



Veeam Backup & Replication for VMware

Version 8.0

Evaluator's Guide

May, 2015

#1 for Virtualization™
Management and Data Protection

Microsoft Partner
Gold Application Development
Gold Management and Virtualization



© 2015 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.

Important! Please read the End User Software License Agreement before using the accompanying software program(s). Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

CONTENTS

GETTING STARTED	4
SYSTEM REQUIREMENTS.....	6
BACKUP INFRASTRUCTURE SETUP	7
INSIGHT INTO VEEAM BACKUP INFRASTRUCTURE	7
EXERCISE LIST	10
Installing Veeam Backup & Replication	11
Connecting Virtual Infrastructure Servers	13
Configuring a Backup Proxy.....	18
Configuring a Backup Repository	24
DATA PROTECTION AND DISASTER RECOVERY TASKS	31
EXERCISE LIST	31
Performing Backup	32
Backing up and Restoring Microsoft SQL Server Databases	52
Performing Full VM Restore.....	59
Restoring Guest OS Files.....	63
Restoring VM Virtual Disk.....	70
Restoring VM Files	73
Creating Backup Copy.....	76
Performing Replication.....	81
Failing Over a VM Replica.....	98
Undoing Failover	101
Failing Back to the Primary VM.....	103
Committing Failback	106
DISTRIBUTED BACKUP INFRASTRUCTURE MANAGEMENT.....	108
EXERCISE LIST	108
Installing and Configuring Veeam Backup Enterprise Manager	109
Searching for Guest OS Files and Performing 1-Click Restore	113
Performing Self-Restore of VM Guest OS Files	115
Cloning and Editing Jobs	120
Restoring Data from Encrypted Backup File Without Password.....	122

GETTING STARTED

About Veeam Backup & Replication

Veeam® Backup & Replication™ is a data protection and disaster recovery solution for virtual environments of any size and complexity. Veeam Backup & Replication provides fast, flexible, and reliable recovery of virtualized applications and data. It unifies backup and replication in a single solution, increases the value of backup and reinvents data protection for VMware vSphere and Microsoft Hyper-V virtual environments. Veeam Backup & Replication supports your entire virtual infrastructure with industry leading features such as instant file-level recovery and streamlined VM recovery, scalability, 2-in-1 backup & replication, built-in de-duplication, centralized management and many more.

About This Guide

This guide will help you become familiar with Veeam Backup & Replication and evaluate its capabilities. This guide explains the primary features of Veeam Backup & Replication and will help you begin using the product, regardless of your previous experience with the product.

Intended Audience

The document is intended for IT professionals who are looking to deploy Veeam Backup & Replication to protect their VMware virtual environment. This guide will be of interest to novices to the product and to VMware administrators, consultants and analysts who have used previous versions of Veeam Backup & Replication and want to evaluate new features in Veeam Backup & Replication.

Document Structure

The guide provides a set of self-guided evaluation exercises that you should follow to familiarize yourself with Veeam Backup & Replication. Each evaluation exercise provides a short feature overview, an evaluation case and evaluation procedure, as well as comments on validation of the exercise results.

The guide is comprised of four major parts:

- **Reference Environment.** This section describes requirements for a test lab that you need to provision to successfully perform evaluation tasks and outlines a sample test lab used in this evaluator's guide.
- **Backup Infrastructure Setup.** This section describes Veeam Backup & Replication installation and deployment steps that need to be completed before you can start evaluating the product.
- **Data Protection and Disaster Recovery.** This section describes evaluation cases covering the most typical data protection and disaster recovery tasks that you can perform with Veeam Backup & Replication.
- **Distributed Backup Infrastructure Management.** This section describes how you can manage a distributed backup infrastructure using Veeam Backup Enterprise Manager.

Help and Support

This guide provides a high-level overview of primary features in Veeam Backup & Replication and should be regarded as a supplement to existing technical documentation. The complete set of documentation can be found on the *Resources* web page at: www.veeam.com/documentation-guides-datasheets.html.

