



Veeam Backup & Replication

Version 12

Quick Start Guide for VMware vSphere

February, 2023

© 2023 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	5
GETTING STARTED	6
ABOUT VEEAM BACKUP & REPLICATION	7
VEEAM BACKUP & REPLICATION UI	8
BACKUP INFRASTRUCTURE	9
Backup Infrastructure Components	10
Deployment Scenarios	11
PLANNING AND PREPARATION	13
System Requirements	14
Used Ports	15
DEPLOYMENT	16
Step 1. Installing Veeam Backup & Replication.....	17
Step 2. Adding Virtual Infrastructure Servers	20
Step 3. Configuring Backup Proxy	23
Step 4. Configuring Backup Repository	28
Step 5. Configuring Object Storage Repositories	34
Step 6. Configuring Scale-Out Backup Repositories	39
VM BACKUP.....	44
Backup Methods	45
Creating Backup Job	47
Monitoring Job Performance in Real Time	52
Start Backup Job Manually	53
Locating Backup Files	54
Creating Application-Aware Backup Job	55
DATA RECOVERY.....	59
Restoring Entire VM.....	60
Restoring VM Files	64
Restoring VM Virtual Disks	67
Restoring Guest OS Files	70
Restoring VM Guest OS Files (FAT, NTFS, ReFS)	71
Restoring VM Guest OS Files (Linux, Unix, etc)	73
Restoring Application Items	76
BACKUP COPY.....	79
VM REPLICATION	83
Creating Replication Job	84
Monitoring Job Performance in Real Time	91

Start Replication Job Manually	92
Replica Failover and Failback	93
Performing Replica Failover	95
Performing Permanent Failover	97
Undoing Failover	98
Performing Failback	99
Committing Failback	101
Undoing Failback	102
ENTERPRISE MANAGER.....	103
Installing Veeam Backup Enterprise Manager	104
Adding Backup Servers.....	107
Managing Jobs	109
Performing 1-Click File Restore	111
Performing Self-Restore of VM Guest OS Files	113
BACKING UP PHYSICAL MACHINES.....	117
Creating Protection Group	118
Creating Veeam Agent Backup Job	122
RESTORING DATA OF PHYSICAL MACHINES	127

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

Getting Started

Document Structure

The guide contains instructions on the following:

- Veeam Backup & Replication functionality: how to deploy Veeam Backup & Replication, perform backup, replication, and restore operations.
- Integration with Veeam Backup Enterprise Manager: a free tool for managing distributed infrastructure.
- Agent management: built-in Veeam Backup & Replication feature to back up physical machines using Veeam Agents.

Help and Support

This guide provides a high-level overview of Veeam Backup & Replication primary features and should be regarded as a supplement to existing technical documentation. The complete set of documentation can be found on the [Veeam Technical Documentation page](#).

For technical support and assistance, use the following resources:

- [Veeam R&D Forums](#)
- [Customer Support Portal](#)

About Veeam Backup & Replication

What is Veeam Backup & Replication?

Veeam Backup & Replication is a data protection and disaster recovery solution for virtual, physical and cloud environments. With Veeam Backup & Replication you can:

- Create crash-consistent and application-consistent backups of virtual and physical machines.
- Quickly restore physical machines, EC2 instances, VMs, VM disks, guest OS files and application items.
- Perform backup health check to verify that backups are not corrupted and are ready for restoring.
- Create VM replicas and switch to them in case of a disaster.
- Automate transferring of backups to tapes and other external repositories.

Note that in this guide, we will not overview all the Veeam Backup & Replication capabilities. You can find them in the [Veeam Backup & Replication User Guide](#).

What Else Can I Do?

Veeam Backup & Replication provides utilities not mentioned in this guide that can help you secure and manage your data:

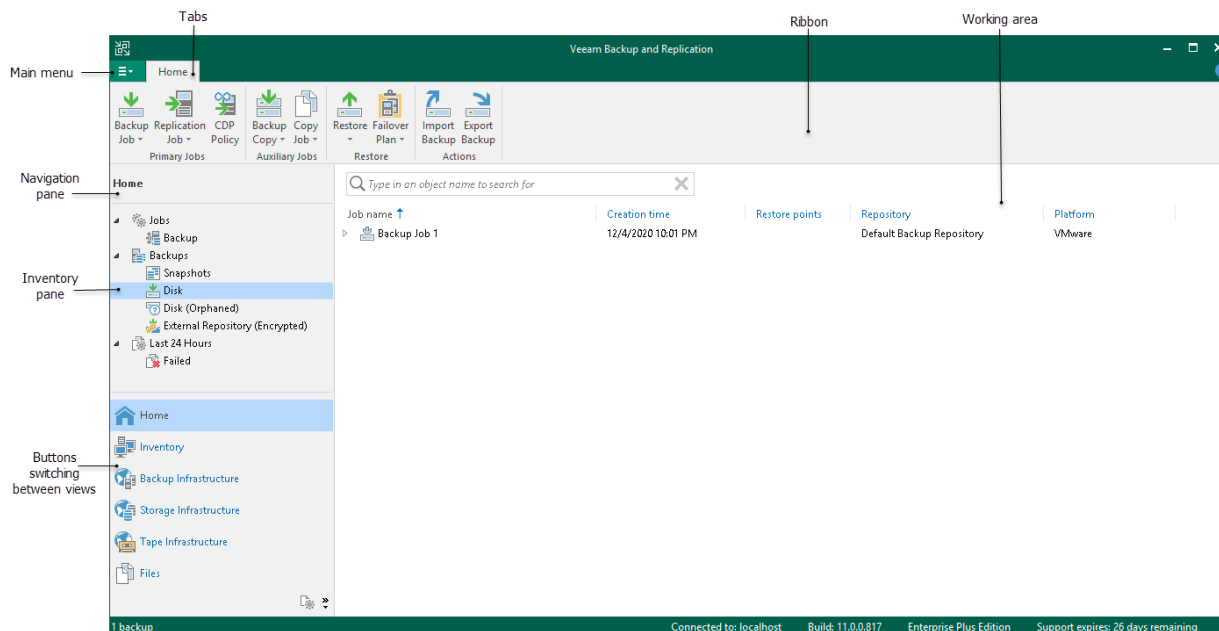
- [Veeam Backup for Microsoft Office 365](#): backup and restore solution for Office 365 data.
- [Veeam Backup for Nutanix AHV](#): backup and restore solution for Nutanix AHV VMs.
- [Plug-in for Veeam Backup for AWS](#): extension for integration with Veeam Backup for AWS.
- [Plug-in for Veeam Backup for Microsoft Azure](#): extension for integration with Veeam Backup for Microsoft Azure.
- [Plug-in for Veeam Backup for Google Cloud](#): extension for integration with Veeam Backup for Google Cloud.

Veeam Backup & Replication also provides the following tools for monitoring and management:

- [Veeam ONE](#): real-time monitoring, reporting, alerting and managing tool for virtual and physical environments.
- [Veeam Management Pack for Microsoft System Center](#): Microsoft System Center extension for managing and monitoring VMware vSphere, Microsoft Hyper-V, and Veeam Backup & Replication.
- [Veeam Availability Orchestrator](#): a tool for automated creation and testing of DR plans that comply regulatory requirements.

Veeam Backup & Replication UI

The user interface of Veeam Backup & Replication is designed to let you quickly find commands that you need and perform data protection and disaster recovery tasks.



TIP:

To open online help, press [F1] in any Veeam Backup & Replication wizard or window. You will be redirected to the corresponding section of the [Veeam Backup & Replication User Guide](#).

Reference

For details, see [Veeam Backup & Replication UI](#) section in the Veeam Backup & Replication User Guide.

Backup Infrastructure

This section describes main Veeam Backup & Replication infrastructure components and deployment scenarios.

Backup Infrastructure Components

To start working with Veeam Backup & Replication, you must set up the backup infrastructure. The basic Veeam Backup & Replication infrastructure consists of the following core components:

- **Backup server**

A Microsoft Windows-based machine on which Veeam Backup & Replication is installed. The backup server performs main management operations: coordinates backup, replication and restore tasks, controls job scheduling and resource allocation.

- **Backup repository**

A server where Veeam Backup & Replication keeps backup files, backup copies and metadata of replicated VMs.

- **Backup proxy**

A component that retrieves data from the source host, processes it and transfers to the backup repository.

- **Infrastructure servers and hosts**

VMware vSphere servers that you plan to use as source and target for backup, replication and other activities. Microsoft Windows and Linux servers for which you plan to assign roles of a backup proxy or backup repository.

Reference

For details on all backup infrastructure components, see the [Backup Infrastructure Components](#) section in the Veeam Backup & Replication User Guide.

Deployment Scenarios

You can use Veeam Backup & Replication in virtual environments of any size and complexity. The architecture of the solution supports onsite and offsite data protection, operations across remote sites and geographically dispersed locations. Veeam Backup & Replication provides flexible scalability and easily adapts to the needs of your virtual environment.

Veeam Backup & Replication supports several deployment scenarios, and each one includes the core infrastructure components: backup server, backup proxy and backup repository. Depending on the size of your virtual environment, you can use one of the following scenarios:

- [Simple deployment](#)

For small virtual environments. In this scenario, the roles of all components required for data protection tasks are assigned to one machine.


- [Advanced deployment](#)

For medium-sized and large-scale virtual environments. In this scenario, the roles of components required for data protection tasks are assigned to dedicated machines.

Veeam Backup & Replication also supports the distributed deployment scenario for large geographically dispersed environments with multiple backup servers. We omit detailed description of this scenario because this guide is aimed for quick overview of basic features. For details on the distributed scenario, see the [Distributed Deployment](#) section in the Veeam Backup & Replication User Guide.

Simple Deployment

In the simple deployment scenario, the roles of the backup server, backup proxy and backup repository are assigned to a single machine. These roles are assigned automatically to the machine where you install Veeam Backup & Replication.



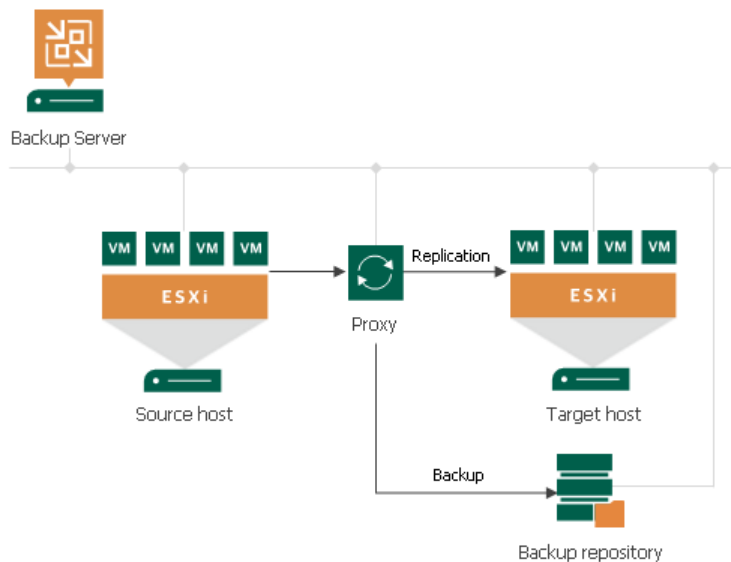
The drawback of the simple deployment scenario is that only the backup server handles and stores all data. For medium-sized or large-scale environments, the capacity of a single backup server may not be enough. To take the load off the backup server and balance it throughout your backup infrastructure, it is recommended to use the [advanced deployment](#) scenario.

Advanced Deployment

In the advanced deployment scenario, the roles of the backup server, backup proxy and backup repository are assigned to different machines. This gives the following advantages:

- The processing load is moved from backup server to backup proxy.
- Increased fault tolerance: you can store data on a separate machine (the backup repository).

Note that this scenario requires that you assign the roles of the proxy and repository manually.



Depending on production environment and backup and replication scenarios you plan to use, the advanced deployment scenario may include multiple backup proxies and backup repositories, both on-site and off-site, controlled by a single backup server.

Planning and Preparation

Before you install Veeam Backup & Replication, you must make sure that the virtual environment and machines that you plan to use as backup infrastructure components meet product hardware recommendations and system requirements.

System Requirements

Make sure that servers that you plan to use as backup infrastructure components meet the system requirements listed in the following sections of the Veeam Backup & Replication User Guide:

- [System Requirements for Backup Server](#)
- [System Requirements for Backup Proxy Server](#)
- [System Requirements for Backup Repository Server](#)
- [System Requirements for Enterprise Manager](#)
- [System Requirements for Supported Applications](#)
- [System Requirements for Veeam Explorer for Microsoft Active Directory](#)
- [System Requirements for Veeam Explorer for Microsoft Exchange](#)
- [System Requirements for Veeam Explorer for Microsoft SharePoint](#)
- [System Requirements for Veeam Explorer for Microsoft SQL Server](#)
- [System Requirements for Veeam Explorer for Oracle](#)

Reference

For the full list of system requirements, see the [System Requirements](#) section in the Veeam Backup & Replication User Guide.

Used Ports

On backup infrastructure components, Veeam Backup & Replication automatically creates firewall rules for the required ports. These rules allow communication between the components.

You can find the lists of the ports in the following sections of the Veeam Backup & Replication User Guide:

- [Backup Server Connections](#)
- [Microsoft Windows Server Connections](#)
- [Linux Server Connections](#)
- [Backup Proxy Connections](#)
- [Backup Repository Connections](#)
- [Mount Server Connections](#)
- [VM Guest OS Connections](#)
- [Veeam Backup Enterprise Manager Connections](#)
- [Veeam Explorer for Microsoft Active Directory Connections](#)
- [Veeam Explorer for Microsoft Exchange Connections](#)
- [Veeam Explorer for Microsoft SharePoint Connections](#)
- [Veeam Explorer for Microsoft SQL Server Connections](#)
- [Veeam Explorer for Oracle Connections](#)
- [Veeam Agent for Microsoft Windows Connections](#)
- [Veeam Agent for Linux Connections](#)

Reference

For the full list of ports, see the [Used Ports](#) section in the Veeam Backup & Replication User Guide.

Deployment

To start using Veeam Backup & Replication, do the following:

1. [Install Veeam Backup & Replication](#)
2. [Add virtual infrastructure servers](#)
3. [Configure backup proxy](#)
4. [Configure backup repository](#)
5. [Configure object storage repository](#) (optional)
6. [Configure scale-out backup repository](#) (optional)

NOTE:

In the simple deployment scenario, Veeam Backup & Replication uses the backup server also as the backup proxy and backup repository. For this reason, you can skip the third and fourth steps.

Step 1. Installing Veeam Backup & Replication

Install Veeam Backup & Replication on a Microsoft Windows-based physical or virtual machine.

Before You Begin

Before you install Veeam Backup & Replication, check the following prerequisites:

- The machine on which you plan to install Veeam Backup & Replication must meet the system requirements for the backup server. For details, see [System Requirements](#).
- A user account that you plan to use for installation must have local Administrator permissions.

Installing Veeam Backup & Replication

To install Veeam Backup & Replication, do the following:

1. Download the latest version of the Veeam Backup & Replication installation image from the [Download Veeam products](#) page.

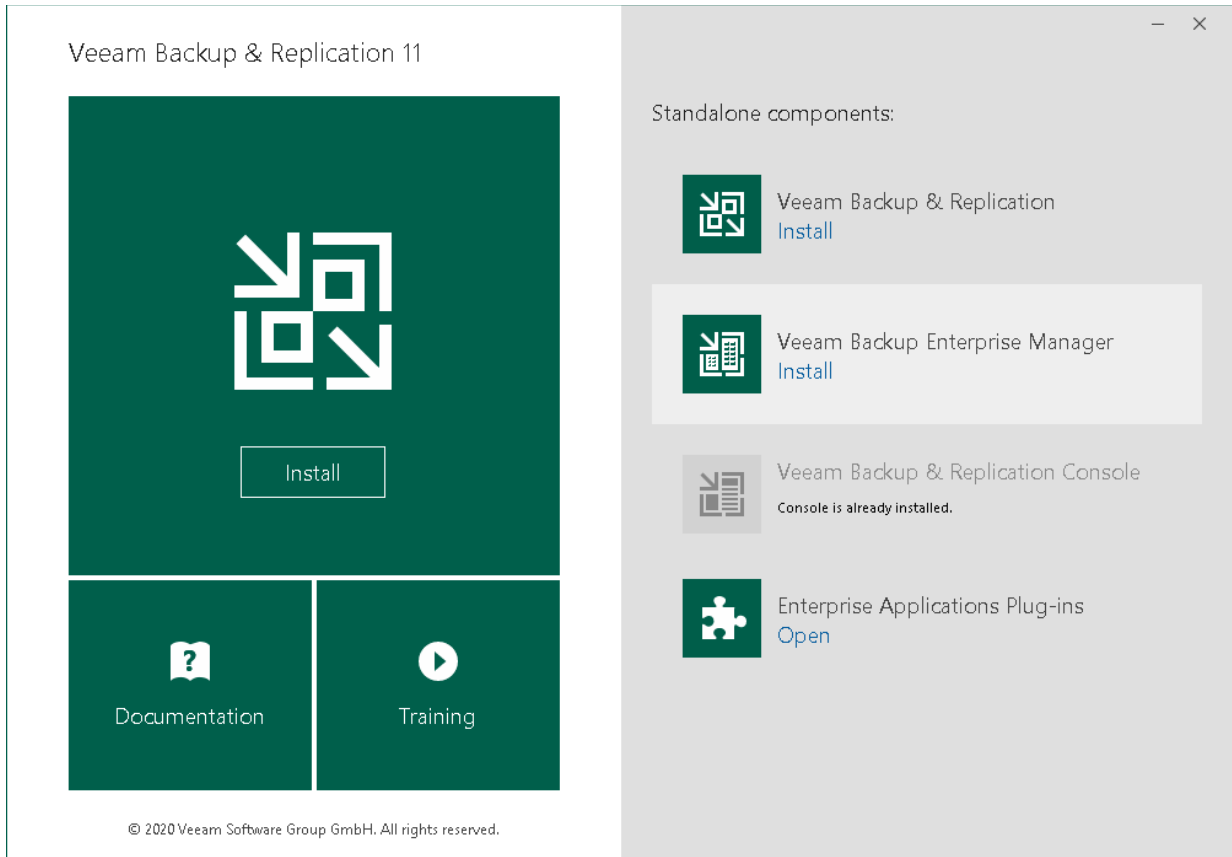
You must sign in with your Veeam account. If you do not have the account, register with your business email address.

2. Mount the installation image to the machine on which you plan to install Veeam Backup & Replication or burn the image file to a flash drive or other removable storage device.

If you plan to install Veeam Backup & Replication on a VM, use built-in tools of the virtualization management software to mount the installation image to the VM.

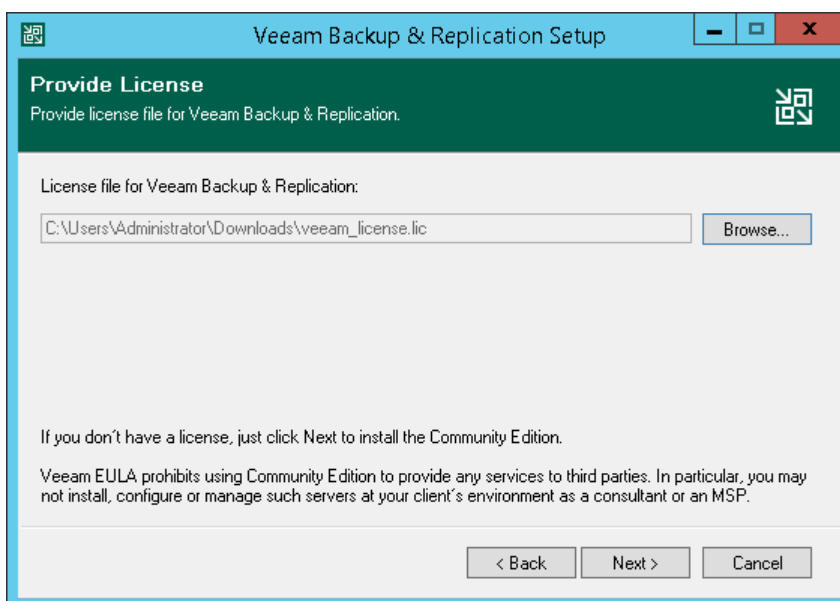
3. Run the `Setup.exe` file from the image or disk to open the splash screen.

4. In the Veeam Backup & Replication section of the splash screen, click **Install**.



5. At the **License Agreement** step of the wizard, read the license agreements and select check boxes to accept the terms.
6. At the **Provide License** step of the wizard, specify the path to the license key.

If you skip this step, Veeam Backup & Replication will operate in the Community edition mode. You can switch to the full version of the product if you install the license. For more information, see [Veeam Backup & Replication Community Edition](#).



7. At the **Program Features** step of the wizard, keep the default settings.

- At the **System Configuration Check** step of the wizard, install missing software components, if any.

NOTE:

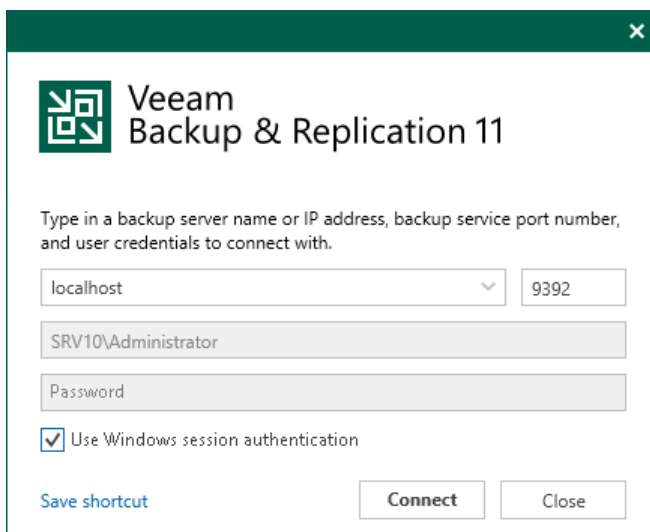
If all required components are already installed on the machine, the **System Configuration Check** step is skipped.

- At the **Default Configuration** step of the wizard, click **Install** to begin the installation.

Starting Veeam Backup & Replication

To start Veeam Backup & Replication, do the following:

- In the Microsoft Windows **Start** menu, select **Apps > Veeam > Veeam Backup & Replication Console**.
- In the authentication window, click **Connect**.



The image shows the Veeam Backup & Replication 11 authentication window. It has a dark green title bar with a close button. The window contains the Veeam logo and the text 'Veeam Backup & Replication 11'. Below this, it says 'Type in a backup server name or IP address, backup service port number, and user credentials to connect with.' There are three input fields: a dropdown menu showing 'localhost', a text box with '9392', and a text box with 'SRV10\Administrator'. Below these is a checkbox labeled 'Use Windows session authentication' which is checked. At the bottom, there is a 'Save shortcut' link and two buttons: 'Connect' and 'Close'.

Step 2. Adding Virtual Infrastructure Servers

To protect virtual machines with Veeam Backup & Replication, you must add the virtual infrastructure servers hosting these machines to the backup infrastructure.

You can add vCenter Servers and ESXi hosts. If an ESXi host is managed by a vCenter Server, it is recommended that you add the vCenter Server, not a standalone ESXi host.

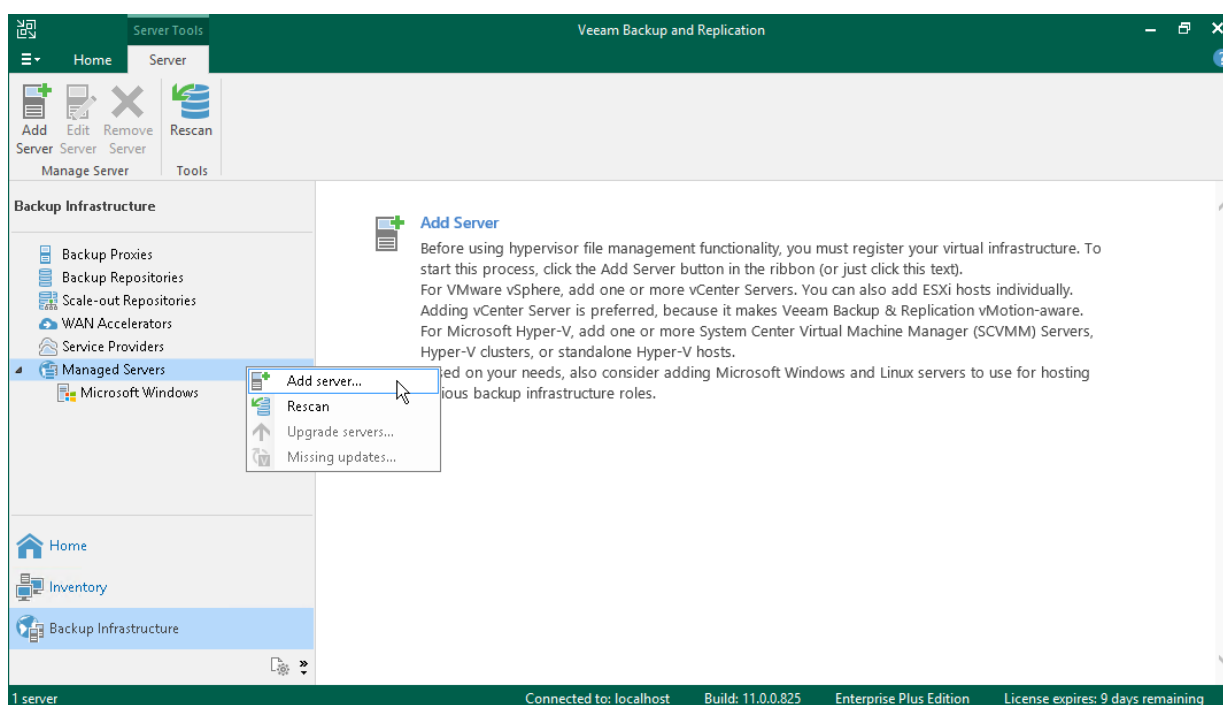
Before You Begin

Make sure that the version of your VMware vSphere platform is supported. For details, see [Platform Support](#) in the Veeam Backup & Replication User Guide.

Adding Infrastructure Server

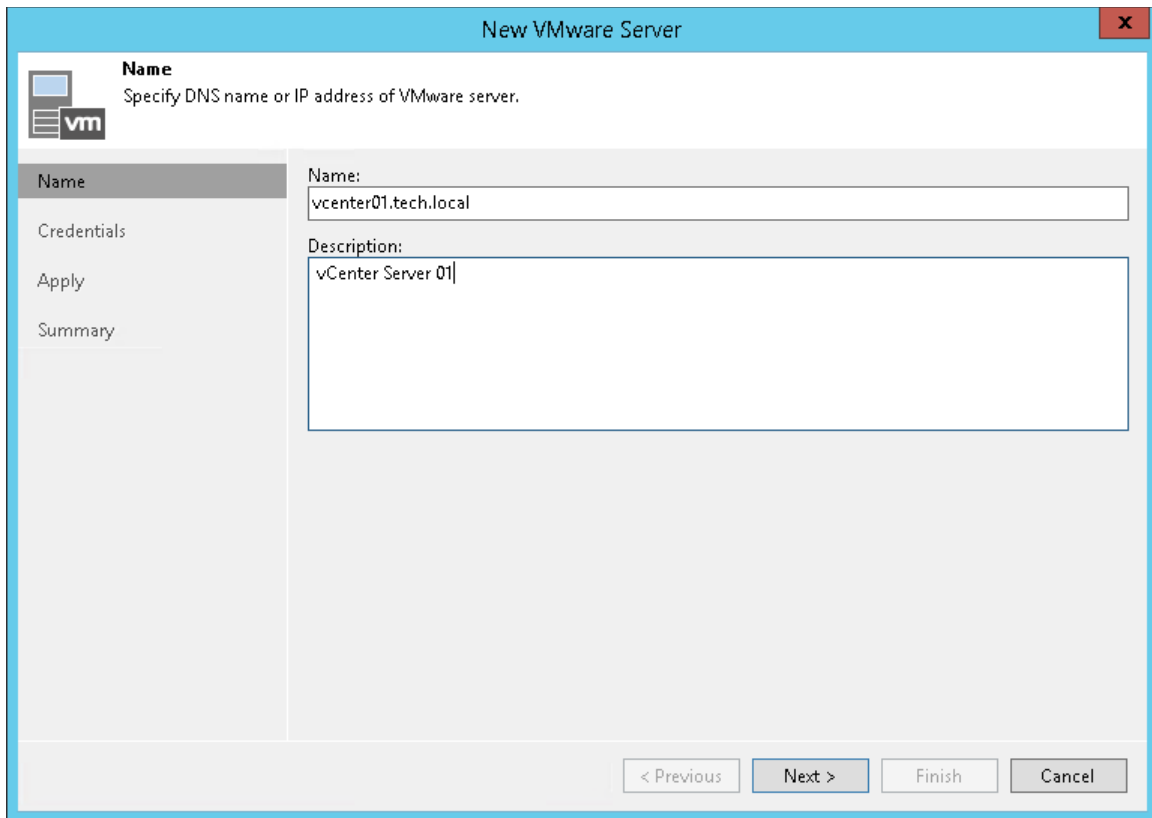
To add the server, do the following:

1. In the inventory pane of the **Backup Infrastructure** view, right-click the **Managed Servers** node and select **Add Server**.



2. In the **Add Server** window, click **VMware vSphere** > **vSphere** to launch the **New VMware Server** wizard.

- At the **Name** step of the wizard, specify the DNS name or IP address of the server.



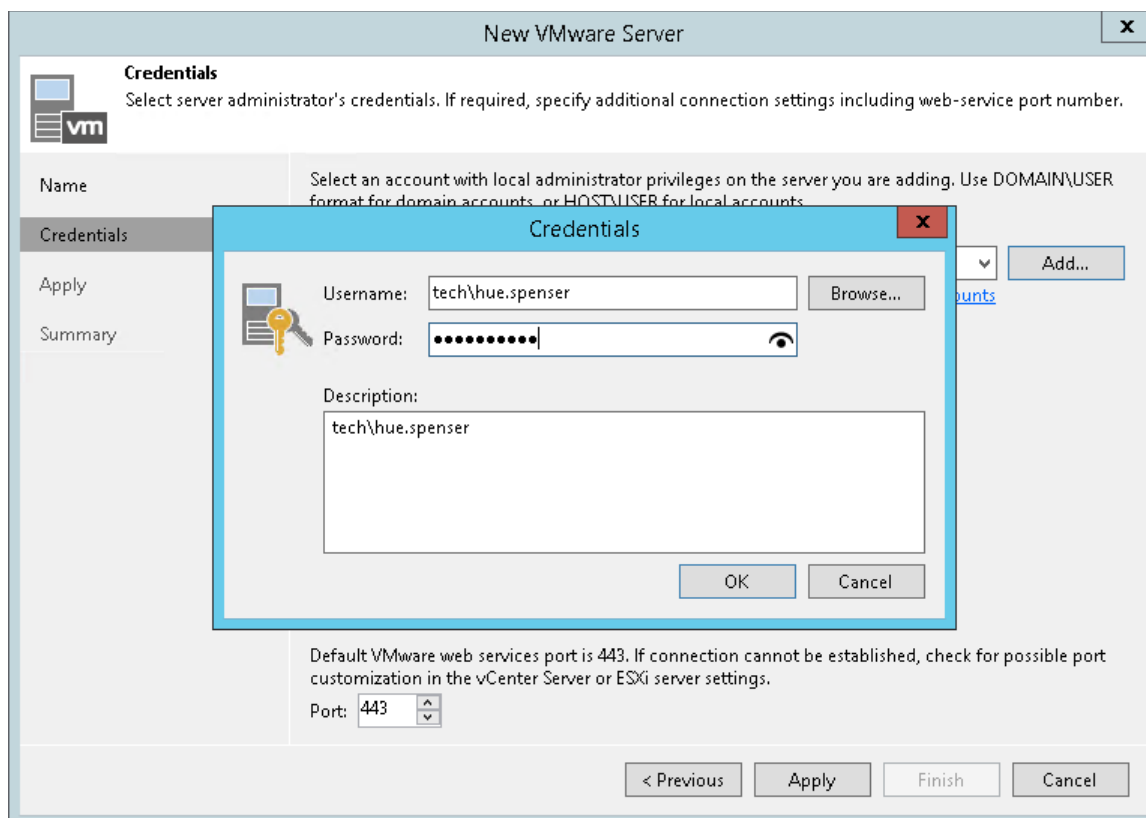
The screenshot shows the 'New VMware Server' wizard window. The title bar is blue with the text 'New VMware Server' and a red close button. The main window has a light blue header with the VMware logo and the text 'Name' and 'Specify DNS name or IP address of VMware server.' On the left, there is a sidebar with four tabs: 'Name' (selected), 'Credentials', 'Apply', and 'Summary'. The main area contains two text input fields. The first is labeled 'Name:' and contains the text 'vcenter01.tech.local'. The second is labeled 'Description:' and contains the text 'vCenter Server 01'. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

- At the **Credentials** step of the wizard, specify credentials for the user account with Administrator permissions on the added server.

To add the account, do the following:

- Click **Add**.
- In the **Credentials** window, specify the username and password used to connect to the added server.

c. Click **OK**.



5. Follow the next steps of the wizard. At the **Summary** step, click **Finish**.
6. Open the **Backup Infrastructure** view and click the **Managed Servers** node. The added server must be available in the working area.

Reference

For details on adding virtual infrastructure servers, see the [Virtualization Servers and Hosts](#) section in the Veeam Backup & Replication User Guide.

Step 3. Configuring Backup Proxy

The backup proxy retrieves data from the production storage, compresses, deduplicates and sends it to the backup repository.

To configure the backup proxy, you must add a Microsoft Windows or Linux server and assign the role of the backup proxy to it. In this section, you will learn how to add a Microsoft Windows proxy.

Before You Begin

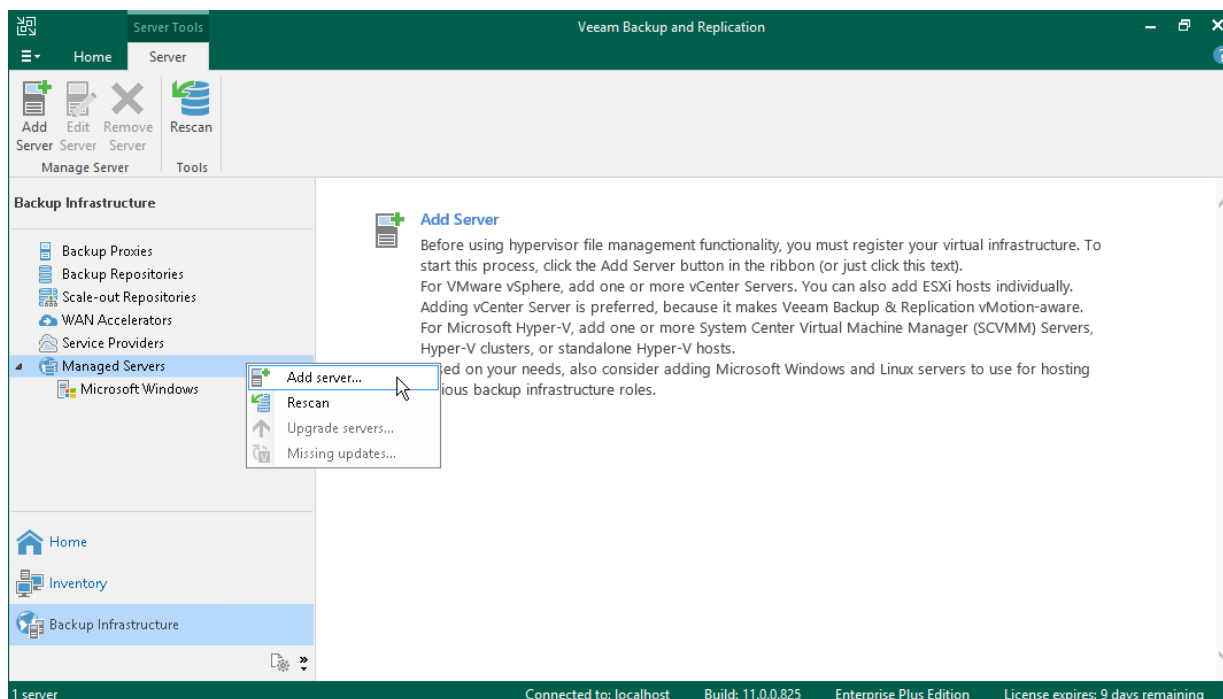
Check the following prerequisites:

- The machine that you plan to use as the backup proxy must meet system requirements. For details, see [System Requirements](#) in the Veeam Backup & Replication User Guide.
- The machine must have access to the backup server, source datastore and backup repository.
- File and printer sharing must be enabled in network connection settings of the added Microsoft Windows machine. On this machine, Veeam Backup & Replication deploys the required components. Without sharing enabled, Veeam Backup & Replication fails to deploy these components.

Adding Server

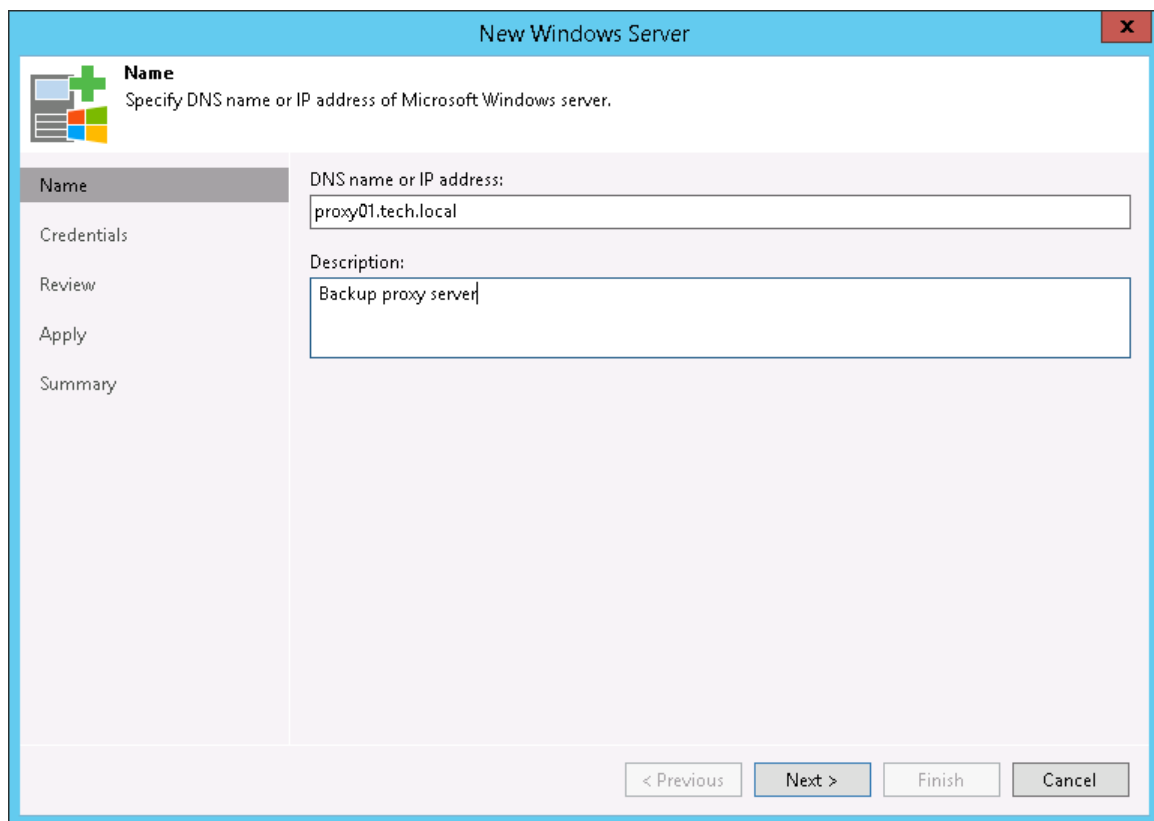
To add a server, do the following:

1. In the inventory pane of the **Backup Infrastructure** view, right-click the **Managed Servers** node and select **Add Server**.



2. In the **Add Server** window, click **Microsoft Windows** to launch the **New Windows Server** wizard.

3. At the **Name** step of the wizard, specify the DNS name or IP address of the server that will perform the role of the backup proxy.



