



Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization

Version 4.1

User Guide

April, 2024

© 2024 Veeam Software.

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

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

NOTE

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

Contents

CONTACTING VEEAM SOFTWARE	6
ABOUT THIS DOCUMENT	7
WELCOME TO VEEAM BACKUP FOR ORACLE LINUX VIRTUALIZATION MANAGER AND RED HAT VIRTUALIZATION.....	8
ARCHITECTURE OVERVIEW	9
PLANNING AND PREPARATION	11
System Requirements	12
Permissions	13
Ports	14
LICENSING	18
DEPLOYMENT	19
Installing oVirt KVM Plug-In	20
Installing oVirt KVM Plug-In in Unattended Mode	22
Upgrading to Veeam Backup for OLVM and RHV 4.1	25
Uninstalling oVirt Plug-In.....	26
CONFIGURING BACKUP INFRASTRUCTURE	27
Configuring Backup Repositories	28
Connecting Virtualization Manager	29
Adding Virtualization Manager to Backup Infrastructure.....	30
Editing Virtualization Manager Properties	36
Rescanning Virtualization Manager	37
Removing Virtualization Manager	38
Managing Backup Appliance	39
Deploying New Backup Appliance	40
Connecting Existing Backup Appliance	50
Editing Backup Appliance	59
Rescanning Backup Appliance.....	61
Removing Backup Appliance	62
Managing Workers	64
Adding Workers	65
Enabling and Disabling Workers	72
Editing Workers	73
Updating Workers	74
Removing Workers	75
PERFORMING CONFIGURATION BACKUP AND RESTORE.....	76
Backing Up Configuration Settings Manually	77
Backing Up Configuration Settings Automatically	79

Restoring Configuration Settings	81
Step 1. Launch Configuration Restore Wizard	82
Step 2. Choose Backup File	83
Step 3. Review Backup Details	84
Step 4. Provide Encryption Password	85
Step 5. Choose Restore Options	86
Step 6. Track Restore Progress	87
Step 7. Finish Working with Wizard	88
PERFORMING BACKUP	89
How Backup Works	90
Backup Chain	91
Retention Policy	96
Creating Backup Jobs	98
Before You Begin	99
Step 1. Launch New Backup Job Wizard	100
Step 2. Specify Job Name and Description	101
Step 3. Configure Backup Source Settings	102
Step 4. Specify Backup Job Settings	106
Step 5. Define Job Schedule	112
Step 6. Finish Working with Wizard	113
Editing Backup Job Settings	114
Starting and Stopping Backup Jobs	115
Analyzing Performance Bottlenecks	116
Cloning Backup Jobs	118
Enabling and Disabling Backup Jobs	119
Deleting Backup Jobs	120
Creating Active Full Backups	121
Creating VeeamZIP Backups	122
MANAGING BACKUPS	124
Viewing Backup Properties	125
Verifying Backups	126
Exporting Backups	127
Copying Backups	128
Copying Backups to Tapes	129
Deleting Backups	130
PERFORMING RESTORE	131
Performing VM Restore	132
Step 1. Launch Full VM Restore to oVirt KVM Wizard	134
Step 2. Select Restore Point	135
Step 3. Choose Restore Mode	136

Step 4. Specify Target Cluster	137
Step 5. Select Storage Domain	138
Step 6. Specify VM Name	139
Step 7. Configure Network Settings	140
Step 8. Specify Restore Reason.....	141
Step 9. Finish Working with Wizard	142
Performing Disk Restore	143
Step 1. Launch Virtual Disk Restore Wizard	144
Step 2. Select Virtual Machine	145
Step 3. Select Restore Point	146
Step 4. Configure Mapping Settings	147
Step 5. Specify Reason for Restore	148
Step 6. Finish Working with Wizard	149
Performing Instant VM Recovery	150
Publishing Disks	151
Performing File-Level Restore	152
Performing Application Item Restore	153
Exporting Disks	155
Performing VM Restore to Amazon Web Services	156
Performing VM Restore to Microsoft Azure	157
Performing VM Restore to Google Cloud	158
UPDATING BACKUP APPLIANCE	159
GETTING TECHNICAL SUPPORT.....	161
APPENDIX. DEPRECATED FUNCTIONALITY	164

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 protect workloads in Red Hat Virtualization or Oracle Linux KVM virtual environment. 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 oVirt KVM environments.

Welcome to Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization

Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization (Veeam Backup for OLVM and RHV) is a solution developed for protection and disaster recovery tasks for oVirt KVM environments. With Veeam Backup for OLVM and RHV, you can perform the following operations:

- Create backups of oVirt VMs and store them in backup repositories.
- Create several instances (copies) of the same backup data in different locations.
- Restore VMs from oVirt VM backups to oVirt KVM environments.
- Restore VMs from oVirt VM backups to Microsoft Azure, Amazon Web Services (AWS) and Google Cloud environments.
- Perform Instant Recovery of oVirt VMs to Nutanix AHV, VMware vSphere and Microsoft Hyper-V environments.
- Restore files and folders of oVirt VM guest OSes.
- Restore oVirt VM disks and attach them to VMs running on oVirt KVM hosts.
- Export disks of backed-up oVirt VMs to VMDK, VHD and VHDX formats.
- Mount disks of backed-up oVirt VMs to any server and access data in the read-only mode.

Architecture Overview

The Veeam Backup for OLVM and RHV architecture comprises the following set of components:

- [Virtualization manager](#)
- [Backup server](#)
- [Backup appliance](#)
- [oVirt KVM Plug-in](#)
- [Backup repositories](#)
- [Workers](#)

Virtualization Manager

Virtualization manager is a Linux-based physical or virtual machine that manages oVirt resources such as VMs, hosts, clusters, storage domains and networks. Veeam Backup for OLVM and RHV uses the Virtualization manager to access oVirt resources while performing backup and restore operations.

Backup Server

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

Backup Appliance

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

oVirt KVM Plug-in

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

Backup Repositories

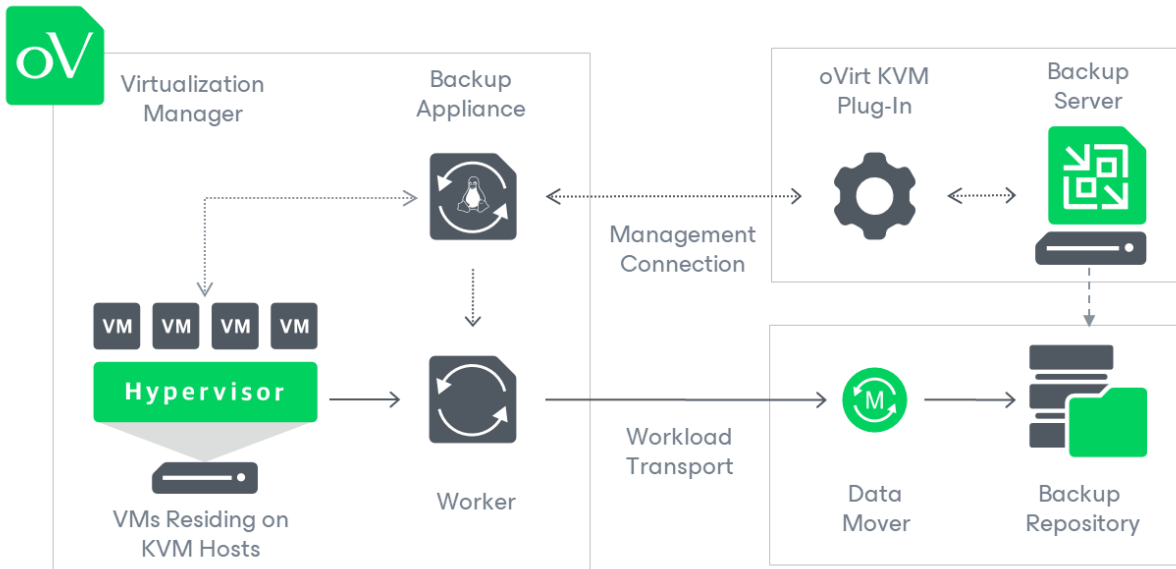
A backup repository is a storage location where Veeam Backup for OLVM and RHV stores backups of protected oVirt VMs.

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

Workers

A worker is an auxiliary Linux-based VM that resides in the cluster and processes backup workloads when transferring data to and from backup repositories.

The backup appliance comes with a preconfigured embedded worker that can be used in small virtual environments. In large environments, it is recommended to deploy dedicated workers that are distributed among the cluster hosts (nodes) and are automatically launched for the duration of a backup or restore process.



Planning and Preparation

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

System Requirements

Before you start deploying Veeam Backup for OLVM and RHV, make sure the virtual environment and the backup infrastructure components meet the following requirements.

Specification	Requirement
Hypervisor	Kernel-based Virtual Machine (KVM) must be installed on x86 hardware that supports virtualization capabilities.
Virtualization Platform	Veeam Backup for OLVM and RHV is supports with the following virtual environments: <ul style="list-style-type: none">• Red Hat Virtualization version 4.4 SP1 only (Red Hat Virtualization Manager version 4.5.0 or later)• Oracle Linux Virtualization version 4.5.4 or later.
Veeam Software	Veeam Backup & Replication version 12.1 with oVirt KVM Plug-in version 12.4.1.45 (or later) must be deployed on the backup server.
Backup Appliance	<p>The backup appliance performs management operations and handles data protection tasks. If you deploy Veeam Backup for OLVM and RHV using the default configuration, the following compute resources will be allocated to the backup appliance:</p> <ul style="list-style-type: none">• <i>CPU</i>: 4 vCPU• <i>Memory</i>: 4 GB RAM• <i>Disk Space</i>: 100 GB for product installation, internal database files and logs <p>With the default configuration, the appliance can handle up to 4 concurrent backup and restore tasks if the embedded worker is enabled. While deploying a new backup appliance or editing settings of an existing one, you can increase the maximum number of concurrent tasks. However, you must allocate 1 vCPU and 1 GB RAM for each additional task. When configuring the maximum number of concurrent tasks, you must also take into account the network traffic throughput in your virtual infrastructure.</p>
Workers	<p>VMs running as dedicated workers must be allocated the following compute resources for each concurrent task:</p> <ul style="list-style-type: none">• CPU: 1 vCPU• Memory: 1 GB RAM

Permissions

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

Backup Server Windows Account Permissions

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

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

Virtualization Manager Permissions

The administrator account that the backup server uses to access the Virtualization manager must have the *SuperUser* privileges. For more information on system permissions, see [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

Ports

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

Backup Appliance

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

From	To	Protocol	Port	Notes
Backup appliance	Backup server	TCP	10006	Used to communicate with the Veeam Backup & Replication server.
	Virtualization manager	TCP/HTTPS	443	Used to communicate with the REST API service running on the Virtualization manager.
	Virtualization manager	TCP	54323	Used to communicate with Virtualization manager (hosted engine).
	KVM host	TCP/HTTPS	443	Used to communicate with the REST API service running on an KVM host.
	KVM host	TCP	54322	Used to communicate with KVM hosts.
	Workers	TCP	19000	Used to communicate with workers.
	Veeam backup repository or gateway server	TCP	2500-3300	Default range of ports used as transmission channels for jobs and restore sessions. For each TCP connection that a job uses, one port from this range is assigned.
	Ubuntu Security and OS Update repository <i>(security.ubuntu.com, archive.ubuntu.com)</i>	TCP/HTTP(S)	80 (443)	Used to get OS security updates.
	.NET Core Update repository	TCP/HTTPS	443	Used to get .NET Core updates.

From	To	Protocol	Port	Notes
	Veeam Updater repository (repository.veeam.com , cloudfront.net)	TCP/HTTPS	443	Used to download backup appliance update packages.
	Nginx repository (nginx.org/packages/ , nginx.org/packages/keys/)	TCP/HTTPS	443	Used to download Nginx packages required for backup appliance web console updates.

Workers

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

From	To	Protocol	Port	Notes
Worker	Backup server	TCP	10006	Used to communicate with the Veeam Backup & Replication server.
	Virtualization manager	TCP/HTTPS	443	Used to communicate with the REST API service running on the Virtualization manager.
	Virtualization manager	TCP	54323	Used to communicate with Virtualization manager (hosted engine).
	KVM host	TCP/HTTPS	443	Used to communicate with the REST API service running on an KVM host.
	KVM host	TCP	54322	Used to communicate with KVM hosts.
	Backup appliance	TCP	19001	Used to communicate with the backup appliance.
	Veeam backup repository or gateway server	TCP	2500-3300	Default range of ports used as transmission channels for jobs and restore sessions. For each TCP connection that a job uses, one port from this range is assigned.

From	To	Protocol	Port	Notes
	Ubuntu Security and OS Update repository <i>(security.ubuntu.com, archive.ubuntu.com)</i>	TCP/HTTP(S)	80 (443)	Used to get OS security updates.
	.NET Core Update repository	TCP/HTTPS	443	Used to get .NET Core updates.
	Veeam Updater repository <i>(repository.veeam.com, cloudfront.net)</i>	TCP/HTTPS	443	Used to download backup appliance update packages.
	Nginx repository <i>(nginx.org/packages/, nginx.org/packages/keys/)</i>	TCP/HTTPS	443	Used to download Nginx packages required for backup appliance web console updates.

Backup Server

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

From	To	Protocol	Port	Notes
Backup appliance, Veeam Backup & Replication console	Backup server	TCP/HTTPS	8544	Used to communicate with the Platform Service REST API.
Backup server	FLR helper appliance	TCP	22	Used to connect to the helper appliance during file-level restore.
	Backup server	TCP/HTTPS	6172	Used by the Platform Service to enable communication with the Veeam Backup & Replication database.
	Virtualization manager	TCP/HTTPS	443	Used to communicate with the REST API service running on the Virtualization manager.

From	To	Protocol	Port	Notes
	Virtualization manager	TCP	54323	Used to communicate with the Virtualization manager (hosted engine).
	Backup appliance	TCP/HTTPS	443	Used by the Platform Service to connect to the backup appliance.

NOTE

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

Licensing

Veeam Backup for OLVM and RHV is licensed by the number of protected oVirt VMs. Each protected oVirt VM consumes one Veeam Universal License instance from the license scope. An oVirt VM is considered protected if it has a restore point created during the past 31 days.

Obtaining New License

You can obtain the following types of licenses for Veeam Backup for OLVM and RHV:

- **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 oVirt VMs. You can check the number of available license instances in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section [Viewing License Information](#).

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

Deployment

To deploy Veeam Backup for OLVM and RHV, do the following:

1. Deploy the backup server as described in the Veeam Backup & Replication User Guide, section [Installing Veeam Backup & Replication](#).

Alternatively, you can use a backup server that already exists in your backup infrastructure if it meets the [Veeam Backup for OLVM and RHV system requirements](#).

2. [Install oVirt KVM Plug-in on the backup server](#).
3. Perform initial configuration of Veeam Backup for OLVM and RHV:
 - a. [Configure backup repositories](#) where Veeam Backup for OLVM and RHV will store backups of oVirt VMs.
 - b. [Add to the backup infrastructure the Virtualization manager](#) that administers oVirt resources you want to protect.
 - c. [Deploy a backup appliance](#) that will process backup and restore operations.
 - d. [Deploy dedicated workers](#) that will transfer backup traffic.

Installing oVirt KVM Plug-In

The default installation package of Veeam Backup & Replication does not provide features that allow you to protect oVirt resources. To be able to add your Virtualization manager and backup appliance to the backup infrastructure, you must install oVirt KVM Plug-in on the backup server.

NOTE

If you use a remote Veeam Backup & Replication console, you do not need to install oVirt KVM Plug-in on the workstation where the remote Veeam Backup & Replication console is deployed. However, you must install oVirt KVM Plug-in on the backup server.

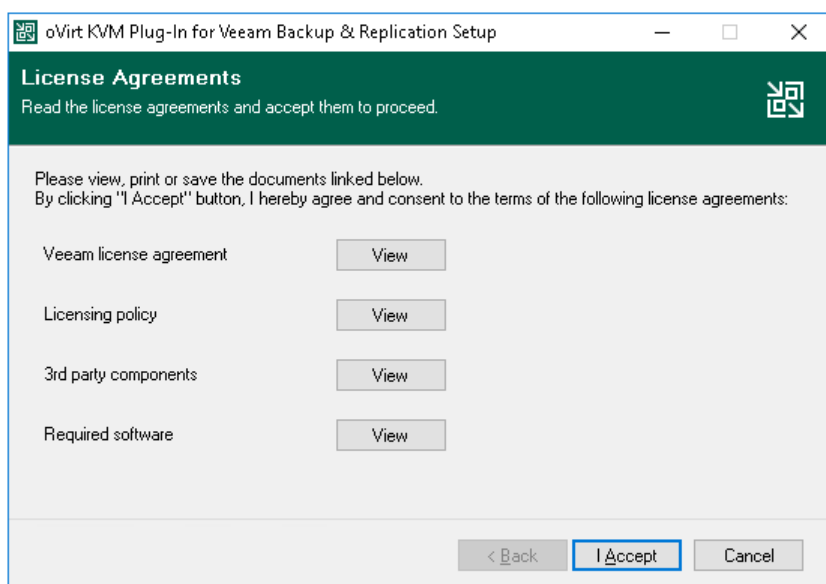
To install oVirt KVM Plug-in, do the following:

1. Log in to the backup server using an account with the local Administrator permissions.
2. Download the product installation file `KVMPlugin_12.4.1.45.zip` from the [Veeam downloads page](#).
3. Open the downloaded archive file and launch the `KVMPlugin_12.4.1.45.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 **oVirt KVM Plug-in for Veeam Backup & Replication Setup** wizard, read and accept both the Veeam license agreement, licensing policy, the 3rd party components and required software license agreement. If you reject the agreements, you will not be able to continue installation.

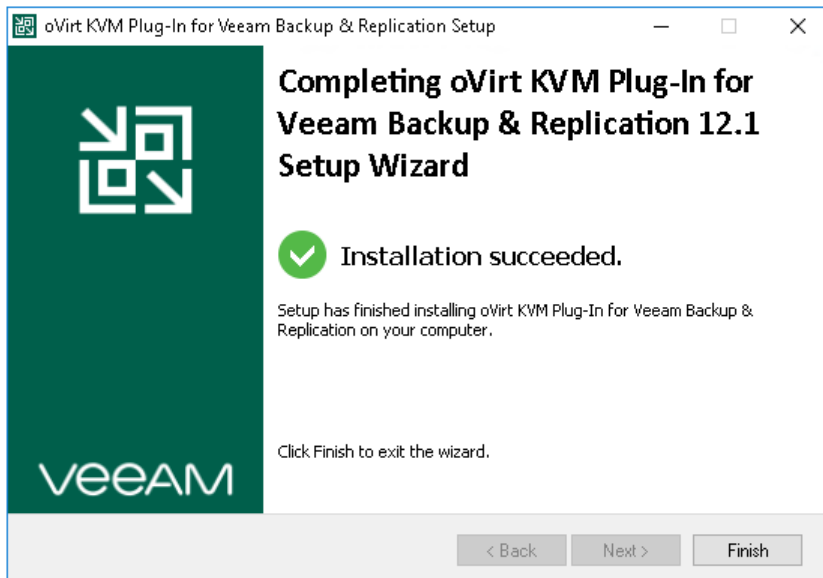
To read the terms of a license agreement, click **View**.



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



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



Installing oVirt KVM Plug-In in Unattended Mode

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

To install oVirt KVM Plug-in in the unattended mode, use either of the following options:

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

Before You Begin

Before you start unattended installation, do the following:

1. Download the `KVMPlugin_12.4.1.45.exe` file as described in section [Installing oVirt KVM Plug-In \(steps 1-3\)](#).
2. Check compatibility of the oVirt KVM Plug-in and Veeam Backup & Replication versions. For more information, see [System Requirements](#).

Installation Command-Line Syntax

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

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

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

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

Parameter	Required	Description
/acceptthirdpartylicenses	Yes	Confirms that you accept the license agreement for 3rd party components that Veeam incorporates.
/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 oVirt KVM Plug-in:

```
KVMPlugin_12.4.1.45.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware
```

The following command repairs oVirt KVM Plug-in:

```
KVMPlugin_12.4.1.45.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware /repair
```

The following command uninstalls oVirt KVM Plug-in:

```
KVMPlugin_12.4.1.45.exe /silent /accepteula /acceptthirdpartylicenses /acceptlicensingpolicy /acceptrequiredsoftware /uninstall
```

Veeam Backup for OLVM and RHV provides the following status codes to report about the installation result:

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

TIP

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

- `VeeamPluginBootstrap.log`
- `RHVPluginSetup.log`
- `RHVPluginUISetup.log`
- `RHVPluginProxySetup.log`

Upgrading to Veeam Backup for OLVM and RHV 4.1

You can upgrade Veeam Backup for Red Hat Virtualization from version 2.0, 2a, 3.0, 3a, 3b or 4.0 to Veeam Backup for OLVM and RHV version 4.1.

IMPORTANT

To upgrade Veeam Backup for Red Hat Virtualization 1.0, you must first upgrade it to version 2a as described in the Veeam Backup for RHV 2.0 User Guide, section [Upgrading to Veeam Backup for RHV 2a](#).

Before you start the upgrade process, do the following:

- [Applies only to upgrading from version 2.0, 2a, 3.0, 3a or 3b] Download Veeam Backup & Replication version 12.1 from the [Veeam downloads page](#).
- Download the latest oVirt KVM Plug-in version from the [Veeam downloads page](#).
- Plan a maintenance period. Typically, the upgrade process takes up to one hour. Make sure there are no jobs currently running or scheduled to run during this period. Wait for the jobs to complete or disable the jobs manually before you start upgrading Veeam Backup for OLVM and RHV.
- Make sure the backup appliance is powered on.
- Back up the configuration database of the backup appliance. For more information, see the following sections:
 - For Veeam Backup for Red Hat Virtualization version 2.0 and 2a, see Backup for Veeam Backup for RHV 2.0 User Guide, section [Performing Configuration Backup](#).
 - For Veeam Backup for Red Hat Virtualization version 3.0 and 3a, see Backup for Veeam Backup for RHV 3.0 User Guide, section [Performing Configuration Backup](#).
 - For Veeam Backup for Red Hat Virtualization version 4.0, see Backup for Veeam Backup for RHV 4.0 User Guide, section [Backing Up Configuration Settings Manually](#).

To upgrade Veeam Backup for Red Hat Virtualization to Veeam Backup for OLVM and RHV 4.1, do the following:

1. [Applies only to upgrading from version 2.0, 2a, 3.0, 3a or 3b] Upgrade your Veeam Backup & Replication server to version 12.1 as described in the Veeam Backup & Replication User Guide, section [Upgrading to Veeam Backup & Replication 12](#). Then, complete the **Components Update** wizard as described in the Veeam Backup & Replication User Guide, section [Server Components Upgrade](#).
Veeam Backup for Red Hat Virtualization will be upgraded to version 4.0.
2. Upgrade oVirt KVM Plug-in to version 4.1. To do that, run the installation file and [complete the oVirt KVM Plug-in for Veeam Backup & Replication wizard](#).
3. Upgrade the backup appliance to version 4.1. To do that:
 - a. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
 - b. Navigate to **Backup Proxies > Out of Date**.
 - c. Select the backup appliance and click **Upgrade Proxy** on the ribbon.
 - d. In the **Components Update** window, click **Apply**.

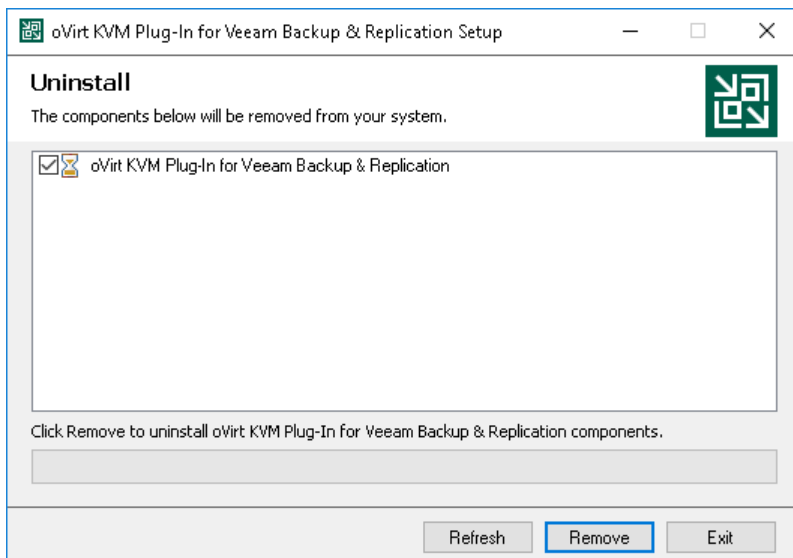
Uninstalling oVirt Plug-In

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

- You will be able to see information on backups of VMs and perform data recovery operations using these backups. However, you will not be able to perform entire VM restore to the oVirt KVM environment.
- You will be able to see information on backup jobs. However, you will only be able to remove these jobs from the Veeam Backup & Replication console.

To uninstall oVirt KVM Plug-in, do the following:

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



Configuring Backup Infrastructure

To set up the backup infrastructure, you must configure backup repositories that will store oVirt VM backups, connect the Virtualization manager that will allow the backup server to access oVirt KVM resources, and add a backup appliance that will process backup and restore operations. For large deployments, it is recommended that you also deploy workers that will transfer backup traffic.

Configuring Backup Repositories

A backup repository is a storage location where Veeam Backup for OLVM and RHV 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.
- **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:** [Amazon S3](#), [Amazon S3 Glacier](#), [AWS Snowball Edge](#), [S3 compatible](#), [Google Cloud](#), [Wasabi Cloud Storage](#), [IBM Cloud](#), [Microsoft Azure Blob](#), [Azure Archive Storage](#) and [Azure Data Box](#)

To combine repositories of different types in one repository, you can also set up a [scale-out backup repository](#).

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

IMPORTANT

Veeam Backup for OLVM and RHV does not support storing backups in [Veeam Cloud Connect](#) repositories. However, you can use Veeam Cloud Connect repositories for [storing copies of backups](#) created with Veeam Backup for OLVM and RHV

Connecting Virtualization Manager

The Virtualization manager allows the backup server to access oVirt resources such as VMs, hosts, clusters, storage domains and networks. After you add the Virtualization manager to the backup infrastructure, you will be able to deploy an backup appliance and to manage data protection tasks for oVirt VMs using the Veeam Backup & Replication console.

Adding Virtualization Manager to Backup Infrastructure

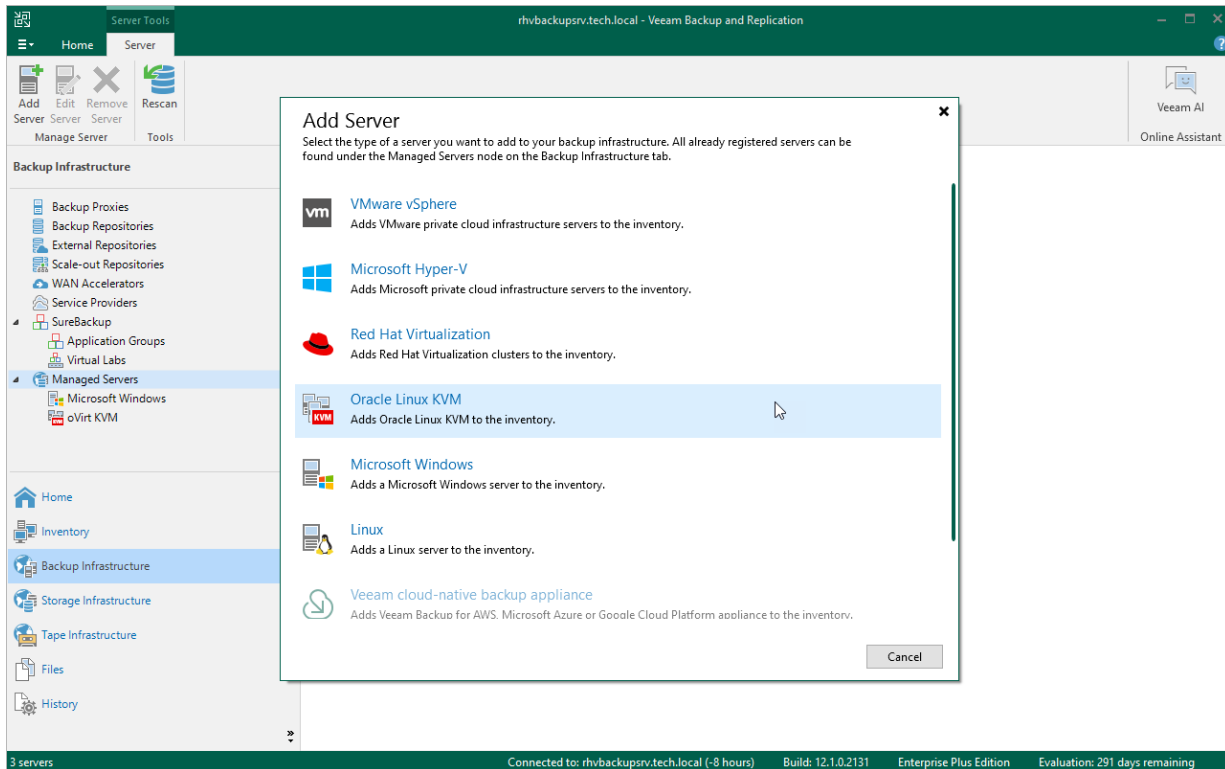
To add the Virtualization manager to the backup infrastructure, do the following:

1. [Launch the New Virtualization Manager wizard.](#)
2. [Specify the Virtualization manager domain name or IP address.](#)
3. [Enter credentials to access the Virtualization manager.](#)
4. [Apply Virtualization manager settings.](#)
5. [Finish working with the wizard.](#)

Step 1. Launch New Virtualization Manager Wizard

To launch the **New Virtualization Manager** 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 **Red Hat Virtualization** or **Oracle Linux KVM** to launch the **New Virtualization Manager** wizard.



Step 2. Specify Domain Name or IP Address of Virtualization Manager

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 Virtualization manager.
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 manager, date and time when the manager was added.

The screenshot shows a window titled "New Virtualization Manager" with a close button (X) in the top right corner. In the top left, there is a KVM icon and the text "Name" followed by the instruction "Specify DNS name or IP address of Virtualization Manager." Below this, a sidebar on the left contains the following options: "Name" (highlighted), "Credentials", "Apply", and "Summary". The main area of the window contains two text input fields. The first field is labeled "DNS name or IP address:" and contains the text "pdcqa387ovirt.robofish.local". The second field is labeled "Description:" and contains the text "Oracle Linux Virtualization Manager". At the bottom of the window, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

Step 3. Enter Credentials

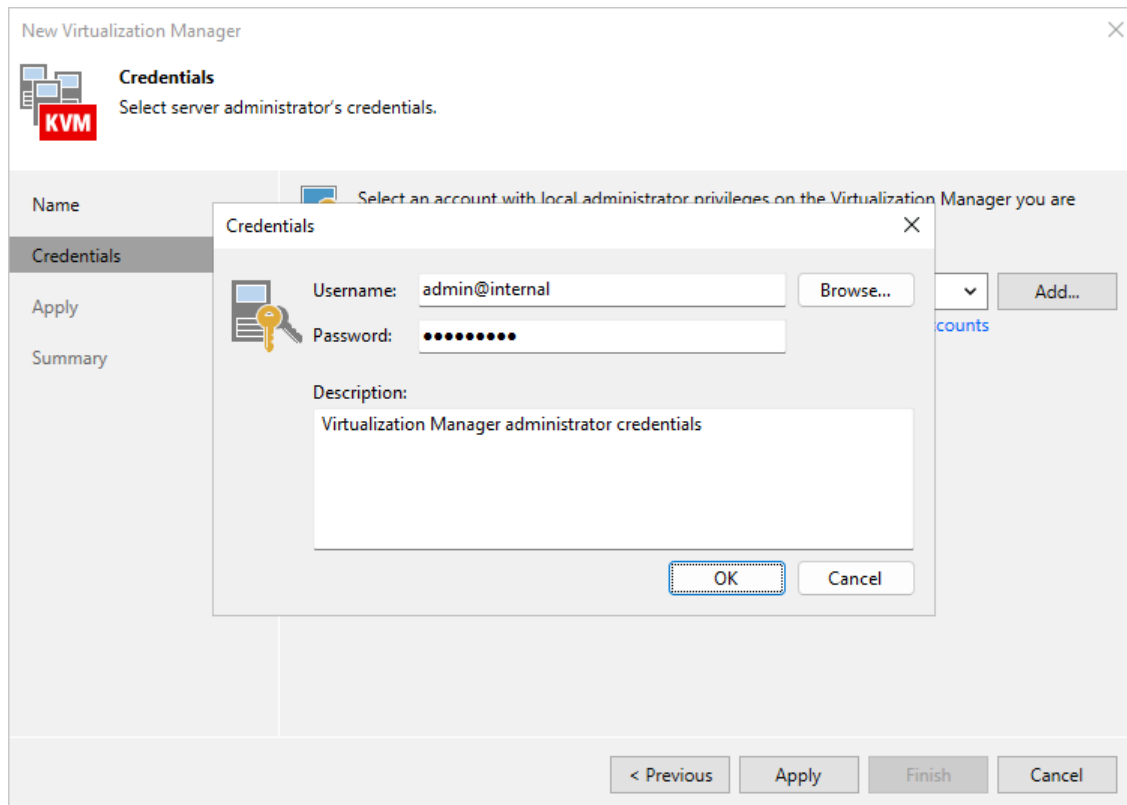
At the **Credentials** step of the wizard, specify credentials for an administrator account with the *SuperUser* role that is used to access the Virtualization manager. For more information on oVirt system administrator roles, see [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager 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 Virtualization Manager** wizard. To add an account, do the following:

1. Click **Add**.
2. In the **Credentials** window, do the following:
 - a. In the **Username** field, enter the name of a user account with administrative privileges and the name of the user domain in the following format:
<username>@<local user domain>, for example, admin@internal.

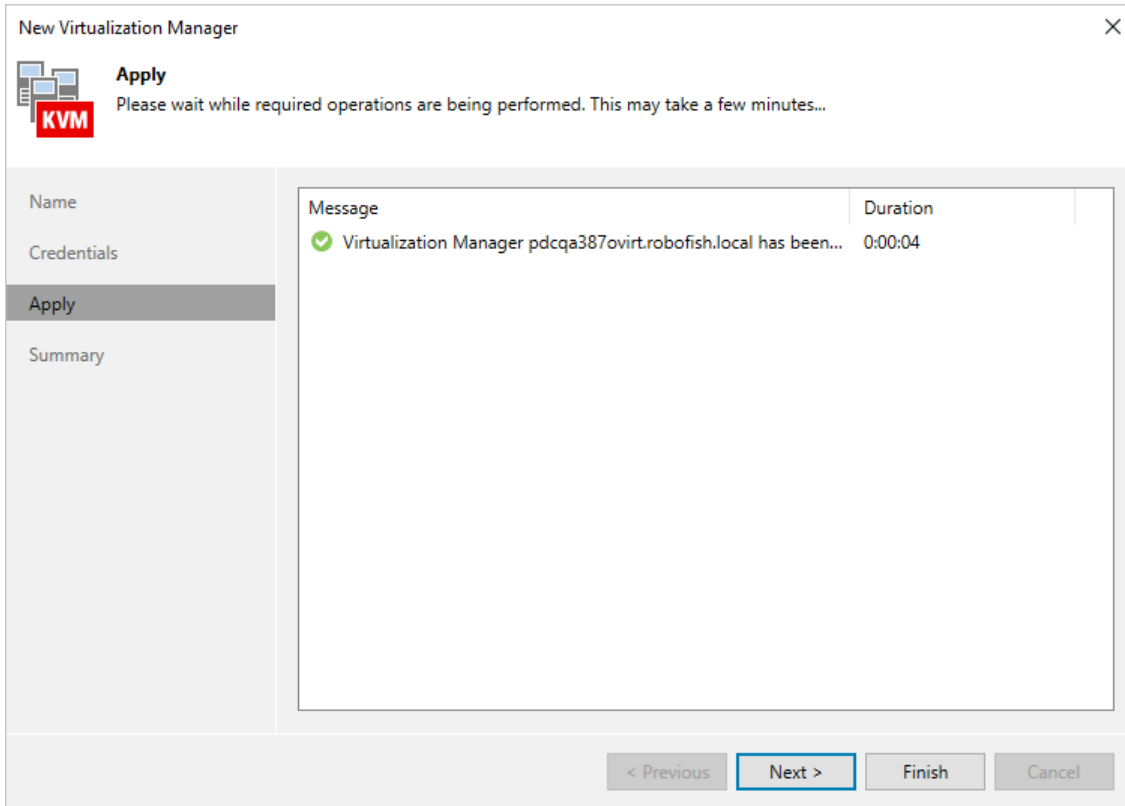
For more information on oVirt user domains, see [Red Hat Product Documentation](#) or [Oracle Linux Virtualization Manager documentation](#).
 - b. In the **Password** field, enter the password for the account.
3. Click **OK**.

The backup server will connect to the Virtualization manager and check its TLS certificate. If the certificate is not trusted on the backup server, 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 Virtualization manager using the certificate, click **Continue**.



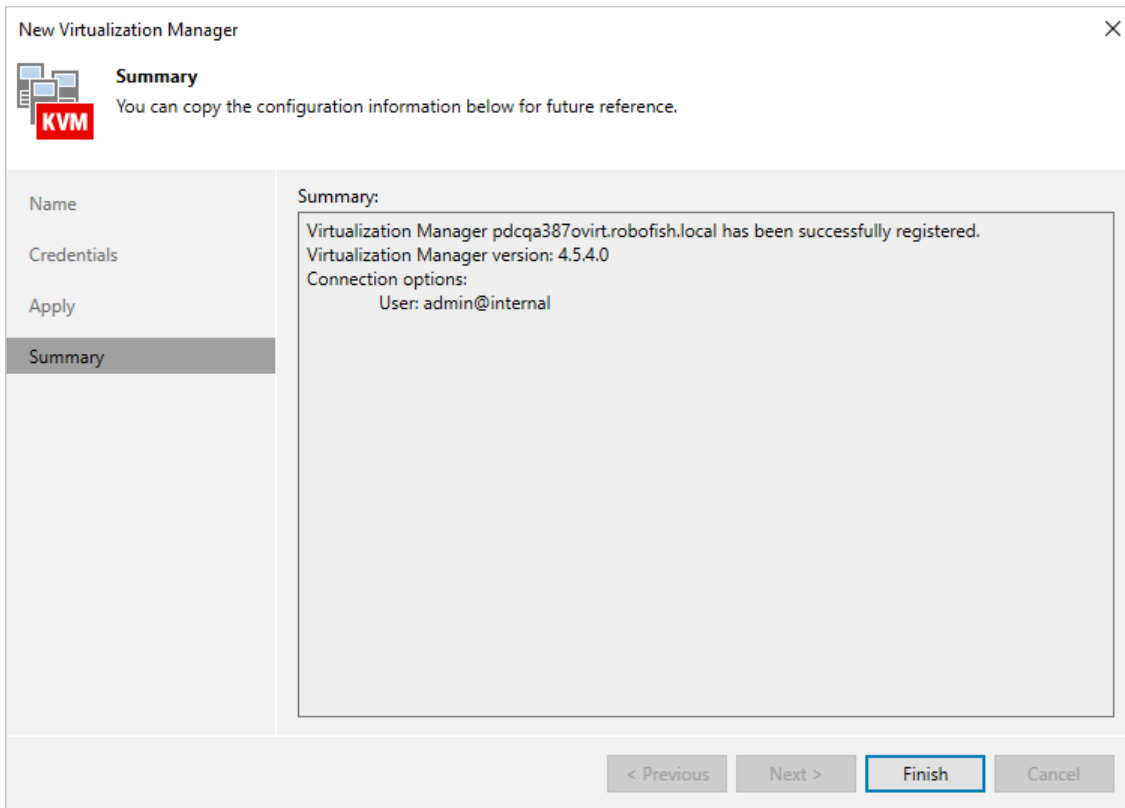
Step 4. Apply Settings

At the **Apply** step of the wizard, wait until the Virtualization manager 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 Virtualization manager has been successfully added and click **Finish**.



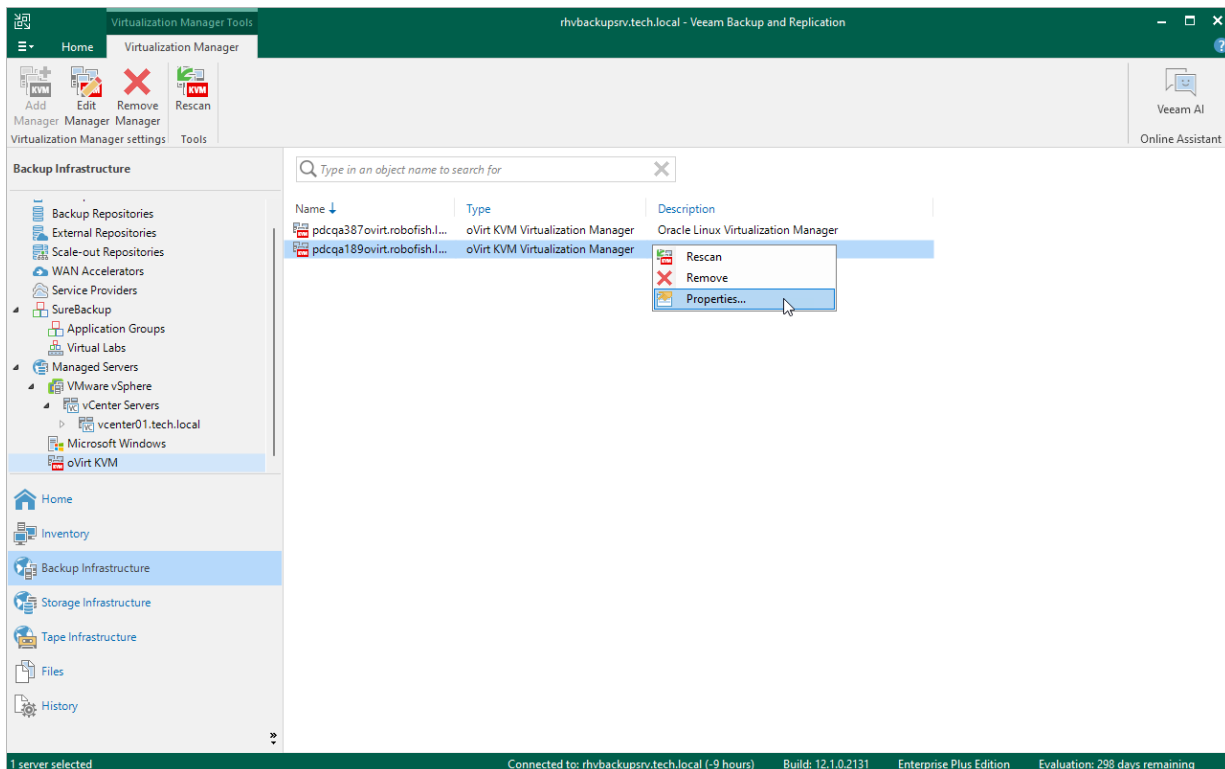
TIP

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

Editing Virtualization Manager Properties

To edit properties of the Virtualization manager added to the backup infrastructure, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > oVirt KVM**.
3. In the working area, select the Virtualization manager and click **Edit Manager** on the ribbon, or right-click the Virtualization manager and select **Properties**.
4. Complete the **Edit Virtualization Manager** wizard as described in section [Adding Virtualization Manager to Backup Infrastructure](#).