For technical support and assistance, use the following resources:

- **Veeam Community Forums:** <http://forums.veeam.com>
- **Customer Support Portal:** cp.veeam.com. Should you have a product issue, suggestion or question, please visit our Customer Center to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Document Revision History

Revision #	Date	Change Summary
Revision 1	06/11/2014	Initial version of the document for the Veeam Backup & Replication 8.0.
Revision 2	06/05/2014	Added scenarios: <ul style="list-style-type: none">▪ Backing up and Restoring Microsoft SQL Server Databases,▪ Performing Self-Restore of VM Guest OS Files and▪ Restoring Data from Encrypted Backup File Without Password. System Requirements section updated.

SYSTEM REQUIREMENTS

To learn about system requirements, required ports and permissions, see [Release Notes](#) for Veeam Backup & Replication 8.0 or <http://helpcenter.veeam.com/backup/80/vsphere/requirements.html> .

BACKUP INFRASTRUCTURE SETUP

This section describes Veeam Backup installation and deployment steps that need to be completed before you can start evaluating the product.

Insight into Veeam Backup Infrastructure

The backup infrastructure is a framework that comprises a set of components needed to perform data protection and disaster recovery tasks. A typical Veeam backup infrastructure includes the following components:

- **Veeam backup server:** a physical or virtual machine running Veeam Backup & Replication. The Veeam backup server performs the role of the main management component: it is the configuration and control center of the backup infrastructure.
- **Virtual infrastructure servers:** ESX(i) hosts used as the source and the target for backup and replication.
- **Backup proxy:** a “data mover” component that retrieves VM data from the source storage, processes it and transfers to the destination.
- **Backup repository:** a location used to store backup files, VM copies and auxiliary replica files.

To perform evaluation exercises, you can use a simple deployment scenario or a distributed deployment scenario (recommended). The choice depends on your test lab environment. If you have a very small test lab and there is no possibility of allocating dedicated servers that will perform the roles of a backup proxy and backup repository, you can use a simple deployment scenario to evaluate the product. However, whenever possible, it is recommended that you use a distributed deployment scenario. This will help you gain insight into the new architecture of the backup infrastructure, evaluate its benefits and increase your experience in deploying backup infrastructure components.


Simple deployment

In a simple deployment scenario, components of the backup infrastructure perform several roles at a time. The backup infrastructure includes the following components:

- Veeam backup server
- Source ESX(i) host
- Target ESX(i) host (used for a replication scenario)

In the simple deployment scenario, the Veeam backup server performs three roles:

- It is the “control center” of the backup infrastructure, coordinating job performance and other administrative activities.
- It is the default backup proxy. When you perform backup, replication or VM copy, VM data is processed directly on the Veeam backup server and then moved to the target. All services necessary for backup proxy functioning are installed locally on the Veeam backup server.
- It is the default backup repository: all backups, VM copies and auxiliary replica backup files are stored on the Veeam backup server, in the `Backup` folder on the volume with the most amount of free disk space.



