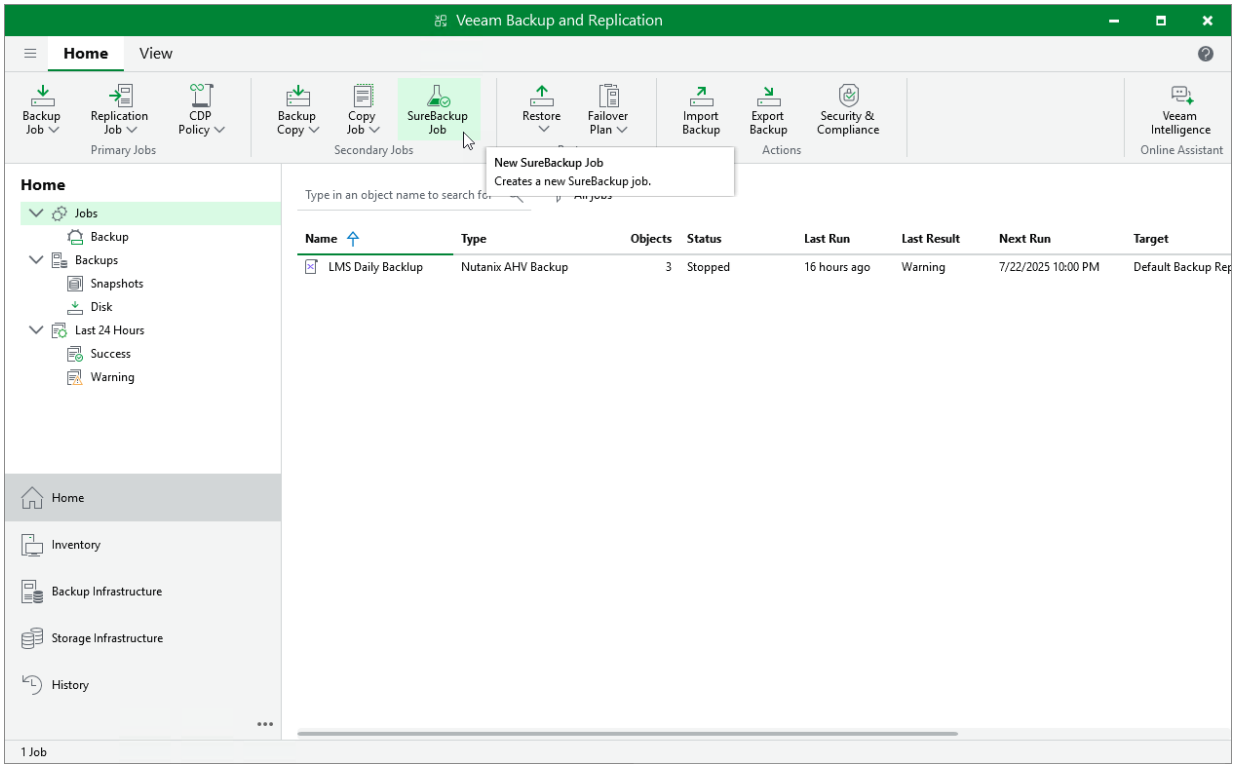

# Exporting Backups

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.

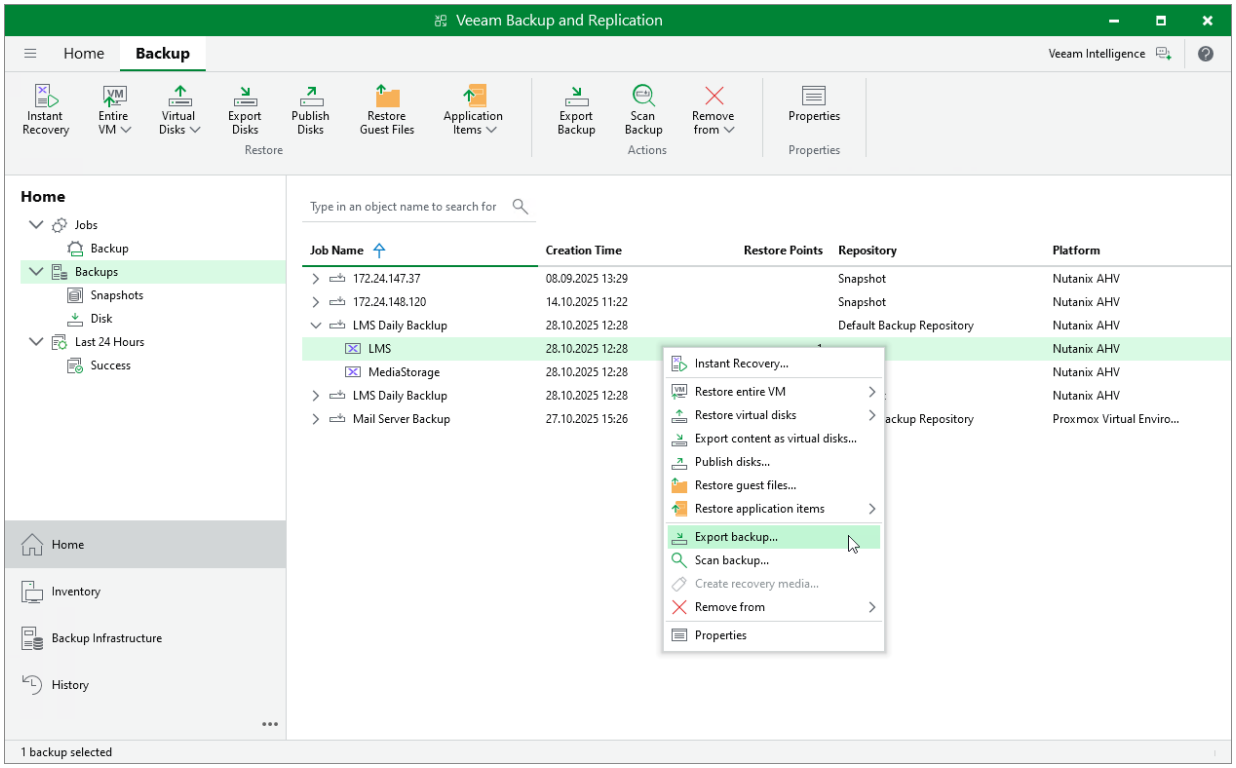

# Copying Backups

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**
>
> Veeam Plug-in for Nutanix AHV copies all backups produced by a source backup job — you cannot select backups of specific VMs.

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.

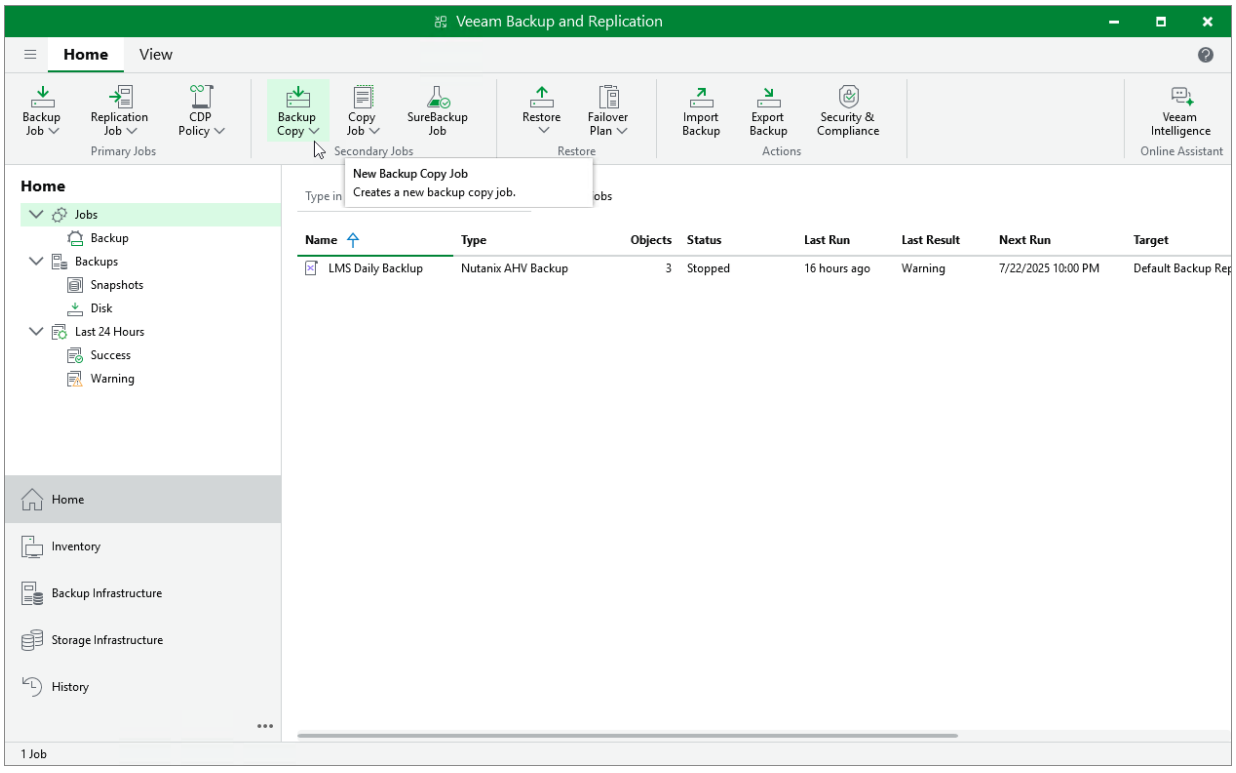

# Copying Backups to Tapes

You can create archives of Nutanix AHV VM backups and copy them to tapes for long-term storage. Veeam Plug-in 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.

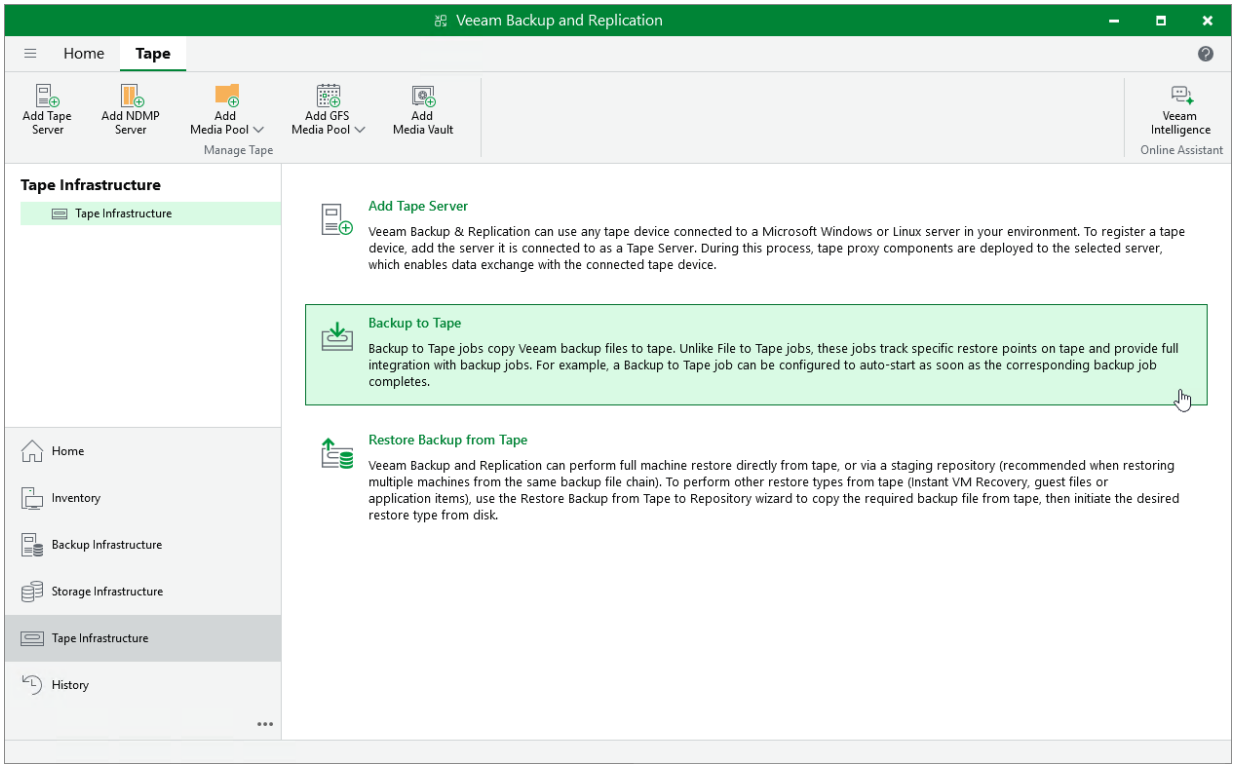

# Deleting Backups

By default, Veeam Plug-in for Nutanix AHV maintains backups in backup repositories according to retention policy settings saved in the backup metadata. If Veeam Plug-in 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 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 **Remove 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.

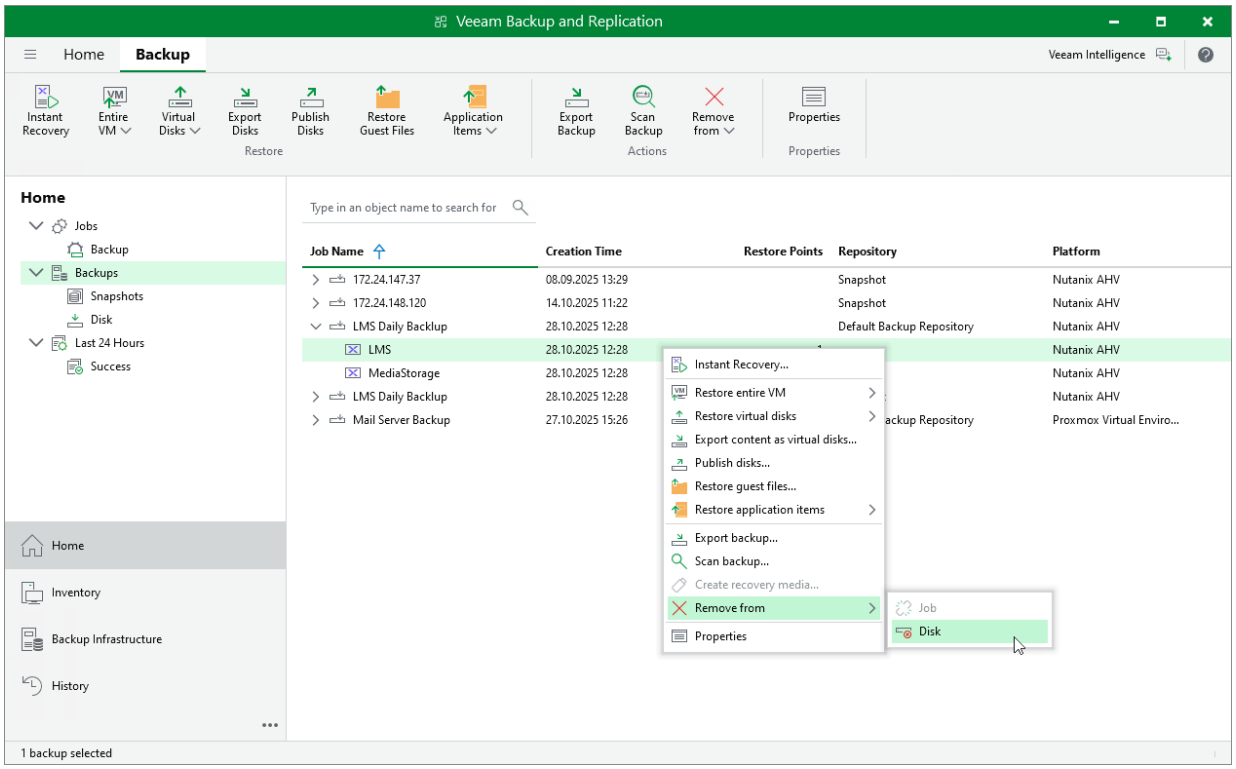

# Performing Restore

In various disaster recovery scenarios, Veeam Plug-in 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 Plug-in 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 manually in Nutanix AHV Prism Central or Prism Element console
- Snapshots of Nutanix AHV VMs created manually in Nutanix AHV Prism Central or Prism Element console
- Backups of Nutanix AHV VMs created by Veeam Plug-in 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 oVirt KVM VMs created by Veeam Plug-in for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups of Proxmox VE VMs created by Veeam Plug-in for Proxmox VE
- Backups of Scale Computing HyperCore VMs created by Veeam Plug-in for Scale Computing HyperCore
- 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 virtual and physical machines created by Veeam Agent for Microsoft Windows and Veeam Agent for Linux

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

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 Plug-in 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 Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV attaches VM disks with the restored data to the target VM disk nodes using their original bus types. Veeam Plug-in 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 Veeam Plug-in for Nutanix AHV 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 Plug-in for Nutanix AHV does not restore the VM affinity policy configuration. Therefore, you must manually configure the affinity policy as described in Nutanix documentation.
- When restoring a Windows 11 VM to the Prism Central running AOS version prior to 7.0, Veeam Plug-in for Nutanix AHV does not restore the Virtual Trusted Platform Module (vTPM) configuration. Therefore, you must manually enable vTPM for the VM 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 job, right-click the VM you want to restore and select **Restore entire VM** > **Nutanix AHV**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Nutanix AHV** on the ribbon.

> **TIP**
>
> To restore a VM from a snapshot taken in the Nutanix AHV Prism Element console, expand the Nutanix AHV cluster where the VM resides, right-click the VM and select **Restore entire VM to Nutanix AHV**.

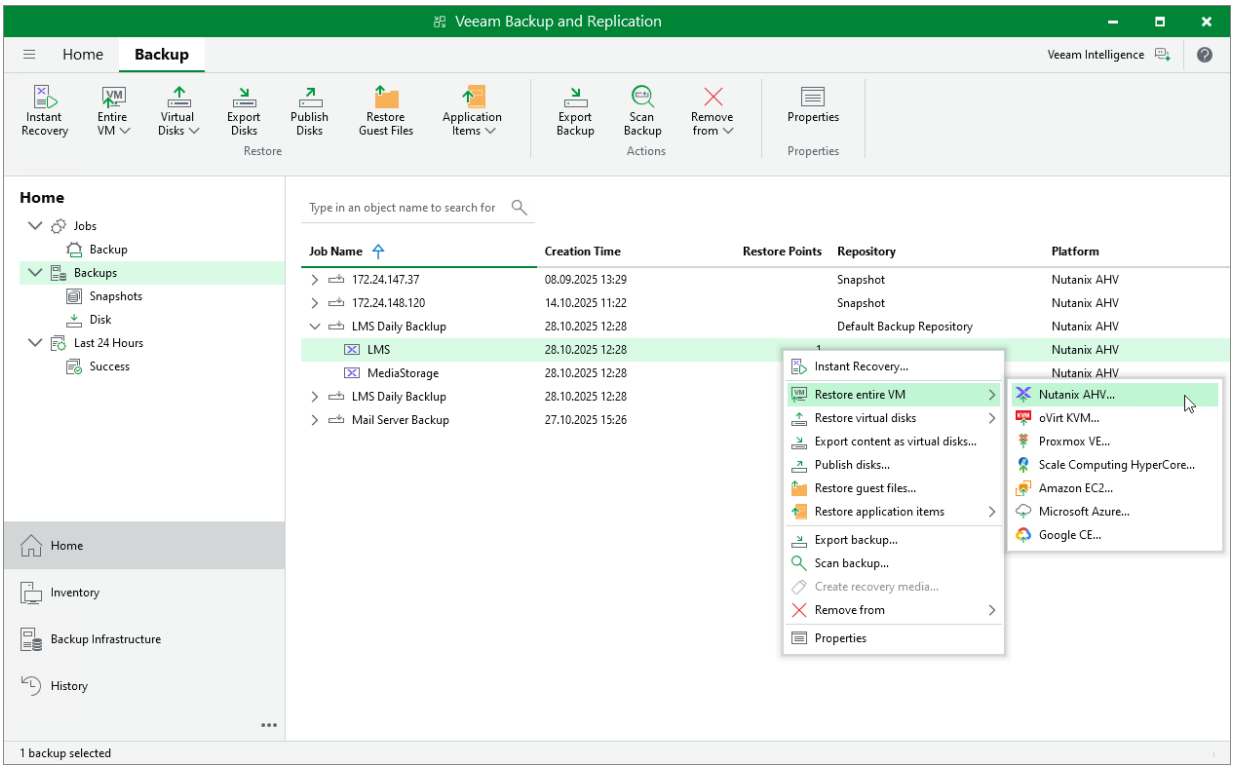

# Step 2. Select Restore Point

At the **Virtual Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Plug-in 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 Plug-in 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.

> **TIP**
>
> 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.