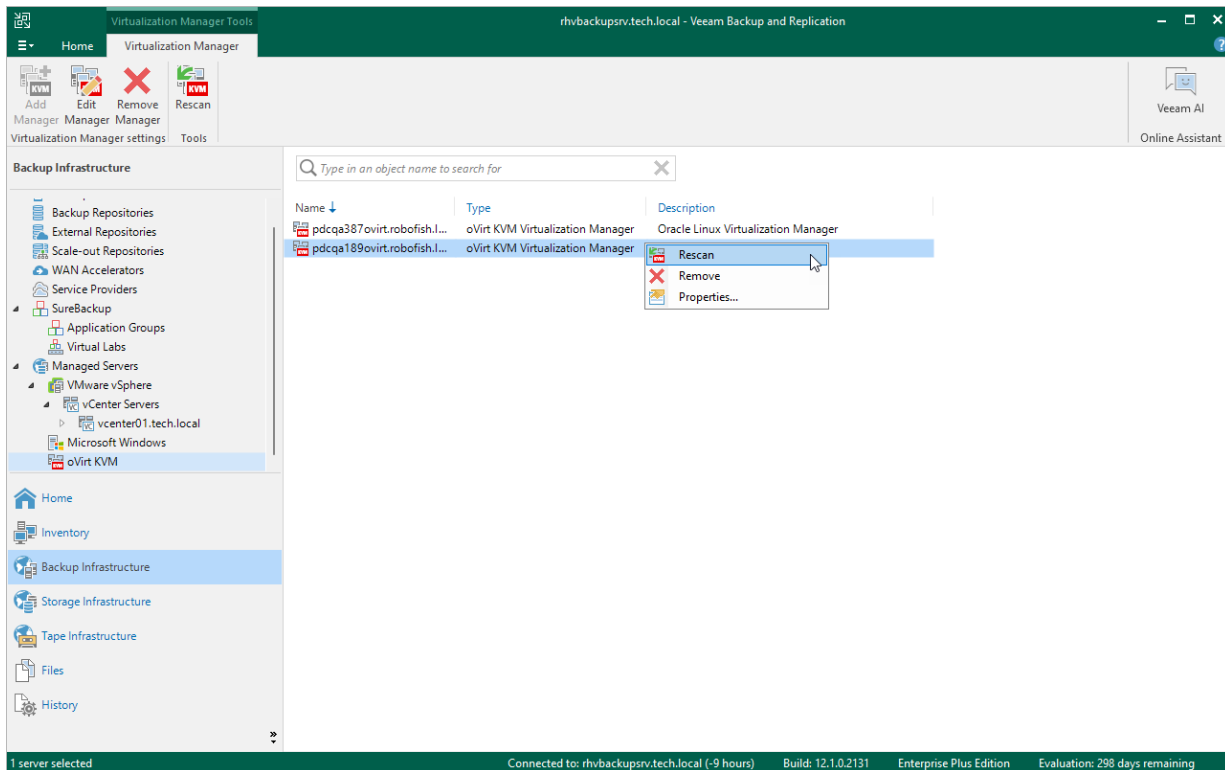


Rescanning Virtualization Manager

Veeam Backup for OLVM and RHV retrieves information about the oVirt KVM environment from the Virtualization manager. However, the data synchronization process may take some time to complete. If you make any changes to the oVirt KVM environment and want the Veeam Backup & Replication console to display the changes immediately, you can rescan the Virtualization manager manually.

To rescan the Virtualization manager, do the following:

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



Removing Virtualization Manager

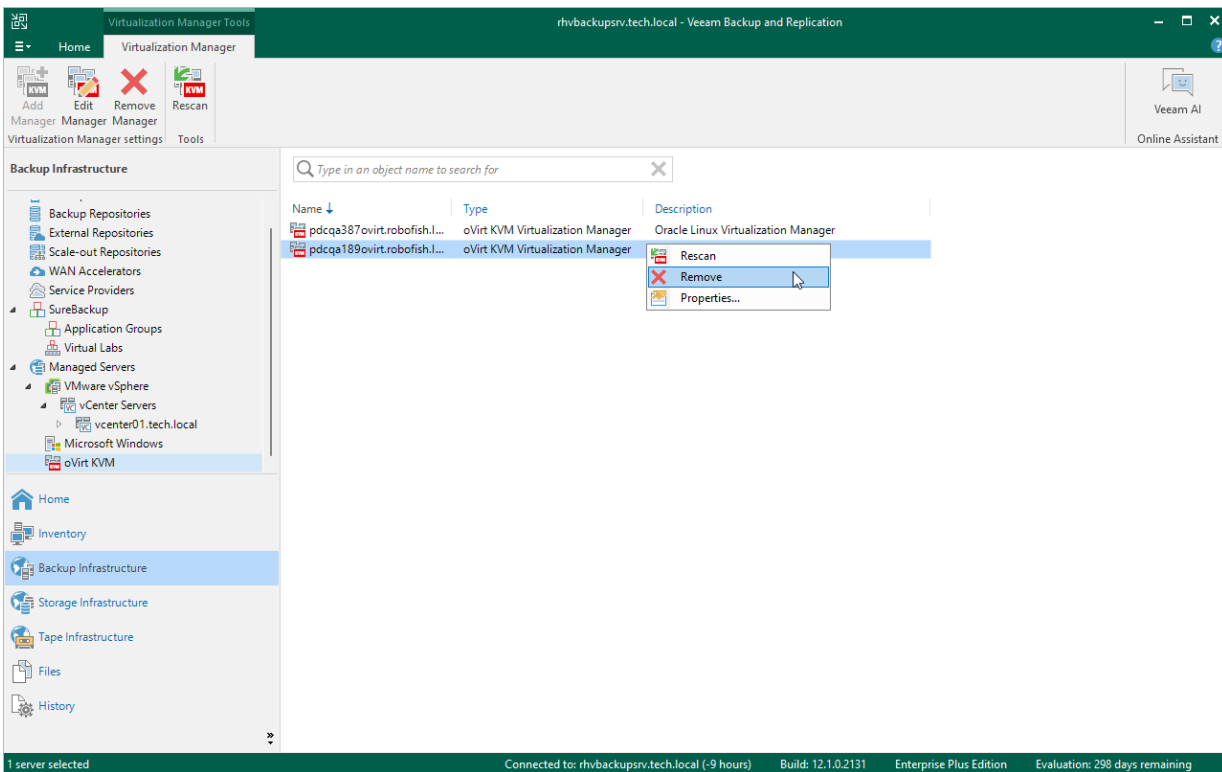
If you do not want to protect resources managed by the connected Virtualization manager anymore, you can remove it from the backup infrastructure.

IMPORTANT

Before you remove the Virtualization manager, you must [remove the backup appliance](#) that processes protection jobs for the oVirt resources managed by the Virtualization manager.

To remove the Virtualization manager from the backup infrastructure:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Managed Servers > oVirt KVM**.
3. In the working area, select the Virtualization manager and click **Remove Manager** on the ribbon, or right-click the Virtualization manager and select **Remove**.



Managing Backup Appliance

To be able to back up VMs residing on hosts that are managed by the Virtualization manager, you must add to the backup infrastructure an backup appliance that will process backup jobs and deliver backup traffic to backup repositories.

To add an backup appliance, you can either deploy a new backup appliance or connect an existing one. Note that you can add only one backup appliance for each Virtualization manager.

Deploying New Backup Appliance

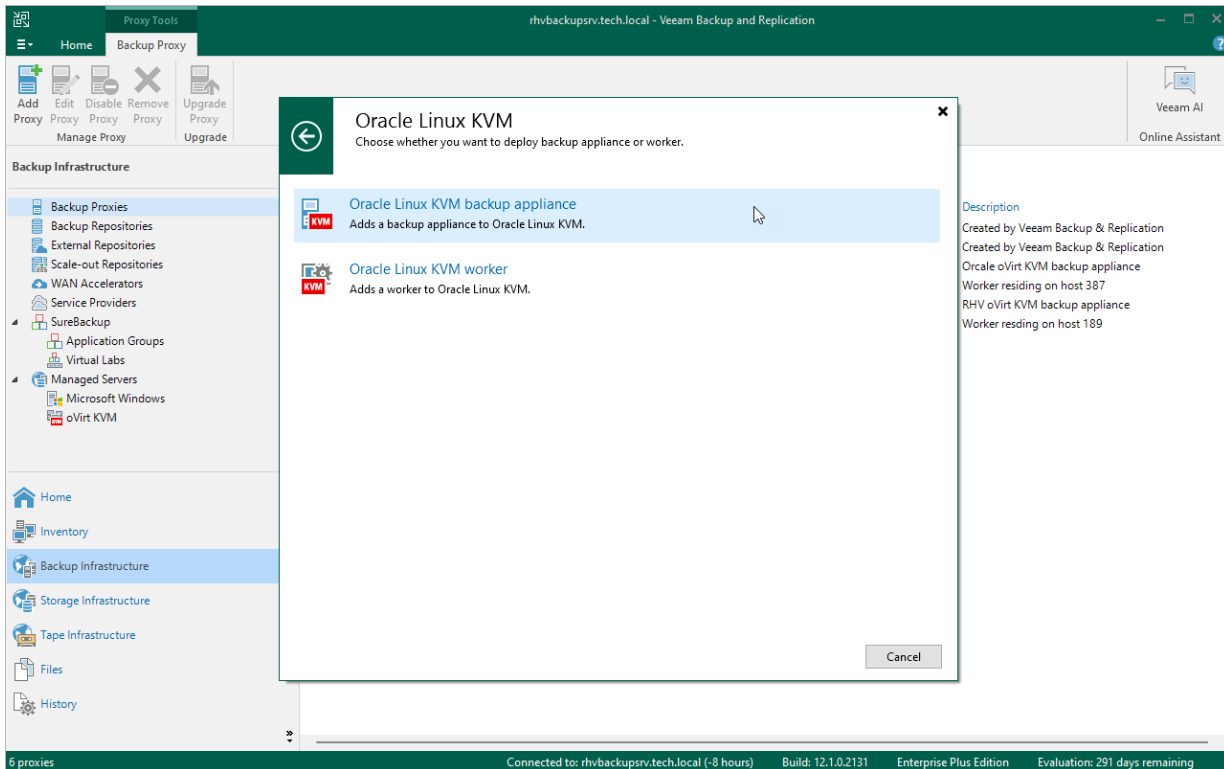
To deploy an backup appliance and to add it to the backup infrastructure, do the following:

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

Step 1. Launch New oVirt KVM Proxy Wizard

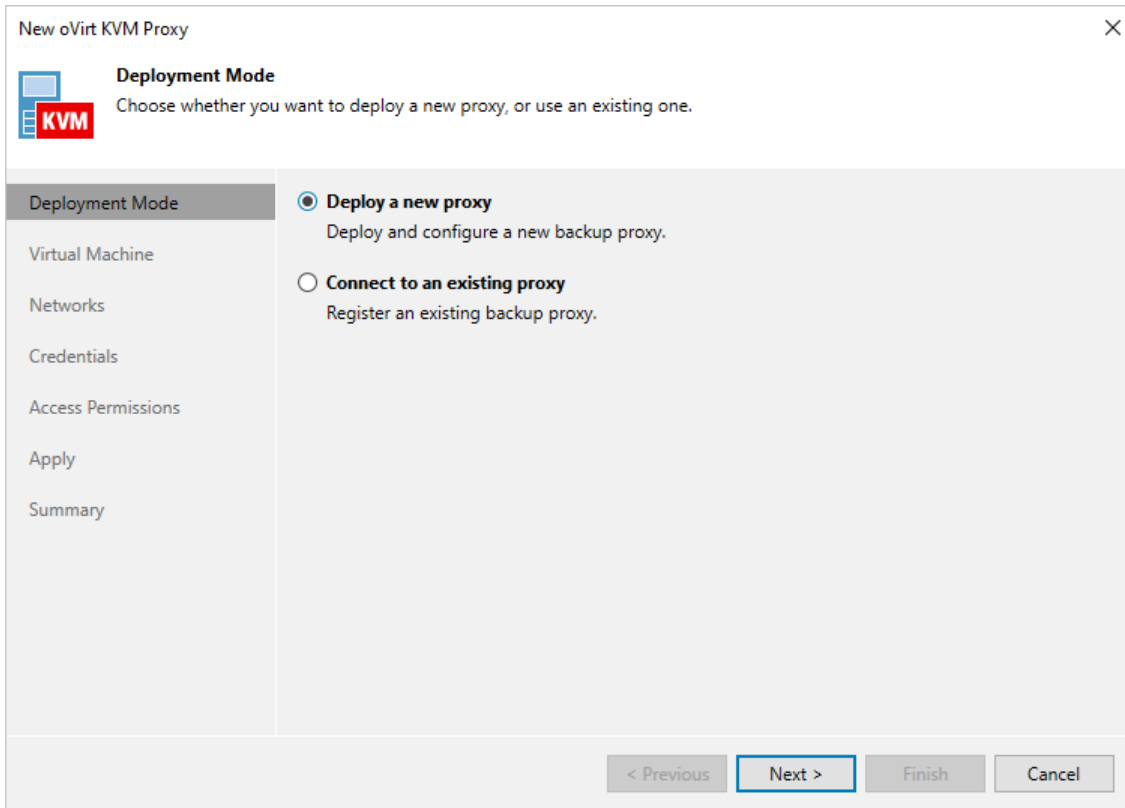
To launch the **New oVirt KVM Proxy** wizard, do the following:

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. Choose the **Red Hat Virtualization** or **Oracle Linux KVM** platform and select the option to deploy a backup appliance.



Step 2. Select Deployment Mode

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



Step 3. Specify VM Configuration

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

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

For a cluster to be displayed in the list of the available clusters, it must be added to the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

2. In the **Name** field, specify a name for the backup appliance.
3. Click **Choose** next to the **Storage Domain** field, and specify a storage domain where backup appliance system files will be stored in the **Select Storage Domain** window.

For a domain to be displayed in the list of the available domains, it must be configured in the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

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

The default number of concurrent tasks is set to 4. When you change this value, the wizard automatically adjusts the amount of resources that will be allocated to the VM running as the backup appliance. If you want to specify the amount of resources manually, click **Advanced**. In the advanced settings, you can also enable warnings that will be added to backup job sessions when CPU or RAM consumption breaks the thresholds you specify.

The screenshot shows the 'New oVirt KVM Proxy' wizard window. The 'Virtual Machine' step is active, with a sidebar on the left containing options: Deployment Mode, Virtual Machine (selected), Networks, Credentials, Access Permissions, Apply, and Summary. The main area contains the following fields and controls:

- Cluster:** A text box containing 'Default' and a 'Choose...' button.
- Name:** A text box containing 'backup-appliance-oracle'.
- Storage Domain:** A text box containing 'hosted_storage' and a 'Choose...' button.
- Proxy description:** A text box containing 'Oracle oVirt KVM backup appliance'.
- Max concurrent tasks:** A spinner box set to '4'.
- Advanced dialog box:** A smaller window with the following settings:
 - Number of vCPU cores: 4
 - Memory size (GB): 4
 - Warn me when free CPU is below 5 %
 - Warn me when free RAM is below 5 %

At the bottom of the wizard, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 4. Specify Network Settings

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

1. Click **Browse** to select a network adapter to which the backup appliance will be connected.

For a network to be displayed in the list of the available networks, it must be configured in the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

2. In the **Hostname** field, specify a hostname (without domain name) that will be assigned to the backup appliance.

The maximum length of the hostname is 64 characters. The hyphen-minus character (-) is supported, but you cannot use it as the first or the last character of the name.

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

If DHCP is disabled for the selected network adapter, or you want to specify an IP address and configure DNS settings manually, click **Configure** and do the following in the **Network settings** window:

- To specify an IP address, select the **Use the following IP address** option and enter the backup appliance IP address, subnet mask and default gateway.
- To configure DNS settings, select the **Use the following DNS server address** option and enter the IP addresses of the preferred and alternate DNS servers.

4. To check for available package updates for the backup appliance and workers, Veeam Backup for OLVM and RHV automatically connects to Veeam repositories over the internet. If the backup appliance and workers are not connected to the internet, you can instruct Veeam Backup for OLVM and RHV to use an HTTP proxy that will provide access to the required resources. To specify HTTP proxy settings, click **Configure** and do the following in the **Network settings** window:

- a. Navigate to the **HTTP proxy** tab.
- b. Select the **Use the following internet proxy settings** check box.
- c. In the **Host** field, enter the IP address or FQDN of the web proxy.
- d. In the **Port** field, enter the port used on the web proxy for HTTP or HTTPS connections.
- e. [Applies only if the HTTP proxy requires authentication] Select the **Use authentication** check box and enter credentials of the account configured on the HTTP proxy to access the internet.

5. To enable SSH access for the purposes of manual management and troubleshooting, select the **Enable SSH server** check box.

You must also select this check box if you want to use the backup appliance as a gateway server that will forward oVirt VM backups to an object storage repository and will process protection tasks related to all backups stored in that repository.

6. If the backup appliance and workers do not have access to the internet and no HTTP proxy is used, clear the **Obtain updates** check box to disable automatic updates. This will help eliminate update failures and session warnings.

The screenshot shows the 'New oVirt KVM Proxy' configuration window, specifically the 'Networks' tab. The window title is 'New oVirt KVM Proxy' with a close button (X) in the top right corner. The 'Networks' tab is selected, and the sub-header is 'Specify network settings for oVirt KVM backup proxy.' The left sidebar contains navigation options: Deployment Mode, Virtual Machine, Networks (selected), Credentials, Access Permissions, Apply, and Summary. The main configuration area includes: Network: 'ovirtmgmt' with a 'Browse...' button; Hostname: 'backup-appliance-oracle'; IP address: 'Obtain automatically'; DNS server: 'Obtain automatically' with a 'Configure...' button; and two checked checkboxes: 'Enable SSH server' and 'Obtain updates', both with warning icons. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

Step 5. Specify Credentials

At the **Credentials** step of the wizard, select credentials for an account that will be created to access the backup appliance.

IMPORTANT

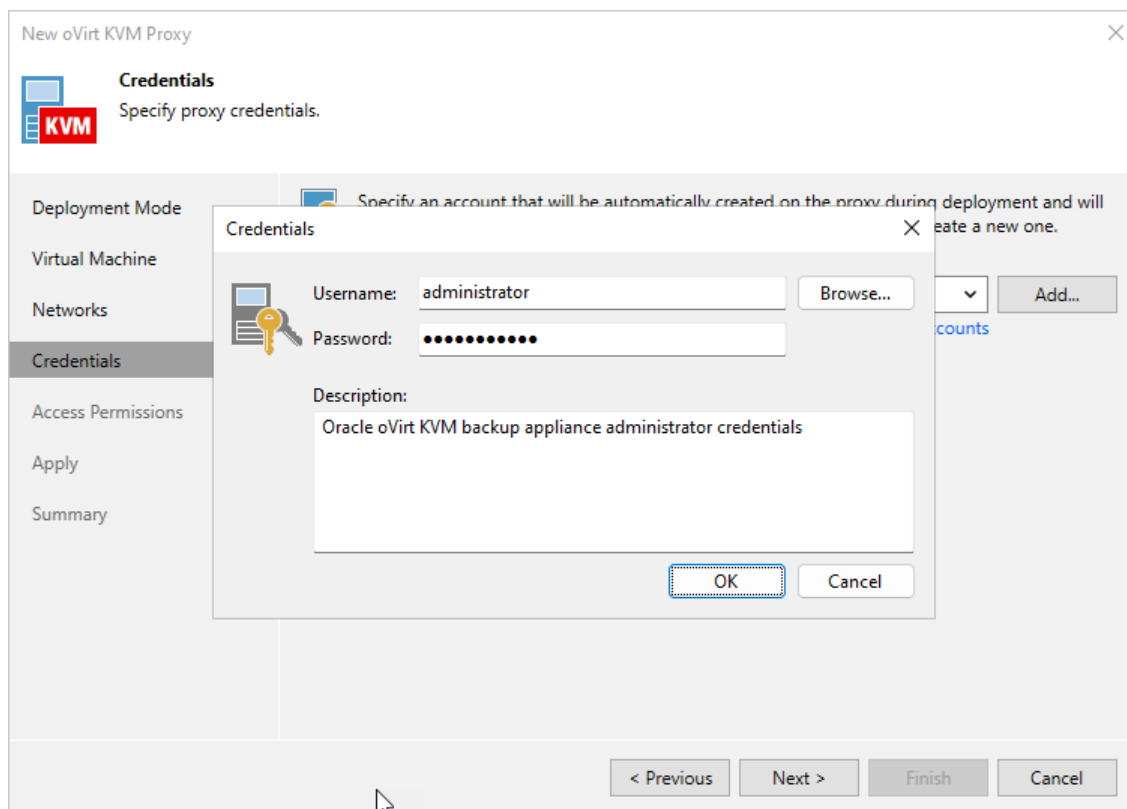
Do not select Active Directory accounts – the backup appliance does not support LDAP integration.

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

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

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

3. Click **OK**.



Step 6. Grant Permissions

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

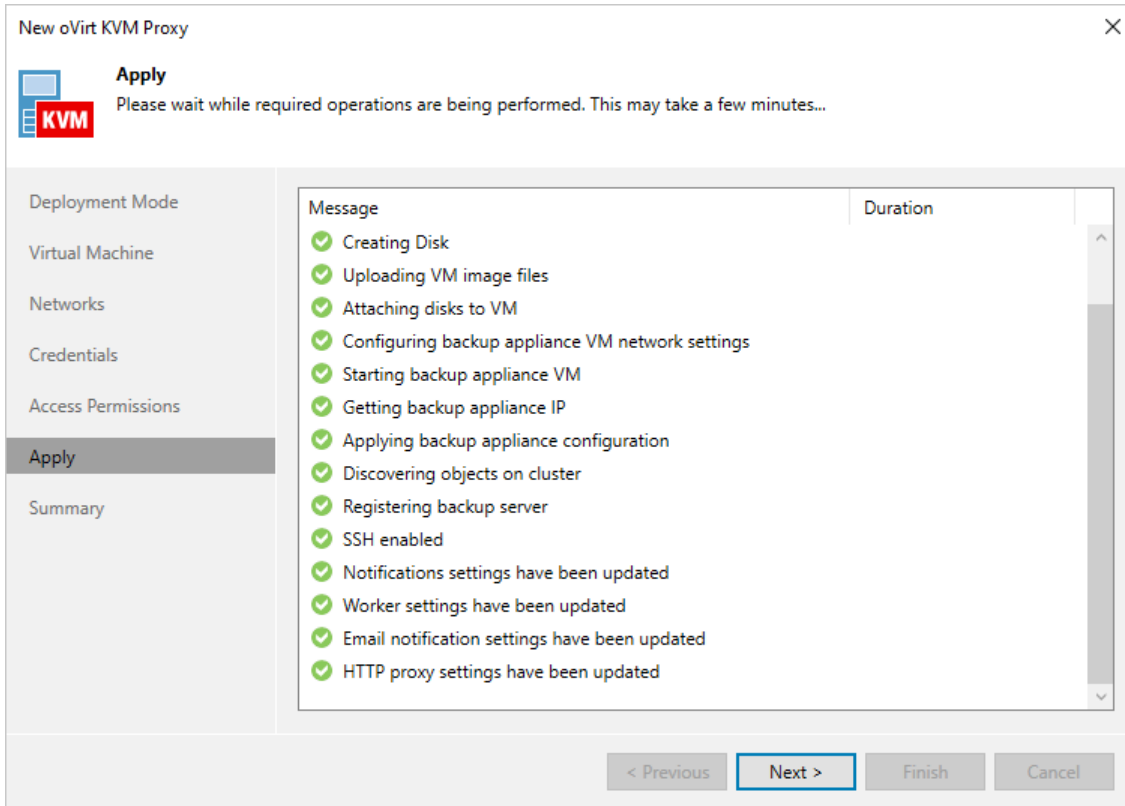
- Select the **Allow access to all backup repositories** option if you want the backup appliance to have access to all backup repositories added to the backup infrastructure.
- Select the **Allow access to the following backup repositories** option if you want the backup appliance to have access to specific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

If you select the **Allow access to the following backup repositories** option, you must also specify backup repositories to which the backup appliance will have access. For a backup repository to be displayed in the **Repository** list, it must be added to the backup infrastructure.

The screenshot shows a window titled "New oVirt KVM Proxy" with a close button (X) in the top right corner. The window contains a sidebar on the left with the following items: "Deployment Mode", "Virtual Machine", "Networks", "Credentials", "Access Permissions" (highlighted), "Apply", and "Summary". The main area is titled "Access Permissions" and includes a sub-header "Specify the backup repositories this oVirt KVM proxy is allowed to access." Below this, there are two radio button options: "Allow access to all backup repositories" (which is selected) and "Allow access to the following backup repositories". To the right of these options are two buttons: "Select All" and "Clear All". Below the options is a list box labeled "Repository" containing one entry: "Default Backup Repository" with an unchecked checkbox. At the bottom of the window, there are four buttons: "< Previous", "Apply" (highlighted with a blue border), "Finish", and "Cancel".

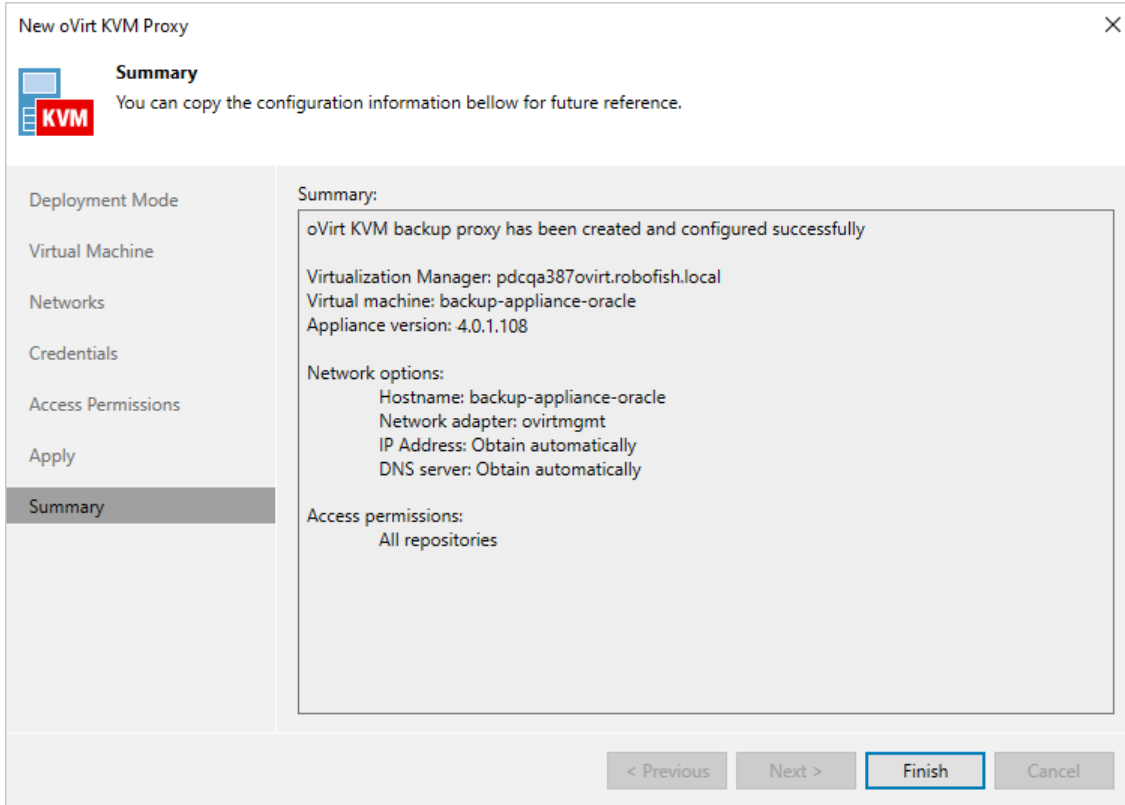
Step 7. Apply Settings

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



Step 8. Finish Working with Wizard

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



Connecting Existing Backup Appliance

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

- To upgrade an backup appliance from version 2.0, 2a, 3.0, 3a, 3b or 4.0 to 4.1.
- To connect an backup appliance that was previously connected to another backup server.

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

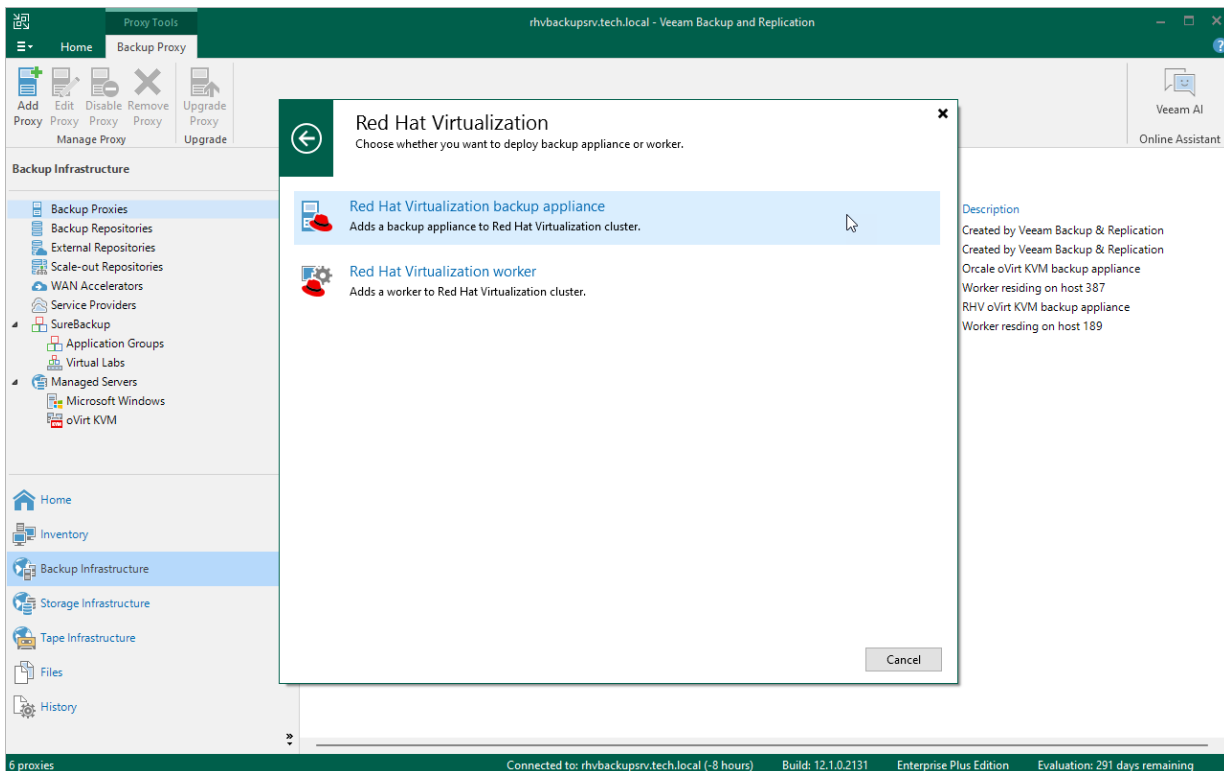
1. [Launch the New oVirt KVM Proxy wizard.](#)
2. [Select the deployment mode.](#)
3. [Specify appliance VM configuration.](#)
4. [Check network settings.](#)
5. [Enter credentials for the appliance account.](#)
6. [Grant permissions to the appliance.](#)
7. [Apply appliance settings.](#)
8. [Finish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

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

Step 1. Launch New oVirt KVM Proxy Wizard

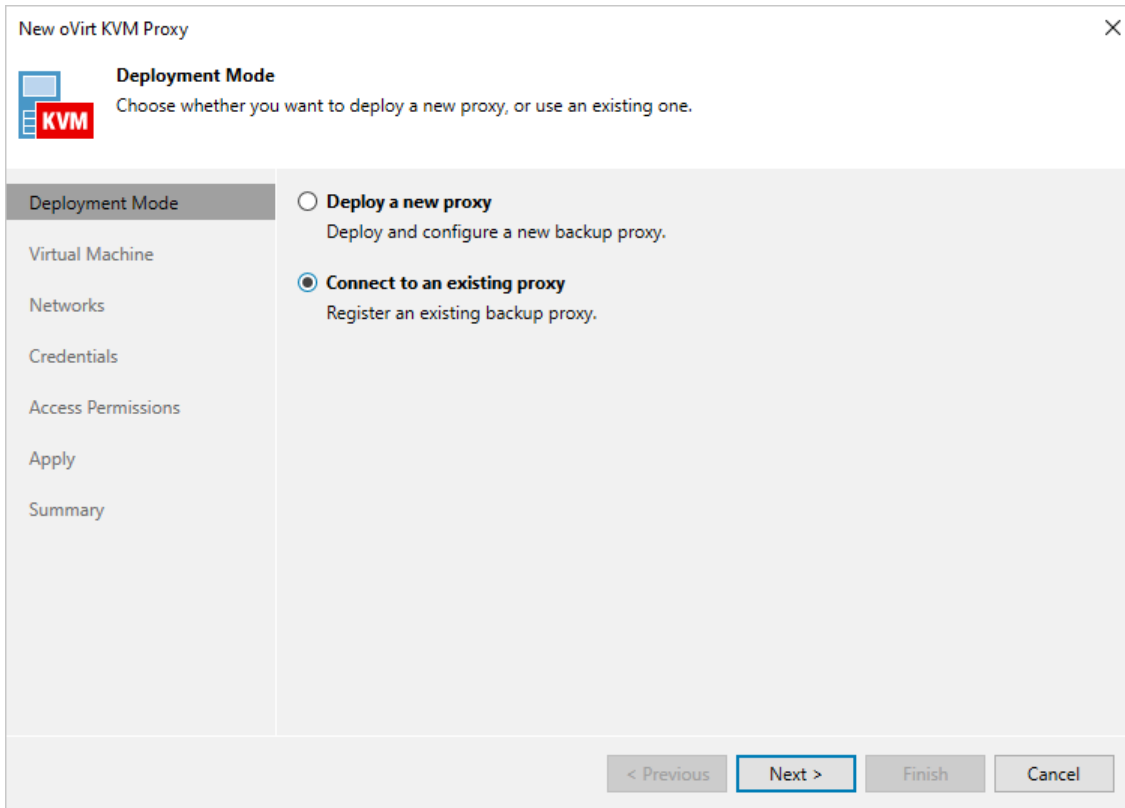
To launch the **New oVirt Proxy** wizard, do the following:

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. On the ribbon, select **Add Proxy**.
5. Choose the **Red Hat Virtualization** or **Oracle Virtualization** platform and select the option to deploy a backup appliance.



Step 2. Select Deployment Mode

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



Step 3. Specify VM Configuration

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

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

For a cluster to be displayed in the list of the available clusters, it must be added to the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

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

NOTE

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

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

The default number of concurrent tasks is set to 4. When you change this value, the wizard automatically adjusts the amount of resources that will be allocated to the VM running as the backup appliance. If you want to specify the amount of resources manually, click **Advanced**. In the advanced settings, you can also enable warnings that will be added to backup job sessions when CPU or RAM consumption breaks the thresholds you specify.

The screenshot shows the 'New oVirt KVM Proxy' wizard window. The 'Virtual Machine' step is active. The configuration fields are as follows:

- Cluster:** Default (Choose... button)
- Name:** backup-appliance-rhv (Choose... button)
- Storage Domain:** hosted_storage
- Proxy description:** RHV oVirt KVM backup appliance
- Max concurrent tasks:** 4 (dropdown)

An 'Advanced' dialog box is open, showing the following settings:

- Number of vCPU cores:** 4 (dropdown)
- Memory size (GB):** 4 (dropdown)
- Warn me when free CPU is below 5 % (dropdown)
- Warn me when free RAM is below 5 % (dropdown)

Buttons at the bottom of the wizard include '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Advanced' dialog has 'OK' and 'Cancel' buttons.

Step 4. Check Network Settings

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

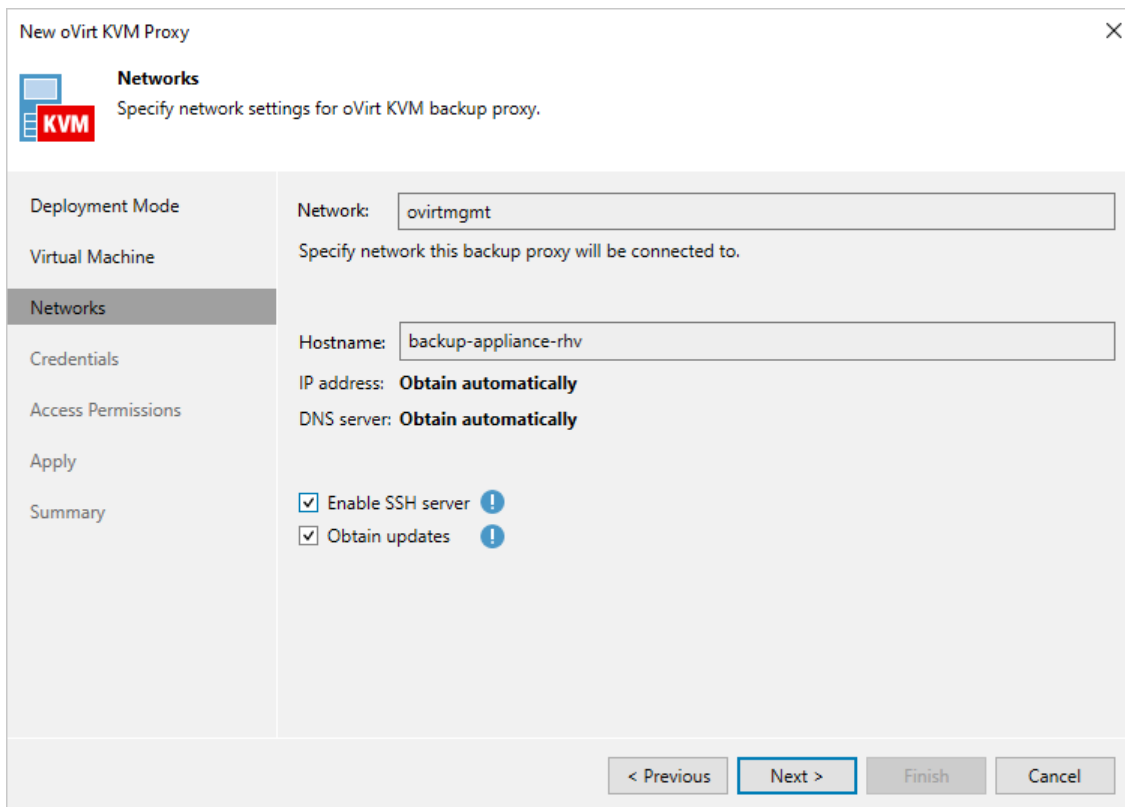
1. Review backup appliance network settings.

You will be able to [change the network settings](#) after you connect the appliance to the backup infrastructure.

2. To enable SSH access for the purposes of manual management and troubleshooting, select the **Enable SSH server** check box.

You must also select the check box if you want to use the backup appliance as a gateway server that will forward oVirt VM backups to an object storage repository and will process protection tasks related to all backups stored in that repository.

3. If the backup appliance and workers do not have access to the internet and no web proxy is used, clear the **Obtain updates** check box to disable automatic updates. This will help eliminate update failures and session warnings.