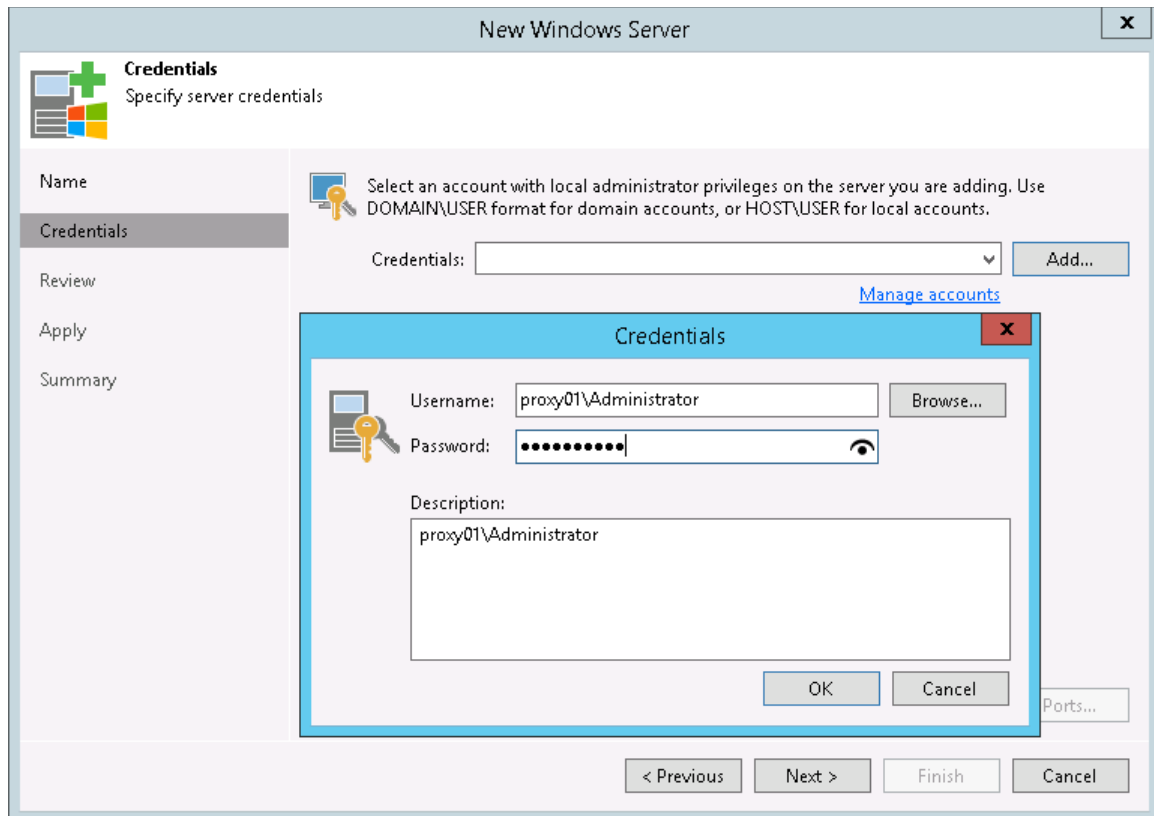
The screenshot shows the 'New Windows Server' wizard window. The title bar is blue with the text 'New Windows Server' and a close button. The main area has a light blue header with a green plus icon and the text 'Name' and 'Specify DNS name or IP address of Microsoft Windows server.' Below this is a sidebar with a list of steps: 'Name' (selected), 'Credentials', 'Review', 'Apply', and 'Summary'. The main content area has two text input fields. The first is labeled 'DNS name or IP address:' and contains the text 'proxy01.tech.local'. The second is labeled 'Description:' and contains the text 'Backup proxy server'. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

4. At the **Credentials** step of the wizard, specify credentials for the user account with Administrator permissions on the added server.

To add the account, do the following:

- a. Click **Add**.
- b. Specify the username and password used to connect to the added server.

c. Click **OK**.

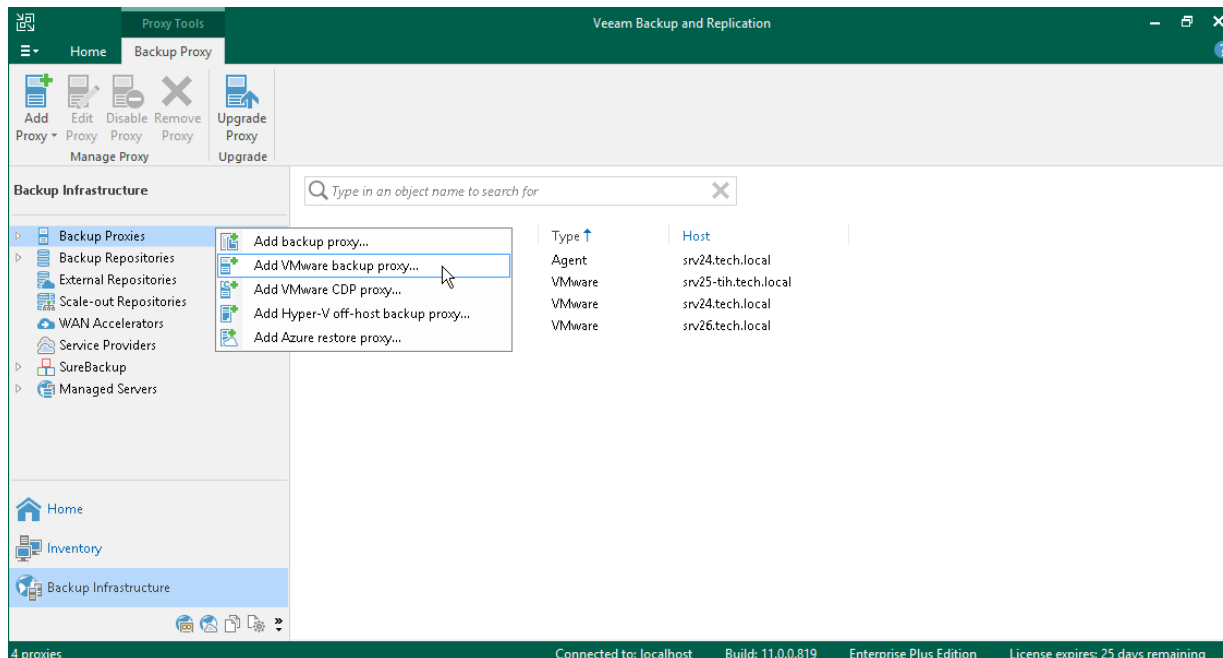


5. Follow the next steps of the wizard. At the **Summary** step of the wizard, click **Finish**.
6. Open the **Backup Infrastructure** view and click the **Managed Servers** node. The added server must be available in the working area.

Assigning Backup Proxy Role to Added Server

To assign the role of the backup proxy:

1. In the inventory pane of the **Backup Infrastructure** view, right-click the **Backup Proxies** node and select **Add VMware Backup Proxy**.



2. At the **Server** step of the wizard, do the following:
 - In the **Choose server** list, select the server that you have added.
 - In the **Transport mode** field, leave the **Automatic selection** option selected.

Veeam Backup & Replication will analyze the backup proxy configuration, define to which datastores it has access and automatically select the best way to retrieve and restore data depending on the type of connection between the backup proxy and the source datastore.

- In the **Connected datastores** field, leave the **Automatic detection** option selected.

Veeam Backup & Replication will detect datastores to which the backup proxy has a direct SAN or NFS connection.

3. At the **Traffic Rules** step of the wizard, keep the default settings.

Network traffic throttling rules help you manage bandwidth usage and minimize impact of data protection and disaster recovery tasks on network performance. Setting up traffic throttling is not described in this guide. For details see Veeam Backup & Replication User Guide, section [Network Traffic Management](#).

4. At the **Apply** step of the wizard, click **Next** and then **Finish** to exit the wizard.
5. Open the **Backup Infrastructure** view and click the **Backup Proxies** node. The added backup proxy must be available in the working area.

Reference

For details on the backup proxy, see the [Backup Proxy](#) section in the Veeam Backup & Replication User Guide.

Step 4. Configuring Backup Repository

The backup repository is a storage where Veeam Backup & Replication keeps backup files and, in case of replication, metadata for replicated VMs. You can use different types of storage as the backup repository. The full list of storage types is available in the [Backup Repository](#) section in the Veeam Backup & Replication User Guide.

In this section, you will learn how to use a Microsoft Windows server as the backup repository. To configure the backup repository, you must add the server to the backup infrastructure and assign the role of the backup repository to it.

Before You Begin

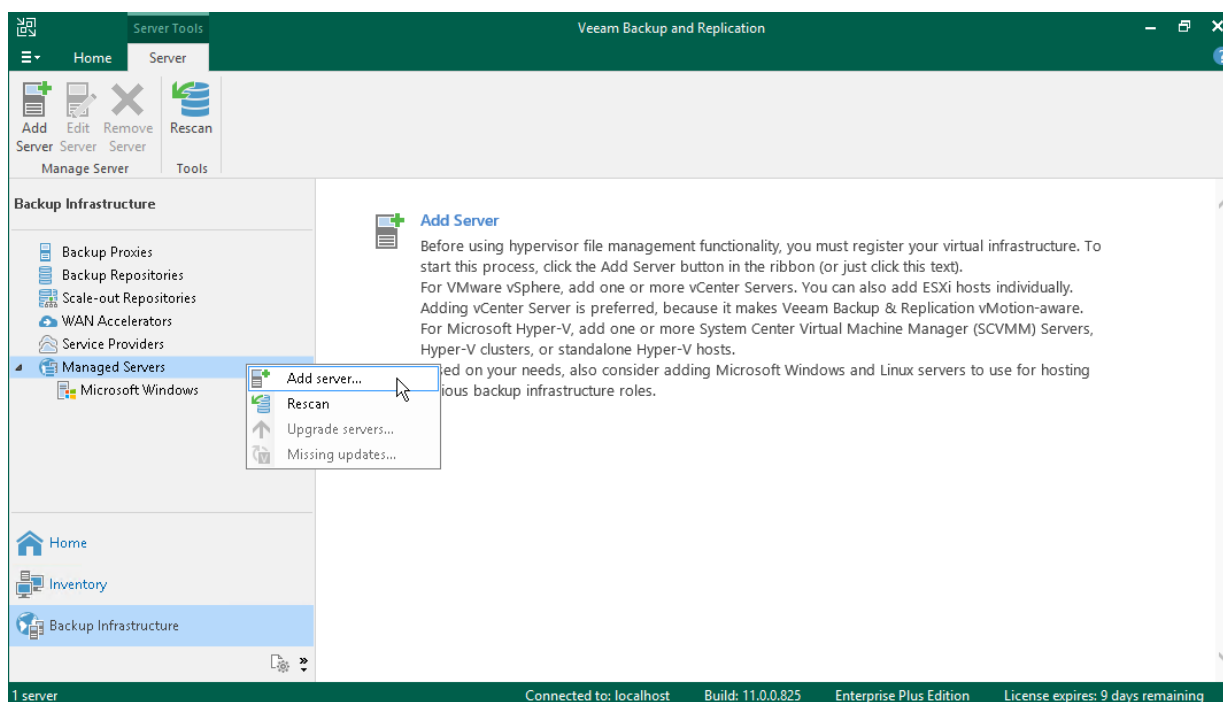
Check the following prerequisites:

- The Microsoft Windows machine that you plan to use as a backup repository must meet system requirements. For details, see [System Requirements](#) in the Veeam Backup & Replication User Guide.
- File and printer sharing must be enabled in network connection settings of the added Microsoft Windows machine. On this machine, Veeam Backup & Replication deploys the required components. Without sharing enabled, Veeam Backup & Replication fails to deploy these components.

Adding a Server

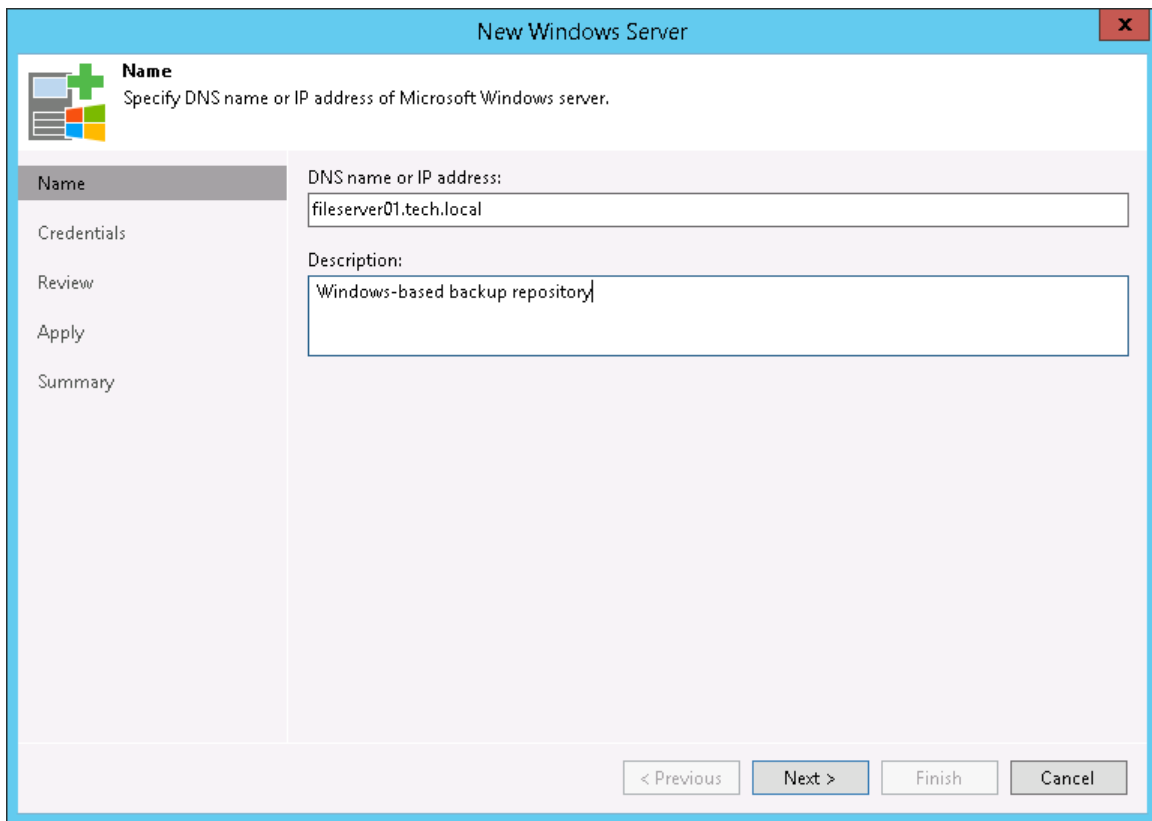
To add a server to the Veeam Backup & Replication infrastructure, do the following.

1. In the inventory pane of the **Backup Infrastructure** view, right-click the **Managed Servers** node and select **Add Server**.



2. In the **Add Server** window, select **Microsoft Windows** to launch the **New Windows Server** wizard.

3. At the **Name** step of the wizard, specify the DNS name or IP address of the server that will perform the role of the backup repository.



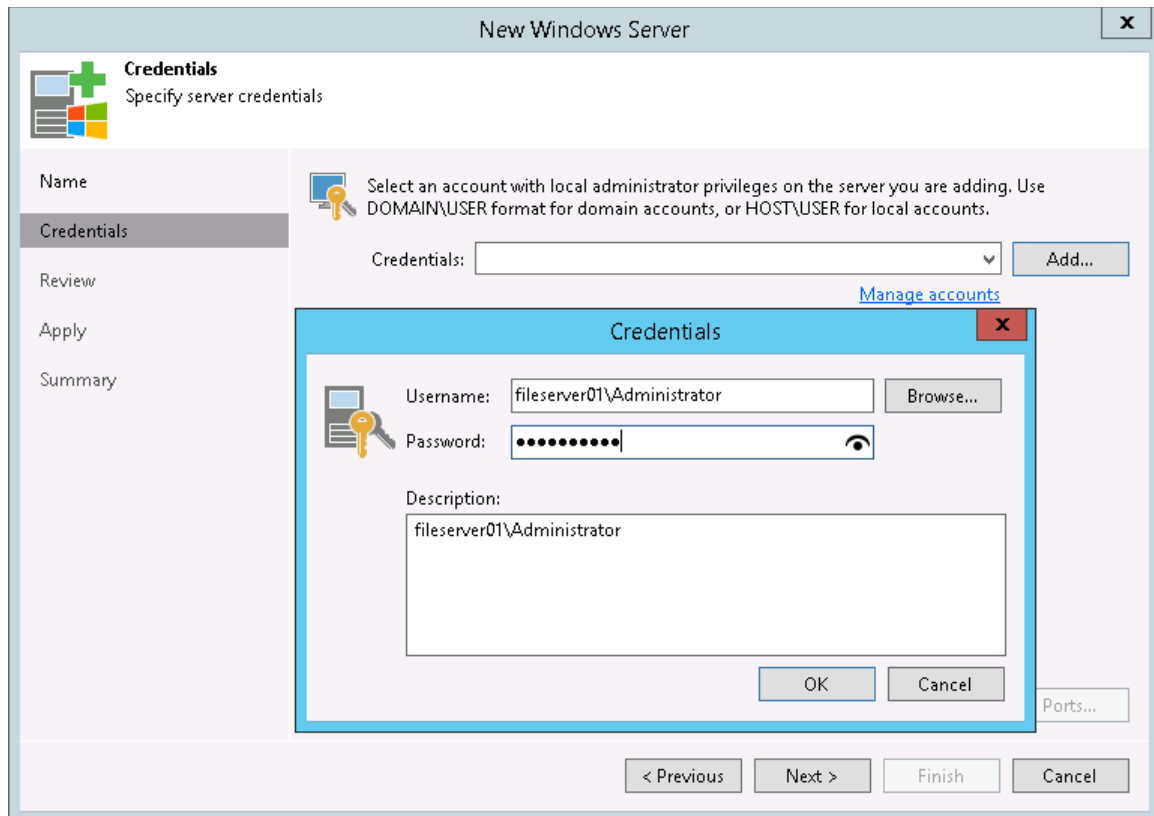
The screenshot shows the 'New Windows Server' wizard window. The title bar is blue with the text 'New Windows Server' and a close button. The main area has a light blue header with a green plus icon and the text 'Name' and 'Specify DNS name or IP address of Microsoft Windows server.' Below this is a sidebar with a list of steps: 'Name' (selected), 'Credentials', 'Review', 'Apply', and 'Summary'. The main content area has two text input fields. The first is labeled 'DNS name or IP address:' and contains the text 'fileserver01.tech.local'. The second is labeled 'Description:' and contains the text 'Windows-based backup repository'. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

4. At the **Credentials** step of the wizard, enter credentials for the user account with local Administrator permissions to the added server.

To add the account, do the following:

- a. Click **Add**.
- b. Specify the username and password used to connect to the added server.

c. Click **OK**.

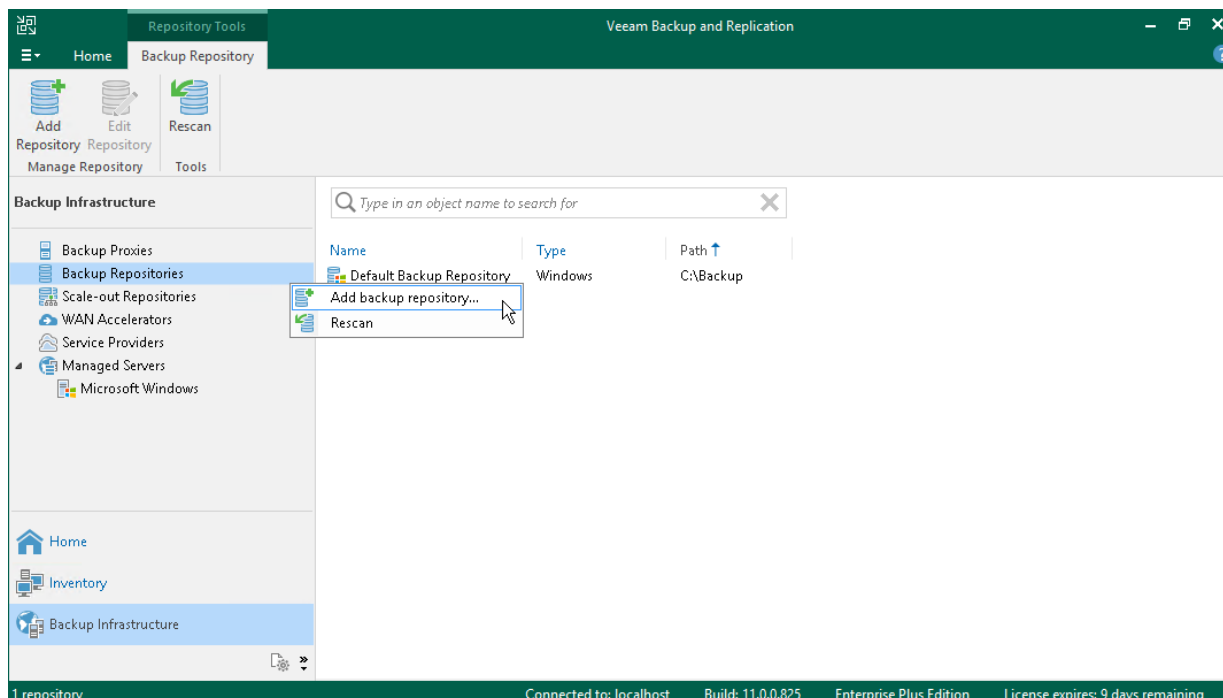


5. Follow the next steps of the wizard. At the last step of the wizard, click **Finish** to add the server.
6. Open the **Backup Infrastructure** view and click the **Managed Servers** node. The added server must be available in the working area.

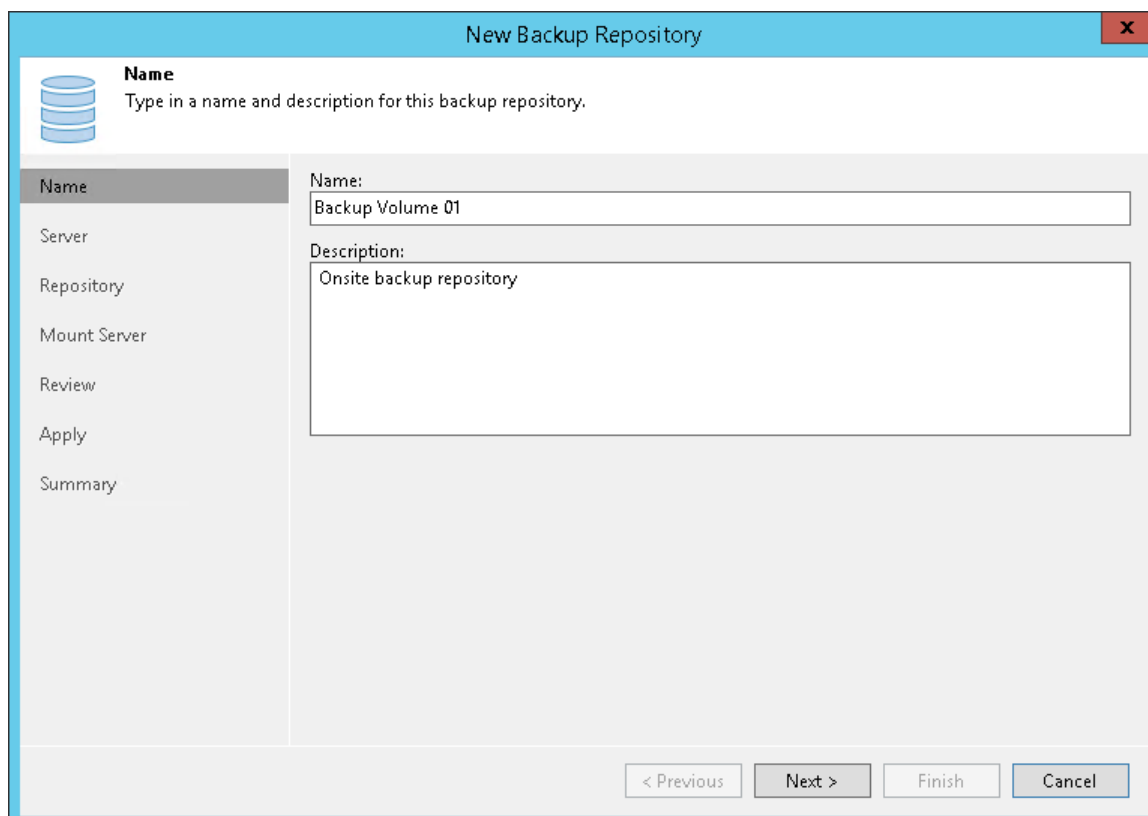
Assigning Backup Repository Role to Added Server

To assign the role of the backup repository:

1. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository** to launch the **New Backup Repository** wizard.



2. In the **Add Backup Repository** window, click **Direct attached storage > Microsoft Windows**.
3. At the **Name** step of the wizard, specify the name for the added backup repository.



- At the **Server** step of the wizard, select the machine that you have added.

The screenshot shows the 'New Backup Repository' wizard at the 'Server' step. The left sidebar contains a list of steps: Name, Server (selected), Repository, Mount Server, Review, Apply, and Summary. The main area is titled 'Server' and contains the instruction: 'Choose repository server. You can select server from the list of managed servers added to the console.' Below this, there is a 'Repository server:' dropdown menu showing 'fileserv01.tech.local (File Server 01)'. To the right of the dropdown are 'Add New...' and 'Populate' buttons. Below the dropdown is a table with three columns: 'Path', 'Capacity', and 'Free'. The first row shows 'C:\' with a capacity of '129.5 GB' and free space of '94.9 GB'. At the bottom of the wizard are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Path	Capacity	Free
C:\	129.5 GB	94.9 GB

- At the **Repository** step of the wizard, specify a path to the folder where backup files will be stored. In addition to them, auxiliary replica files will be placed in this folder.

The screenshot shows the 'New Backup Repository' wizard at the 'Repository' step. The left sidebar contains a list of steps: Name, Server, Repository (selected), Mount Server, Review, Apply, and Summary. The main area is titled 'Repository' and contains the instruction: 'Type in path to the folder where backup files should be stored, and set repository load control options.' Below this, there is a 'Location' section with a 'Path to folder:' text box containing 'C:\Backups' and a 'Browse...' button. Below the text box are 'Capacity:' and 'Free space:' labels with a 'Populate' button. Below the 'Location' section is a 'Load control' section with the text: 'Running too many concurrent tasks against the same repository may reduce overall performance, and cause I/O operations to timeout. Control storage device saturation with the following settings:'. There are two checkboxes: 'Limit maximum concurrent tasks to:' with a value of '4' and 'Limit read and write data rates to:' with a value of '0' MB/s. At the bottom of the wizard are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. There is also an 'Advanced...' button at the bottom right.

- At the **Mount Server** step of the wizard, keep the default settings.

7. Follow the next steps of the wizard. At the **Summary** step, click **Finish**.
8. Open the **Backup Infrastructure** view and click the **Backup Repositories** node. The added backup repository must be available in the working area.

Reference

For details on the backup repository, see the [Backup Repository](#) section in the Veeam Backup & Replication User Guide.

Step 5. Configuring Object Storage Repositories

An object storage repository is a repository intended for long-term data storage. It can be based on either a cloud solution or an S3 compatible storage solution. Configuring an object storage repository is an optional step.

Veeam Backup & Replication supports the following types of object storage repositories:

- S3 compatible
- Amazon S3, Amazon S3 Glacier and Amazon Snowball Edge
- Google Cloud
- IBM Cloud
- Microsoft Azure Blob, Azure Archive Storage and Azure Data Box

In this section, you will learn how to configure Amazon S3 as a backup repository.

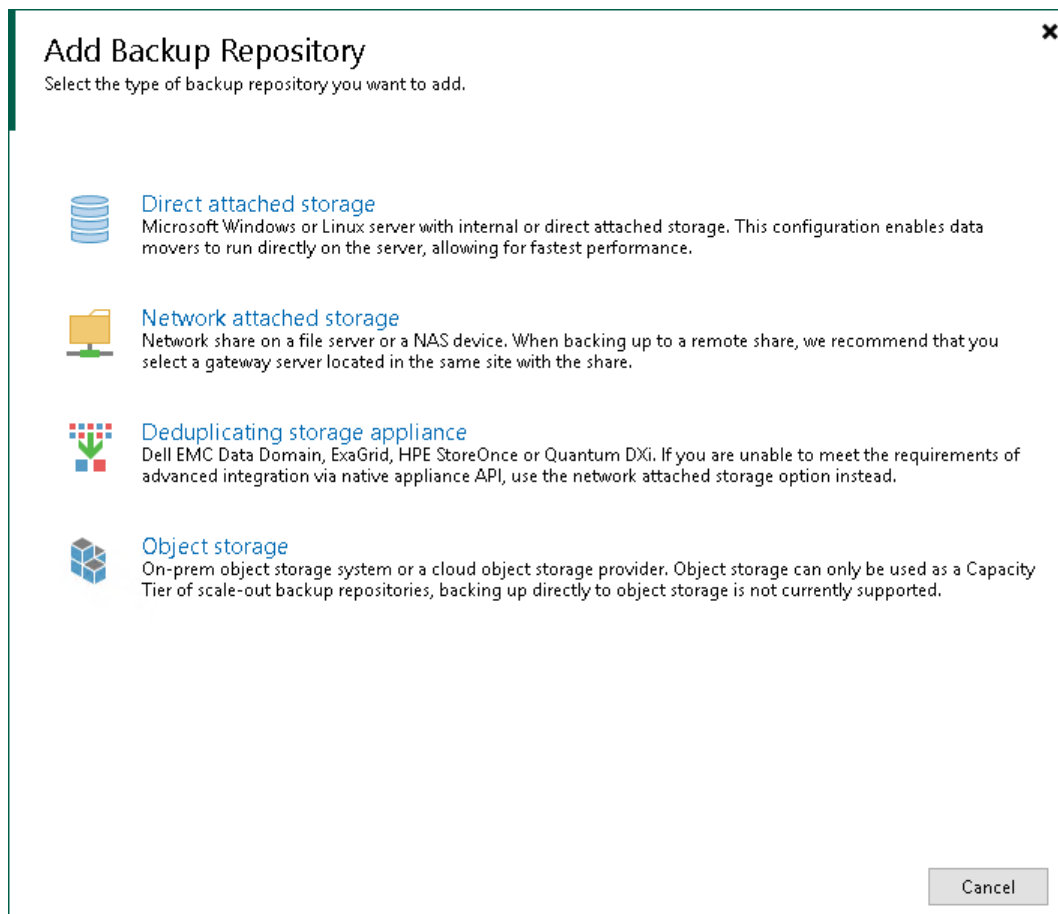
Before You Begin

Before you add an object storage repository, check limitations. To learn about limitations for different storage repositories, see [Considerations and Limitations](#) in the Veeam Backup & Replication User Guide.

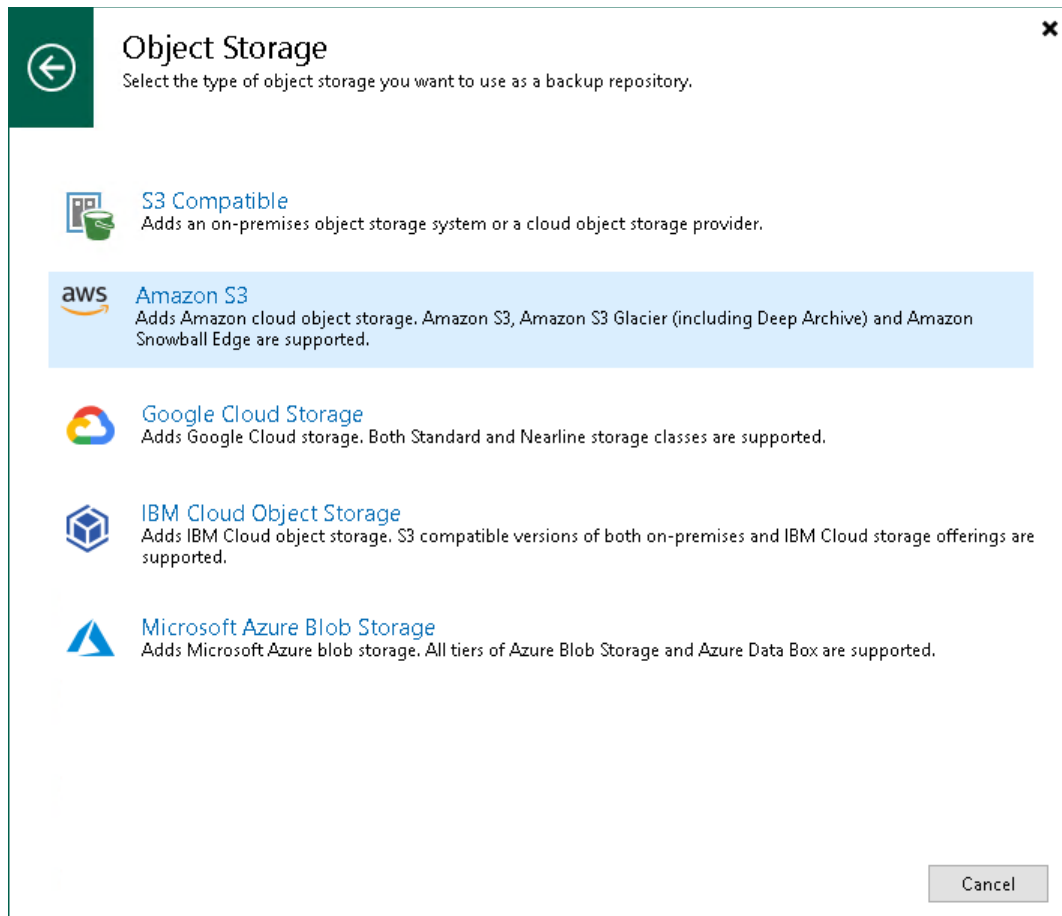
Configuring Object Storage Repository

To add an object storage repository to the Veeam Backup & Replication infrastructure, do the following:

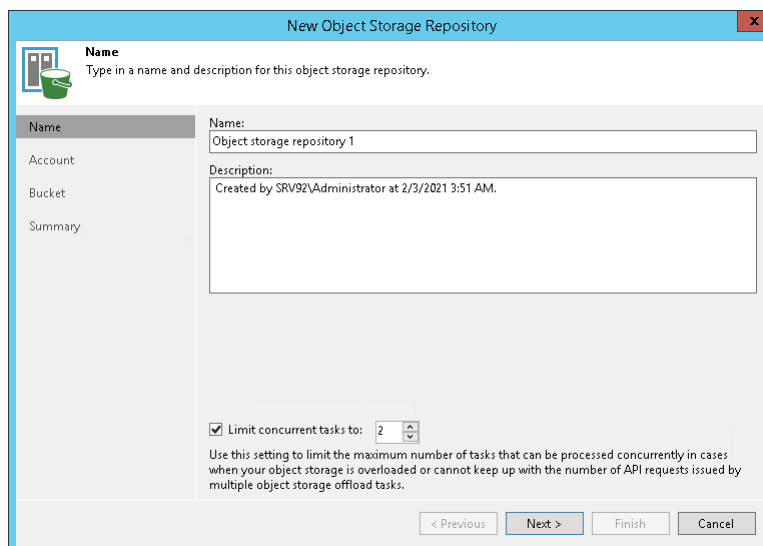
1. Open the **Backup Infrastructure** view. In the inventory pane, right-click the **Backup Repositories** node and select **Add Backup Repository**. In the **Add Backup Repository** dialog, select **Object Storage**.



2. In the **Object Storage** dialog, select **Amazon S3**.



3. At the **Name** step of the wizard, specify the name and description for the object storage repository.



4. At the **Account** step of the wizard, specify the connection settings:
 - a. [For AWS Snowball Edge, Azure Data Box, S3 Compatible, IBM Cloud] In the **Service point / Service endpoint** field, specify a service point address of your object storage.
 - b. From the **Credentials** drop-down list, select user credentials to access your object storage. If you already have a credentials record that was configured in advance, select such a record from the drop-down list. Otherwise, click **Add** and provide your access and secret keys. For more information, see [Cloud Credentials Manager](#) in the Veeam Backup & Replication User Guide.

- c. In the **Region / Data center region** drop-down list, select a region type.

If your organization has NAT or different types of firewalls and your access to the internet is limited, you may want to use a gateway server. To do so, select the **Use the following gateway server** check box and choose a server from the list.

The screenshot shows the 'New Object Storage Repository' wizard, specifically the 'Account' step. The window title is 'New Object Storage Repository'. The subtitle is 'Specify account to use for connecting to S3 compatible storage system.' The left sidebar has four tabs: 'Name', 'Account' (selected), 'Bucket', and 'Summary'. The main area contains the following fields and controls:

- Service point:** A text field containing 'https://172.17.53.37:9000'.
- Region:** A text field containing 'us-east-1'.
- Credentials:** A text field containing a long string of 'x' characters. To the right of the field is an 'Add...' button and a link labeled 'Manage cloud accounts'.
- Use the following gateway server:** A checked checkbox. Below it is a dropdown menu showing 'srv92.tech.local (Backup server)'. A note below the dropdown states: 'Select a gateway server to proxy access to the object storage system. If no gateway server is specified, all scale-out backup repository extents must have direct network access to the storage system.'

At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

5. At the **Bucket** step of the wizard, specify how your data will be stored.

- a. [For Amazon S3] From the **Data center region** drop-down list, select a region.
- b. From the **Bucket** drop-down list, select a bucket. Make sure that the bucket you want to use to store your backup data was created in advance.
- [For Azure Blob, Azure Data Box] From the **Container** drop-down list, select a container. Make sure that the container you want to use to store your backup data was created in advance.
- c. In the **Folder / Select Folder** field, select a cloud folder to which you want to map your object storage repository. To do it, click **Browse** and either select an existing folder or create a new one by clicking **New Folder**.

To define a soft limit for your object storage consumption that can be exceeded temporarily, select the **Limit object storage consumption to** check box and provide the value in TB or PB.

[For Amazon S3, S3 Compatible] To prohibit deletion of blocks of data from object storage, select the **Make recent backups immutable for** check box and specify the immutability period.

[For Amazon S3] If you plan to access your backup data rarely, select the **Use infrequent access storage class** check box. To enable Amazon S3 One Zone-Infrequent Access, select the **Store backups in a single availability zone only** check box. For more information, see [this Amazon article](#).

The screenshot shows the 'New Object Storage Repository' wizard in the 'Bucket' step. The left sidebar has tabs for 'Name', 'Account', 'Bucket' (selected), and 'Summary'. The main area is titled 'Bucket' with the instruction 'Specify object storage system bucket to use.' It contains a 'Bucket:' dropdown menu with 'osrv10' selected, a 'Folder:' text input with 'OSRv10' and a 'Browse...' button, and two checked options: 'Limit object storage consumption to: 10 TB' and 'Make recent backups immutable for: 30 days (increases costs)'. The bottom of the wizard has four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

6. At the **Summary** step of the wizard, review the settings and click **Finish**.
7. Open the **Backup Infrastructure** view and click the **Backup Repositories** node. The added object storage repository must be available in the working area.

Reference

For details on adding object storage repositories, see the [Adding Object Storage Repositories](#) section in the Veeam Backup & Replication User Guide.

Step 6. Configuring Scale-Out Backup Repositories

A scale-out backup repository is a repository system for multi-tier storage of data, where the capacities of all the added storage devices and systems are summarized. A scale-out backup repository consists of one or more backup repositories called performance extents, and can be expanded with an object repository called capacity extent.

Configuring a scale-out backup repository is an optional step. For more information on scale-out backup repositories, see the [Scale-Out Backup Repository](#) section in the Veeam Backup & Replication User Guide.

Before You Begin

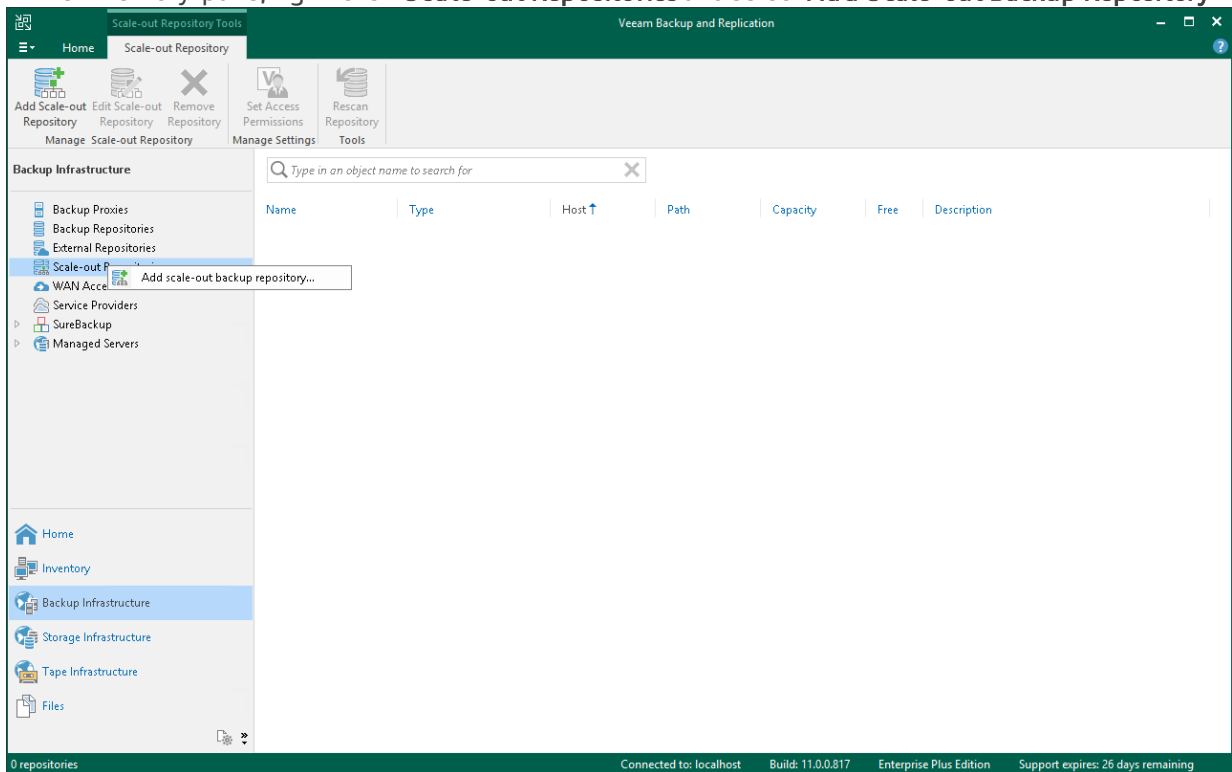
Before you add a scale-out backup repository to the backup infrastructure, check the following prerequisites:

- Backup repositories that you plan to add as performance extents to the scale-out backup repository must be added to the backup infrastructure. For more information, see [Configuring Backup Repository](#).
- If you wish to use the capacity tier option of the scale-out backup repository, an object storage repository that you plan to add as a capacity extent to the scale-out backup repository must be added to the backup infrastructure. For more information, see [Configuring Object Storage Repository](#).
- Check limitations for scale-out backup repositories. For more information, see [Limitations for Scale-Out Backup Repositories](#) in the Veeam Backup & Replication User Guide.