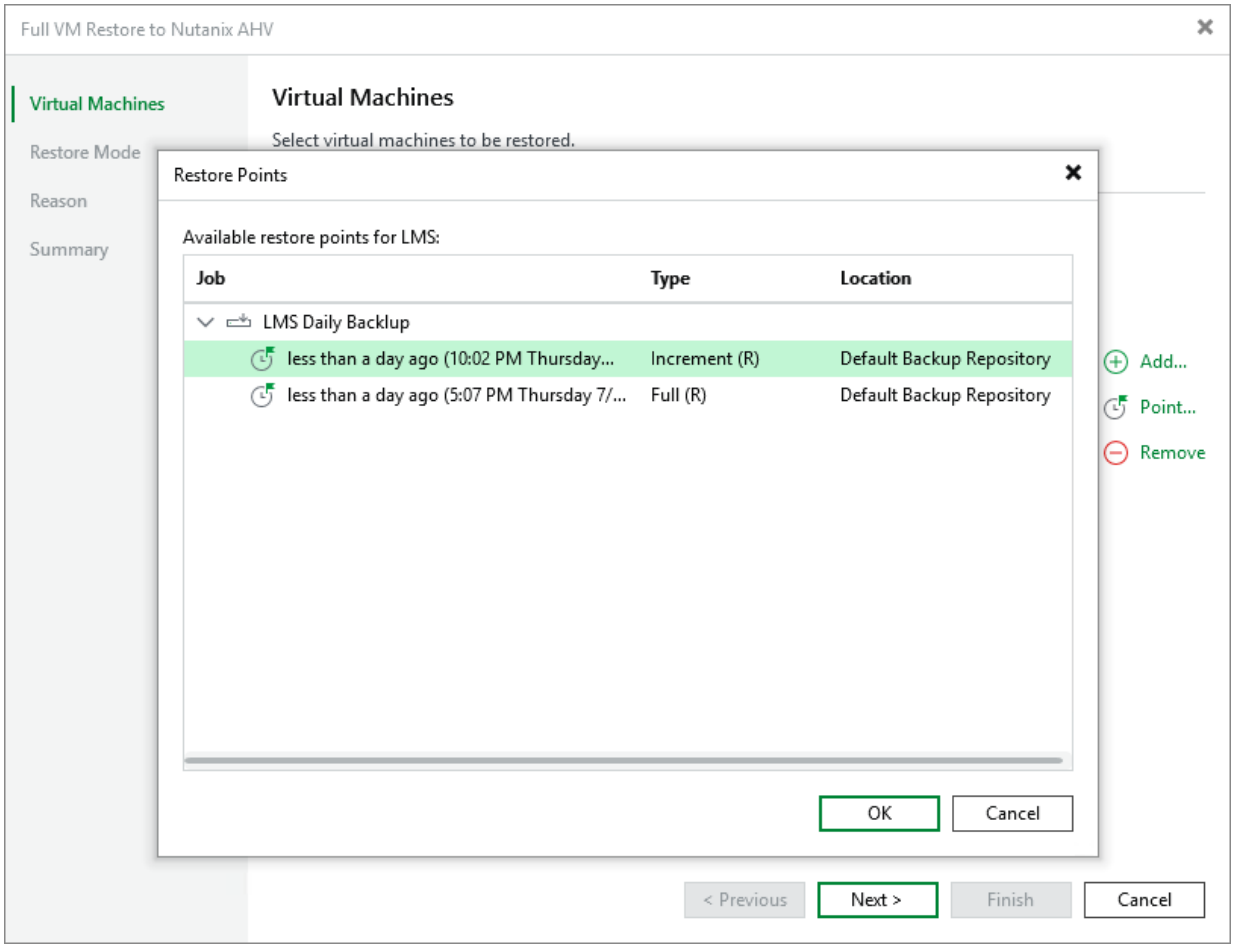

# Step 3. Choose Restore Mode

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.

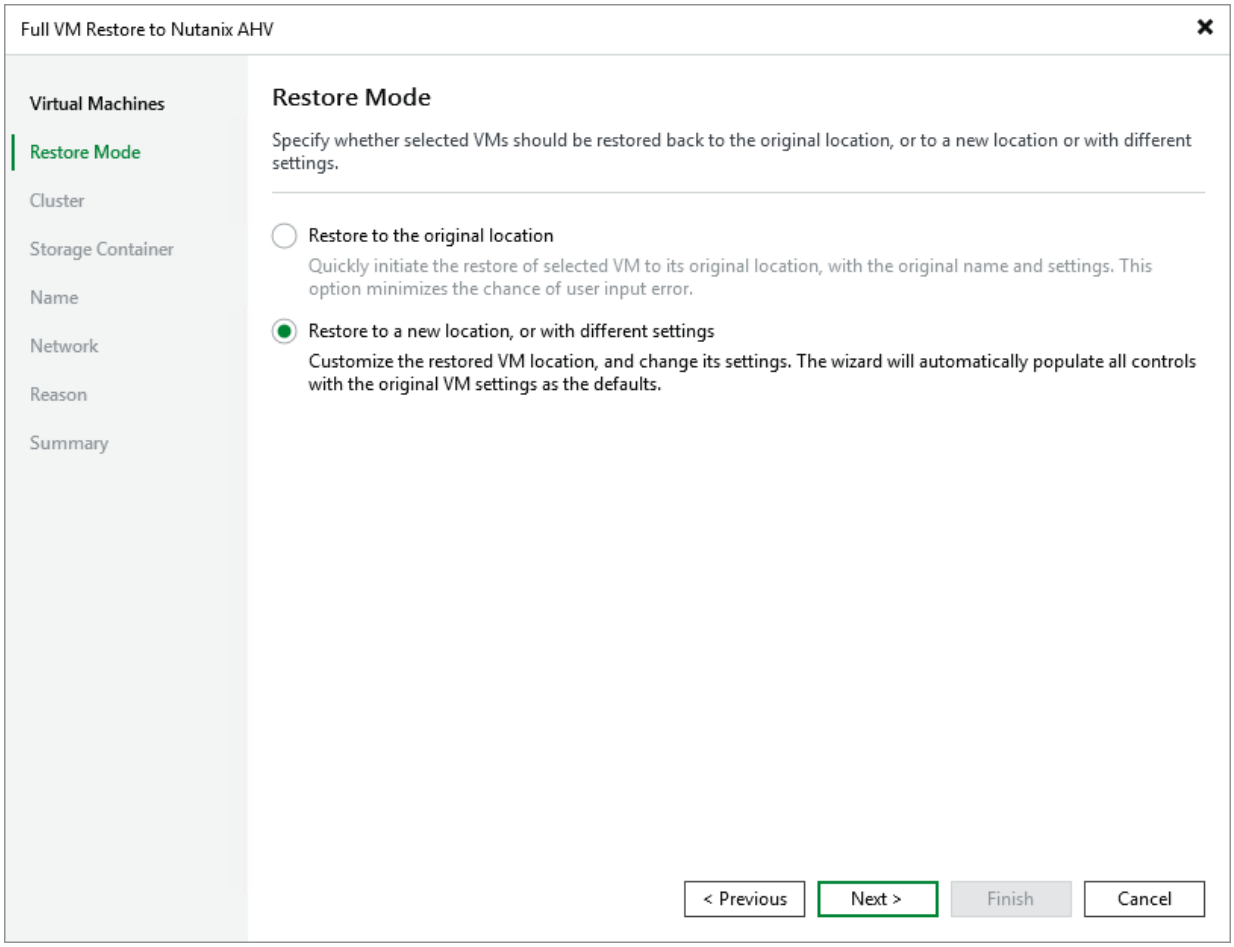

# Step 4. Specify Target Cluster

[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 to Backup Infrastructure.

> **NOTE**
>
> The **Cluster** step of the **Full VM Restore to Nutanix AHV** wizard is only available when you restore the VM from a backup.

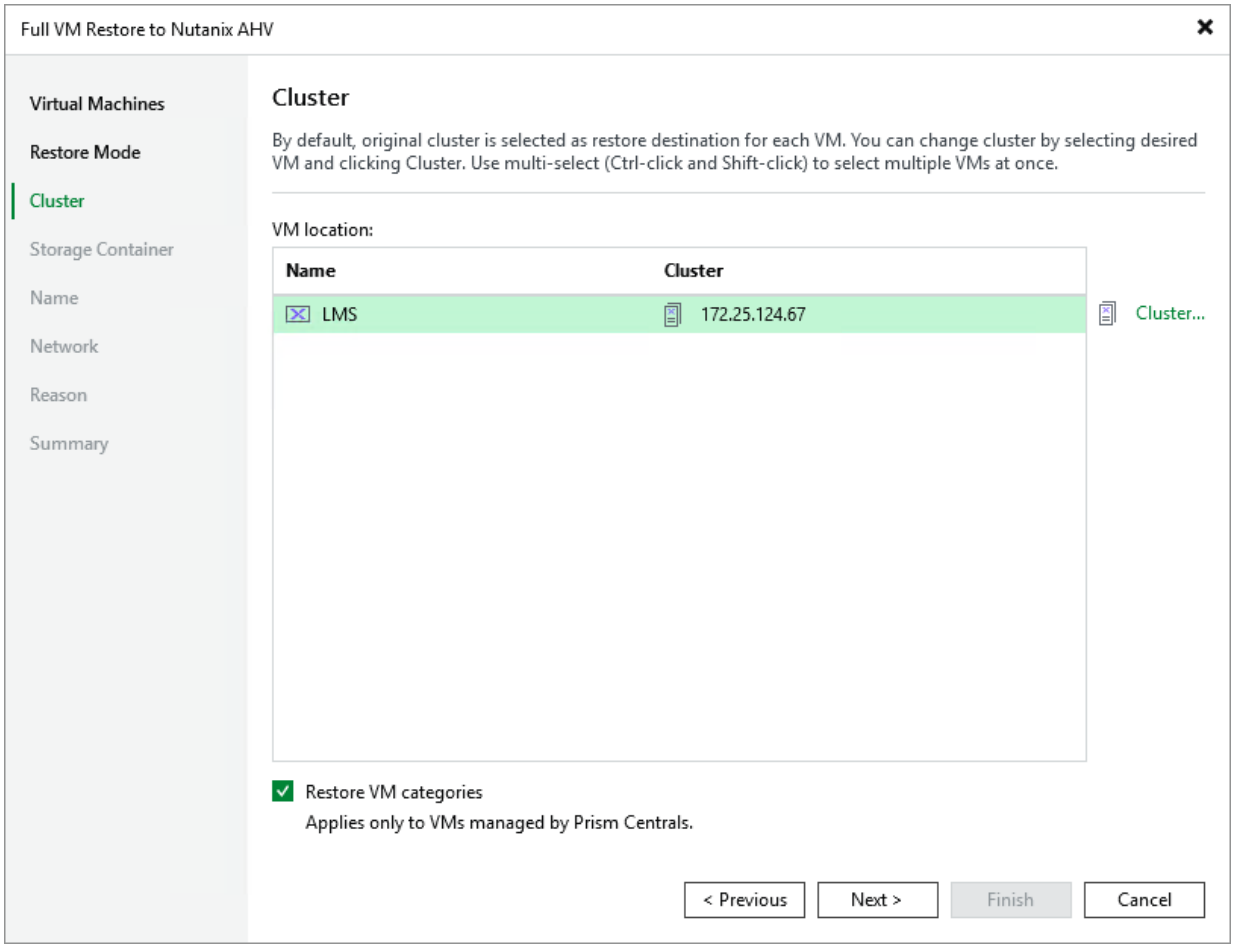

# Step 5. Select Storage Container

[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.

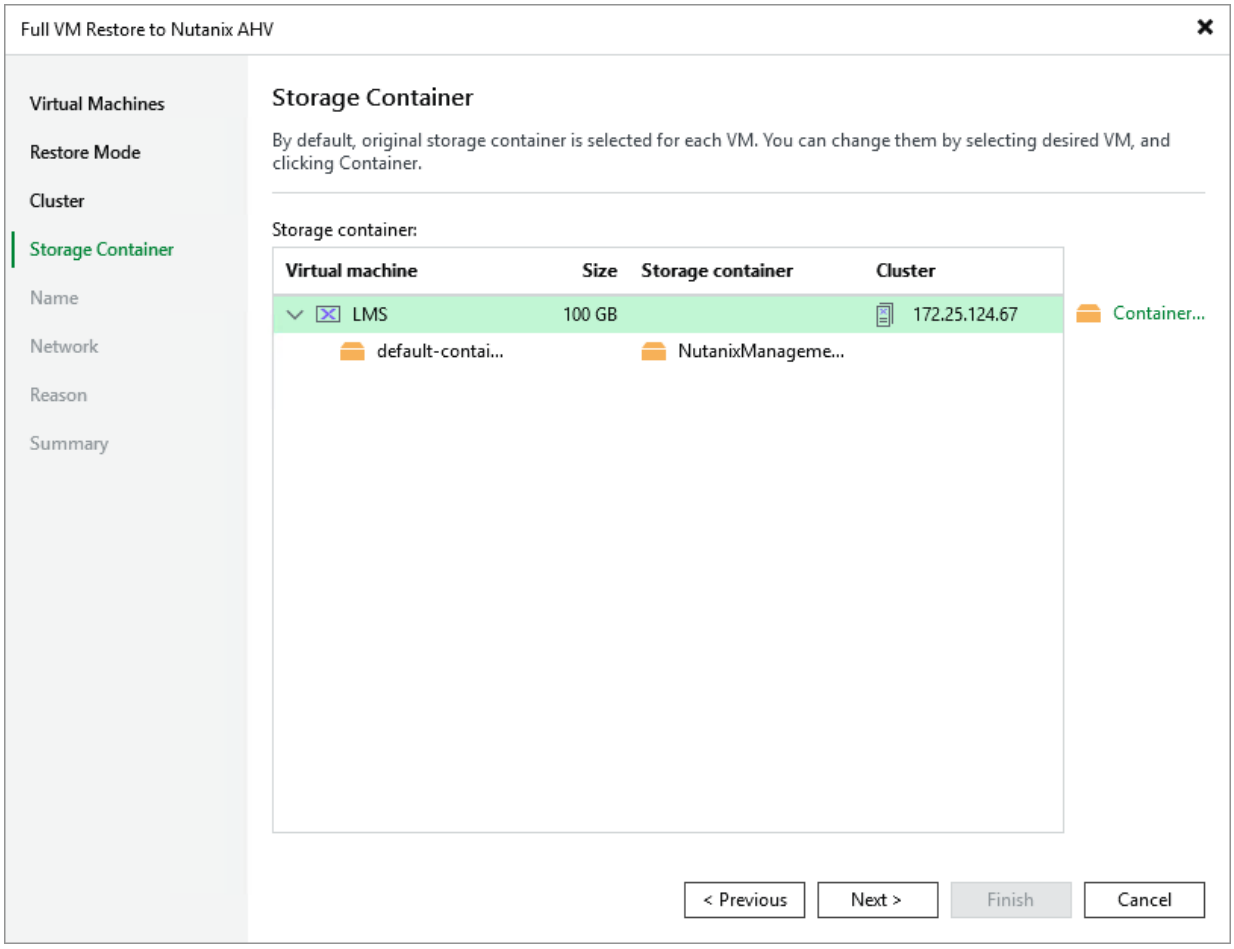

# Step 6. Specify VM Name

[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.

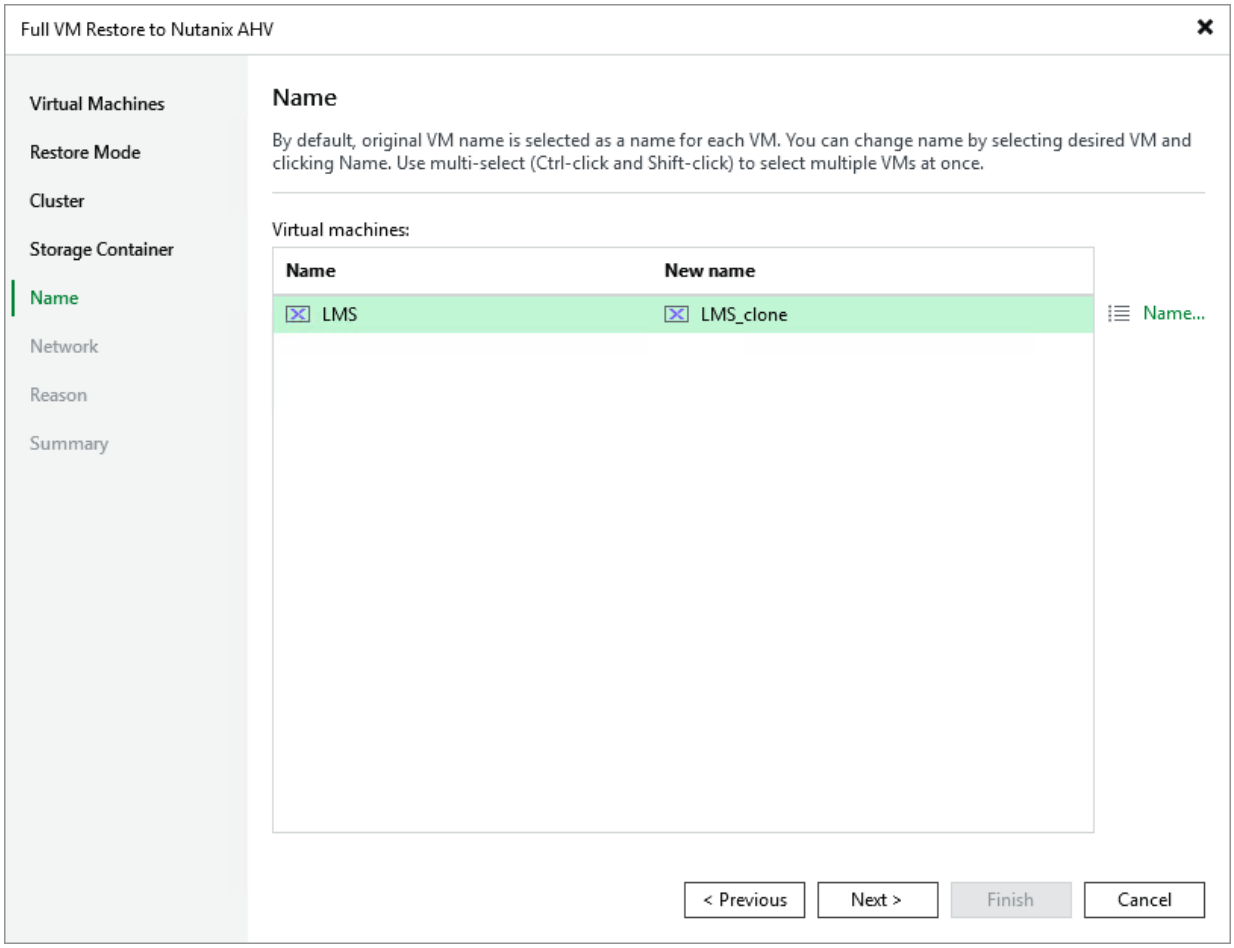

# Step 7. Configure Network Settings

[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, when restoring the VM from a user snapshot, you can choose to disconnect the original network.

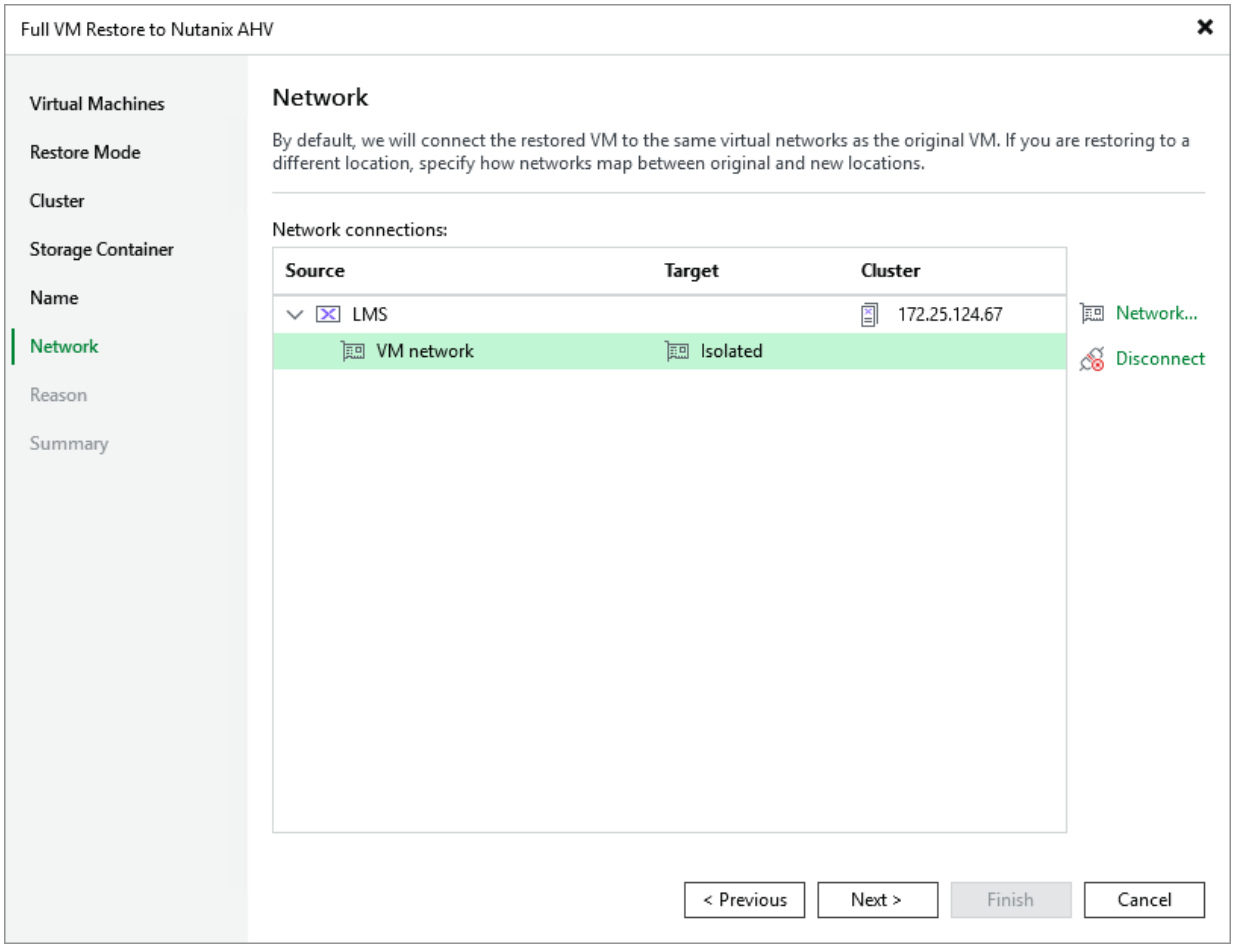

# Step 8. Specify Restore Reason

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.