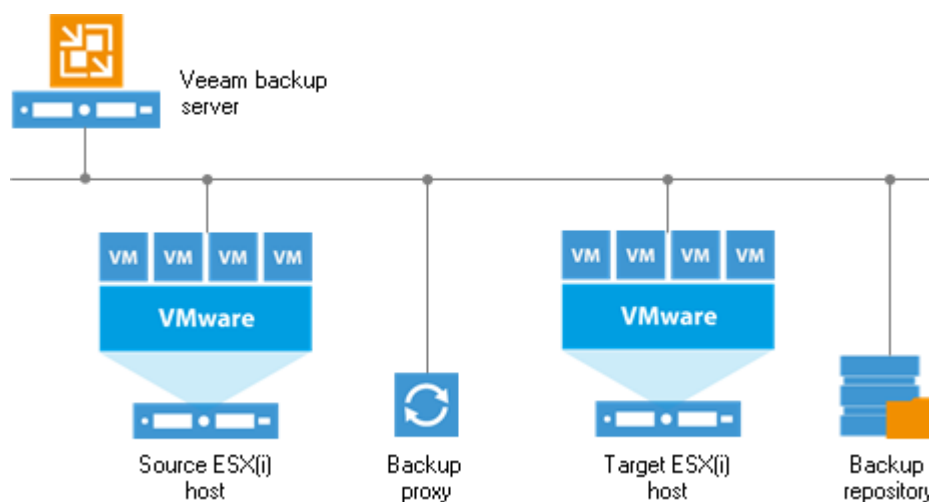
Distributed deployment

In the distributed deployment scenario, Veeam backup server, backup proxy and backup repository roles are assigned to dedicated servers. This type of deployment enables efficient data transfer and use of resources in your backup infrastructure because you move the data processing load from the Veeam backup server to a dedicated backup proxy and store data to a dedicated backup repository.

Depending on your production environment and the backup and replication scenarios you plan to use, the distributed backup infrastructure may include a number of dedicated backup proxies and backup repositories, both onsite and offsite, controlled by a single Veeam backup server. Use of multiple backup proxies and repositories allows for dynamic assignment of jobs and intelligent load balancing across the backup infrastructure.

To evaluate the product using this guide, it is sufficient to deploy one backup proxy and one backup repository locally in your test lab. Therefore, the backup infrastructure will consist of the following components:

- Veeam backup server
- Backup proxy
- Backup repository
- Virtual infrastructure servers



Exercise List

To deploy the backup infrastructure, perform the following exercises:

Exercise	Description	Time Estimates
Installing Veeam Backup & Replication	Install Veeam Backup & Replication on a physical or virtual machine.	5-10 minutes
Connecting virtual infrastructure servers	Add to the Veeam Backup & Replication console the hosts that you want to use as the source and target for backup and replication.	5-7 minutes
Configuring a backup proxy	Add to the Veeam Backup & Replication console the server that you want to use as a backup proxy and assign the role of the backup proxy to it. This exercise should be performed only if you decide to follow the distributed deployment scenario.	5-7 minutes
Configuring a backup repository	Add to the Veeam Backup & Replication console the server or that you want to use as a backup repository and assign the role of a backup repository to it. This exercise should be performed only if you decide to follow the distributed deployment scenario.	5-7 minutes

Installing Veeam Backup & Replication

You should install Veeam Backup & Replication on a Windows-based machine, either physical or virtual, that meets the system requirements.

Evaluation Case

In this exercise, you will install Veeam Backup & Replication. By installing Veeam Backup & Replication, you configure the Veeam backup server — the core component in the backup infrastructure that controls all other components.