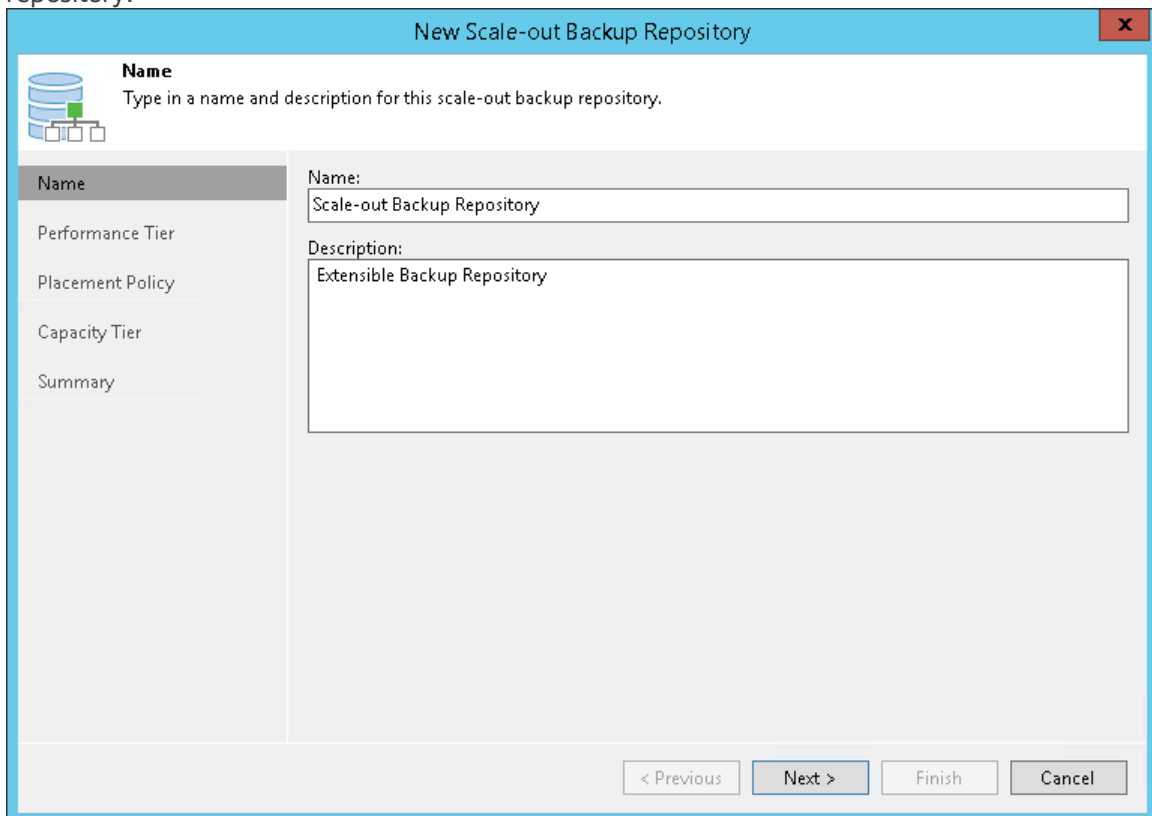
Configuring Scale-Out Backup Repository

To add a scale-out backup repository to the Veeam Backup & Replication infrastructure, do the following:

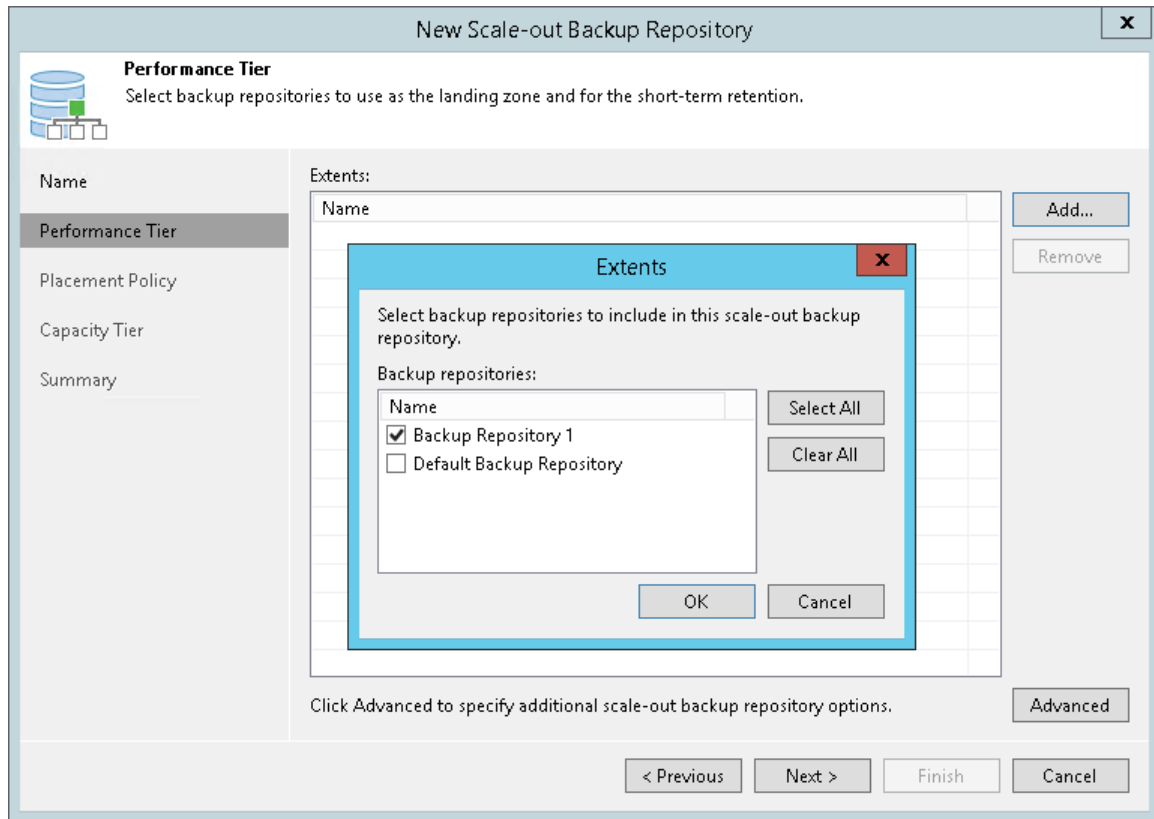
1. In the inventory pane, right-click **Scale-out Repositories** and select **Add Scale-out Backup Repository**.



2. At the **Name** step of the wizard, specify a name and an optional description for the scale-out backup repository.



- At the **Performance Tier** step of the wizard, specify which backup repositories you want to add as performance extents, and configure options for the scale-out backup repository. To do it, on the right side of the **Extents** list, click **Add**. In the **Extents** window, select check boxes next to backup repositories that you want to add as performance extents. Afterwards, click **OK**.



- At the **Policy** step of the wizard, specify how you want to store backup files at the performance extents of the scale-out backup repository:
 - Select **Data locality** if you want to store backup files that belong to the same backup chain at the same performance extent.

- Select **Performance** if you want to store full and incremental backup files at different performance extents of the scale-out backup repository.

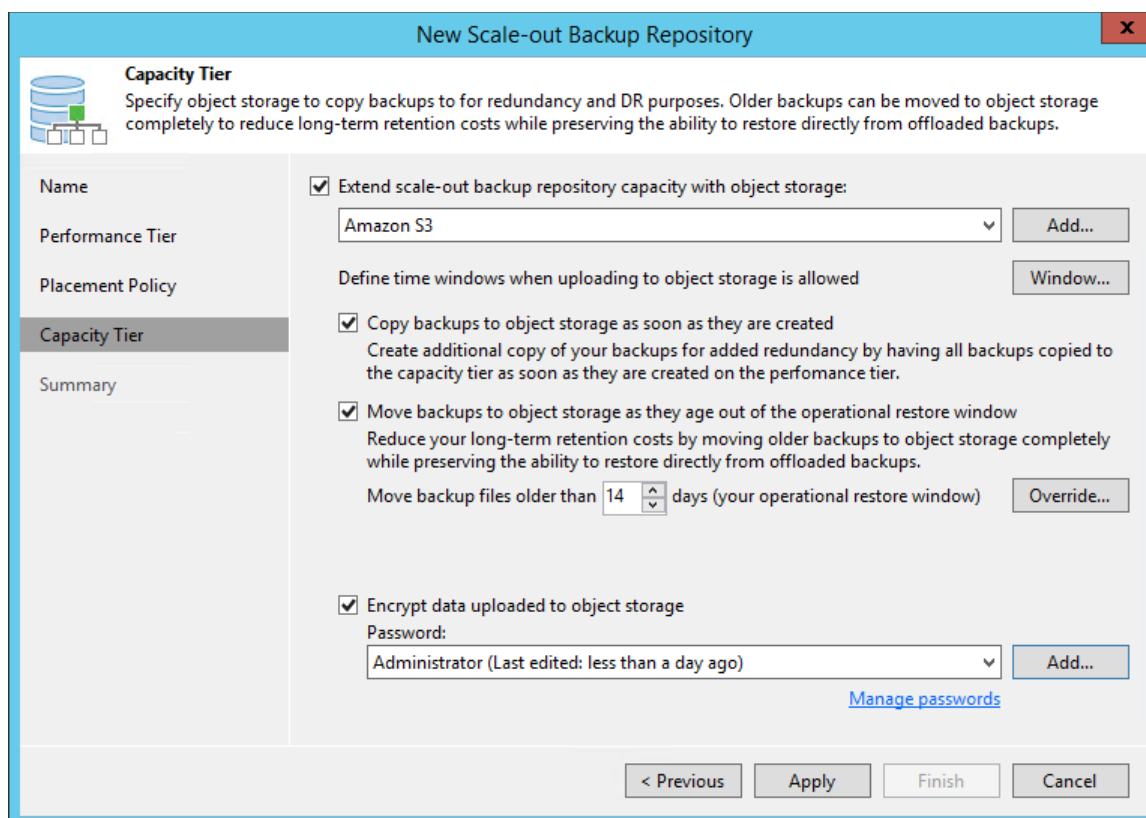
The screenshot shows the 'New Scale-out Backup Repository' wizard window. The title bar is blue with the text 'New Scale-out Backup Repository' and a red close button. The main content area has a light blue header with a database icon and the title 'Placement Policy'. Below the header, there is a description: 'Choose a backup files placement policy for this performance tier. When more than one extent matches the placement policy, backup job will chose extent with the most free disk space available.' On the left, there is a sidebar with a tree view containing 'Name', 'Performance Tier', 'Placement Policy' (which is selected and highlighted), 'Capacity Tier', and 'Archive Tier'. The main area displays two radio button options: 'Data locality' (unselected) and 'Performance' (selected). The 'Performance' option has a detailed description: 'Incremental backup files are placed on a different extent from the corresponding full backup file, providing for better backup file transformation performance with raw storage devices. Note that losing an extent with a full backup makes restoring from increments impossible.' Below this, there is a text field with the placeholder 'Specify the placement policy for full and incremental backup files.' and a 'Customize...' button. At the bottom, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

5. At the **Capacity Tier** step of the wizard, select an object storage repository that you want to add as a capacity extent and specify when to move and/or copy data. This is an optional step.

To configure the capacity tier, do the following:

- a. Select the **Extend scale-out backup repository capacity with object storage** check box.
- b. From the drop-down list, select an object storage repository to which you want to offload your data.
- c. Click **Window** and specify when it is allowed or prohibited to move or copy data to object storage.
- d. Select the **Copy backups to object storage as soon as they are created** check box to copy new backups as soon as they are created.
- e. Select the **Move backups to object storage as they age out of the operational restores window** check box to move inactive backup chains to the capacity extent.

- f. To offload data encrypted, select the **Encrypt data uploaded to object storage** check box and provide a strong password. If you have not created the password beforehand, click **Add** or use the **Manage passwords** link to specify a new password.



The screenshot shows the 'New Scale-out Backup Repository' wizard, specifically the 'Capacity Tier' step. The left sidebar contains a tree view with 'Name', 'Performance Tier', 'Placement Policy', 'Capacity Tier' (selected), and 'Summary'. The main area has a title bar 'New Scale-out Backup Repository' and a close button. Below the title bar is a description: 'Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.' The main configuration area includes several options: 'Extend scale-out backup repository capacity with object storage:' (checked) with a dropdown set to 'Amazon S3' and an 'Add...' button; 'Define time windows when uploading to object storage is allowed' with a 'Window...' button; 'Copy backups to object storage as soon as they are created' (checked) with a description; 'Move backups to object storage as they age out of the operational restore window' (checked) with a description and a 'Move backup files older than' spinner set to '14' days, with an 'Override...' button; and 'Encrypt data uploaded to object storage' (checked) with a 'Password:' dropdown set to 'Administrator (Last edited: less than a day ago)' and an 'Add...' button. A 'Manage passwords' link is also present. At the bottom are buttons for '< Previous', 'Apply', 'Finish', and 'Cancel'.

6. At the **Summary** step of the wizard, review the settings and click **Finish**.
7. Open the **Backup Infrastructure** view and click the **Scale-Out Repositories** node. The added scale-out backup repository must be available in the working area.

Reference

For details on adding scale-out backup repositories, see the [Adding Scale-Out Backup Repositories](#) section in the Veeam Backup & Replication User Guide.

VM Backup

When you perform a backup, Veeam Backup & Replication retrieves VM data from the source storage, compresses and deduplicates it. After that, Veeam Backup & Replication writes data to the backup repository in Veeam proprietary format.

In Veeam Backup & Replication, backup is a job-driven process. To perform the backup, you need to configure a backup job. For details, see [Creating Backup Job](#).

During runs of backup jobs, Veeam Backup & Replication creates backup chains. The chain consists of the following backup files:

- Full backup file (.Vbk) that contains a copy of the entire VM.
- Incremental backup file (.Vib or .Vrb) that contains only those data blocks that have changed since the last backup job session.
- Metadata file (.Vbm) that contains information on the backup job, VMs in the backup, number and structure of backup files, restore points and so on.

The amount of these files and how Veeam Backup & Replication places them depend on the chosen backup method. For details, see [Backup Methods](#).

After you performed backups, you can use them to restore the following instances: entire VM, VM files or VM disks, guest OS files and application items. For details on restore, see [Data Restore](#).

Reference

For details, see the [About Backup](#) section in the Veeam Backup & Replication User Guide.

Backup Methods

Veeam Backup & Replication provides three methods for creating backup chains:


- Forever forward incremental backup
- Forward incremental backup
- Reverse incremental backup

Forever Forward Incremental Backup Method

The forever forward incremental backup method produces a backup chain that consists of the first full backup file (.VBK) and a set of forward incremental backup files (.VIB) following it.

During the first session of a backup job, Veeam Backup & Replication creates a full backup file on the backup repository. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session and saves these blocks as an incremental backup file in the backup chain.

After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy for the job and deletes outdated restore points. For details, see [Forever Forward Incremental Backup Retention Policy](#) in the Veeam Backup & Replication User Guide.



Forward Incremental Backup Method

The forward incremental backup method produces a backup chain that consists of the first full backup file (.VBK) and a set of forward incremental backup files (.VIB) following it.

Additionally, the forward incremental backup chain contains full backup files that “split” the backup chain into shorter series. The subsequent full backup files can be the following:

- **Active**


The active full backup contains the copy of a VM. This backup is similar to the full backup created when you run a job for the first time. Veeam Backup & Replication retrieves data for the whole VM from the source, compresses and deduplicates it and stores it to the active full backup file.

- **Synthetic**

The synthetic full backup also contains the copy of a VM. However, this copy is created from the backup files that you already have on the backup repository. Veeam Backup & Replication does not retrieve VM data from the source datastore.

During the first backup job session, Veeam Backup & Replication creates a full backup file. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session and saves these blocks as an incremental backup file in the backup chain. On a day when the synthetic or active full backup is scheduled, Veeam Backup & Replication creates a full backup file and adds it to the backup chain.

After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy and deletes outdated restore points. For details, see [Forward Incremental Backup Retention Policy](#) in the Veeam Backup & Replication User Guide.




Reverse Incremental Backup Method

The reverse incremental backup method produces a backup chain that consists of the last full backup file (.Vbk) and a set of reverse incremental backup files (.VRB) preceding it.

During the first backup job session, Veeam Backup & Replication creates a full backup file on the backup repository. During subsequent backup job sessions, Veeam Backup & Replication copies only VM data blocks that have changed since the last backup job session. Veeam Backup & Replication "injects" copied data blocks into the full backup file to rebuild it to the most recent state of the VM. Additionally, Veeam Backup & Replication saves the changed block in the reverse incremental backup file and places this file before the full backup file.

After adding a new restore point to the backup chain, Veeam Backup & Replication checks the retention policy and deletes outdated restore points. For details, see [Reverse Incremental Backup Retention Policy](#) in the Veeam Backup & Replication User Guide.



Reference

For details, see the following section in the Veeam Backup & Replication User Guide:

- [Retention Policy](#)
- [Active Full Backup](#)
- [Synthetic Full Backup](#)
- [Backup Methods](#)

Creating Backup Job

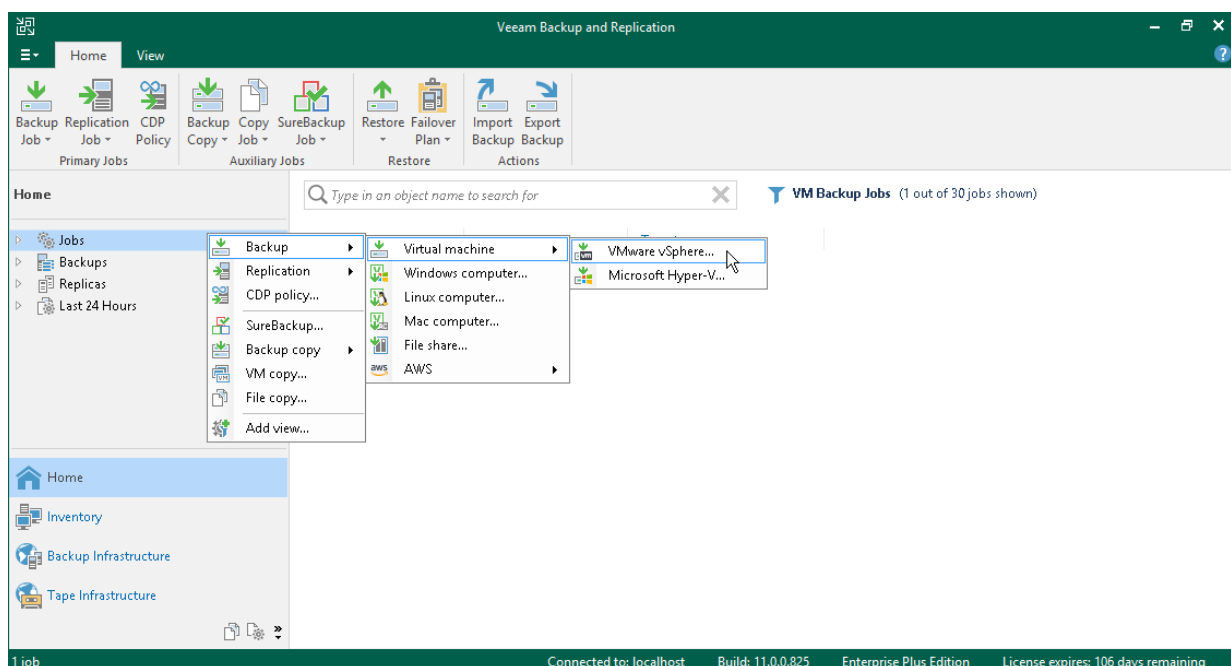
Before You Begin

Make sure that all backup infrastructure components that take part in the backup process are added to the backup infrastructure. These components include ESXi hosts on which VMs are registered, backup proxy and backup repository.

Creating Backup Job

To back up VMs, do the following:

1. In the inventory pane of the **Home** view, right-click **Jobs** and select **Backup > Virtual Machine > VMware vSphere** to launch the **New Backup Job** wizard.



- At the **Name** step of the wizard, specify a name and description for the backup job.

The screenshot shows the 'New Backup Job' wizard in Veeam Backup & Replication. The 'Name' step is active, indicated by a green arrow and the 'vm' icon. The wizard title is 'New Backup Job'. Below the title, it says 'Type in a name and description for this backup job.' On the left, there is a sidebar with steps: Name, Virtual Machines, Storage, Guest Processing, Schedule, and Summary. The 'Name' step is selected. The main area has two text boxes: 'Name:' with the value 'DB Backup' and 'Description:' with the value 'Daily Backup Job'. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

- At the **Virtual Machines** step of the wizard, click **Add**. From the list, select VMs that you want to back up.

You can also back up VM containers: folders, resource pools, clusters, vApps, datastores and so on. If you add a new VM to the container after the backup job is created, Veeam Backup & Replication automatically updates the job to include the new VM.

The screenshot shows the 'Add Objects' dialog box in Veeam Backup & Replication. The dialog is titled 'Add Objects'. It has a sidebar on the left with steps: Name, Virtual Machines, Storage, Guest Processing, Schedule, and Summary. The 'Virtual Machines' step is selected. The main area is titled 'Select objects:' and shows a tree view of the backup infrastructure. The tree structure is: Hosts and Clusters > vcenter01.tech.local > Atlanta > esx01.tech.local > vApp01 > Alpha > DB > db01. The 'db01' object is selected. Below the tree, there is a search bar with the text 'Type in an object name to search for'. At the bottom, there are 'Add' and 'Cancel' buttons. On the right side of the dialog, there are buttons for 'Add...', 'Remove', 'Exclusions...', 'Up', 'Down', 'Recalculate', and 'Total size: 0 B'.

4. At the **Storage** step of the wizard, do the following:

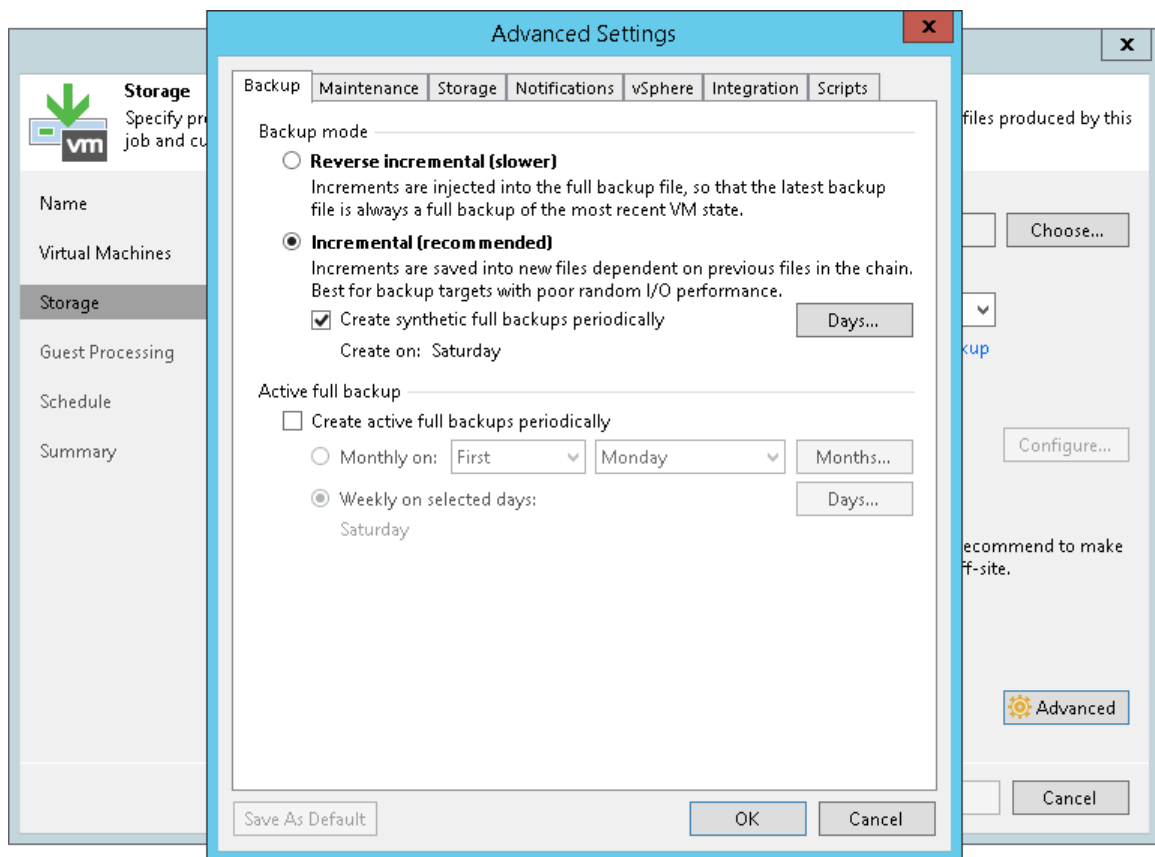
- From the **Backup repository** list, select the backup repository that you configured in the [Configure Backup Repository](#) section.
- In the **Restore points to keep on disk**, define the number of restore points to keep.

When the number of restore points exceeds the allowed value, Veeam Backup & Replication automatically removes the earliest restore point from the backup chain. For more information, see [Retention Policy](#) in the Veeam Backup & Replication User Guide.

The screenshot shows the 'New Backup Job' wizard in Veeam Backup & Replication, specifically the 'Storage' step. The window has a blue title bar and a sidebar on the left with navigation options: Name, Virtual Machines, Storage (selected), Guest Processing, Schedule, and Summary. The main area is titled 'Storage' and contains the following elements:

- Backup proxy:** A text box with 'Automatic selection' and a 'Choose...' button.
- Backup repository:** A dropdown menu showing 'Backup Volume 01' with a 'Map backup' link to its right.
- Retention policy:** A numeric spinner set to '14' and a dropdown menu set to 'restore points', with an information icon to the right.
- ☐ **Keep some periodic full backups longer for archival purposes** (with a 'Configure...' button and the note 'GFS retention policy is not configured').
- ☐ **Configure secondary destinations for this job** (with the text: 'Copy backups produced by this job to another backup repository, or to tape. Best practices recommend maintaining at least 2 backups of production data, with one of them being off-site.').
- Advanced job settings:** A link to view settings for backup mode, compression, deduplication, block size, notifications, and post-job activity.
- Buttons:** '< Previous', 'Next >', 'Finish', and 'Cancel' at the bottom.

- Click the **Advanced** button and, in the **Backup** tab, specify the backup method or leave the default settings. For details, see [Backup Methods](#).



5. At the **Guest Processing** step of the wizard, leave the default settings.

The settings of this step are detailed in the [Creating Application-Aware Backup Job](#) section.

6. At the **Schedule** step of the wizard, do the following:
 - a. Select the **Run the job automatically** check box. If you do not select this check box, you will have to launch the job manually. For details, see [Start Backup Job Manually](#).
 - b. Select the schedule type: daily, monthly or periodically.
 - c. Make sure the **Retry failed VM processing** check box is selected.

d. Click **Apply**.

The screenshot shows the 'New Backup Job' wizard in Veeam Backup & Replication, specifically the 'Schedule' step. The left sidebar contains a tree view with the following items: Name, Virtual Machines, Storage, Secondary Target, Guest Processing, **Schedule** (highlighted), and Summary. The main area is titled 'Schedule' and includes a sub-header: 'Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.' Below this, there are several configuration sections: 1. 'Run the job automatically' (checked): Includes radio buttons for 'Daily at this time:', 'Monthly at this time:', 'Periodically every:', and 'After this job:'. The 'Daily' option is selected, showing a time of '10:00 PM' and a frequency of 'Everyday'. 2. 'Automatic retry': Includes a checked checkbox for 'Retry failed VMs processing:' with a value of '3' times, and a 'Wait before each retry attempt for:' of '10' minutes. 3. 'Backup window': Includes an unchecked checkbox for 'Terminate job if it exceeds allowed backup window' and a 'Window...' button. A note below states: 'If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.' At the bottom right, there are four buttons: '< Previous', 'Apply' (highlighted), 'Finish', and 'Cancel'.

7. At the **Summary** step of the wizard, select the **Run the job when I click Finish** check box and click the **Finish** button.
8. In the inventory pane of the **Home** view, expand the **Last 24 Hours** node to see the created job.

Reference

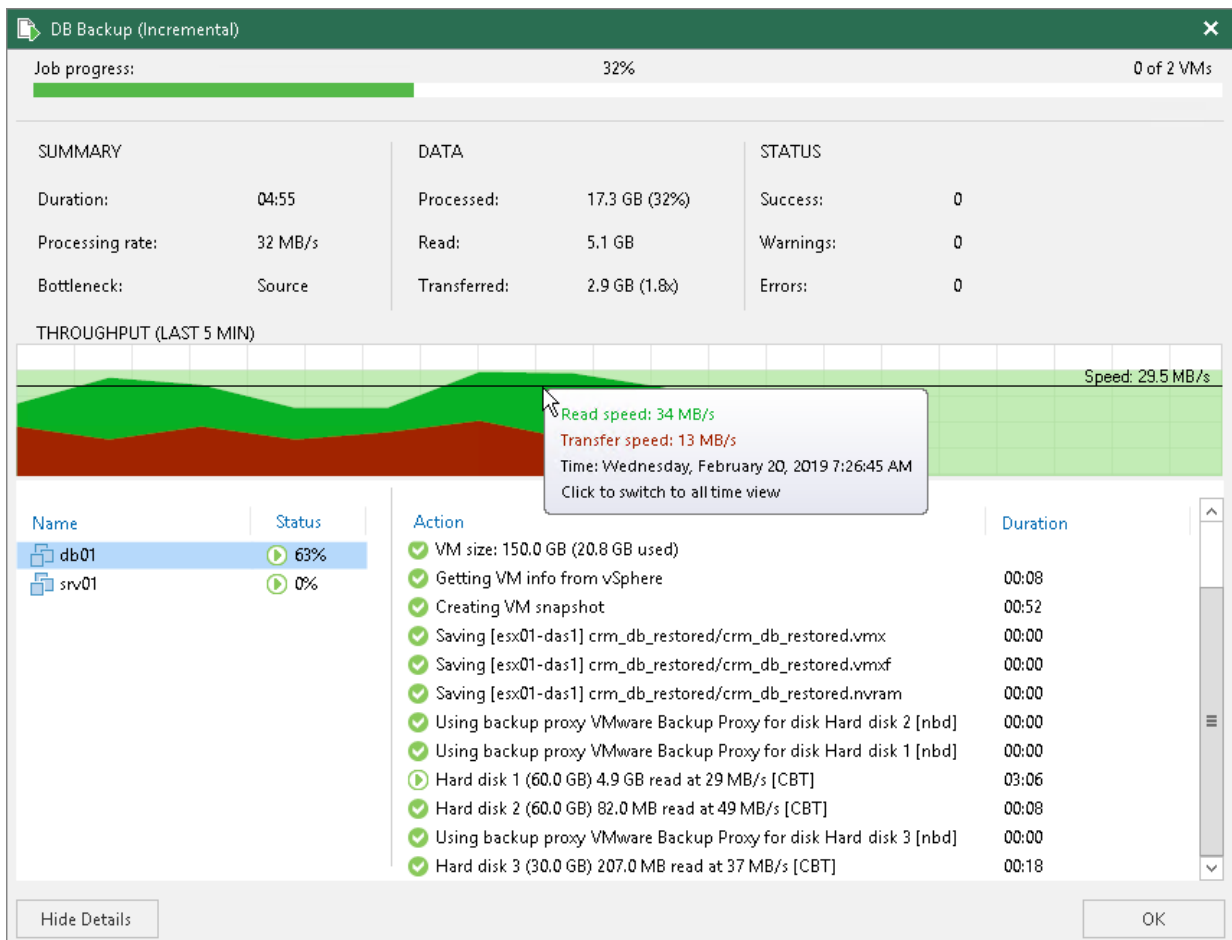
For more information on backup creation, see [Creating Backup Jobs](#) in the Veeam Backup & Replication User Guide.

Monitoring Job Performance in Real Time

When the job is running, you can view job statistics in real time. Statistics include job progress, duration, processing rate, performance bottlenecks, the amount of read and transferred data, and other details of the job performance.

To view the job statistics, do the following:

1. In the inventory pane of the **Home** view, select the **Jobs** node.
2. In the working area, right-click a running job and click **Statistics**.
3. In the opened window, select a VM to view its statistics.



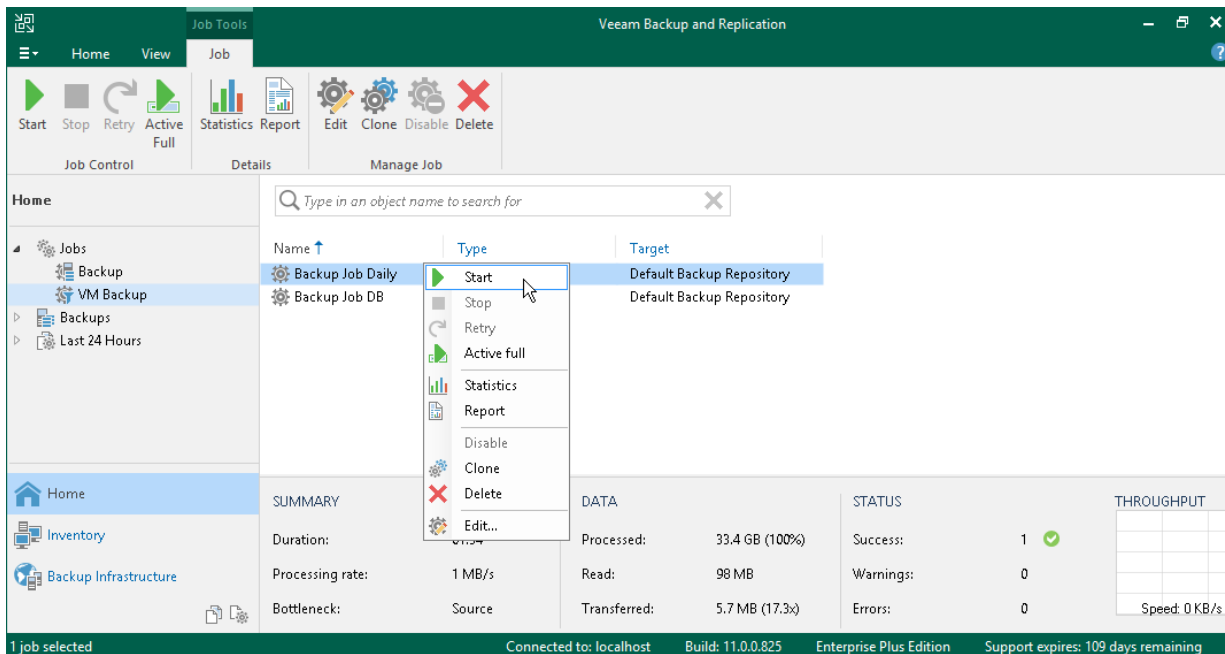
Note that the job must complete with the *Success* or *Warning* status. If the job completes with the *Failed* status, Veeam Backup & Replication does not create the backup file is not able to perform restore operations.

You can configure email notifications to get job results. For details, see [Configuring Global Email Notification Settings](#) in the Veeam Backup & Replication User Guide.

Start Backup Job Manually

If you do not schedule a backup job, you must start it manually. To start the job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Jobs** node.
3. In the working area, right-click the job and select **Start**. Wait for the job to complete. Note that the job must complete with the *Success* or *Warning* status.

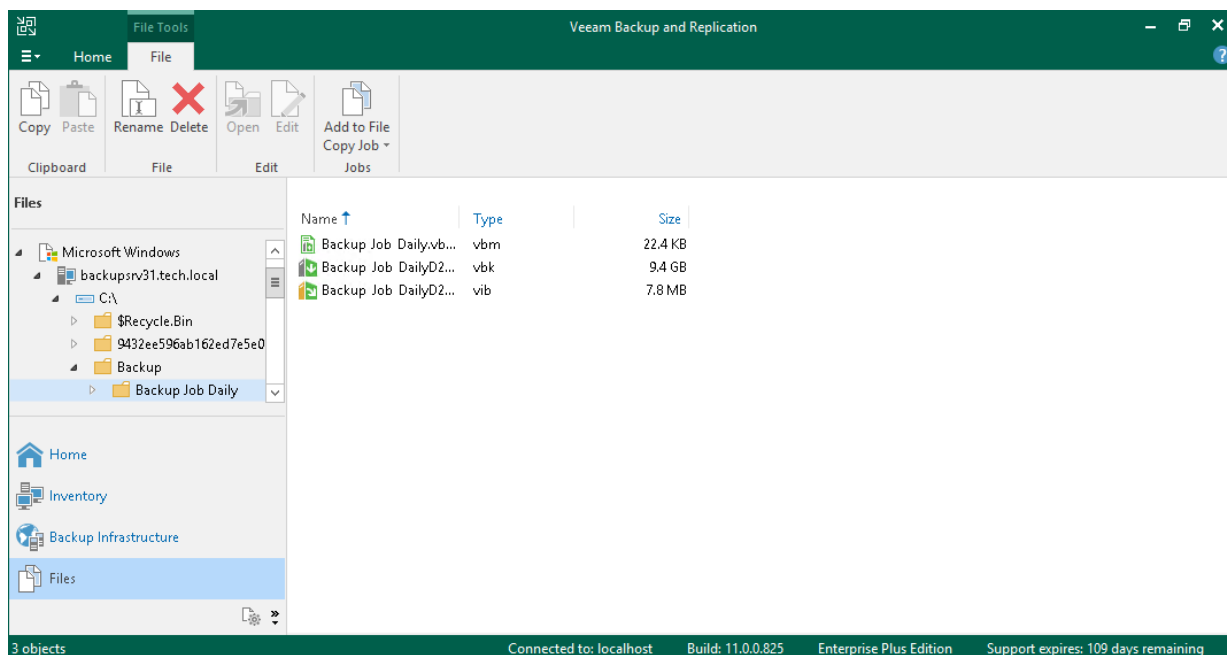


Locating Backup Files

When a backup job finishes, Veeam Backup & Replication saves backup files in the backup repository that you specified as a backup target.

To locate backup files using the Veeam Backup & Replication console, do the following:

1. Open the **Files** view.
2. In the inventory pane, expand the backup repository file tree and open the **Backup** folder.
3. In the **Backup** folder, find the subfolder with the backup job name and open it. It must contain a .VBK and .VBM files. If the job was run several times, the subfolder also contains .VIB or .VRB files.



Creating Application-Aware Backup Job

Application-aware processing allows you to create transactionally consistent backups. These backups allow you to further restore application items: emails for mail agents, tables for DB servers, accounts for domain controllers.

Veeam Backup & Replication can create transactionally consistent backups of VMs that run the following applications:

- Microsoft Exchange
- Active Directory
- SharePoint
- SQL Server
- Oracle Database

In this section, you will learn how to create the application-aware backup job for a Microsoft SQL Server.

IMPORTANT!

Application-aware processing is supported only for VSS-aware applications and applications listed above. If an application that you want to back up is not supported, you can use VMware Tools quiescence with pre-freeze and post-thaw scripts. For more information, see [VMware Tools Quiescence](#) and [Pre-Freeze and Post-Thaw Scripts](#) in the Veeam Backup & Replication User Guide.

Before You Begin

Make sure that the version of your Microsoft SQL Server is supported. For details, see the [System Requirements](#) section in the Veeam Backup & Replication User Guide.

Creating Application-Aware Backup

To create the application-aware backup job for the Microsoft SQL Server, do the following:

1. In the inventory pane of the **Home** view, right-click **Jobs** and select **Backup > Virtual Machine > VMware vSphere** to launch the **New Backup Job** wizard.
2. At the **Name** step of the wizard, specify a name and description for the backup job.
3. At the **Virtual Machines** step of the wizard, select the VM.
4. At the **Storage** step of the wizard, select a backup repository or keep the default settings.
5. At the **Guest Processing** step of the wizard, do the following:
 - Select the **Enable application-aware processing** check box.
 - Select the **Enable guest file system indexing** check box.