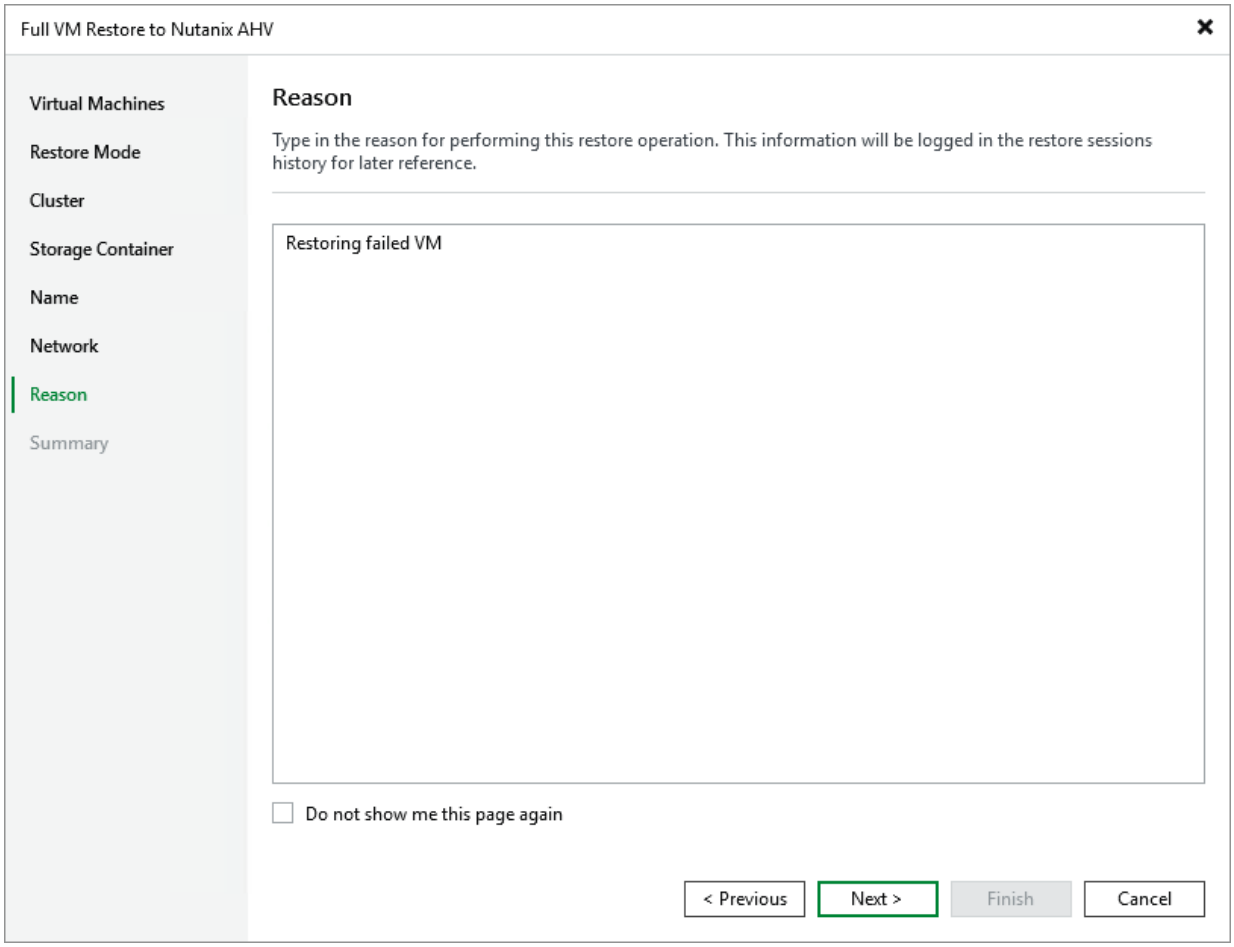

# Step 9. Finish Working with Wizard

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.

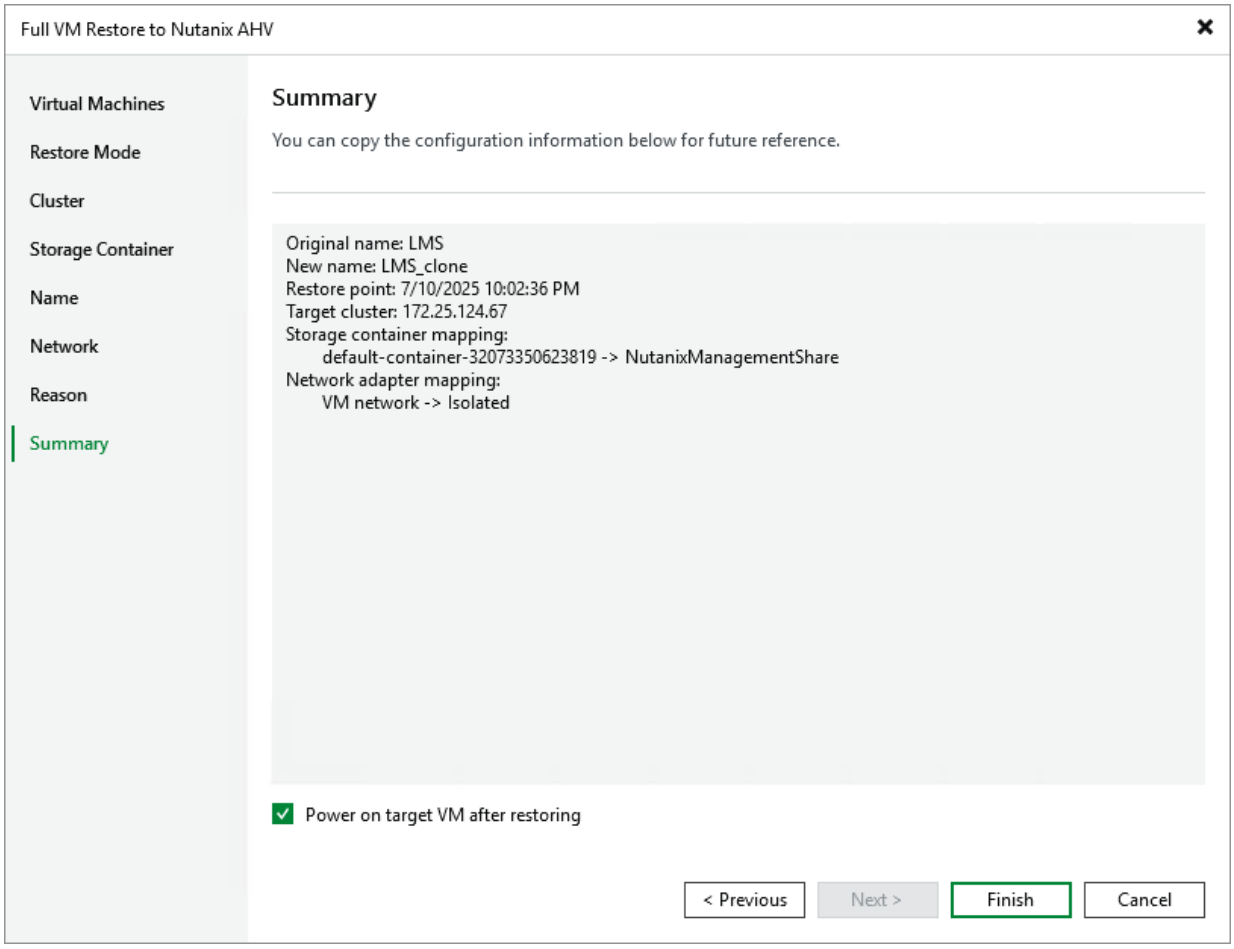

# Performing Disk Restore

In case a disaster strikes, you can restore disks of a Nutanix AHV VM from a backup or backup snapshot. Veeam Plug-in 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.

To restore disks attached to a protected VM, do the following:

1. Launch the Virtual Disk Restore wizard.
2. Select a VM.
3. Select a restore point.
4. Configure mapping settings.
5. Specify a reason for the restore.
6. Finish working with the wizard.

# Step 1. Launch Virtual Disk Restore Wizard

To launch the **Virtual Disk Restore** 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 job, right-click the VM you want to restore and select **Restore virtual disks** > **Nutanix AHV**.

   Alternatively, expand the necessary backup job, select the VM and click **Virtual Disks** > **Nutanix AHV** on the ribbon.

> **TIP**
>
> To restore a VM from a backup snapshot, expand the job that contains the snapshot of the VM, right-click the VM and select **Restore entire VM to Nutanix AHV**.

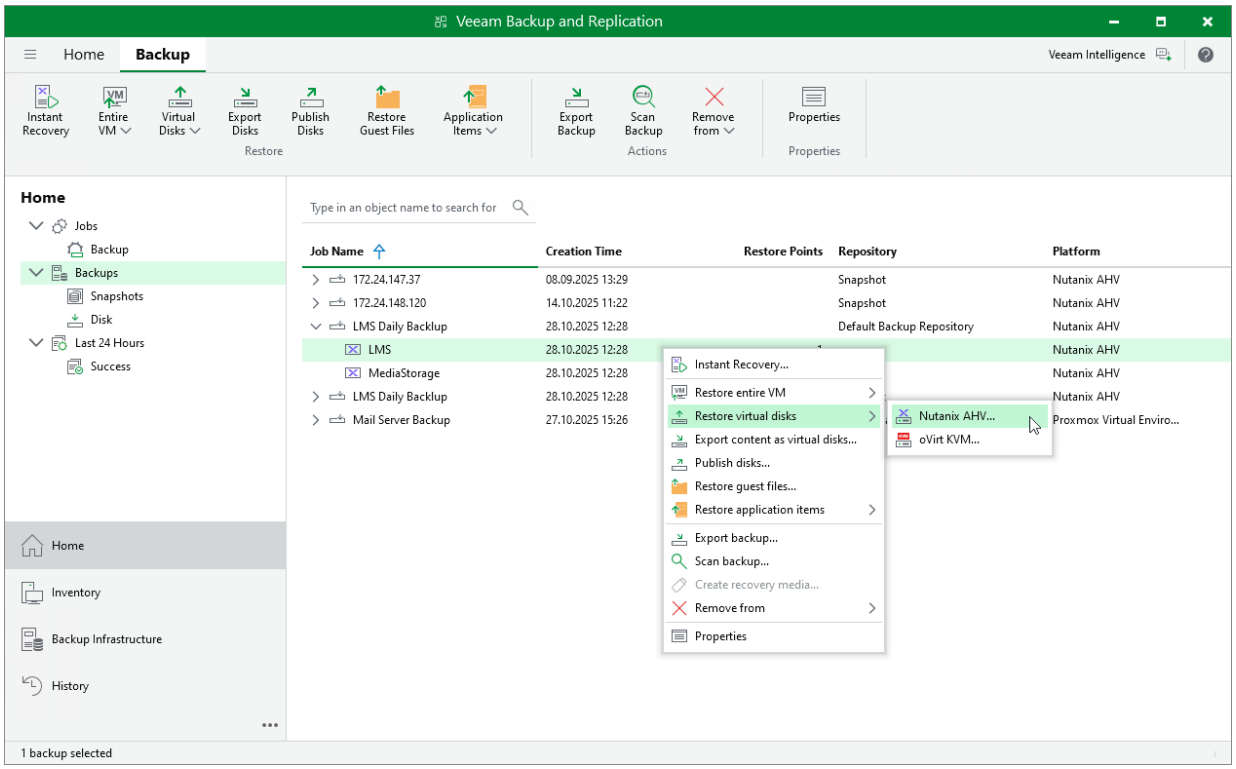

# Step 2. Select Virtual Machine

At the **Virtual Machine** step of the wizard, expand the backup job tree and select the VM whose virtual disks you want to restore.

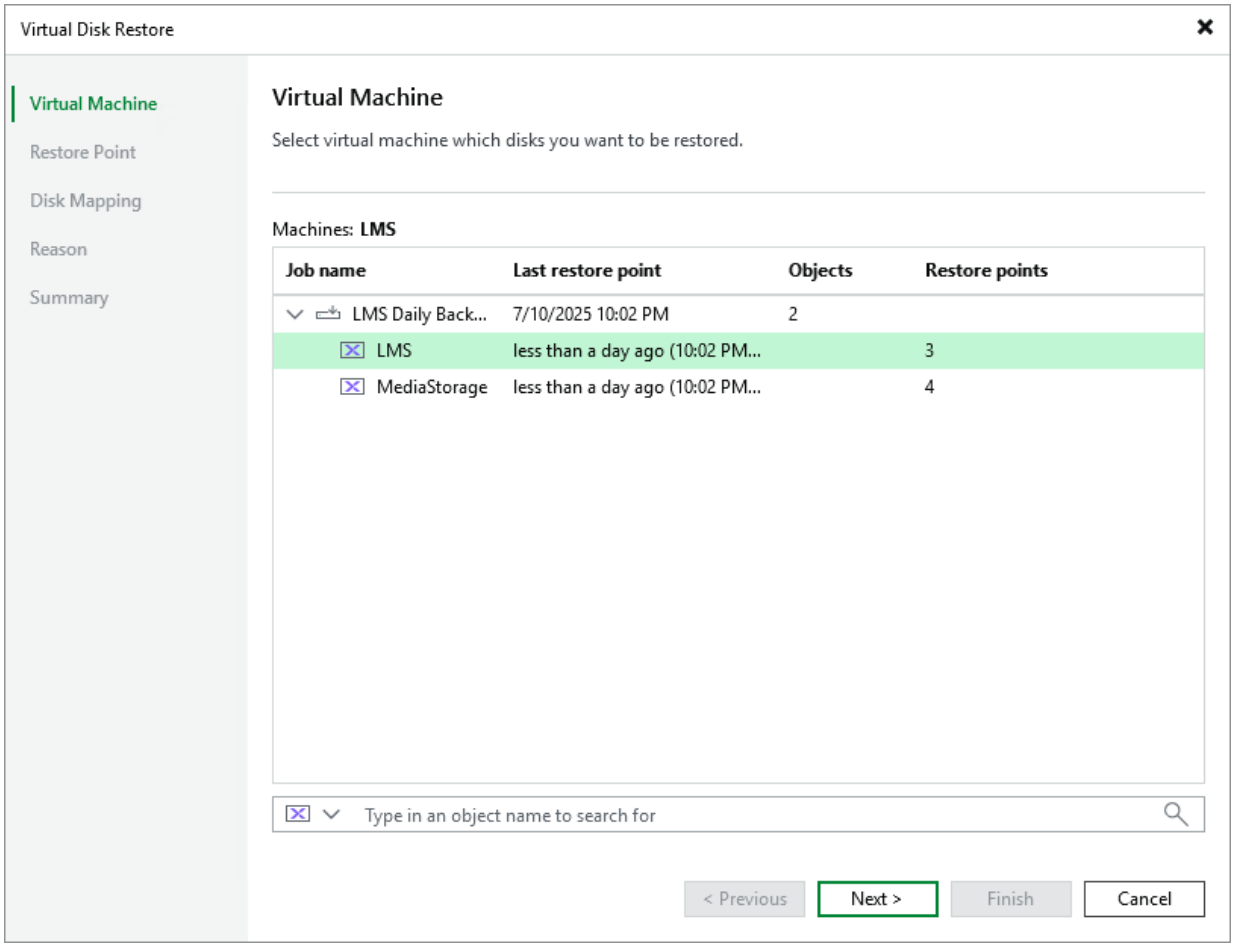

# Step 3. Select Restore Point

At the **Restore Point** step of the wizard, select a restore point that will be used to restore data. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the data to an earlier state.

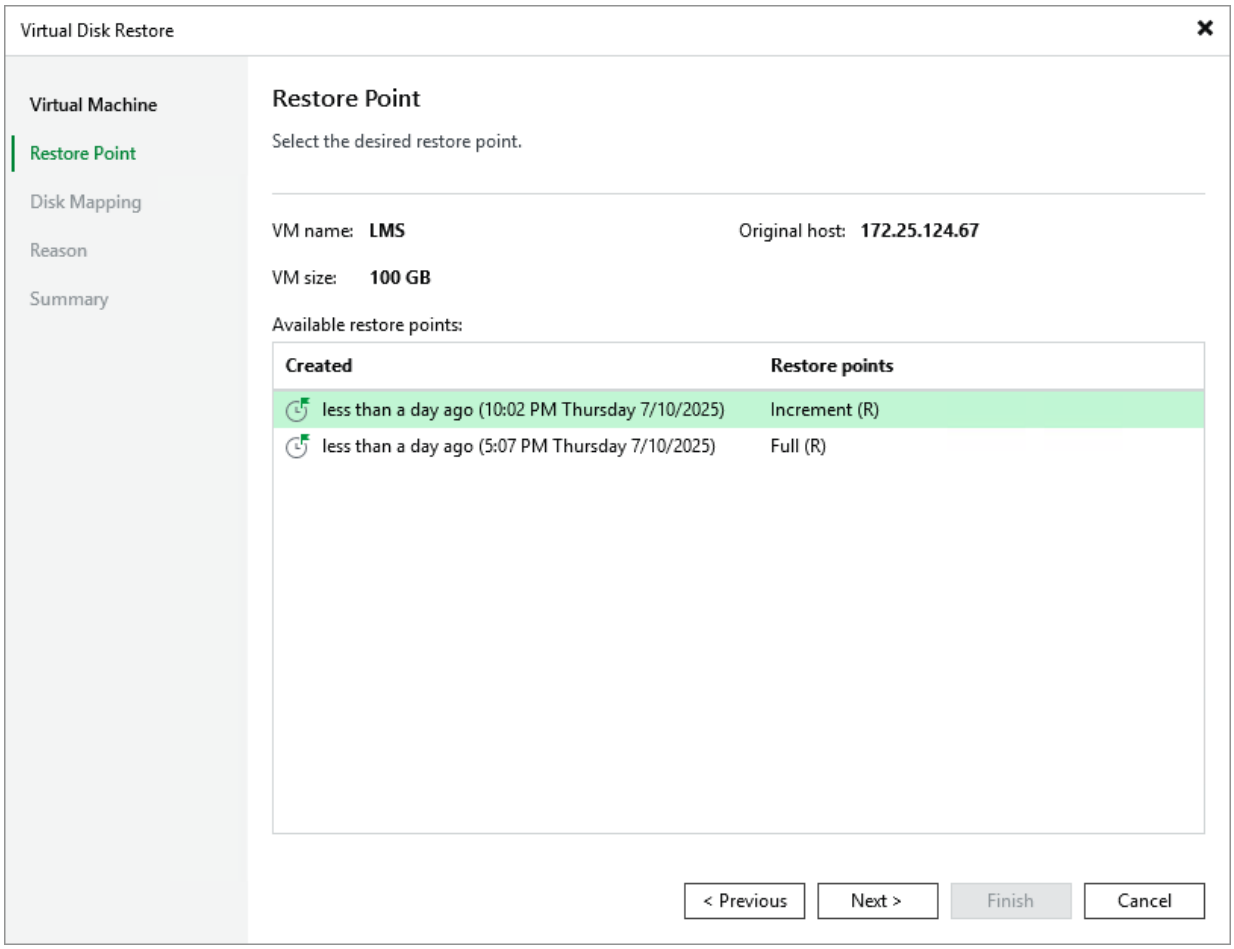

# Step 4. Configure Mapping Settings

At the **Disk Mapping** step of the wizard, do the following:

1. Choose a target VM to which you want to attach the restored disks.

   By default, Veeam Backup & Replication attaches the restored disks to the original VM. To attach the disks to another VM, click **Choose**.