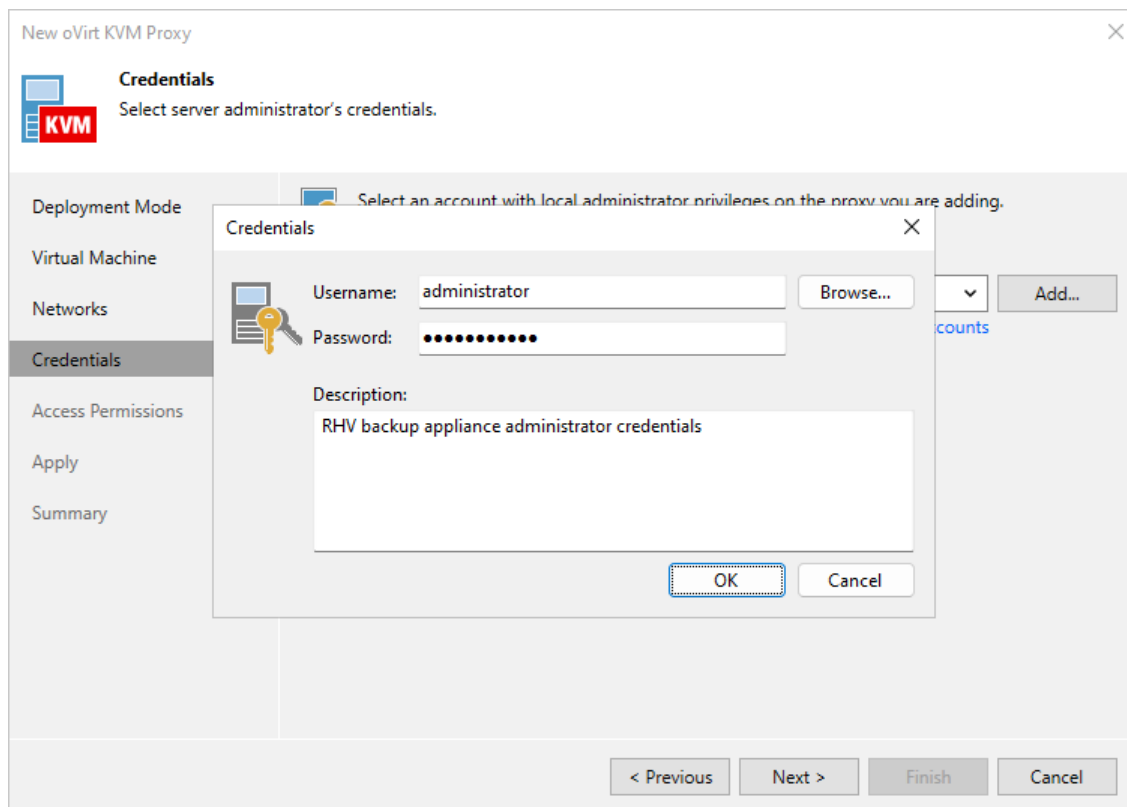
The screenshot shows the 'New oVirt KVM Proxy' wizard, specifically the 'Networks' step. The window title is 'New oVirt KVM Proxy' with a close button (X) in the top right corner. The 'Networks' step is highlighted in the left sidebar, which also includes 'Deployment Mode', 'Virtual Machine', 'Credentials', 'Access Permissions', 'Apply', and 'Summary'. The main content area is titled 'Networks' and contains the instruction 'Specify network settings for oVirt KVM backup proxy.' Below this, there are several fields and options: 'Network:' with a text box containing 'ovirtmgmt'; 'Specify network this backup proxy will be connected to.'; 'Hostname:' with a text box containing 'backup-appliance-rhv'; 'IP address:' with the option 'Obtain automatically'; and 'DNS server:' with the option 'Obtain automatically'. At the bottom of the main content area, there are two checked checkboxes: 'Enable SSH server' and 'Obtain updates', each with a warning icon (exclamation mark in a circle). At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

Step 5. Enter Credentials

At the **Credentials** step of the wizard, select credentials for the account that you are used to access the backup appliance.

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 oVirt KVM Proxy** wizard. To add credentials, do the following:

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



Step 6. Grant Permissions

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

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

If you select the **Allow access to the following backup repositories** option, you must also specify backup repositories to which the backup appliance will have access. For a backup repository to be displayed in the **Repository** list, it must be added to the backup infrastructure.

The screenshot shows a window titled "New oVirt KVM Proxy" with a close button (X) in the top right corner. The window contains a sidebar on the left with the following items: "Deployment Mode", "Virtual Machine", "Networks", "Credentials", "Access Permissions" (highlighted), "Apply", and "Summary". The main area is titled "Access Permissions" and includes a sub-header "Specify the backup repositories this oVirt KVM proxy is allowed to access." Below this, there are two radio button options: "Allow access to all backup repositories" (which is selected) and "Allow access to the following backup repositories". To the right of these options are two buttons: "Select All" and "Clear All". Below the options is a list box labeled "Repository" containing one entry: "Default Backup Repository" with an unchecked checkbox. At the bottom of the window, there are four buttons: "< Previous", "Apply" (highlighted with a blue border), "Finish", and "Cancel".

Step 7. Apply Settings

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

New oVirt KVM Proxy

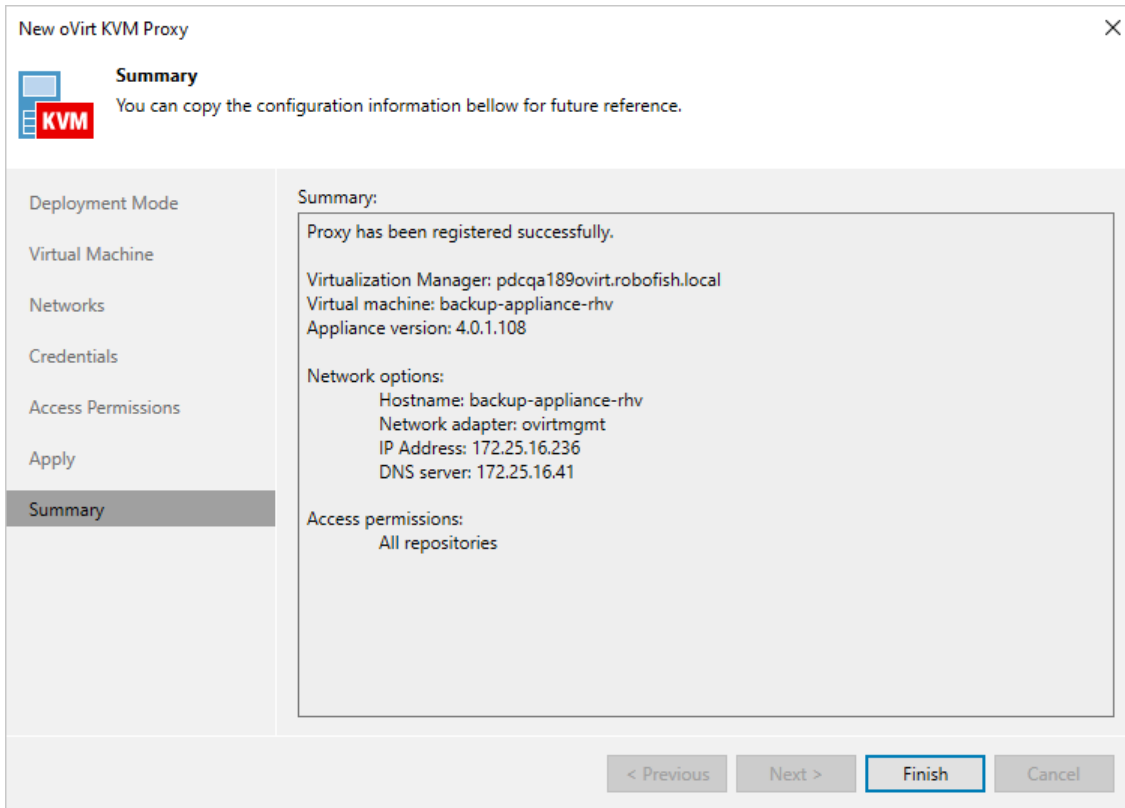
Apply
Please wait while required operations are being performed. This may take a few minutes...

Message	Duration
✓ oVirt KVM backup proxy 172.25.16.236 has been registered su...	
✓ Connecting to backup appliance	
✓ Updating cluster registration data	
✓ Registering backup server	
✓ Updating existing backup chains	
✓ Preparing	
✓ Successfully Completed	
✓ SSH enabled	
✓ Notifications settings have been updated	
✓ Worker settings have been updated	
✓ Email notification settings have been updated	0:00:06

< Previous **Next >** Finish Cancel

Step 8. Finish Working with Wizard

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



Editing Backup Appliance

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

To edit backup appliance settings, do the following:

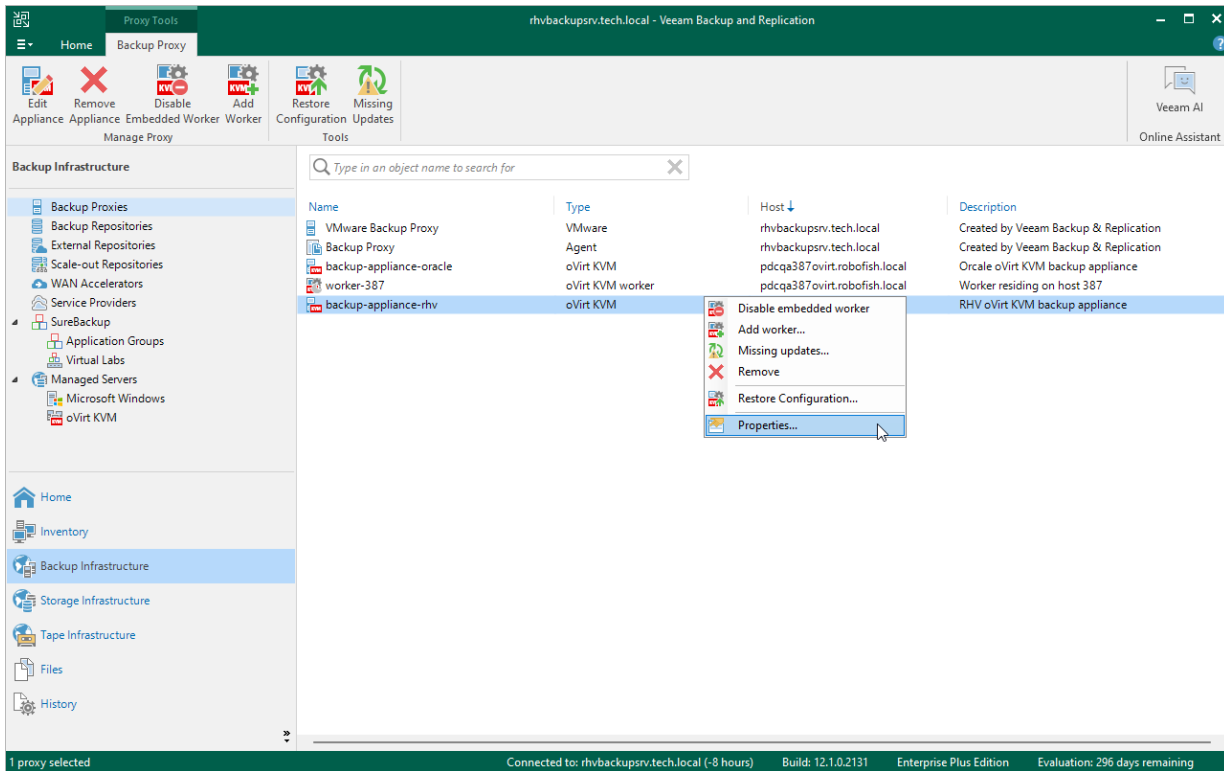
1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the backup appliance and click **Edit Proxy** on the ribbon, or right-click the backup appliance and select **Properties**.
4. Complete the **Edit oVirt KVM Proxy** wizard:
 - a. To provide a new description for the backup appliance and change the number of tasks that the embedded worker is able to handle in parallel, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 3).
 - b. To manage SSH access to the appliance, to enable or disable worker updates, to change the backup appliance network settings, or to configure an HTTP proxy for accessing Veeam update repositories, follow the instructions provided in section [Deploying New Backup Appliance](#) (step 4).
 - c. To change credentials that are used to access the backup appliance, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 5).

NOTE

The user name and password must be updated in the record of the Credential manager that is already selected at the **Credentials** step of the wizard. If you create a new record and select it, Veeam Backup for OLVM and RHV will not be able to update credentials and will show you an authorization error.

- d. To specify backup repositories the backup appliance can access, follow the instructions provided in section [Connecting Existing Backup Appliance](#) (step 6).

e. To save changes made to the appliance settings, click **Finish**.

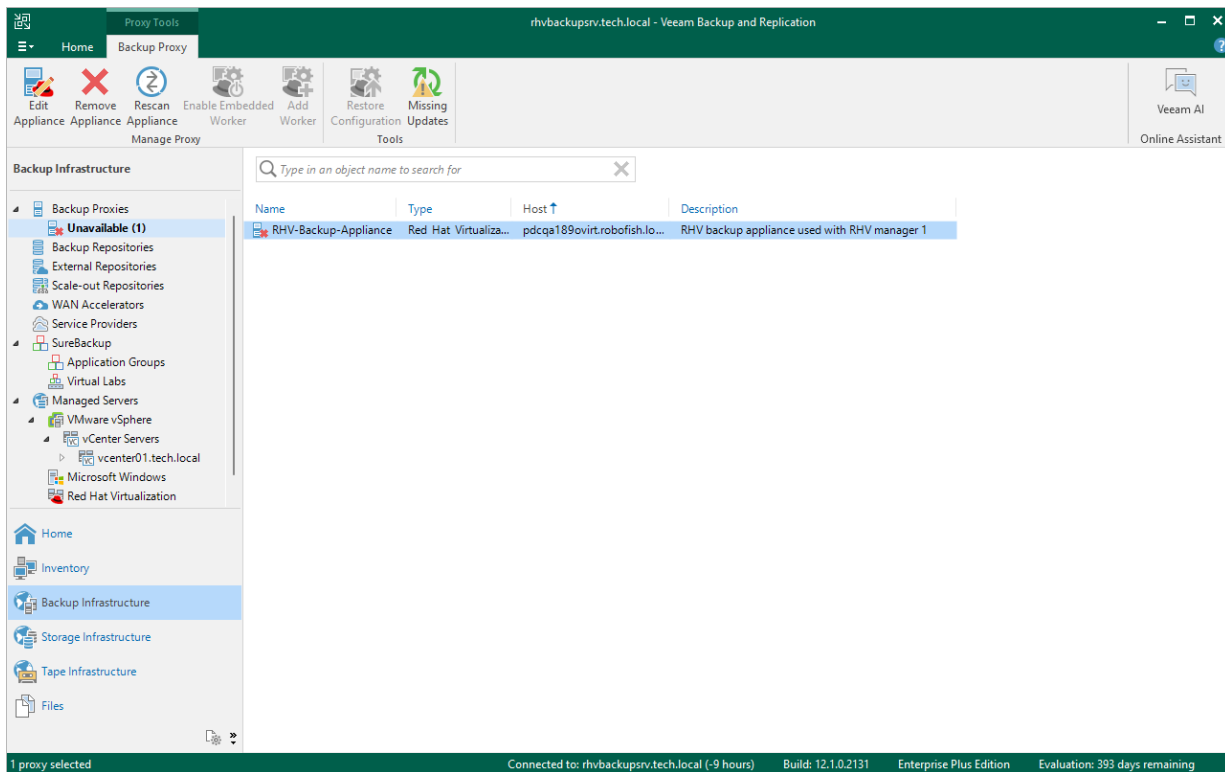


Rescanning Backup Appliance

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

To rescan the backup appliance, do the following:

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



Removing Backup Appliance

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

IMPORTANT

After you remove the backup appliance:

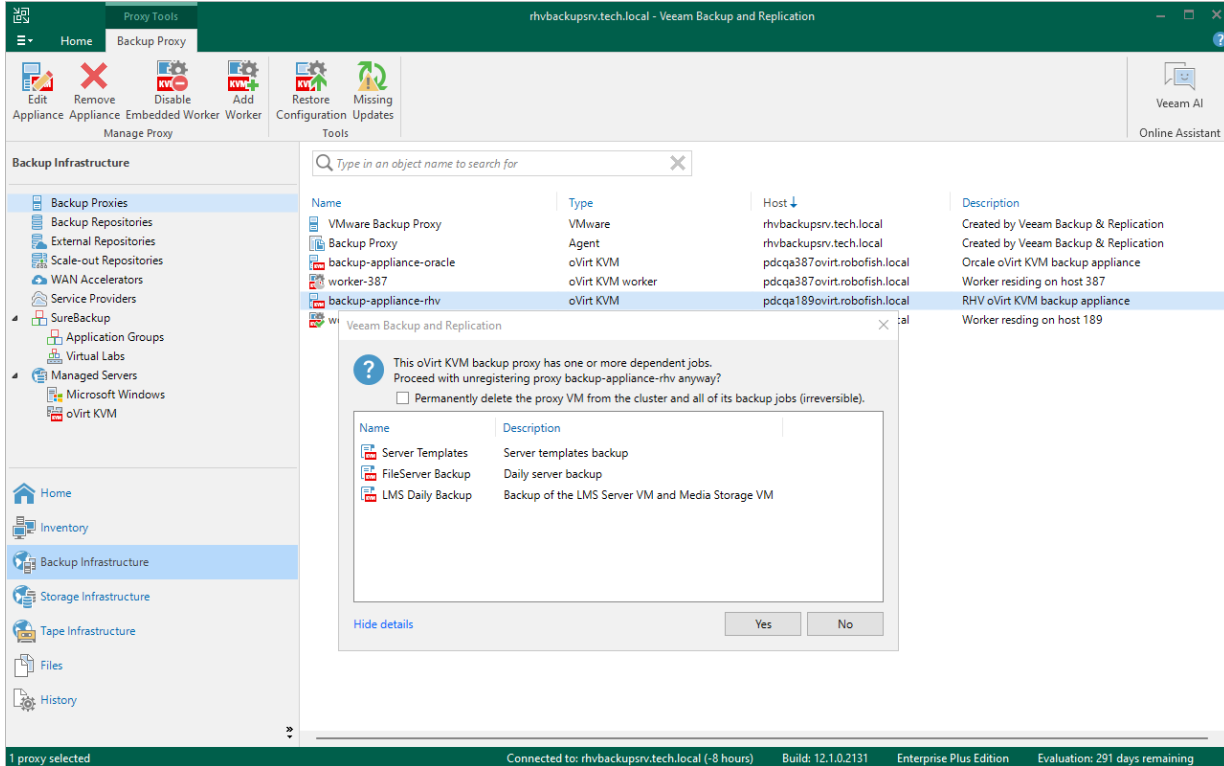
- You will not be able to perform [oVirt VM backup](#), [entire VM restore](#) and [VM disk restore](#) operations unless you deploy a new backup appliance. However, you will still be able to [manage oVirt VM backups](#) and perform all other restore operations described in section [Performing Restore](#).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

To remove the backup appliance, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the backup appliance and click **Remove Appliance** on the ribbon, or right-click the backup appliance and select **Remove**.
4. In the **Veeam Backup & Replication** window, choose whether you want to permanently remove from the host the VM running as the backup appliance.

TIP

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



Managing Workers

To perform most data protection and disaster recovery operations, Veeam Backup for OLVM and RHV uses workers. Workers are Linux-based VMs that are responsible for the interaction between the backup appliance and other Veeam Backup for OLVM and RHV components. Workers process backup workload and distribute backup traffic when transferring data to backup repositories.

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

Each dedicated worker is launched on a specific host for the duration of a backup or restore operation. While configuring the worker, you can manually select the host or instruct Veeam Backup for OLVM and RHV to choose a host automatically. Manual selection may be helpful if you want to avoid launching workers on specific hosts (for example, production ones), while automatic selection allows Veeam Backup for OLVM and RHV to optimize data transfer and to balance the load on the hosts in the cluster.

Worker Lifecycle

As soon as a backup or restore session starts, Veeam Backup for OLVM and RHV launches a worker and test its configuration. Veeam Backup for OLVM and RHV checks host affinity settings specified for the worker and chooses a host where the worker VM will run. Then, Veeam Backup for OLVM and RHV powers on the worker VM and installs system updates (if available). When the backup or restore session completes, Veeam Backup for OLVM and RHV shuts down the worker VM so that it can be used for other sessions later.

Adding Workers

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

1. [Check prerequisites and limitations.](#)
2. [Launch the New oVirt KVM worker wizard.](#)
3. [Specify worker VM configuration.](#)
4. [Specify worker network settings.](#)
5. [Apply worker settings.](#)
6. [Finish working with wizard.](#)

Before You Begin

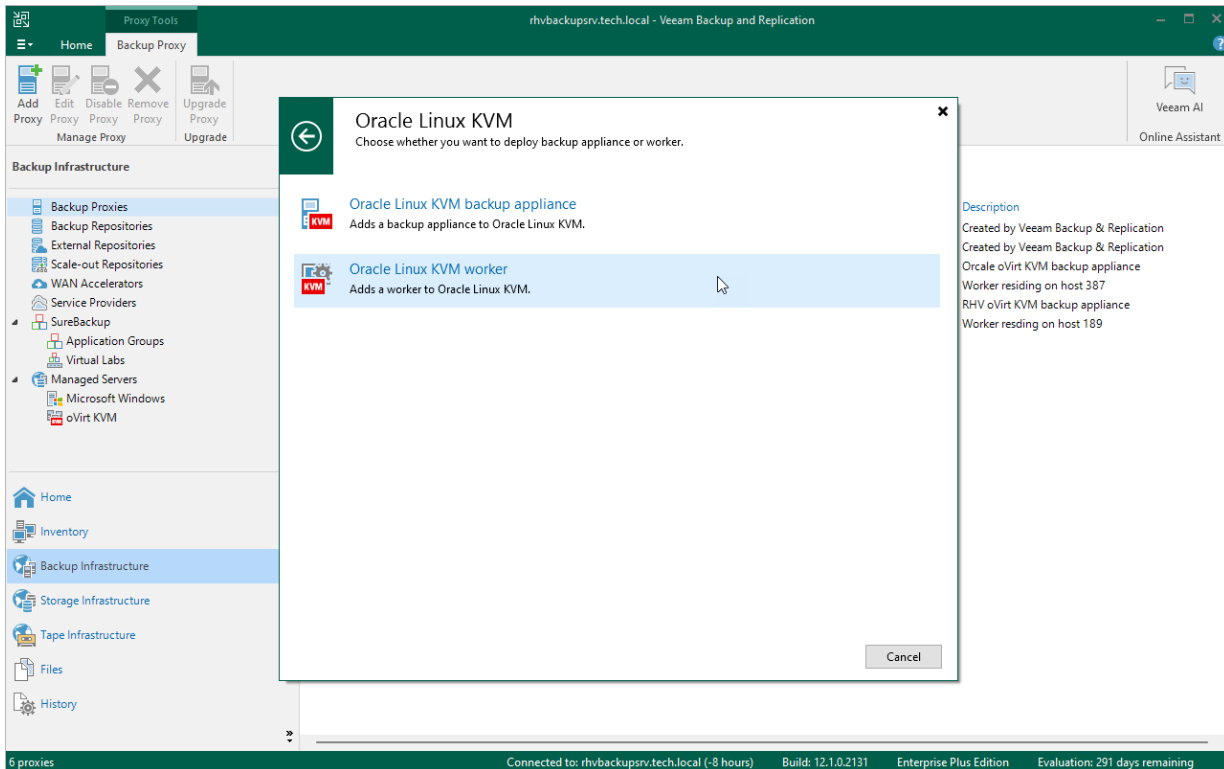
Before you add a dedicated worker to the backup infrastructure, consider the following:

- 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.
- You can change the maximum number of concurrent tasks (the best practice is to allocate 1 vCPU and 1 GB RAM for each additional task) while deploying a new worker or editing settings of an existing one.
- If you plan to use dedicated workers, you can [disable the embedded worker.](#)

Step 1. Launch New oVirt KVM Worker

To launch the **New oVirt KVM 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. Choose the **Red Hat Virtualization** or **Oracle Linux KVM** platform and select the option to deploy a worker.



Step 2. Specify Worker VM Settings

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

1. Click **Choose** next to the **Cluster** field, and specify a host where the worker will be launched in the **Select Cluster or Host** window. If you select the whole cluster, Veeam Backup for OLVM and RHV will automatically define the host to launch the worker.

For a cluster to be displayed in the list of the available clusters, it must be added to the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

2. In the **Name** field, specify a name for the worker.

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

3. Click **Choose** next to the **Storage Domain** field, and specify a storage domain where worker system files will be stored in the **Select Storage Domain** window.

For a domain to be displayed in the list of the available domains, it must be configured in the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

4. In the **Worker description** field, provide a description for future reference. The field already contains a default description with information about the user who added the worker, date and time when the worker was added.

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 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 VM. If you want to specify the amount of resources manually, click **Advanced**.

New oVirt KVM Worker

Virtual Machine
Configure a virtual machine for oVirt KVM worker placement.

Virtual Machine
Networks
Apply
Summary

Cluster or host: !
Default Choose...

Name:
worker-387

Storage Domain:
hosted_storage Choose...

Description:
Worker residing on host 387

Max concurrent tasks:
4

Advanced

Number of vCPU cores: 4

Memory size (GB): 4

OK Cancel

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

< Previous Next > Finish Cancel

Step 3. Configure Network Settings

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

1. Click **Browse** to select a network adapter to which the worker will be connected.

For a network to be displayed in the list of the available networks, it must be configured in the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

2. If DHCP is enabled for the selected network adapter, the IP address and DNS settings of the worker can be obtained automatically.

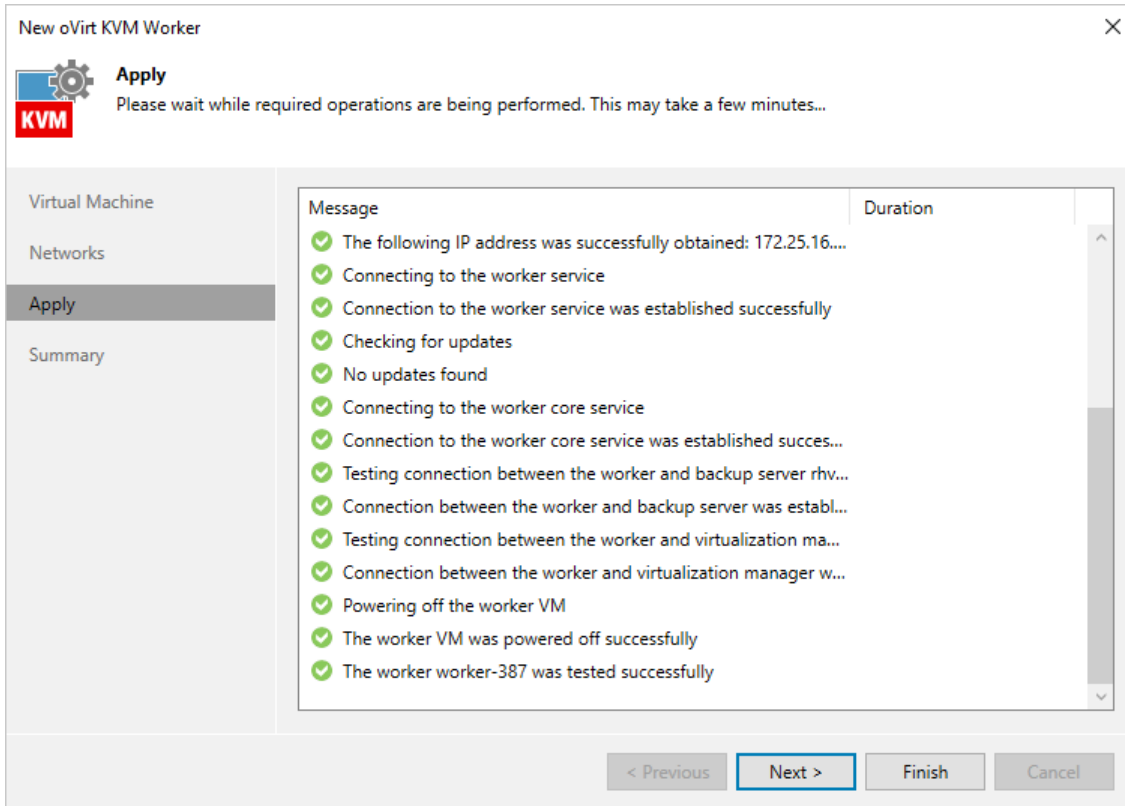
If DHCP is disabled for the selected network adapter, or you want to specify an IP address and configure DNS settings manually, click **Configure** and do the following in the **Network settings** window:

- To specify an IP address, select the **Use the following IP address** option and enter the worker IP address, subnet mask and default gateway.
- To configure DNS settings, select the **Use the following DNS server address** option and enter the IP addresses of the preferred and alternate DNS servers.

The screenshot shows the 'New oVirt KVM Worker' wizard window, specifically the 'Networks' step. The window title is 'New oVirt KVM Worker' with a close button (X) in the top right corner. Below the title bar, there is a gear icon and the text 'Networks' and 'Specify network settings for oVirt KVM worker.' A sidebar on the left contains four items: 'Virtual Machine', 'Networks' (which is highlighted), 'Apply', and 'Summary'. The main area of the window is divided into two sections. The top section is labeled 'Network:' and contains a text input field with the value 'ovirtmgmt' and a 'Browse...' button to its right. Below this, the text 'Specify network this worker will be connected to.' is displayed. The bottom section is labeled 'IP address:' and 'DNS server:', both with the value 'Obtain automatically'. To the right of these labels is a 'Configure...' button. At the bottom of the window, there are four buttons: '< Previous', 'Apply', 'Finish', and 'Cancel'.

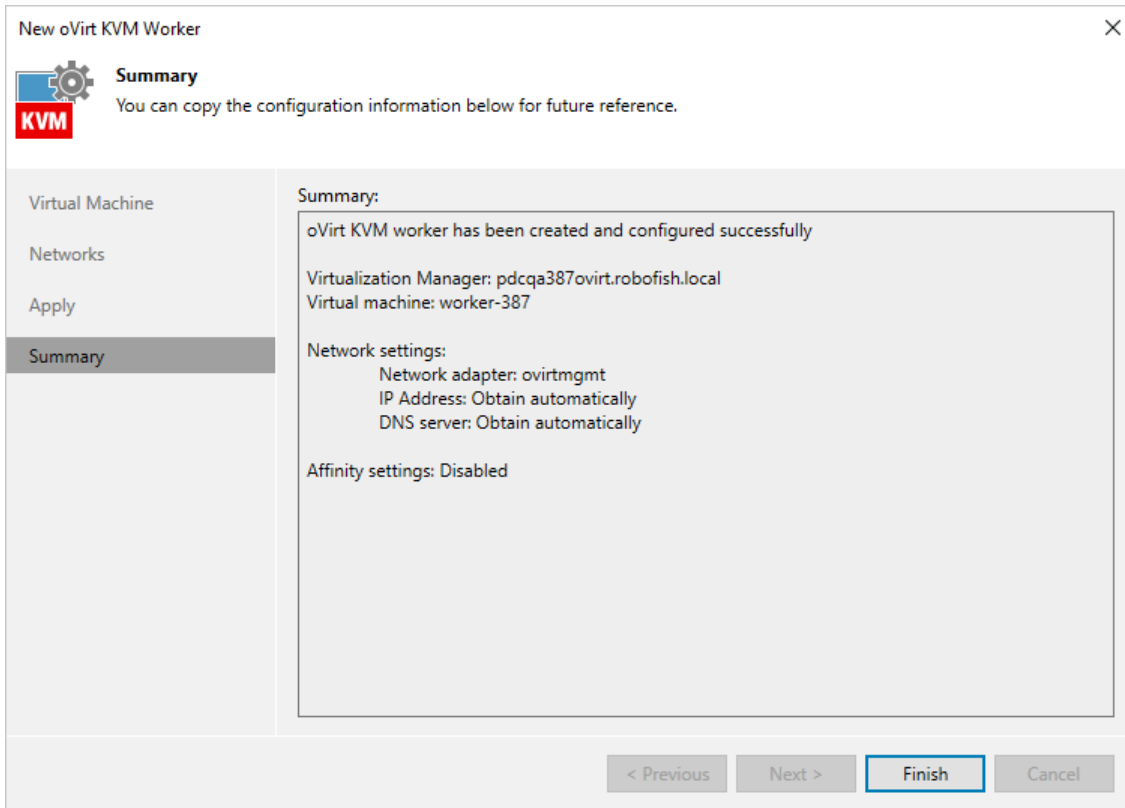
Step 4. Apply Worker Settings

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



Step 5. Finish Working with Wizard

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



Enabling and Disabling Workers

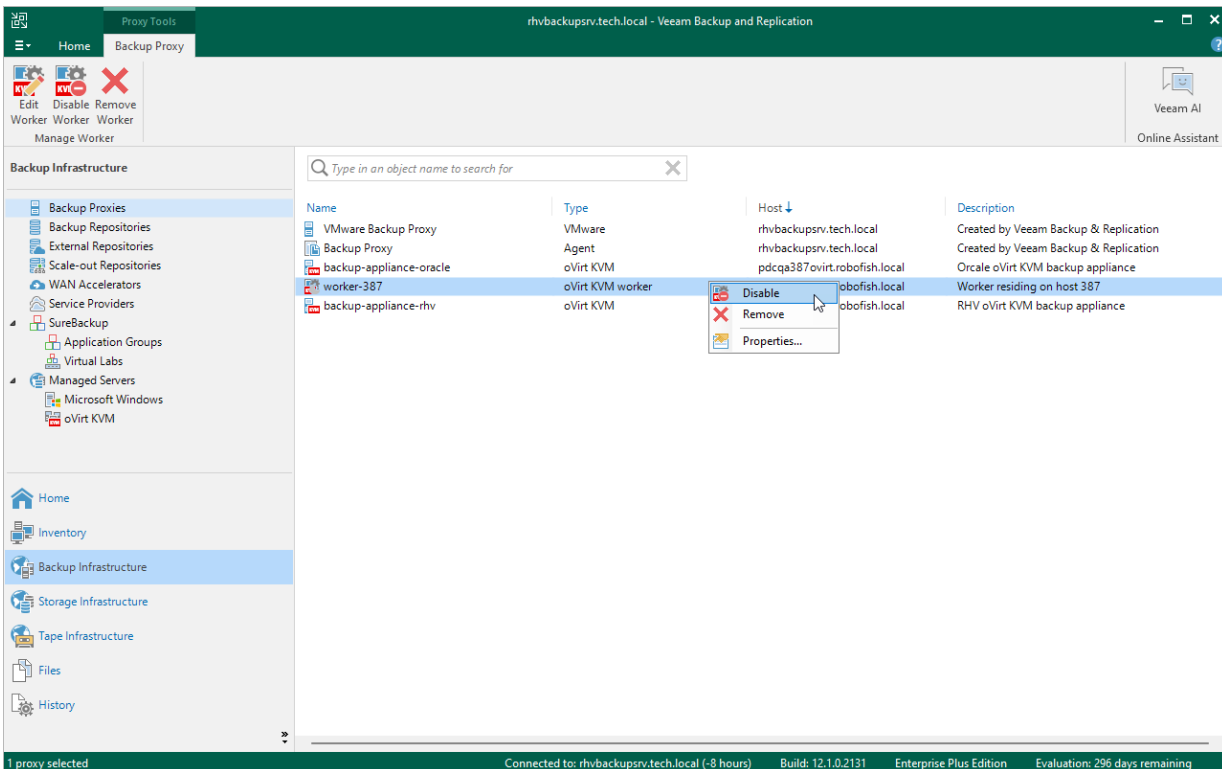
By default, workers are launched when jobs or restore sessions start. However, you can temporarily disable a worker – this may be helpful when you reconfigure a worker and you do not want it to be used for a backup or restore operation. You will still be able to enable the disabled worker at any time 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 worker and click **Disable Worker** or **Enable Worker** on the ribbon, or right-click the worker and select **Disable** or **Enable**.

TIP

If you use dedicated workers, it is recommended that you disable the embedded worker. To do that, in the working area, select the backup appliance and click **Disable Embedded Worker** on the ribbon.



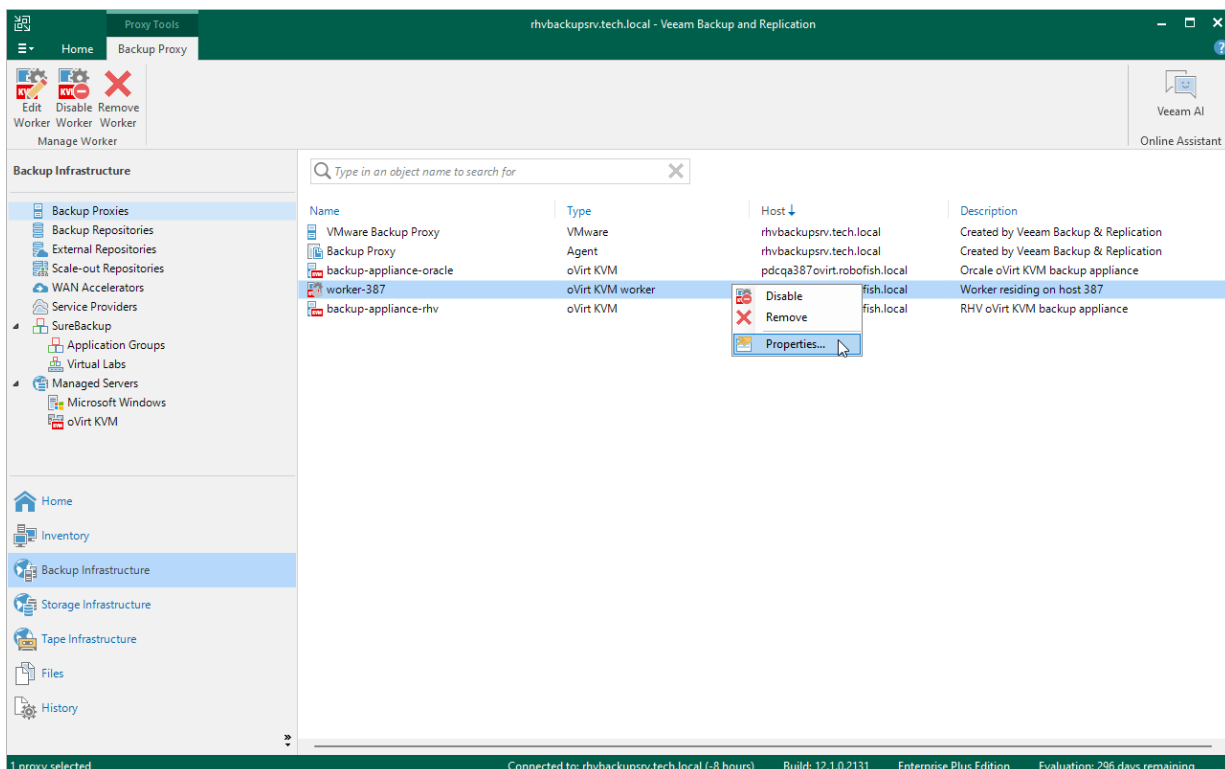
Editing Workers

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 worker and click **Edit Worker** on the ribbon, or right-click the worker and select **Properties**.
4. Complete the **Edit oVirt KVM Worker** wizard:
 - a. To provide a new name and description for the worker, to change the storage domain where worker system files are stored, to specify a host where the worker is launched or to modify the number of tasks that the worker is able to handle in parallel, follow the instructions provided in section [Adding Workers](#) (step 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 storage domain, decrease the amount of allocated resources, adjust the affinity settings or modify the network settings while the worker is currently transferring data. In this case, Veeam Backup for OLVM and RHV will terminate the related operations, power off the worker and update the settings immediately.

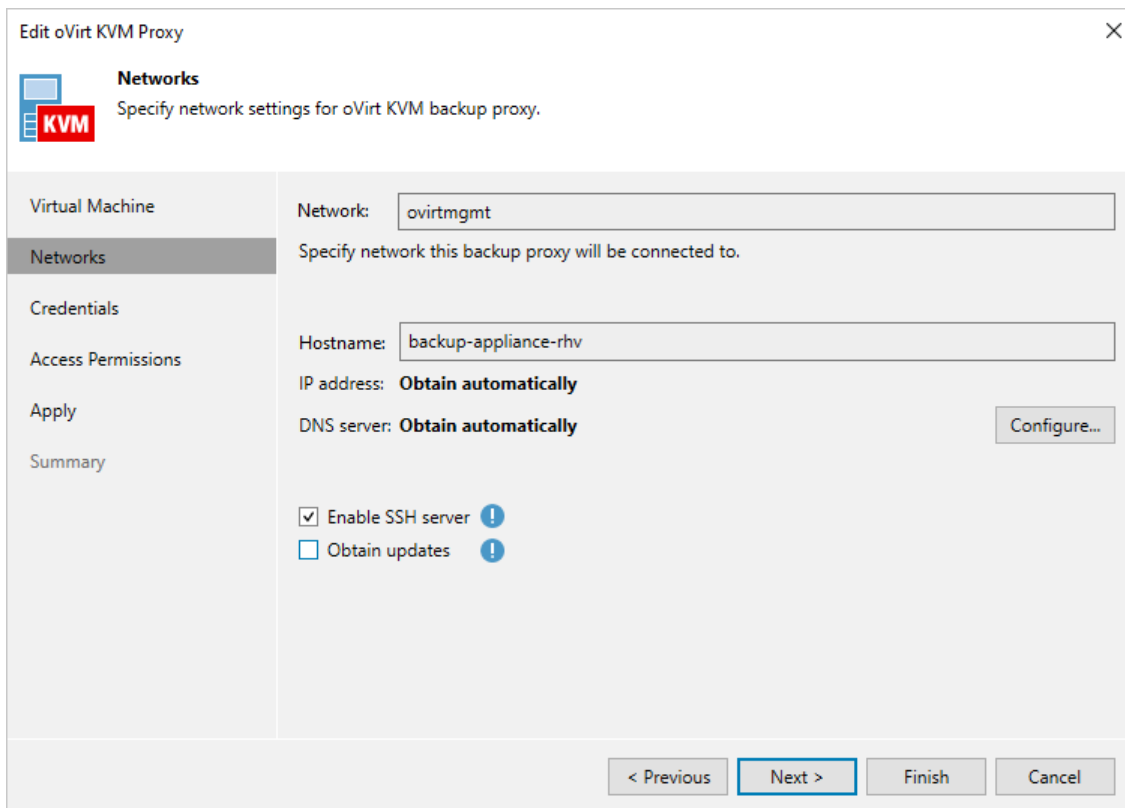


Updating Workers

While starting workers, Veeam Backup for OLVM and RHV automatically downloads updates from Veeam repositories and installs them. If workers are not connected to the internet, you can instruct Veeam Backup for OLVM and RHV to [use an HTTP proxy](#) that will provide access to the required resources.

If workers do not have access to the internet and no HTTP proxy is used, disable automatic updates to eliminate update failures and session warnings:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the backup appliance and click **Edit Appliance** on the ribbon, or right-click the backup appliance and select **Properties**.
4. At the **Networks** step of the **Edit oVirt KVM Proxy** wizard, clear the **Obtain updates** check box.