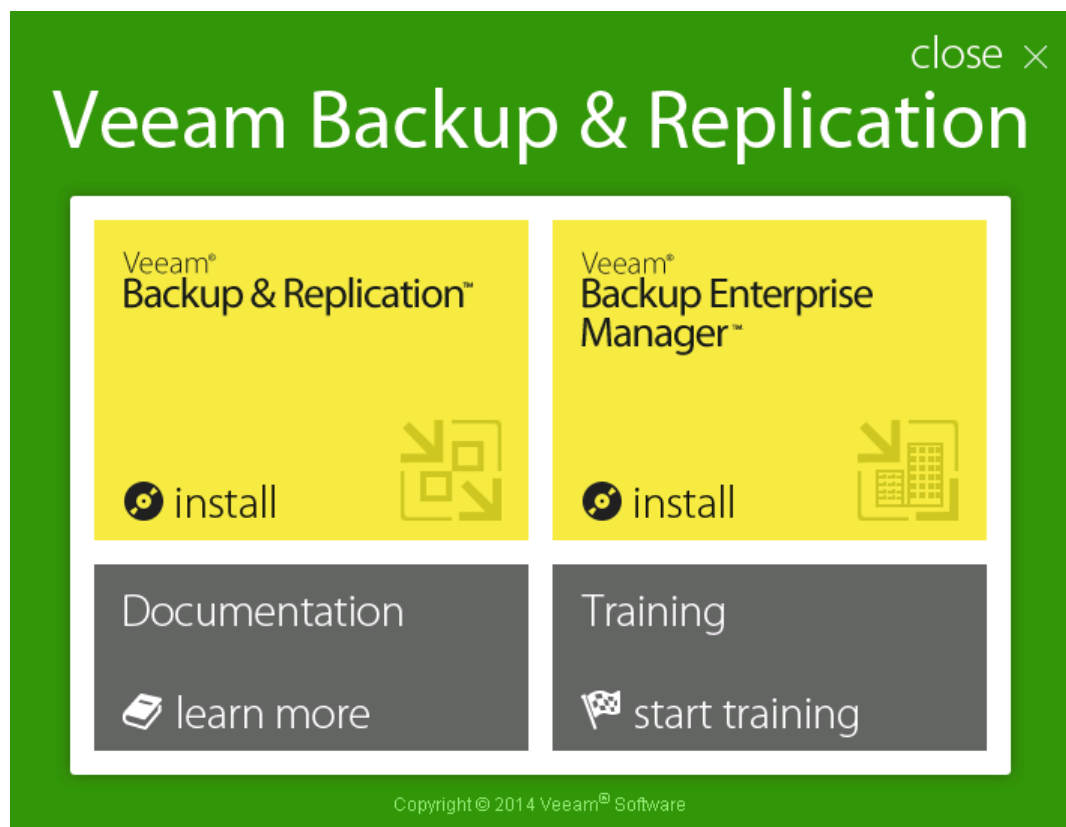
Prerequisites

- Your test lab must meet the system requirements.
- You must have *Local Administrator* permissions on the server where Veeam Backup & Replication will be installed.
- You must have a valid trial license or full paid license for Veeam Backup & Replication.
- Veeam Backup & Replication uses a Microsoft SQL Server instance installed locally or remotely and requires .NET Framework 4.0. In case you do not have these, the setup wizard will automatically install Microsoft SQL Server 2012 Express and .NET Framework 4.0.

Procedure

To install Veeam Backup & Replication:

1. Download the latest version of Veeam Backup & Replication for your version of Microsoft Windows from www.veeam.com/downloads.
2. Mount the installation image using disk image emulation software, or burn the downloaded ISO image file to a blank CD/DVD or. If you are installing Veeam Backup & Replication on a virtual machine, use built-in tools of the virtualization management software to mount the installation image to the virtual machine.
3. After you mount or insert the disk with Veeam Backup & Replication setup, Autorun will open a splash screen with installation options. If Autorun is not available or disabled, run the `Setup.exe` file from the CD/DVD disk. Alternatively, you can right-click the new disk in **My Computer** and select **Execute Veeam Backup & Replication Autorun** or simply double-click the new disk to launch the splash screen.



4. Click the **Install** link in the **Veeam Backup & Replication** section of the splash screen. On the **Welcome** step of the wizard, click **Next** to start the installation.
5. At the **Provide License** step of the wizard, specify the path to the license file you obtained after you downloaded the product from the web site.
6. At the **Default Configuration** step of the wizard, specify if you want to use default installations settings or specify installation settings on your own. In the latter case, the setup wizard will include additional steps for configuring the necessary settings.
7. If you have selected to specify custom installation settings, pass through the next wizard steps and specify the required settings manually. If you have selected to use default installations settings, click **Install** to begin the installation process.
8. When the installation process is complete, click **Finish** to close the wizard.

Validation

Select **Programs > Veeam > Veeam Backup & Replication** from the **Start** menu to make sure that Veeam Backup & Replication has been installed successfully.

Connecting Virtual Infrastructure Servers

To connect virtual infrastructure hosts to the Veeam backup server, you need to add them to the Veeam Backup & Replication console.

You can connect ESX(i) hosts that will be used as source and target for backup and replication. If you plan to use an ESX(i) host that is a part of the vCenter Server hierarchy, it is recommended that you add the vCenter Server instead of separate hosts as this will provide more flexibility at work.

Evaluation Case

In this exercise, you will connect to the Veeam backup server vCenter Server and/or ESX(i) hosts that you plan to use as the source and target for backup and replication.