> **IMPORTANT**
>
> During disk restore, Veeam Backup & Replication turns off the target VM to reconfigure its settings and attach the restored disks. It is recommended that you stop all activities on the target VM till the restore session completes.

2. Select virtual disks to restore.

   By default, Veeam Backup & Replication attaches the restored disks to the target VM as new disks. If you want the restored disks to replace the existing disks, or if you want to change the disk bus type and to specify a storage domain for the restored disks, click **Change**.

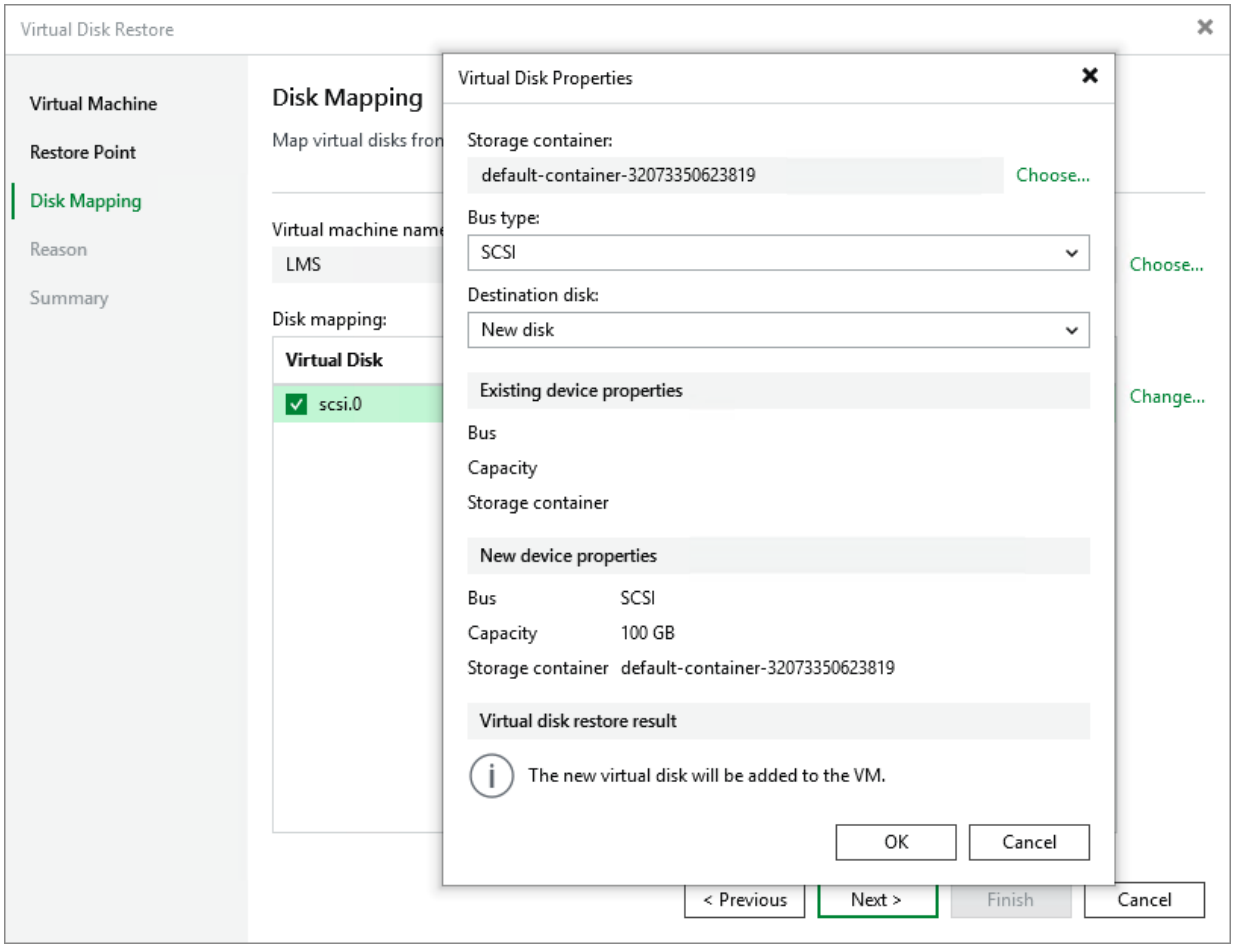

# Step 5. Specify Reason for Restore

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.

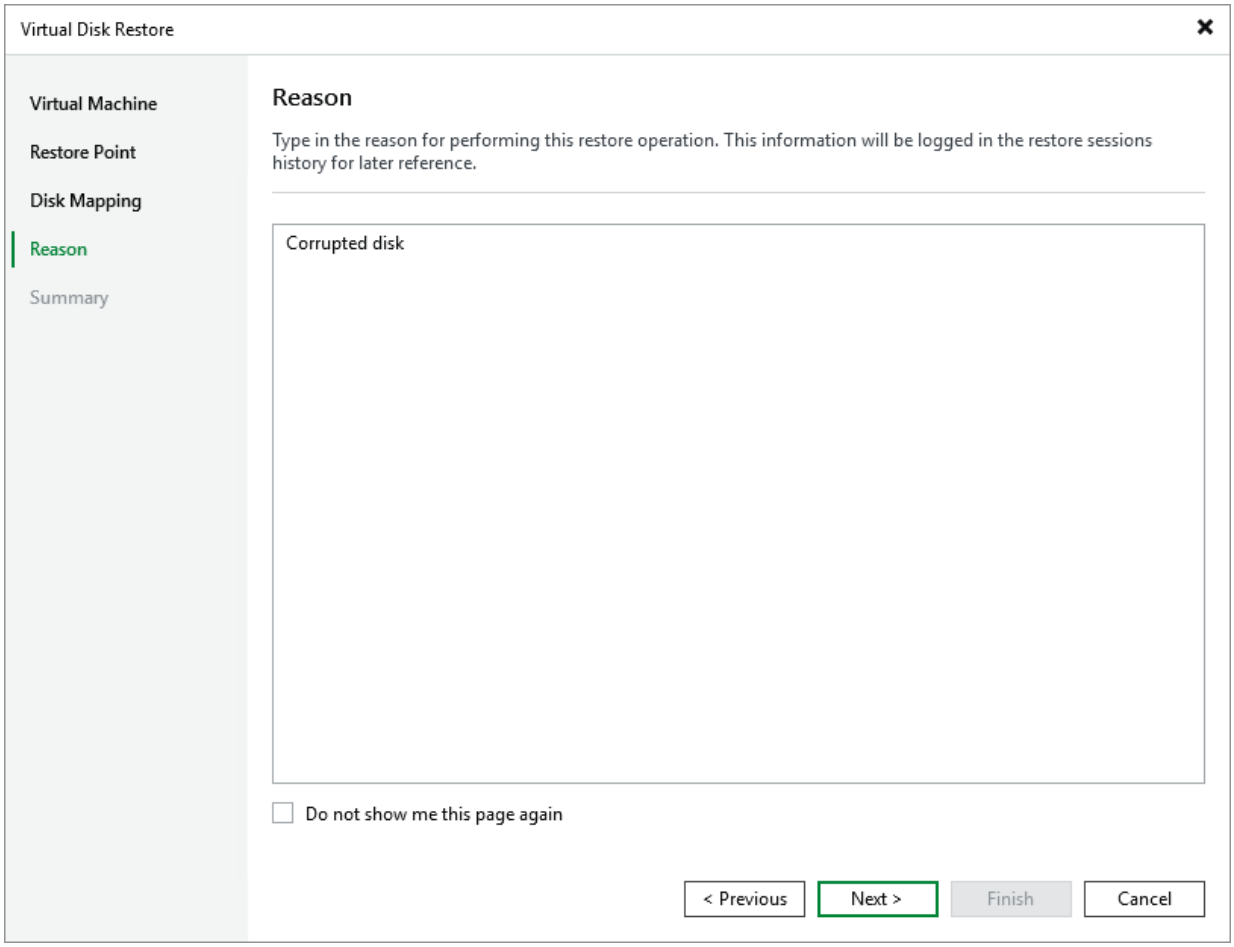

# Step 6. Finish Working with Wizard

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.

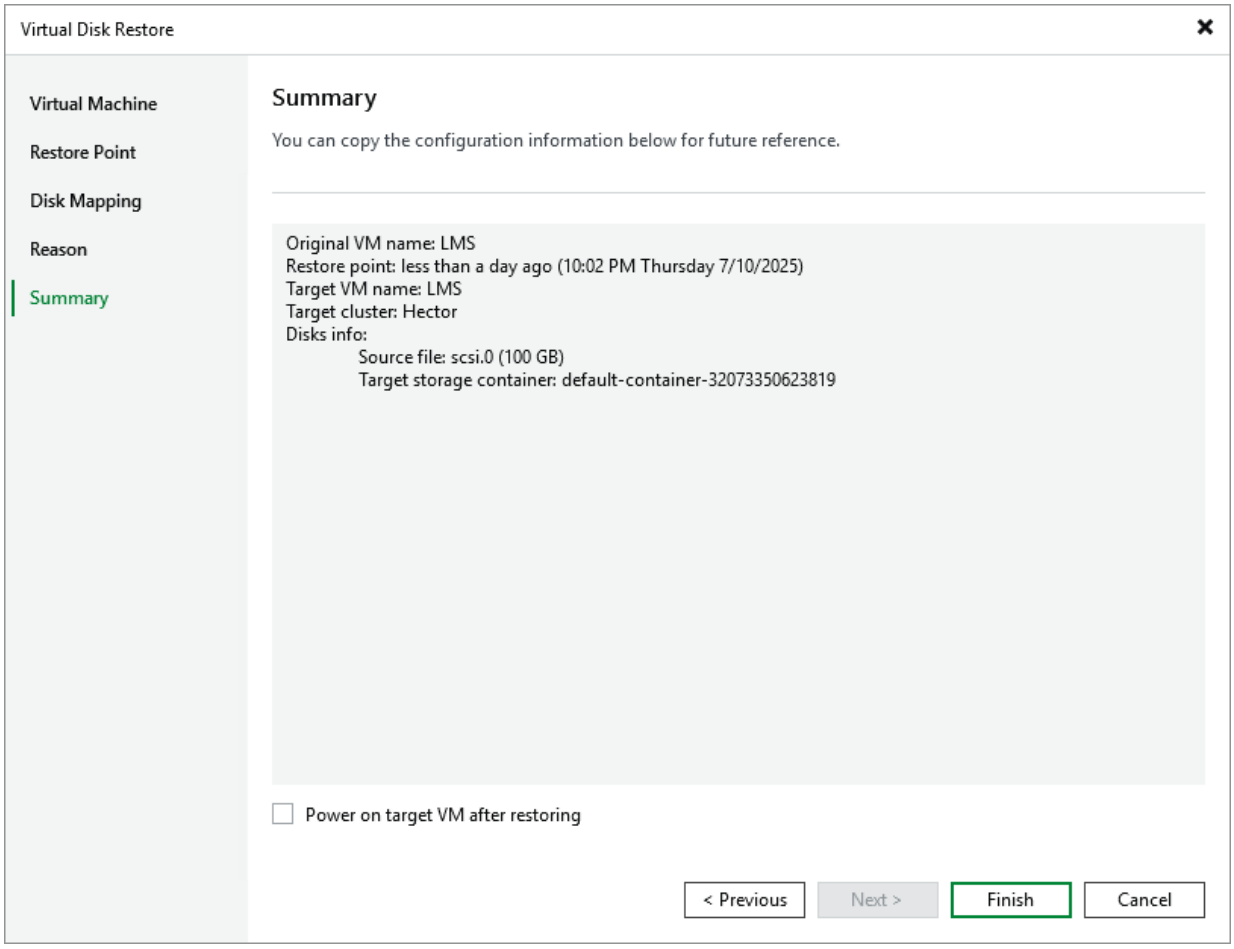

# Instant Recovery

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 Plug-in 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 Plug-in for Oracle Linux Virtualization Manager and Red Hat Virtualization
- Backups of Proxmox VE VMs created by Veeam Plug-in 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 Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV] Veeam Plug-in for Nutanix AHV attaches VM disks with the restored data to the target VM disk nodes using their original bus types. Veeam Plug-in 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 Veeam Plug-in for Nutanix AHV 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 job, right-click the VM you want to restore and select **Instant recovery**.

   Alternatively, expand the necessary backup job, select the VM and click **Instant Recovery** on the ribbon.

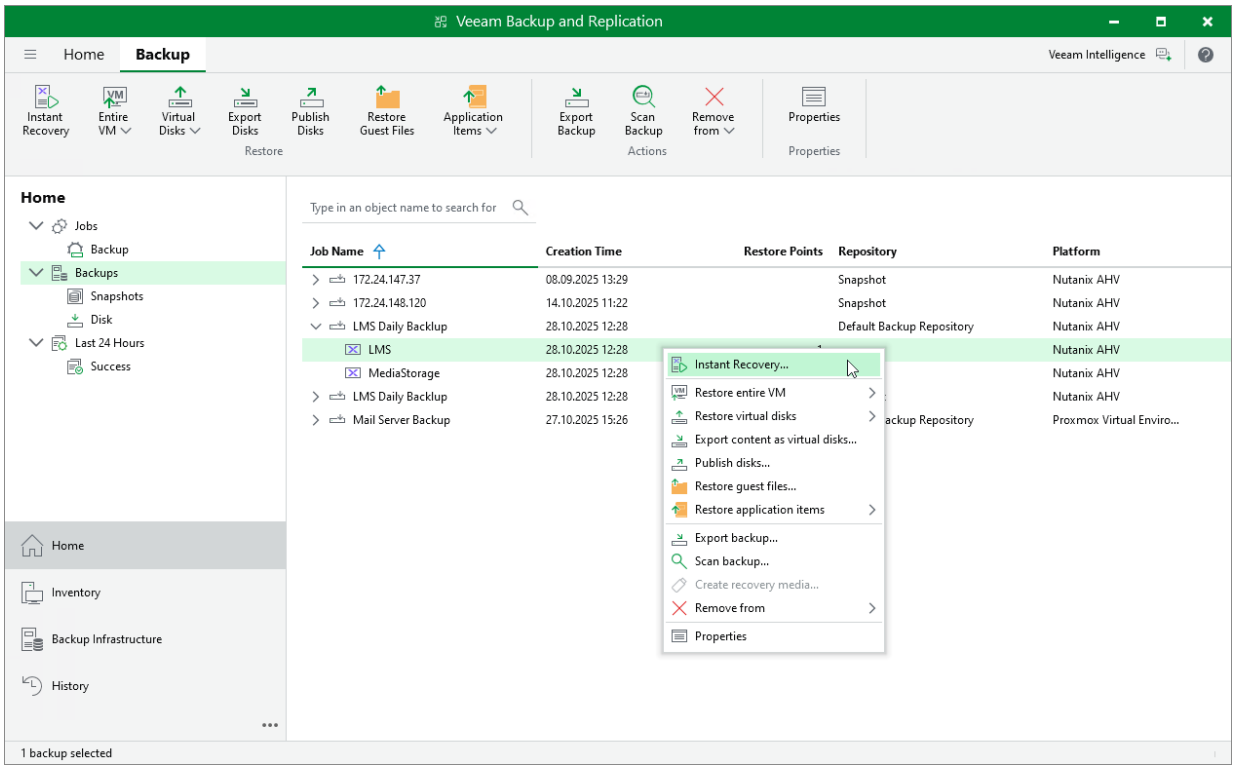

# Step 2. Select Restore Point

At the **Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Plug-in 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 Plug-in 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.

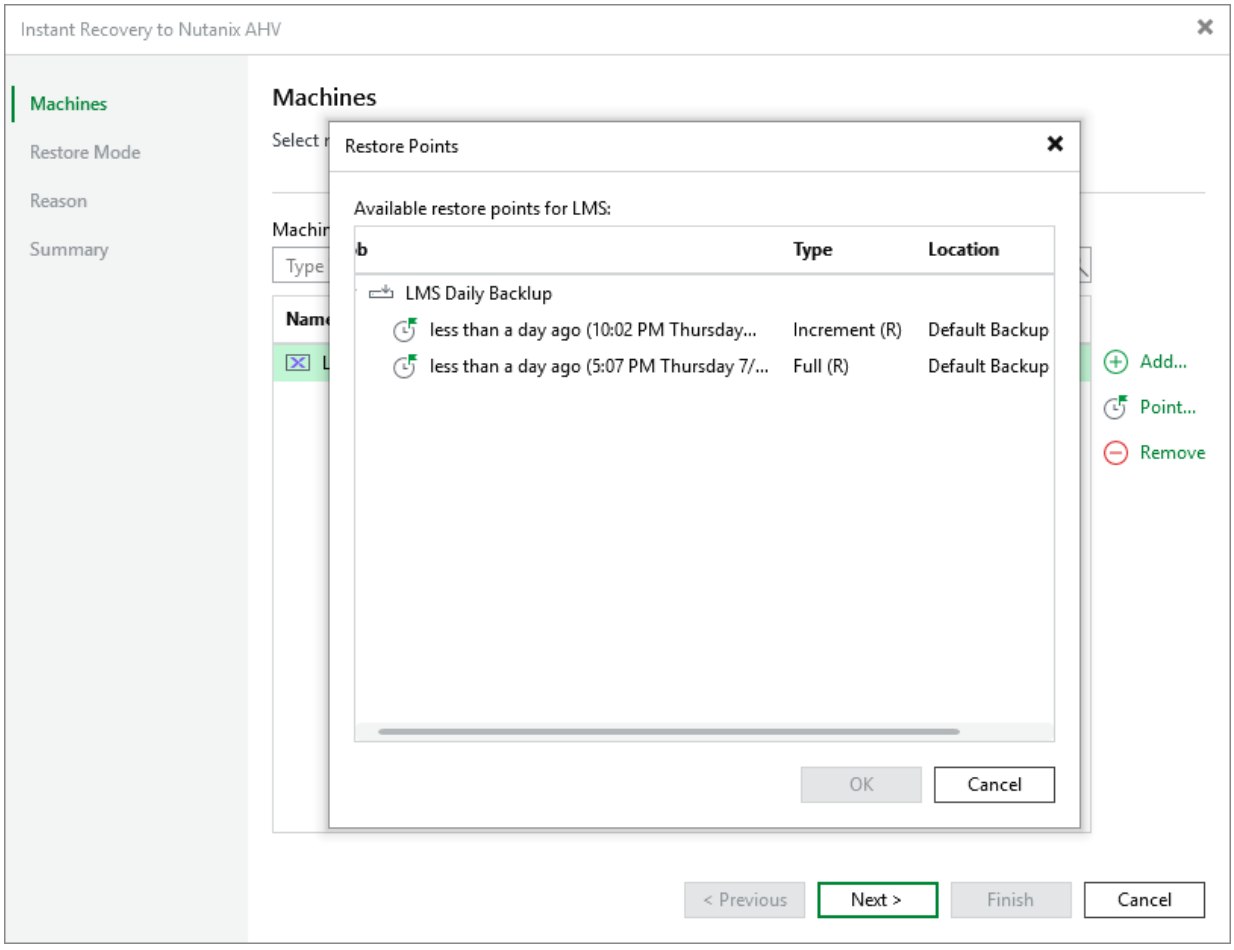

# Step 3. Choose Restore Mode

[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 Plug-in 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.