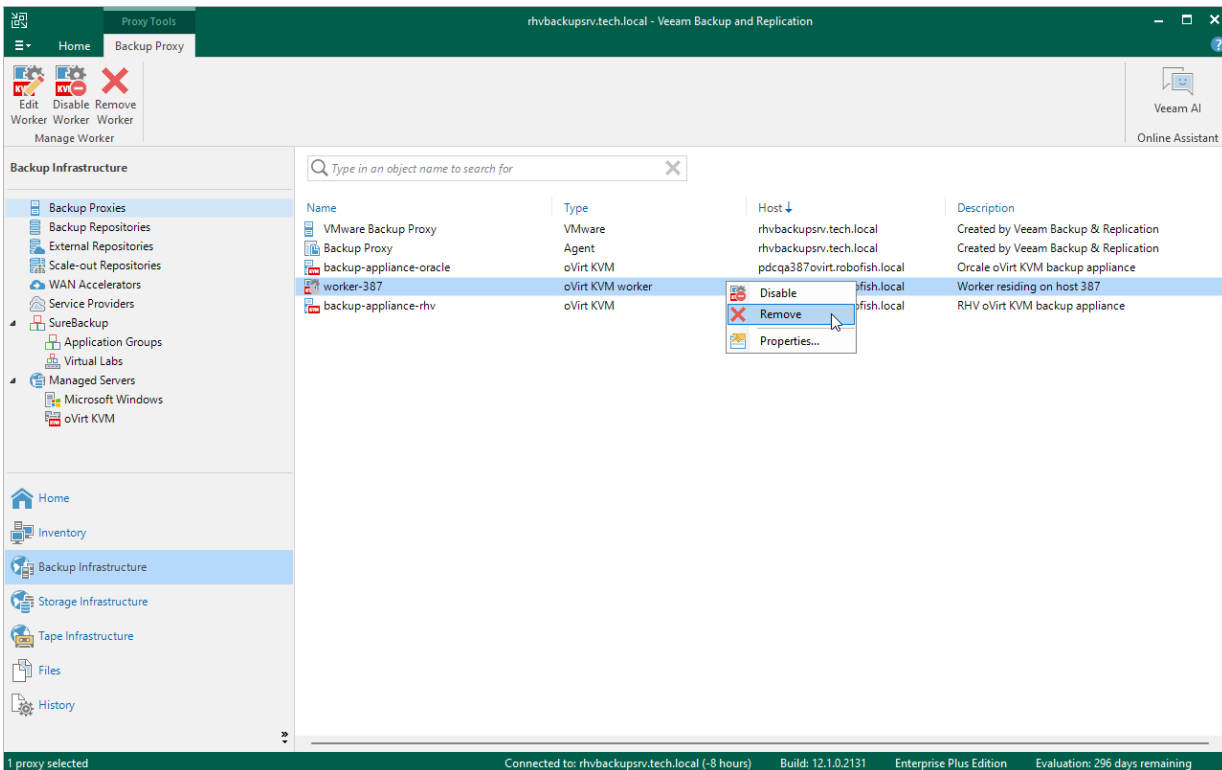
The screenshot shows the 'Edit oVirt KVM Proxy' wizard window, specifically the 'Networks' step. The window title is 'Edit oVirt KVM Proxy' with a close button (X) in the top right corner. Below the title bar, there is a 'KVM' logo and the text 'Specify network settings for oVirt KVM backup proxy.' The left sidebar contains a navigation pane with the following items: 'Virtual Machine', 'Networks' (selected), 'Credentials', 'Access Permissions', 'Apply', and 'Summary'. The main area of the wizard is divided into two columns. The left column contains the navigation pane. The right column contains the configuration fields: 'Network:' with a text box containing 'ovirtmgmt'; 'Specify network this backup proxy will be connected to.'; 'Hostname:' with a text box containing 'backup-appliance-rhv'; 'IP address:' with the text 'Obtain automatically'; 'DNS server:' with the text 'Obtain automatically' and a 'Configure...' button to its right; and two checkboxes: 'Enable SSH server' (checked) and 'Obtain updates' (unchecked). At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Removing Workers

Veeam Backup for OLVM and RHV allows you to permanently remove workers if you no longer need them. Note that you can remove a worker only when it is not processing 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 worker and click **Remove Worker** on the ribbon, or right-click the worker and select **Remove**.
4. In the **Veeam Backup & Replication** window, confirm that you want to permanently delete the worker.



Performing Configuration Backup and Restore

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

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

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

Backing Up Configuration Settings Manually

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

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. In the **Configuration Backup Settings** window, do the following:
 - a. Select the **Enable configuration backup to the following repository** check box and choose a repository where the configuration backup will be stored. Note that you cannot store configuration backups in scale-out backup repositories and external repositories.

For a backup repository to be displayed in the list of available repositories, it must be added to the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Adding Backup Repositories](#).
 - b. In the **Restore points to keep** field, specify the number of configuration backups you want to keep.
 - c. Select the **Enable backup file encryption** check box.
 - d. From the **Password** drop-down list, select a password.

IMPORTANT


If you do not specify the password, the Veeam Backup for OLVM and RHV configuration database will not be backed up.

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

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

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

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



Backing Up Configuration Settings Automatically

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

1. From the main menu of the Veeam Backup & Replication console, select **Configuration Backup**.
2. In the **Configuration Backup Settings** window, do the following:
 - a. Select the **Enable configuration backup to the following repository** check box and choose a repository where the configuration backup will be stored. Note that you cannot store configuration backups in scale-out backup repositories and external repositories.

For a backup repository to be displayed in the list of available repositories, it must be added to the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Adding Backup Repositories](#).
 - b. In the **Restore points to keep** field, specify the number of configuration backups you want to keep.
 - c. Click **Schedule** and choose whether configuration backups will be created every day or monthly on specific days.
 - d. Select the **Enable backup file encryption** check box.
 - e. From the **Password** drop-down list, select a password.

IMPORTANT

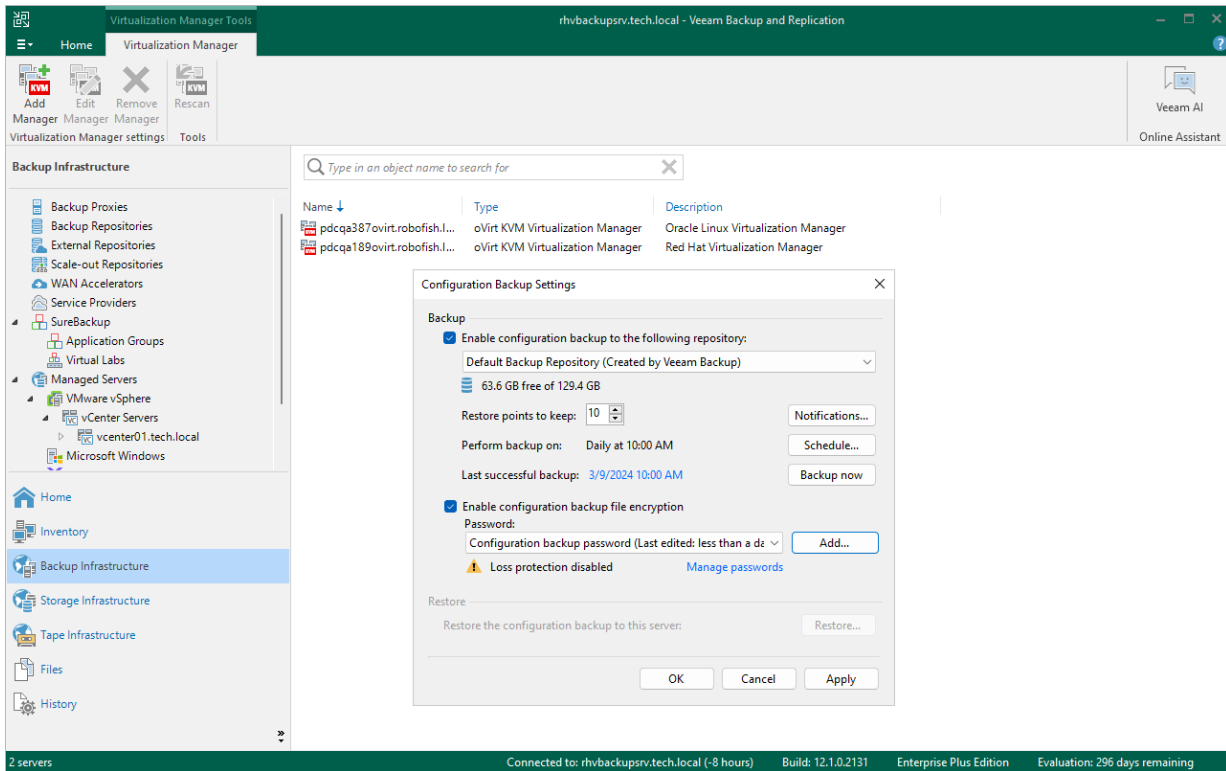
If you do not specify the password, the Veeam Backup for OLVM and RHV configuration database will not be backed up.

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

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

- f. Click **OK**.

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



Restoring Configuration Settings

Veeam Backup for OLVM and RHV offers restore of the configuration database that can be helpful in the following situations:

- The configuration database got corrupted, and you want to recover data from a configuration backup.
- The backup appliance got corrupted, and you want to recover its configuration from a configuration backup.
- The backup appliance went down, and you want to apply its configuration to a new backup appliance.
- You want to roll back the configuration database to a specific point in time.
- You want to apply the backed-up configuration of a backup appliance version 2.0 (or later) to a newly deployed backup appliance.

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

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

IMPORTANT

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

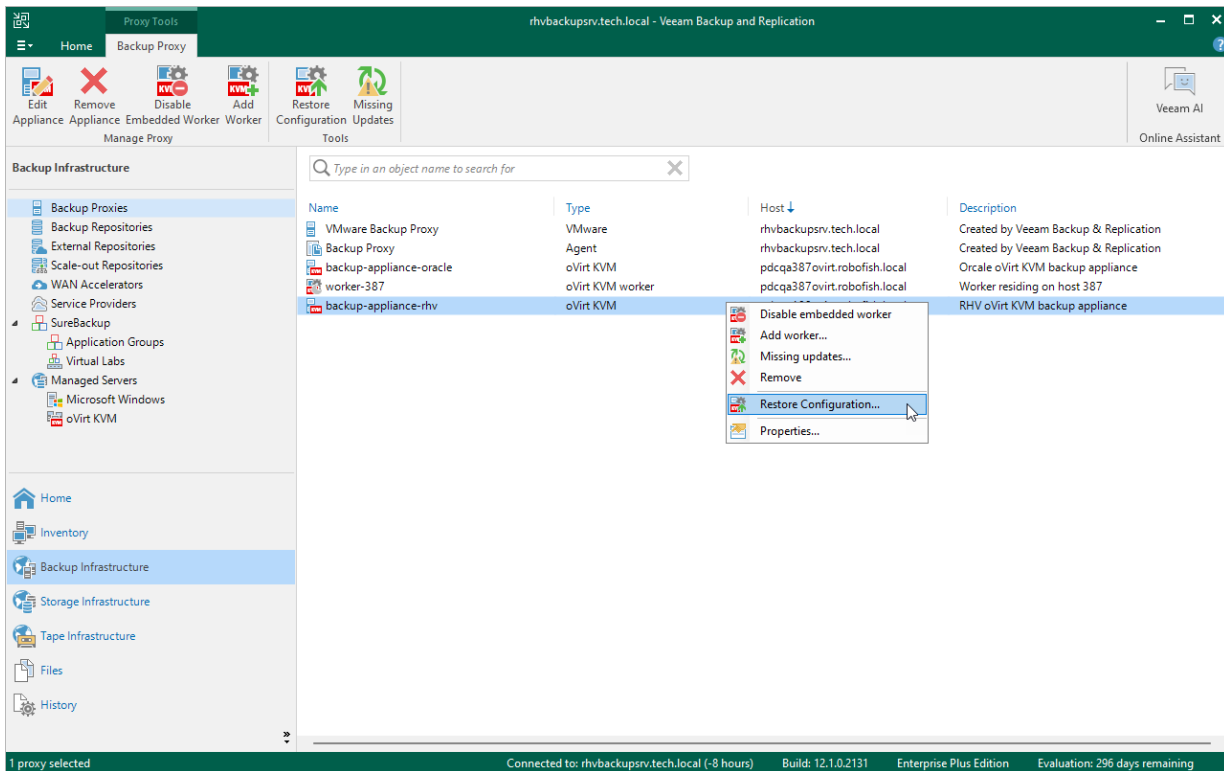
To restore the configuration database, do the following:

1. [Launch the Configuration Restore wizard](#).
2. [Choose a backup file](#).
3. [Review the backup file information](#).
4. [Provide the encryption password](#).
5. [Choose restore options](#).
6. [Track the restore progress](#).
7. [Finish working with the wizard](#).

Step 1. Launch Configuration Restore Wizard

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

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



Step 2. Choose Backup File

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

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

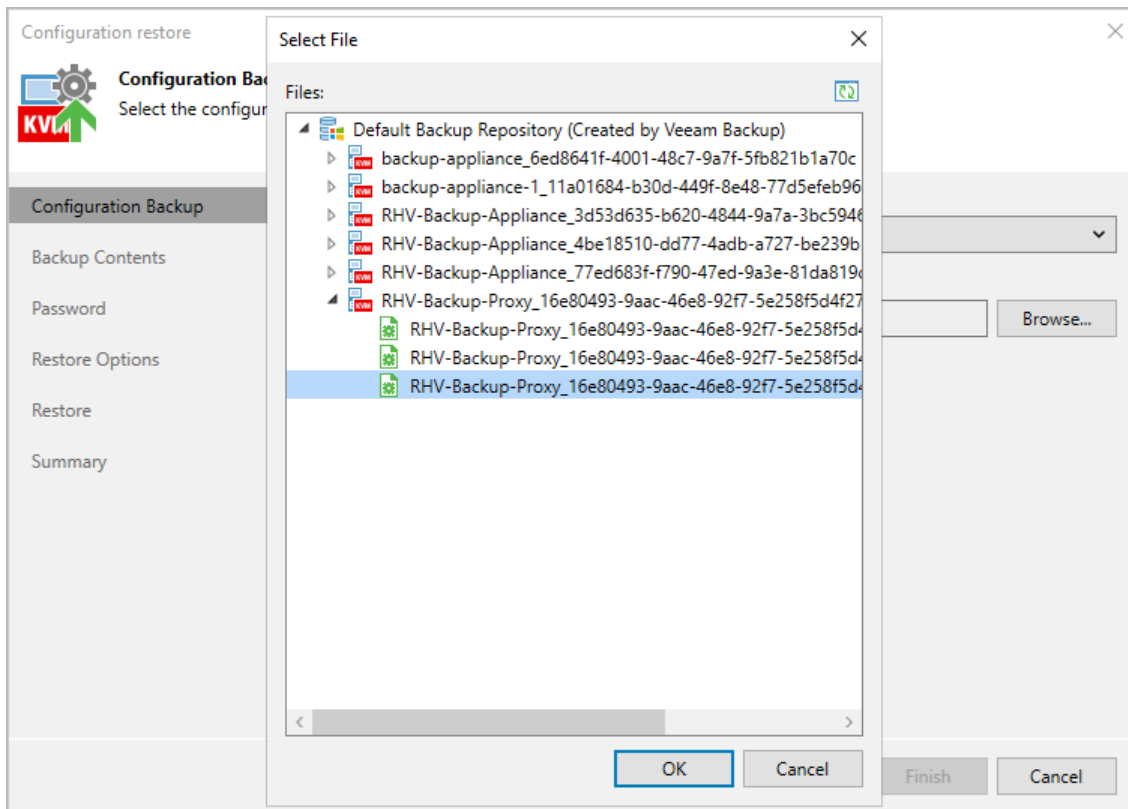
For a backup repository to be displayed in the **Backup repository** list, it must be added to the backup infrastructure as described in the Veeam Backup & Replication User Guide, section [Backup Repository](#). Note that the repository list does not include scale-out backup repositories and external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

NOTE

Configuration restore is supported for backup appliances version 2.0, 2a, 3.0. 3a or 3b only if their configuration backup files are stored on the backup server.

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

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



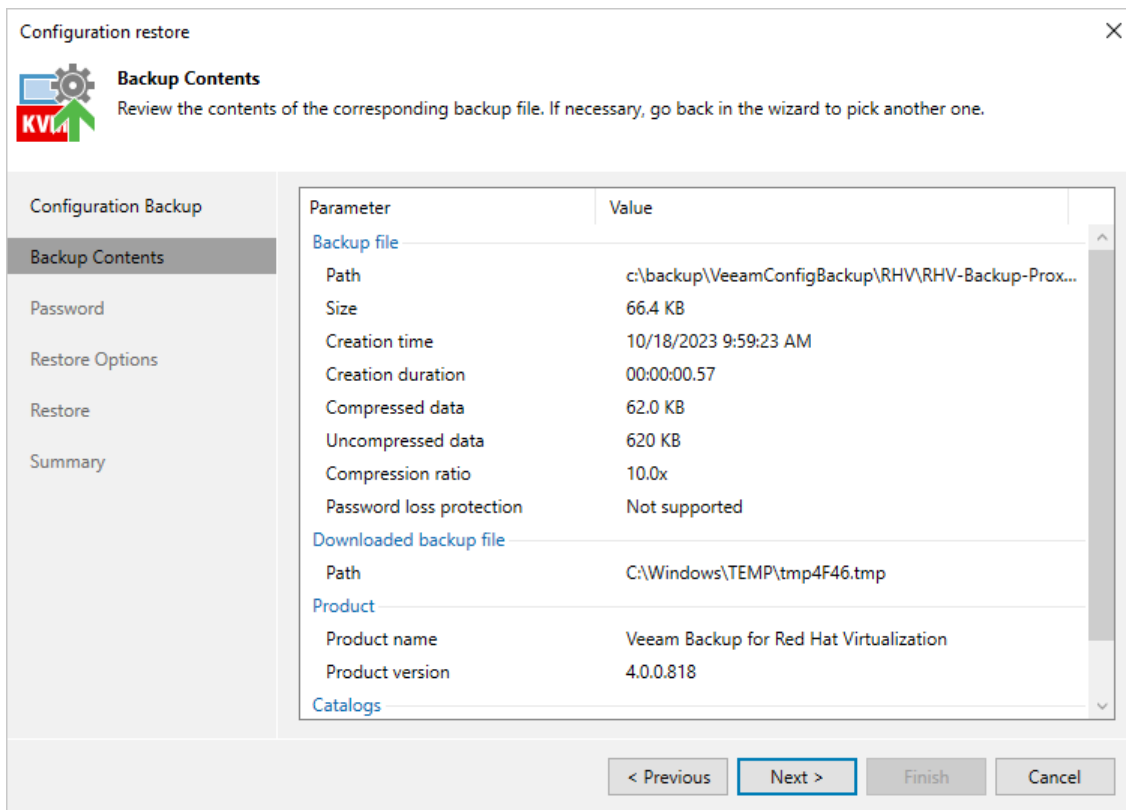
Step 3. Review Backup Details

[This step applies only if you restore the backup configuration database of a backup appliance version 4.0 or later]

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

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

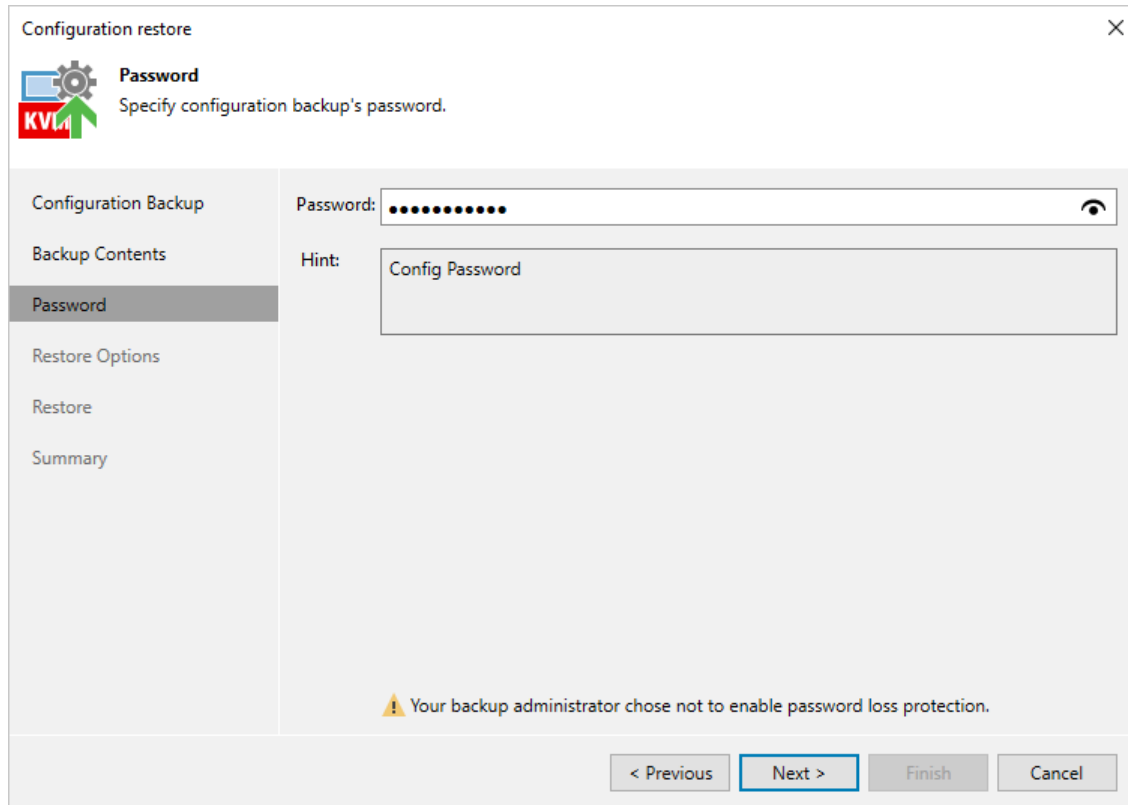
At the **Backup Content** step of the wizard, review the provided information and click **Next** to confirm that you want to use the selected file to restore the configuration data.



Step 4. Provide Encryption Password

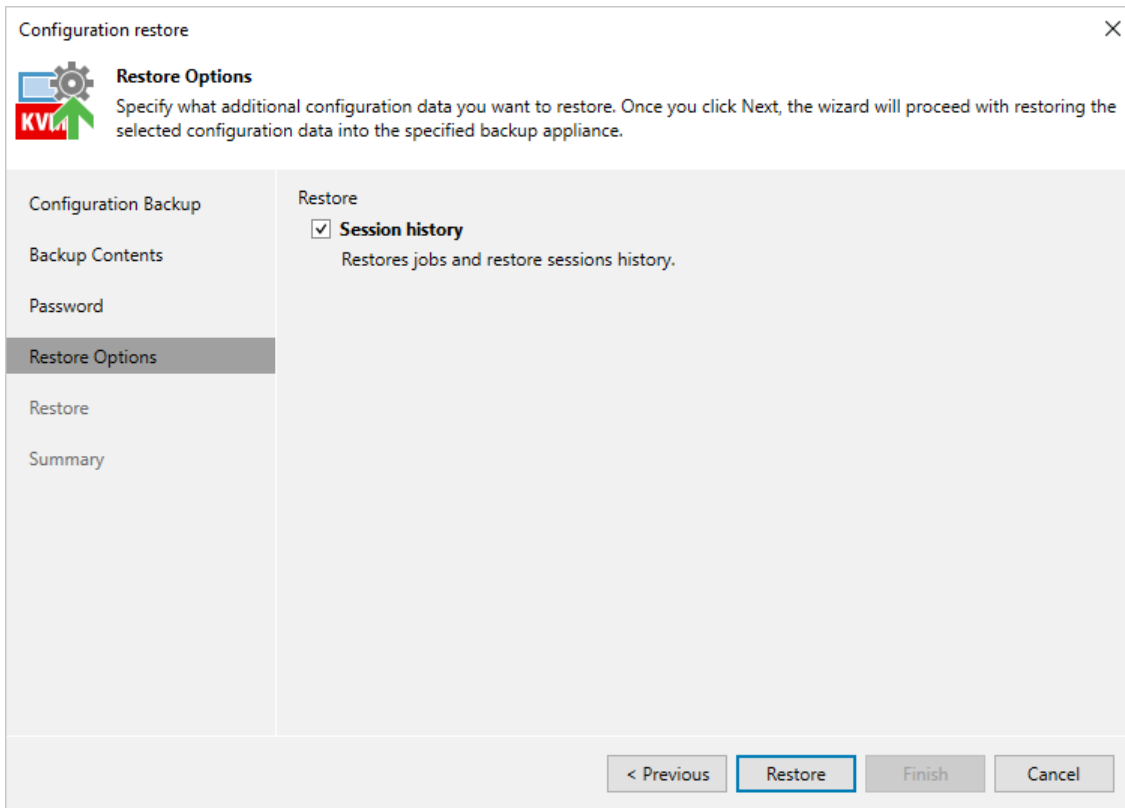
At the **Password** step of the wizard, provide a password that was used to encrypt the file while creating configuration backup.

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



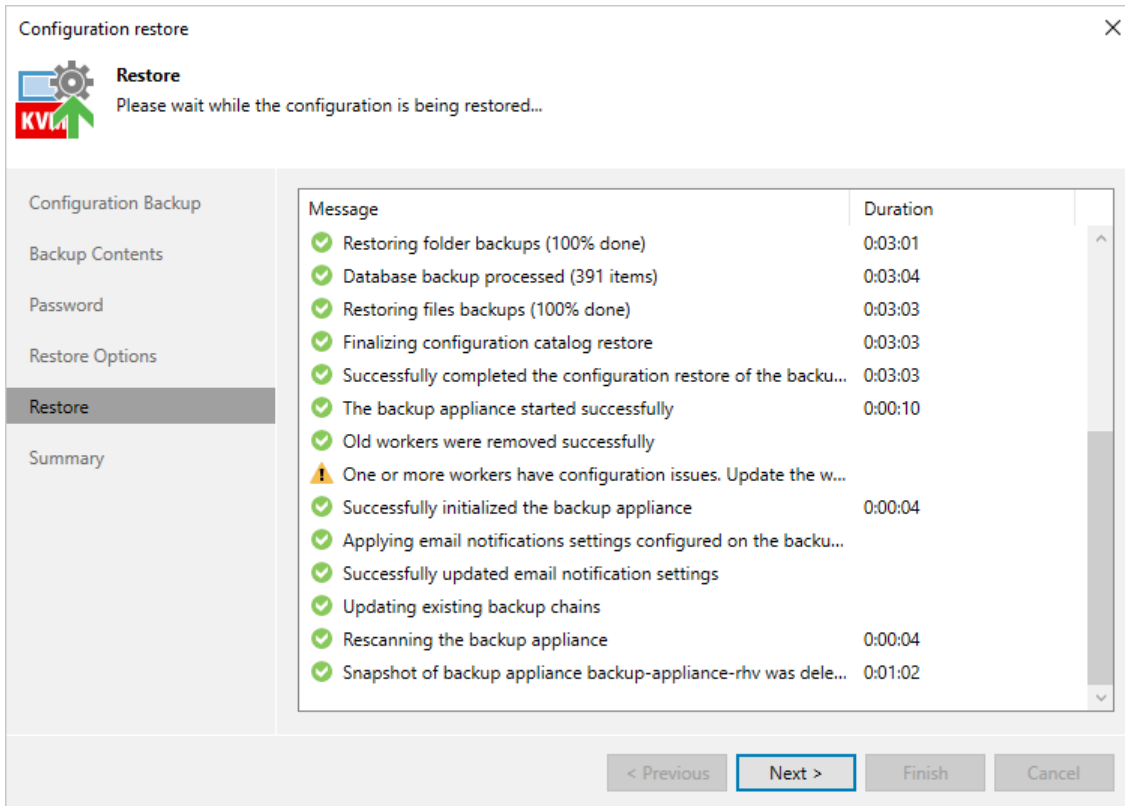
Step 5. Choose Restore Options

At the **Restore Options** step of the wizard, you can choose whether you want to restore jobs and session logs.



Step 6. Track Restore Progress

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



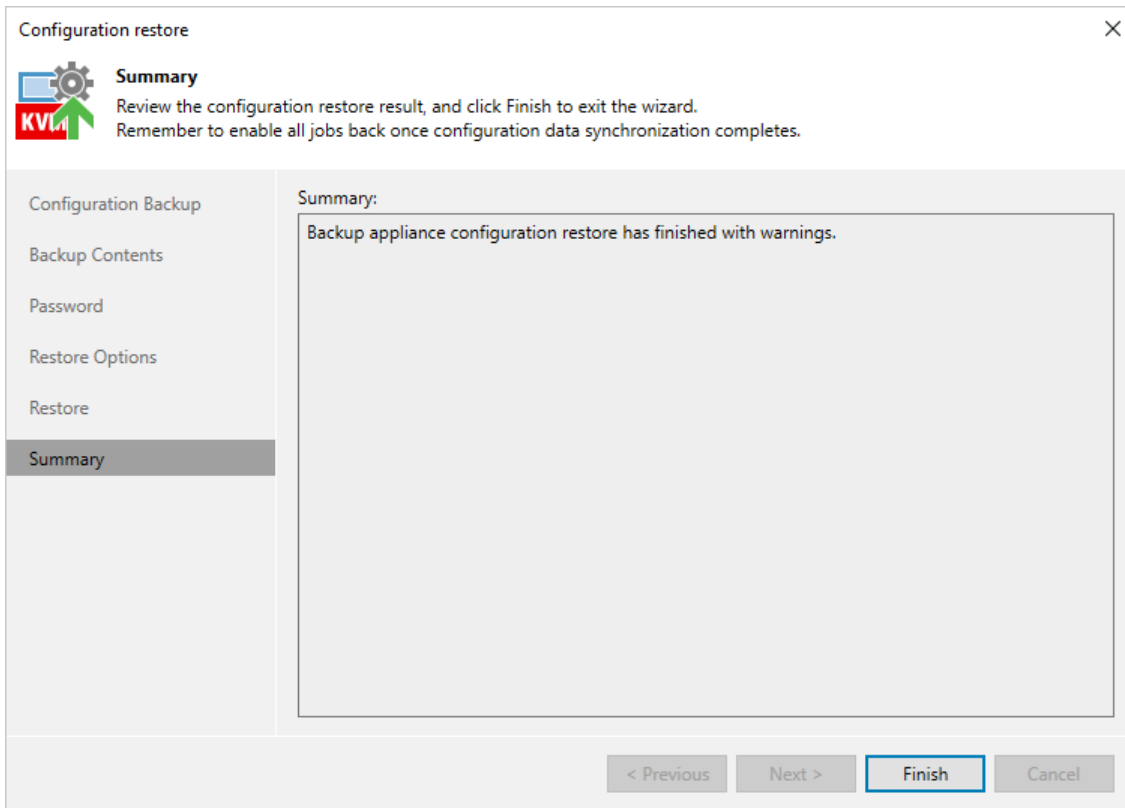
The screenshot shows a 'Configuration restore' window with a sidebar on the left containing the following steps: Configuration Backup, Backup Contents, Password, Restore Options, **Restore**, and Summary. The main area displays a list of messages with their durations:

Message	Duration
✓ Restoring folder backups (100% done)	0:03:01
✓ Database backup processed (391 items)	0:03:04
✓ Restoring files backups (100% done)	0:03:03
✓ Finalizing configuration catalog restore	0:03:03
✓ Successfully completed the configuration restore of the backu...	0:03:03
✓ The backup appliance started successfully	0:00:10
✓ Old workers were removed successfully	
⚠ One or more workers have configuration issues. Update the w...	
✓ Successfully initialized the backup appliance	0:00:04
✓ Applying email notifications settings configured on the backu...	
✓ Successfully updated email notification settings	
✓ Updating existing backup chains	
✓ Rescanning the backup appliance	0:00:04
✓ Snapshot of backup appliance backup-appliance-rhv was dele...	0:01:02

At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

Step 7. Finish Working with Wizard

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



Performing Backup

To produce backups of oVirt VMs, Veeam Backup for OLVM and RHV 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.

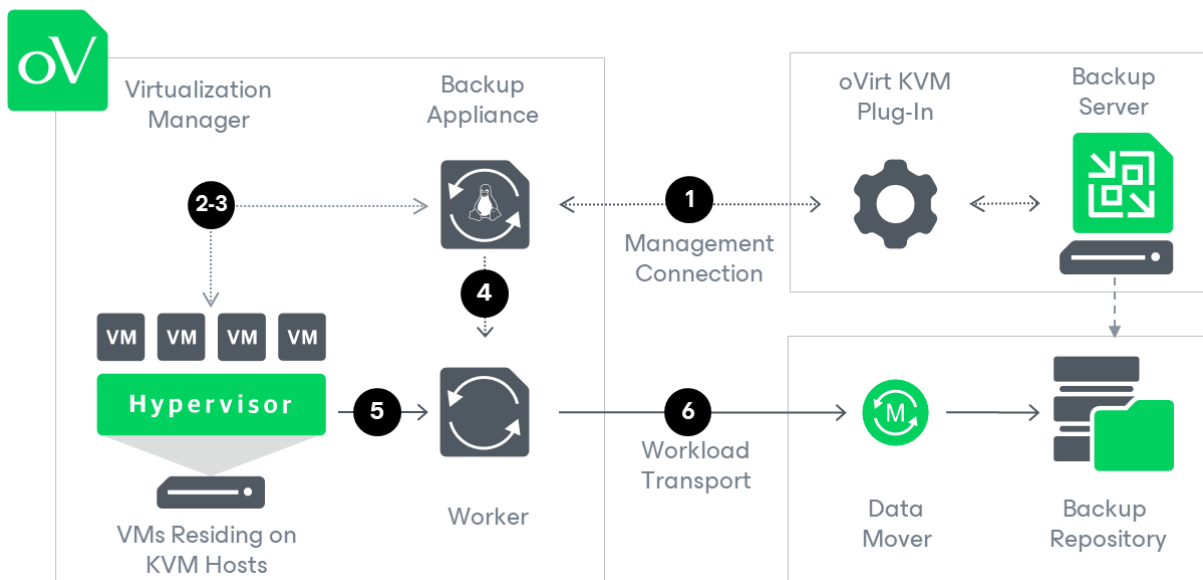
One backup job can be used to process multiple VMs, but you can back up each VM with one backup job at a time. If a VM is added to more than one backup job, it will be processed only by the backup job that started earlier.

You can instruct the Veeam Backup for OLVM and RHV to run jobs automatically according to a specified schedule or start them manually.

How Backup Works

While creating image-level backups, Veeam Backup for OLVM and RHV does not install agent software inside VMs to retrieve data. Veeam Backup for OLVM and RHV uses [native oVirt capabilities](#) to take VM snapshots and produces backups in the following way:

1. The backup server starts a backup job and forwards the backup session data to the backup appliance.
2. The backup appliance connects to the Virtualization manager over REST API and creates snapshots of all VMs added to the job.
3. The backup appliance sends a REST API request to the Virtualization manager to create an image transfer session and to provide its URL.
4. The backup appliance launches a worker.
5. The worker retrieves the VM data using the provided URL.
6. The worker compresses and deduplicates the VM data and forwards it to the target backup repository in the native Veeam format.



Backup Chain

Veeam Backup for OLVM and RHV creates a new backup file in a backup repository during every backup session. A sequence of backup files created during a set of backup sessions makes up a backup chain. Each backup chain contains data for one VM only. If a backup job includes several VMs, Veeam Backup for OLVM and RHV 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.

Backup Methods

Veeam Backup for OLVM and RHV provides the following methods for creating backup chains:

- **Forever forward incremental**

When the forever forward incremental backup method is used, Veeam Backup for OLVM and RHV 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 section [Forever Forward Incremental Backup](#).

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

- **Forward incremental**


When the forward incremental backup method is used, Veeam Backup for OLVM and RHV 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 section [Forward Incremental Backup](#).

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

Forever Forward Incremental Backup

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


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




Forward Incremental Backup

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

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



3. On a day when the synthetic full or active full backup is scheduled, Veeam Backup for OLVM and RHV 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.



Changed Block Tracking

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

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

To detect unallocated and changed data blocks, CBT relies on the [oVirt REST API](#):

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

IMPORTANT

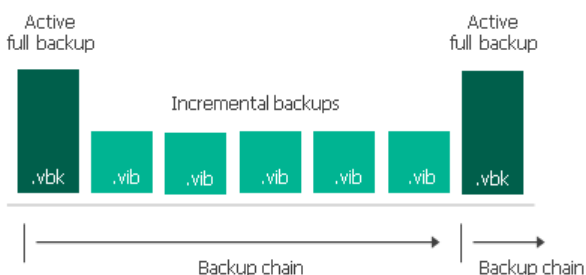
If Veeam Backup for OLVM and RHV is unable to use CBT while creating incremental backups, you may get the following warnings in backup session logs:

- *"Unable to enable oVirt incremental backups for disk. Full scan backups will be performed"*. To resolve the issue, follow the [Veeam KB article](#).
- *"The Disk id=<disk id> has RAW format and can be backed up only in full scan mode"*. To resolve the issue, take a VM snapshot in the Administration Portal as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#). Alternatively, [back up the VM](#) and [restore its VM disks](#) with the **Restore all VM disks to QCOW2 format** option selected at the **Configure Mapping Settings** step of the **Virtual Disk Restore** wizard.

Active Full Backup

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

When creating an active full backup, Veeam Backup for OLVM and RHV 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.



The active full backup session starts at the same time when the backup job is scheduled. For example, if you schedule the backup job to run at 12:00 AM Sunday through Friday, and schedule active full backup to be created on Saturday, Veeam Backup for OLVM and RHV will start a backup job session that will produce an active full backup at 12:00 AM on Saturday.


If the backup job is not scheduled to run automatically or is disabled, Veeam Backup for OLVM and RHV will not perform active full backup. If a regular backup session and an active full backup session are scheduled on the same day, Veeam Backup for OLVM and RHV will produce an active full backup – an incremental backup that should have been created by the regular backup session will not be added to the backup chain. However, if you run the backup job again on the same day manually, Veeam Backup for OLVM and RHV will perform incremental backup in a regular manner.

Synthetic Full Backup


In some situations, running active full backups periodically may not be an option. Active full backups are resource-intensive and consume considerable amount of network bandwidth. As an alternative, you can create synthetic full backups that also produce VBK files and contain data of the whole VM. However, while creating synthetic full backups, Veeam Backup for OLVM and RHV does not retrieve VM data from the cluster but processes the data that is already stored in the backup repository.

To create a synthetic full backup, Veeam Backup for OLVM and RHV performs the following operations:


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



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



3. Veeam Backup for OLVM and RHV deletes the created incremental backup as its data is already incorporated in the synthetic full backup.



When creating a synthetic full backup, Veeam Backup for OLVM and RHV 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.

NOTE

The synthetic full backup session starts only on the day 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 synthetic full backup to be created on Saturday, Veeam Backup for OLVM and RHV will never start a backup job session that will produce a synthetic full backup.

If the backup job is not scheduled to run automatically or is disabled, Veeam Backup for OLVM and RHV will not perform synthetic full backup. If a regular backup session and a synthetic full backup session are scheduled on the same day, Veeam Backup for OLVM and RHV will produce a synthetic full backup – an incremental backup that should have been created by the regular backup session will not be added to the backup chain. However, if you run the backup job again on the same day manually, Veeam Backup for OLVM and RHV will perform incremental backup in a regular manner.

Retention Policy

Backups created by jobs are not kept forever – they are removed according to retention policy settings specified while creating the jobs. Depending on the data protection scenario, retention policy can be specified:

- **In days**

Restore points in the backup chain can be stored only for the allowed period of time. If a restore point is older than the specified time limit, Veeam Backup for OLVM and RHV removes it from the backup chain. Since for retention policy specified in days, the backup chain must contain at least 3 restore points, Veeam Backup for OLVM and RHV may retain restore points for a longer period than configured in the retention policy settings.

- **In restore points**

The chain can contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for OLVM and RHV removes the earliest restore point from the chain.

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


NOTE

To backup chains created by backup jobs that no longer exist, Veeam Backup for OLVM and RHV applies [background retention](#).

Forever Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forever forward incremental backup chain, Veeam Backup for OLVM and RHV performs the following actions once a day:


1. Veeam Backup for OLVM and RHV checks the configuration database to detect backup chains where the number of allowed restore points is exceeded.
 - If retention policy is specified in days, Veeam Backup for OLVM and RHV detects backup chains with restore points that are older than the specified time limit.
 - If retention policy is specified in restore points, Veeam Backup for OLVM and RHV detects backup chains where the number of allowed restore points is exceeded.
2. If a redundant restore point exists in a backup chain, Veeam Backup for OLVM and RHV transforms the backup chain in the following way:
 - a. Rebuilds the full backup to include the data of the incremental backup that follows the full backup. To do that, Veeam Backup for OLVM and RHV 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 Backup for OLVM and RHV repeats step 2 for all other redundant restore points found in the backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full backup, Veeam Backup for OLVM and RHV ensures that the backup chain is not broken and that you will be able to recover your data when needed.



Forward Incremental Backup Retention Policy

To track and remove redundant restore points from a forward incremental backup chain, Veeam Backup for OLVM and RHV performs the following actions once a day:

1. Veeam Backup for OLVM and RHV 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 Backup for OLVM and RHV checks the following:
 - o If retention policy is specified in days, Veeam Backup for OLVM and RHV checks whether the period to keep restore points in the new chain fragment has reached the allowed time limit.
 - o If retention policy is specified in restore points, Veeam Backup for OLVM and RHV checks whether the number of restore points in the new chain fragment has reached the number of allowed restore points.
3. If the new backup chain fragment has reached the limit of allowed restore points, Veeam Backup for OLVM and RHV removes all restore points of the older backup chain fragment.