VM guest OS file indexing allows you to search for VM guest OS files inside VM backups and perform 1-click restore in Veeam Backup Enterprise Manager. For details, see [VM Guest OS File Indexing](#) in the Veeam Backup & Replication User Guide.

- In the **Guest OS credentials** section, specify credentials of a user account to connect to the VM guest OS. The user account must have Administrator permissions on the Microsoft SQL Server.
- Click the **Applications** button at the top of the window.

New Backup Job

Guest Processing
Choose guest OS processing options available for running VMs.

Left Sidebar:

- Name
- Virtual Machines
- Storage
- Secondary Target
- Guest Processing**
- Schedule
- Summary

Main Content Area:

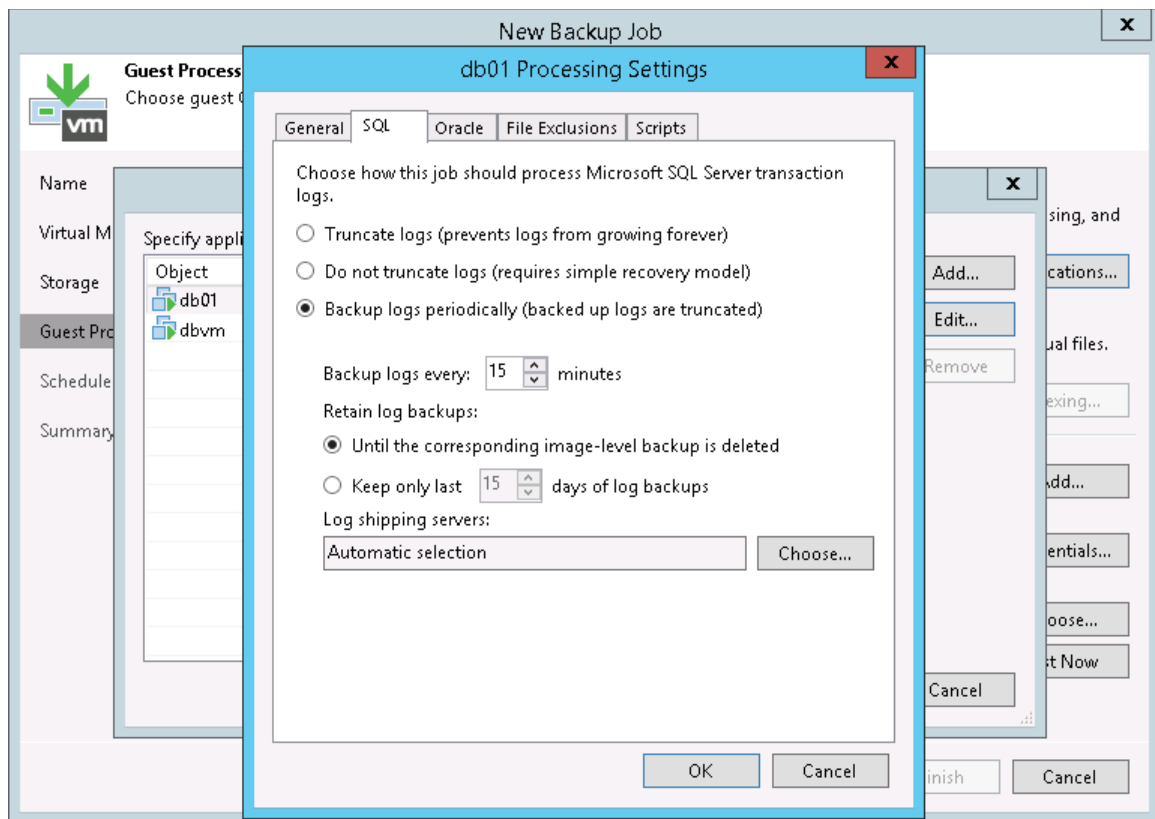
- ☒ **Enable application-aware processing**
Detects and prepares applications for consistent backup, performs transaction logs processing, and configures the OS to perform required application restore steps upon first boot.
Customize application handling options for individual VMs and applications [Applications...](#)
- ☒ **Enable guest file system indexing**
Creates catalog of guest files to enable browsing, searching and 1-click restores of individual files. Indexing is optional, and is not required to perform instant file level recoveries.
Customize advanced guest file system indexing options for individual VMs [Indexing...](#)
- Guest OS credentials**
 - TECH\Administrator (TECH\Administrator, last edited: 2 days ago) [Add...](#)
 - [Manage accounts](#)
 - Customize guest OS credentials for individual VMs and operating systems [Credentials...](#)
- Guest interaction proxy:**
 - Automatic selection [Choose...](#)
 - [Test Now](#)

Bottom Buttons: < Previous, Next >, Finish, Cancel

6. In the opened window, select the Microsoft SQL Server from the list and click **Edit**.
7. In the **Processing Settings** windows, do the following:
 - In the **Transaction logs** section of the **General** tab, check that the **Process transaction logs with this job** option is selected.

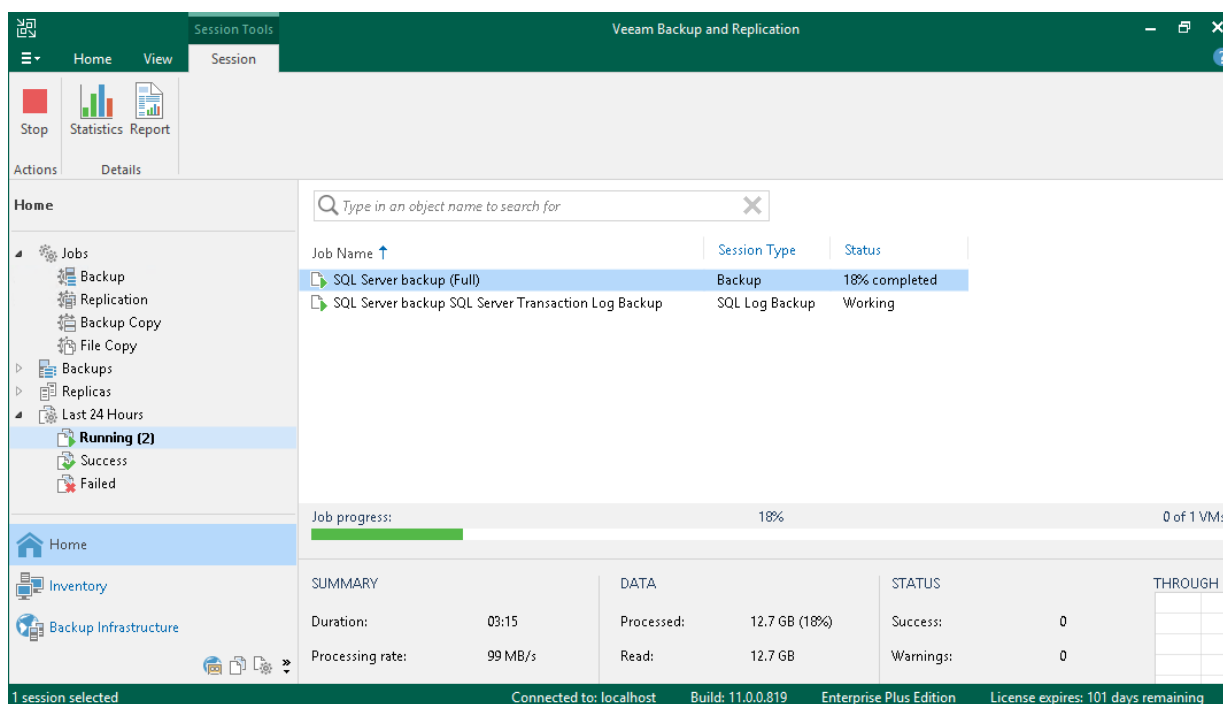
- On the **SQL** tab, select **Backup log periodically** option.

Veeam Backup & Replication will create an auxiliary job that runs continuously and ships database transaction logs. Transaction logs are shipped to the backup repository and saved in .VLB files next to other backup job files. Thus, you have a chain of restore points and a set of transaction logs that cover intervals between these restore points.



8. At the **Schedule** step of the wizard, define scheduling settings for the job.
9. At the **Summary** step of the wizard, select the **Run the job when I click Finish** check box and click the **Finish** button.

10. In the inventory pane of the **Home** view, expand the **Last 24 Hours** node to see the created jobs. You must see two jobs: one that processes the Microsoft SQL Server and the other one that ships transaction logs.



Reference

For more information on application-aware backups, see the [Application-Aware Processing](#) section in the Veeam Backup & Replication User Guide.

Data Recovery

Veeam Backup & Replication allows you to restore the following instances:

- [Entire VM](#)
- [Guest OS files](#)
- [VM disks](#)
- [VM files](#)
- [Application items](#)

Restoring Entire VM

If a VM fails, you can restore it from a backup file. You can restore a single VM or multiple VMs to the original or new location.

In this section, you will learn how to restore a VM to the original location. For more information on how to restore the VM to another location, see [Restoring Entire VM](#) in the Veeam Backup & Replication User Guide.

Before You Begin

Before you restore a VM from a backup, consider the following:

- You can restore the VM from a backup that has at least one successfully created restore point.

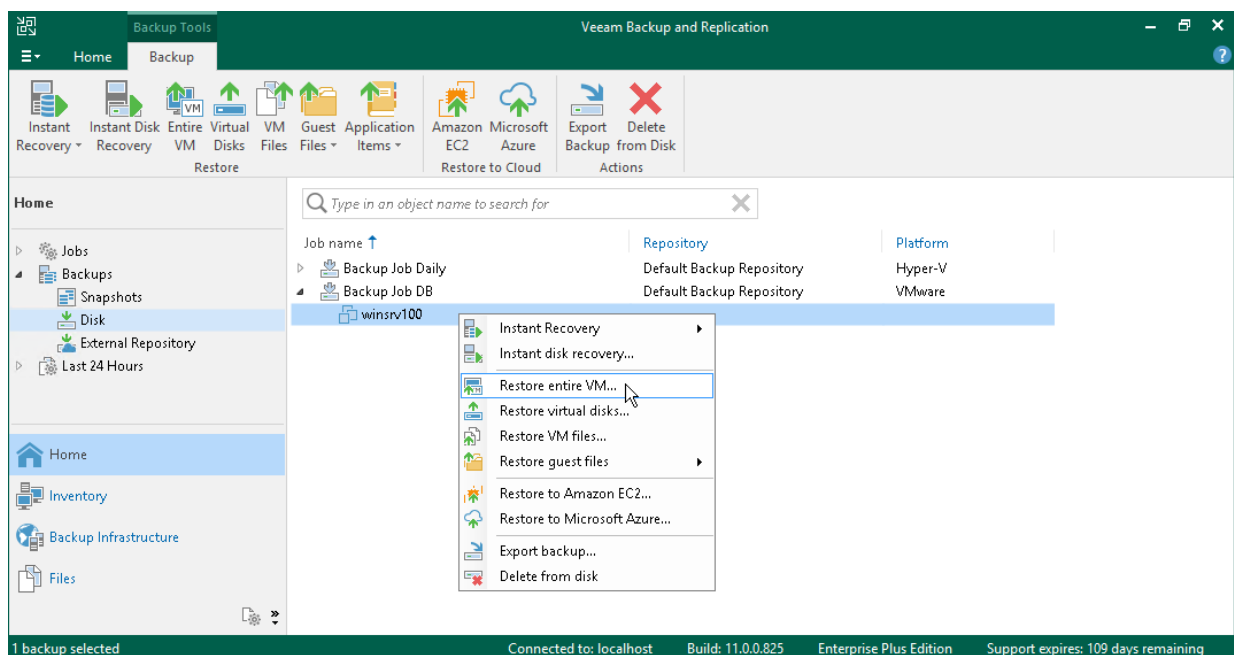
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available for the VM.

- When you restore the VM to its original location, and the original VM is still running, Veeam Backup & Replication powers off the original VM and restores only those disks that are included in the backup. All other disks remain unchanged.

Restoring Entire VM

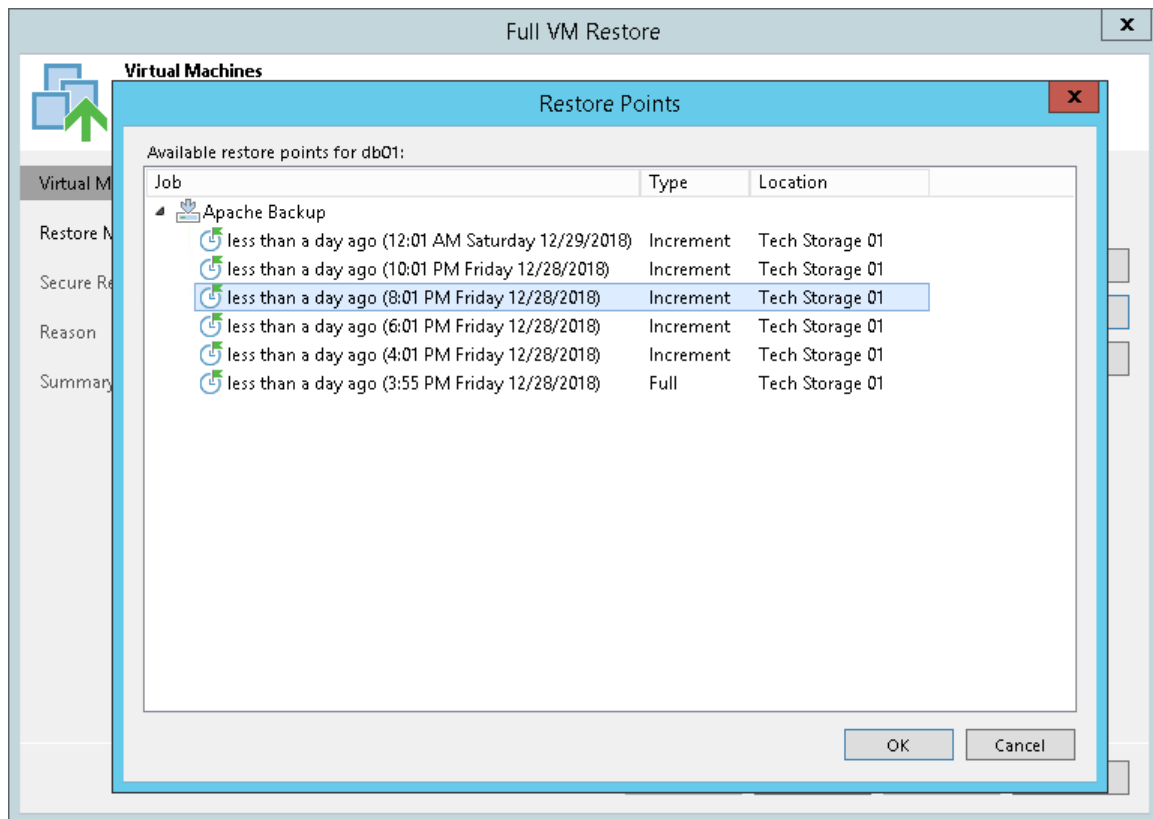
To restore an entire VM to its original location, do the following.

- Open the **Home** view.
- In the inventory pane, select the **Backups > Disk** node. Expand the backup job in the working area, right-click a VM in a backup job and select **Restore entire VM** to launch the **Full VM Restore** wizard.



3. At the **Virtual Machines** step of the wizard, select the VM from the list, click the **Point** button and choose a restore point.

If you select an incremental restore point, Veeam Backup & Replication automatically restores data blocks from the full backup file and the chain of incremental backup files.

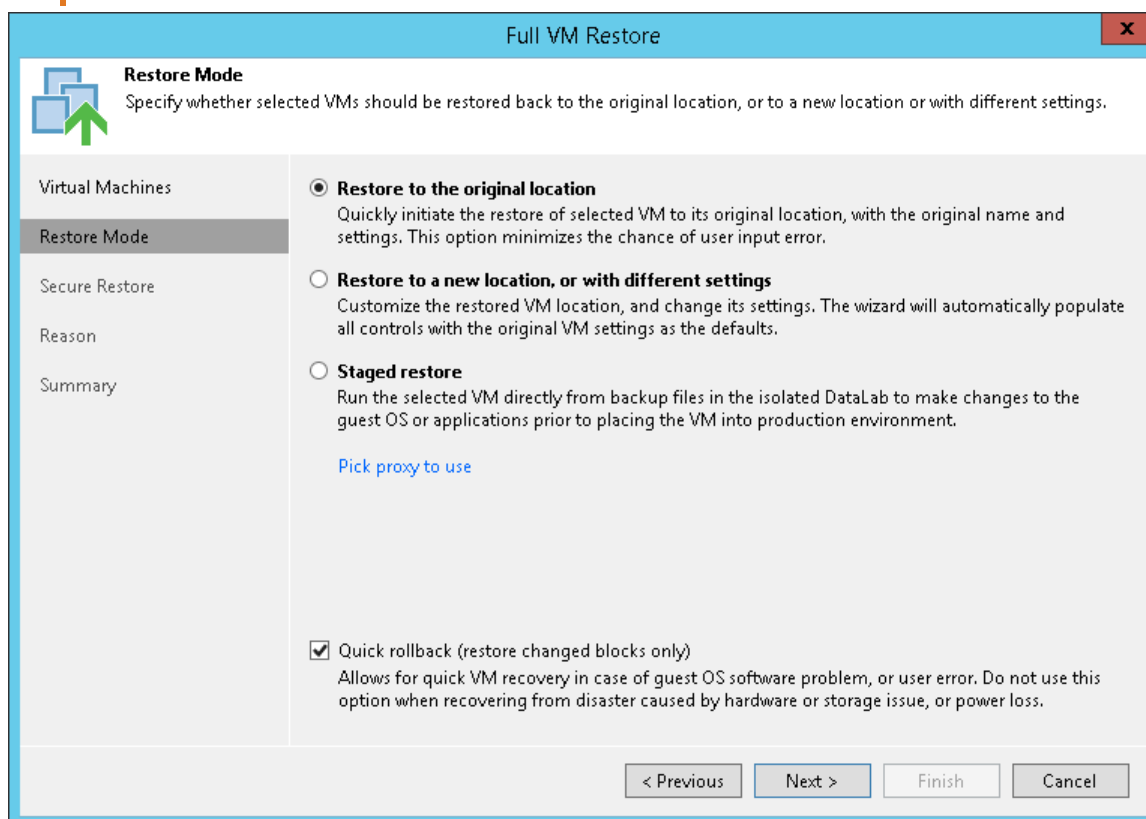


4. At the **Restore mode** step of the wizard, do the following:
 - Select the **Restore to the original location** option.
 - Select the **Quick rollback** check box.

Veeam Backup & Replication will get data blocks that are necessary to revert the VM to an earlier point in time and will restore only these data blocks from the backup. Quick rollback significantly reduces the restore time.

IMPORTANT!

Do not enable the **Quick rollback** option if the problem occurred at the VM hardware level, storage level or due to a power loss.



The screenshot shows the 'Full VM Restore' wizard window. The title bar is 'Full VM Restore' with a close button. The main content area is titled 'Restore Mode' with a subtitle '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 tree view containing 'Virtual Machines', 'Restore Mode' (selected), 'Secure Restore', 'Reason', and 'Summary'. The main area contains three radio button options: 'Restore to the original location' (selected), 'Restore to a new location, or with different settings', and 'Staged restore'. Below these is a checkbox for 'Quick rollback (restore changed blocks only)' which is checked. At the bottom are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Full VM Restore

Restore Mode
Specify whether selected VMs should be restored back to the original location, or to a new location or with different settings.

Virtual Machines

Restore Mode

Secure Restore

Reason

Summary

☒ **Restore to the original location**
Quickly initiate the restore of selected VM to its original location, with the original name and settings. This option minimizes the chance of user input error.

☐ **Restore to a new location, or with different settings**
Customize the restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the defaults.

☐ **Staged restore**
Run the selected VM directly from backup files in the isolated DataLab to make changes to the guest OS or applications prior to placing the VM into production environment.

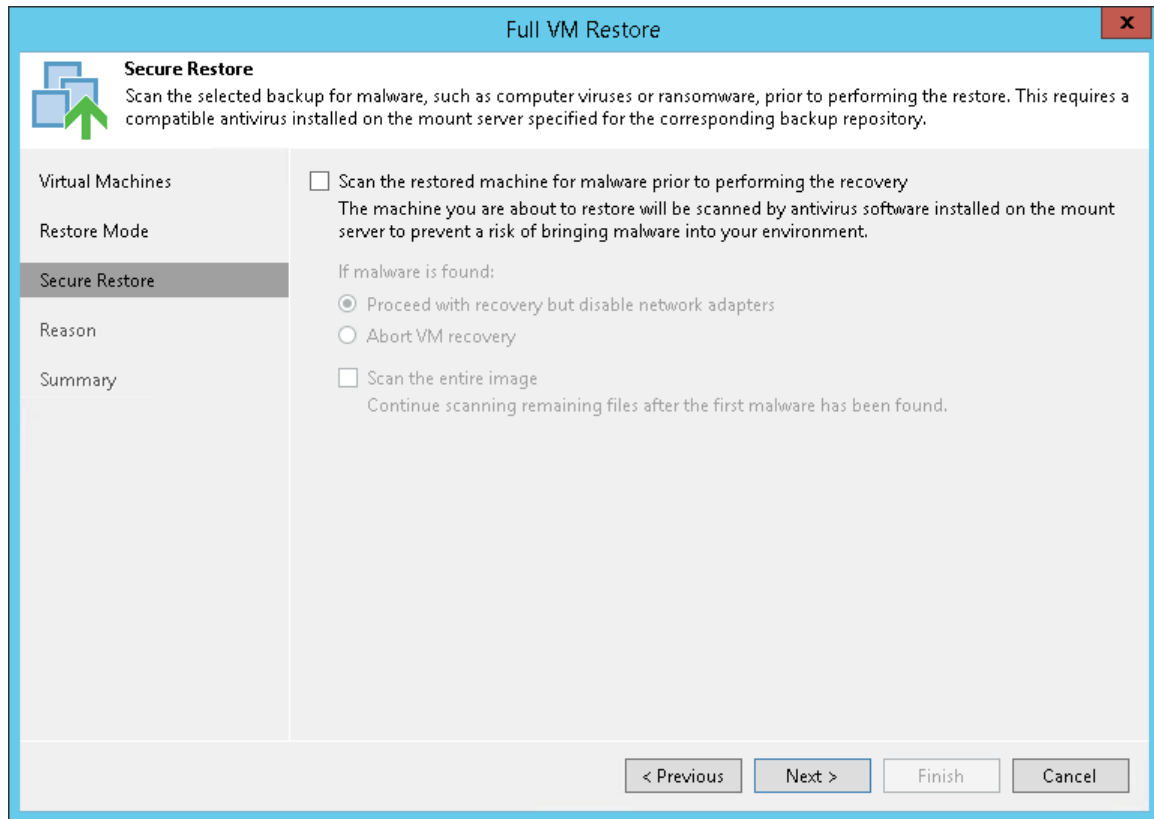
[Pick proxy to use](#)

☒ **Quick rollback (restore changed blocks only)**
Allows for quick VM recovery in case of guest OS software problem, or user error. Do not use this option when recovering from disaster caused by hardware or storage issue, or power loss.

< Previous Next > Finish Cancel

- At the **Secure Restore** step of the wizard, enable scanning of the machine or leave the default settings.

If secure restore is enabled, Veeam Backup & Replication uses antivirus software to scan machine data before restoring the machine to the production environment. For details, see [Secure Restore](#) in the Veeam Backup & Replication User Guide.



- At the **Reason** step of the wizard, specify the reason for restoring the VM.
- At the **Summary** step of the wizard, select the **Power on VM after restoring** check box and click **Finish**.

Restoring VM Files

Veeam Backup & Replication can help you restore VM files: VMX, VMXF, NVRAM and VMDK including flat files. For example, your VM configuration file is missing and you need to restore it. Instead of restoring the entire VM image, you can restore a single VM file.

You can restore VM files to the latest state or any valid point in time. You can also restore them to the original or new location.

Before You Begin

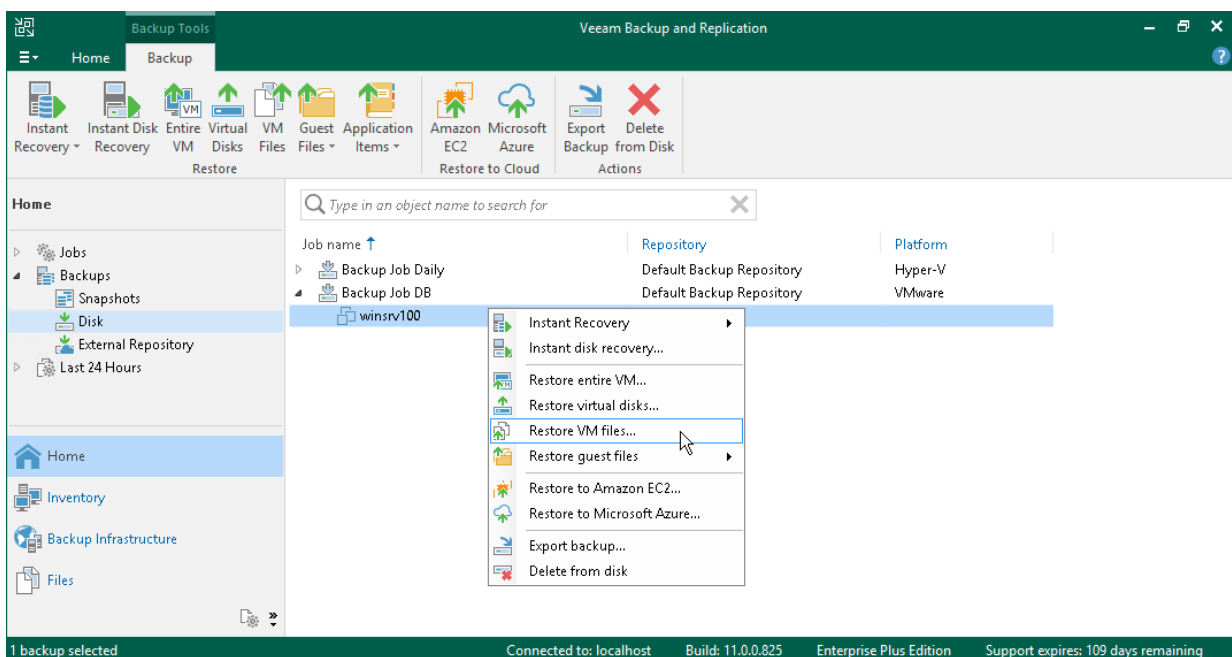
Before you restore VM files from a backup, consider the following:

- You can restore VM files a backup that has at least one successfully created restore point.
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available for the VM.
- The server on which you plan to save the restored VM files must be added to the backup infrastructure.

Restoring VM Files

To recover VM files, do the following.

- Open the **Home** view.
- In the inventory pane, select the **Backups > Disk** node. Expand a backup job in the working area, right-click a VM and choose **Restore VM files** to launch the **Virtual Machines File Restore** wizard.



3. At the **Restore Point** step of the wizard, select a restore point.

If you select an incremental restore point, Veeam Backup & Replication automatically restores data blocks from the full backup file and the chain of incremental backup files.

The screenshot shows the 'Virtual Machine Files Restore' wizard window. The 'Restore Point' step is active, indicated by a green arrow icon and the 'Restore Point' tab being selected in the left sidebar. The main area displays the VM name 'abc_srv01' and size '116.6 GB'. Below this, a table lists 'Available restore points' with columns for 'Created', 'Type', and 'Location'. The table contains five rows of restore points, all from the 'Default Backup Repository'. The fourth row, 'less than a day ago (6:00 PM Tuesday ...)', is highlighted. At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Created	Type	Location
less than a day ago (12:00 AM Wednes...	Increment	Default Backup Repository
less than a day ago (10:00 PM Tuesday...	Increment	Default Backup Repository
less than a day ago (8:00 PM Tuesday ...	Increment	Default Backup Repository
less than a day ago (6:00 PM Tuesday ...	Increment	Default Backup Repository
less than a day ago (4:09 PM Tuesday ...	Full	Default Backup Repository

4. At the **Destination** step of the wizard, do the following:
 - a. In the **Server** list, select the server to which you want to restore VM files.
 - b. In the **Path to folder** field, specify the path to the folder where you want to restore files.

- c. In the **VM files to restore** section, select the required files.

The screenshot shows the 'Virtual Machine Files Restore' wizard in the 'Destination' step. The left sidebar contains a tree view with 'Destination' selected. The main area has a title bar 'Virtual Machine Files Restore' and a subtitle 'Choose server and folder where VM files should be restored, and pick files to restore.' Below this, there are fields for 'Server:' (esx01.tech.local) and 'Path to folder:' ([esx01-das3]). A table titled 'VM files to restore:' lists four files with checkboxes. The first three files are selected. To the right of the table are 'Select All' and 'Clear All' buttons. At the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

Name	Size
<input checked="" type="checkbox"/> abc_srv01.vmx	2.7 KB
<input checked="" type="checkbox"/> abc_srv01.nvram	8.5 KB
<input checked="" type="checkbox"/> abc_srv01-000002.vmdk	618.0 B
<input type="checkbox"/> abc_srv01-000002-flat.vmdk	50.0 GB

5. At the **Reason** step of the wizard, specify the reason for restoring files.
6. At the **Summary** step of the wizard, click **Finish** to restore the VM files.

Reference

For more information on restoring VM files, see [VM Files Restore](#) in the Veeam Backup & Replication User Guide.

Restoring VM Virtual Disks

Veeam Backup & Replication allows you to recover individual virtual disks of a VM. Recovered virtual disks can be attached to the original VM or to any other VM. This recovery option can be helpful if a VM virtual disk becomes corrupted.

You can restore VM virtual disk to the latest state or any valid restore point. You can preserve the format of a recovered virtual disk or convert it to the thin or thick provisioned disk format.

Before You Begin

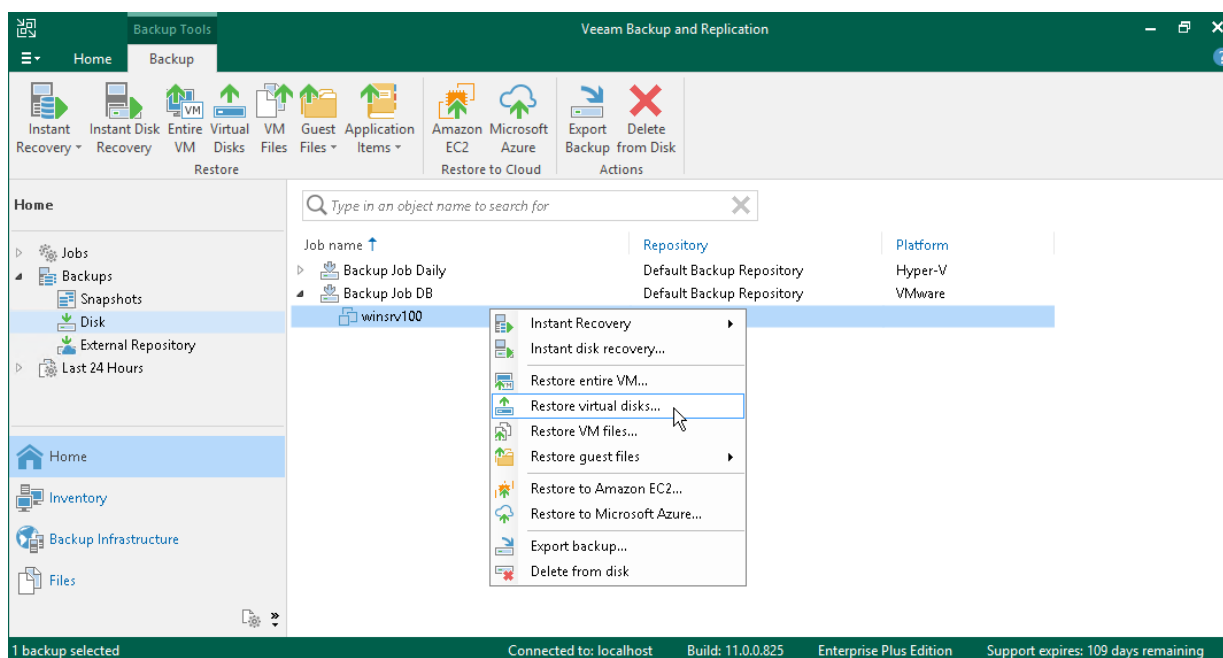
Before you restore VM virtual disks from a backup, consider the following:

- You can restore virtual disks from a backup that has at least one successfully created restore point.
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available for the VM.
- During the virtual disk restore, Veeam Backup & Replication turns off the target VM (the VM to which you plan to attach the restored virtual disk) to reconfigure its settings and connect restored disks. It is recommended that you stop all activities on the target VM for the restore period.

Restoring Virtual Disks

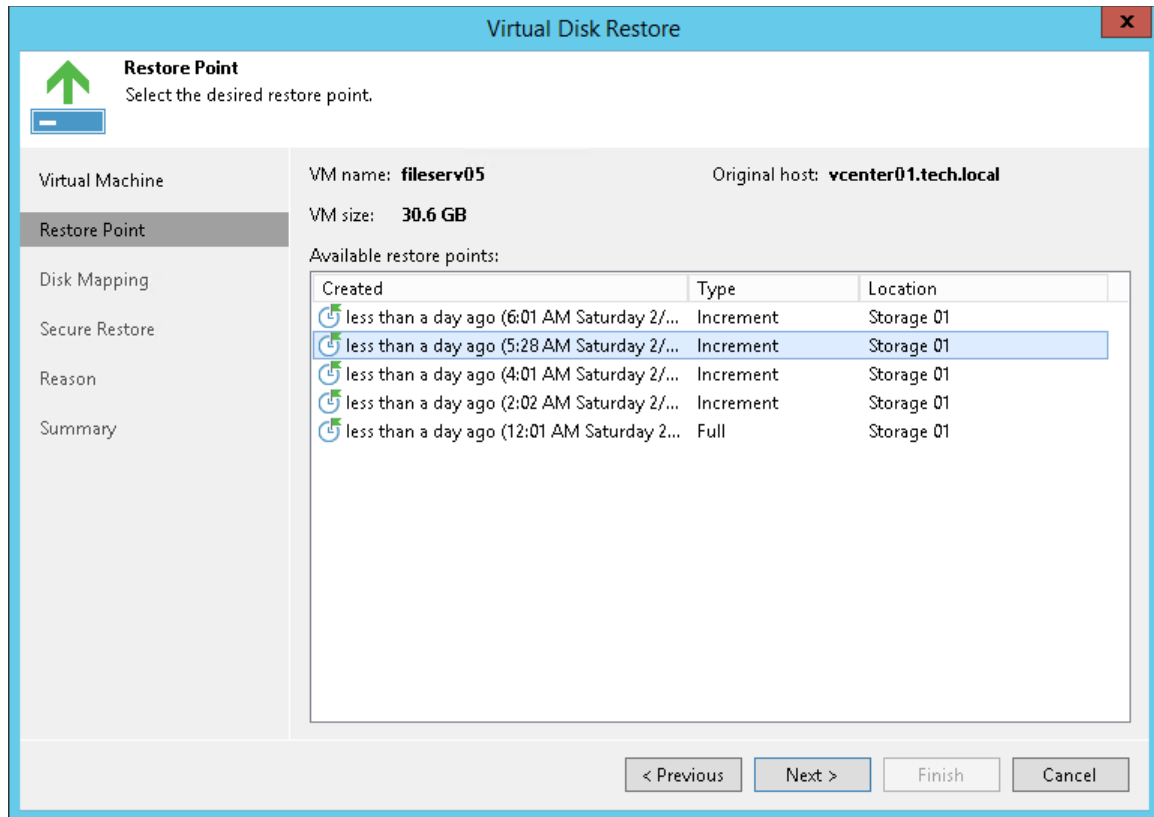
To restore a VM virtual disk and attach it to another VM as a new drive, do the following:

1. Open the **Home** view.
2. In the inventory pane, select the **Backups > Disk** node. Expand the backup job in the working area, right-click a VM and select **Restore virtual disks** to launch the **Virtual Disk Restore** wizard.



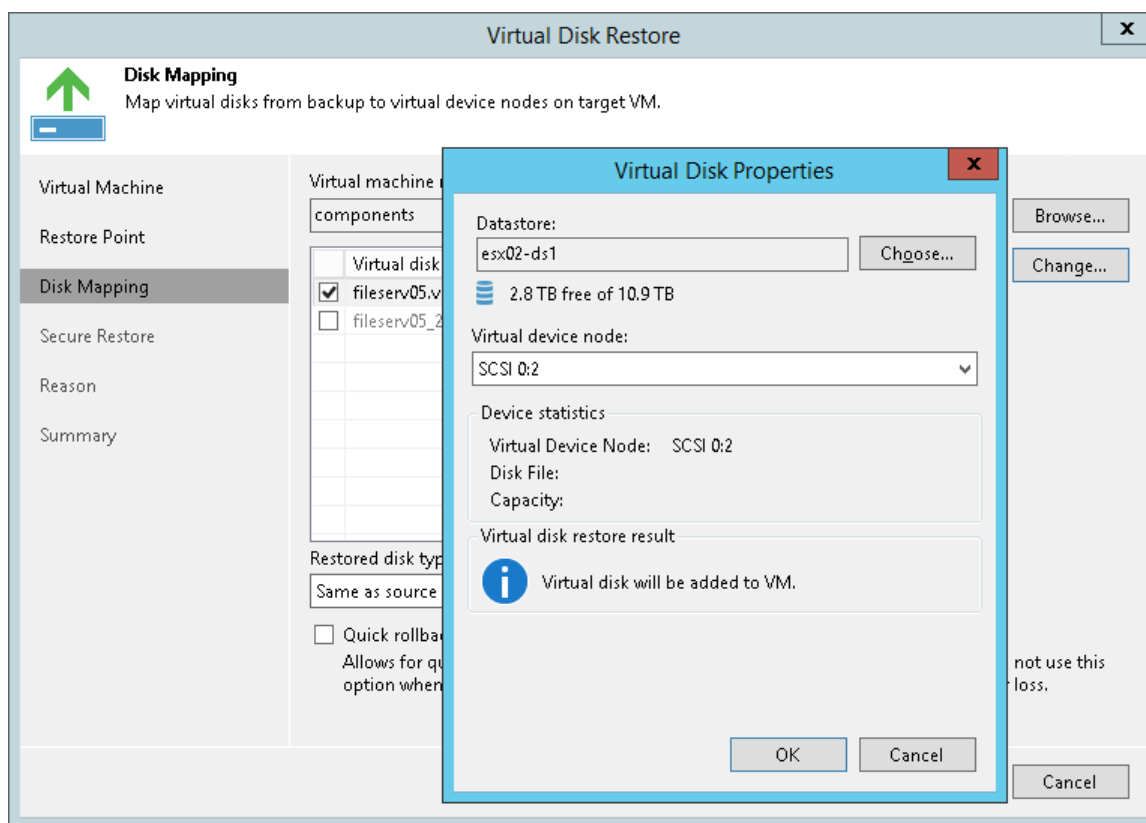
3. At the **Restore Point** step of the wizard, select a restore point.

If you select an incremental restore point, Veeam Backup & Replication automatically restores data blocks from the full backup file and the chain of incremental backup files.



4. At the **Disk Mapping** step of the wizard, do the following:
 - a. Click **Browse** and select the VM to which the restored hard disk must be attached.
 - b. Select the virtual hard disk that you want to restore.
 - c. To change the disk format, select the required option from the **Restored disk type** list: same as the original disk, thin or thick (lazy or eager zeroed).

- d. Select the VM disk in the list and click **Change**. From the **Virtual device node** list, select a node that is not occupied yet. Click **OK**.



5. At the **Secure Restore** step of the wizard, enable scanning of the machine or leave the default setting.
- If secure restore is enabled, Veeam Backup & Replication uses antivirus software to scan machine data before restoring the machine to the production environment. For details, see [Secure Restore](#) in the Veeam Backup & Replication User Guide.
6. At the **Reason** step of the wizard, specify the reason for restoring.
7. At the last step of the wizard, select the **Power on VM after restoring** check box and click **Finish**.

Reference

For more information on restoring virtual disks, see [Virtual Disks Restore](#) in the Veeam Backup & Replication User Guide.

Restoring Guest OS Files

Veeam Backup & Replication allows you to recover individual guest OS files and folders. You can restore files and folders directly from a backup. This makes the restore process fast and does not require additional storage resources.

Veeam Backup & Replication supports recovering files for the following file systems:

- [Microsoft Windows file systems \(FAT, NTFS and ReFS\)](#)
- [File systems of Linux-based OSes](#)
- Other file systems

In this guide, we omit restoring files from other file systems. This is an advanced scenario that requires additional actions. For details, see the [Restore from Other File Systems](#) section in the Veeam Backup & Replication User Guide.

Restoring VM Guest OS Files (FAT, NTFS, ReFS)

You can restore individual files from the backup of a Microsoft Windows VM. For this purpose, Veeam Backup & Replication provides the **File-Level Restore** wizard.