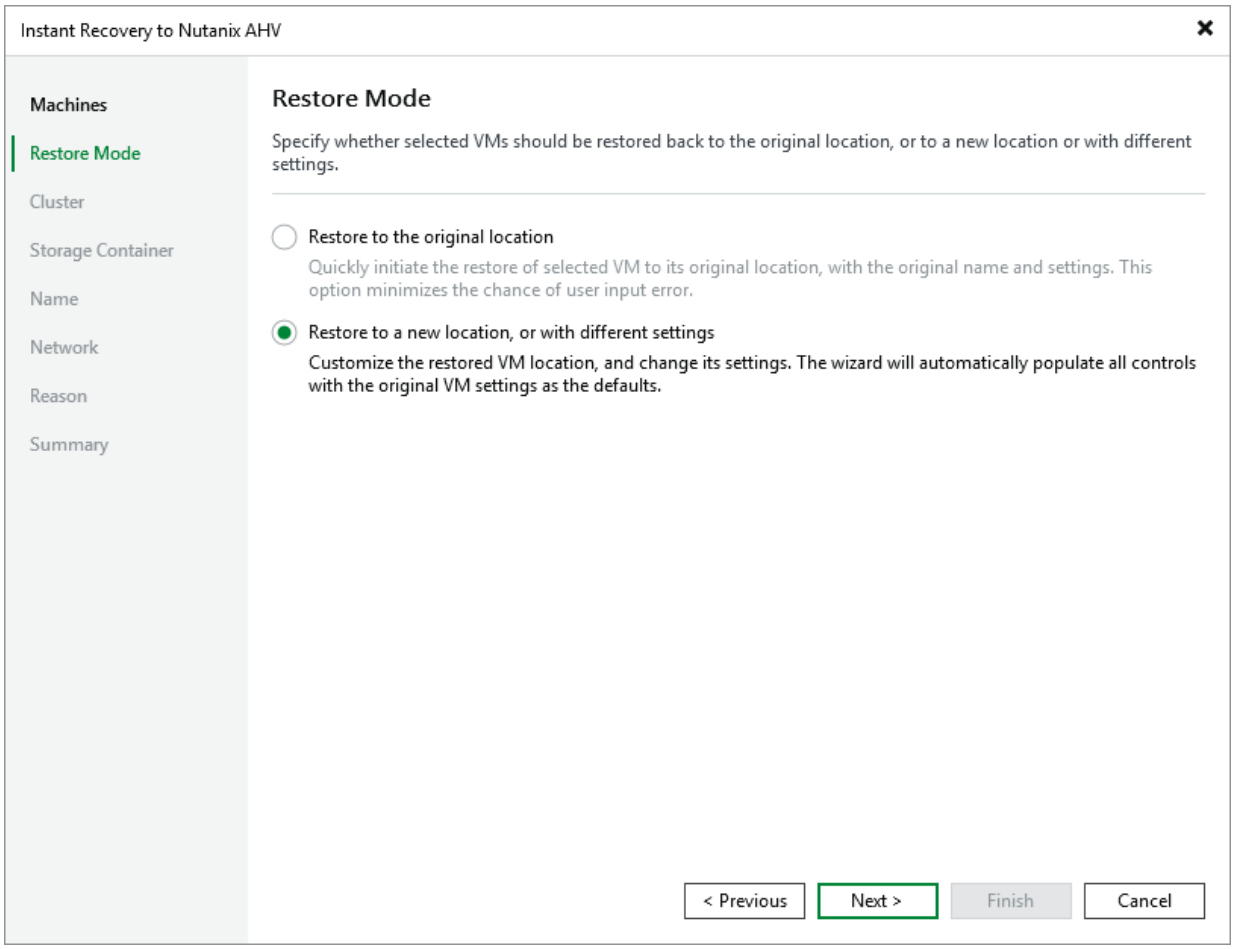

# Step 4. Specify Target Cluster

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 to Backup Infrastructure.

> **IMPORTANT**
>
> If a selected VM has an attached volume group, the disks of the volume group will not be restored.

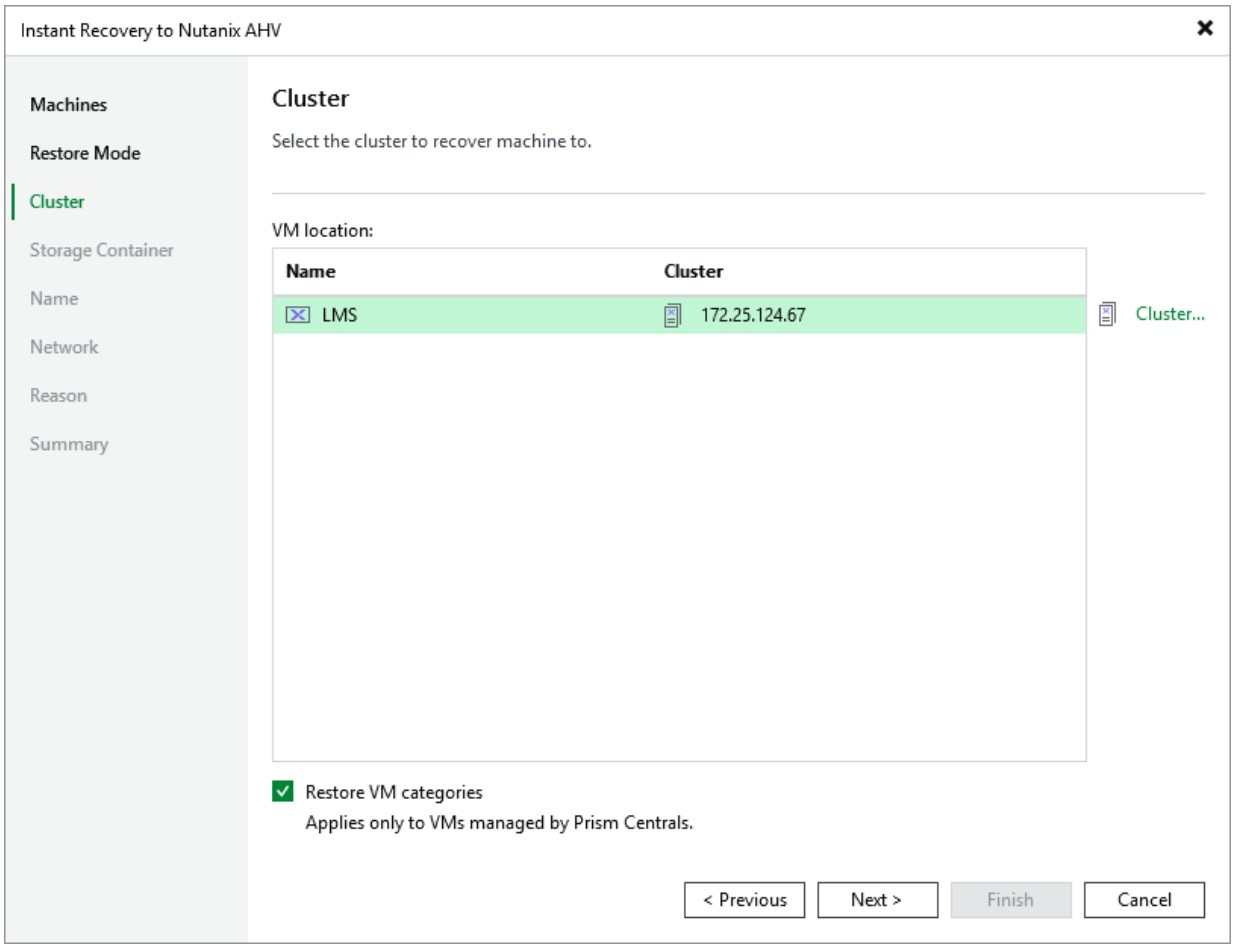

# Step 5. Select Storage Container

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.

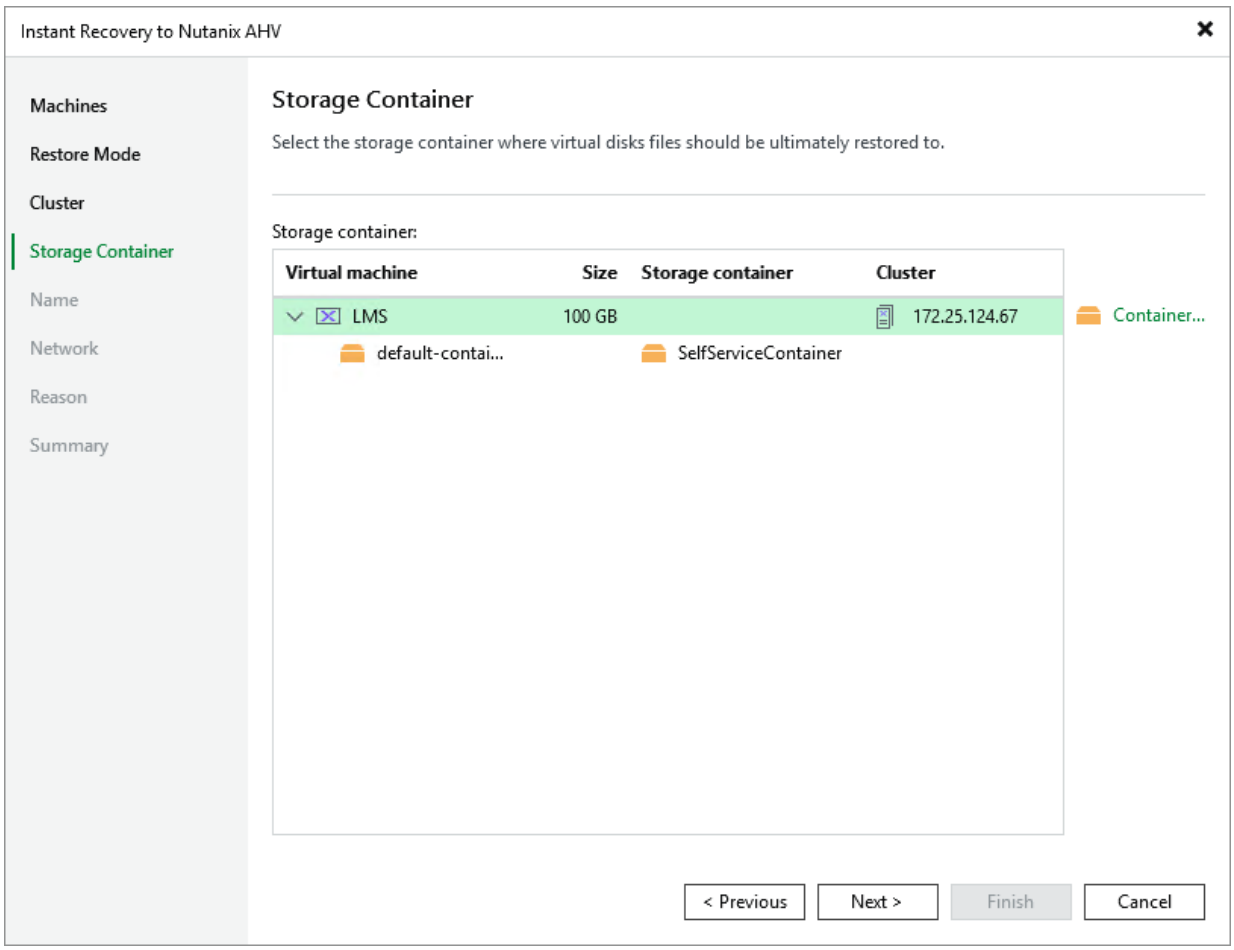

# Step 6. Specify VM Name

At the **Name** step of the wizard, you can specify a new name for the recovered VM.

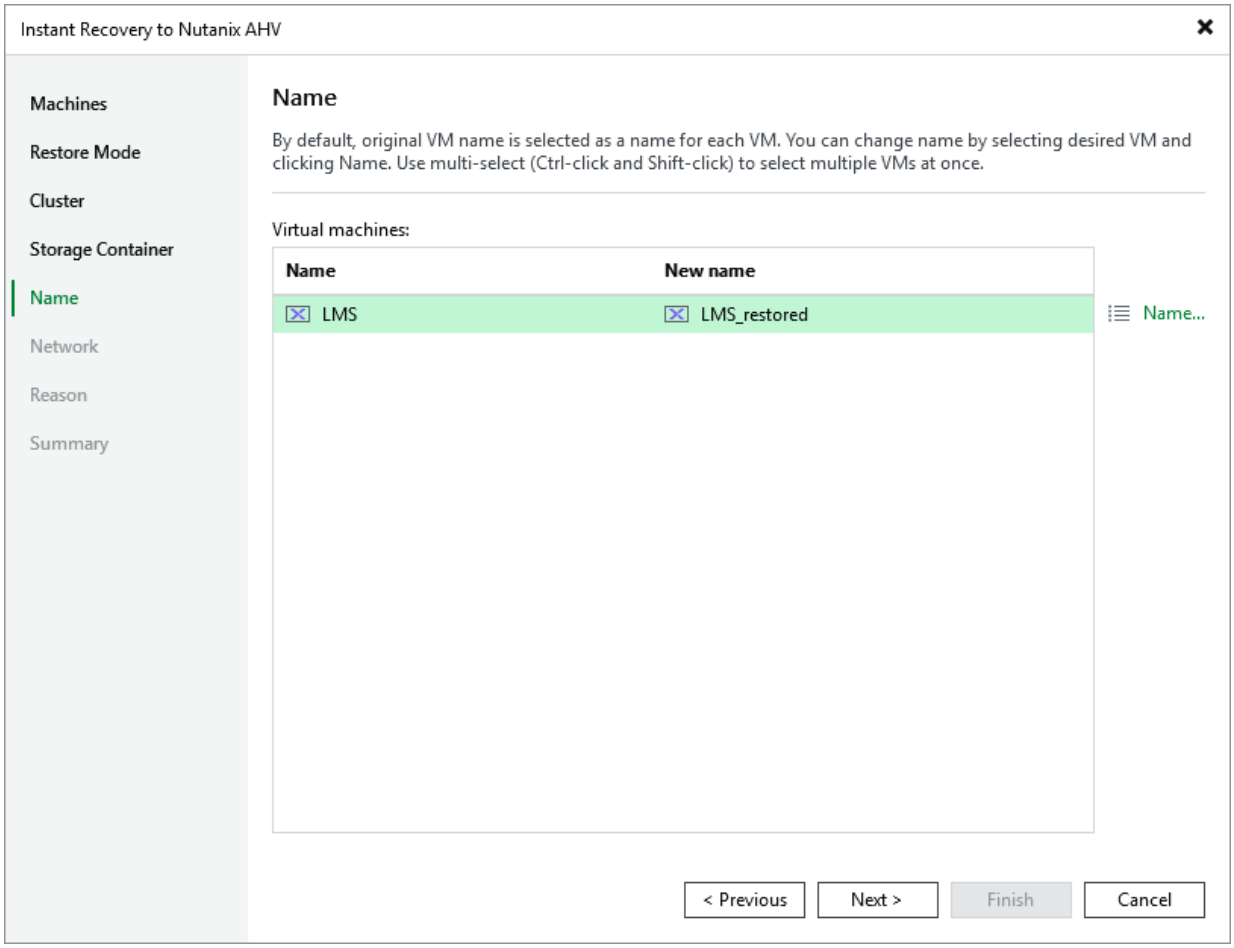

# Step 7. Configure Network Settings

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.

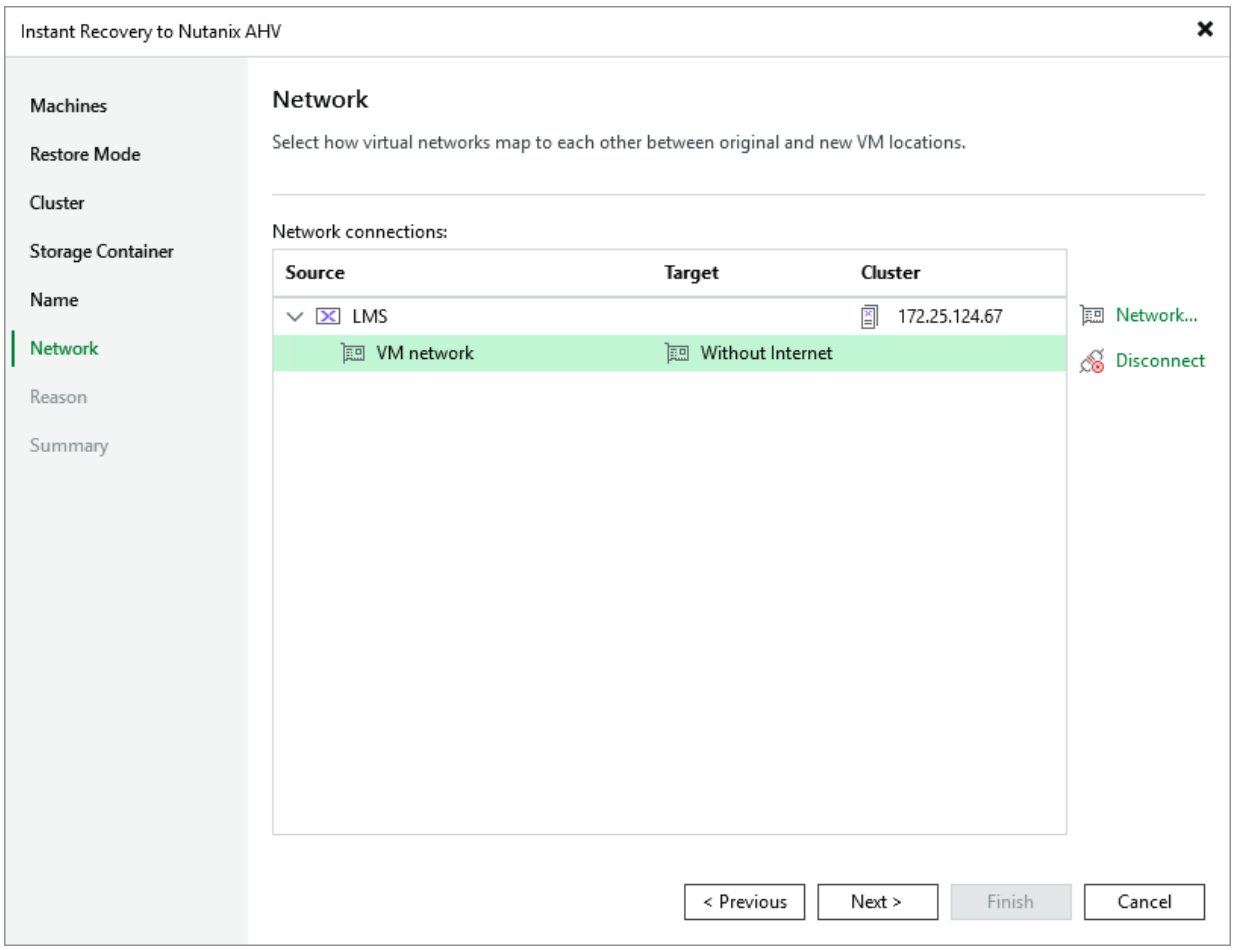

# Step 8. Specify Restore Reason

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.

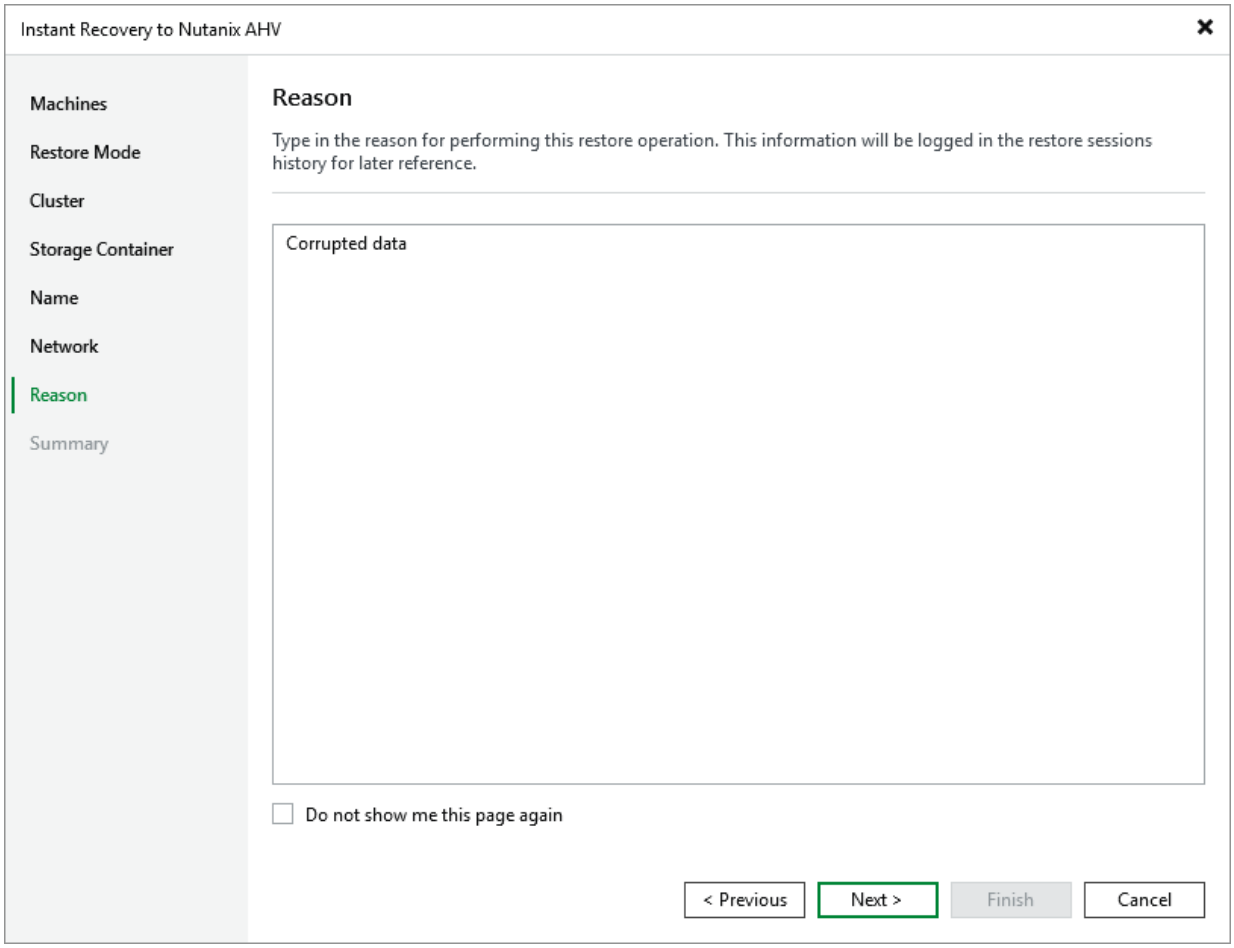

# Step 9. Review Configured Settings

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.