Creating Backup Jobs

To create a backup job, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the New Backup Job wizard.](#)
3. [Specify a job name and description.](#)
4. [Selects VMs to backup.](#)
5. [Specify a backup repository where backups will be stored and configure backup settings.](#)
6. [Create a schedule for the backup job.](#)
7. [Finish working with the wizard.](#)

Before You Begin

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

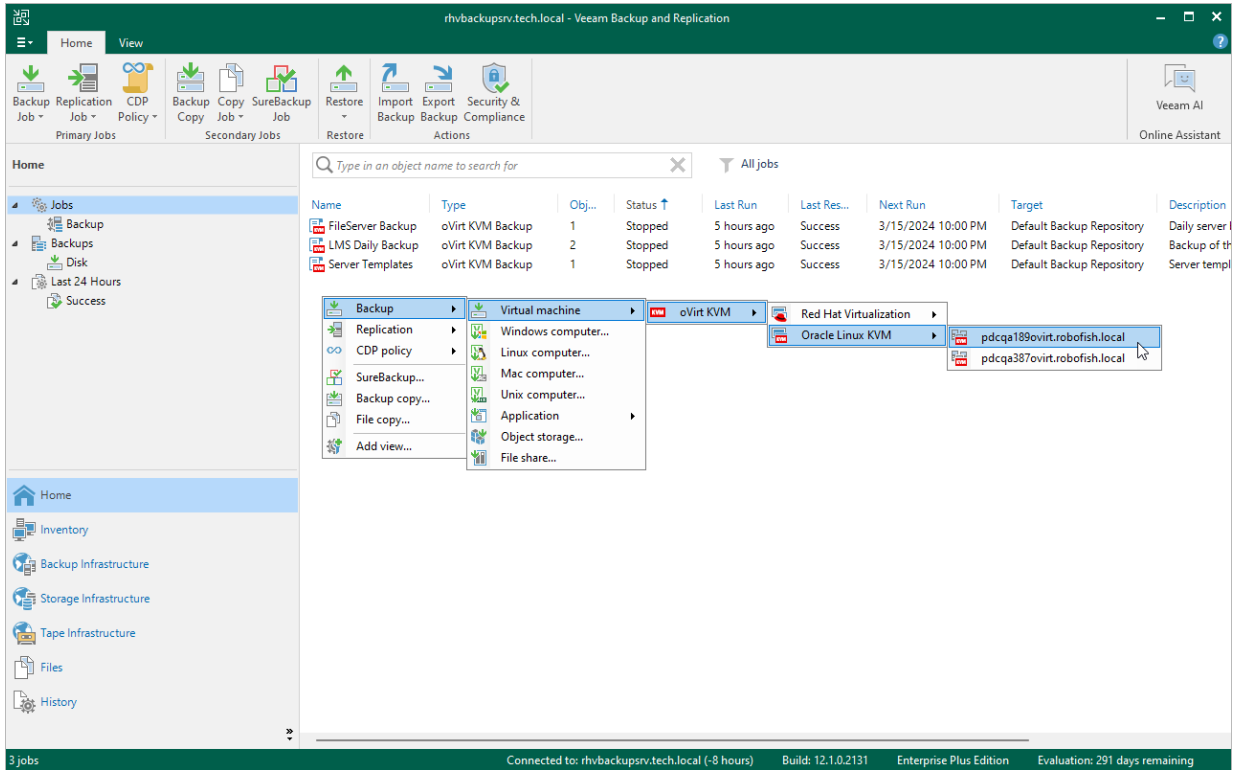
- You can back up each VM with one backup job at a time. If a VM is already being processed by a backup job, another backup job will not start processing this VM until the currently running backup operation completes.
- You cannot back up a VM being restored. Wait for the restore process to complete, and then start the backup job.
- You cannot back up hosted-engine VMs. However, you can create a backup of the oVirt configuration. For more information, see [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).
- You cannot back up a VM while previewing its snapshot. For more information, see [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).
- You cannot back up a VM that has [shareable disks](#) or [direct LUN disks](#) attached.
- You cannot include into a backup job a VM that is being backed up by 3rd party software or an backup appliance connected to another backup server. Wait for the backup process to complete or stop the currently running job manually, and then add the VM to the necessary backup job.
- By default, Veeam Backup for OLVM and RHV applies the following [deduplication and compression settings](#) to backed-up data:
 - Deduplication: *Enabled*
 - Data compression level: *Optimal*
 - Storage optimization: *Local target* (1024 KB block size)

Due to technical limitations, you cannot change these 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 Veeam Backup & Replication User Guide, section [Access Permissions](#).
- [VM guest OS file indexing](#) is not supported for backups created with Veeam Backup for OLVM and RHV.
- Since Veeam Backup & Replication does not allow you to assign [information about locations](#) to the Virtualization manager and backup appliance, job statistics do not include information on the oVirt VM data migration between different geographic regions.
- If you want to back up a VM that has been configured with a [oVirt KVM Virtualization Cloud-Init custom script](#), first remove the script from the VM since it may contain secure data (such as credentials and authorized keys) that will appear in Veeam Backup for OLVM and RHV backup logs.

Step 1. Launch New Backup Job Wizard

In the inventory pane, select **Jobs** and navigate to **Backup > Virtual Machine>oVirt KVM** and choose **Red Hat Virtualization** or **Oracle Linux KVM**.



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 for OLVM and RHV.

The maximum length of the name is 40 characters; the following characters are not supported: ~ " # % & * : < > ! ? / \ { | } . ' ` \$. The maximum length of the description is 1024 characters.

New Backup Job

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

Name

Virtual Machines

Storage

Schedule

Summary

Name:
FileServer Backup

Description:
Daily server backup

< Previous Next > Finish Cancel

Step 3. Configure Backup Source Settings

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

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

Step 3a. Choose Resources

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

1. Click **Add**.
2. In the **Add Objects** window, choose whether you want to back up specific VMs or groups of VMs arranged by tags:
 - If you click the **VM** icon, you must specify the machines explicitly.

NOTE

If any of the selected VMs have disks in the RAW format attached, Veeam Backup for OLVM and RHV will display the following warning: "*There are some VM disks that do not support oVirt incremental backup. The policy will do a full scan backup for those disks*". Due to technical limitations, Veeam Backup for OLVM and RHV is only able to apply the **CBT mechanism** to disks in the QCOW2 format while performing incremental backup.

You can proceed with the wizard and resolve the issue later by using one of the following workarounds:

- Take VM snapshots in the Administration Portal as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).
- **Back up the VMs** and **restore their disks** with the **Restore all VM disks to QCOW2 format** option selected at the **Configure Mapping Settings** step of the **Virtual Disk Restore** wizard.
- If you click the **Tag** icon and add a tag to the backup scope, Veeam Backup for OLVM and RHV will regularly check for new VMs assigned the added tag and automatically update the backup job settings to include these VMs in the scope. For a tag to be displayed in the list, it must be created in the Administration Portal and assigned to a VM. For more information on tags, see [oVirt Product Documentation](#).

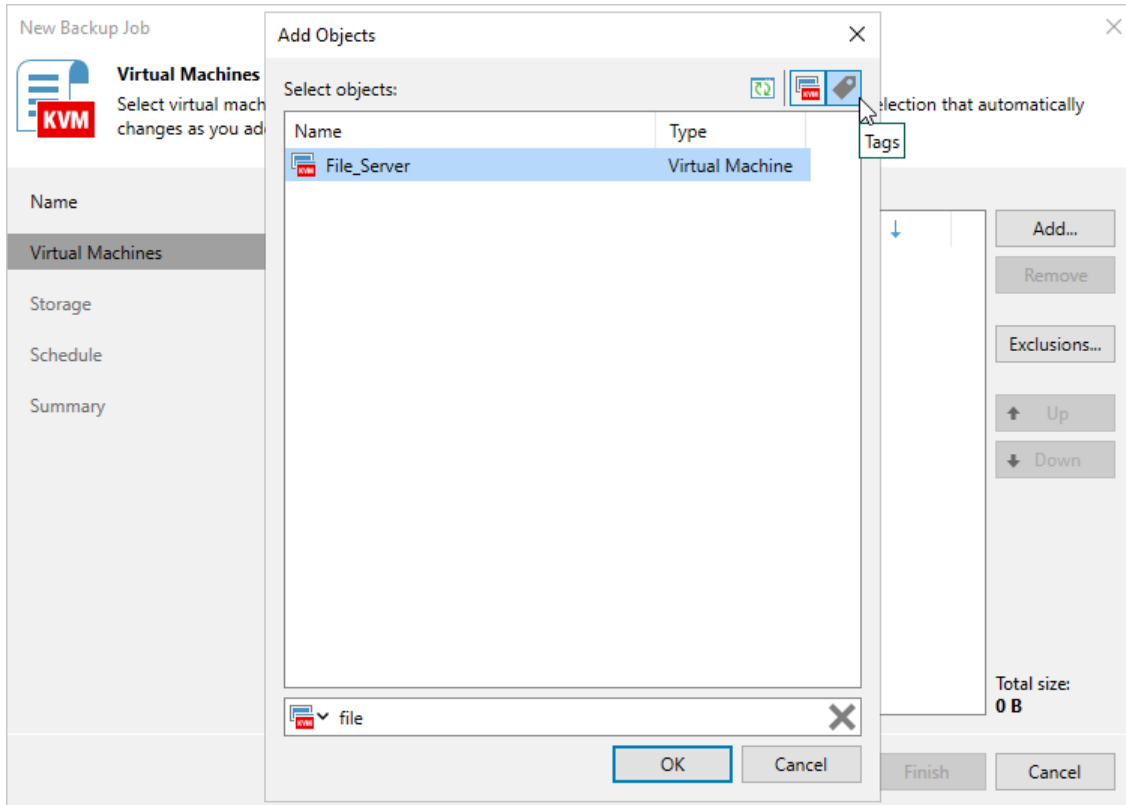
By default, backup jobs process all VMs to which the added tags are assigned. If you want to exclude specific VMs from the backup scope, click **Exclusions** and specify the VMs that you do not want to back up – the procedure is the same as described for including VMs in the backup scope.

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

NOTE

If you include a tag into the backup scope, VMs assigned this tag are processed at random. To ensure that the VMs are processed in a specific order, you must add them as standalone VMs.

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

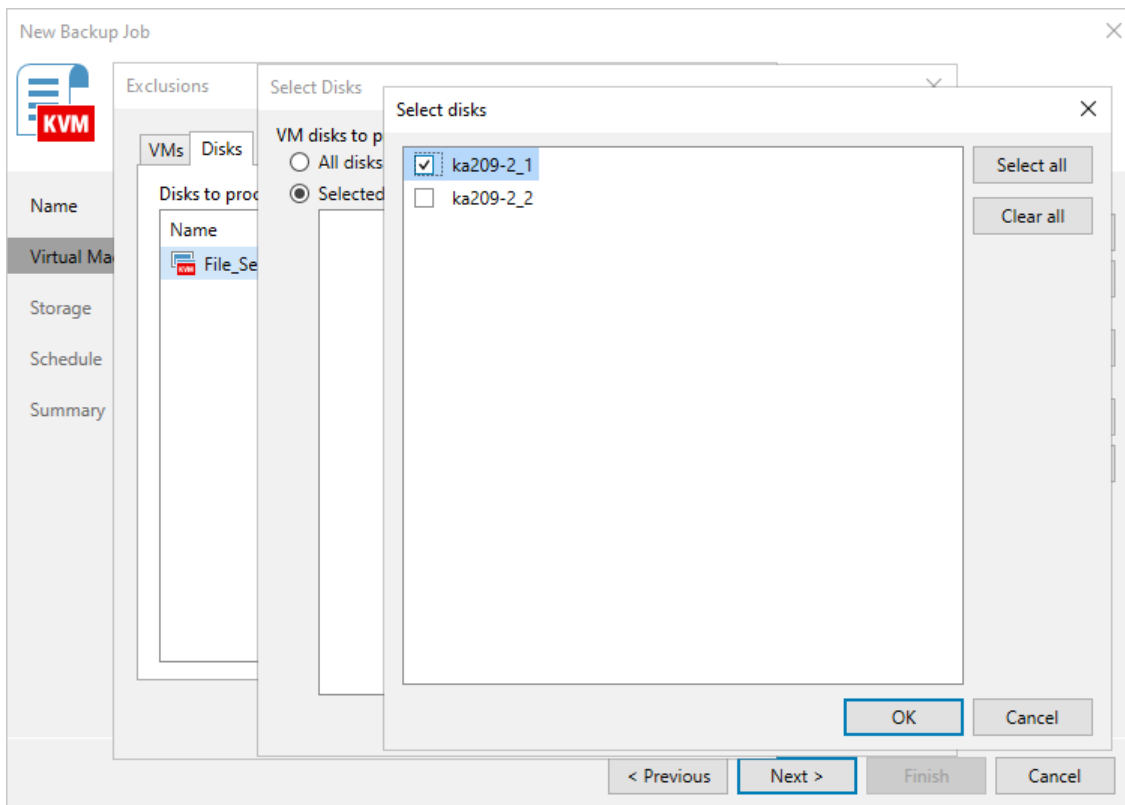


Step 3b. Choose Disks

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

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

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



Step 4. Specify Backup Job Settings

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

1. In the **Backup repository** drop-down list, select a backup repository where you want to store backups.

For a backup repository to be displayed in the list of the available repositories, it must be [added to the backup infrastructure](#), and the backup appliance must [have access to the repository](#).

NOTE

You can back up oVirt VMs to object storage repositories. However, it is recommended that you use the backup appliance as a gateway server to transfer backed-up data. To do that, [enable SSH access](#) on the backup appliance and [add the appliance as a Linux server](#) to the backup infrastructure. Then, edit your object storage repository configuration to choose the **Through gateway server** connection type and select the appliance. For more information on object storage repository configuration, see the Veeam Backup & Replication User Guide, section [Adding Object Storage Repositories](#).

2. In the **Retention policy** section, choose a retention policy that Veeam Backup for OLVM and RHV will apply to backups created by the job:
 - Select *days* if you want to keep restore points in a backup chain for the allowed period of time. If a restore point is older than the specified limit, Veeam Backup for OLVM and RHV removes it from the chain.
 - Select *restore points* if you want a backup chain to contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for OLVM and RHV removes the earliest restore point from the chain.

When the restore point limit is exceeded, Veeam Backup for OLVM and RHV removes the earliest restore point from the chain. For more information, see section [Retention Policy](#).

If the UUID of a VM changes (for example, if the VM was migrated to another cluster), Veeam Backup for OLVM and RHV will be unable to continue the backup chain for this VM. After you re-add the VM to the backup job, Veeam Backup for OLVM and RHV 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.

IMPORTANT

If you use [hardened repositories](#) to store oVirt VM backups, you must consider the following requirements:

- Active full backups must be scheduled in the backup job settings.
- The backup job retention period must be longer than the backup repository immutability period.

For example, if the backup repository immutability period is set to 25 days, you can configure a one-month retention period: specify 4 as the number of restore points, [schedule one backup per week](#) and schedule active full backup to run on the last day of the month.

To help you implement a comprehensive backup strategy, Veeam Backup for OLVM and RHV allows you to [enable long-term retention policy for backups](#) and to [configure backup job advanced settings](#) (such as backup maintenance, health check, active and synthetic full backups).

New Backup Job

Storage
Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

Name

Virtual Machines

Storage

Schedule

Summary

Backup repository:
Default Backup Repository

62.7 GB free of 129 GB

Retention policy: 7 days

Keep certain full backups longer for archival purposes

Configure...

GFS retention policy is not configured

Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings.

Advanced...

< Previous Next > Finish Cancel

Configuring GFS Policy Schedules

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

Veeam Backup for OLVM and RHV re-uses full backups created according to the backup job schedule to achieve the desired retention for a GFS policy schedule (weekly, monthly and yearly). Each full backup is marked with a flag of a specific GFS policy schedule type: the (W) flag is used to mark full backups for the weekly schedule, (M) – monthly, and (Y) – yearly. Veeam Backup for OLVM and RHV 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 Veeam Backup & Replication User Guide, section [Long-Term Retention Policy \(GFS\)](#).

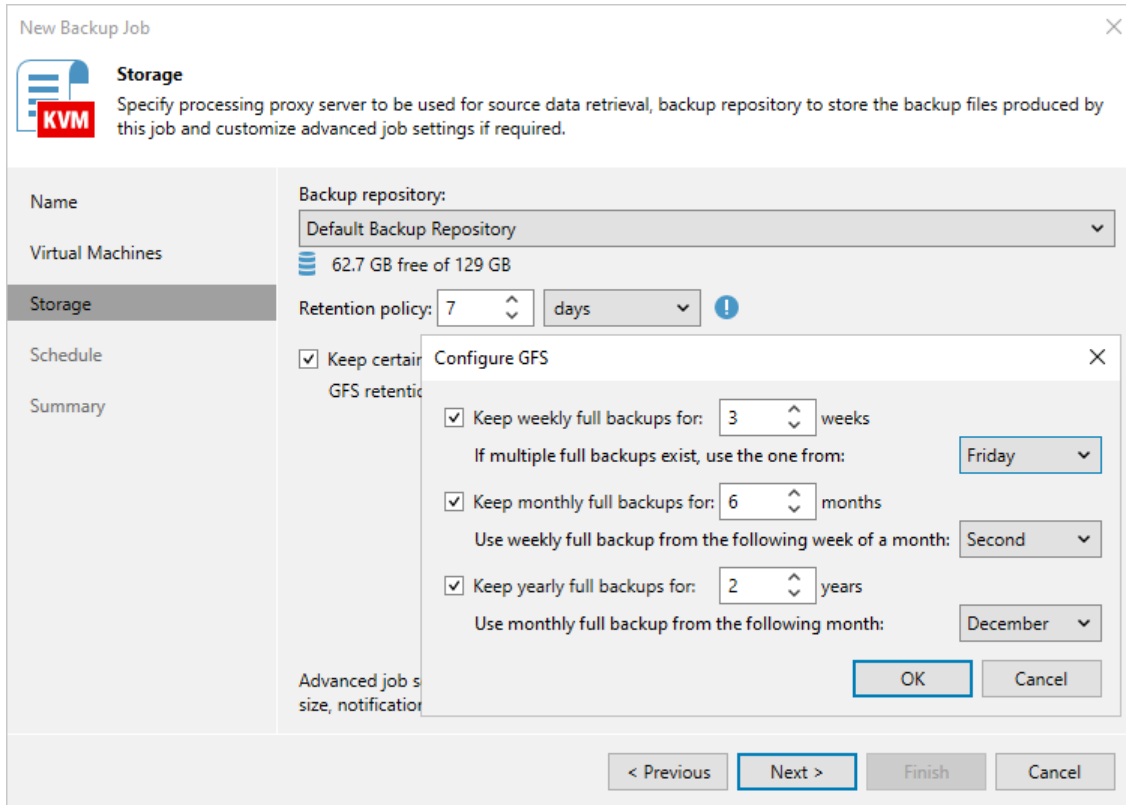
To configure a GFS policy schedule, select the **Keep certain full backups longer for archival purposes** check box and click **Configure**. Then specify the following options in the **Configure GFS** window:

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

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

NOTE

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



Configuring Advanced Settings

To configure backup job advanced settings, do the following:

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

IMPORTANT

- Synthetic full backups cannot be scheduled if an object storage repository is selected as the target location for backups.
- Schedule synthetic full backups to run on days when the backup job is scheduled. Otherwise, no synthetic full backup will be created.

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

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

IMPORTANT

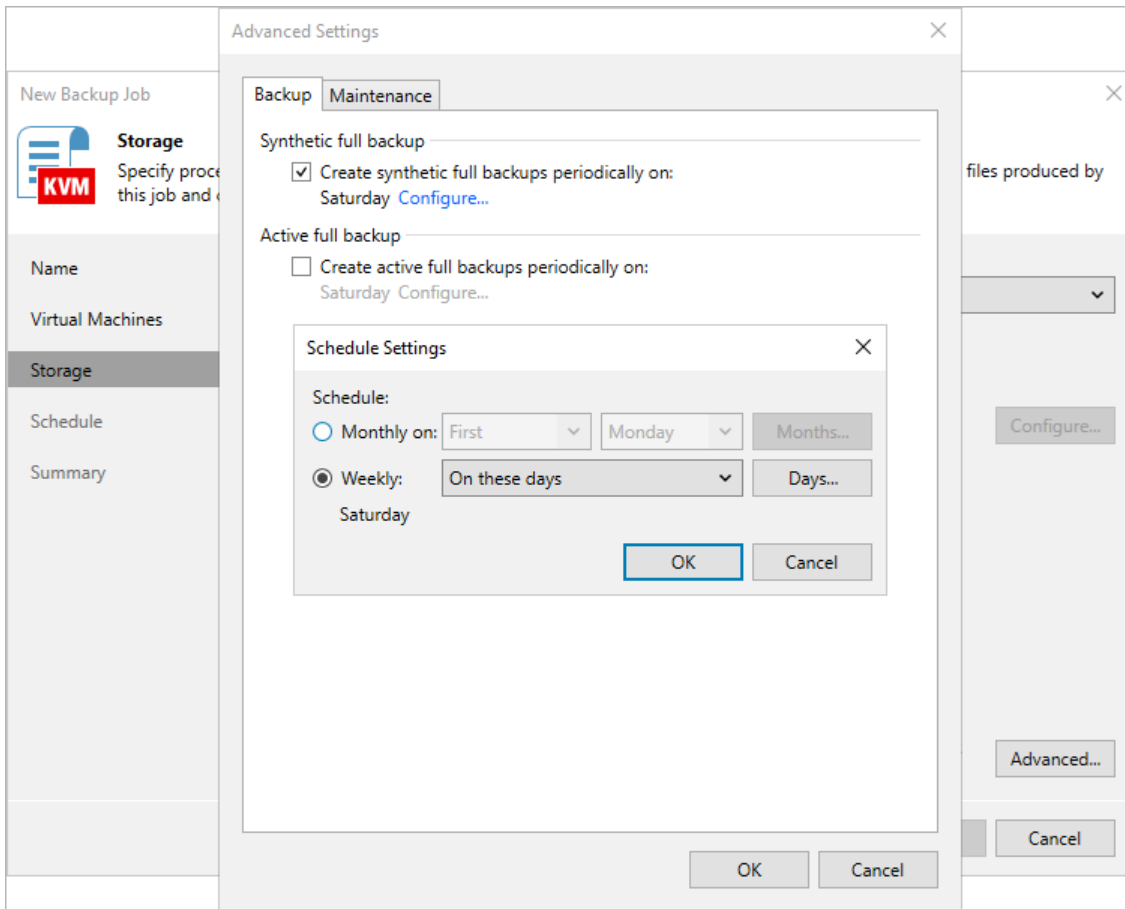
Do not schedule synthetic and active full backups to run at the same time. Due to technical limitations, Veeam Backup for OLVM and RHV will be unable to create synthetic full backups according to the specified schedule.

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

IMPORTANT

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

- To configure retention settings for backups of VMs that are no longer processed by the backup job, on the **Maintenance** tab of the **Advanced settings** window, select the **Remove deleted items data after** check box, and specify the number of days during which Veeam Backup for OLVM and RHV will keep backups of VMs excluded from the job.



How Health Check Works

When Veeam Backup for OLVM and RHV saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup & Replication verifies the availability of data blocks and uses the saved values to ensure that the restore points being verified are consistent.

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

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

- If the health check detects corrupted metadata in a full or incremental restore point, Veeam Backup & Replication marks the backup chain as corrupted in the configuration database. During the next backup job session, Veeam Backup for OLVM and RHV 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 full or an incremental restore point, Veeam Backup for OLVM and RHV marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup job session, Veeam Backup for OLVM and RHV 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. Define Job Schedule

At the **Schedule** step of the wizard, you can instruct Veeam Backup for OLVM and RHV 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.

Veeam Backup for OLVM and RHV allows you to create schedules of the following types:

- **Daily at this time** – the backup job will create restore points at a specific time on specific days.
- **Monthly at this time** – the backup job will create restore points once a month on a specific day.
- **Periodically every** – the backup job will create restore points repeatedly with a specific time interval every day.

TIP

You can instruct Veeam Backup for OLVM and RHV 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 backup job and the time interval between retries. When retrying backup jobs, Veeam Backup for OLVM and RHV processes only those VMs that failed to be backed up during the previous attempt.

New Backup Job [Close]

Schedule
Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.

Run the job automatically

Daily at this time: 10:00 PM [Dropdown] Everyday [Dropdown] [Days...]

Monthly at this time: 10:00 PM [Dropdown] Fourth [Dropdown] Saturday [Dropdown] [Months...]

Periodically every: 1 [Dropdown] Hours [Dropdown]

Automatic retry

Retry failed items processing: 3 [Dropdown] times

Wait before each retry attempt for: 10 [Dropdown] minutes

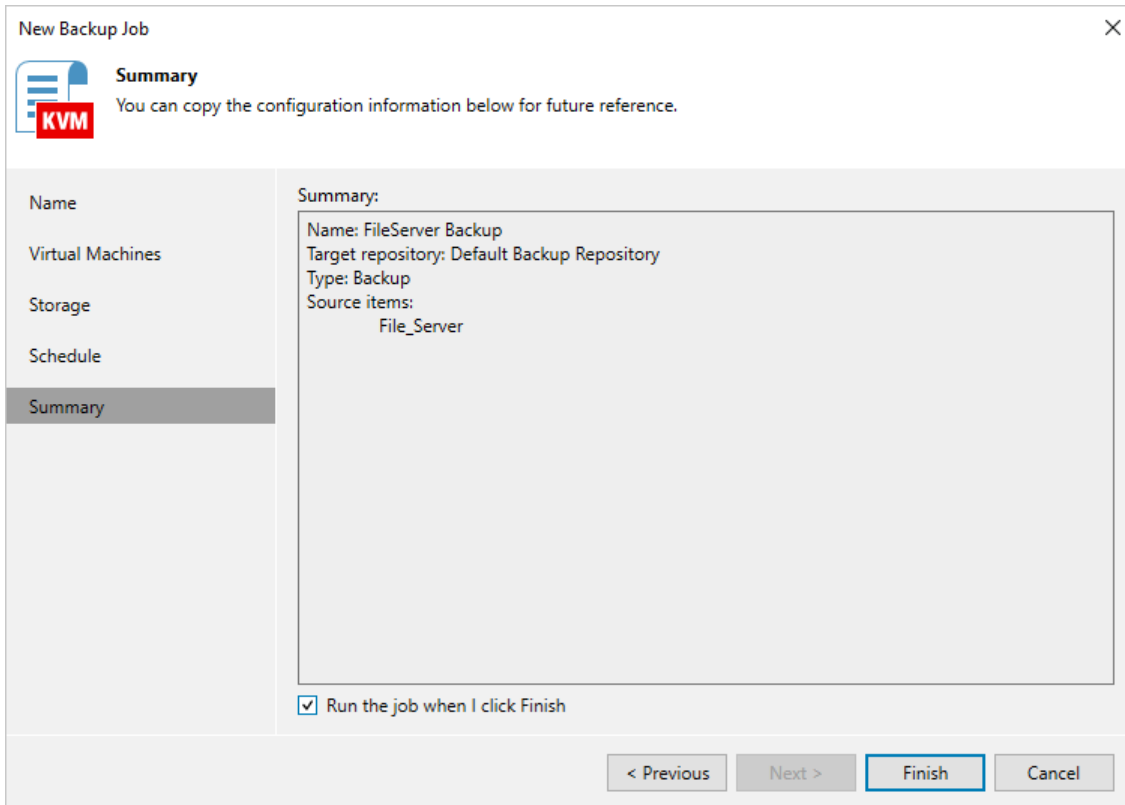
< Previous Apply Finish Cancel

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. As soon as Veeam Backup for OLVM and RHV starts the job, the backup progress will be displayed in the working area when you navigate to **Jobs > Backups** in the inventory pane of the **Home** view.

TIP

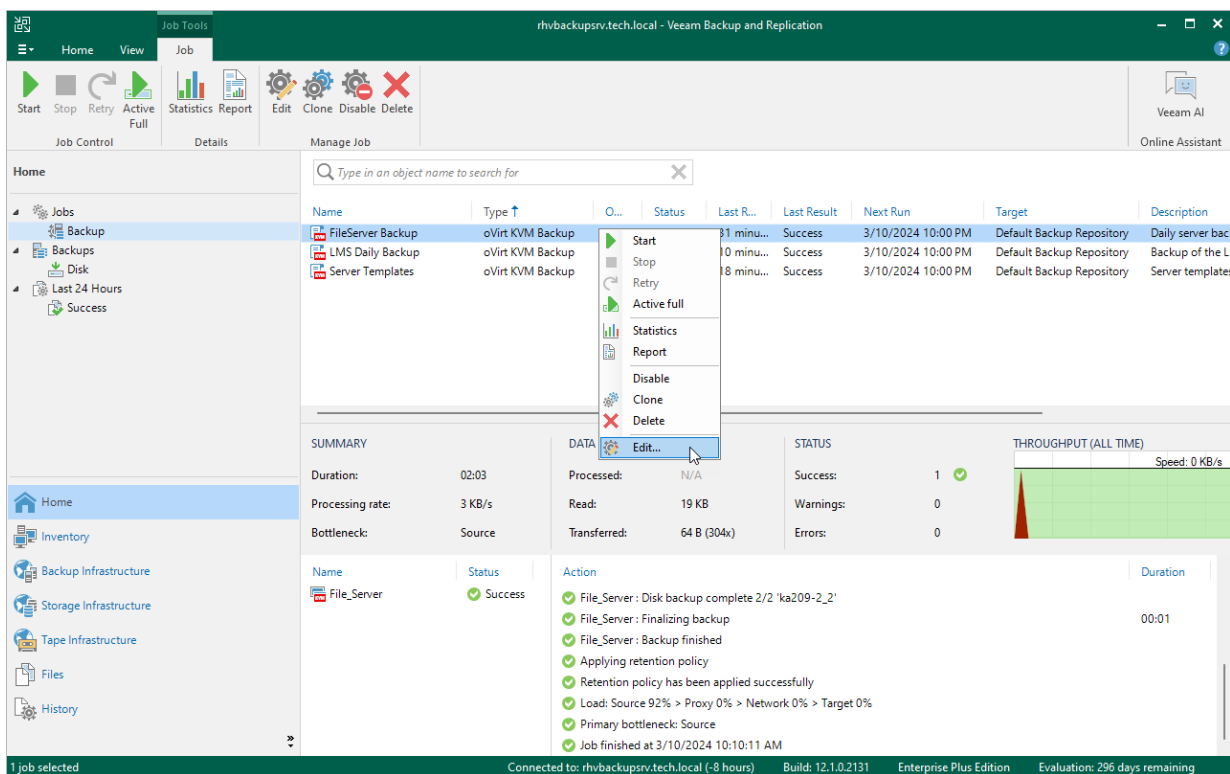
If you want to start the job immediately, select the **Run the job when I click Finish** check box and then click **Finish**.



Editing Backup Job Settings

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

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Edit** on the ribbon, or 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, and to configure health checks, follow the instructions provided in section [Creating Backup Jobs](#) (step 4).
 - d. To modify the job schedule and configure automatic retry settings, follow the instructions provided in section [Creating Backup Jobs](#) (step 5).
 - e. At the **Summary** step of the wizard, review configuration information and click **Finish**.

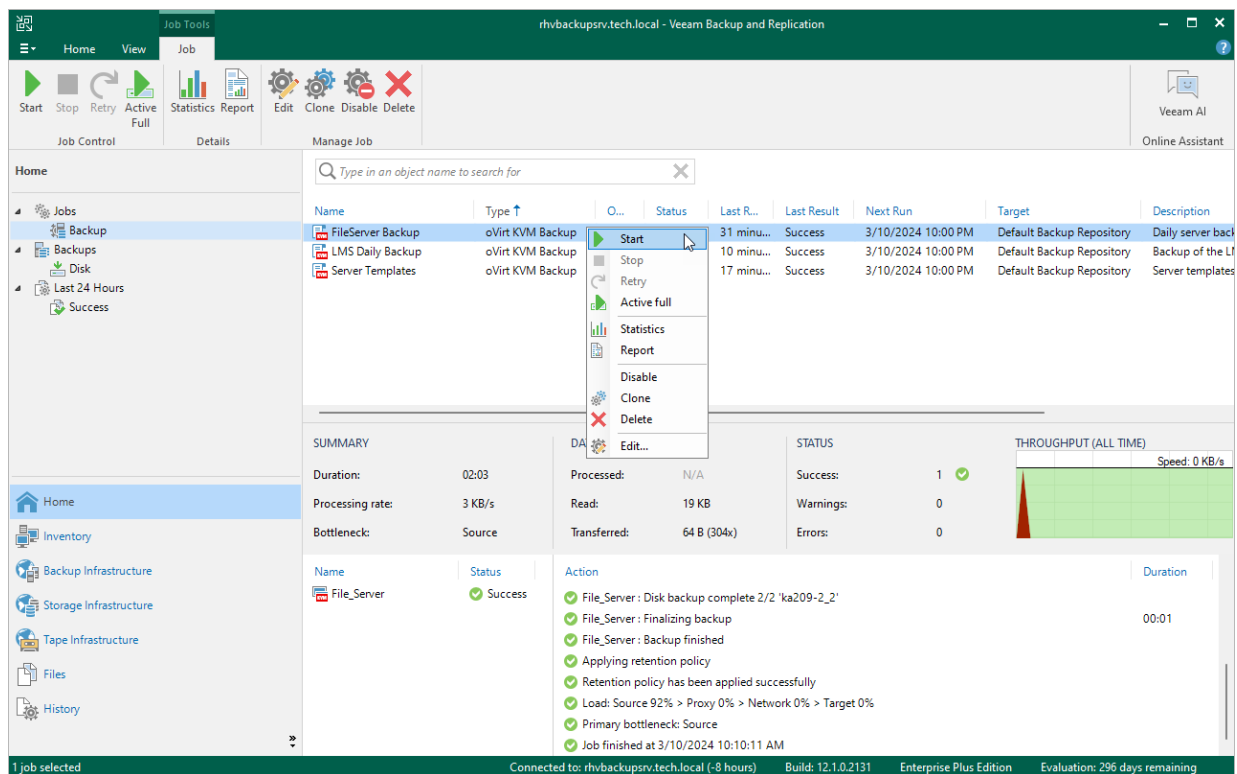


Starting and Stopping Backup Jobs

You can start a backup 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 backup job manually if processing of an oVirt VM is about to take too long, and you do not want the job to have an impact on the production environment during business hours. When you stop a running job, Veeam Backup for OLVM and RHV 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 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 **Start** or **Stop** on the ribbon, or right-click the job and select **Start** or **Stop**.



The screenshot displays the Veeam Backup and Replication console interface. The top ribbon shows the 'Job' view with buttons for Start, Stop, Retry, Active Full, Statistics, Report, Edit, Clone, Disable, and Delete. The main area shows a list of backup jobs. The 'FileServer Backup' job is selected, and a context menu is open over it, with the 'Start' option highlighted. Below the job list, a summary and action log are visible. The summary shows a duration of 02:03, a processing rate of 3 KB/s, and a bottleneck at the source. The action log shows the job completed successfully.

Name	Type	O...	Status	Last R...	Last Result	Next Run	Target	Description
FileServer Backup	oVirt KVM Backup		Start	31 minu...	Success	3/10/2024 10:00 PM	Default Backup Repository	Daily server back
LMS Daily Backup	oVirt KVM Backup		Stop	10 minu...	Success	3/10/2024 10:00 PM	Default Backup Repository	Backup of the L
Server Templates	oVirt KVM Backup		Retry	17 minu...	Success	3/10/2024 10:00 PM	Default Backup Repository	Server templates

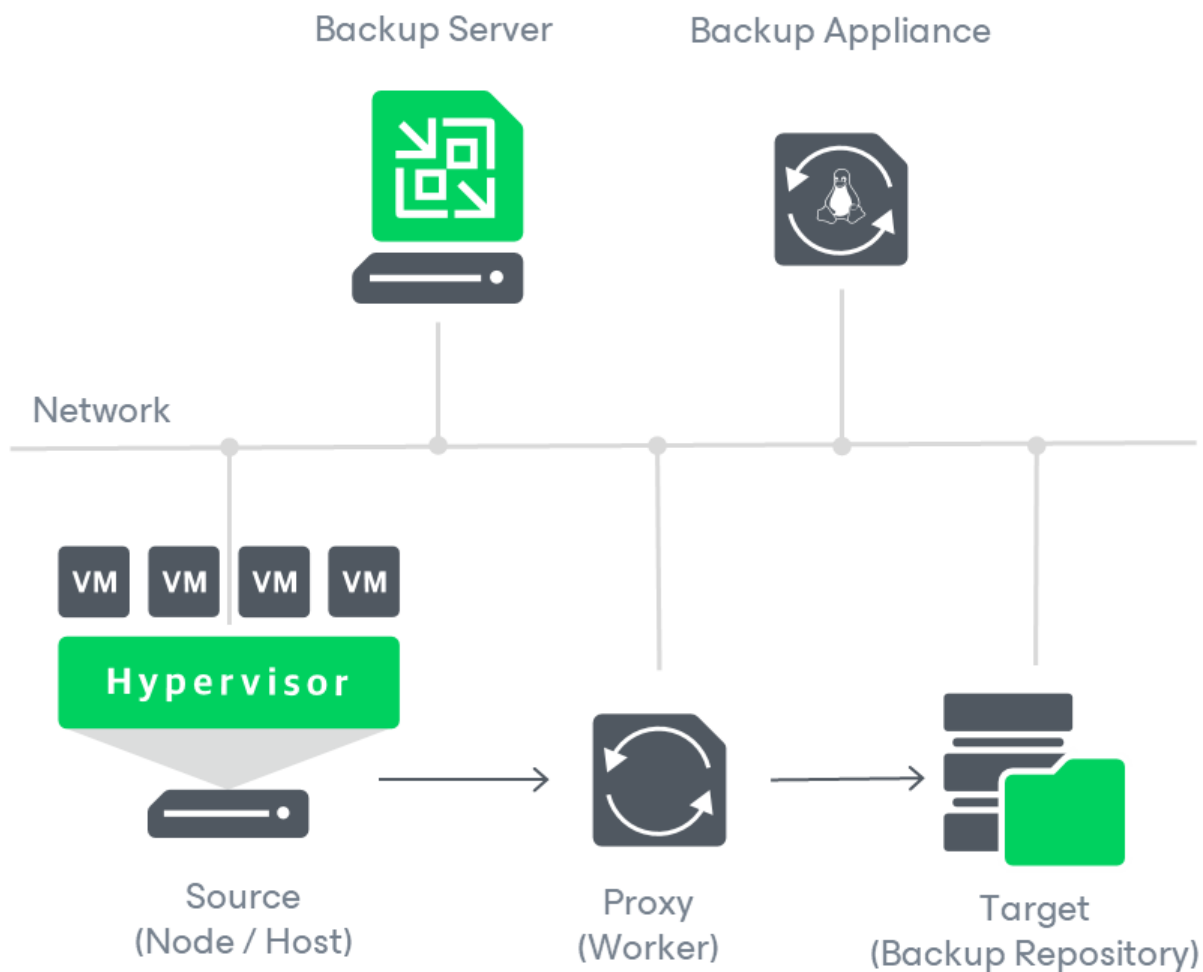
Summary	Processed	Status	Throughput
Duration: 02:03	N/A	Success: 1	Speed: 0 KB/s
Processing rate: 3 KB/s	Read: 19 KB	Warnings: 0	
Bottleneck: Source	Transferred: 64 B (304x)	Errors: 0	

Name	Status	Action	Duration
File_Server	Success	File_Server : Disk backup complete 2/2 'ka209-2_2'	
		File_Server : Finalizing backup	
		File_Server : Backup finished	00:01
		Applying retention policy	
		Retention policy has been applied successfully	
		Load: Source 92% > Proxy 0% > Network 0% > Target 0%	
		Primary bottleneck: Source	
		Job finished at 3/10/2024 10:10:11 AM	

Analyzing Performance Bottlenecks

As any backup application handles a great amount of data, it is important to make sure the data flow is efficient and all resources engaged in the backup process are optimally used. For backup jobs, Veeam provides advanced statistics about the data flow efficiency and lets you identify bottlenecks at the following stages of the data transmission process:

1. Reading VM data blocks from the source.
2. Processing VM data on a worker.
3. Transporting data over the network.
4. Writing data to the target.



While evaluating the data transmission process, Veeam Backup for OLVM and RHV leverages the Veeam Backup & Replication functionality to analyze performance of all the data flow components:

- **Source** – the source disk reader component responsible for retrieving data from the source node.
- **Proxy** – the worker component responsible for processing VM data.
- **Network** – the network queue writer component responsible for getting processed VM data from the worker and sending it over the network to the Target (directly or through the Gateway Server).
- **Target** – the gateway server component responsible for processing VM data, or the target disk writer component responsible for storing data in the backup repository.

To see the bottleneck statistics for a job or a specific VM processed by the job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select a backup job for which you want to see the bottleneck statistics and click **Statistics** on the ribbon, or right-click the job and select **Statistics**.
4. In the job session details window, check the **Bottleneck** field in the **SUMMARY** column.

TIP

To see the bottleneck statistics for a specific VM, click **Show Details**, select the VM name in the **Name** column and check the **Load** record in the **Action** column. To learn how to analyze the statistics, see Veeam Backup & Replication User Guide, section [Performance Bottlenecks](#).

The screenshot displays the Veeam Backup & Replication console interface. The main window shows the 'LMS Daily Backup' job details. The job progress is 100% complete for 2 of 2 VMs. The SUMMARY section includes:

SUMMARY		DATA		STATUS	
Duration:	02:05	Processed:	93 MB (100%)	Success:	2 ✓
Processing rate:	1 MB/s	Read:	93 MB	Warnings:	0
Bottleneck:	Source	Transferred:	3 KB (>999x)	Errors:	0

The THROUGHPUT (ALL TIME) section shows a speed of 9 MB/s. Below this is a table of actions:

Name	Status	Action	Duration
LMS_Server	Success	✓ LMS_Server : Creating incremental KVM backup	
Media_Storage	Success	✓ Media_Storage : Backup in progress	01:22
		✓ Media_Storage : Backing up disk 1/2 'KA209_1'	
		✓ Media_Storage : Disk backup complete 1/2 'KA209_1'	
		✓ Media_Storage : Backing up disk 2/2 'KA209_2'	
		✓ Media_Storage : Disk backup complete 2/2 'KA209_2'	
LMS_Server	Success	✓ LMS_Server : Backup in progress	01:19
		✓ LMS_Server : Backing up disk 1/2 'ka209-2_1'	
		✓ LMS_Server : Disk backup complete 1/2 'ka209-2_1'	
		✓ LMS_Server : Backing up disk 2/2 'ka209-2_2'	
		✓ LMS_Server : Disk backup complete 2/2 'ka209-2_2'	
Media_Storage	Success	✓ Media_Storage : Finalizing backup	00:00
Media_Storage	Success	✓ Media_Storage : Backup finished	
LMS_Server	Success	✓ LMS_Server : Finalizing backup	00:01
LMS_Server	Success	✓ LMS_Server : Backup finished	

The interface also shows a 'Hide Details' button and an 'OK' button. The status bar at the bottom indicates '1 job selected', 'Connected to: rhvbackupsrv.tech.local (-8 hours)', 'Build: 12.1.0.2131', 'Enterprise Plus Edition', and 'Evaluation: 296 days remaining'.

Cloning Backup 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.