Before You Begin

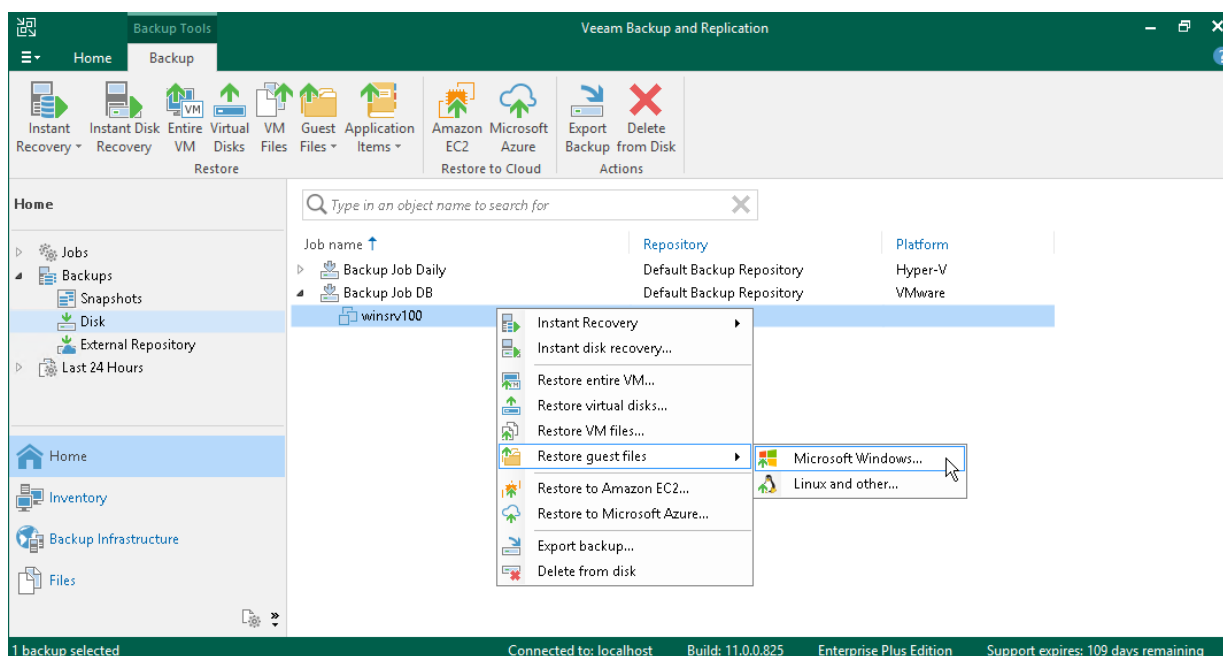
Consider the following:

- You can restore guest OS files from a backup that has at least one successfully created restore point.
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available for the VM.
- You cannot restore files from a backup created in the reverse incremental mode if the backup job is being performed. If the backup is created in the incremental backup mode and the backup job is being performed, you can restore files from any available restore point.
- You cannot restore VM guest OS files from a running replica or if the replication job with the necessary VM is being performed.

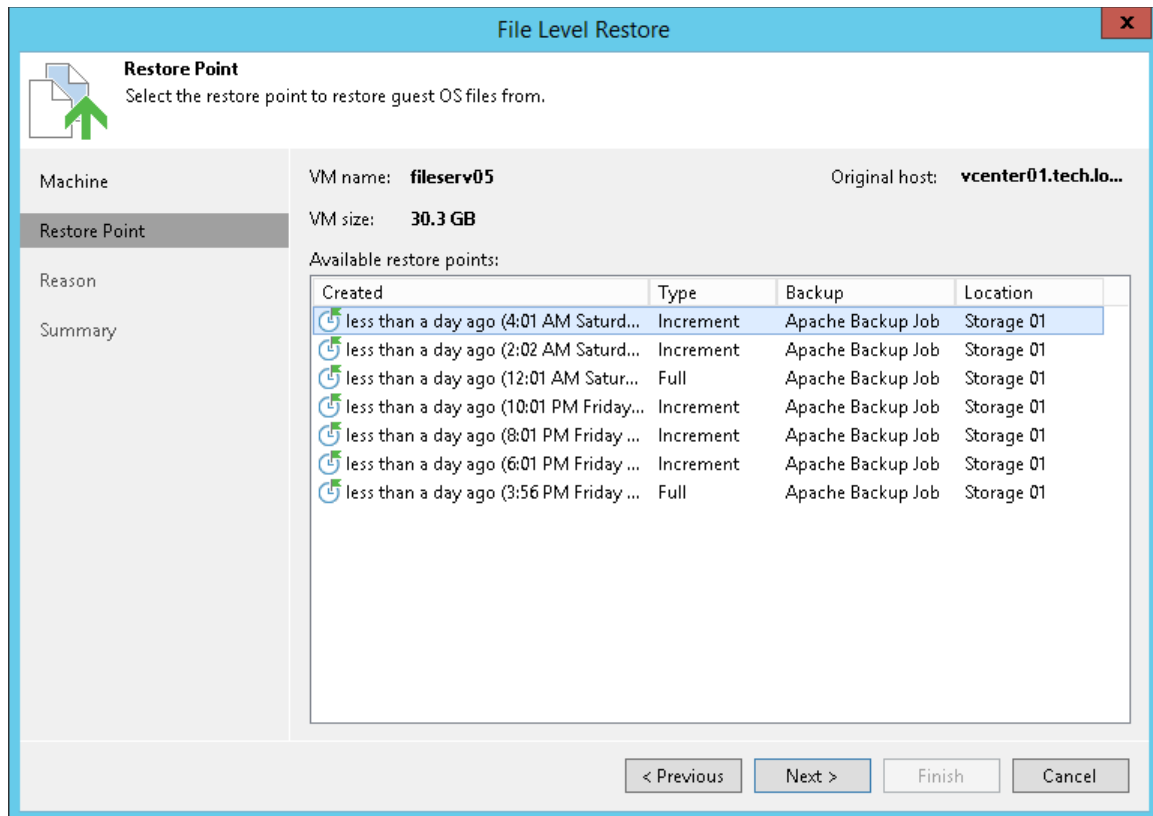
Restoring Guest OS Files

To restore guest OS files from a Microsoft Windows VM:

- Open the **Home** view.
- In the inventory pane, select the **Backups > Disk** node. Expand a backup job in the working area, right-click a VM and choose **Restore guest files > Microsoft Windows** to launch the **File Level Restore** wizard.



- At the **Restore Point** step of the wizard, select the necessary restore point.



- At the **Reason** step of the wizard, specify the reason for restoring VM guest OS files.
- At the last step of the wizard, click **Finish**.
- Veeam Backup & Replication will display the **Backup Browser** window with the file system tree of the VM. Right-click the necessary file or folder, select **Copy To**.
- In the opened window, specify the location to which you want to restore files or folders. This location is a network shared folder or folder on the backup server.

Reference

For more information on restoring guest OS files, see [Restore from FAT, NTFS or ReFS](#) in the Veeam Backup & Replication User Guide.

Restoring VM Guest OS Files (Linux, Unix, etc)

You can restore individual files and folders from file systems of Linux-based OSes. For this purpose, Veeam Backup & Replication provides the multi-OS **File-Level Restore** wizard. The multi-OS restore wizard allows you to restore guest OS files for such OSes as Linux, Unix, BSD, macOS and others.

To restore files from VM guest OS, Veeam Backup & Replication uses a helper appliance. The helper appliance is a helper VM running a stripped-down Linux kernel that has a minimal set of components. When you perform file-level restore, Veeam Backup & Replication automatically starts the appliance and mounts VM disks to the helper appliance as virtual hard drives. Virtual disks are mounted directly from backup files, without prior extraction of the backup content. This makes the restore process much faster.

Before You Begin

Consider the following:

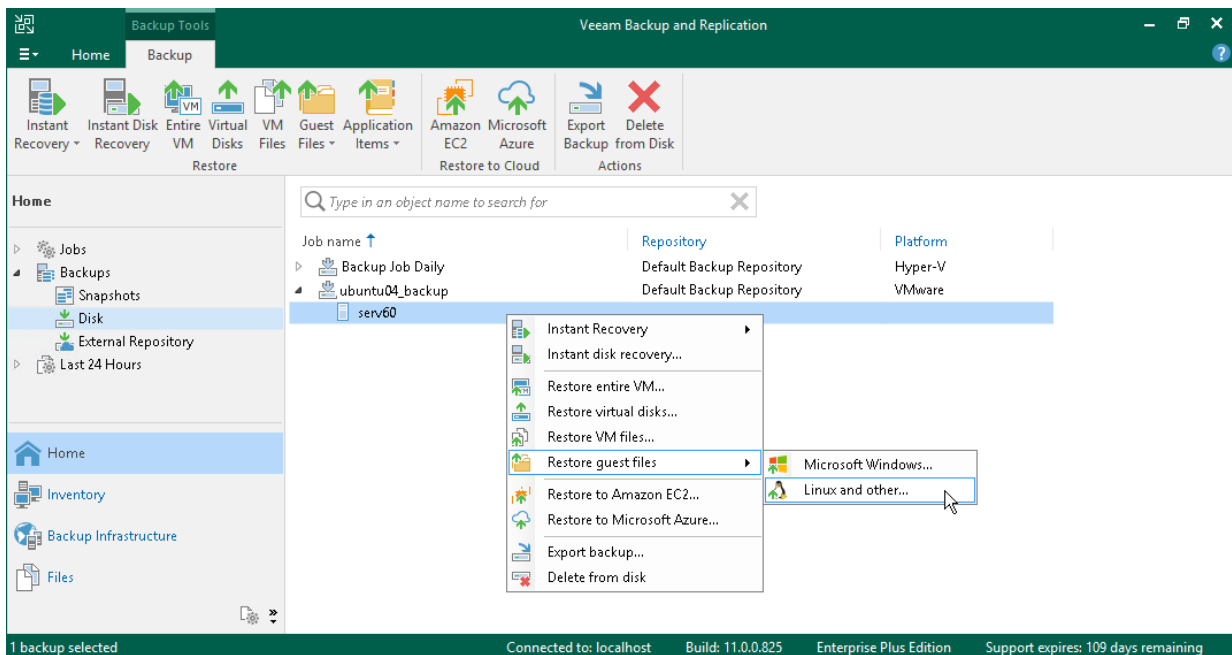
- Check the supported file systems. For details, see the [File-Level Restore](#) section in the Veeam Backup & Replication User Guide.
- You can restore guest OS files from a backup that has at least one successfully created restore point.
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available for the VM.
- You cannot restore files from reverse incremental backups.
- You cannot restore files from a VM being currently backed up or replicated.

Restoring Guest OS Files

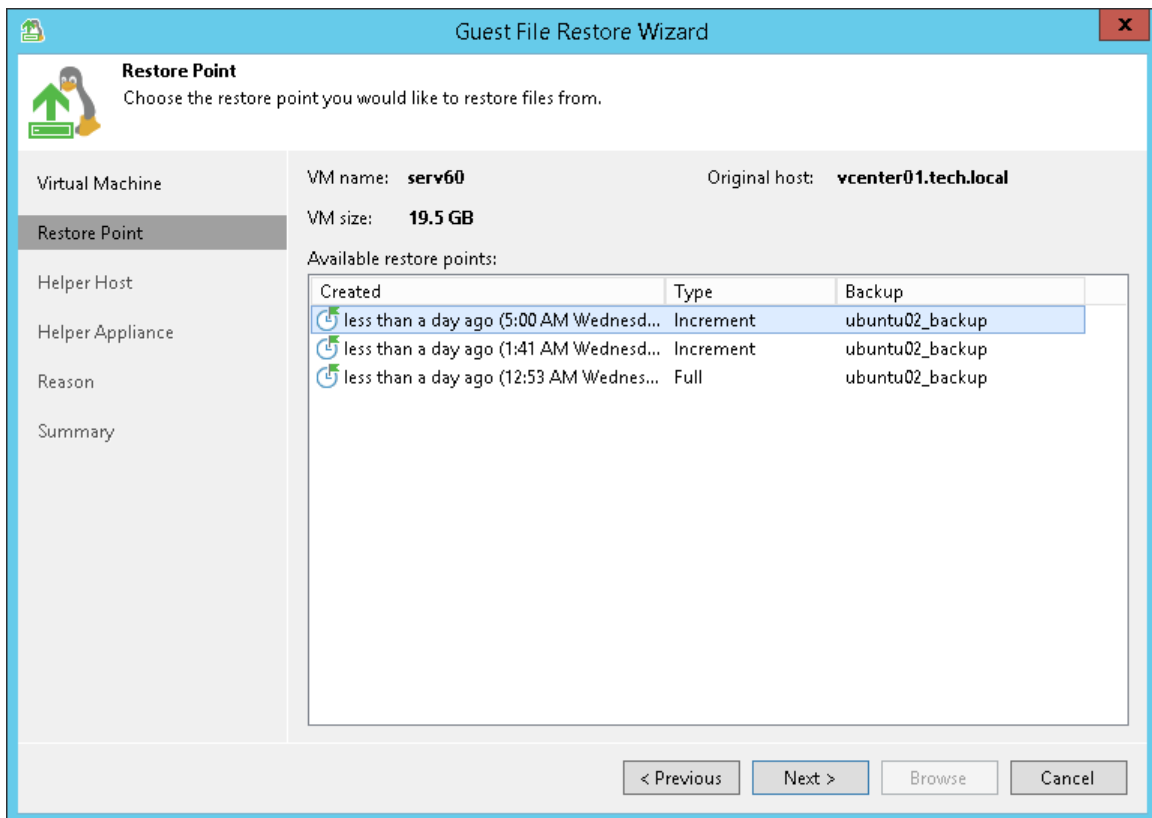
To restore guest OS files from a Linux-based VM, do the following.

1. Open the **Home** view.

2. In the inventory pane, select the **Backups > Disk** node. Expand a backup job in the working area, right-click a VM and select **Restore guest files > Linux and other** to launch the **Guest File Restore** wizard.



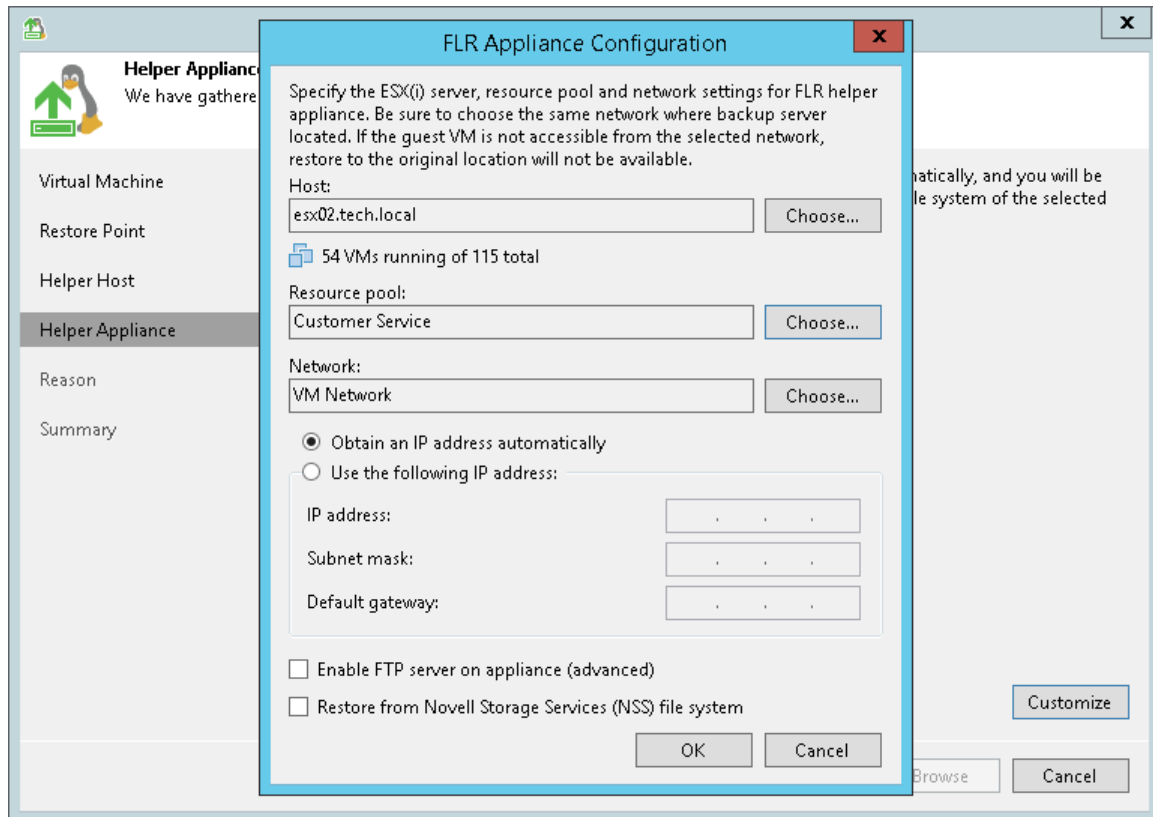
3. At the **Restore Point** step of the wizard, select a restore point.



4. At the **Helper Host** leave the default settings.

- At the **Helper Appliance** step of the wizard, click **Customize** to specify settings for the helper appliance. Select the ESXi host, resource pool and network on which the helper appliance will run.

If you are restoring files from the NSS file system, select the **Restore from Novell Storage Services (NSS) file system** check box.



- At the **Reason** step of the wizard, specify the reason for restoring.
- At the last step of the wizard, click **Finish**. Note that the helper appliance can boot about 20 seconds.
- Veeam Backup & Replication will display the **Backup Browser** with the file system tree of the VM. Right-click a file or folder, select **Copy To**.
- In the **Select Destination** window, do the following:
 - In the **Server** field, select the server to which you want to restore files.
 - In the **Path to folder** field, specify a destination folder.
- Click **Restore**.

Reference

For more information on restoring guest OS files, see [Restore from Linux, Unix and Other File Systems](#) in the Veeam Backup & Replication User Guide.

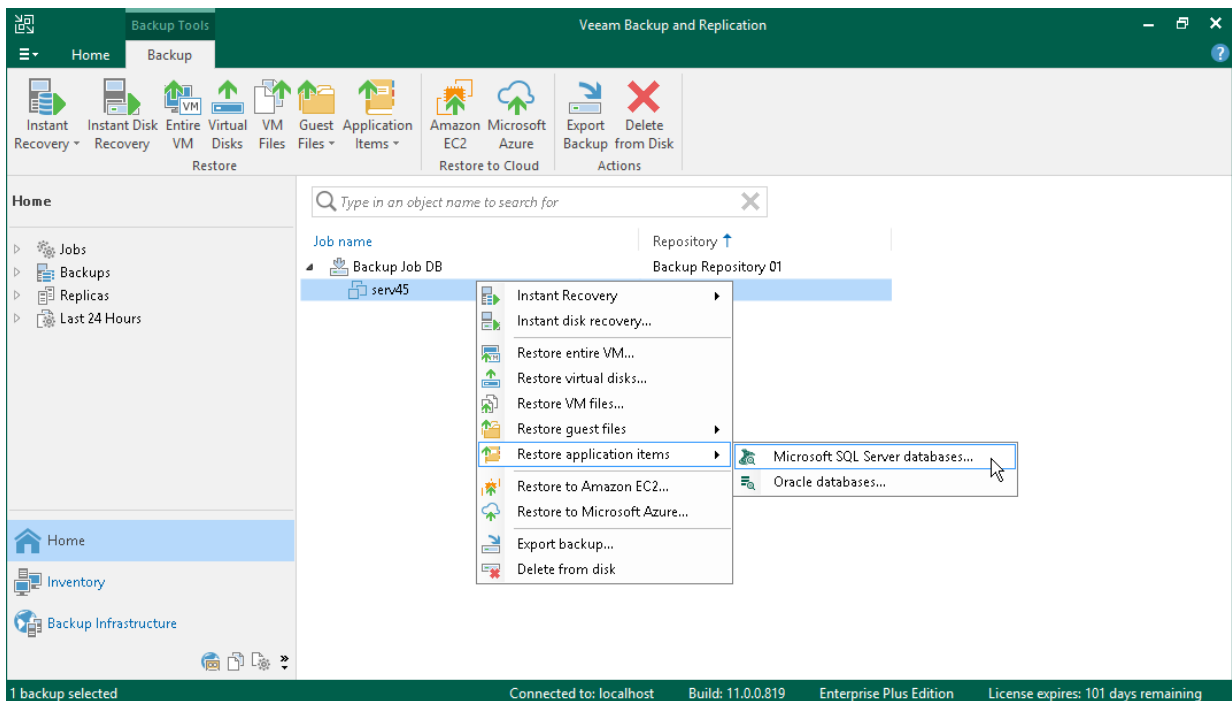
Restoring Application Items

If you have an application-aware backup, you can restore application items for Microsoft SQL Server, Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, and Oracle Database. To restore application items, Veeam Backup & Replication uses special built-in tools – Veeam Explorers.

Veeam Explorers mount the file system of the backed up VM, detect available applications and display their content in the convenient interface. You can then browse for necessary application items and restore them to the original or new location. For details, see [Veeam Backup Explorers User Guide](#).

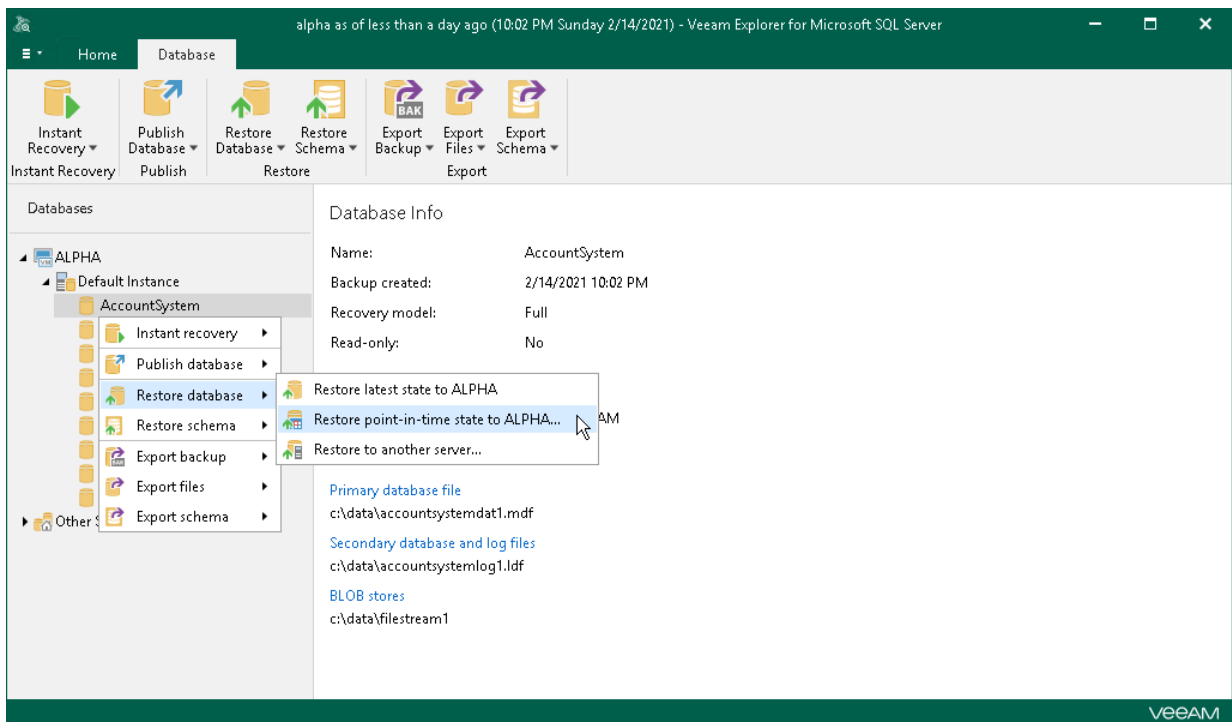
In this section, you will learn how to restore application items for the Microsoft SQL Server. For this purpose, you will use the backup created in the [Creating Application-Aware Backup Job](#) section.

1. In the inventory pane of the **Home** view, click the **Backups** node.
2. In the working area, expand the backup job that processes the VM with Microsoft SQL Server. Right-click the VM and select **Restore application items > Microsoft SQL Server databases** to open the **Microsoft SQL Server Database Restore** wizard.

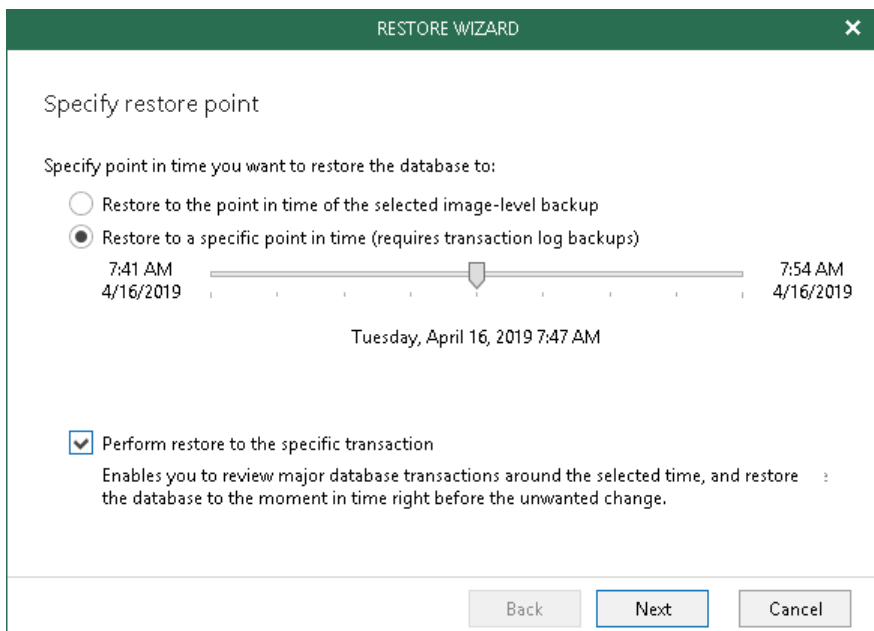


3. At the **Restore Point** step of the wizard, select the required restore point.
4. At the **Reason** step of the wizard, specify the reason for restoring.
5. At the last step of the wizard, click **Finish** to start the recovery process.
6. Veeam Backup & Replication will display the **Veeam Explorer for Microsoft SQL Server** window with available databases.

7. In the **Databases** pane of the window, right-click a database and select **Restore point-in-time state to <Microsoft SQL Server\Instance Name>**.



8. At the **Specify restore point** step of the wizard:
 - a. Select the **Restore to a specific point in time option**.
 - b. Use the slider to define the exact point in time to which you want to restore the database.
 - c. Select the **Perform restore to the specific transaction** check box and click **Next**.



9. At the **Fine-tune the restore point** step of the wizard, select the transaction to which you want to restore the database and click **Restore**.

Veeam Backup & Replication will start restoring the database to the selected transaction. When the restore process is complete, Veeam Explorer for Microsoft SQL Server will display a popup message with the results of the restore operation.

Reference

For more information on restoring application items, see [Application Items Restore](#) in the Veeam Backup & Replication User Guide.

Backup Copy

Backup copy allows you to create several instances of the same backup data in different locations. This is the mechanism that Veeam Backup & Replication provides to help you follow the 3-2-1 rule:

- 3: You must have at least three copies of your data: the original production data and two backups.
- 2: You must use at least two different types of media to store copies of your data, for example, local disk and cloud.
- 1: You must keep at least one backup offsite (for example, in the cloud or in a remote site).

In Veeam Backup & Replication, backup copy is a job-driven process. When the backup copying job starts, Veeam Backup & Replication accesses backup files on the source backup repository, retrieves data blocks for a specific machine from the backup file, copies them to the target backup repository, and composes copied blocks into a backup file on the target backup repository. This backup file has the same format as the primary backup file.

Before You Begin

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

- The source and target backup repositories that take part in the backup copy process must be added to the backup infrastructure.
- You must have a backup that has been successfully run at least once.

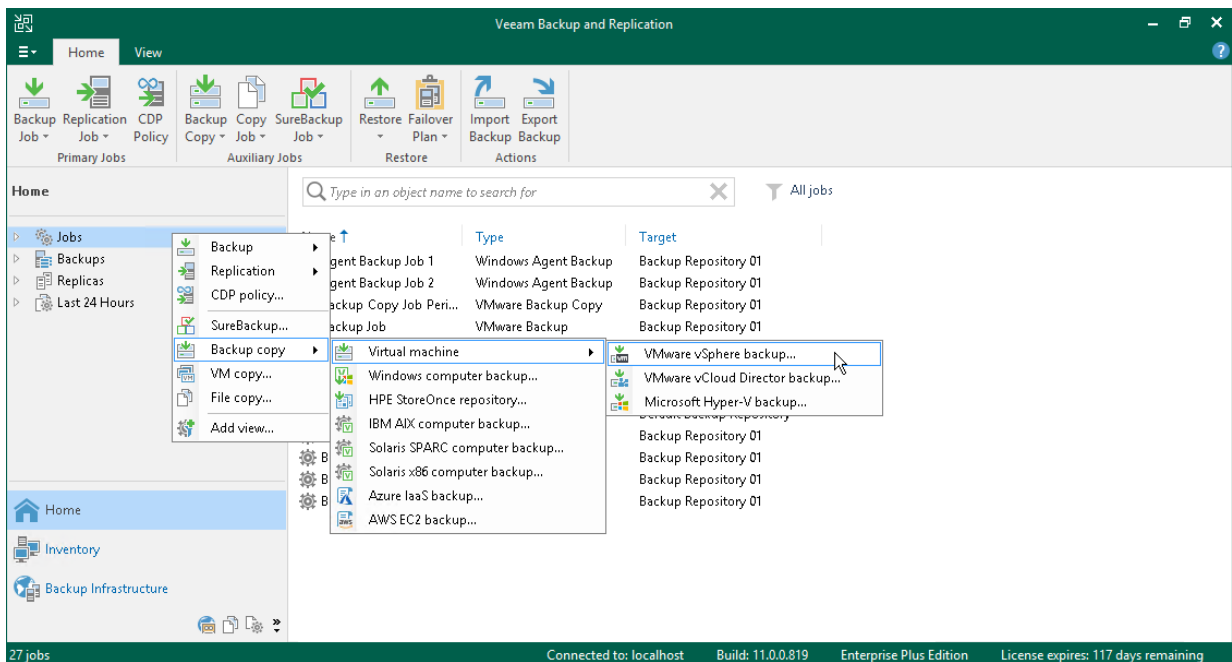
To check whether restore points are created, open the inventory pane of the **Home** view and select the **Backups** node. Then, expand the backup job and verify that there is at least one restore point available.

Creating Backup Copy Job

To create a backup copy job, do the following:

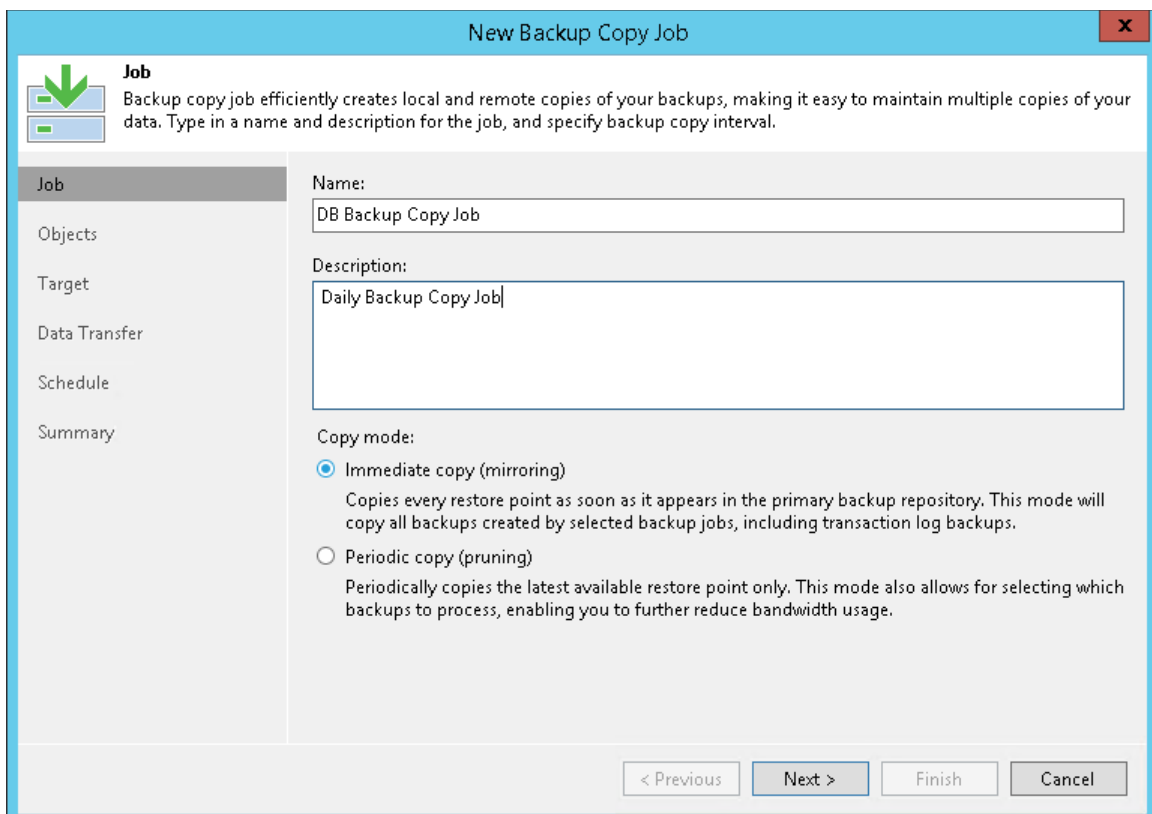
1. Open the **Home** view.

2. In the inventory pane, right-click **Jobs** and select **Backup Copy > Virtual machine > VMware vSphere backup** to launch the **New Backup Copy Job** wizard.

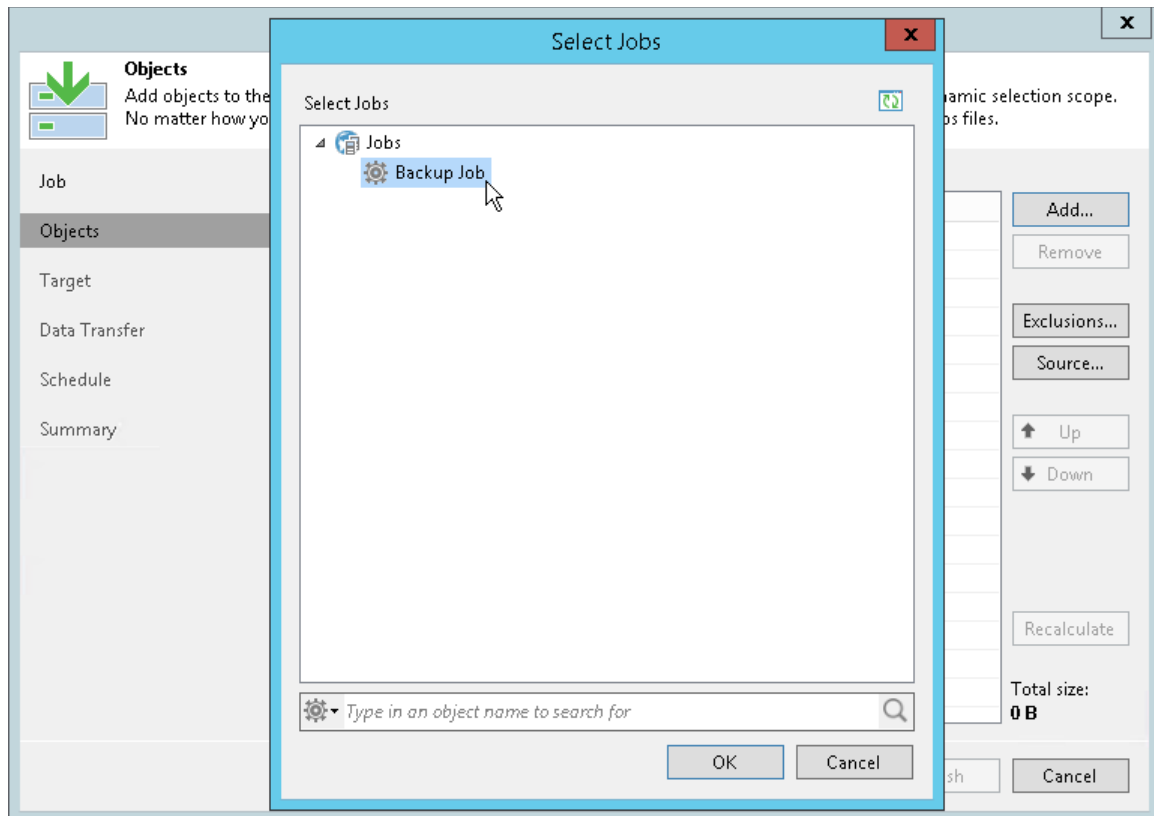


3. At the **Job** step of the wizard, do the following:
 - Specify a name and description for the backup copy job.
 - In the **Copy mode** field, check that the **Immediate copy** is selected.

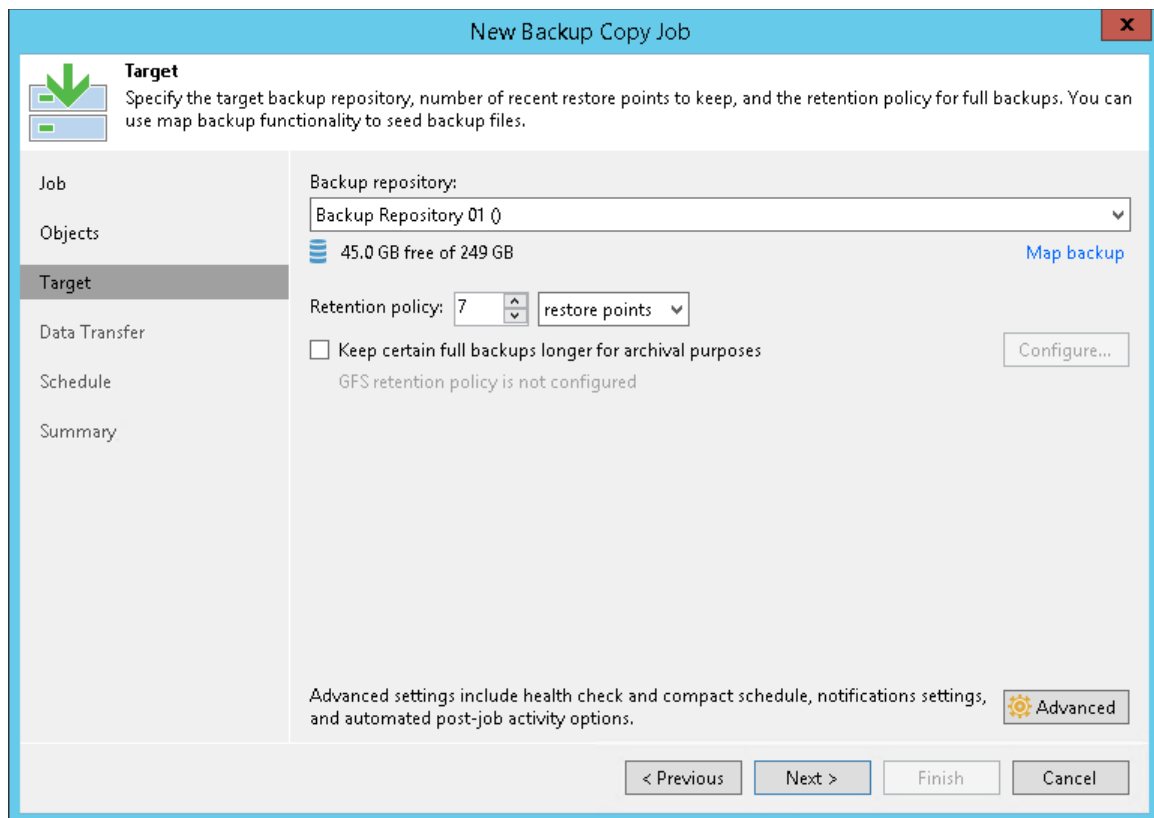
In the immediate copy mode, Veeam Backup & Replication copies new data as soon as it appears on the source repository. For more information, see [Backup Copy Modes](#) in the Veeam Backup & Replication User Guide.



- At the **Objects** step of the wizard, click **Add** and select backup jobs that you want to copy.