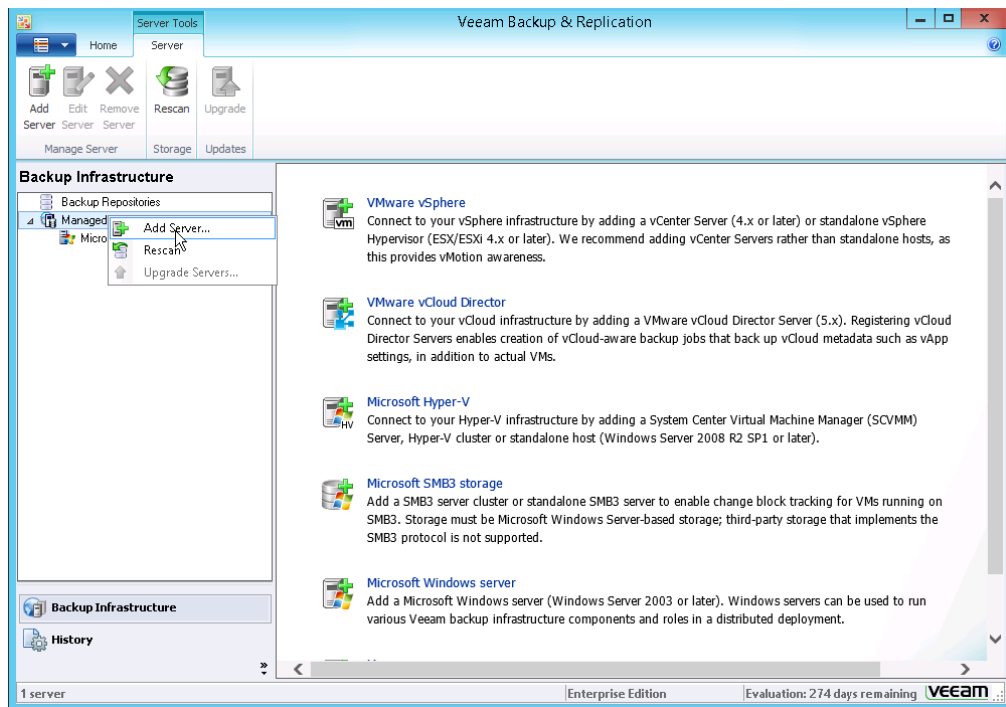
Prerequisites

- You must decide which hosts you want to use as source and target for backup and replication.
- The Veeam backup server must have access to all hosts you plan to use for backup and replication.
- Make sure that all necessary ports are open.