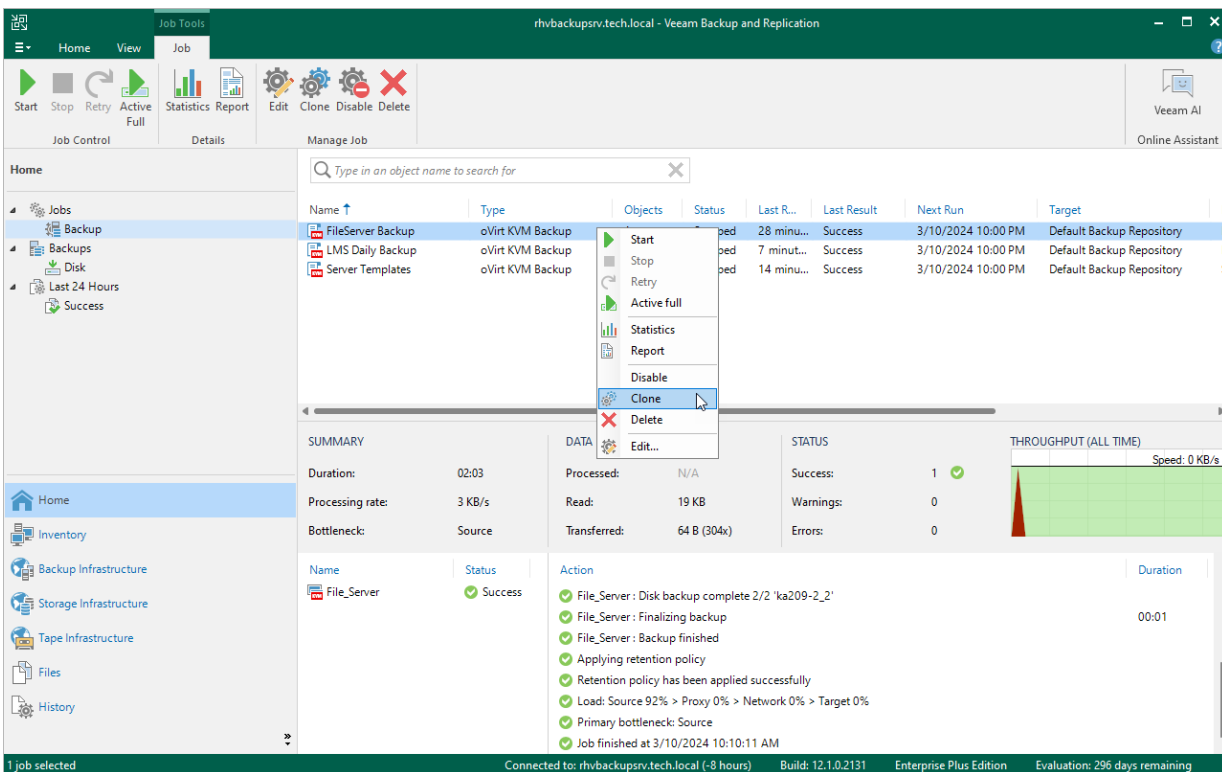
To clone a job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Clone** on the ribbon, or 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 Backup for OLVM and RHV disables the cloned job. To enable the cloned job, select it in the job list and click **Enable**.



The screenshot shows the Veeam Backup & Replication console interface. The top navigation bar includes 'Home', 'View', and 'Job' tabs. The 'Job' tab is active, showing a ribbon with options: Start, Stop, Retry, Active Full, Statistics, Report, Edit, Clone, Disable, and Delete. The 'Clone' option is highlighted. Below the ribbon, a table lists jobs with columns for Name, Type, Objects, Status, Last Run, Last Result, Next Run, and Target. A context menu is open over the 'FileServer Backup' job, with 'Clone' selected. The bottom section of the console displays job details for 'FileServer Backup', including a SUMMARY table with Duration (02:03), Processing rate (3 KB/s), and Bottleneck (Source). It also shows a STATUS section with Success (1), Warnings (0), and Errors (0). A THROUGHPUT (ALL TIME) graph is visible on the right, showing a speed of 0 KB/s. The bottom status bar indicates '1 job selected', 'Connected to: rhvbackupsrv.tech.local (-8 hours)', 'Build: 12.1.0.2131', 'Enterprise Plus Edition', and 'Evaluation: 296 days remaining'.

Enabling and Disabling Backup Jobs

By default, all created backup 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. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs**.
3. In the working area, select the job and click **Enable** or **Disable** on the ribbon, or right-click the job and select **Enable** or **Disable**.

The screenshot shows the Veeam Backup & Replication console interface. The top ribbon includes buttons for Start, Stop, Retry, Active Full, Statistics, Report, Edit, Clone, Disable, and Delete. The left sidebar shows the 'Jobs' view selected. The main area displays a table of jobs with columns for Name, Type, Objects, Status, Last Run, Last Result, Next Run, and Target. A context menu is open over the 'FileServer Backup' job, with the 'Disable' option highlighted. Below the table, there is a SUMMARY section with fields for Duration, Processing rate, and Bottleneck. A DATA section shows Processed, Read, and Transferred values. A STATUS section shows Success, Warnings, and Errors counts. A THROUGHPUT (ALL TIME) graph is also visible. The bottom status bar indicates '1 job selected' and provides connection and build information.

Name	Type	Objects	Status	Last R...	Last Result	Next Run	Target
FileServer Backup	oVirt KVM Backup		Success	27 minut...	Success	3/10/2024 10:00 PM	Default Backup Repository
LMS Daily Backup	oVirt KVM Backup		Success	6 minut...	Success	3/10/2024 10:00 PM	Default Backup Repository
Server Templates	oVirt KVM Backup		Success	14 minut...	Success	3/10/2024 10:00 PM	Default Backup Repository

SUMMARY

Duration: 02:03
Processing rate: 3 KB/s
Bottleneck: Source

DATA

Processed: N/A
Read: 19 KB
Transferred: 64 B (304x)

STATUS

Success: 1
Warnings: 0
Errors: 0

THROUGHPUT (ALL TIME)

Speed: 0 KB/s

Summary Table

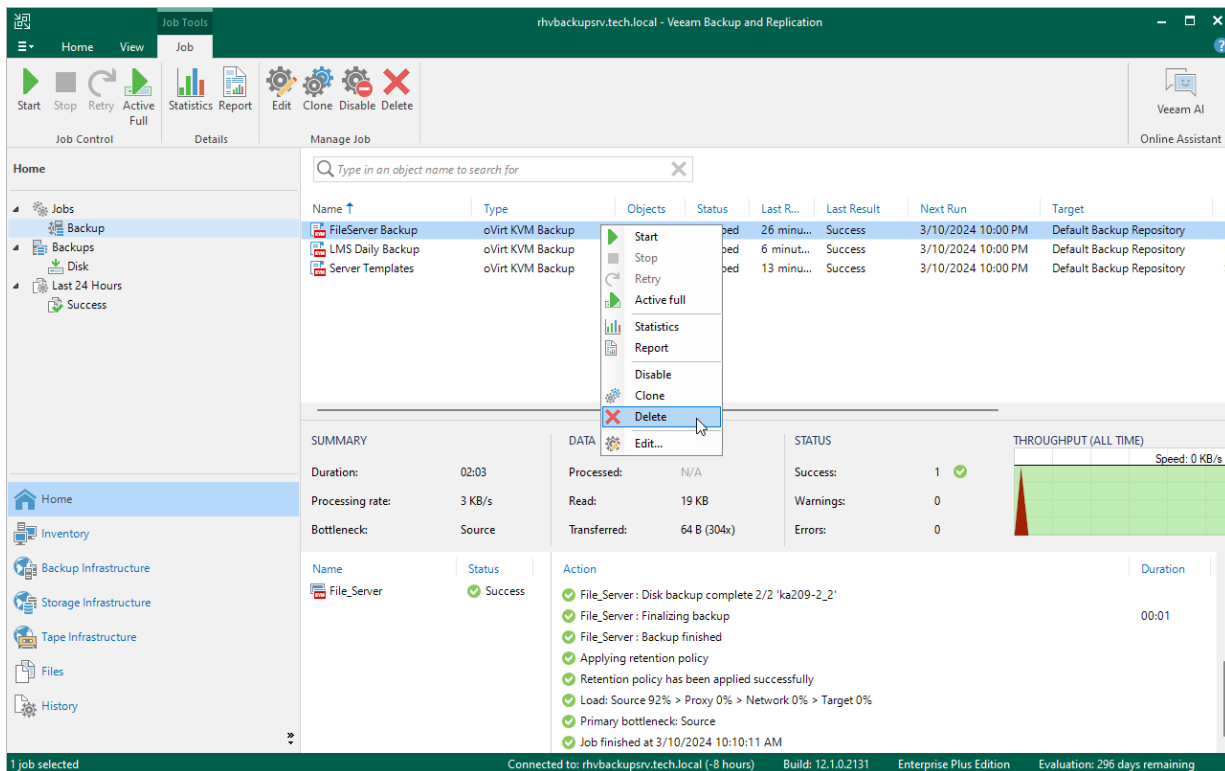
Name	Status	Action	Duration
File_Server	Success	File_Server : Disk backup complete 2/2 'ka209-2_2'	
		File_Server : Finalizing backup	
		File_Server : Backup finished	
		Applying retention policy	
		Retention policy has been applied successfully	
		Load: Source 92% > Proxy 0% > Network 0% > Target 0%	
		Primary bottleneck: Source	
		Job finished at 3/10/2024 10:10:11 AM	00:01

Deleting Backup Jobs

You can permanently delete a backup job from the Veeam Backup for OLVM and RHV 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. In the Veeam Backup & Replication console, 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, or right-click the job and select **Delete**.



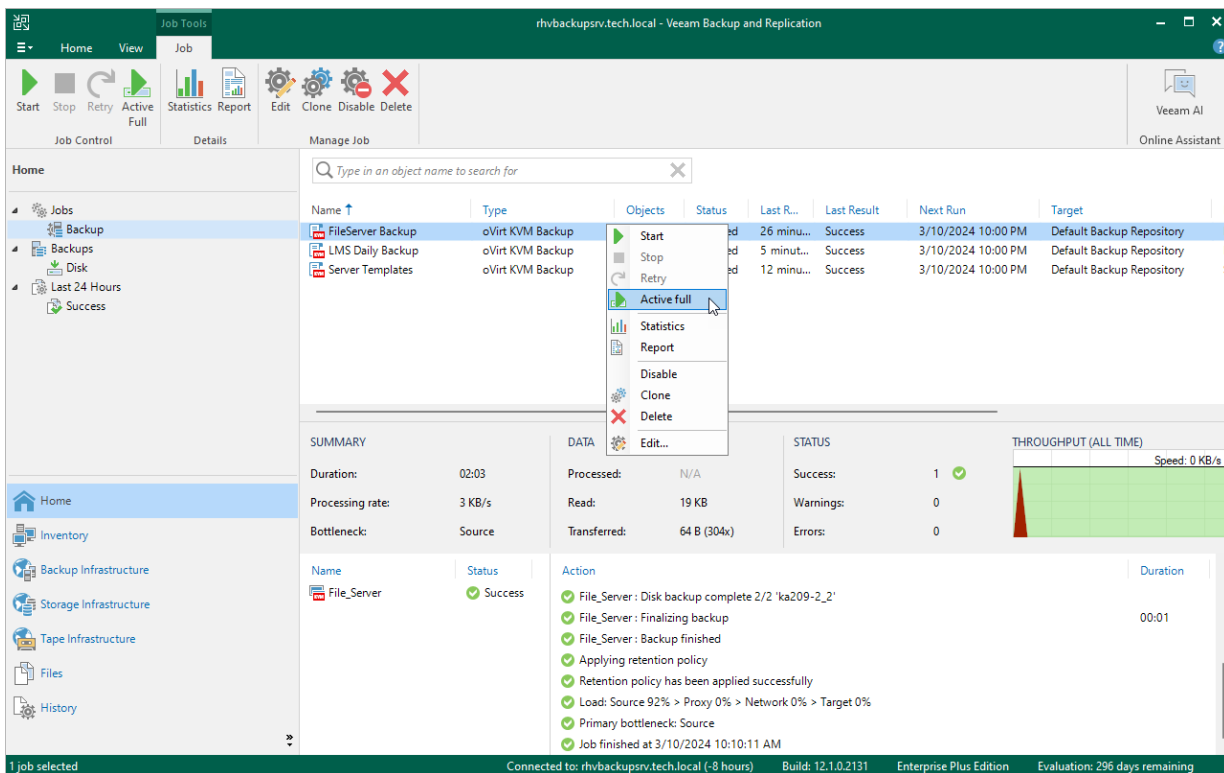
Creating Active Full Backups

You can manually create an [active full backup](#) for all VMs added to a backup job:

1. In the Veeam Backup & Replication console, 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, or 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).



The screenshot displays the Veeam Backup & Replication console interface. The top ribbon includes 'Job Control' (Start, Stop, Retry, Active Full), 'Details' (Statistics, Report), and 'Manage Job' (Edit, Clone, Disable, Delete). The 'Home' view is active, showing a tree view on the left with 'Jobs' selected. The main area shows a table of backup jobs. A context menu is open over the 'FileServer Backup' job, with 'Active full' highlighted. Below the table, a 'SUMMARY' section shows job details: Duration: 02:03, Processing rate: 3 KB/s, Bottleneck: Source. A 'DATA' section shows: Processed: N/A, Read: 19 KB, Transferred: 64 B (304x). A 'STATUS' section shows: Success: 1, Warnings: 0, Errors: 0. A 'THROUGHPUT (ALL TIME)' graph is visible on the right. At the bottom, a 'Job History' section shows a list of actions for the 'File_Server' job, including 'File_Server: Disk backup complete 2/2 'ka209-2_2'', 'File_Server: Finalizing backup', 'File_Server: Backup finished', 'Applying retention policy', 'Retention policy has been applied successfully', 'Load: Source 92% > Proxy 0% > Network 0% > Target 0%', 'Primary bottleneck: Source', and 'Job finished at 3/10/2024 10:10:11 AM'.

Creating VeeamZIP Backups

You can back up one or multiple oVirt 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.

NOTES

- You cannot store VeeamZIP backups in [Veeam Cloud Connect repositories](#).
- Veeam Backup & Replication does not apply network traffic throttling rules to VeeamZIP backup sessions. For more information, see the Veeam Backup & Replication User Guide, section [Configuring Network Traffic Rules](#).

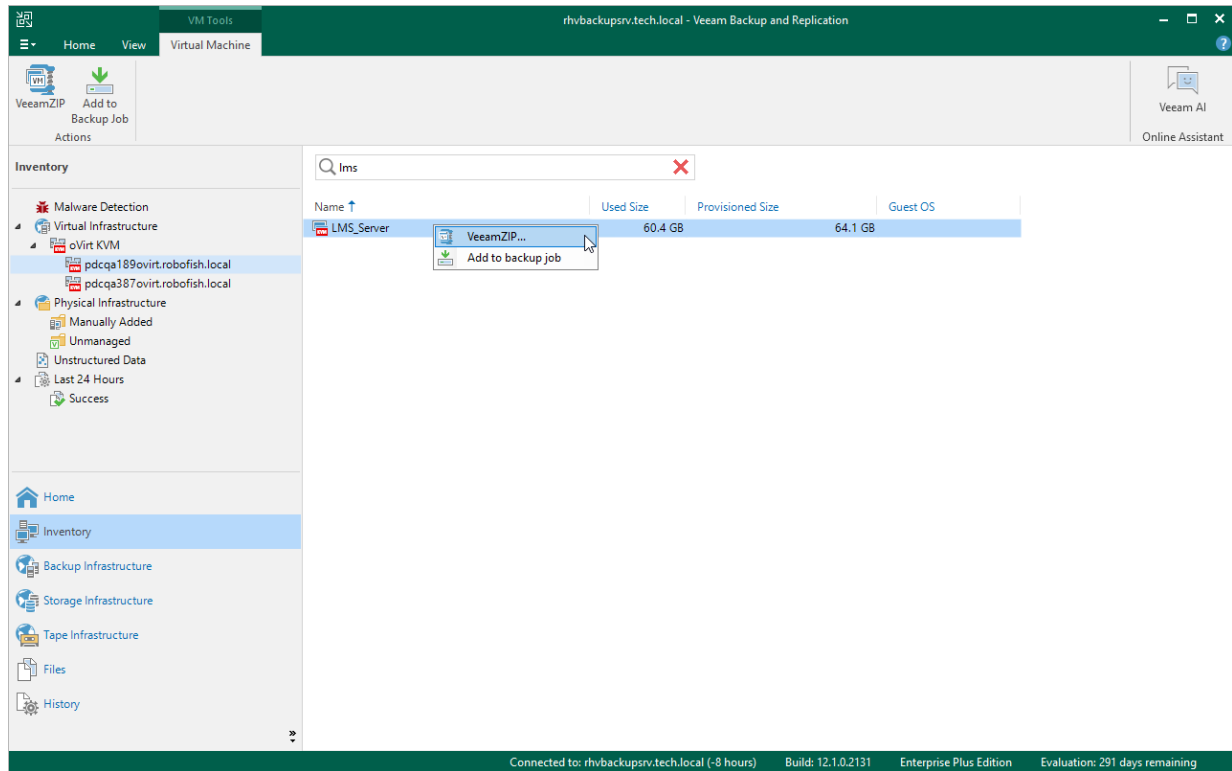
To create a VeeamZIP backup, do the following:

1. In the Veeam Backup & Replication console, open the **Inventory** view.
2. In the inventory pane, select **Virtual Infrastructure > oVirt KVM**.
3. In the working area, select the VM that you want to back up and click **VeeamZIP** on the ribbon, or 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

Veeam Backup for OLVM and RHV stores information on all protected oVirt VMs in the configuration database. Even if a VM is no longer protected by any configured backup job and even if the VM no longer exists in the oVirt KVM environment, records about created backups will not be deleted from the database until Veeam Backup for OLVM and RHV automatically removes all restore points associated with this VM according to the retention settings saved in the backup metadata. You can manage oVirt VM backups as long as their records are present in the configuration database.

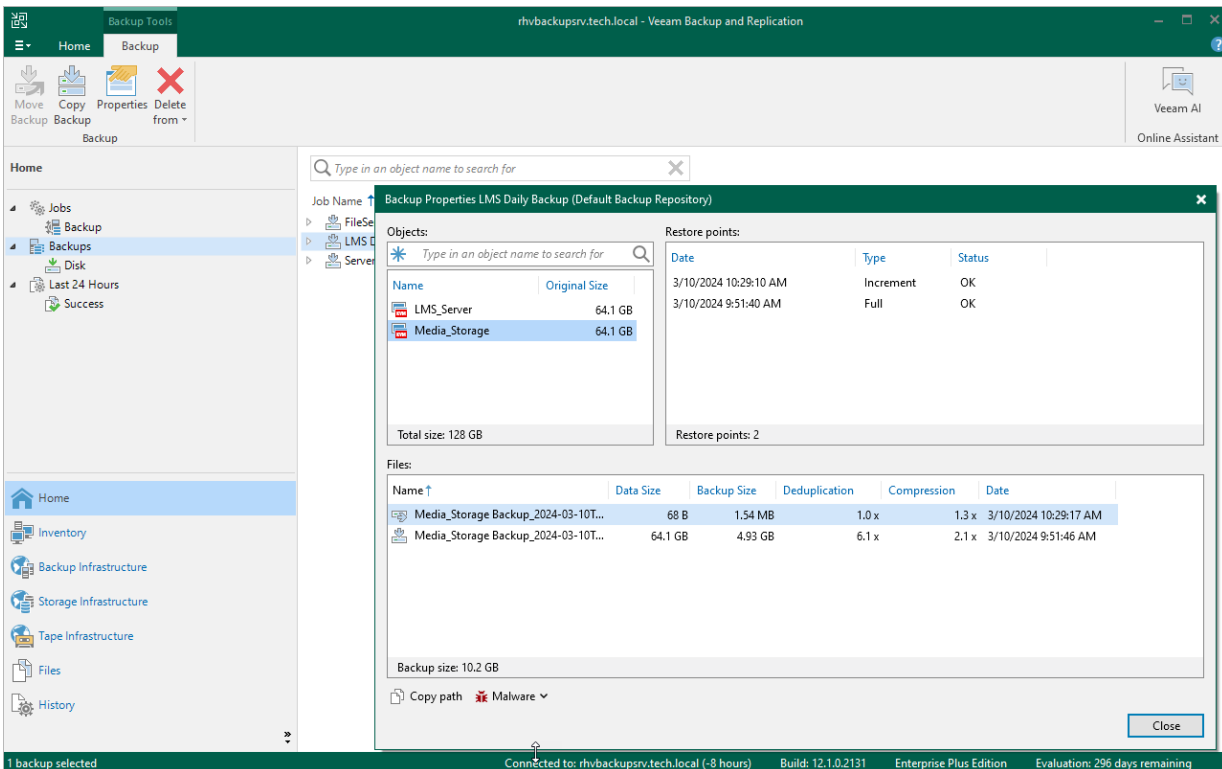
Viewing Backup Properties

After a backup job successfully creates a backup of an oVirt VM according to the specified schedule, or after you create an active full backup of a VM manually, the backup is displayed under the **Backups** node in the **Home** view of the Veeam Backup & Replication console. Each backup is represented with a set of properties, such as:

- **Objects** – the names and sizes of backed-up VMs.
- **Restore Points** – the date and time of all restore points created for a VM.
- **Files** – the size of processed VM data, the size of backed-up VM data, the ratio of [data deduplication](#) and the ratio of [data compression](#).

To view backup properties, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the backup and select **Properties**.



Verifying Backups

To perform an integrity check of oVirt VM backups, Veeam Backup & Replication offers the SureBackup technology that allows you to ensure that the created restore points are not corrupted. For backups of Windows VMs, you can also scan the restore points with antivirus software installed on the backup server, and run YARA rules to detect malware and sensitive data.

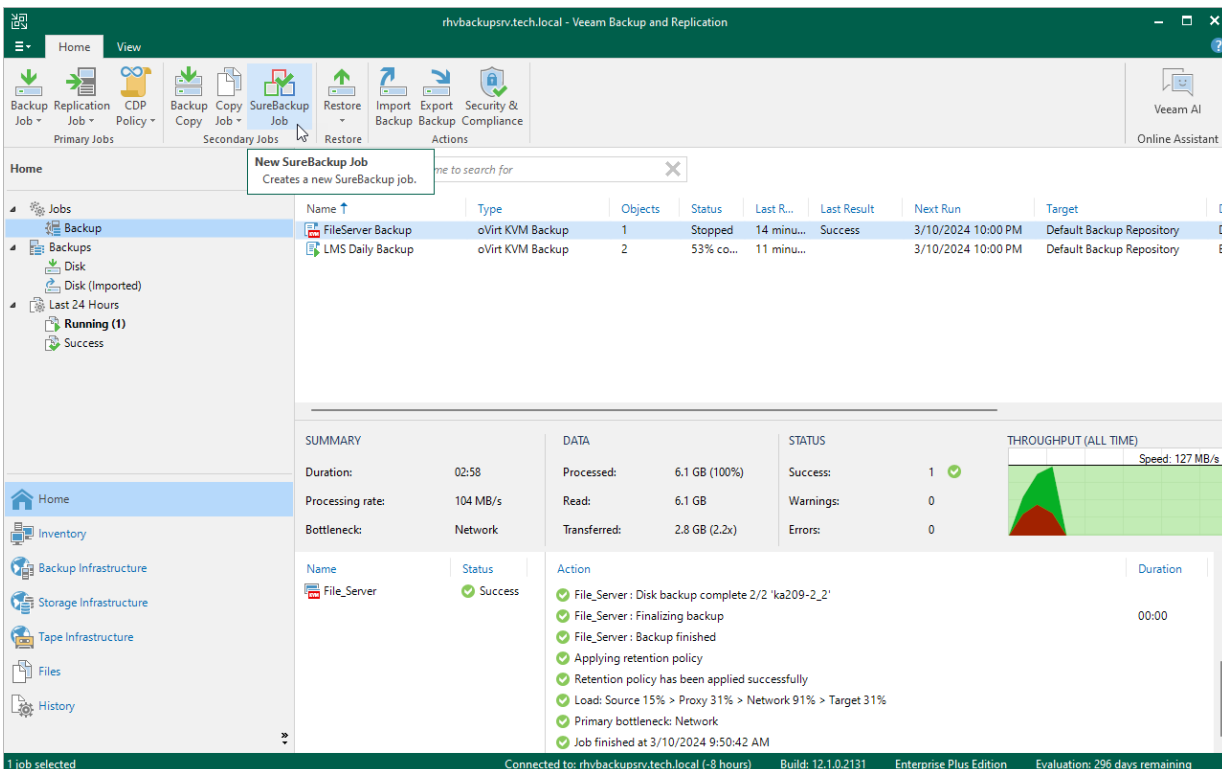
To create a SureBackup job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Jobs > Backup** and click **SureBackup Job** on the ribbon.
3. At the **Name** step of the **New SureBackup Job** wizard, select the **Backup verification and content scan only verification** mode, and then complete the wizard as described in the Veeam Backup & Replication User Guide, section [Creating SureBackup Jobs](#).

If any of the verification checks fail for a restore point, Veeam Backup & Replication will mark both this restore point and all subsequent points in the backup chain as *Infected*. To learn how to manage infected restore points, see Veeam Backup & Replication User Guide, section [Managing Malware Status](#).

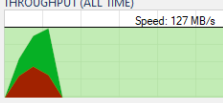
TIP

You can scan backups of Windows VMs manually on demand, without creating a SureBackup job. To learn how to do that, see the Veeam Backup & Replication User Guide, section [Scan Backup](#).



The screenshot displays the Veeam Backup & Replication console interface. The main window shows a list of jobs under the 'Backup' category. A 'New SureBackup Job' dialog box is open, indicating the creation of a new job. Below the job list, a summary and details for a specific job are shown.

Name	Type	Objects	Status	Last R...	Last Result	Next Run	Target
FileServer Backup	oVirt KVM Backup	1	Stopped	14 minu...	Success	3/10/2024 10:00 PM	Default Backup Repository
LMS Daily Backup	oVirt KVM Backup	2	53% co...	11 minu...		3/10/2024 10:00 PM	Default Backup Repository

SUMMARY		DATA		STATUS		THROUGHPUT (ALL TIME)
Duration:	02:58	Processed:	6.1 GB (100%)	Success:	1	
Processing rate:	104 MB/s	Read:	6.1 GB	Warnings:	0	
Bottleneck:	Network	Transferred:	2.8 GB (2.2x)	Errors:	0	

Name	Status	Action	Duration
File_Server	Success	File_Server : Disk backup complete 2/2 'ka209-2_2'	
		File_Server : Finalizing backup	
		File_Server : Backup finished	00:00
		Applying retention policy	
		Retention policy has been applied successfully	
		Load: Source 15% > Proxy 31% > Network 91% > Target 31%	
		Primary bottleneck: Network	
		Job finished at 3/10/2024 9:50:42 AM	

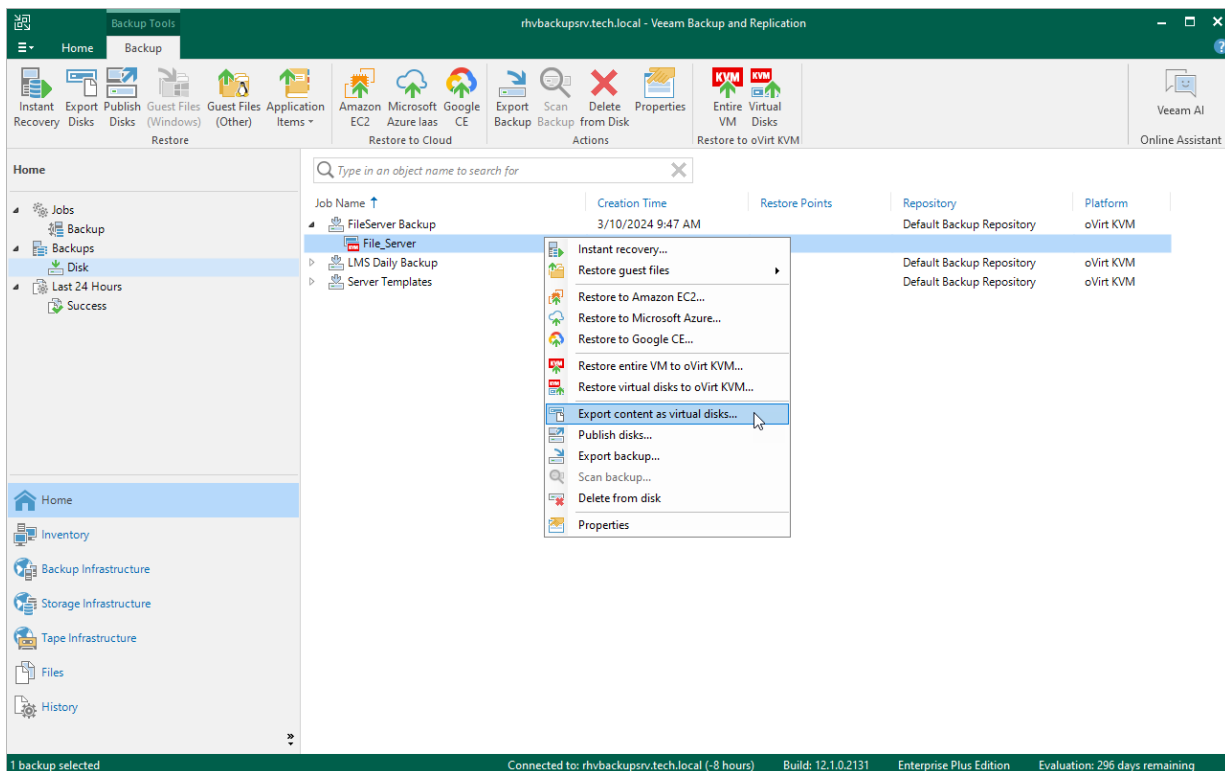
Exporting Backups

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 a repository or folder.

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, expand the job that created the backup, 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 > Backup** and click **Backup Copy** on the ribbon.
3. Create a backup copy job as described in the Veeam Backup & Replication User Guide, section [Creating Backup Copy Jobs](#).

Note that for backup copies, you can also use [Veeam Cloud Connect repositories](#) if a service provider is added to Veeam Backup & Replication.

TIP

Alternatively, you can create a copy of a backup without configuring a job as described in the Veeam Backup & Replication User Guide, section [Copying Backups](#).

The screenshot displays the Veeam Backup & Replication console interface. The top ribbon shows the 'Backup Copy' option selected. The main area shows a table of jobs with the following data:

Name	Type	Objects	Status	Last R...	Last Result	Next Run	Target
FileServer Backup	oVirt KVM Backup	1	Stopped	15 minu...	Success	3/10/2024 10:00 PM	Default Backup Repository
LMS Daily Backup	oVirt KVM Backup	2	56% co...	11 minu...		3/10/2024 10:00 PM	Default Backup Repository

Below the table, a summary of the job is shown:

SUMMARY	DATA	STATUS	THROUGHPUT (ALL TIME)
Duration: 02:58	Processed: 6.1 GB (100%)	Success: 1	Speed: 127 MB/s
Processing rate: 104 MB/s	Read: 6.1 GB	Warnings: 0	
Bottleneck: Network	Transferred: 2.8 GB (2.2x)	Errors: 0	

The bottom section shows a list of actions for the File_Server job:

Name	Status	Action	Duration
File_Server	Success	File_Server : Disk backup complete 2/2 'ka209-2_2'	
		File_Server : Finalizing backup	00:00
		File_Server : Backup finished	
		Applying retention policy	
		Retention policy has been applied successfully	
		Load: Source 15% > Proxy 31% > Network 91% > Target 31%	
		Primary bottleneck: Network	
		Job finished at 3/10/2024 9:50:42 AM	

Copying Backups to Tapes

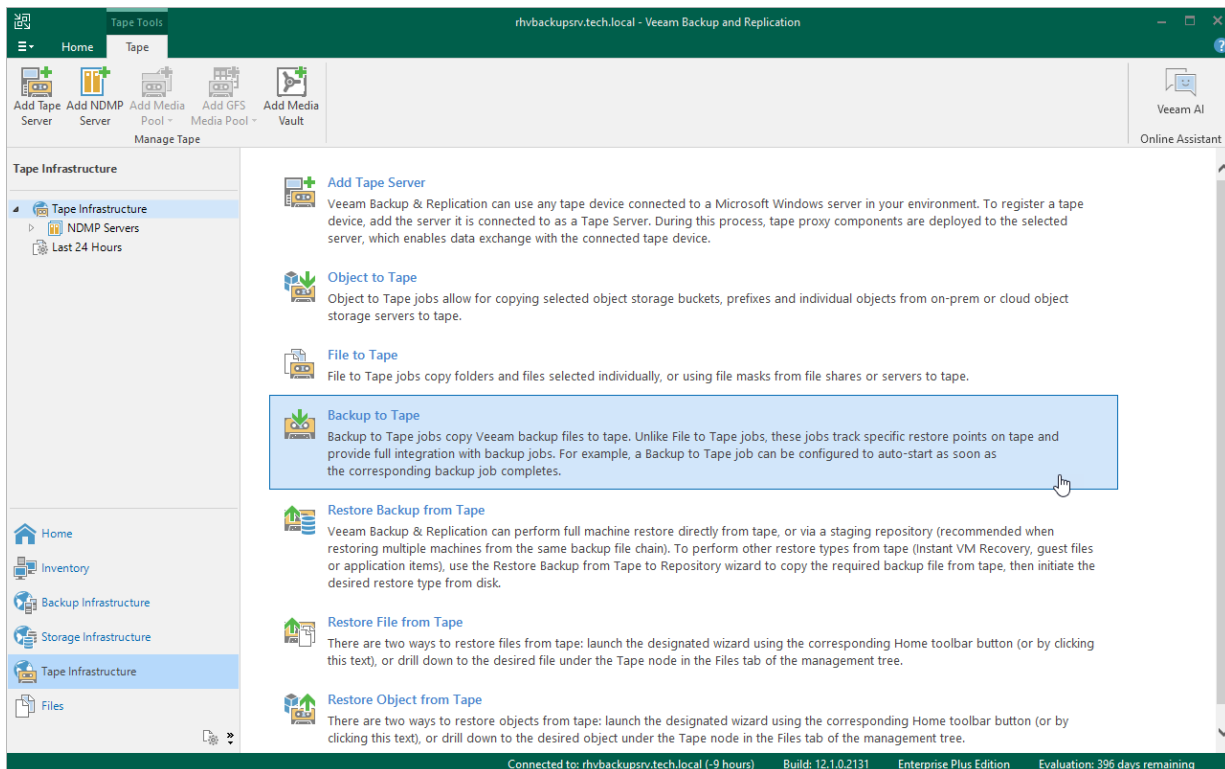
You can create archives of oVirt VM backups and copy them to tapes for long-term storage. Veeam Backup for OLVM and RHV allows you to manage tape archives the same way you manage backups in backup repositories. However, it usually takes more time to access archived data on tapes than to access backed-up data in repositories. For more information on tapes, see the Veeam Backup & Replication User Guide, section [Tape Devices Support](#).

To archive oVirt 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 oVirt VMs directly from tapes. To restore an oVirt VM, you must first restore its backups to a repository as described in the Veeam Backup & Replication User Guide, section [Backup Restore from Tape to Repository](#).



Deleting Backups

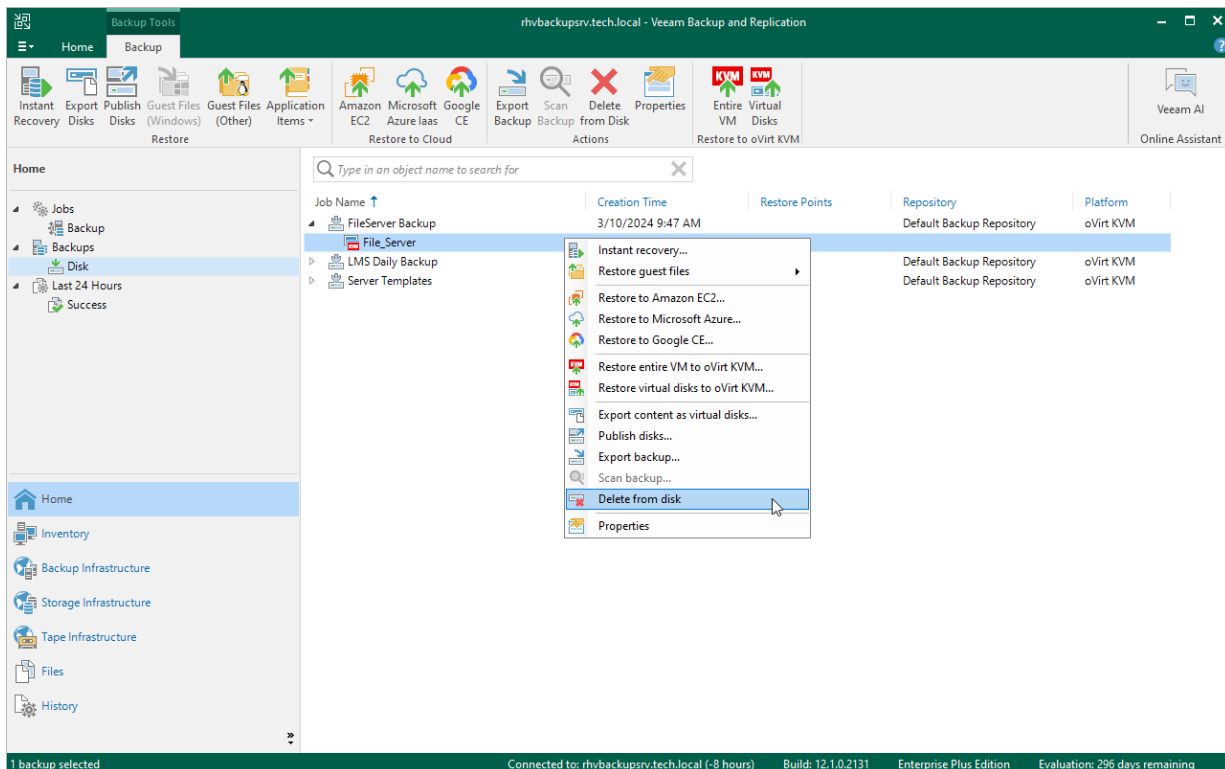
By default, Veeam Backup for OLVM and RHV maintains backups stored in backup repositories according to retention policy settings saved in the backup metadata. If Veeam Backup for OLVM and RHV detects that the number of restore points in the backup chain exceeds the allowed number, it automatically removes obsolete backups. You can also delete backup files from backup repositories manually if you no longer need them.

To delete backup files created for an oVirt VM, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane of the **Home** view, select **Backups**.
3. In the working area, expand the job that created the backup, right-click the VM name and select **Delete from disk**.

NOTE

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 Backup for OLVM and RHV allows you to perform the following operations using backed-up data:

- [Entire VM restore](#) – recover oVirt 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 an oVirt VM directly from a backup.
- [Disk publishing](#) – mount specific disks of a backed-up oVirt 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 oVirt VMs to Amazon Web Services as EC2 instances.
- [Restore to Microsoft Azure](#) – restore oVirt VMs to Microsoft Azure as Azure VMs.
- [Restore to Google Cloud](#) – restore oVirt VMs to Google Cloud as VM instances.

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

Performing VM Restore

In case of a disaster, you can restore an entire oVirt VM from a backup. Veeam Backup for OLVM and RHV allows you to restore one or more VMs at a time, to the original location or to a new location.

VM restore is supported only for backups stored in backup repositories, object storage repositories, Veeam Cloud Connect 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).

NOTE

You cannot restore VMs from backups stored in external repositories and on tapes. However, you can copy backups to a supported repository and then use them to restore VMs.

How VM Restore Works

During the VM restore process, the following steps are performed:

1. The Veeam Backup & Replication console sends the restore session configuration data to the backup appliance.
If multiple VMs are added to the restore session, these VMs are processed in parallel.
2. [This step applies only if you perform restore to the original location and if the source VM is still present in the location] The backup appliance powers off the source VM and removes it from the oVirt KVM environment.
3. The backup appliance launches a worker.
4. The worker connects to the Virtualization manager over REST API and creates a VM in the target location.
5. The worker creates empty virtual disks in the target location. The number of empty disks equals the number of disks attached to the source VM.
6. The worker connects to the backup repository and restores backed-up data to the empty disks.
If multiple disks are attached to the source VM, the worker restores these disks sequentially, one disk at a time.
7. The worker attaches the created disks with the restored data to the VM.

How to Perform VM Restore

To restore a protected VM, do the following:

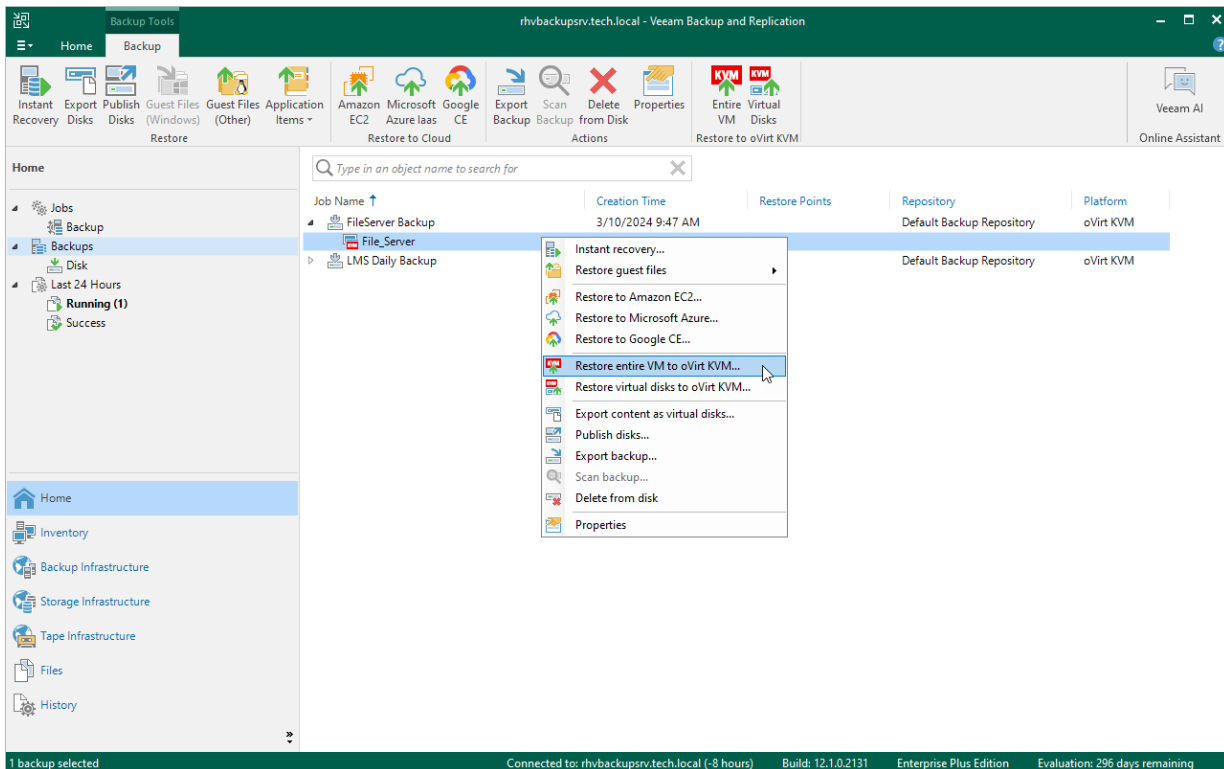
1. [Launch the Full VM Restore to oVirt KVM wizard.](#)
2. [Select a restore point.](#)
3. [Choose a restore mode.](#)
4. [Specify a target cluster.](#)
5. [Select a storage domain where VM virtual disks will be stored.](#)
6. [Specify a name for the restored VM.](#)
7. [Configure network settings.](#)

8. Specify a restore reason.
9. Verify restore settings.

Step 1. Launch Full VM Restore to oVirt KVM Wizard

To launch the **Full VM Restore to oVirt KVM** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup, select the VM that you want to restore and click **Entire VM** on the ribbon, or right-click the VM and select **Restore entire oVirt KVM**.