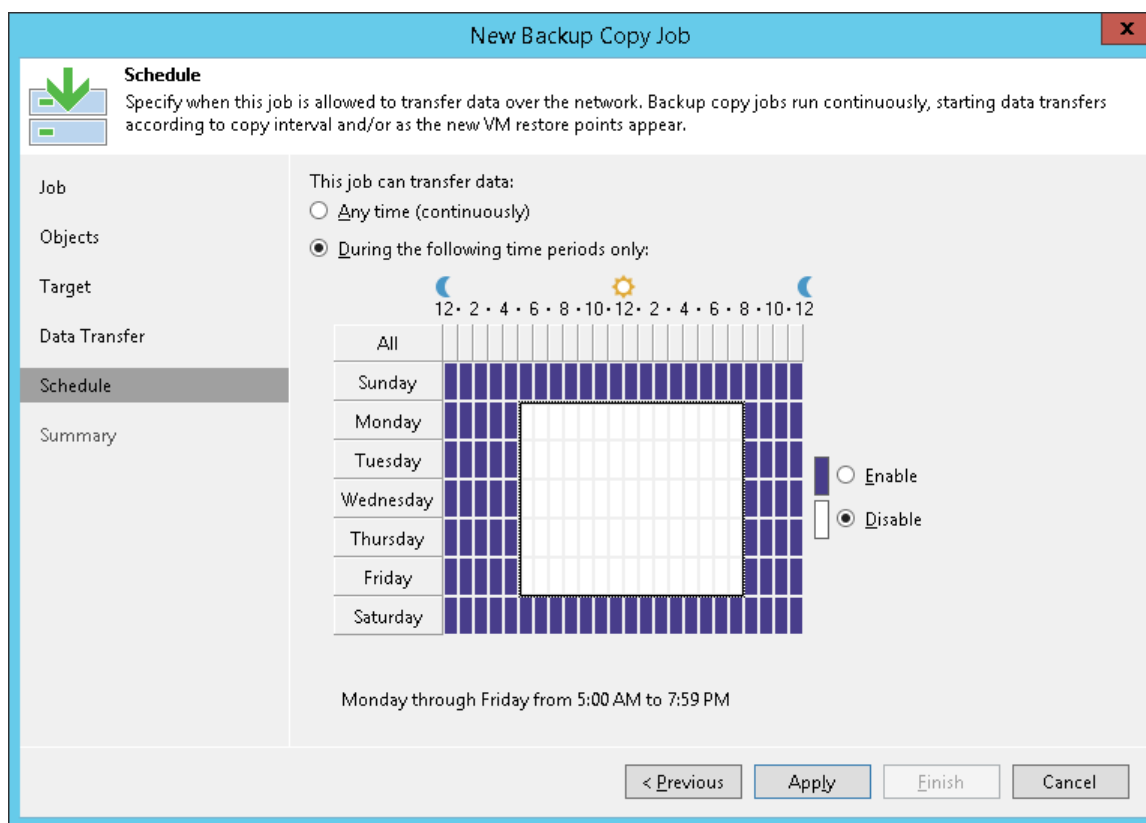


- At the **Target** step of the wizard, select the backup repository where you want to store the backup copy. For other settings, keep the default values.



- At the **Data Transfer** step of the wizard, keep the default settings.

- At the **Schedule** step of the wizard, define the period of time when the backup copy job is allowed to transport data over the network.



- At the last step of the wizard, select the **Enable the job when I click Finish** check box and click **Finish**. The job will start running in the continuous mode.

Reference

For more information on the backup copy, see the [Backup Copy](#) section in the Veeam Backup & Replication User Guide.

VM Replication

When you replicate a VM, Veeam Backup & Replication creates the exact copy of the VM in the native VMware vSphere format on a spare ESXi host and keeps this copy synchronized with the original VM.

Replication provides the best recovery time objective (RTO) values. You actually have the copy of a VM in the ready-to-start state. That is why replication is recommended for VMs running most critical applications.

Replication is a job-driven process. During the first run of a replication job, Veeam Backup & Replication copies data of the original VM running on the source host and creates its full replica on the target host. During next job runs, Veeam Backup & Replication copies only those data blocks that have changed since the last replication job session. Veeam Backup & Replication writes these changes to restore points, so that you can further publish this replica in the required state.

Veeam Backup & Replication supports several replication scenarios. Depending on the location of the host where you plan to store replicas, you can choose the following scenarios:

- Onsite replication
The target host is located in the same site as the source host.
- Offsite replication
The target host is located on another site.

In this section, you will learn how to work with onsite replicas. For more information on offsite replication, see [Replication Scenarios](#) in the Veeam Backup & Replication User Guide.

Reference

For more information on replication, see [Replication](#) in the Veeam Backup & Replication User Guide.

Creating Replication Job

Before You Begin

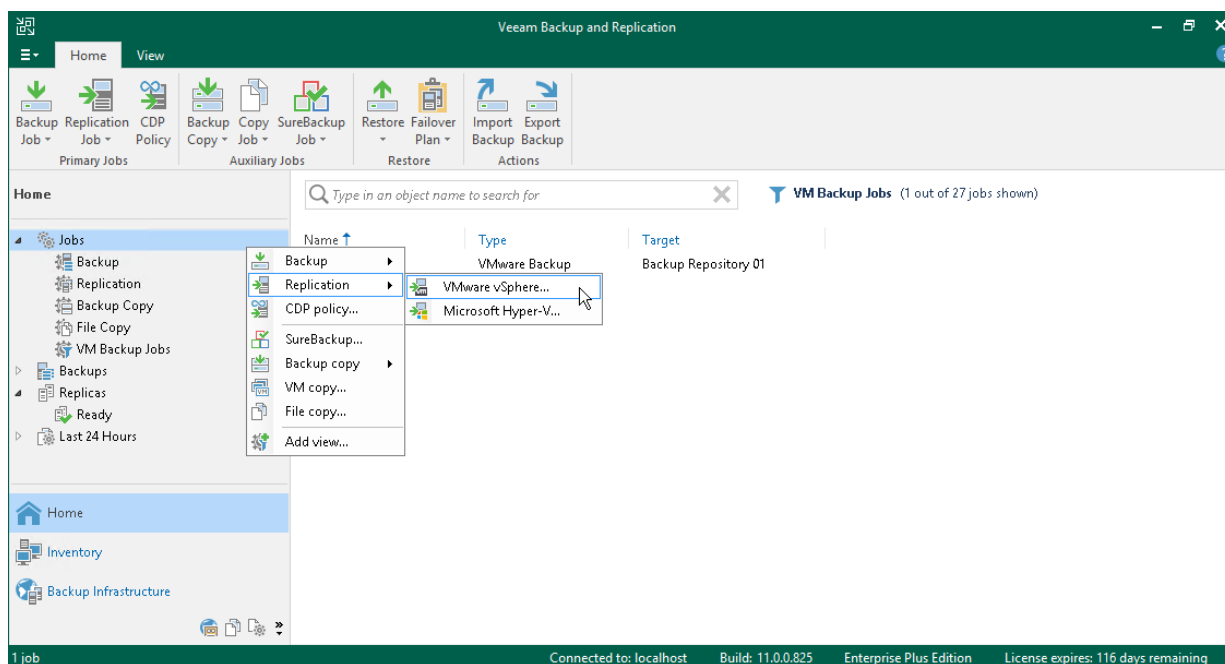
Before you replicate a VM, consider the following:

- You must add all components needed for the replication process to the backup infrastructure. These components are source and target ESXi hosts, backup proxy and backup repository.
- Due to VMware vSphere limitations, if you change the size of VM disks on the source VM, Veeam Backup & Replication deletes all available restore points on the VM replica during the next replication job session. For more information, see [this Veeam KB article](#).

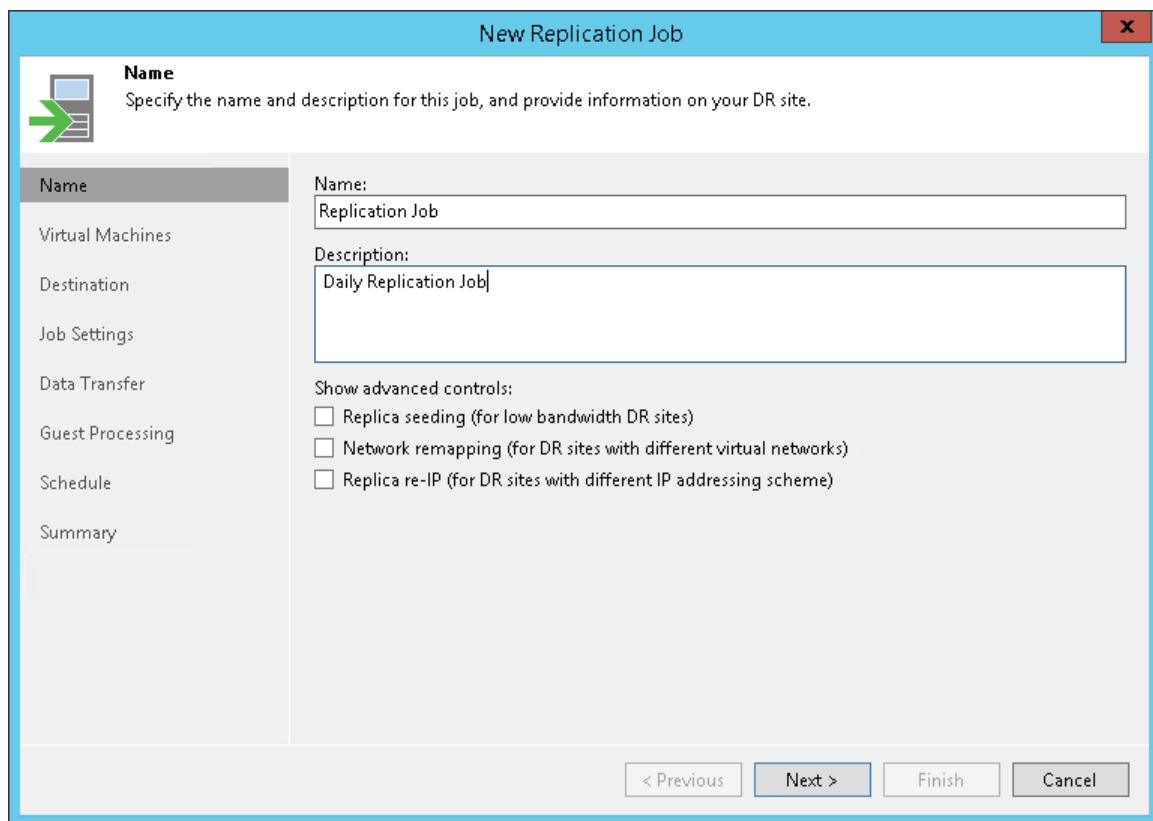
Creating Replication Job

To replicate a VM, do the following:

1. Open the **Home** view.
2. In the inventory pane of the **Home** view, right-click the **Jobs** node and select **Replication > Virtual machine > VMware vSphere** to launch the **New Replication Job** wizard.



- At the **Name** step of the wizard, specify a name and description for the replication job. For other settings, leave the default values.



The screenshot shows the 'New Replication Job' wizard window. The title bar is blue with the text 'New Replication Job' and a close button. The main area has a light blue header with a green arrow icon and the text 'Name' and 'Specify the name and description for this job, and provide information on your DR site.' Below this is a sidebar with a list of steps: 'Name' (selected), 'Virtual Machines', 'Destination', 'Job Settings', 'Data Transfer', 'Guest Processing', 'Schedule', and 'Summary'. The main content area has two text input fields: 'Name:' with the value 'Replication Job' and 'Description:' with the value 'Daily Replication Job'. Below these fields are three checkboxes under the heading 'Show advanced controls:': 'Replica seeding (for low bandwidth DR sites)', 'Network remapping (for DR sites with different virtual networks)', and 'Replica re-IP (for DR sites with different IP addressing scheme)'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

New Replication Job

Name
Specify the name and description for this job, and provide information on your DR site.

Name:
Replication Job

Description:
Daily Replication Job


Show advanced controls:

- ☐ Replica seeding (for low bandwidth DR sites)
- ☐ Network remapping (for DR sites with different virtual networks)
- ☐ Replica re-IP (for DR sites with different IP addressing scheme)

< Previous Next > Finish Cancel

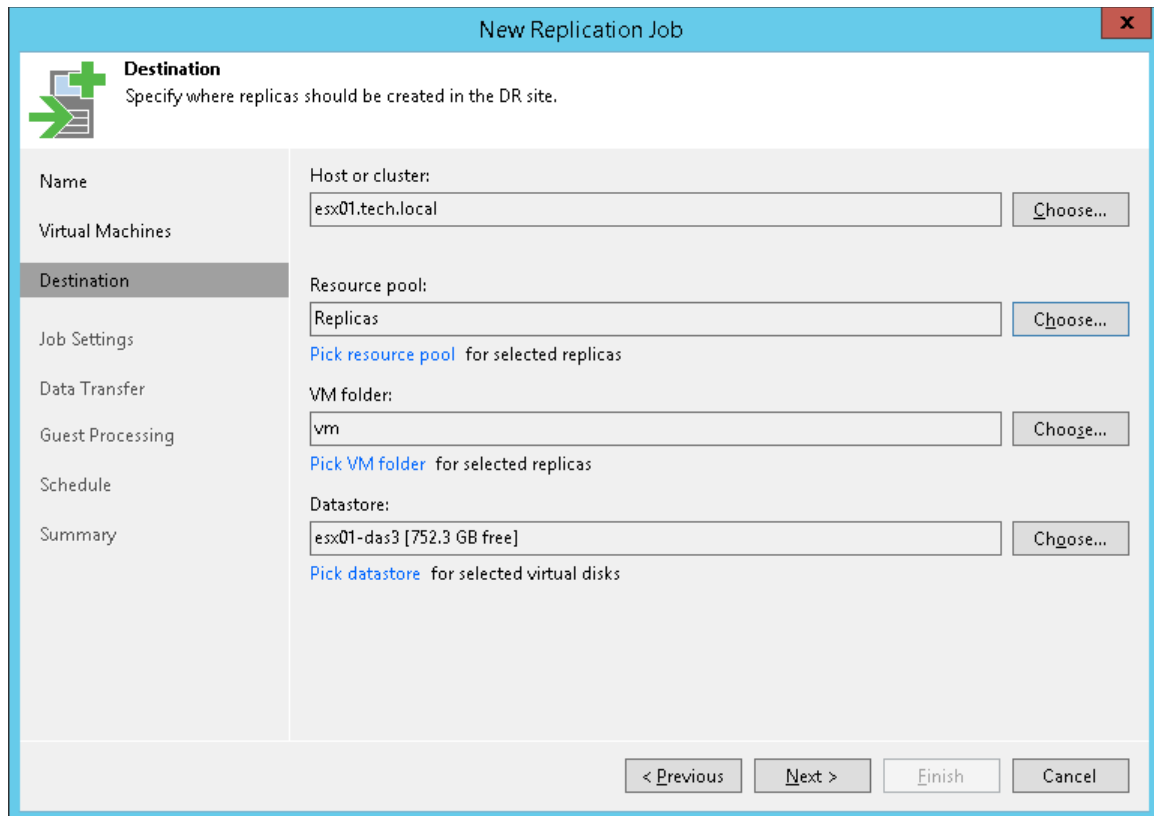
4. At the **Virtual Machines** step of the wizard, click **Add**. From the list, select VMs that you want to replicate.

You can also replicate VM containers: folders, resource pools, clusters, vApps, datastores and so on. If you add a new VM to the container after the replication job is created, Veeam Backup & Replication automatically updates the job to include the new VM.



5. At the **Destination** step of the wizard, do the following:
- Click **Choose** next to the **Host or cluster** field and select a host on which the VM replica must be registered.
 - Click **Choose** next to the **Resource pool** field and select the target resource pool.
 - Click **Choose** next to the **VM folder** field and select the folder in which the replica must be placed.

- Click **Choose** next to the **Datastore** field and select a datastore where VM replica files must be stored.



New Replication Job

Destination
Specify where replicas should be created in the DR site.

Virtual Machines

Destination

Job Settings

Data Transfer

Guest Processing

Schedule

Summary

Host or cluster:
esx01.tech.local [Choose...](#)

Resource pool:
Replicas [Choose...](#)
[Pick resource pool](#) for selected replicas

VM folder:
vm [Choose...](#)
[Pick VM folder](#) for selected replicas

Datastore:
esx01-das3 [752.3 GB free] [Choose...](#)
[Pick datastore](#) for selected virtual disks

< Previous Next > Finish Cancel

6. At the **Job Settings** step of the wizard, do the following:

- In the **Repository for replica metadata** list, select the backup repository where you want to store the metadata file.
- In the **Replica name suffix** field, specify the suffix that will be appended to the name of the original VM.

- In the **Restore points to keep** field, define the number of restore points to keep.

When this number is exceeded, the earliest restore point is removed. Due to VMware restrictions, the maximum number of restore points for VM replicas is limited to 28.

The screenshot shows the 'New Replication Job' wizard in the Veeam Backup & Replication console. The 'Job Settings' step is active, showing a sidebar with steps: Name, Virtual Machines, Destination, Job Settings (selected), Data Transfer, Guest Processing, Schedule, and Summary. The main area contains the following settings:

- Repository for replica metadata:** A dropdown menu showing 'Default Backup Repository (Created by Veeam Backup)' with a status bar indicating '100 GB free of 119 GB'.
- Replica settings:**
 - Replica name suffix:** A text box containing '_replica'.
 - Restore points to keep:** A spinner box set to '7'.
- Advanced settings:** A note at the bottom states 'Advanced job settings include traffic compression, block size, notification settings, automated post-job activity and other options.' with an 'Advanced' button (gear icon).

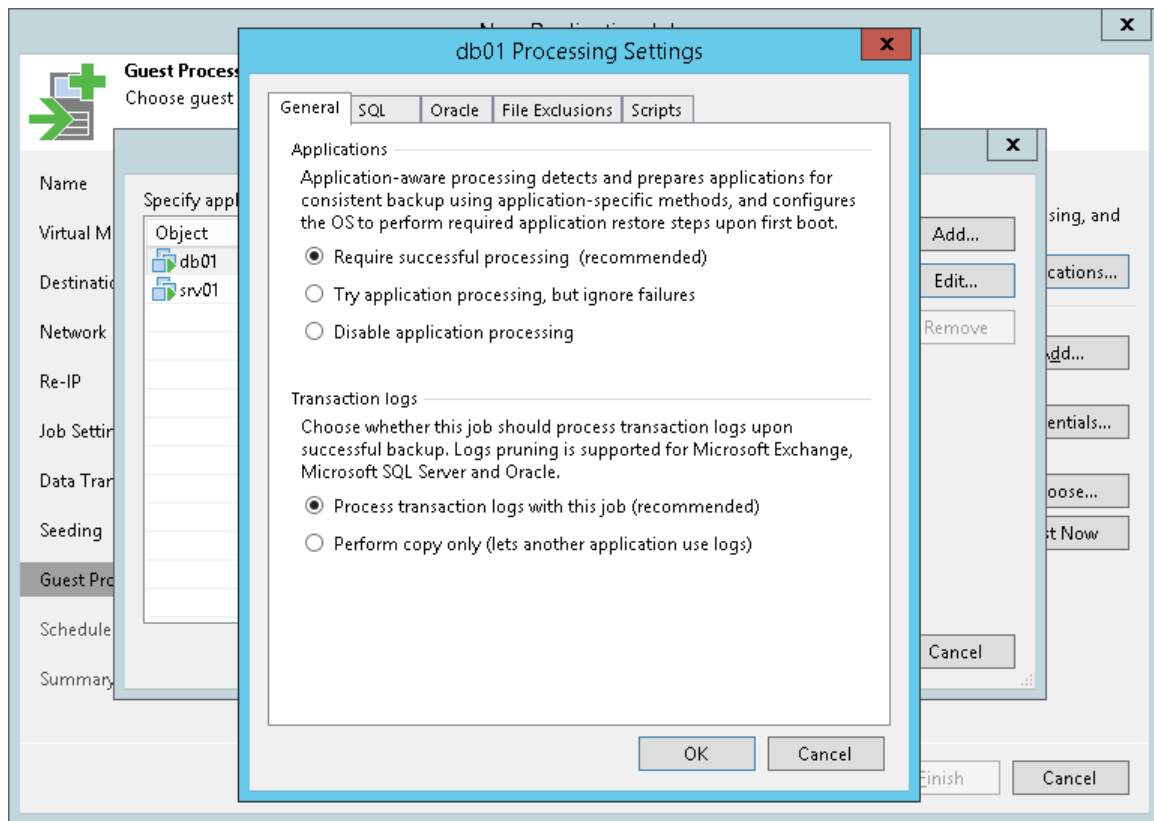
At the bottom of the wizard are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

7. At the **Data Transfer** step of the wizard, leave the default settings.
8. At the **Guest Processing** step of the wizard, leave the default settings if you do not need transactionally consistent replicas. Otherwise, select the **Enable application-aware processing** check box and specify credentials of a user account to connect to the VM guest OS. The user account must have Administrator permissions.

To specify advanced options for VSS processing, click **Applications**. Select a VM in the list and click **Edit**. In the opened window on the **General** tab, do the following:

- In the **Applications** section, select **Try application processing, but ignore failures** to continue the replication job even if VSS errors occur. If VSS processing fails, the created replica will not be transactionally consistent but crash consistent.

- In the **Transaction logs** section, check that the **Process transaction logs with this job** option is selected.



9. At the **Schedule** step of the wizard, do the following:

- Select the **Run the job automatically** check box. If you do not select this check box, you will have to launch the job manually. For details, see [Start Replication Job Manually](#).
- Select the schedule type: daily, monthly or periodically.

In the **Periodically every** field, you can select **Continuously** to run the job in a non-stop manner. A new session of the job will start as soon as the previous job session completes.

- c. Make sure the **Retry failed VM processing** check box is selected.

The screenshot shows the 'New Replication Job' wizard in the Veeam Backup & Replication console, specifically the 'Schedule' step. The left-hand navigation pane lists the steps: Name, Virtual Machines, Destination, Job Settings, Data Transfer, Guest Processing, Schedule (which is currently selected and highlighted), and Summary. The main area of the wizard is titled 'Schedule' and includes a sub-header: 'Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.' Below this, there are several configuration options. The 'Run the job automatically' checkbox is checked. Under this, four scheduling options are available: 'Daily at this time:' (selected with a radio button), 'Monthly at this time:', 'Periodically every:', and 'After this job:'. The 'Daily at this time:' option is configured with a time of '10:00 PM' and a frequency of 'Everyday'. Below these options, the 'Automatic retry' section is visible, with the 'Retry failed VMs processing:' checkbox checked, set to '3' times, and a wait time of '10' minutes before each retry attempt. The 'Backup window' section at the bottom has the 'Terminate job if it exceeds allowed backup window' checkbox unchecked. At the bottom right of the wizard, there are four buttons: '< Previous', 'Apply', 'Finish', and 'Cancel'.

10. At the **Summary** step of the wizard, select the **Run the job when I click Finish** check box and click the **Finish** button.
11. In the inventory pane of the **Home** view, expand the **Last 24 Hours** node to see the created job.
12. Open **vSphere Client** and make sure that the replica appeared on the target host.

Reference

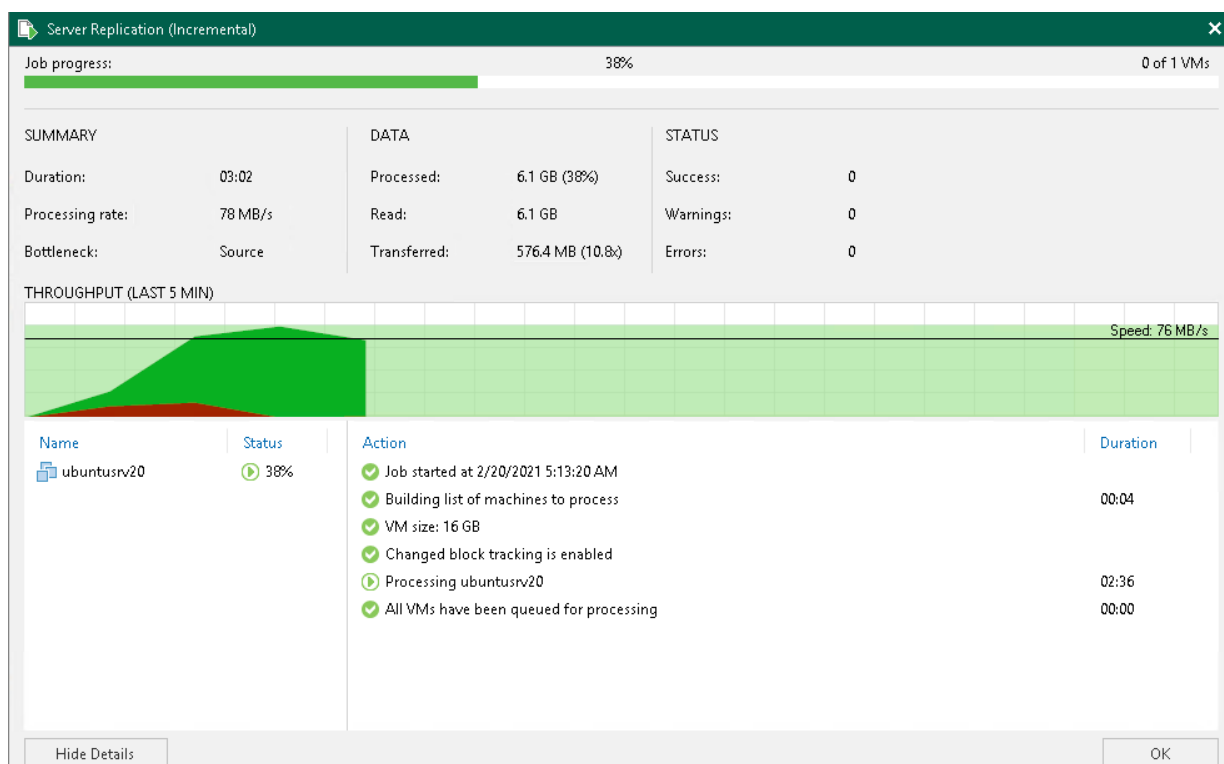
For more information on replica creation, see [Creating Replication Jobs](#) in the Veeam Backup & Replication User Guide.

Monitoring Job Performance in Real Time

When the job is running, you can view job statistics in real time. Statistics include job progress, duration, processing rate, performance bottlenecks, the amount of read and transferred data, and other details of the job performance.

To view the job statistics, do the following:

1. In the inventory pane of the **Home** view, select the **Jobs > Replication** node.
2. In the working area, right-click a job and click **Statistics**.
3. In the opened window, select a VM to view its statistics.



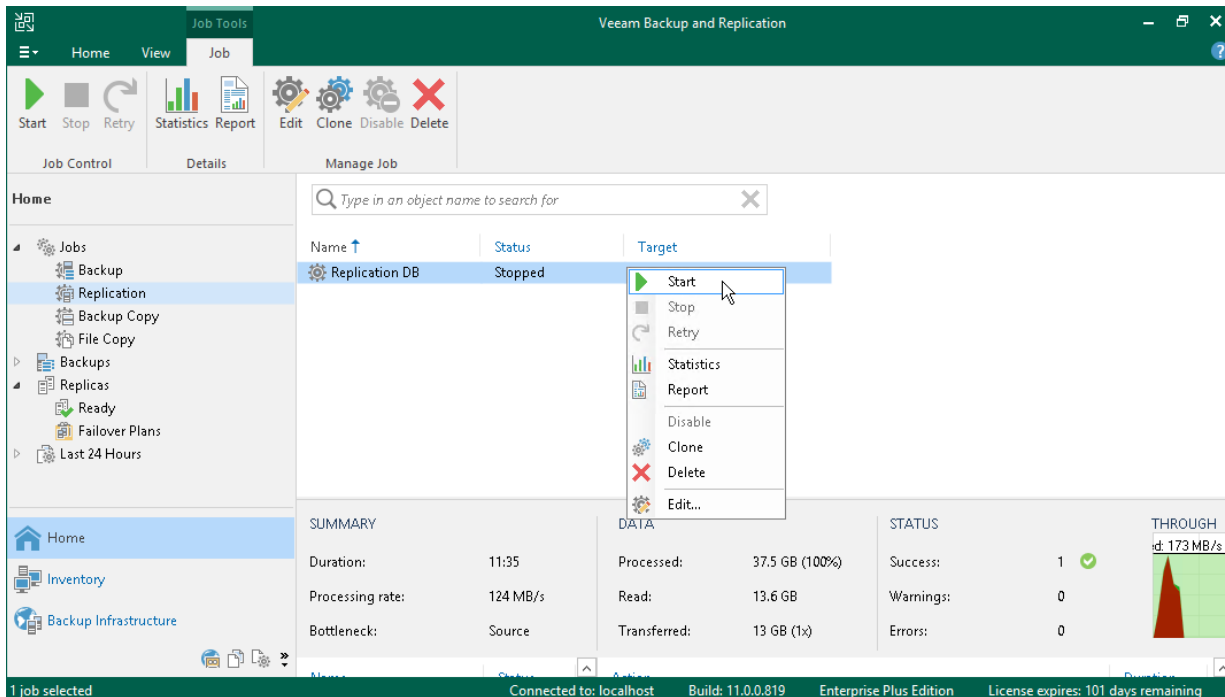
Note that the job must complete with the *Success* or *Warning* status. If the job completes with the *Failed* status, Veeam Backup & Replication does not create a replica and is not able to perform failover and failback operations.

You can configure email notifications to get job results. For details, see [Configuring Global Email Notification Settings](#) in the Veeam Backup & Replication User Guide.

Start Replication Job Manually

If you do not schedule a replication job, you must start it manually. To start the job, do the following:

1. Open the **Home** view.
2. In the inventory pane, select **Jobs > Replication**.
3. In the working area, right-click the job and select **Start**. Wait for the job to complete. Note that the job must complete with the *Success* or *Warning* status.
4. Open **vSphere Client** and make sure that a VM replica is created.



Replica Failover and Failback

If the original VM in the production site becomes unavailable, you can quickly restore services by failing over to its replica. When you perform failover, the VM replica takes the role of the original VM. All processes shift from the original VM on the production host to the VM replica on the secondary host. You can fail over to the latest state of a replica or to any of its restore points.

When you fail over to the VM replica, Veeam Backup & Replication changes the replica state from *Normal* to *Failover*.

Failover is an intermediate step that needs to be finalized. Depending on a disaster recovery scenario, you can do one of the following:

- **Undo failover**

When you undo failover, you switch back to the original VM and discard all changes made to the VM replica while it was running. The state of the VM replica gets back to *Normal*. You can use the undo failover scenario if you have failed over to the VM replica for testing and troubleshooting purposes and you do not need the changes made to the VM replica.

- **Perform failback**

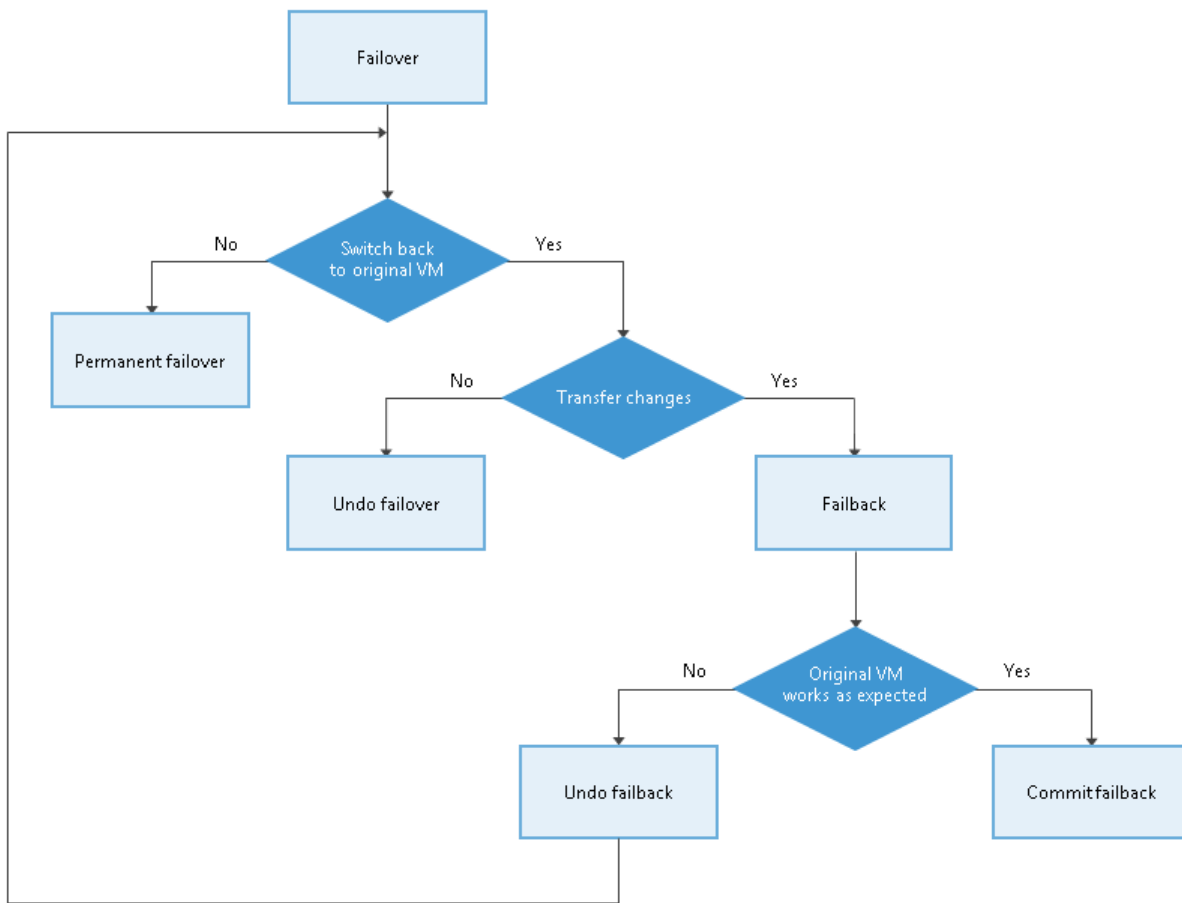
When you perform failback, you switch back to the original VM and transfer all changes that took place while the VM replica was running to the original VM. If the source host is not available, you can restore the original VM to a new location and switch back to it.

When you perform failback, changes are only transferred but not published. You must test whether the original VM works with these changes. Depending on the test results, you can do the following:

- **Commit failback.** When you commit failback, you confirm that the original VM works as expected and you want to get back to it. The state of the VM replica gets back to *Normal*.
- **Undo failback.** If the original VM is not working as expected, you can undo failback and get back to the VM replica. In this case, the state of the VM replica returns to *Failover*.

- **Perform permanent failover**

When you perform failover, you permanently switch from the original VM to a VM replica and use this replica as the original VM. This scenario is acceptable if the original VM and VM replica are located in the same site and are nearly equal in terms of resources.



Veeam Backup & Replication supports failover and failback operations for several VMs simultaneously. In case one or several hosts fail, you can use batch processing to restore operations with minimum downtime.

Performing Replica Failover

Before You Begin

Before you perform failover, consider the following:

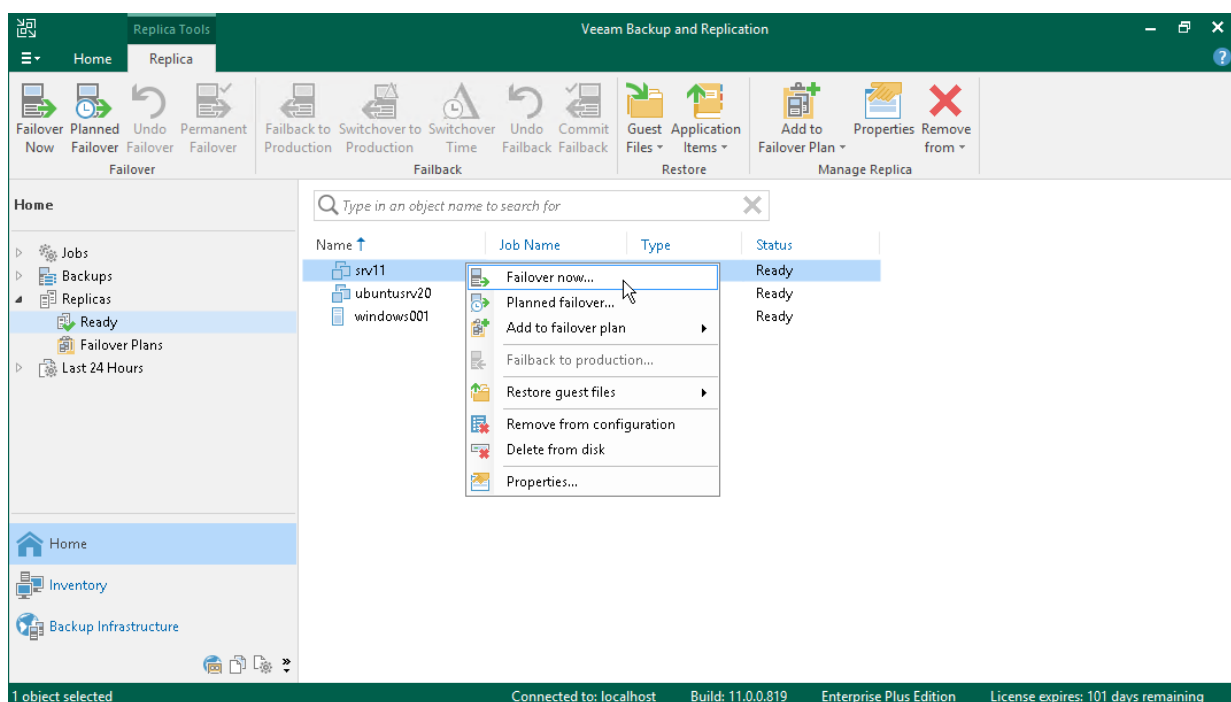
- For original VMs and replica VMs located in the same network. If you plan to perform replica failover while the original VM is running, consider temporarily disconnecting the original VM from the network to avoid IP addresses and/or machine names conflicts.
- To successfully fail over to a VM replica, make sure that this replica has at least one successfully created restore point.

To check whether restore points are created, open the inventory pane of the **Home** view and select the **Replicas** node. Then, select the VM and verify that there is at least one restore point available for the VM.


Performing Failover

To fail over to a VM replica, do the following.

1. In the inventory pane of the **Home** view, select the **Replicas** node.
2. Right-click the replicated VM and select **Failover Now** to launch the **VMware Failover Wizard**.



- At the **Virtual Machines** step of the wizard, select the VM from the list, click **Point** and choose the restore point to which you want to fail over.



- At the **Reason** step of the wizard, specify the reason for failover.
- At the **Summary** step of the wizard, click **Finish** to fail over to the VM replica.

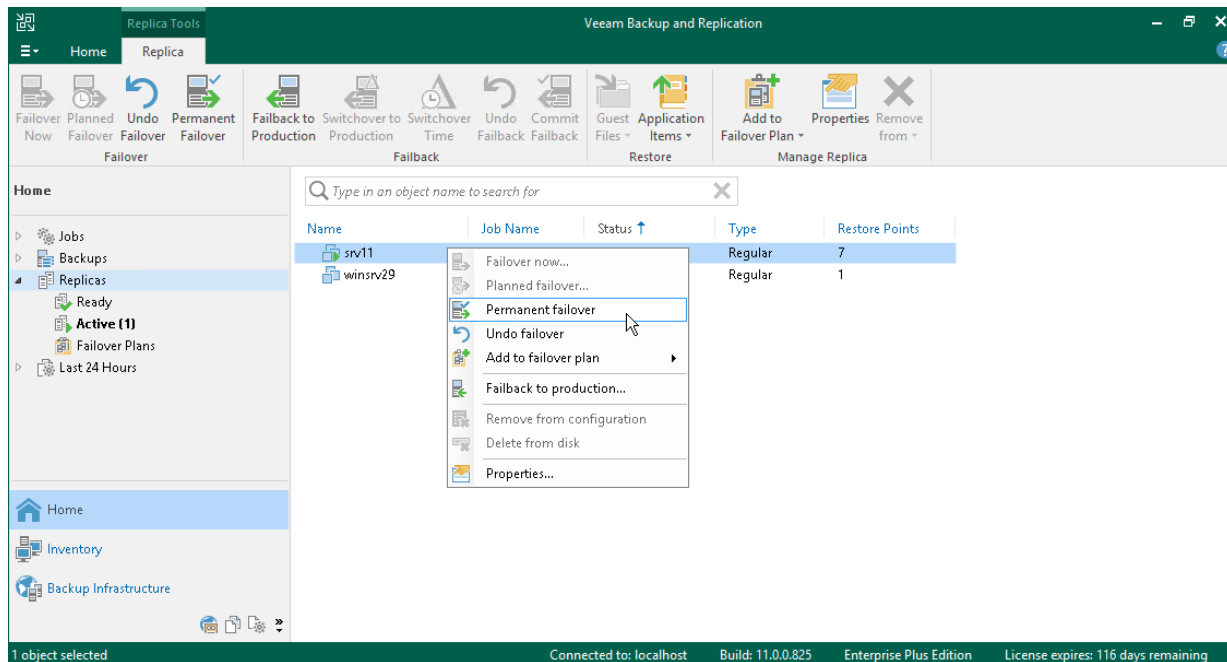
Reference

For more information on failover, see the [Replica Failover](#) section in the Veeam Backup & Replication User Guide.

Performing Permanent Failover

To perform permanent failover, do the following:

1. In the inventory pane of the **Home** view, click the **Replicas** node.
2. In the working area, right-click the VM replica and select **Permanent Failover**.
3. In the opened window, click **Yes** to confirm the operation.