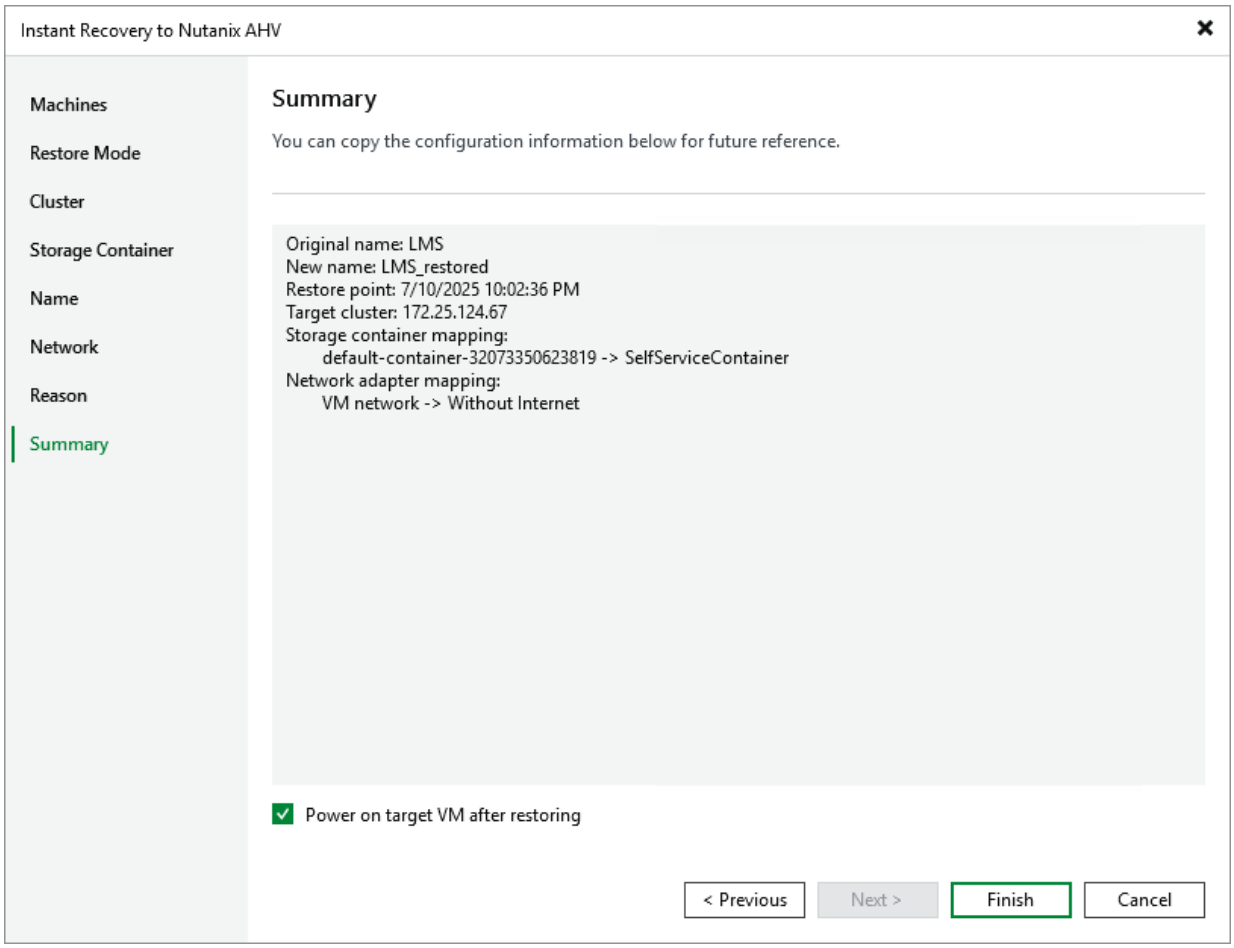

# Step 10. Finalize Instant Recovery

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 Plug-in 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.

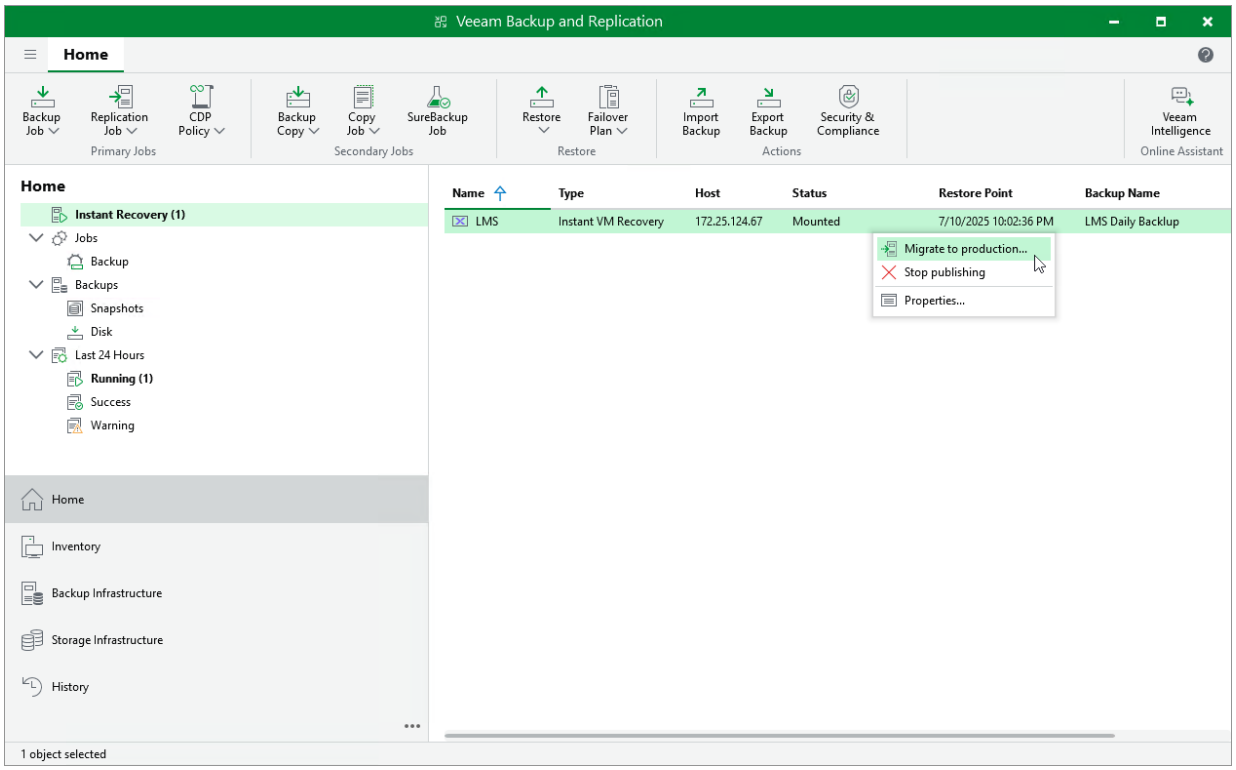

# Performing Instant Recovery of Workloads to VMware vSphere

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, expand the necessary backup job, right-click the VM you want to restore and select **Instant recovery**.

   Alternatively, expand the necessary backup job, select the VM and click **Instant Recovery** on the ribbon.
4. Complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide, section Performing Instant VM Recovery of Workloads to VMware vSphere VMs.

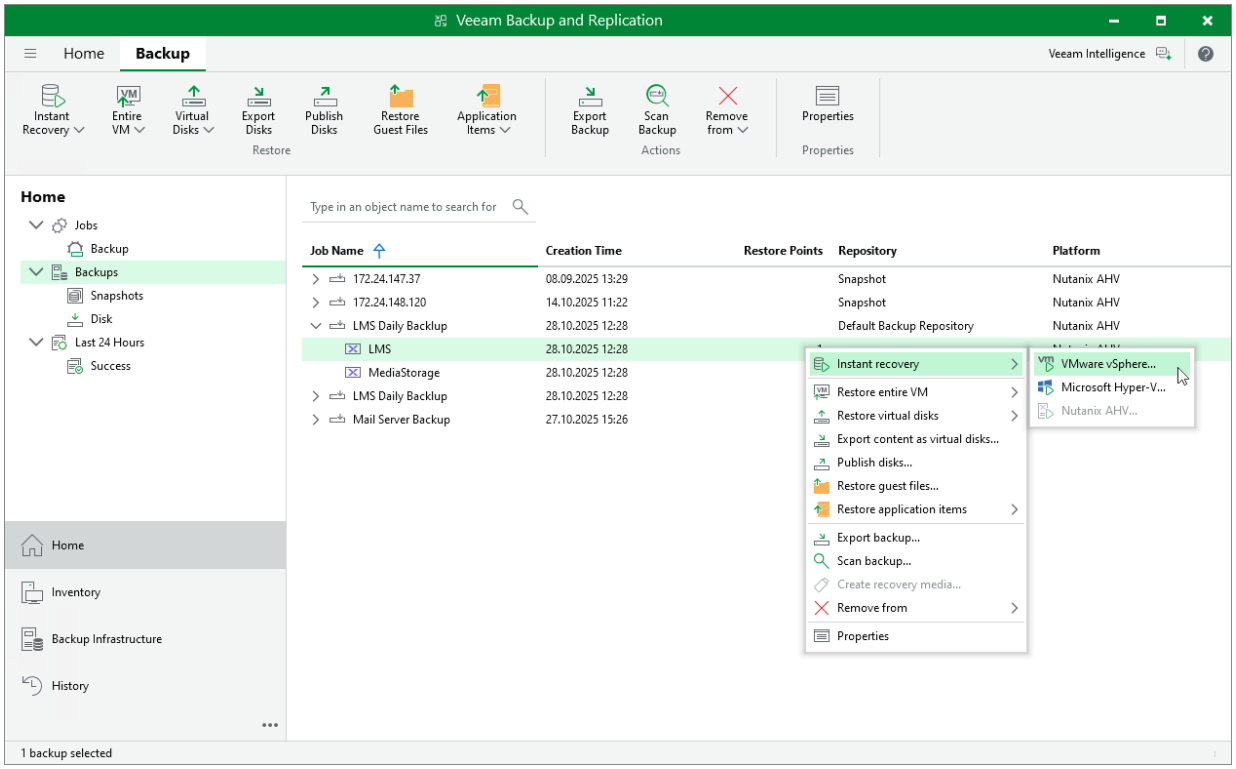

# Performing Instant Recovery of Workloads to Hyper-V

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, expand the necessary backup job, right-click the VM you want to restore and select **Instant recovery**.

   Alternatively, expand the necessary backup job, select the VM and click **Instant Recovery** on the ribbon.
4. Complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide, section Performing Instant VM Recovery of Workloads to Hyper-V VMs.

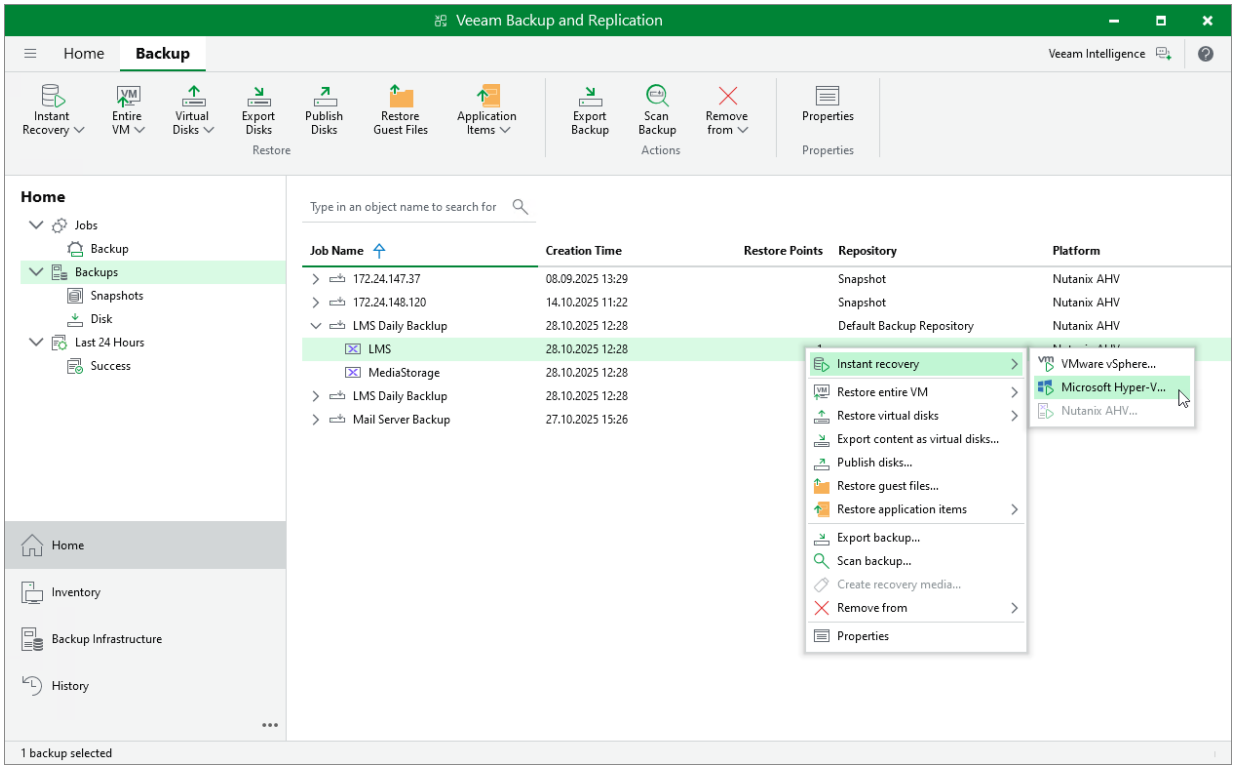

# Publishing Disks

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**.

   Alternatively, expand the necessary backup job, select the VM and click **Publish disks** on the ribbon.
4. Complete the **Publish Disk** wizard as described in the Veeam Backup & Replication User Guide, section Publishing Disks.

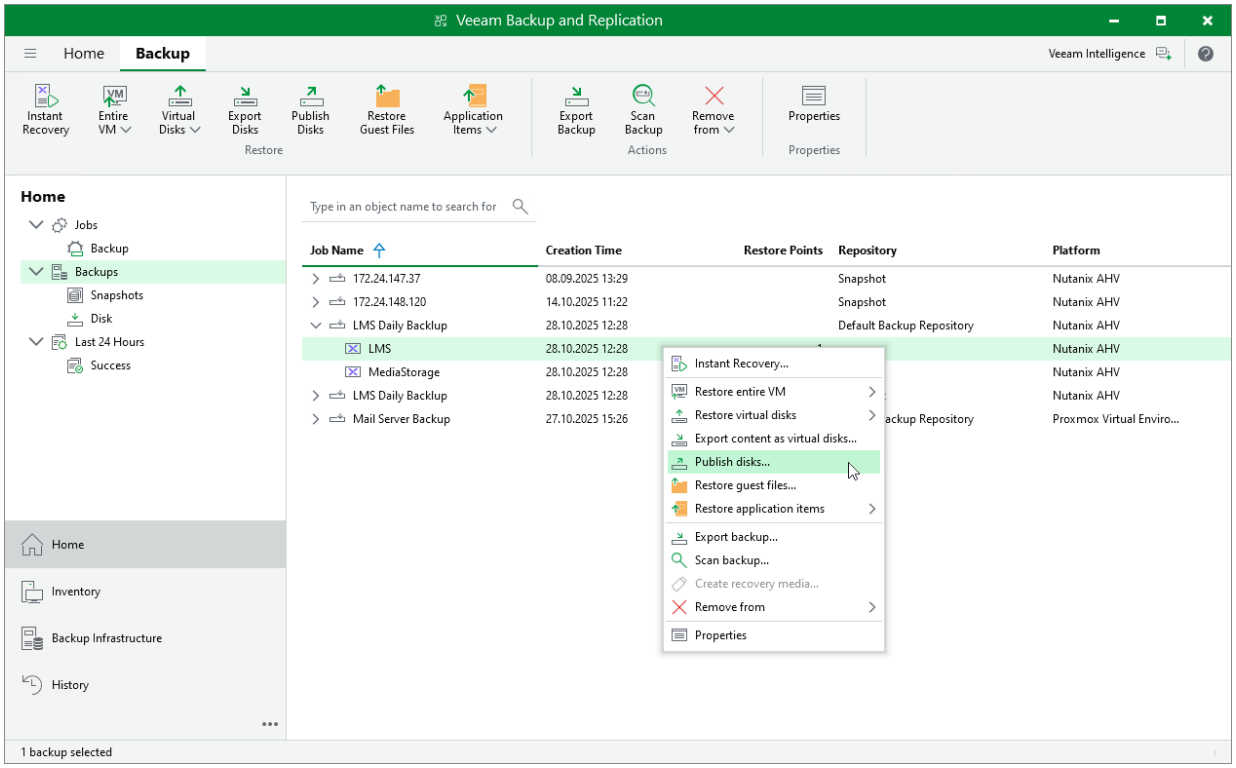

# Performing File-Level Restore

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 Plug-in 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.

> **NOTE**
>
> Veeam Plug-in for Nutanix AHV does not support restore from Novell Storage Service (NSS) file systems.

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 job, right-click the VM that contains files you want to restore and select **Restore guest files**.

   Alternatively, expand the necessary backup job, select the VM and click **Restore Guest Files** on the ribbon.

> **TIP**
>
> To restore VM guest OS files and folders from a backup snapshot, expand the job that contains the snapshot of the VM, right-click the VM and select **Restore guest files**.

4. Complete the **File Level Restore** wizard as described in the Veeam Backup & Replication User Guide, section Recovering Guest OS Files Using Console.

> **NOTE**
>
> Depending on the operating system of a VM whose files and folders you want to restore, Veeam Backup & Replication may require a mount host — a server that will be used to mount VM disks. While completing the **Guest File Restore** wizard, you will be able either to choose a server already added to the backup infrastructure or to specify connection settings of a new server that will used as the mount host. For more information on how Veeam Backup & Replication selects mount hosts, see the Veeam Backup & Replication User Guide, section Mount Host Automatic Selection.

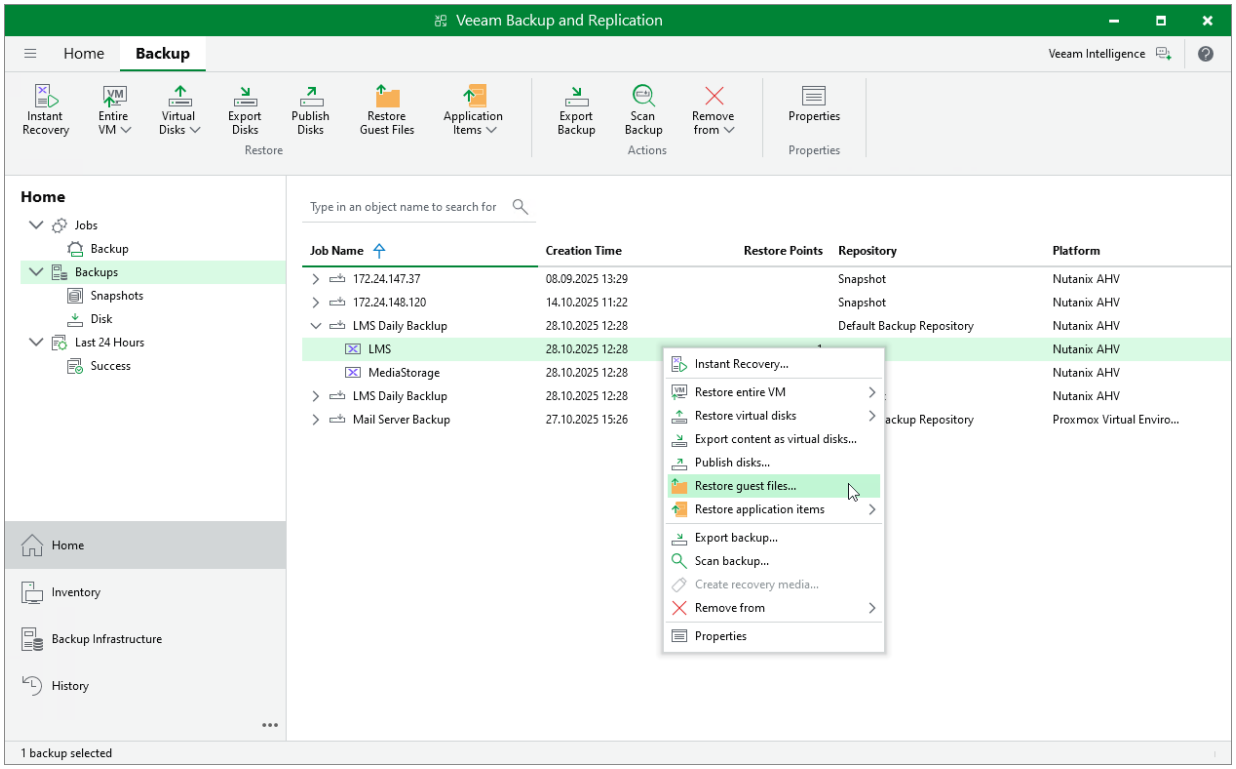


## TIP

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 and 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 databases
- Oracle databases
- PostgreSQL databases

> **NOTE**
>
> It is recommended that you use application-consistent backups for application item restore.