Step 2. Select Restore Point

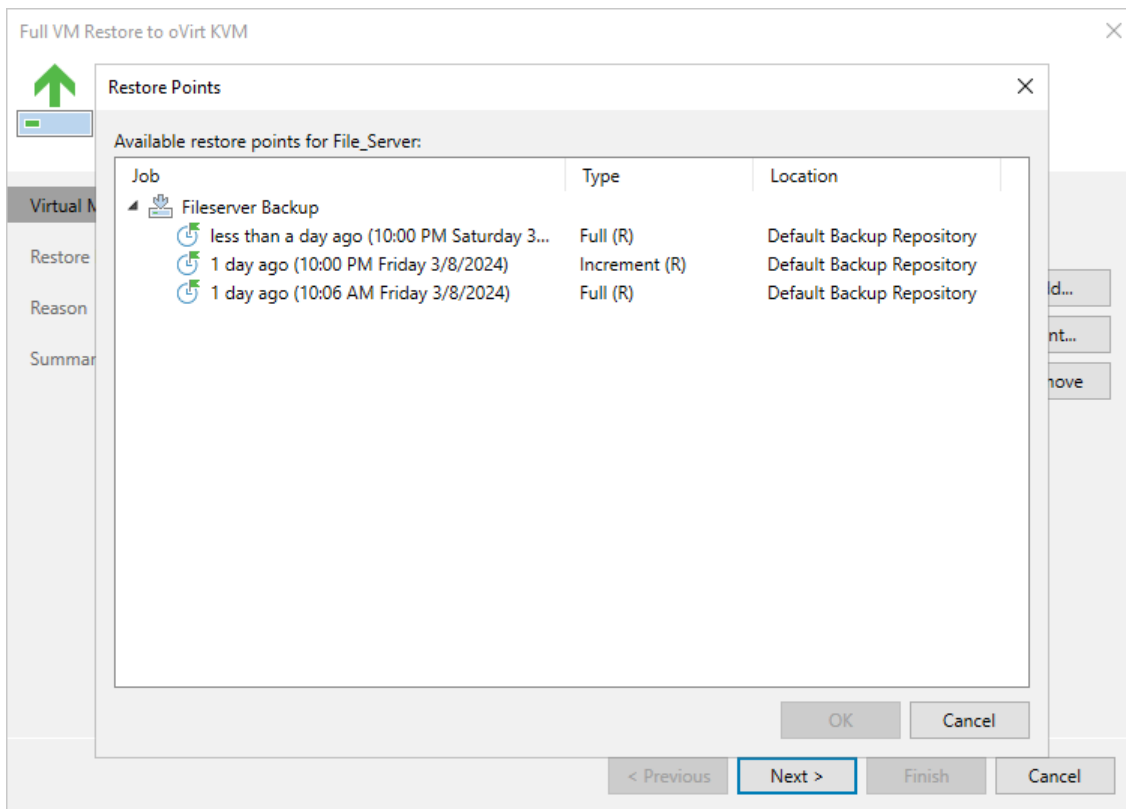
At the **Virtual Machines** step of the wizard, select a restore point that will be used to restore the selected VM. By default, Veeam Backup for OLVM and RHV uses the most recent valid restore point. However, you can restore the VM data to an earlier state.

To select a restore point, do the following:

1. Select the VM.
2. Click **Point**.
3. In the **Restore Points** window, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup for OLVM and RHV 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.

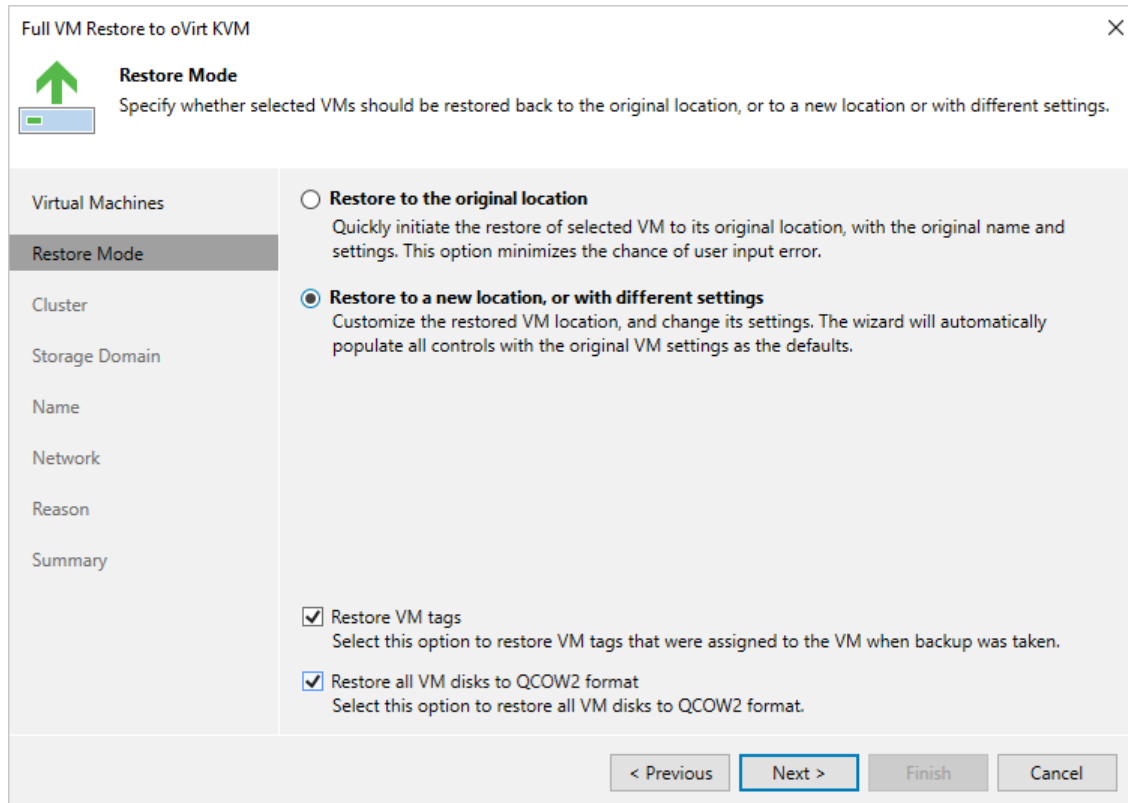


Step 3. Choose Restore Mode

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. You can also choose whether you want the recovered VM to have the same tags as the original VM.

TIP

You can instruct Veeam Backup for OLVM and RHV to restore disks attached to the recovered VM in the QCOW2 format. This will [increase speed and efficiency of incremental backups](#) further created for the VM.



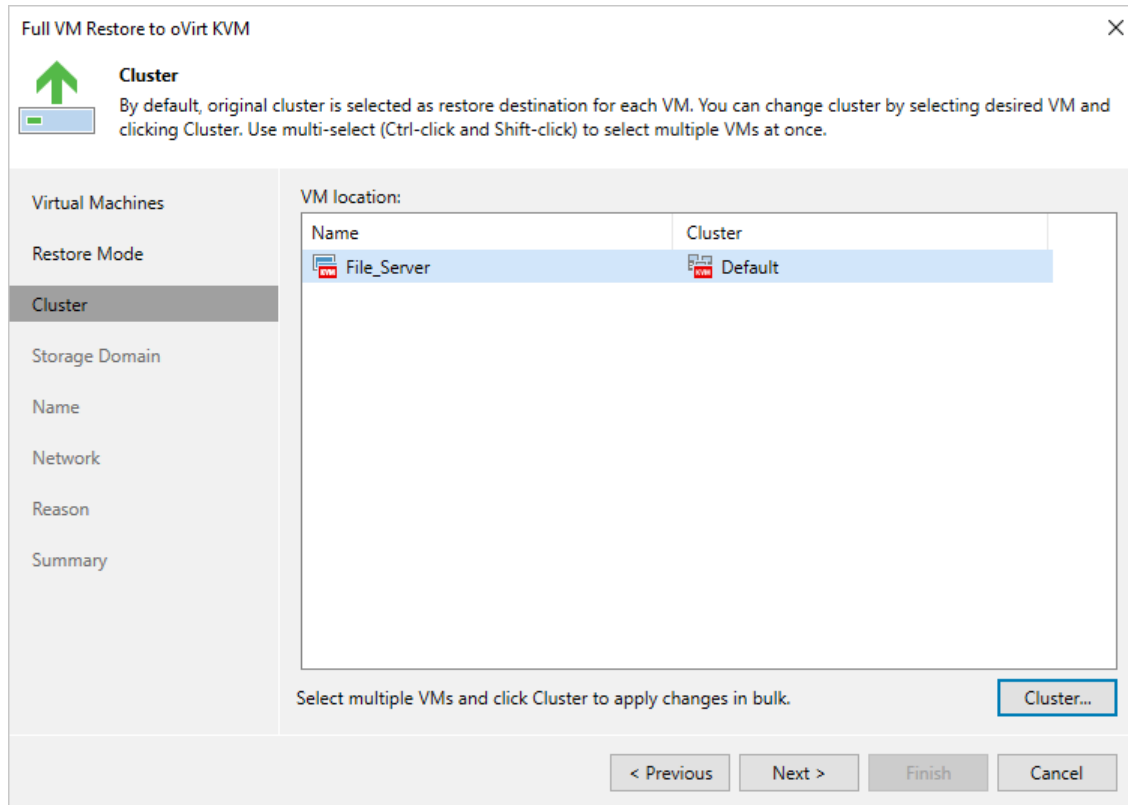
The screenshot shows a wizard window titled "Full VM Restore to oVirt KVM" with a close button (X) in the top right corner. The main heading is "Restore Mode" with a green upward arrow icon and a sub-heading: "Specify whether selected VMs should be restored back to the original location, or to a new location or with different settings." On the left is a sidebar with a list of steps: "Virtual Machines", "Restore Mode" (highlighted), "Cluster", "Storage Domain", "Name", "Network", "Reason", and "Summary". The main area contains two radio button options: "Restore to the original location" (described as quickly initiating the restore to the original location) and "Restore to a new location, or with different settings" (described as customizing the location and settings). Below these are two checked checkboxes: "Restore VM tags" and "Restore all VM disks to QCOW2 format". At the bottom are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 4. Specify Target Cluster

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Cluster** step of the wizard, choose the cluster to which the recovered VM will belong.

For a cluster to be displayed in the list of the available clusters, it must be added to the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).

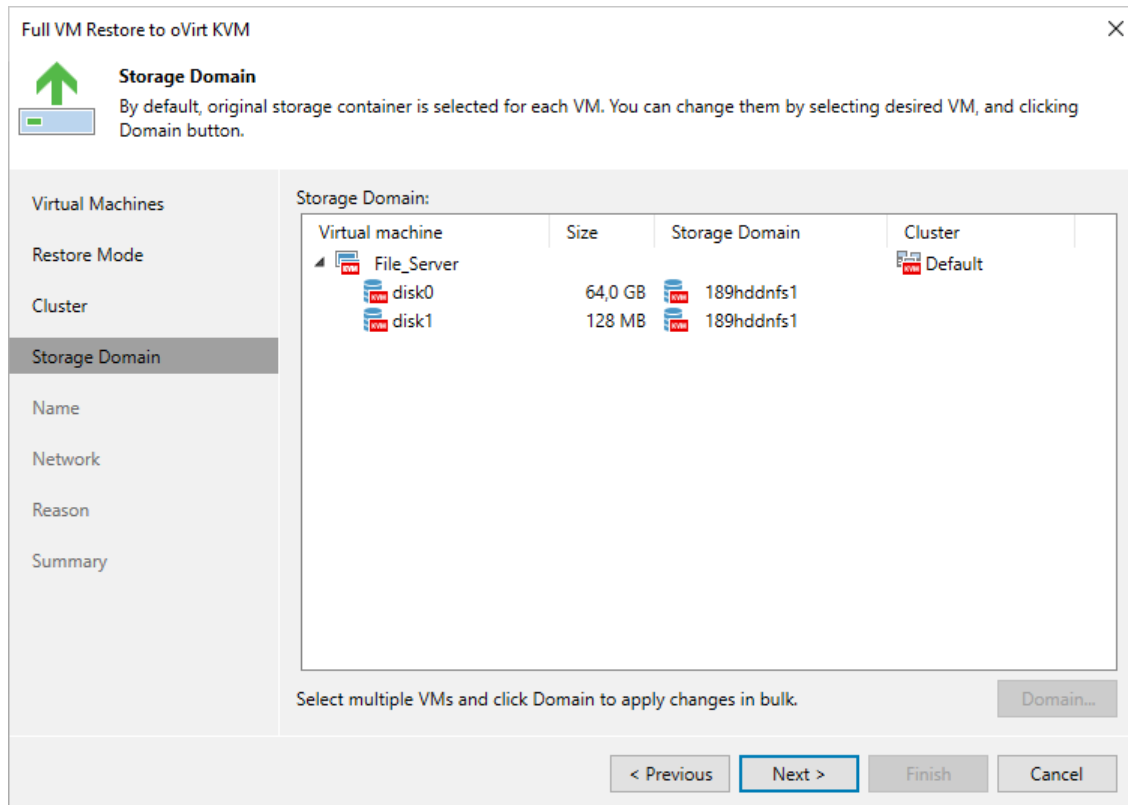


Step 5. Select Storage Domain

[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 Domain** step of the wizard, choose the storage domain where virtual disks of the recovered VM will be stored.

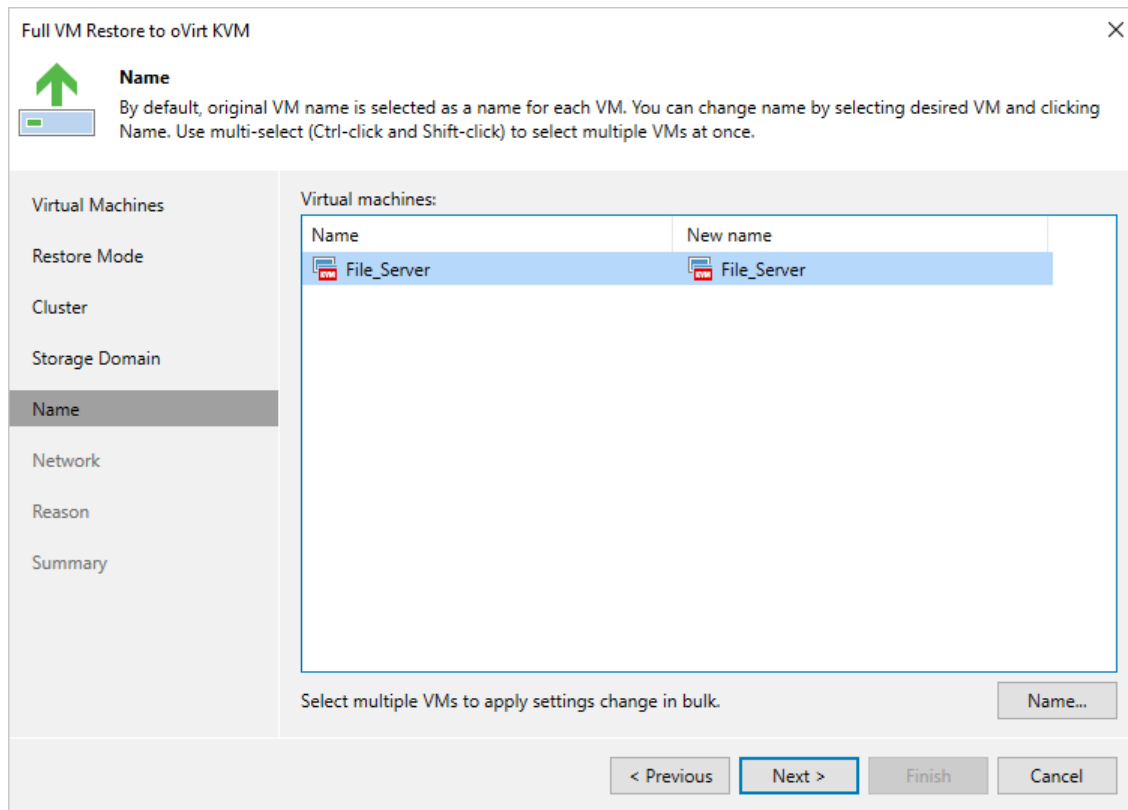
For a domain to be displayed in the list of the available domains, it must be configured in the virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).



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.

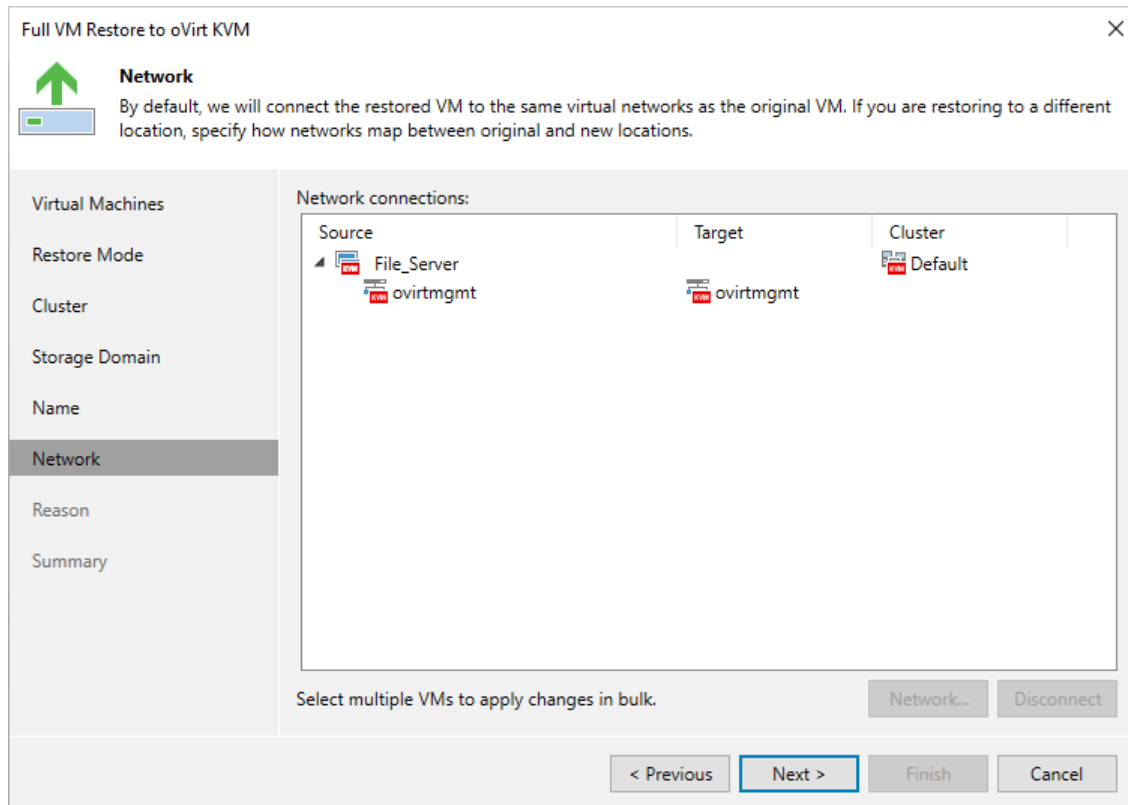


Step 7. Configure Network Settings

[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 virtual environment as described in [Red Hat Virtualization documentation](#) or [Oracle Linux Virtualization Manager documentation](#).



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM. This information will be saved to the session history, and you will be able to reference it later.

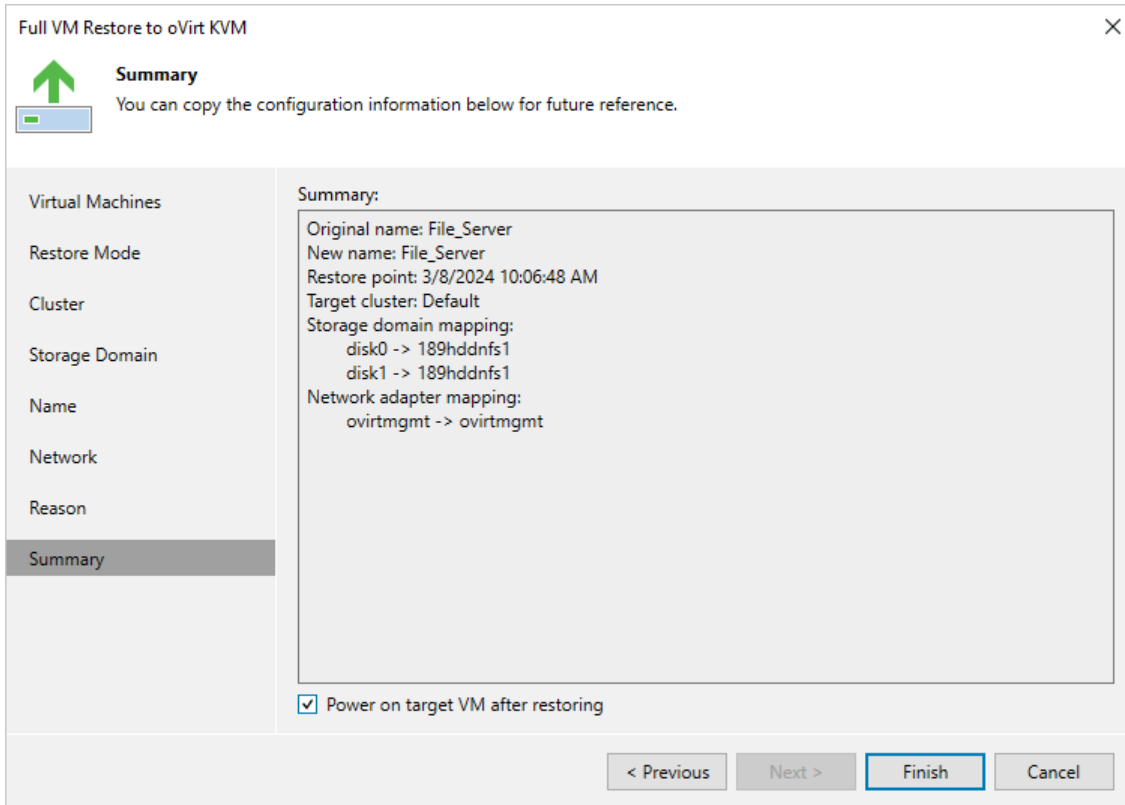
The screenshot shows a wizard window titled "Full VM Restore to oVirt KVM" with a close button (X) in the top right corner. On the left side, there is a vertical navigation pane with the following items: "Virtual Machines", "Restore Mode", "Cluster", "Storage Domain", "Name", "Network", "Reason" (which is highlighted with a dark grey background), and "Summary". Above the navigation pane, there is a green upward-pointing arrow icon and the word "Reason" in bold. Below this, a text box contains the instruction: "Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference." The main area of the wizard is a large text input field labeled "Restore reason:" containing the text "Corrupted disks". Below the text field is a checkbox labeled "Do not show me this page again". At the bottom of the window, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

Step 9. Finish Working with Wizard

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

TIP

If you want to start the recovered VM as soon as the restore process completes, select the **Power on target VM after restoring** check box.



Performing Disk Restore

In case a disaster strikes, you can restore a disk of an oVirt VM from a backup. Veeam Backup for OLVM and RHV allows you to attach the restored disk to the original VM or any other oVirt VM in the oVirt KVM environment.

How Disk Restore Works

During the VM disk restore process, the following steps are performed:

1. The Veeam Backup & Replication console sends the restore session configuration data to the backup appliance.
2. The backup appliance powers off the target VM.
3. The backup appliance launches a worker.
4. The worker creates an empty virtual disk in the oVirt KVM environment.
5. The worker connects to the backup repository and restores backed-up data to the empty disk.
6. [This step applies only if you restore the disk to the original VM and if you choose to replace the existing disk] The worker detaches the original disk from the VM and removes it from the oVirt KVM environment.
7. The worker attaches the created disk with the restored data to the target VM.

How to Perform Disk Restore

To restore a disk 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.](#)

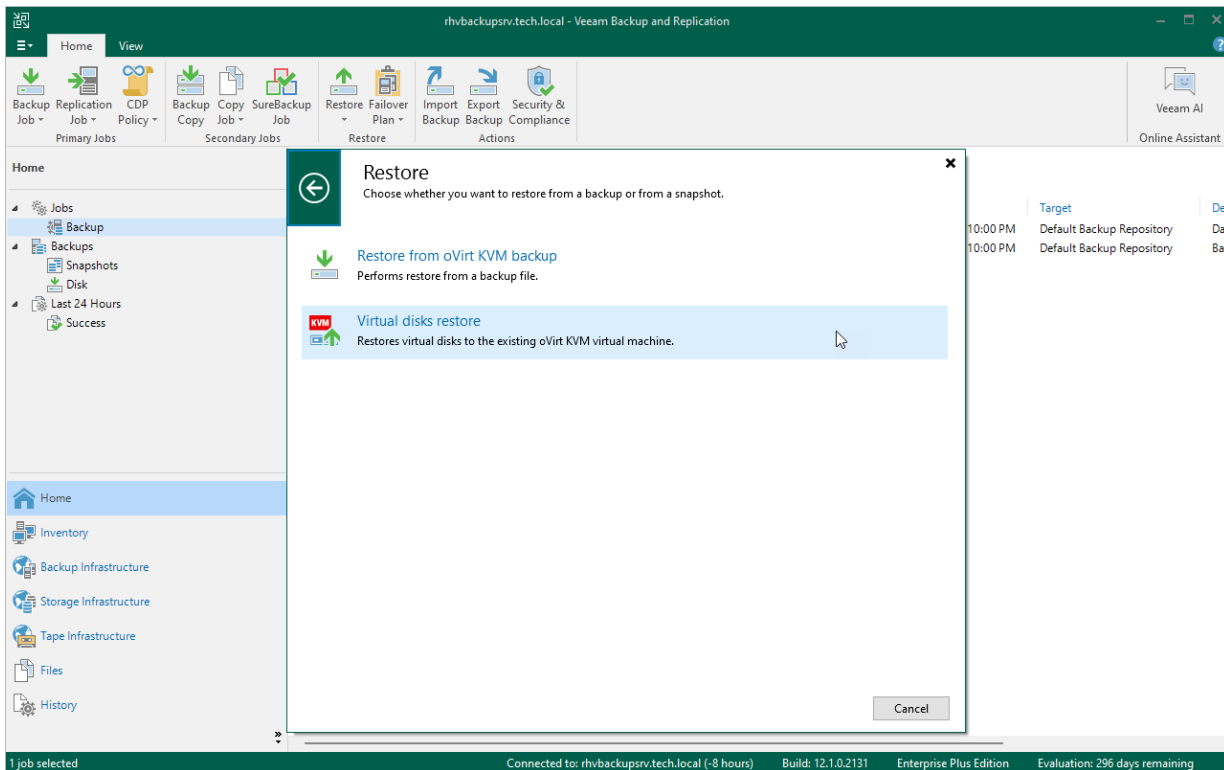
Step 1. Launch Virtual Disk Restore Wizard

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 **Jobs > Backup**.
3. On the ribbon, click **Restore > oVirt KVM**.
4. Click **Entire machine restore**.
5. Click **Restore to oVirt KVM**.
6. Click **Virtual disks restore**.

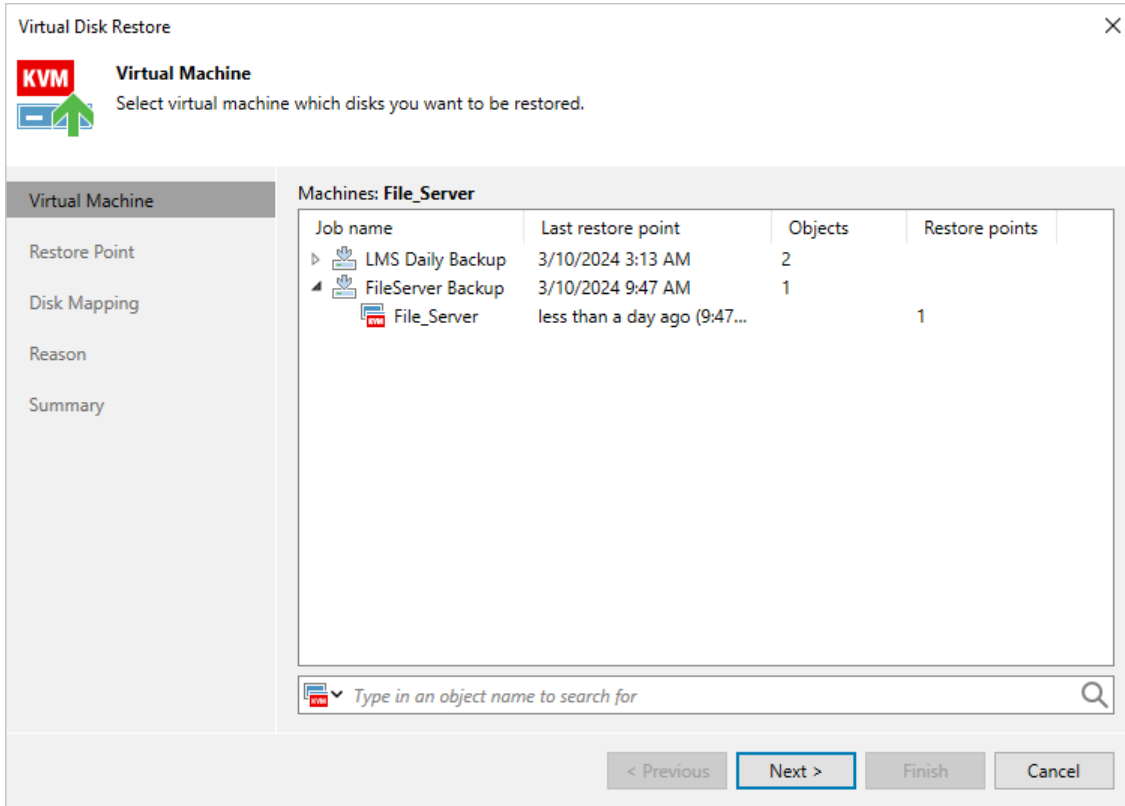
TIP

Alternatively, you can expand the necessary backup in the working area, right-click the VM and select **Restore virtual disks to oVirt KVM**.



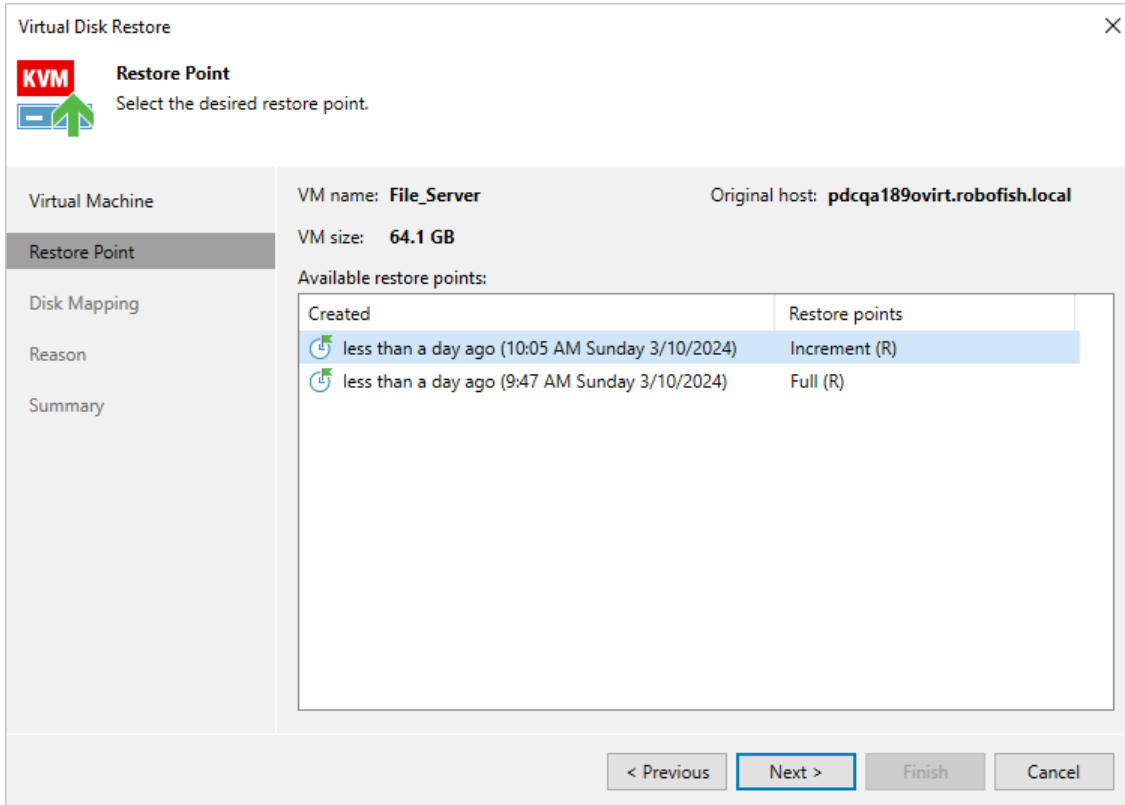
Step 2. Select Virtual Machine

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 for OLVM and RHV 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 for OLVM and RHV attaches the restored disks to the original VM. To attach the disks to another VM, click **Choose**.

IMPORTANT

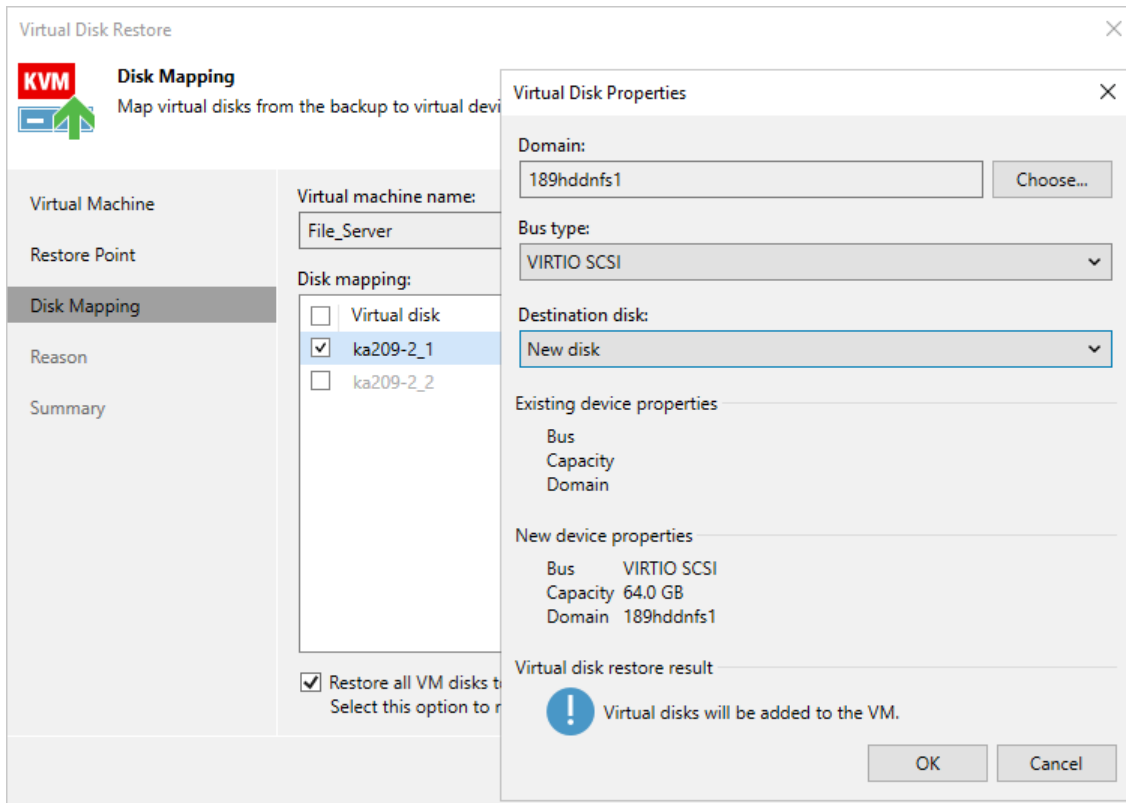
During disk restore, Veeam Backup for OLVM and RHV turns off the target VM to reconfigure its settings and attach the restored disks. It is recommended that you stop all activities on the target VM till the restore session completes.

2. Select virtual disks to restore.

By default, Veeam Backup for OLVM and RHV 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**.

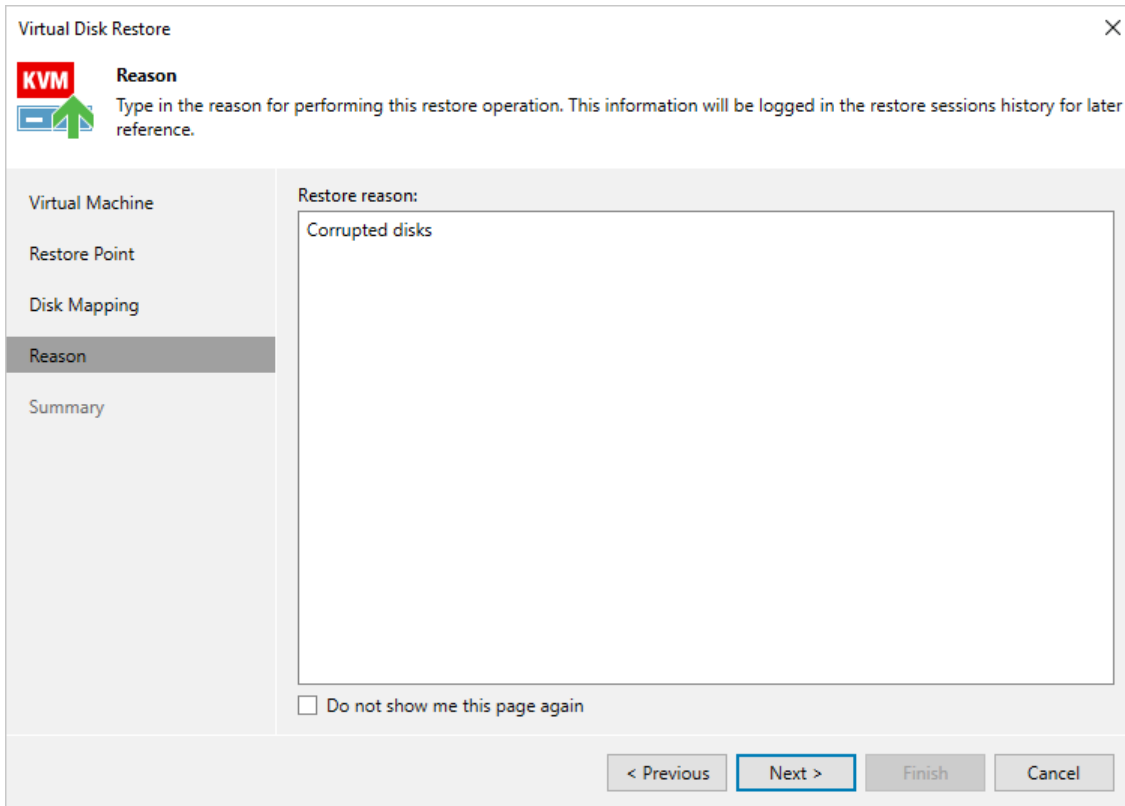
TIP

You can instruct Veeam Backup for OLVM and RHV to restore the disks in the QCOW2 format. This will [increase speed and efficiency of incremental backups](#) further created for the VM.



Step 5. Specify Reason for Restore

At the **Reason** step of the wizard, specify a reason for restoring the disks. This information will be saved to the session history, and you will be able to reference it later.