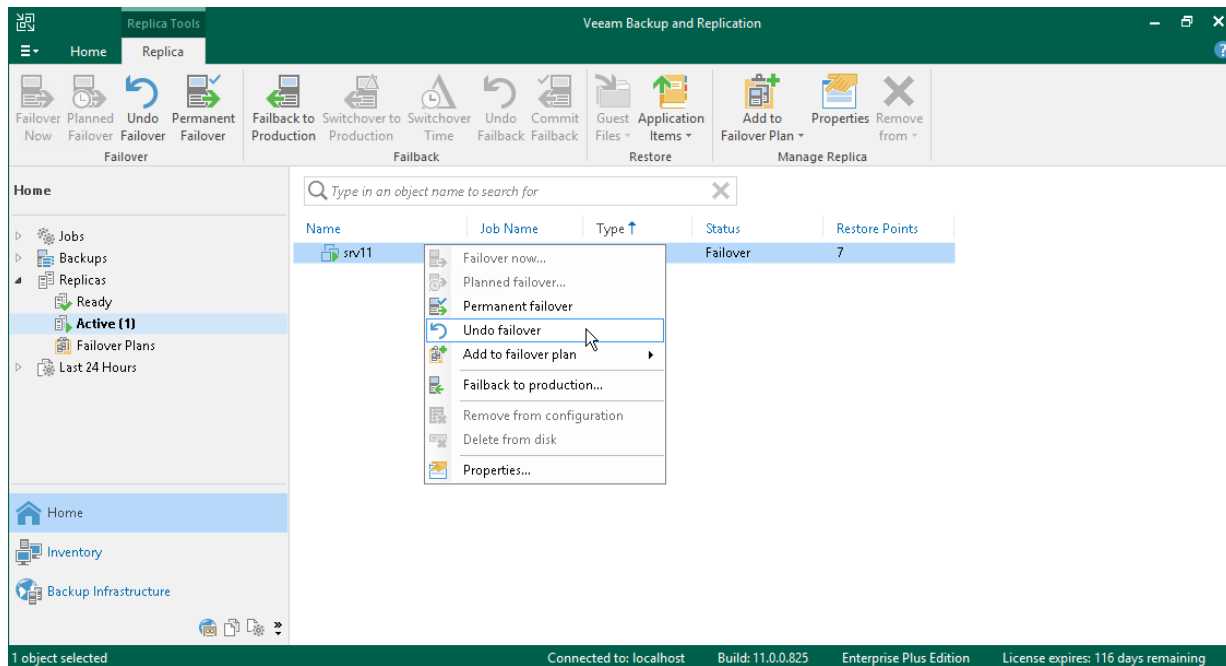
Reference

For more information on permanent failover, see the [Permanent Failover](#) section in the Veeam Backup & Replication User Guide.

Undoing Failover

To undo failover, do the following:

1. In the inventory pane of the **Home** view, select the **Replicas** node.
2. In the working area, right-click the VM replica and select **Undo Failover**.
3. In the opened window, click **Yes** to confirm the operation.



Reference

For more information on undoing failover, see the [Undo Failover](#) section in the Veeam Backup & Replication User Guide.

Performing Failback

You can fail back to a VM in the original or new location. In this section, you will learn how to fail back to the original VM on the source host. For more information on how to do this on another host, see the [Performing Failback](#) section in the Veeam Backup & Replication User Guide.

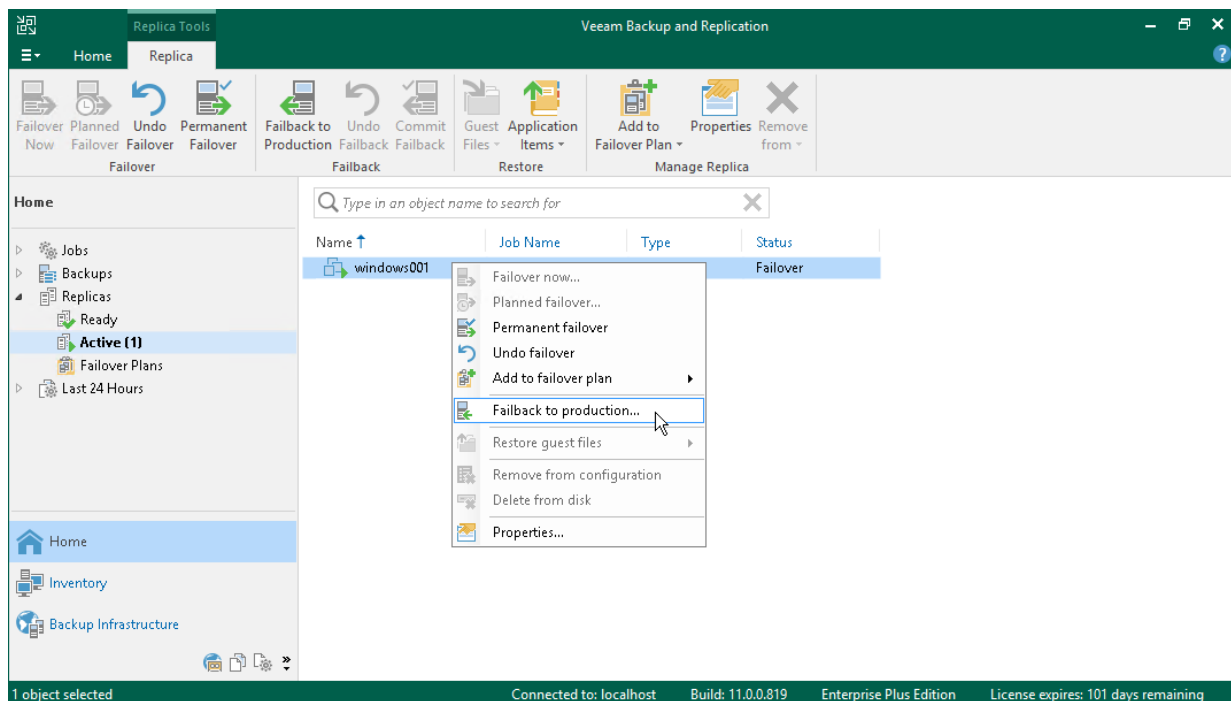
Before You Begin

Make sure that the VM replica for which you want to perform failback is in the *Failover* state. The replica gets into this state after you [perform replica failover](#).

Performing Failback

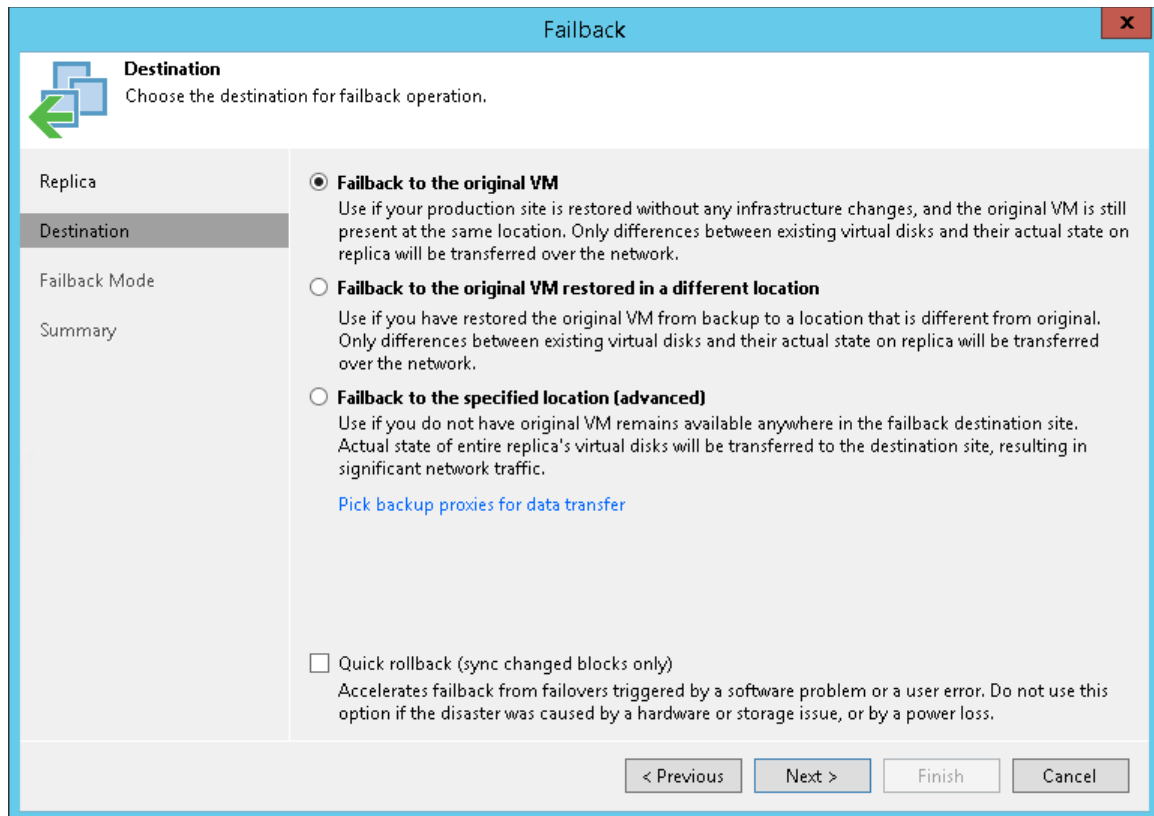
To fail back from VM replica to the original VM on the source host, do the following:

1. In the inventory pane of the **Home** view, select the **Replicas** node.
2. In the working area, right-click the VM replica and select **Failback to production** to launch the **Failback Wizard**.



3. At the **Replica** step of the wizard, click **Next**.

- At the **Destination** step of the wizard, select **Failback to the original VM**.



- At the **Failback Mode** step of the wizard, select **Auto**. In this case, failback will be performed as soon as VMs are ready.
- At the **Summary** step of the wizard, select the **Power on VM after restoring** check box and click **Finish**.

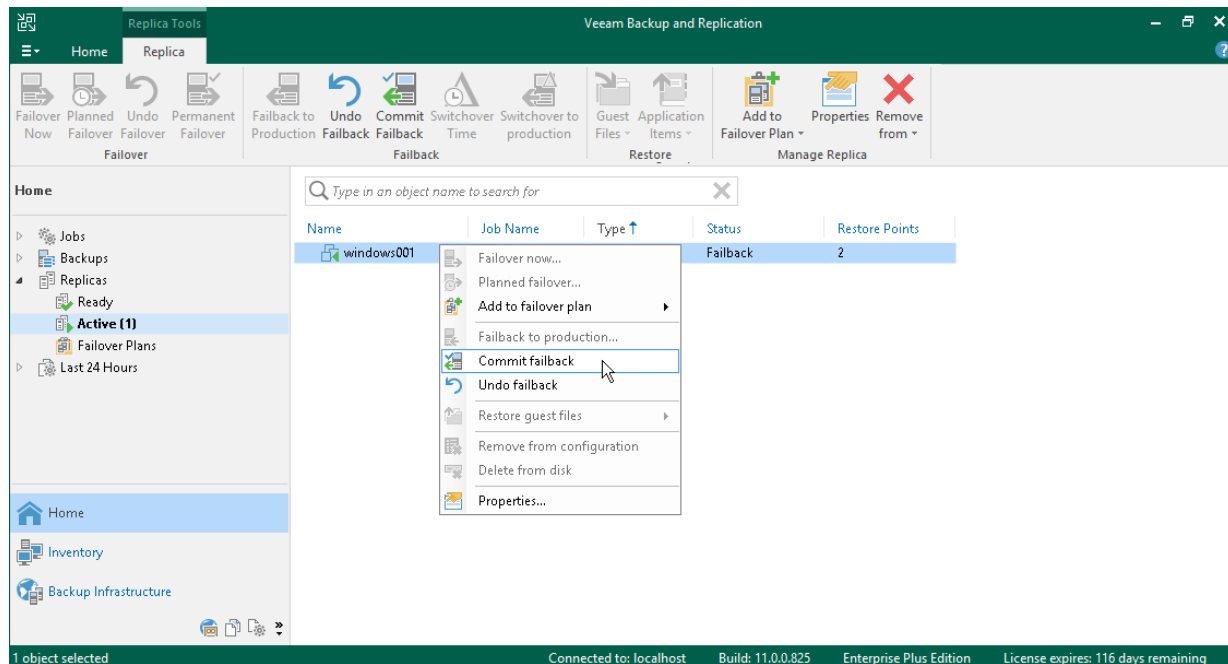
Reference

For more information on failback, see the [Replica Failback](#) section in the Veeam Backup & Replication User Guide.

Committing Failback

To commit a failback, do the following:

1. In the inventory pane of the **Home** view, select the **Replicas** node.
2. In the working area, right-click the VM replica and select **Commit Failback**.
3. In the opened window, click **Yes** to confirm the operation.



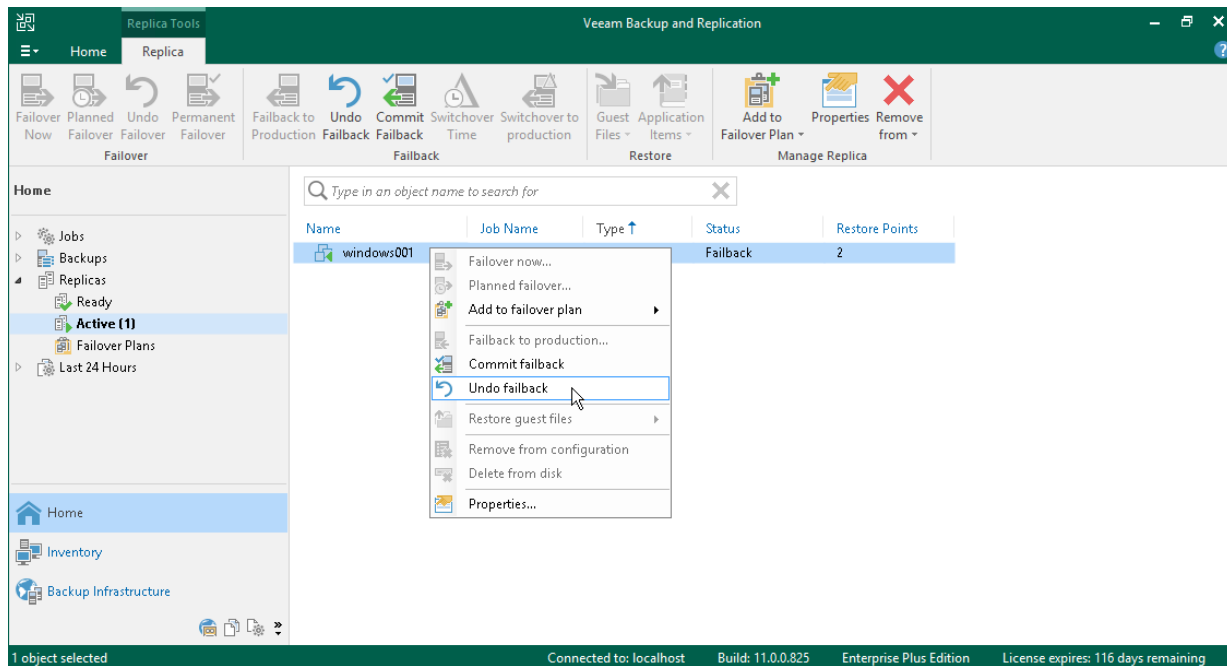
Reference

For more information on committing failback, see the [Commit Failback](#) section in the Veeam Backup & Replication User Guide.

Undoing Failback

To undo a failback, do the following:

1. In the inventory pane of the **Home** view, select the **Replicas** node.
2. In the working area, right-click the VM replica and select **Undo Failback**.
3. In the opened window, click **Yes** to confirm the operation.



Reference

For more information on undoing failback, see the [Undo Failback](#) section in the Veeam Backup & Replication User Guide.

Enterprise Manager

If you have a geographically dispersed virtual environment with multiple Veeam Backup & Replication servers, you can use Veeam Backup Enterprise Manager. Veeam Backup Enterprise Manager is a solution that helps manage multiple backup servers from a single web UI.

You can use Veeam Backup Enterprise Manager to perform the following tasks:

- Manage jobs configured on different backup servers from a single web console
- Edit and clone jobs
- Monitor the state of jobs
- Generate reports on jobs and backup servers
- Search for guest OS files in all backups and restore these files in one click

For the full list of the Veeam Backup Enterprise Manager capabilities, see [About Veeam Backup Enterprise Manager](#) in the Enterprise Manager User Guide.

NOTE:

Veeam Backup Enterprise Manager is not shipped with the Community edition. For more information, see [Editions Comparison](#).

Installing Veeam Backup Enterprise Manager

Before You Begin

Consider the following:

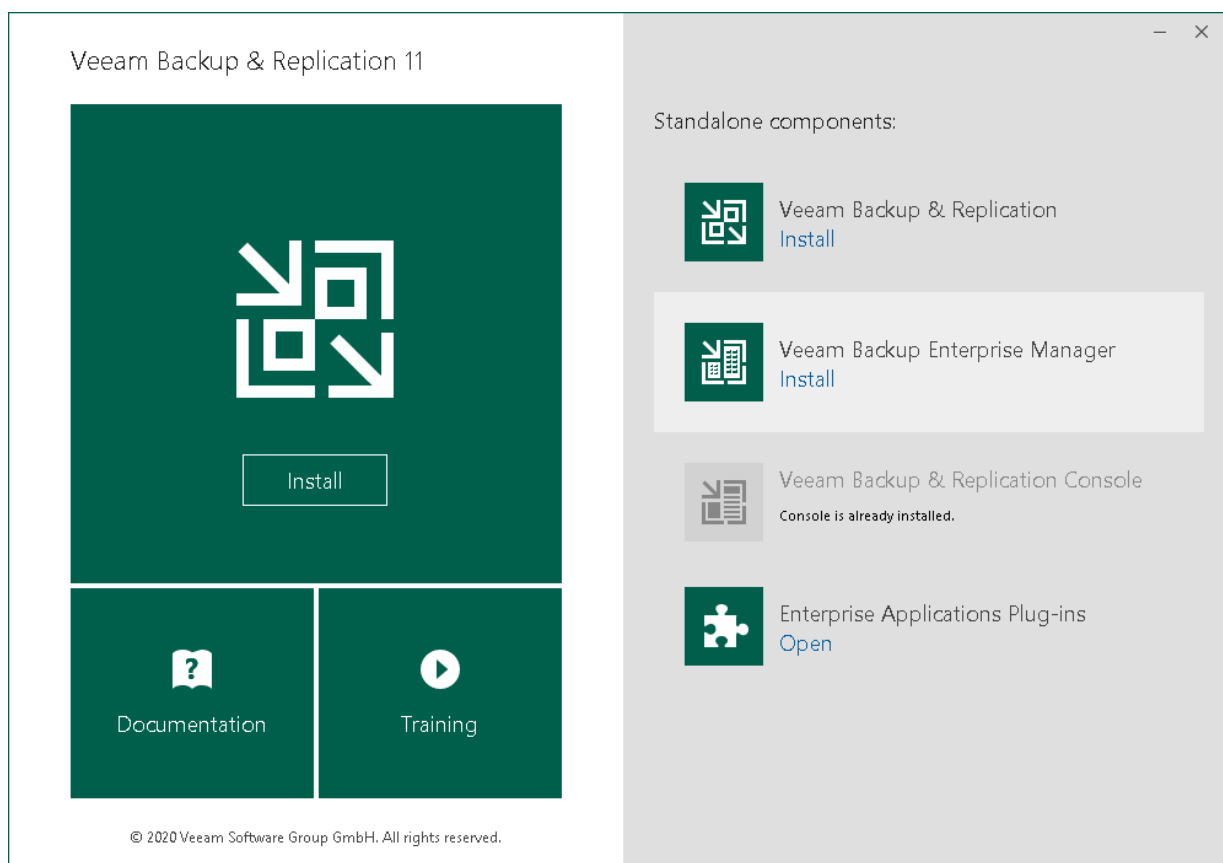
- The machine on which you plan to install Veeam Backup Enterprise Manager must meet the system requirements. For details, see [System Requirements](#) in the Enterprise Manager User Guide.
- It is recommended to install the same product version on the Veeam Backup Enterprise Manager server and Veeam Backup & Replication backup servers.
- If you plan to install Veeam Backup Enterprise Manager on the same machine where the backup server runs, you must disable all backup and replication jobs and close the Veeam Backup & Replication console.
- Make sure that all necessary ports are opened. For details, see [Used Ports](#) in the Enterprise Manager User Guide.

Installing Enterprise Manager

To install Veeam Backup Enterprise Manager, perform the following:

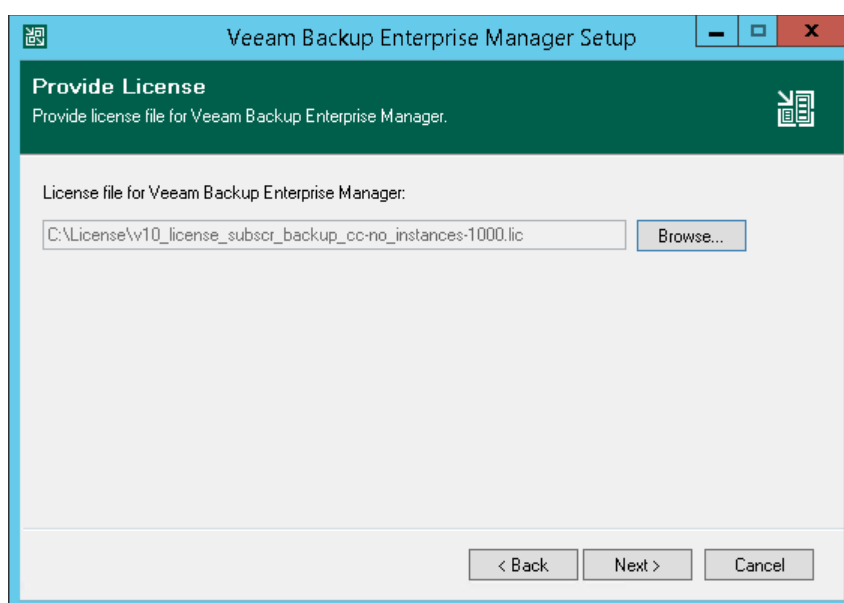
1. Download the latest version of Veeam Backup & Replication installation image from the [Download Veeam products](#) page.
2. Mount the installation image to the machine on which you plan to install Veeam Backup Enterprise Manager or burn the image file to a flash drive or other removable storage device.
3. Run the `Setup.exe` file from the image or disk to open the splash screen.

- On the splash screen, click the **Veeam Backup Enterprise Manager** tile in the **Standalone components** section to launch the **Veeam Backup Enterprise Manager Setup** wizard.



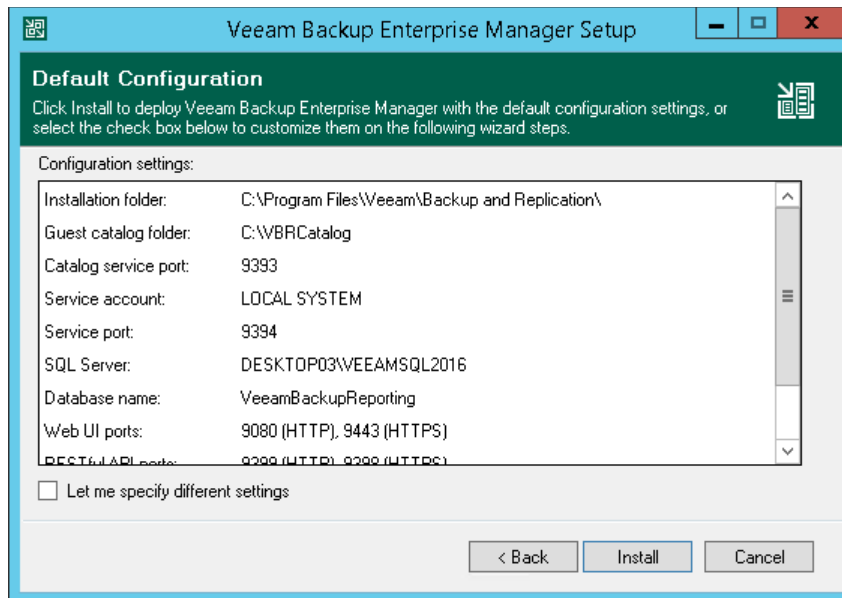
- At the **License Agreement** step of the wizard, read the license agreements and select check boxes to accept the terms.
- At the **Provide License** step of the wizard, specify the path to the license key.

If you install Veeam Backup Enterprise Manager on the backup server, you can proceed without providing a license file. In this case, Veeam Backup Enterprise Manager will use the license that is already installed on the backup server.



- At the **Program Features** step of the wizard, leave default settings.

8. At the **System Configuration Check** step of the wizard, install missing software components and enable missing features, if any.
9. At the **Default Configuration** step of the wizard, click **Install** to begin installation.



10. When the installation process completes, click **Finish** to close the wizard.

Adding Backup Servers

To manage backup servers from a single web console, you must add them to Veeam Backup Enterprise Manager.

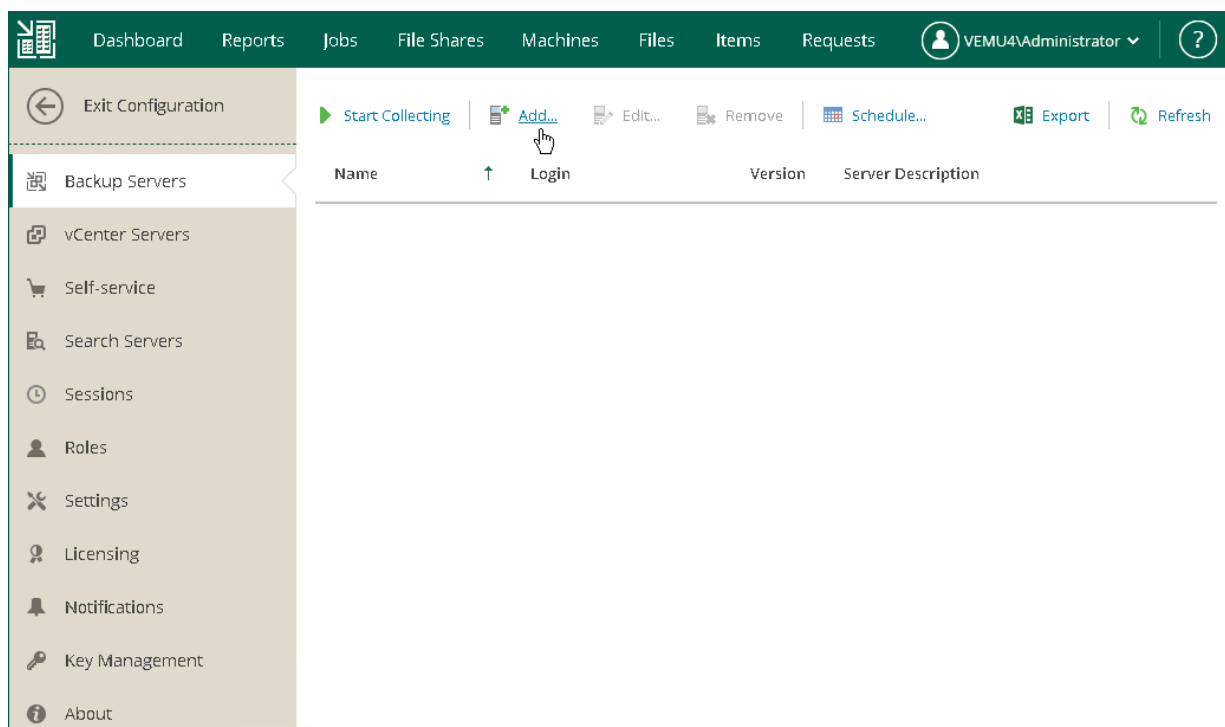
To add a backup server to Veeam Backup Enterprise Manager, do the following:

1. From the Microsoft Windows **Start** menu, select **Programs > Veeam > Veeam Backup Enterprise Manager to launch** Veeam Backup Enterprise Manager.

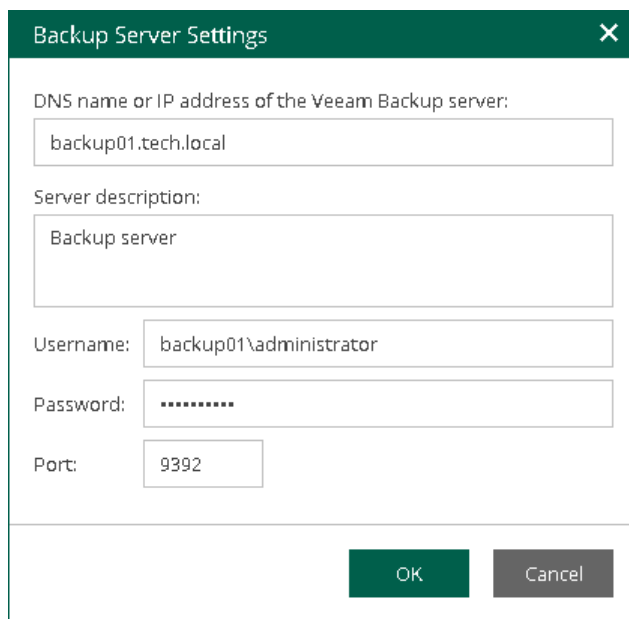
To access Veeam Backup Enterprise Manager remotely, use the following address:

`https://enterprise_manager_server_address:9443`

2. In the **Username** and **Password** fields, specify credentials of the user with local Administrator rights or the user who installed Veeam Backup Enterprise Manager.
3. Click **Login**.
4. At the top right corner of the opened window, click **Configuration** to open the **Configuration** view.
5. Select the **Backup Servers** tab. In the working area, click **Add** to open the **Backup Server Settings** window.



6. In the opened window, specify the DNS name or IP address of the backup server you want to add. Provide the name and password of the user account with local Administrator rights on the added backup server.



Backup Server Settings

DNS name or IP address of the Veeam Backup server:

backup01.tech.local

Server description:

Backup server

Username: backup01\administrator

Password:

Port: 9392

OK Cancel

7. Click **OK**.

Veeam Backup Enterprise Manager will start collecting data about all backup and replication jobs on the added backup server.

Reference

For more information on adding backup servers, see [Managing Veeam Backup Servers](#) in the Enterprise Manager User Guide.

Managing Jobs

Veeam Backup Enterprise Manager allows you to manage jobs that were configured on backup servers: start, stop, retry, edit and clone jobs.

In this section, you will learn how to clone and then edit jobs. When you clone a job, you create its exact copy. Configuration details of the created job copy are written to the same Microsoft SQL database where details of the original job are stored. You can work with the created job both in Veeam Backup Enterprise Manager and in the Veeam Backup & Replication console on the backup server.

Before You Begin

Consider the following:

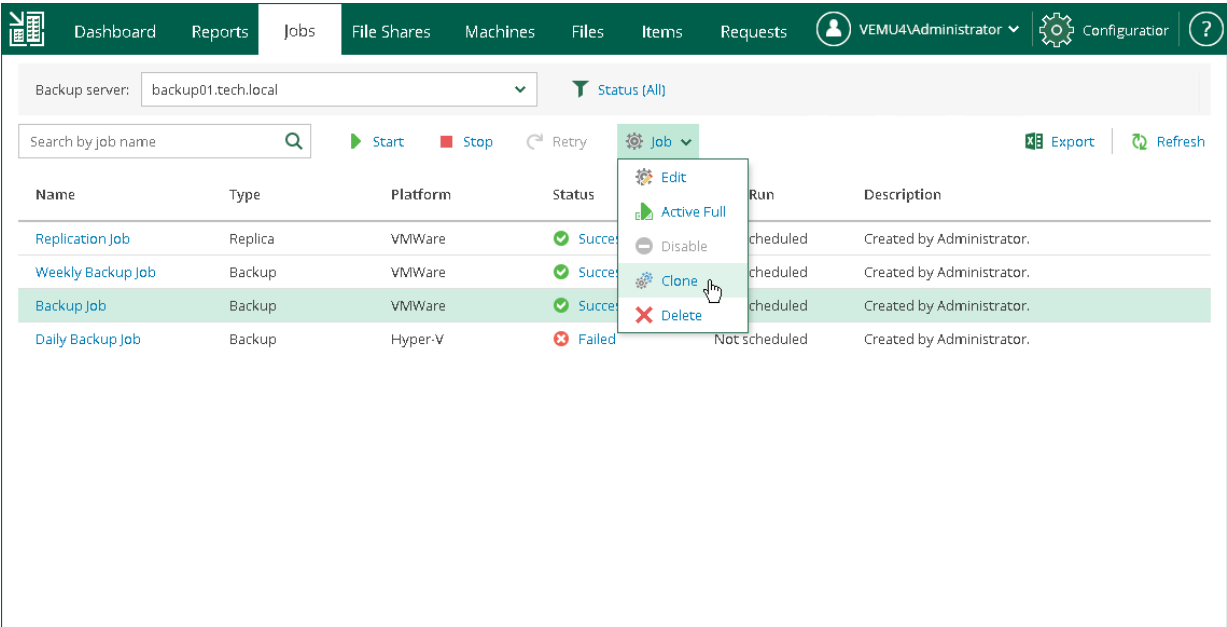
- Make sure that you have successfully connected backup servers to Veeam Backup Enterprise Manager and collected data from them. For details, see [Adding Veeam Backup Server](#).
- You have created jobs on the backup server.

Cloning and Editing Job

To clone a job, perform the following.

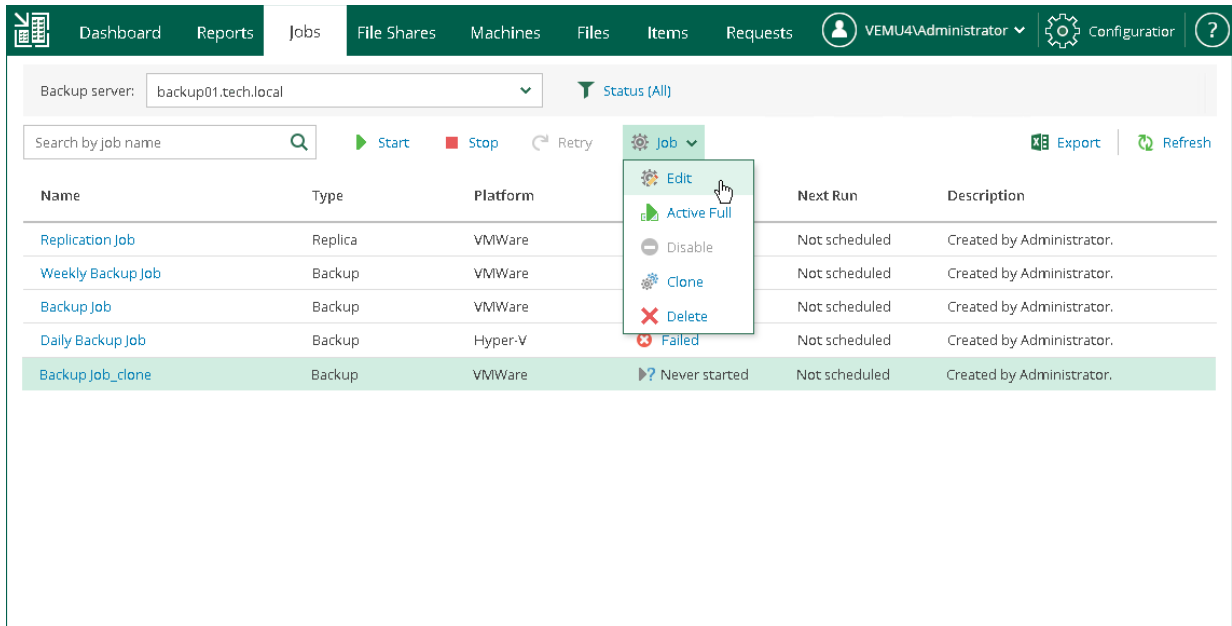
1. In Veeam Backup Enterprise Manager, click the **Jobs** tab.
2. Select the required job from the list, click **Job** at the top of the working area and click **Clone**.

The cloned job has the same name as the original job plus the `_cloned` suffix.

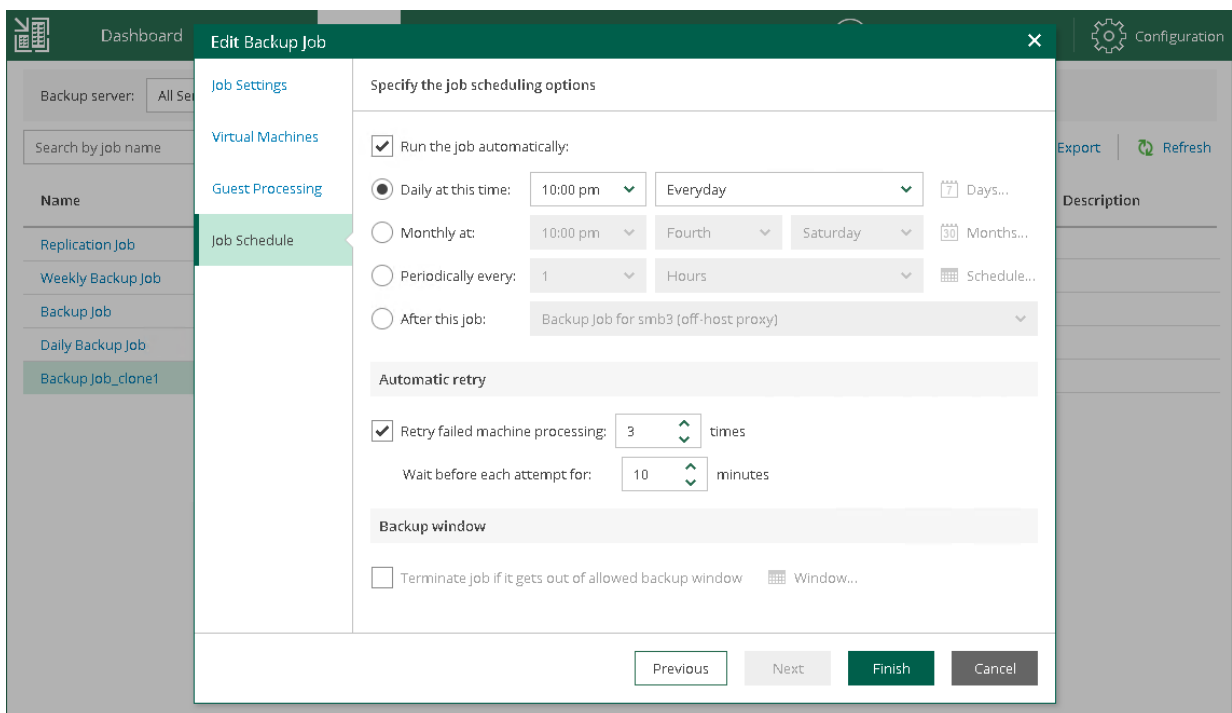


3. Select the cloned job from the list.

4. At the top of the working area, click **Job** and select **Edit**.



5. Follow the steps of the wizard and edit the job settings as required.



6. At the last step of the wizard, click **Finish**.

Reference

For more information on managing jobs, see [Managing Backup Jobs in Veeam Enterprise Manager](#) in the Enterprise Manager User Guide.

Performing 1-Click File Restore

Veeam Backup Enterprise Manager allows you to search through Microsoft Windows and Linux guest files in backed up VMs. Once you find the required file, you can immediately restore it using 1-Click file restore capability. The file can be restored to its original location or saved to a local machine.

Before You Begin

Consider the following:

- The Enterprise or Enterprise Plus license is installed on the Veeam Backup Enterprise Manager server.
- Make sure that you have successfully connected backup servers to Veeam Backup Enterprise Manager and collected data from them. For details, see [Adding Veeam Backup Server](#).
- You can search files on machines that have at least one successfully created backup with guest file indexing enabled. For details, see [Creating Application-Aware Backup Job](#).