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 job, right-click the VM that contains an application you want to restore, select **Restore application items** and select the application.

   Alternatively, expand the necessary backup job, select the VM, click **Application Items** on the ribbon and select the application.

> **TIP**
>
> To restore application items from a backup snapshot, expand the job that contains the snapshot of the VM, right-click the VM and select **Restore application items**.

4. In the restore wizard, select a restore point that will be used to restore the application, specify a restore reason and click **Browse.**
5. 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.

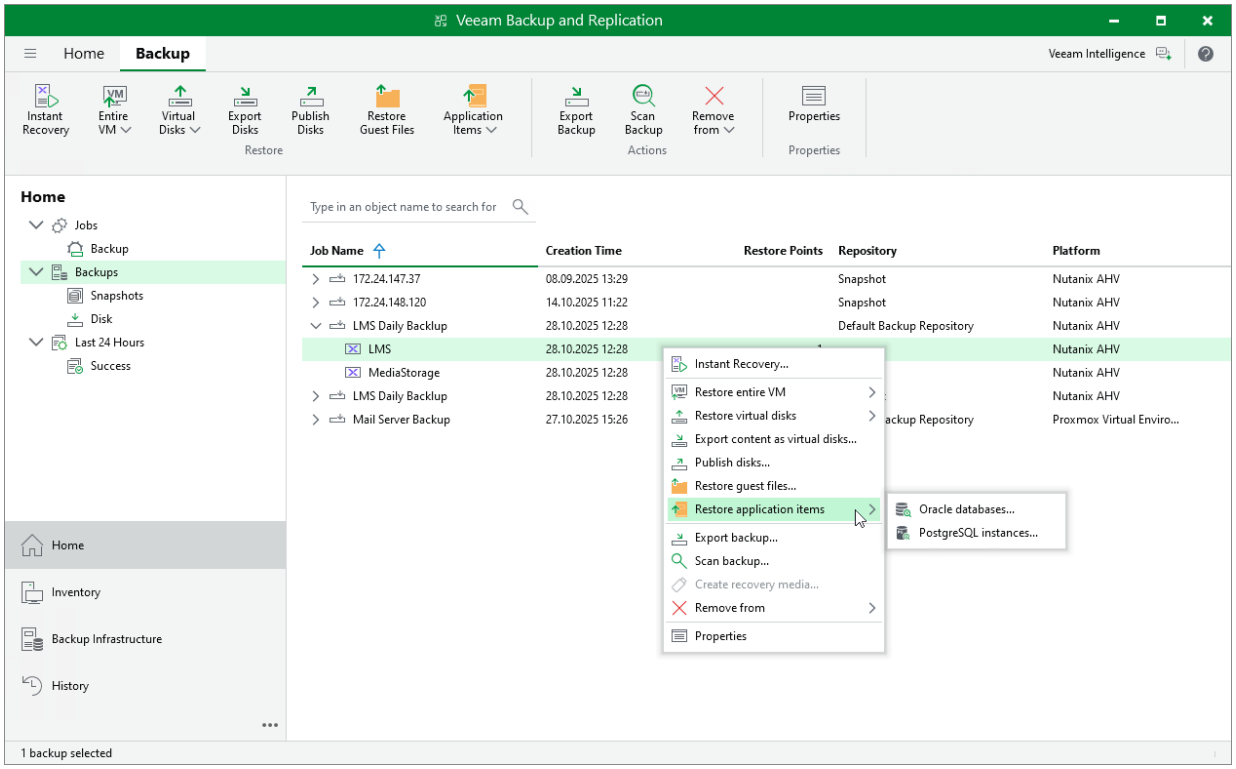
Veeam Backup and Replication
Home
Backup
Veeam Intelligence
Instant Recovery
Entire VM
Virtual Disks
Export Disks
Publish Disks
Restore Guest Files
Application Items
Restore
Export Backup
Scan Backup
Remove from
Actions
Properties
Properties
Home
Jobs
Backup
Backups
Snapshots
Disk
Last 24 Hours
Success
Type in an object name to search for
Job Name
Creation Time
Restore Points
Repository
Platform
172.24.147.37
08.09.2025 13:29
Snapshot
Nutanix AHV
172.24.148.120
14.10.2025 11:22
Snapshot
Nutanix AHV
LMS Daily Backup
28.10.2025 12:28
Default Backup Repository
Nutanix AHV
LMS
28.10.2025 12:28
Nutanix AHV
MediaStorage
28.10.2025 12:28
Nutanix AHV
LMS Daily Backup
28.10.2025 12:28
Nutanix AHV
Mail Server Backup
27.10.2025 15:26
Proxmox Virtual Enviro...
Instant Recovery...
Restore entire VM
Restore virtual disks
Export content as virtual disks...
Publish disks...
Restore guest files...
Restore application items
Export backup...
Scan backup...
Create recovery media...
Remove from
Properties
Oracle databases...
PostgreSQL instances...
Home
Inventory
Backup Infrastructure
History
1 backup selected

# Exporting Disks

Veeam Plug-in 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**.

   Alternatively, expand the necessary backup job, select the VM and click **Export Disks** on the ribbon.
4. Complete the **Export Disk** wizard as described in the Veeam Backup & Replication User Guide, section Exporting Disks.

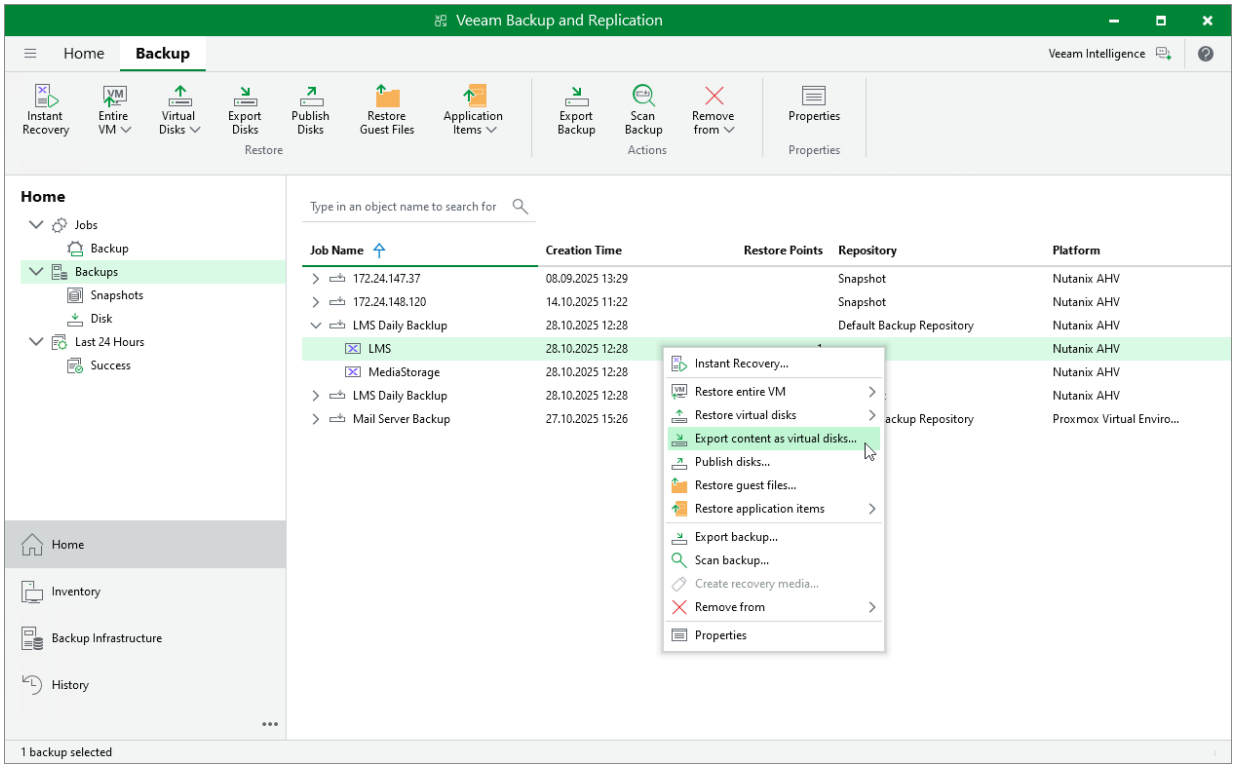

# Performing VM Restore to Amazon Web Services

Veeam Plug-in 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 **Amazon EC2**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Amazon EC2** on the ribbon.
4. Complete the **Restore to Amazon EC2** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Amazon EC2.

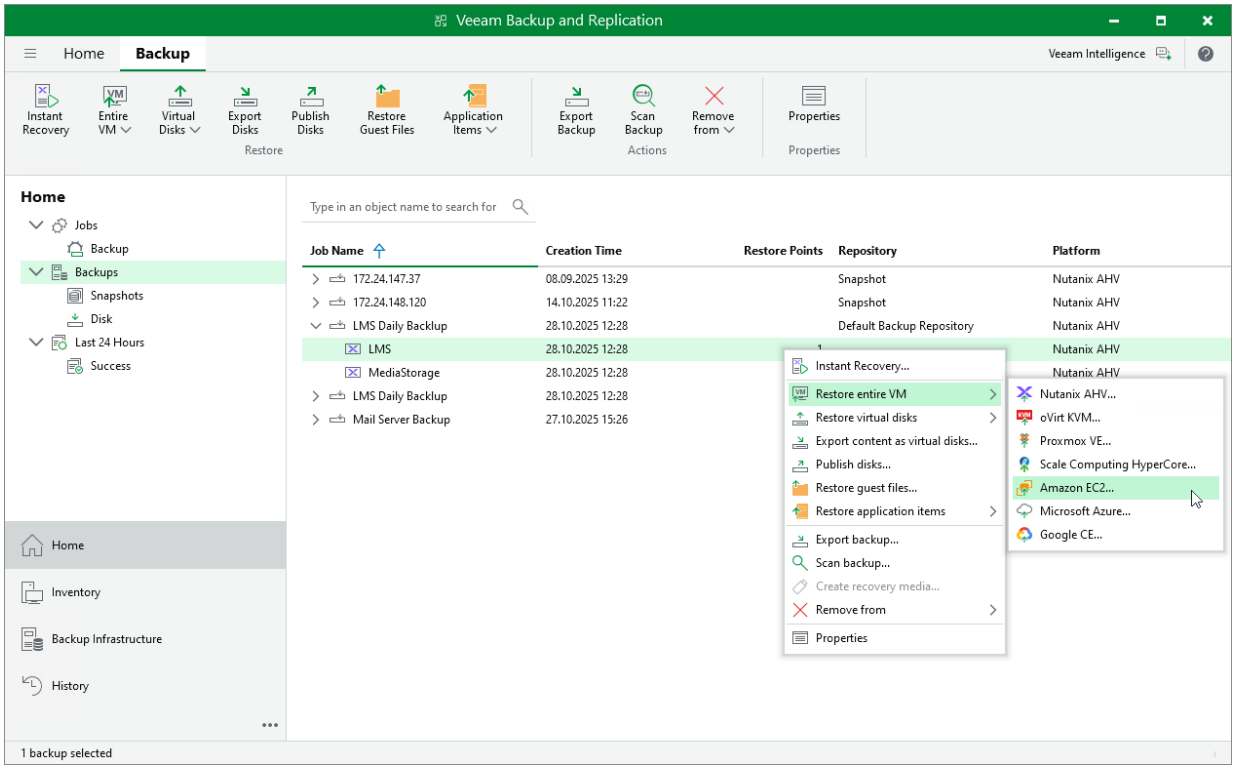

# Performing VM Restore to Microsoft Azure

Veeam Plug-in 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 **Microsoft Azure**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Microsoft Azure** on the ribbon.
4. Complete the **Restore to Microsoft Azure** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Microsoft Azure.

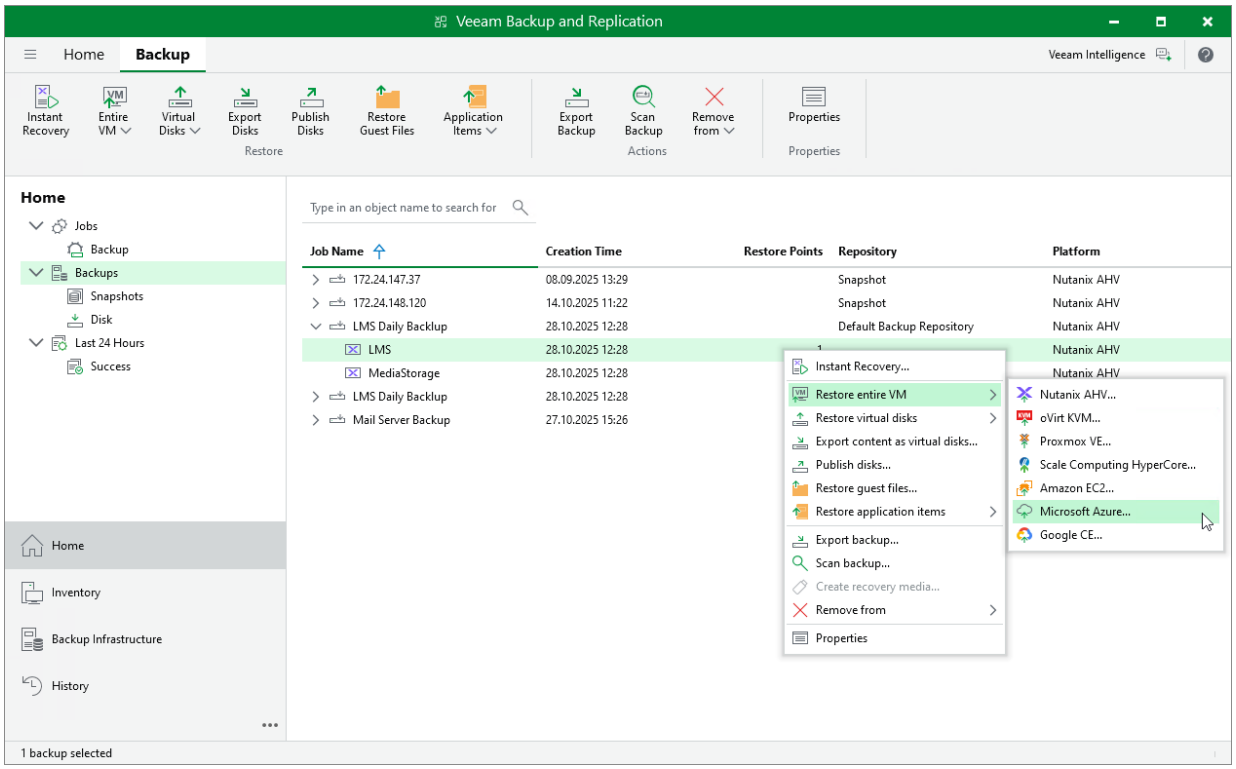

# Performing VM Restore to Google Cloud

Veeam Plug-in 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 **Google CE**.

   Alternatively, expand the necessary backup job, select the VM and click **Entire VM** > **Google CE** on the ribbon.
4. Complete the **Restore to Google Compute Engine** wizard as described in the Veeam Backup & Replication User Guide, section Restoring to Google Compute Engine.

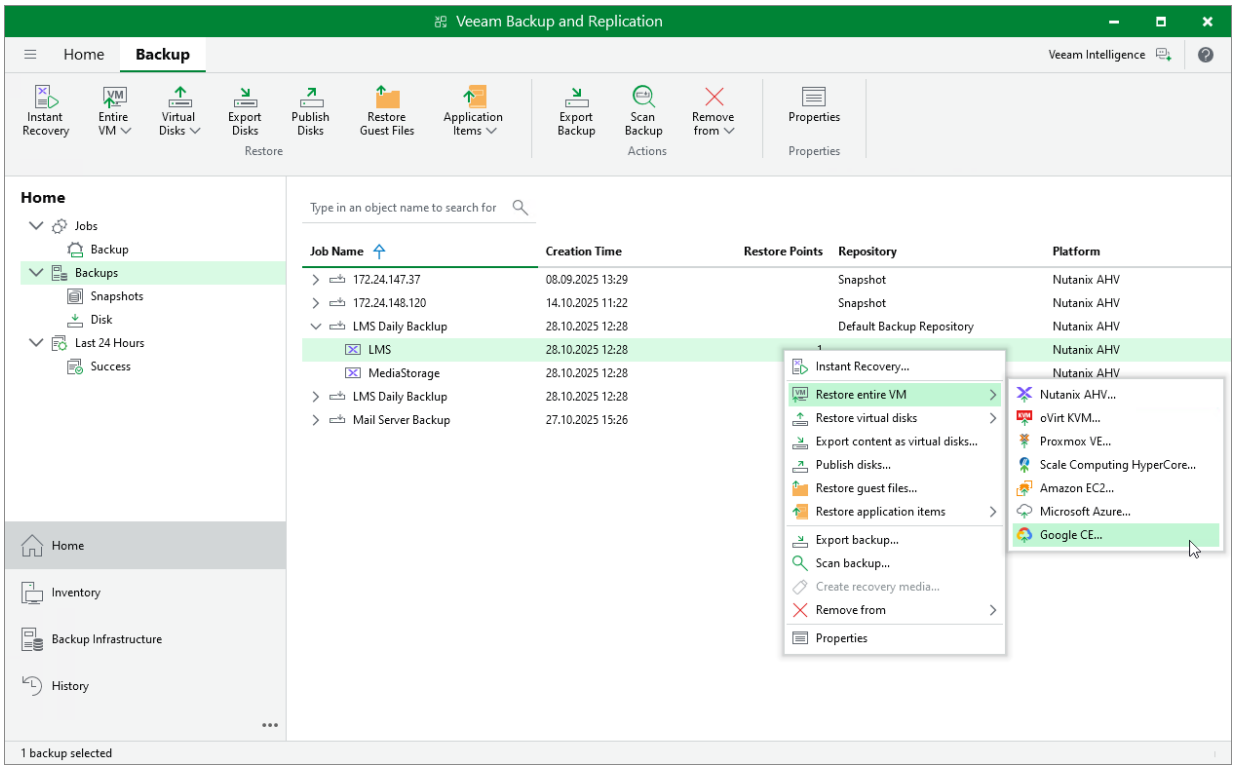

# Getting Technical Support

If you have any questions or issues with Veeam Plug-in 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 infrastructure components
- The error message or an accurate description of the problem you are facing
- Log files

## Viewing Product Details

To view the product details, do the following:

1. On the machine where the Veeam Backup & Replication console is installed, navigate to the **Control Panel**.
2. In the **Control Panel** window, navigate to **Programs** > **Programs and Features**.
3. In the program list, check the version of **Veeam Plug-in for Nutanix AHV**.