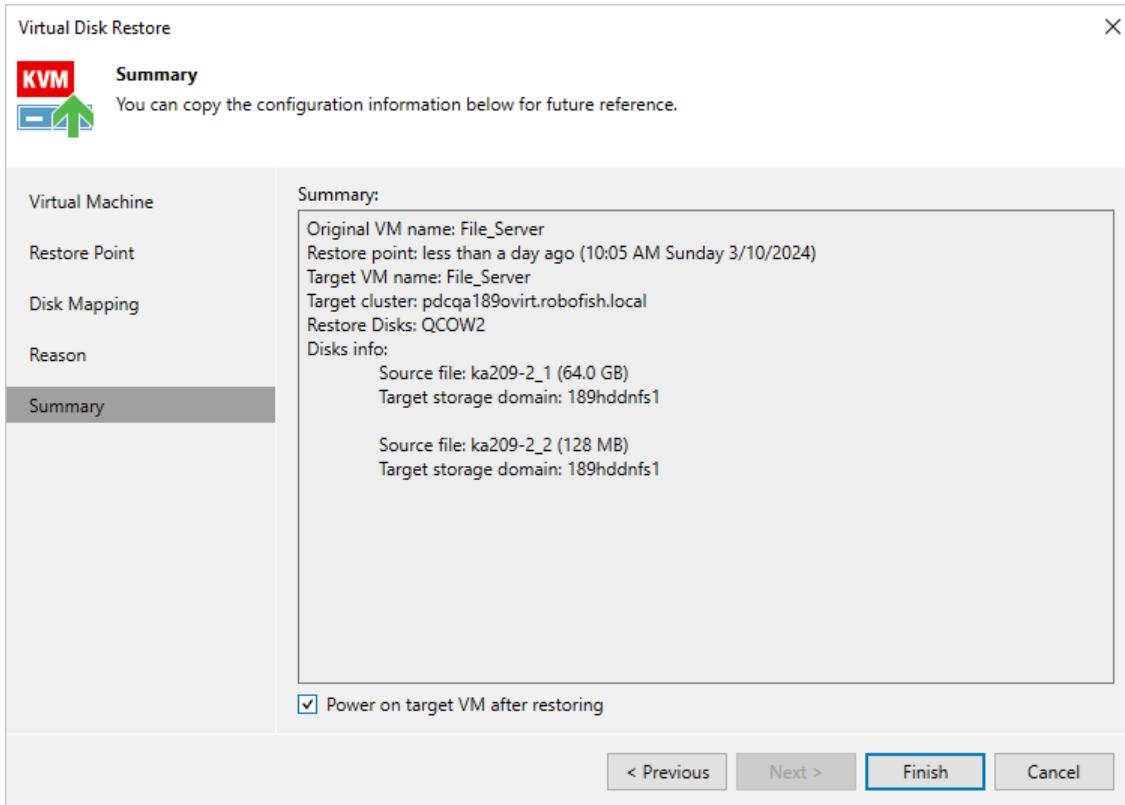


Step 6. Finish Working with Wizard

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 restoring** check box.

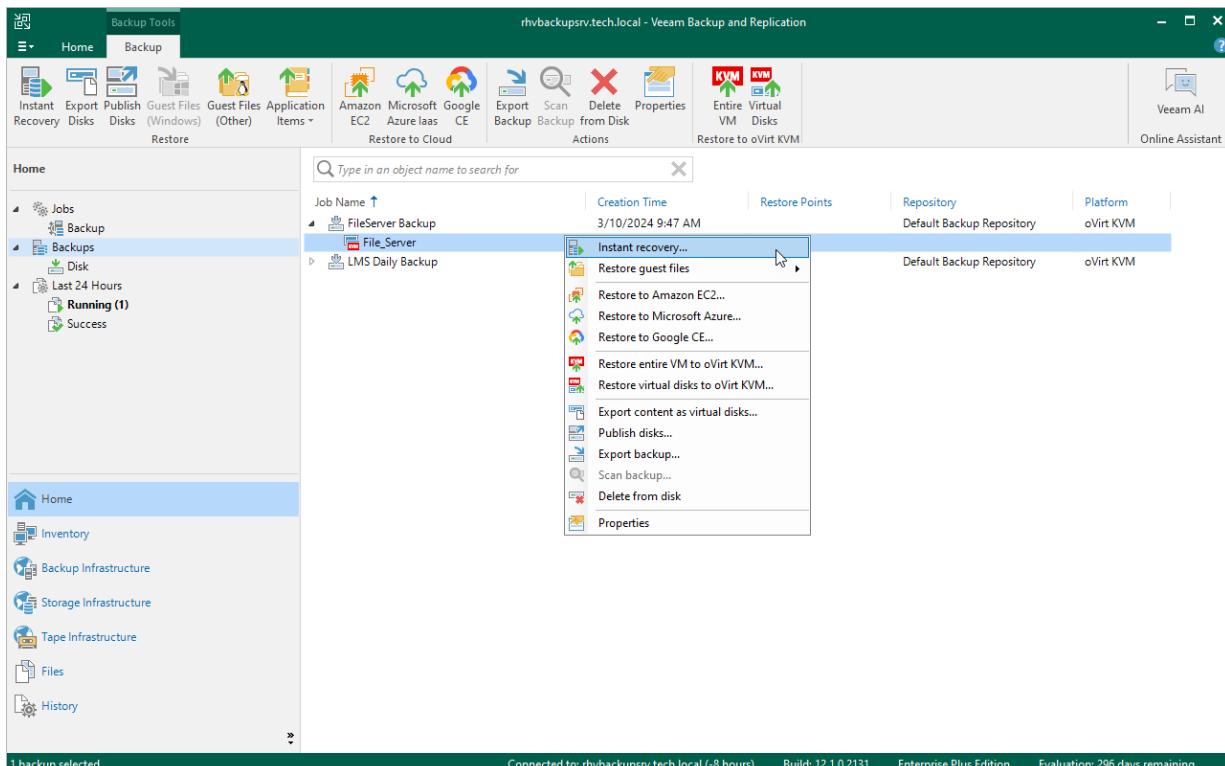


Performing Instant VM Recovery

With Instant VM Recovery, you can immediately restore oVirt VMs as VMware vSphere, Microsoft Hyper-V or Nutanix AHV VMs to your production environment by running them directly from their backups. Instant VM Recovery helps you improve recovery time objectives and minimize disruption and downtime of production workloads. For more information on Instant VM Recovery, see the Veeam Backup & Replication User Guide, section [VM Recovery](#).

To perform Instant VM Recovery, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, right-click the VM you want to restore, and select **Instant Recovery**.
 - To restore the VM to VMware vSphere, complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for VMware vSphere, section [Performing Instant VM Recovery of Workloads to VMware vSphere VMs](#).
 - To restore the VM to Microsoft Hyper-V, complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide for Microsoft Hyper-V, section [Performing Instant VM Recovery of Workloads to Hyper-V VMs](#).
 - To restore the VM to Nutanix AHV, complete the **Instant Recovery** wizard as described in the Veeam Backup for Nutanix AHV User Guide, section [Performing Instant VM Recovery of Workloads to Nutanix AHV](#).

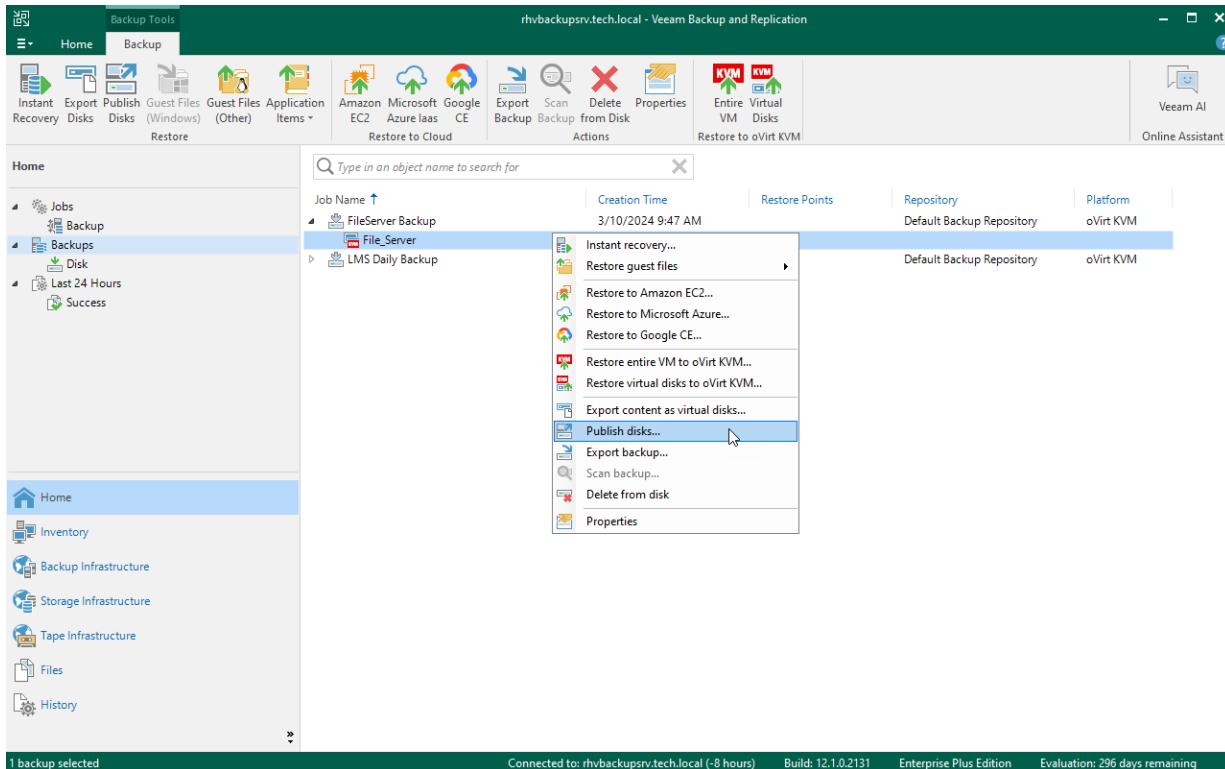


Publishing Disks

Veeam Backup & Replication allows you to mount specific disks of backed-up oVirt 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 an oVirt VM, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to mount and select **Publish disks**.
4. Complete the **Publish Disk** wizard as described in the Veeam Backup & Replication User Guide, section [Publishing Disks](#).

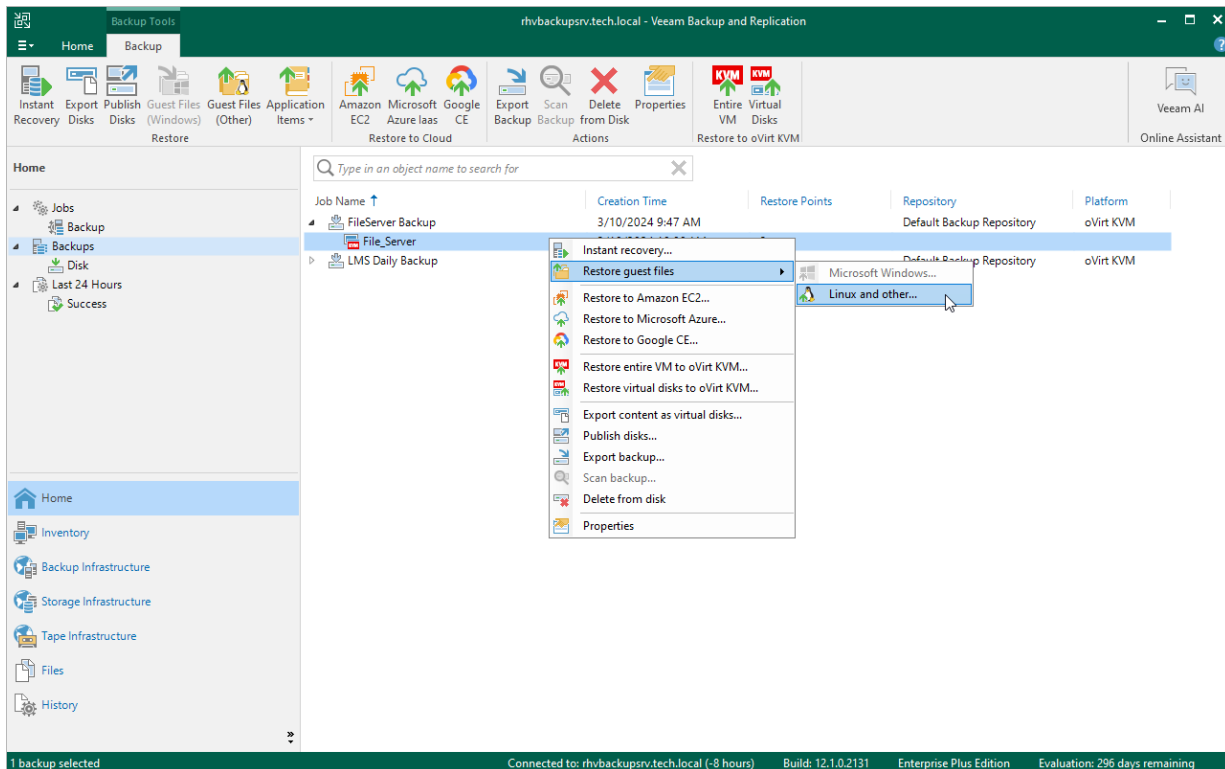


Performing File-Level Restore

With guest OS file recovery (file-level restore), you can restore individual guest OS files and folders from oVirt VM backups created with Veeam Backup for OLVM and RHV. 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. For more information on VM guest OS file restore, see the Veeam Backup & Replication User Guide, section [Guest OS File Recovery](#).

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 do the following:
 - If you want to restore files of a Microsoft Windows machine, select **Restore guest files > Microsoft Windows** and complete the **Guest File Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(FAT, NTFS or ReFS\)](#).
 - If you want to restore files of a Linux, Solaris, BSD, Novell Storage Services, Unix or Mac machine, select **Restore guest files > Linux and other** and complete the **Guest File Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(Multi-OS\)](#).



TIP

Alternatively, you can use Veeam Backup Enterprise Manager to restore guest OS files and folders as described in the Veeam Backup Enterprise Manager Guide, section [Restoring VM Guest OS Files](#).

Performing Application Item Restore

With application item restore, you can use Veeam Backup for OLVM and RHV backups to restore the following data:

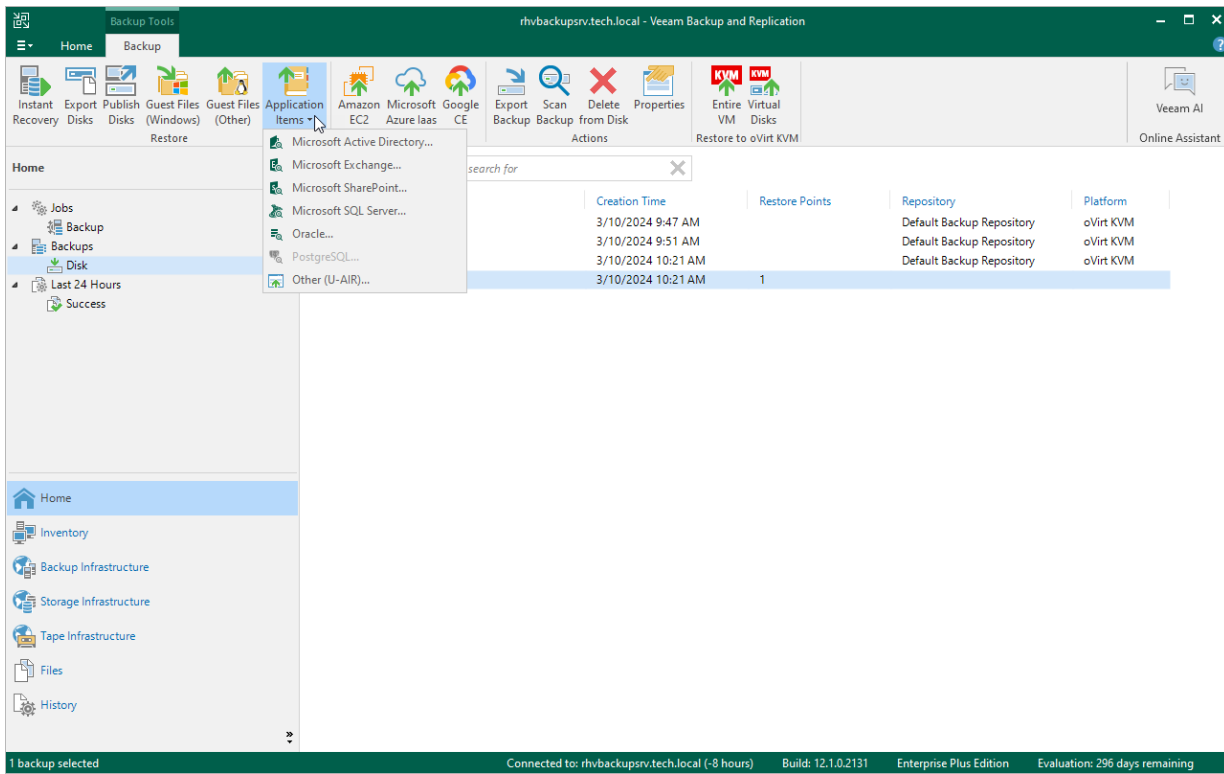
- Microsoft Active Directory objects and containers
- Microsoft Exchange mailboxes, folders and messages
- Microsoft SharePoint sites and lists
- Microsoft SQL Server
- Oracle databases
- PostgreSQL instances and databases residing on Linux VMs

To restore application items from a Veeam Backup for OLVM and RHV VM 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, expand the necessary backup job, select the VM that contains an application you want to restore.
4. Click **Application Items** on the ribbon and then select the application.
5. In the restore wizard, select a restore point that will be used to restore the application, specify a restore reason and click **Browse**.
6. In the Veeam Explorer application, perform the steps described in the [Veeam Explorers User Guide](#).

TIP

As an alternative to application item restore, you can also [perform file-level restore](#) to recover standalone databases using Veeam Explorers.

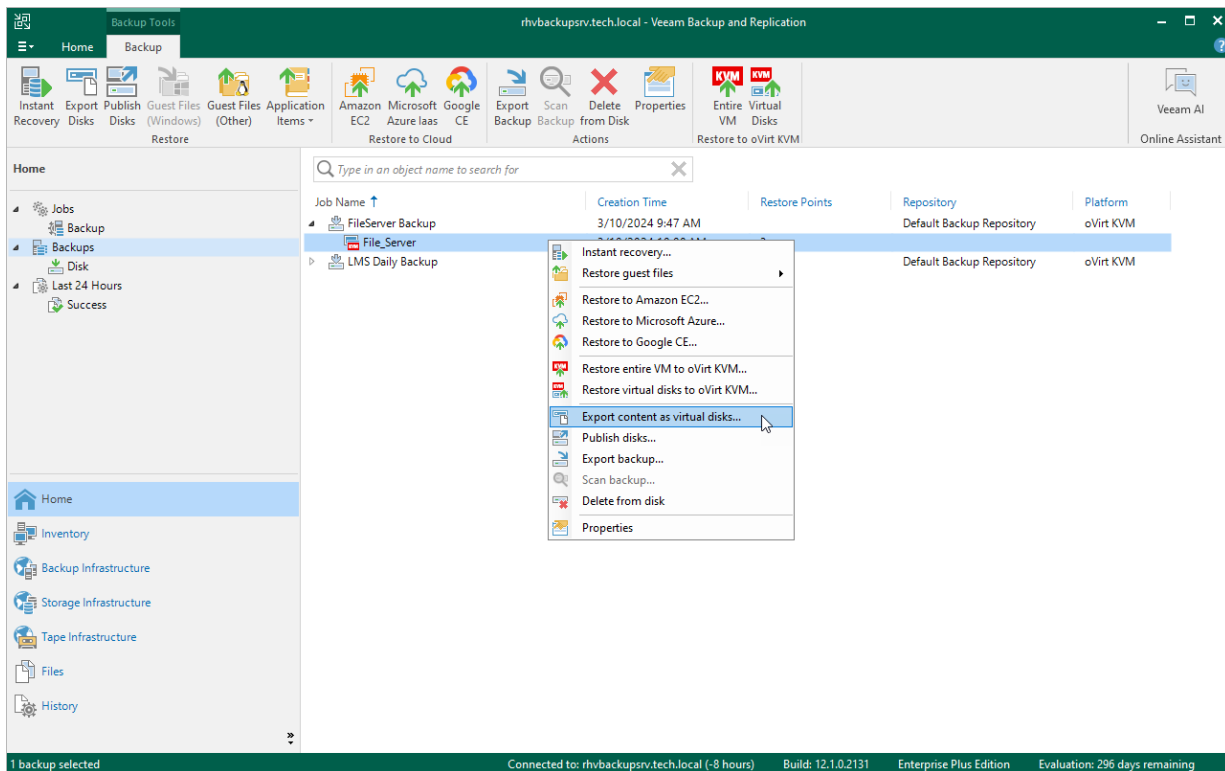


Exporting Disks

Veeam Backup for OLVM and RHV allows you to export disks, that is, restore disks from oVirt 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 an oVirt VM, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that contains disks you want to export and select **Export content as virtual disks**.
4. Complete the **Export Disk** wizard as described in the Veeam Backup & Replication User Guide, section [Exporting Disks](#).

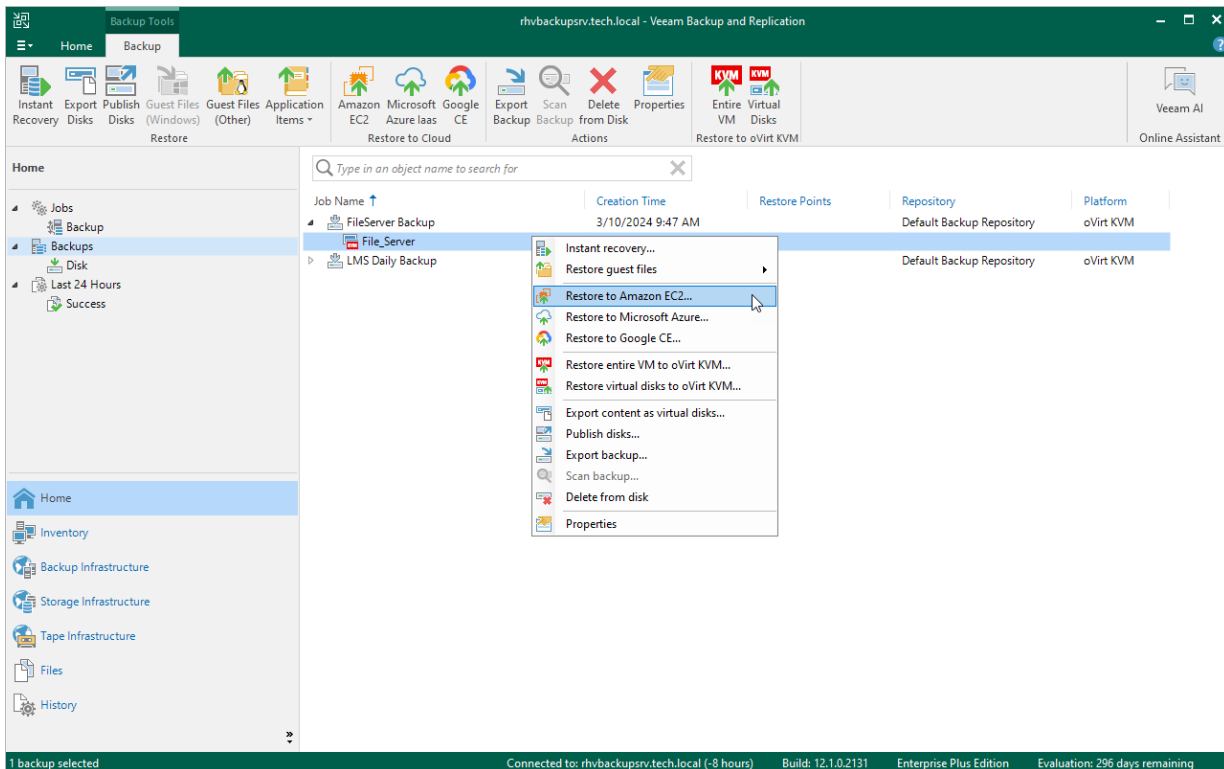


Performing VM Restore to Amazon Web Services

Veeam Backup for OLVM and RHV allows you to restore oVirt VMs to Amazon Web Services (AWS) as EC2 instances. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Amazon EC2](#).

To restore a VM to Amazon EC2, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Amazon EC2**.
4. Complete the **Restore to Amazon EC2** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Amazon EC2](#).

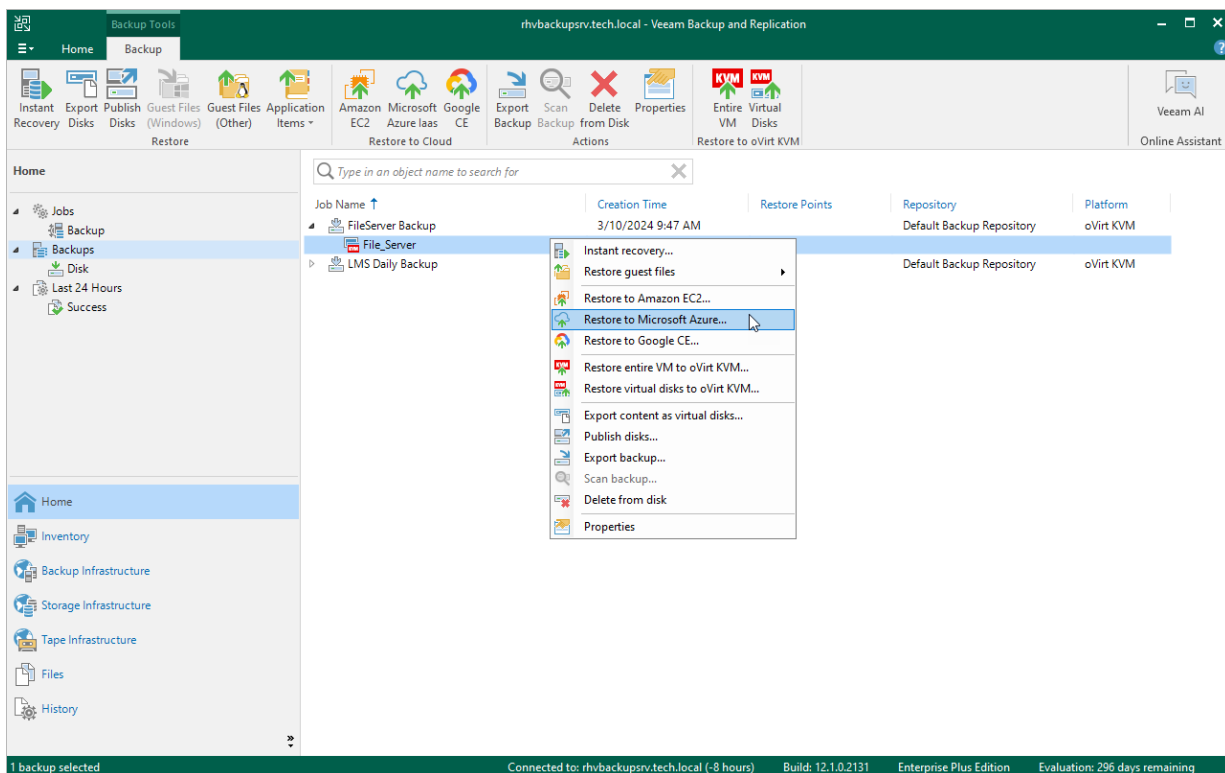


Performing VM Restore to Microsoft Azure

Veeam Backup for OLVM and RHV allows you to restore oVirt VMs to Microsoft Azure as Azure VMs. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Microsoft Azure](#).

To restore a VM to Microsoft Azure, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Microsoft Azure**.
4. Complete the **Restore to Microsoft Azure** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Microsoft Azure](#).

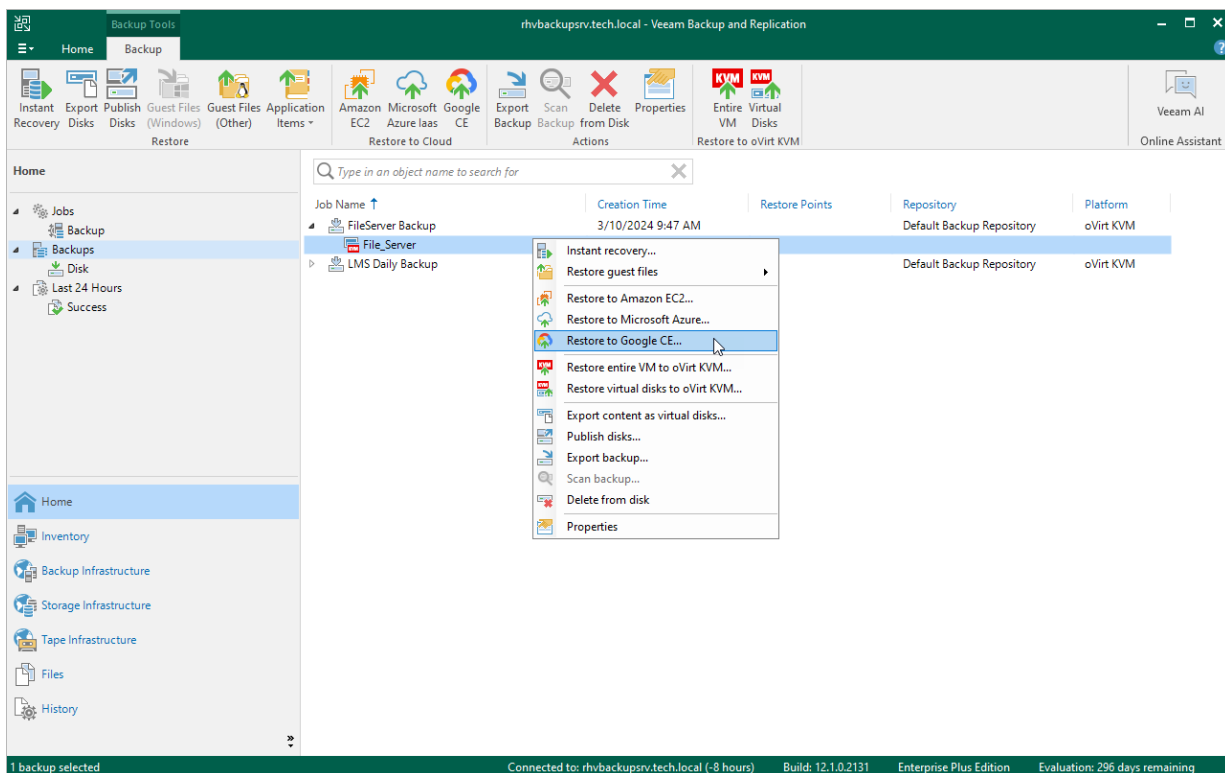


Performing VM Restore to Google Cloud

Veeam Backup for OLVM and RHV allows you to restore oVirt VMs to Google Cloud as VM instances. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Google Compute Engine](#).

To restore a VM to Google Cloud, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. In the inventory pane, select **Backups**.
3. In the working area, expand the necessary backup job, right-click the VM that you want to restore and select **Restore to Google CE**.
4. Complete the **Restore to Google Compute Engine** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Google Compute Engine](#).



Updating Backup Appliance

Veeam Backup for OLVM and RHV allows you to check for new product versions and available package updates, download and install them right from the Veeam Backup & Replication console. It is recommended that you timely install available updates to avoid issues while working with the product. For example, timely installed security updates may help you prevent potential security issues and reduce the risk of compromising sensitive data.

IMPORTANT

Before you install a product update, make sure all backup jobs are stopped and restore tasks are finished. Otherwise, the update process will interrupt the running activities, which may result in data loss.

To download and install available product and package updates, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the backup appliance and click **Missing Updates** on the ribbon, or right-click the backup appliance and select **Missing Updates**.
4. In the **Missing Updates** window, select check boxes next to the necessary updates.

You can view detailed information on an update in the **Description** field.

5. Click **Install**.

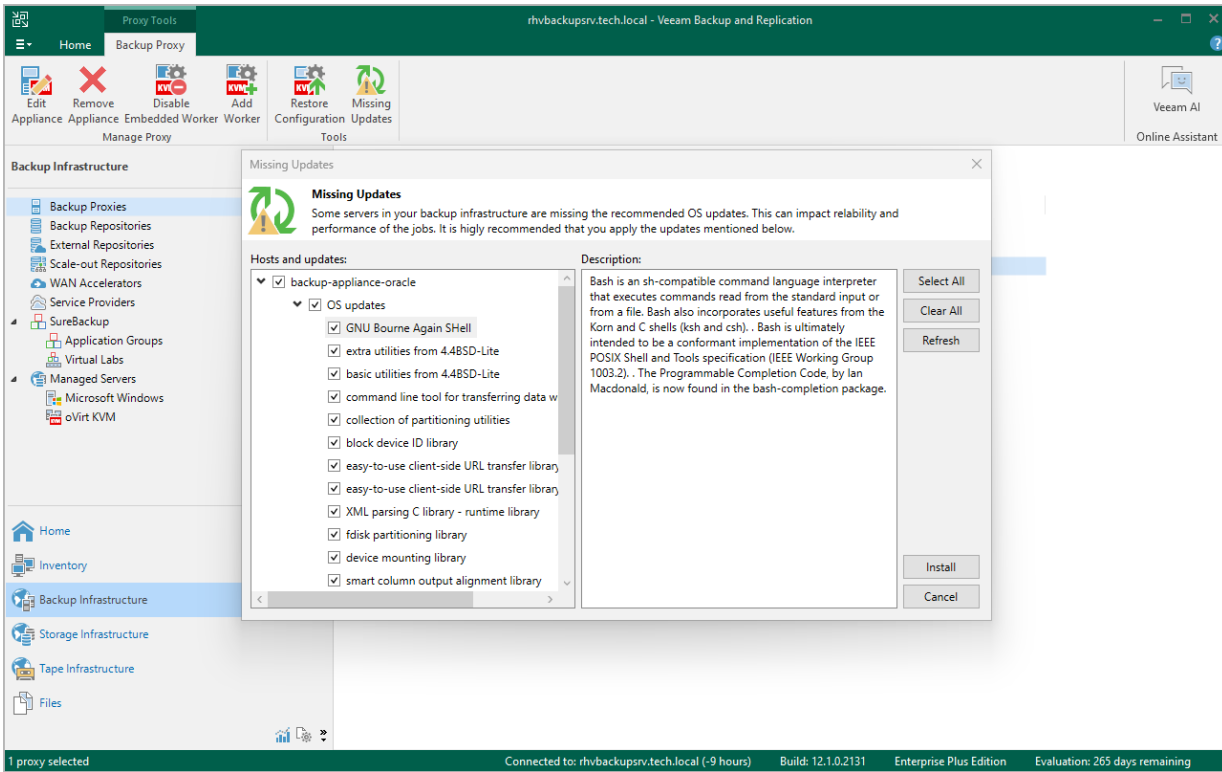
The updater may require you to read and accept the Veeam license agreement and the 3rd party components license agreement. If you reject the agreements, you will not be able to continue installation.

6. Select the **Allow proxy appliance to perform automatic reboot if required** check box and click **Install Updates**.

Veeam Backup for OLVM and RHV will download and install the updates from the internet; it may take several minutes for the installation process to complete.

TIP

If the backup appliance is not connected to the internet, you can instruct Veeam Backup for OLVM and RHV to use an HTTP proxy.



Getting Technical Support

If you have any questions or issues with Veeam Backup for OLVM and RHV, 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

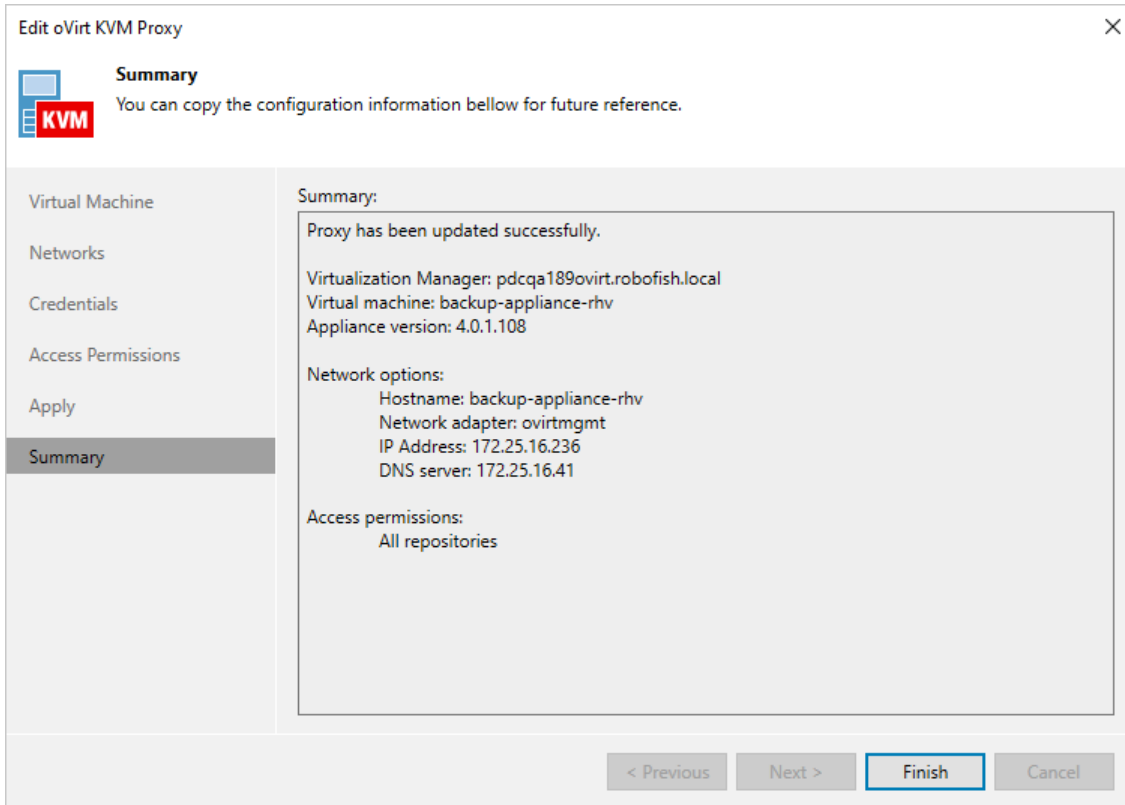
To view the product details, do the following:

1. Open the **Backup Infrastructure** view.
2. In the inventory pane, select **Backup Proxies**.
3. In the working area, select the backup appliance and click **Edit Proxy** on the ribbon, or right-click the backup appliance and select **Properties**.
4. In the **Edit oVirt KVM Proxy** wizard, click **Finish**.
5. Wait for Veeam Backup for OLVM and RHV to complete the backup appliance configuration check and click **Next**.

At the **Summary** step of the wizard, the following information will be displayed:

- Virtualization manager hostname or IP address
- Name of the VM running as the backup appliance
- Currently installed version of Veeam Backup for OLVM and RHV
- Backup appliance network settings

- Repositories to which the backup appliance has access

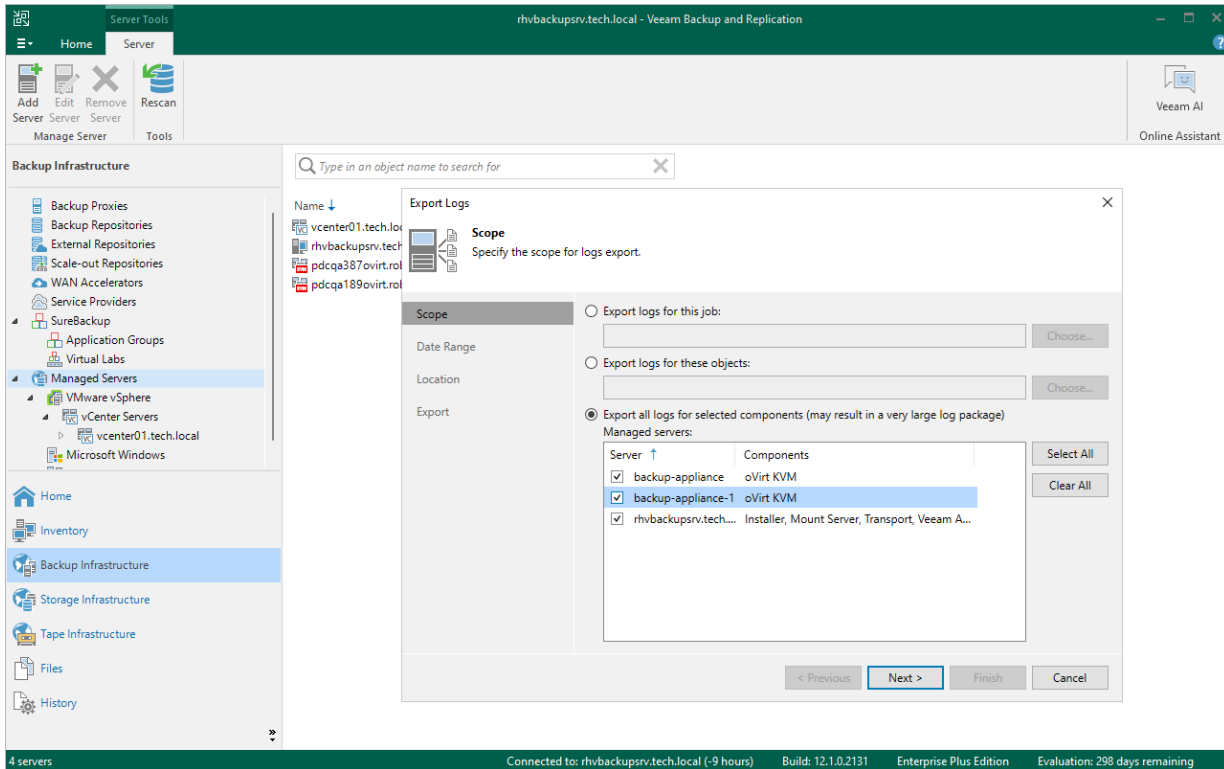


Downloading Logs

To download the product logs, do the following:

1. From the main menu of the Veeam Backup & Replication console, select **Help > Support Information**.

- At the **Scope** step of the **Export Logs** wizard, select the **Export all logs for selected components** option. Then, in the **Managed servers** list, select the backup server and the VM running as the backup appliance. Complete the wizard as described in the Veeam Backup & Replication User Guide, section [Exporting Logs](#).



Appendix. Deprecated Functionality

Starting from version 4.0, Veeam Backup for OLVM and RHV comes without the web console that was previously used to manage the backup appliance. The functionality of the web console is now integrated into the Veeam Backup & Replication console, which allows you to perform the following tasks:

- [Create backup jobs](#)
- [Perform VM restore](#)
- [Perform disk restore](#)
- [Enable SSH on the backup appliance VM](#)
- [Edit backup appliance network settings](#)
- [Edit the Administrator account](#)
- [Configure settings for CPU and RAM usage notifications](#)
- [Back up and restore appliance configuration](#)
- [Update the backup appliance](#)
- [Configure notification settings for automated delivery of backup job results](#)

NOTE

Veeam Backup for OLVM and RHV supports mail servers with SMTP basic authentication only.

You can also use the Veeam Backup & Replication console to track real-time statistics of all running and completed operations and to generate reports with statistics data. For more information, see the Veeam Backup & Replication User Guide, section [Reporting](#).