Performing 1-Click Restore

To restore a guest OS file, perform the following:

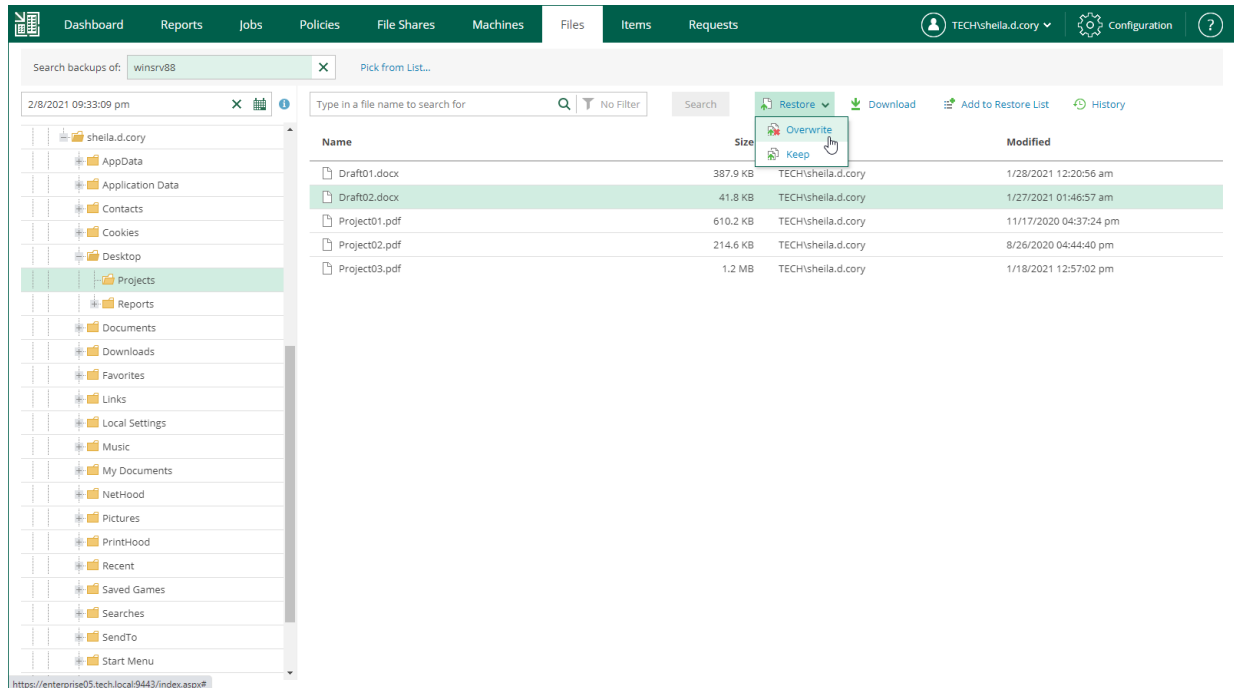
1. In the main view of Veeam Backup Enterprise Manager, click the **Files** tab.
2. In the **Type in machine name** field, specify the name of the backed up VM whose file system you want to browse.
3. In the field with the calendar icon, enter or choose a date and time of a restore point from which you want to restore files.
4. Click the **Mount** link and wait for Veeam Backup & Replication to mount the content of the backup file to the backup server.

After the backup is mounted, you can browse the guest OS files.

5. Select the necessary file from the list.

- At the top of the working area, click **Restore > Keep**.

The original and restored files will be kept. The restored file will have the *Restored* prefix.



- Click **Yes** to confirm the operation.

Reference

For more information on 1-click file restore, see [Performing 1-Click File Restore](#) in the Enterprise Manager User Guide.

Performing Self-Restore of VM Guest OS Files

Self-restore allows you to delegate guest OS file restore from backup administrators to users with local Administrator privileges on VMs. Users do not have to wait for backup administrators to recover deleted or modified files and folders.

For self-restore, Veeam Backup Enterprise Manager provides the Self Service File Restore portal. When users log on to the portal, they see only those VMs where they are members of local Administrators group. Other VMs are not visible to the users.

Before You Begin

Consider the following:

- The Enterprise Plus license is installed on the Veeam Backup Enterprise Manager server. You can use a valid trial license or paid license.
- The user account under which you plan to perform self-service restore belongs to a trusted domain or the same domain as the Veeam Backup Enterprise Manager server. Users from untrusted domains cannot use the self-restore capability.
- The user must be a member of local Administrators group on the VM whose guest OS files you plan to restore.

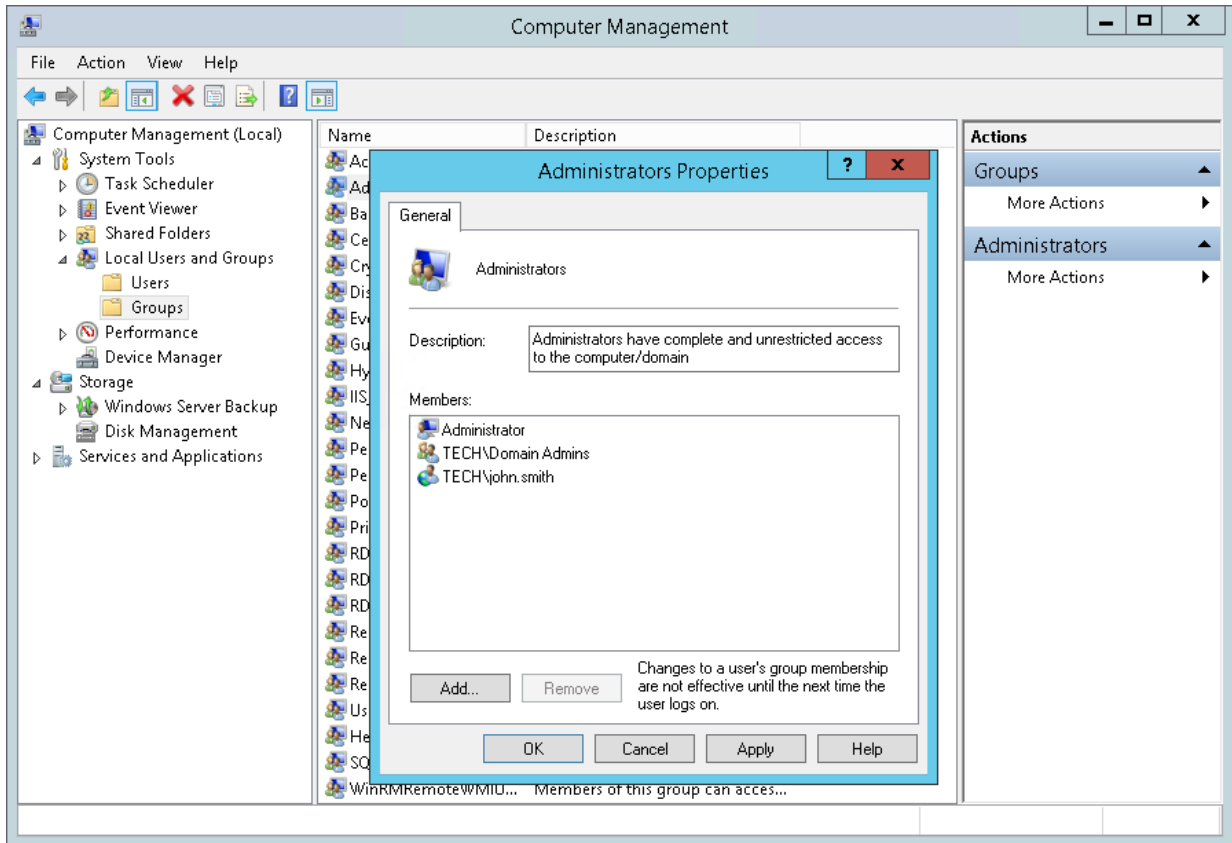
The user has access to restore points created after this user gets local Administrator privileges.

- You have successfully connected backup servers to Veeam Backup Enterprise Manager and collected data from them. For details, see [Adding Veeam Backup Server](#).
- You have at least one successfully created backup with guest file indexing enabled. For details, see [Creating Application-Aware Backup Job](#).

Performing Self-Restore

To restore VM guest OS files using the Veeam Self-Service File Restore portal:

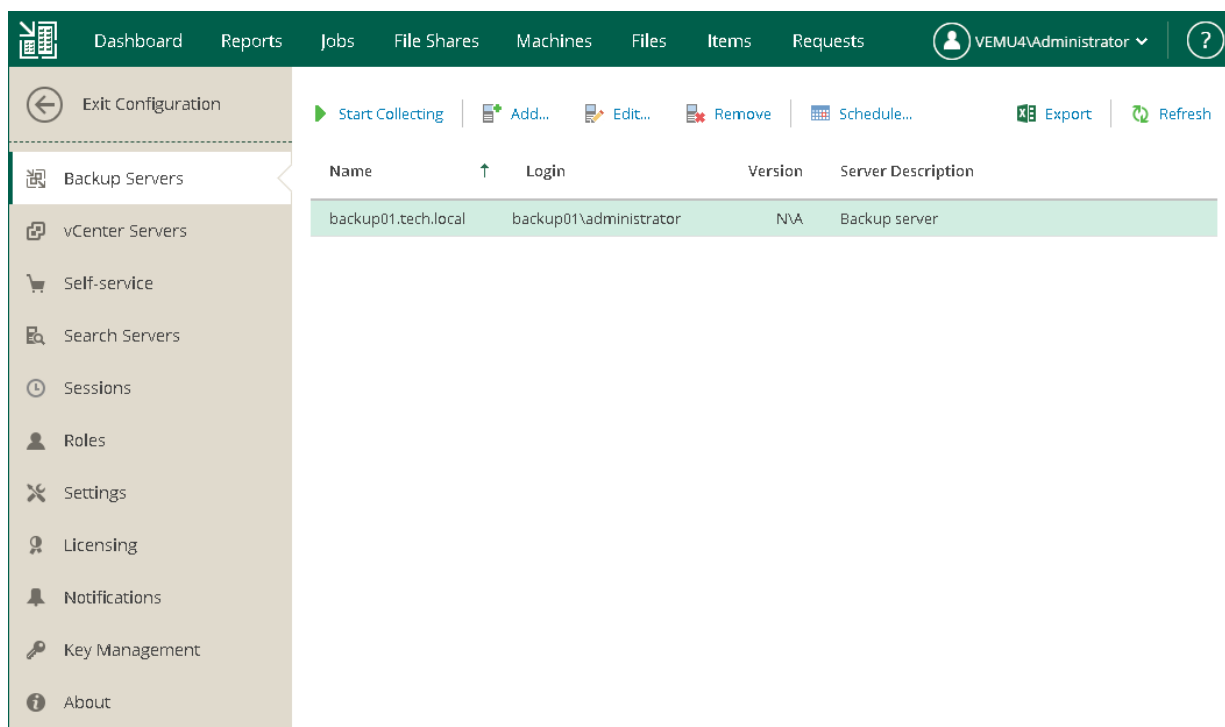
1. Log in to the VM whose guest OS files you plan to restore. Make sure that the user account under which you plan to perform self-service restore is added to the local Administrators group on this VM.



2. Log in to Veeam Backup Enterprise Manager using the user account under which you installed the program.
3. Open the **Jobs** tab and run the job with guest file indexing enabled. You can run the job several times to produce several restore points.
4. At the top right corner, click **Configuration** to open the **Configuration** view.

5. Select the **Backup Servers** tab. In the working area, click **Start Collecting**.

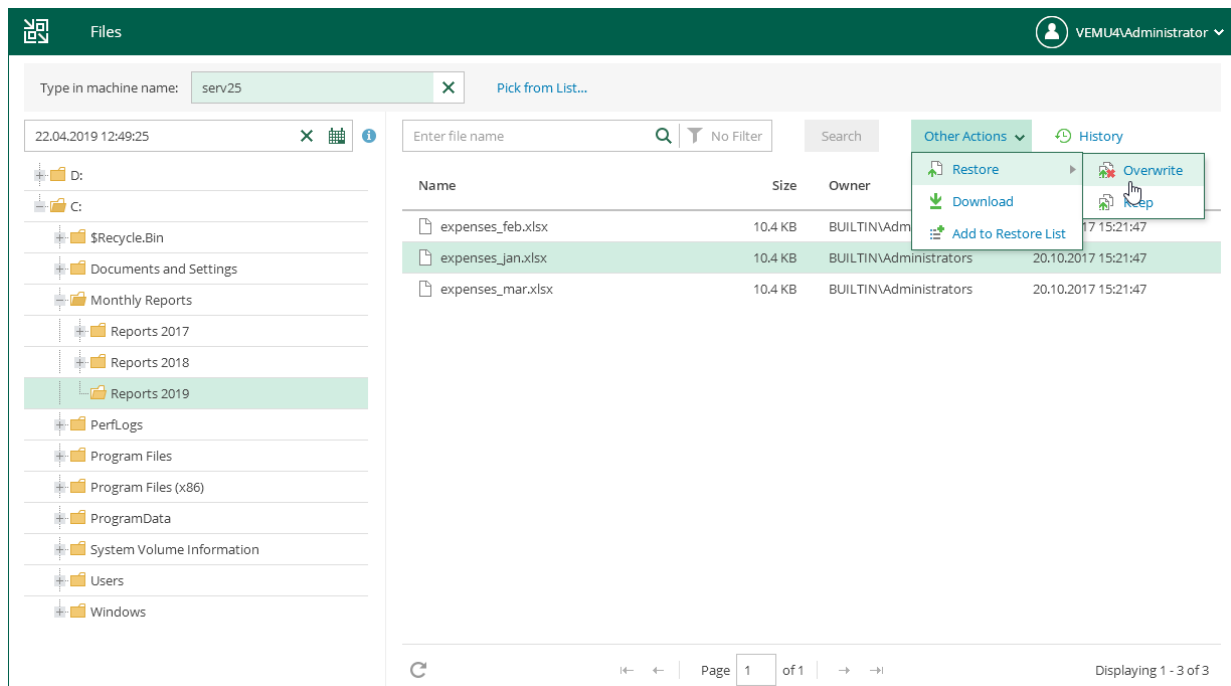
Veeam Backup Enterprise Manager will collect data about jobs from the backup server. To check whether collecting data finished, click the **Sessions** tab. Make sure that the data collection session has completed with the *Success* status.



6. On another machine, open a browser and use the following address to access Veeam Self-Service File Restore portal remotely: https://enterprise_manager_server_IP_address:9443/selfrestore
7. Log in to the portal. Specify the user account with local Administrators privileges on the machine for which you plan to restore files. See the first step.
8. The portal will display only one tab – **Files**. At the top of the working area, click the **pick different machine** link. Select the required VM.
9. Click the field with the calendar icon and choose the restore point from which you want to restore data.
10. At the top of the working area, click **Mount**.

11. Find the necessary file or folder, select it. At the top-right corner, click **Other Actions > Restore > Overwrite**.

The original file will be overwritten.



Reference

For more information on self restore, see [Using Self-Service File Restore Portal to Restore Machine Guest Files](#) in the Enterprise Manager User Guide.

Backing Up Physical Machines

To back up physical machines, Veeam Backup & Replication uses Veeam Agents: Veeam Agent for Microsoft Windows and Veeam Agent for Linux.

You do not need to install, set up and operate Veeam Agent on every machine whose data you want to protect. Instead, you can perform the whole set of deployment, administration, data protection and disaster recovery tasks on computers remotely from the Veeam Backup & Replication console.

How to back up physical machines

To back up physical machines using Veeam Backup & Replication, you must do the following:

1. [Create a protection group](#)

When you create a protection group, you add individual machines or Active Directory containers to the protection group. Veeam Backup & Replication automatically installs agents and other required components on the machines included in the protection group.

2. [Create an Agent backup job](#)

In the Veeam Backup & Replication console, create an agent job that will back up machines included in the protection group.

Reference

For more information on agents, see the following topics:

- [Licensing Requirements](#)
- [Veeam Agent Management User Guide](#)
- [Veeam Agent for Linux 3.0 User Guide](#)
- [Veeam Agent for Microsoft Windows User Guide](#)

Creating Protection Group

In Veeam Backup & Replication, protection groups are logical containers that pool protected computers of a specific type into groups. For example, you can create a protection group for computers of the same type (laptops, workstations or servers) or computers running the same OS type to simplify their management.

You can add individual machines or Active Directory objects that include several machines to the protection group. In this section, you will learn how to create the protection group with Active Directory objects.

TIP:

If you plan to manage only a small number of computers, you can add the necessary computers directly to a Veeam Agent backup job. Veeam Backup & Replication will automatically add such computers to the **Manually Added** protection group. For details, see [Protection Groups](#) in the Veeam Agent Management Guide.

Before You Begin

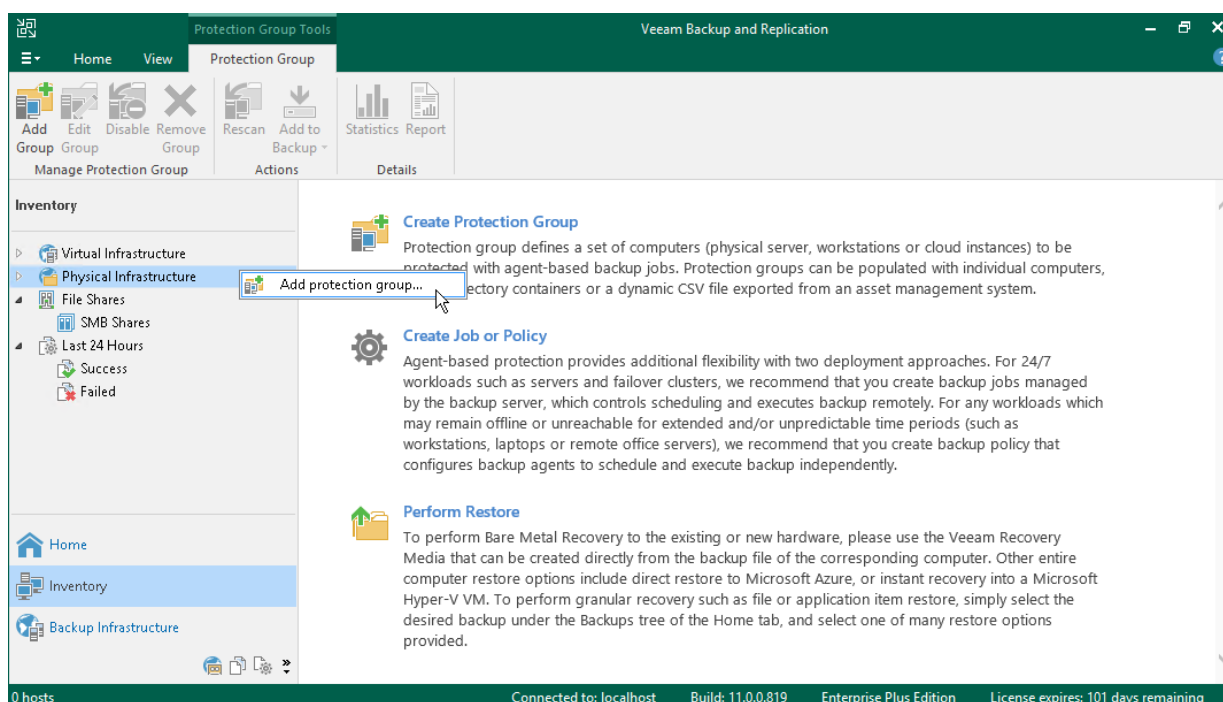
Consider the following:

- Make sure that all computers added to the protection group are powered on and can be accessed over the network.
- If you add an Active Directory container to a protection group, it is not recommended to add a computer that exists in this container to another protection group.

Creating Protection Group

To create a protection group, do the following:

1. In the inventory pane of the **Inventory** view, right-click the **Physical & Cloud Infrastructure** node and select **Add protection group** to launch the **New Protection Group** wizard.



- At the **Name** step of the wizard, specify a name and description for the protection group.

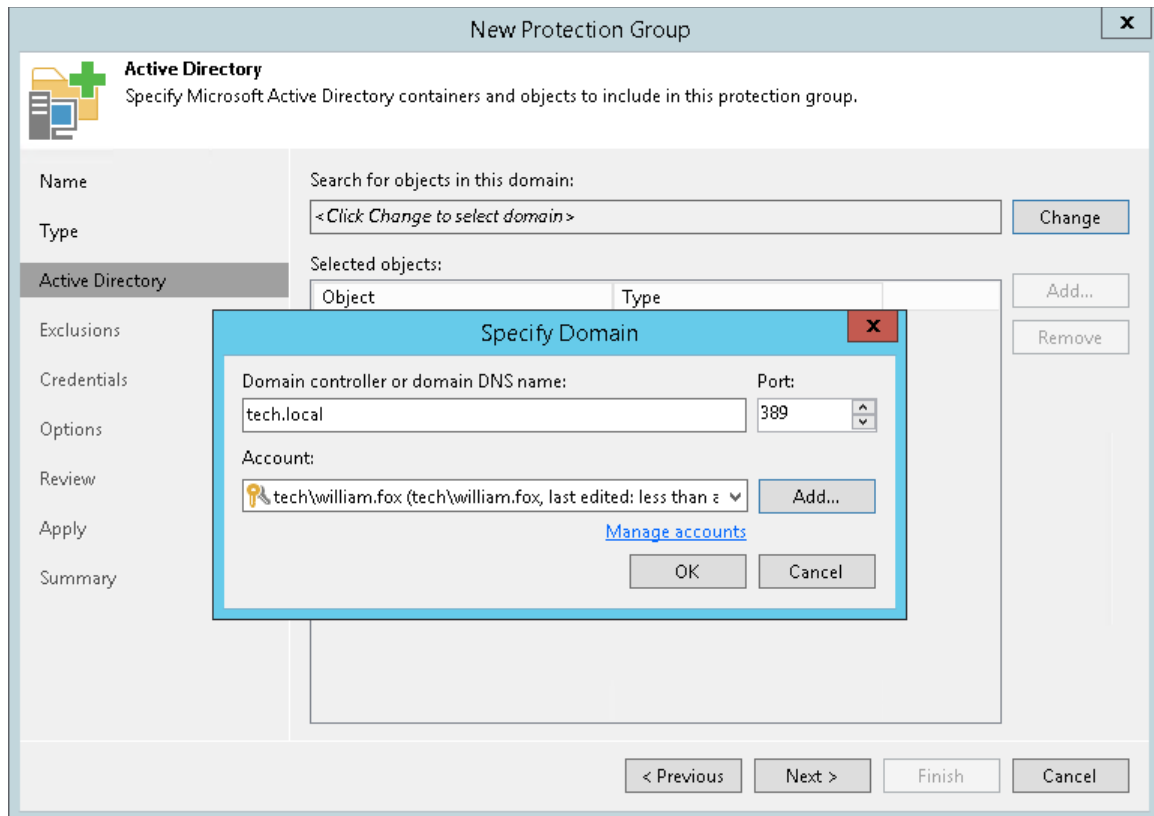
The screenshot shows the 'New Protection Group' wizard window. The title bar is blue with the text 'New Protection Group' and a close button. The main area has a light blue header with a folder icon and the text 'Name' and 'Type in a name and description for this protection group.' Below this is a sidebar with a list of steps: Name, Type, Active Directory, Exclusions, Credentials, Options, Review, Apply, and Summary. The 'Name' step is selected. The main content area has two text boxes: 'Name:' with the text 'Windows Servers' and 'Description:' with the text 'Protection group for Windows based servers'. At the bottom right are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

- At the **Type** step of the wizard, select **Microsoft Active Directory objects**.

Active Directory objects can be the following: entire domain, container, organization unit, group, computer or cluster. Protection groups that include Active Directory objects are dynamic. Veeam Backup & Replication discovers these computers and deploy Veeam Agent on them during the next rescan session.

- At the **Active Directory** step of the wizard, click **Change** near the **Search for objects in this domain field**. In the opened window, do the following:
 - In the **Domain controller or domain DNS name** field, type a name of the domain controller or domain whose objects you want to include in the protection group.
 - In the **Port** field, leave the default value.
 - Near the **Account** field, click **Add** and specify user credentials. This user must be a member of the *DOMAIN\Administrators* group. Click **OK**.

- Click **OK**.



5. Click **Add** near the **Selected objects** field. In the **Add Objects** window, select the necessary Active Directory object and click **OK**.
6. At the **Exclusions** step of the wizard, leave the default settings.
7. At the **Credentials** step of the wizard, specify credentials to connect to computers included in the protection group.

If you want to use the same credentials for all computers in the protection group, select the necessary user account from the **Master account** list. The account must have administrative permissions on all computers that you have added to the protection group.

You can also specify credentials for individual computes. For details, see [Specify Credentials](#) in the Veeam Agent Management Guide.

Credentials

Specify the master account for all hosts in this protection group. You can also customize credentials for individual computers. The specified account must have Local Administrator privileges on the protected computers.

Name: _____

Type: _____

Active Directory: _____

Exclusions: _____

Credentials

Options: _____

Review: _____

Apply: _____

Summary: _____

Master account:
 tech\william.fox (tech\william.fox, last edited: less than a day ago)

[Manage accounts](#)

☐ Use custom credentials for the following objects:

Object	Account
Computers	<Master account>

Click Test Now to validate the specified credentials.

< Previous Next > Finish Cancel

- At the **Options** step of the wizard, leave the default settings.
- At the **Review** step of the wizard, review the components that will be installed. Click **Apply**.
- At the **Apply** step of the wizard, Veeam Backup & Replication creates the configured protection group. Wait for the operation to complete and click **Next**.
- At the **Summary** step of the wizard, select the **Run discovery when I click Finish** check box and click **Finish**.

Reference

For more information on creating protection groups, see [Creating Protection Groups](#) in the Veeam Agent Management Guide.

Creating Veeam Agent Backup Job

To back up physical machines, you must configure a Veeam Agent backup job in the Veeam Backup & Replication console. In Veeam Backup & Replication, you can create Veeam Agent backup jobs of the following types:

- **Backup job**

The backup job runs on the backup server, like VM backup jobs. The backup job is intended for computers that are connected to the backup server. This connection must not be interrupted.

- **Backup policy**

The backup policy describes configuration of individual Veeam Agent backup jobs that run on protected computers. The backup policy is intended for computers that are connected to the backup server. The connection may be interrupted for short periods of time. For example, when you move your laptop from one location to another. Veeam Backup & Replication uses the backup policy as a template and applies settings from the backup policy to Veeam Agents that run on computers specified in the backup policy.

In this guide, we do not detail backup policy. For more information, see [Creating Veeam Agent Backup Policy](#).

Veeam Backup & Replication lets you create backup jobs for Microsoft Windows and Linux computers. In this section, you will learn how to create a Veeam Agent backup job for Microsoft Windows computers. For details on how to create backup jobs for Linux computers, see [Creating Agent Backup Job for Linux Computers](#) in the Veeam Agent Management Guide.

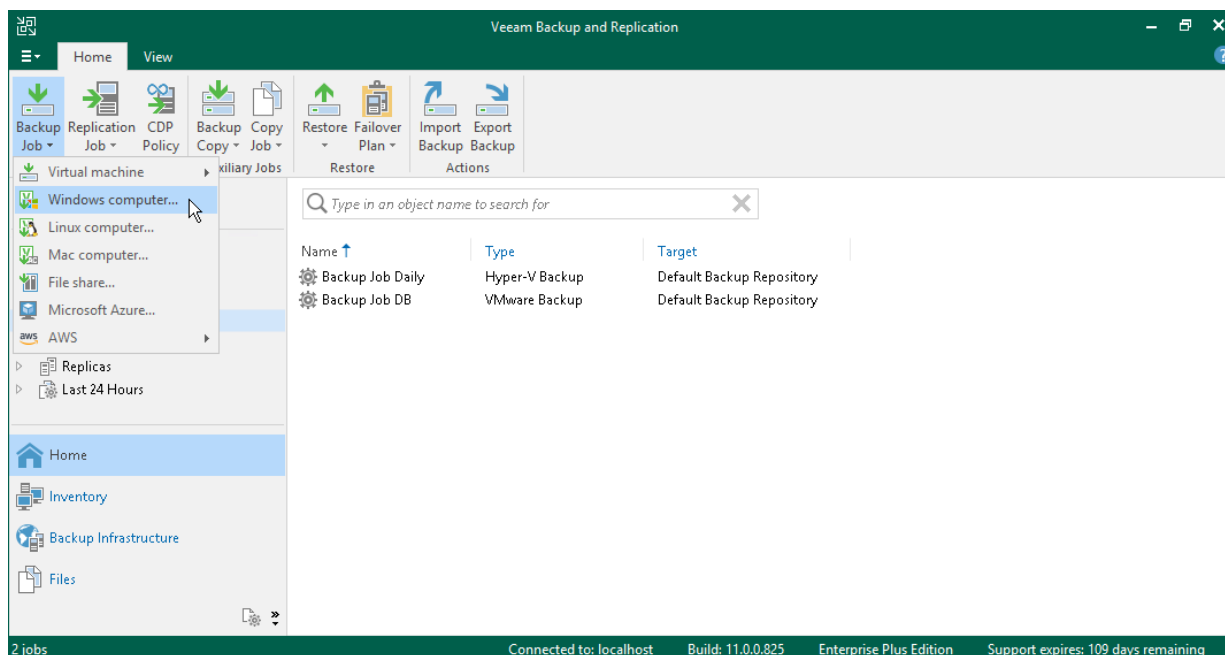
Before You Begin

Consider the following:

- You must have at least one protection group. For details, see [Creating Protection Group](#).
- You can create Veeam Agent backups on a Veeam backup repository only. Other types of target locations are not supported.
- Veeam Agent for Microsoft Windows does not back up data to which symbolic links are targeted. It only backs up the path information that the symbolic links contain. After restore, identical symbolic links are created in the restore destination.

Creating Veeam Agent Backup Job

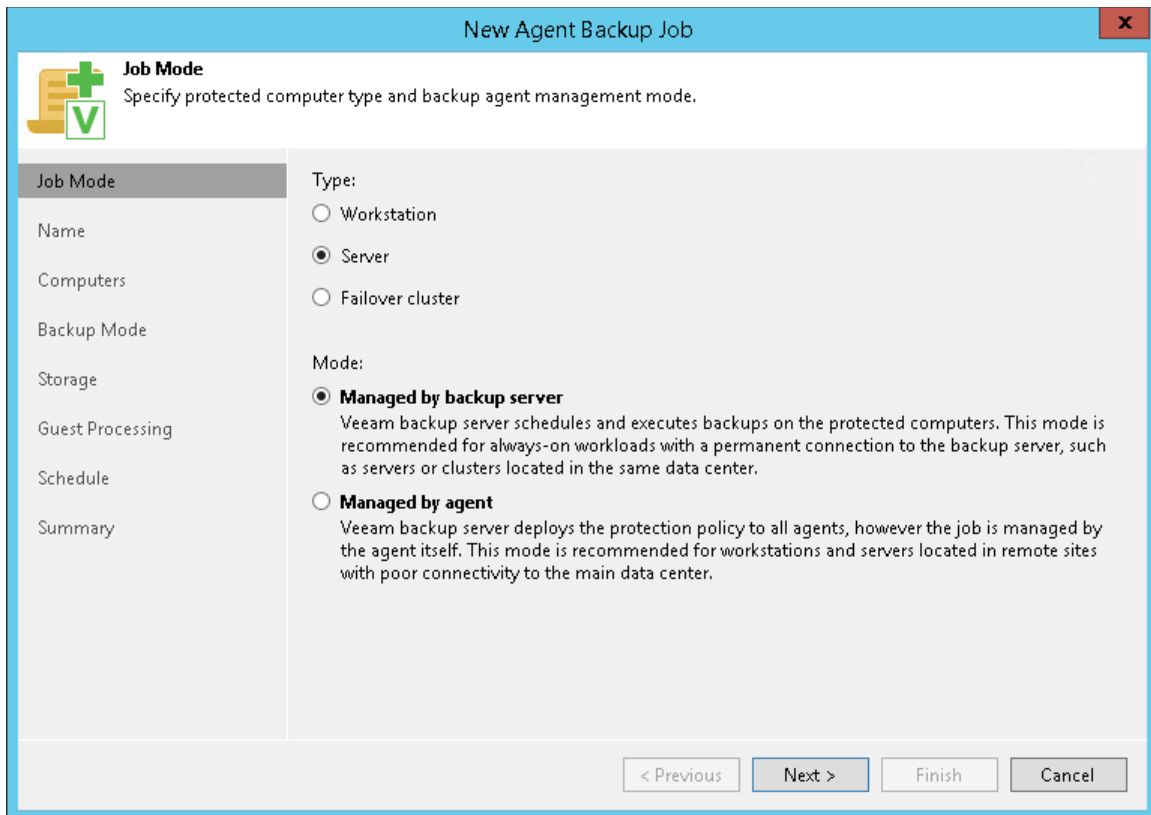
1. On the **Home** tab, click **Backup Job > Windows computer** to open the **New Agent Backup Job** wizard.



2. At the **Job Mode** step of the wizard, specify protection settings for the backup job:
 - In the **Type** list, select **Server** to add to the backup job standalone servers that have permanent connection to the backup server.

- In the **Mode** list, select **Managed by backup server**.

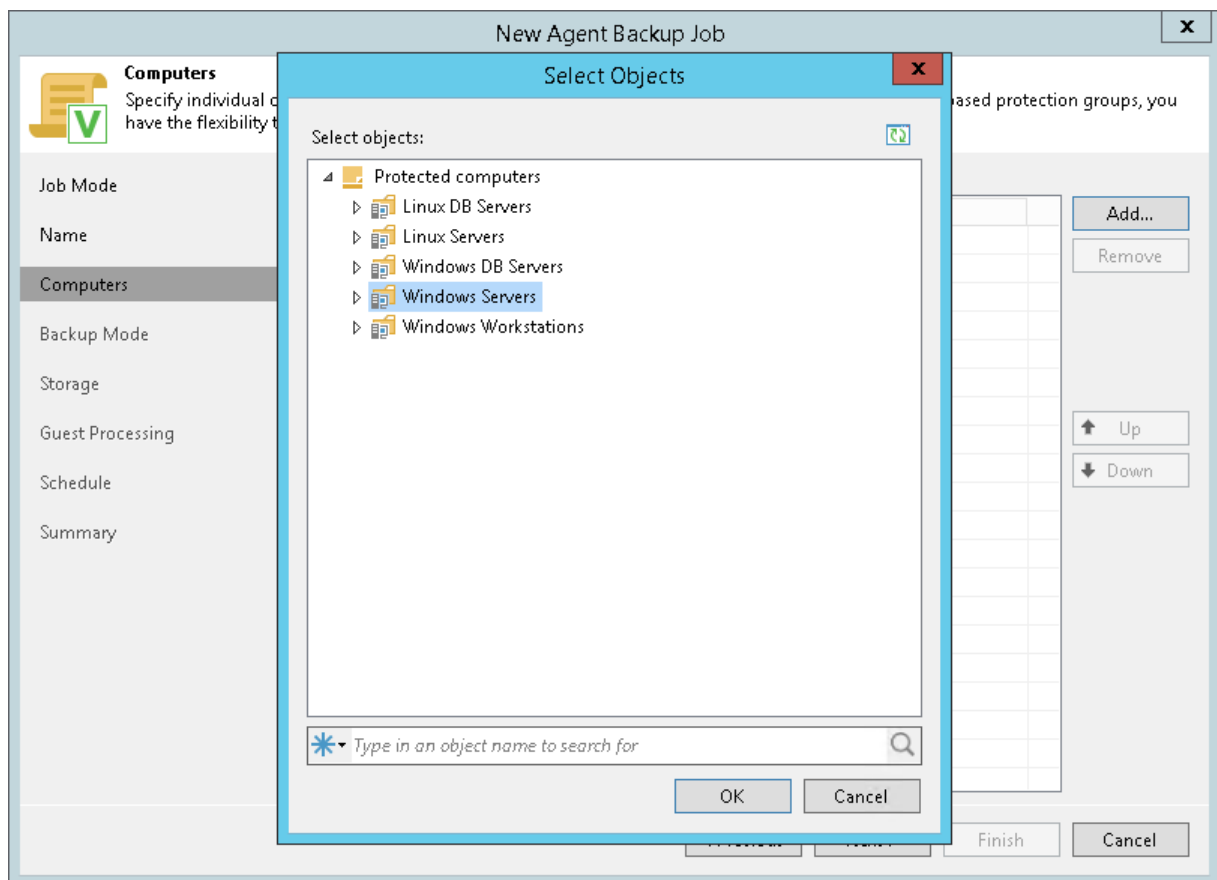
When you create a Veeam Agent backup job managed by the backup server, Veeam Backup & Replication saves the job settings in its database. Veeam Backup & Replication performs all management tasks for the Veeam Agent backup job: starts a job upon the defined schedule, allocates backup infrastructure resources and so on.



The screenshot shows the 'New Agent Backup Job' wizard window. The title bar is blue with the text 'New Agent Backup Job' and a close button. The main area has a light blue header with a green plus icon and the text 'Job Mode' and 'Specify protected computer type and backup agent management mode.' Below this is a list of steps on the left: 'Job Mode' (selected), 'Name', 'Computers', 'Backup Mode', 'Storage', 'Guest Processing', 'Schedule', and 'Summary'. The main content area is divided into two sections: 'Type:' and 'Mode:'. Under 'Type:', there are three radio buttons: 'Workstation', 'Server' (selected), and 'Failover cluster'. Under 'Mode:', there are two radio buttons: 'Managed by backup server' (selected) and 'Managed by agent'. The 'Managed by backup server' option has a description: 'Veeam backup server schedules and executes backups on the protected computers. This mode is recommended for always-on workloads with a permanent connection to the backup server, such as servers or clusters located in the same data center.' The 'Managed by agent' option has a description: 'Veeam backup server deploys the protection policy to all agents, however the job is managed by the agent itself. This mode is recommended for workstations and servers located in remote sites with poor connectivity to the main data center.' At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

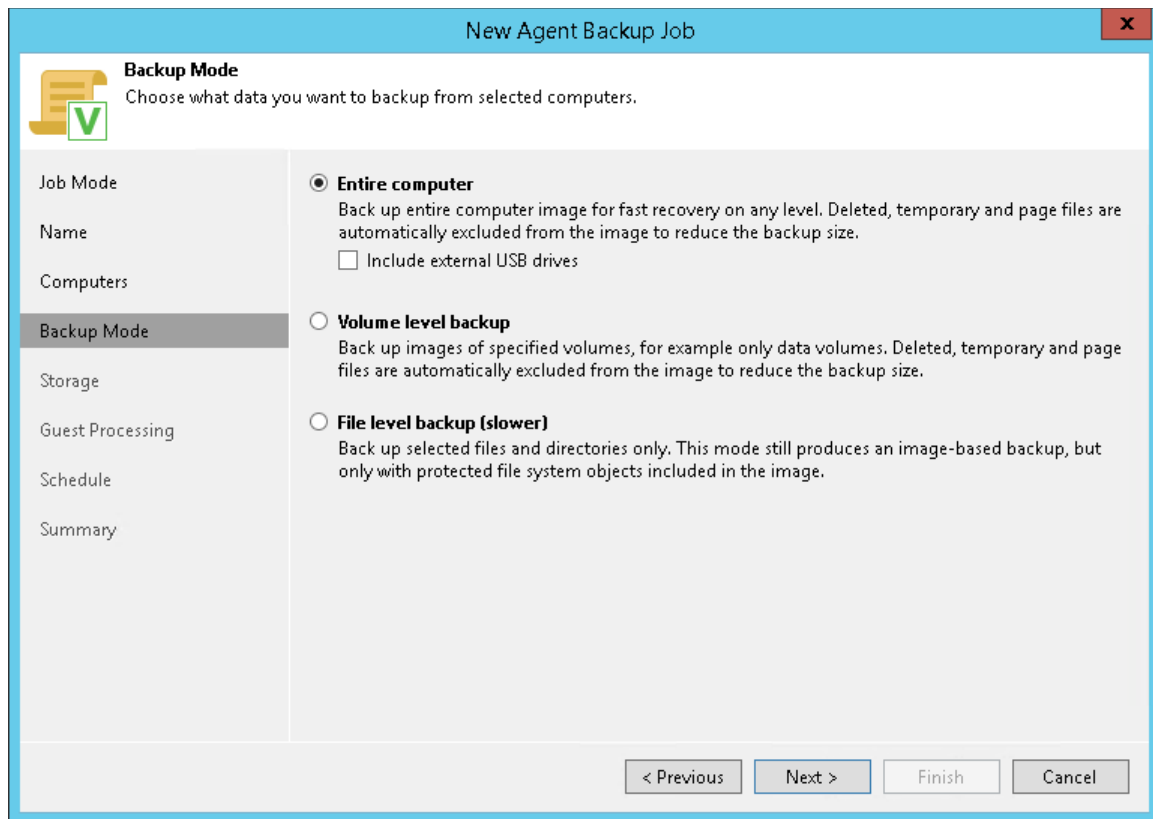
3. At the **Name** step of the wizard, specify a name and description for the backup job.

4. At the **Computers** step of the wizard, click **Add** and select one or several protection groups and/or computers in the list. Click **OK**.



- At the **Backup Mode** step of the wizard, select **Entire computer**.

When you restore data from such a backup, you are able to recover the entire computer image as well as data on specific computer volumes: files, folders and application data.



- At the **Storage** step of the wizard, select the backup repository where you want to store your backups. For other settings, leave the default values.
- At the **Guest Processing** step of the wizard, leave the default settings.
- At the **Schedule** step of the wizard, define scheduling settings for the job.
- At the **Summary** step of the wizard, select the **Run the job when I click Finish** check box and click the **Finish** button.
- In the inventory pane of the **Home** view, expand the **Last 24 Hours** node to see the created job.

Reference

For more information on creating Veeam Agent backup jobs for Microsoft Windows computers, see [Creating Agent Backup Job for Windows Computers](#) in the Veeam Agent Management Guide.

Restoring Data of Physical Machines

Veeam Backup & Replication allows you to restore data of physical machines. You can perform the following data restore tasks with Veeam Agent backups:

- [Restoring Veeam Agent Backup to vSphere VM](#)
- [Restoring Veeam Agent Backup to Hyper-V VM](#)
- [Restoring to Microsoft Azure](#)
- [Restoring to Amazon EC2](#)
- [Restoring to Google Cloud Platform](#)
- [Restoring Volumes](#)
- [Restoring Files and Folders](#)
- [Restoring Application Items](#)

In this section, you will learn how to restore computer files and folders. For more information on other restore processes, follow the links in the list.

Restoring Files and Folders

The procedure of file-level restore from a Veeam Agent backup is similar to the same procedure for a VM backup. The difference is that you select a Veeam Agent backup instead of a VM backup in the **File Level Restore** wizard. To learn more, see [Restoring VM Files](#).