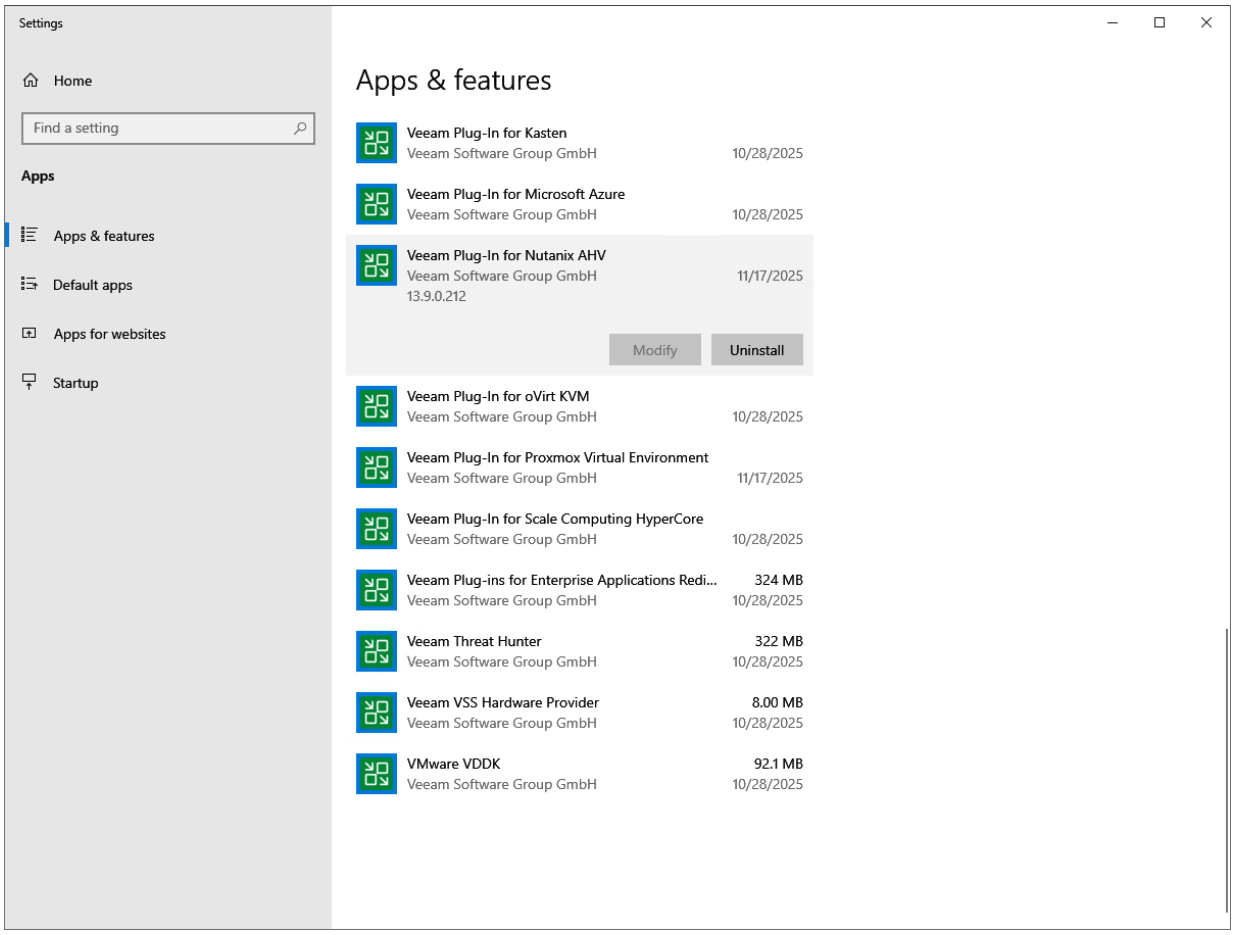

> **TIP**
>
> [Applies to Linux-based backup servers only] Alternatively, you can view product details in the Veeam Host Management Console as described in the Veeam Backup & Replication User Guide, section Performing Maintenance Tasks.

# Downloading Logs

To download the product logs, 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.

   Complete the wizard as described in the Veeam Backup & Replication User Guide, section Exporting Logs.

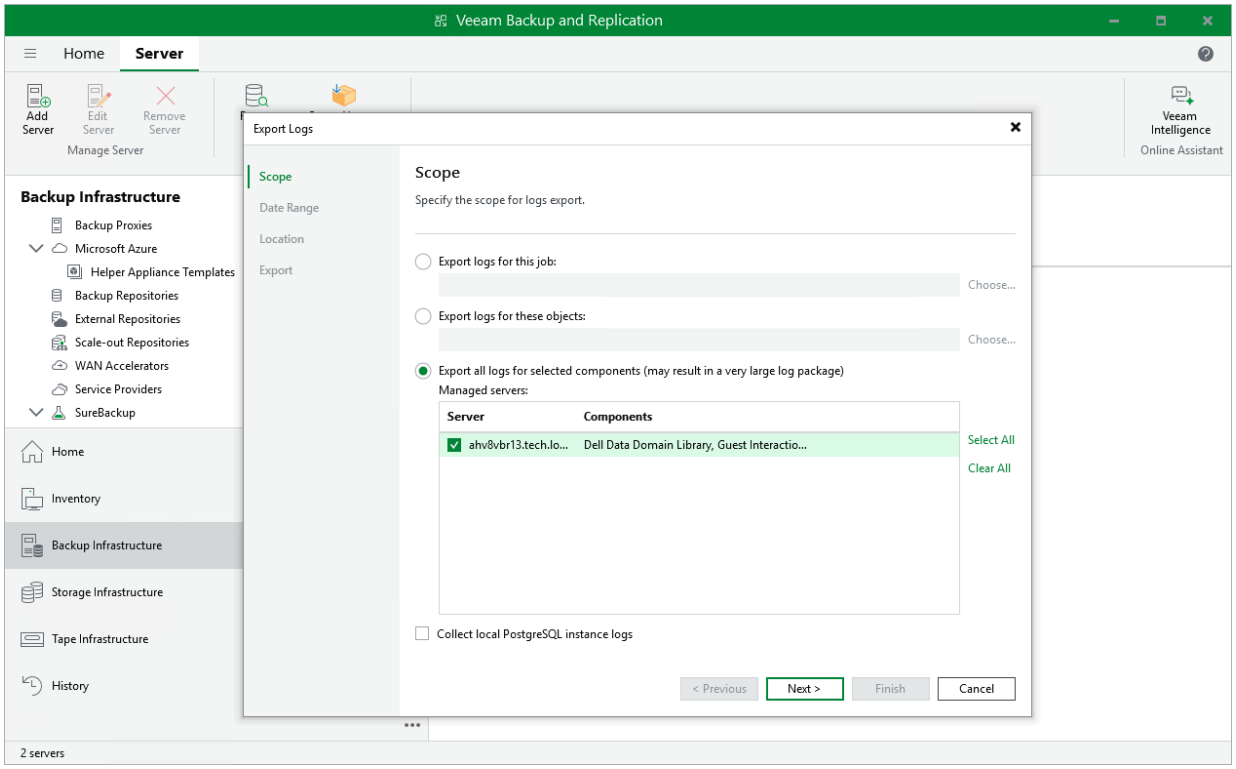

# Appendices

See in this section:

# Appendix A. Deprecated Functionality

Starting from version 8, Veeam Plug-in for Nutanix AHV comes without the Nutanix AHV backup appliance, the following changes have been introduced to the product functionality:

| Functionality | Availability in Versions Prior to 8 | Availability in Versions 8 and 9 |
|---|---|---|
| Managing the backup appliance | Veeam Backup & Replication console | – |
| Managing backup appliance users | Backup appliance web console | – |
| Backing up and restoring the backup appliance configuration | Veeam Backup & Replication console | – |
| Reviewing the backup appliance dashboard | Backup appliance web console | – |
| Managing workers | Backup appliance web console | Veeam Backup & Replication console |
| Creating backup jobs | Backup appliance web console<br>Veeam Backup & Replication console | Veeam Backup & Replication console |
| Creating snapshot jobs | Backup appliance web console<br>Veeam Backup & Replication console | – |
| Creating PD snapshot jobs | Backup appliance web console<br>Veeam Backup & Replication console | – |
| Deleting snapshots | Backup appliance web console | – |
| Configuring notification settings | Backup appliance web console | Veeam Backup & Replication console |
| Viewing session statistics | Backup appliance web console | Veeam Backup & Replication console |

| Functionality | Availability in Versions Prior to 8 | Availability in Versions 8 and 9 |
|---|---|---|
| Managing backups | Backup appliance web console<br>Veeam Backup & Replication console | Veeam Backup & Replication console |
| Performing VM restore | Backup appliance web console<br>Veeam Backup & Replication console | Veeam Backup & Replication console |
| Performing disk restore | Backup appliance web console | Veeam Backup & Replication console |
| Instant recovery | Veeam Backup & Replication console | Veeam Backup & Replication console |

# Appendix B. Configuring Bus Type Restore Priority

When restoring a VM that originally resided on a platform other than Nutanix AHV, Veeam Plug-in 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 Plug-in 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 Plug-in for Nutanix AHV will process disks that cannot be attached using their original bus types. You can also instruct Veeam Plug-in for Nutanix AHV to ignore the original bus types of VM disks. In the latter case, Veeam Plug-in 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 Plug-in 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 Plug-in 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 | • 6 SATA (originally, 6 SATA)<br>• 24 SCSI (originally, 24 SATA)<br>• 2 IDE (originally)<br>• 0 PCI |
| SATA, IDE, PCI, SCSI | False | • 6 SATA (originally, 6 SATA)<br>• 4 IDE (originally, 2 IDE and 2 SATA)<br>• 7 PCI (originally, 7 SATA)<br>• 15 SCSI (originally, 15 SATA) |
| SCSI, IDE, PCI, SATA | False | • 24 SCSI (originally, 24 SATA)<br>• 2 IDE (originally, 2 IDE)<br>• 0 PCI<br>• 6 SATA (originally, 6 SATA) |
| SCSI, IDE, PCI, SATA | True | • 32 SCSI (originally, 30 SATA and 2 IDE)<br>• 0 IDE<br>• 0 PCI<br>• 0 SATA |

# Configuring Bus Type Priority on Linux Server

To modify the default bus type restore priority on a Linux-based backup server, do the following:

1. In a web browser, log in to the **Host Management** web console as described in the Veeam Backup & Replication User Guide, section Accessing Host Management Console.
2. Navigate to **Logs and Services > Host Configuration**.
3. In the **Configuration Files** section, select the `/etc/veeam/platform-service-ahv/appsettings.json` file and click **Export**.

   The file will be downloaded to your local machine.
4. Use a plain text editor to open the .JSON file.
5. Locate the **RestoreDefaults** configuration section.

   To instruct Veeam Plug-in 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",
   ```
6. Save the `appsettings.json` file.
7. Back to the web browser, select the `/etc/veeam/platform-service-ahv/appsettings.json` file, click **Import**, choose the updated `appsettings.json` file and click **Open**.
8. Restart the Veeam Nutanix AHV Platform Service. To do that, navigate to **Logs and Services > Services**, select *veeam-platform-service-ahv.service* and click **Restart**.

# Configuring Bus Type Priority on Windows Server

To modify the default bus type restore priority on a Windows-based backup server, 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 Plug-in 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 C. Configuring Multiple Networks

Starting from version 6.0, Nutanix AHV allows you to connect 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 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 have been added first while configuring workers, which means 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 workers:

- If you want 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 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 workers, the backup server, the repository and the Nutanix AHV cluster are connected to Network B, while 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.

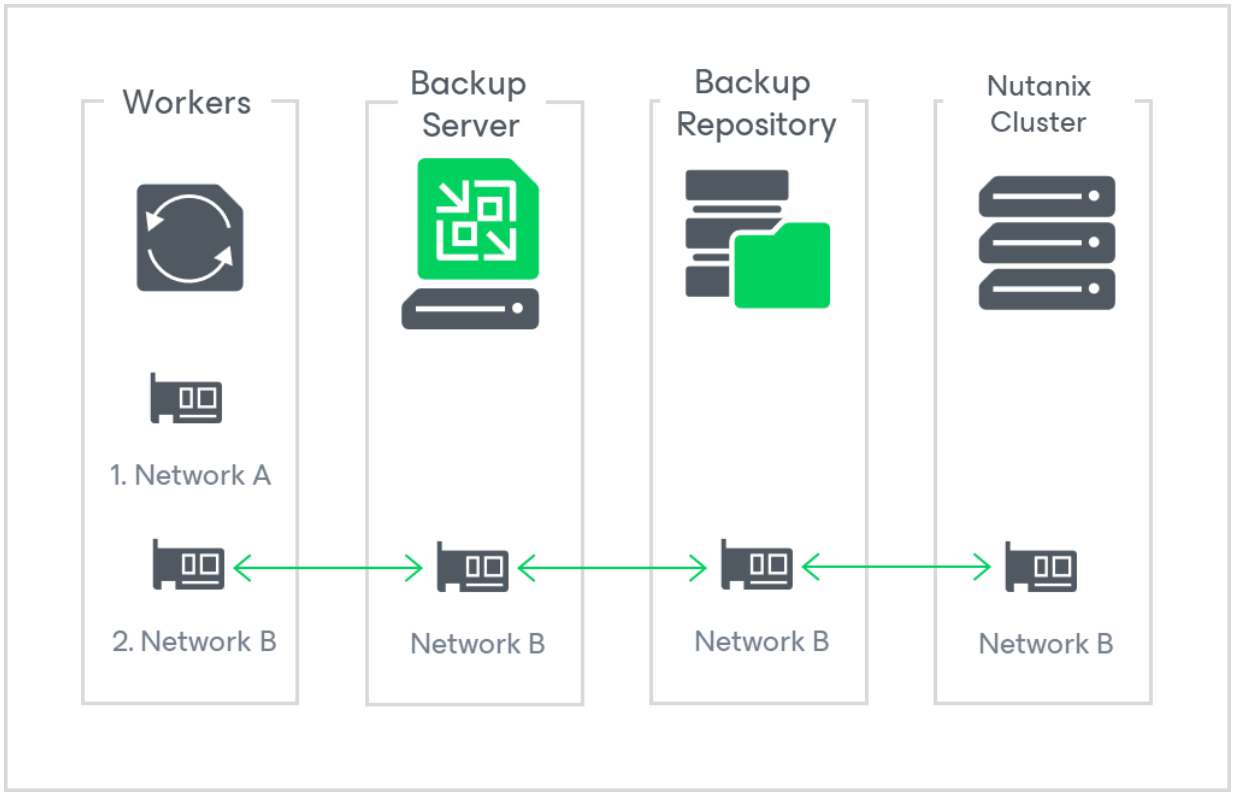

# Example 2. Valid Configuration

In this example, the 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 workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.

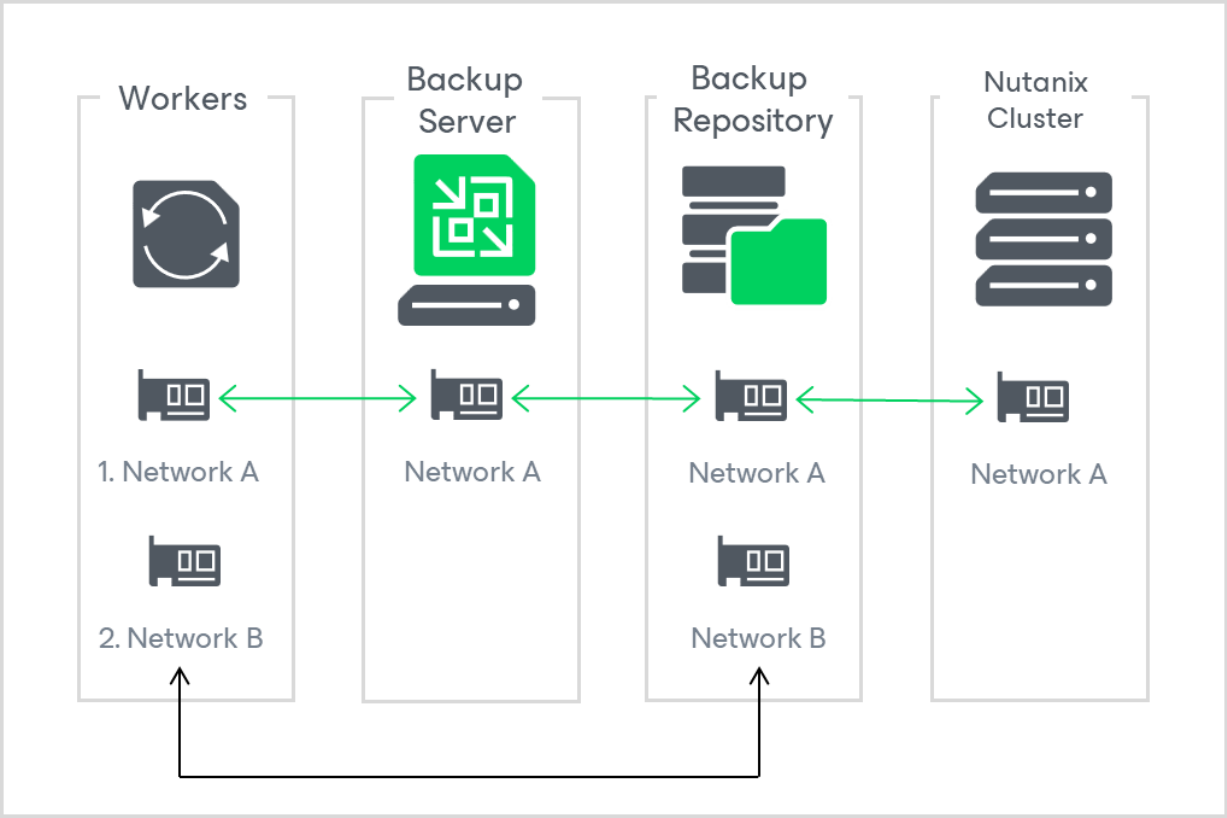


# Example 3. Valid Configuration

In this example, the 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 workers can use Network A to communicate with other backup infrastructure components though the router.

> **NOTE**
>
> The workers will be able to obtain updates from online Veeam repositories only if Network A is configured to allow inbound and outbound internet traffic.

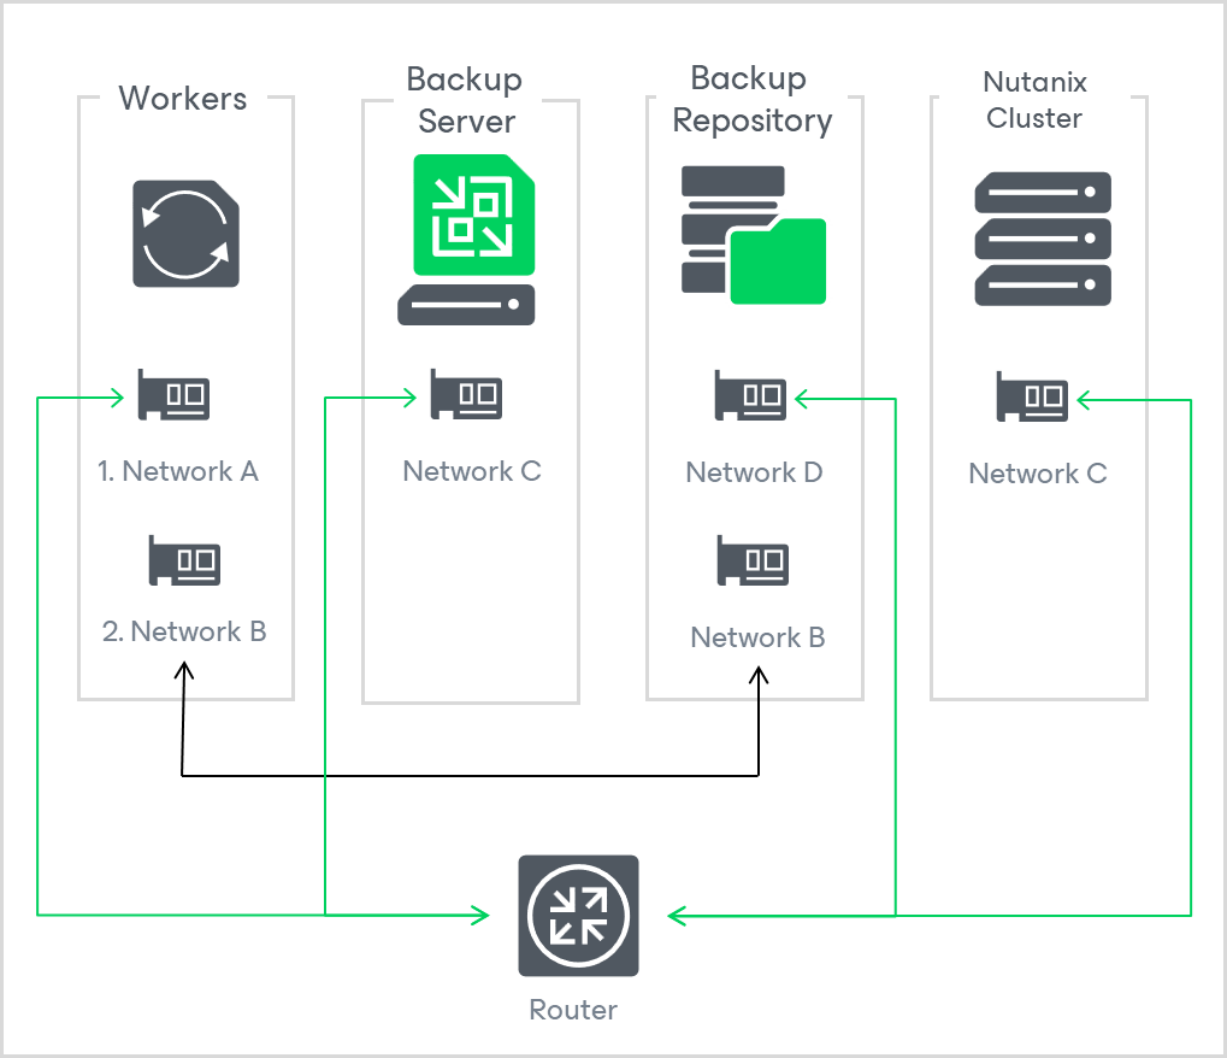


# Example 4. Invalid Configuration

In this example, the 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 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 workers. Then, connect these vNICs to Network C.