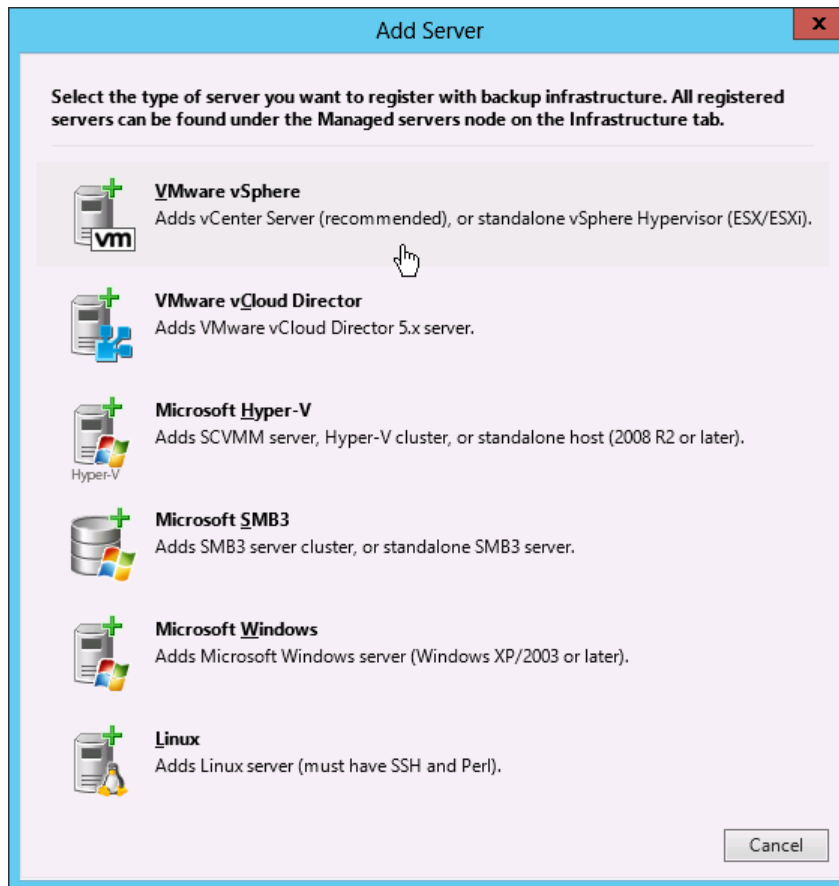
Procedure

To add vCenter Server or ESX(i) hosts, do the following:

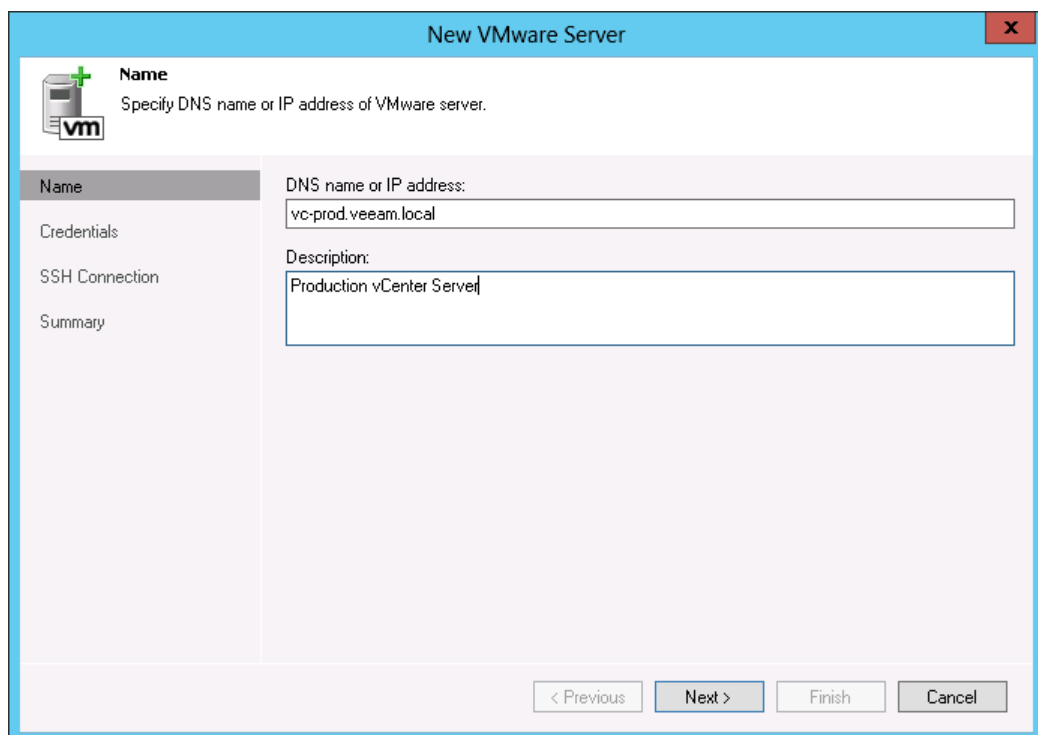
1. Run Veeam Backup & Replication: from the **Start** menu, select **Programs > Veeam > Veeam Backup & Replication**.
2. Open the **Backup Infrastructure** view.
3. Right-click the **Managed servers** node in the inventory pane and select **Add Server**.



4. In the **Add Server** window, select **VMware vSphere**.



5. Specify a DNS name or the IP address of the server.



6. Enter credentials for the user account with administrator access permissions to the added server: click **Add** on the right of the **Credentials** field and specify the user name and password to connect to the added server. For vCenter Server, credentials must be specified in the *DOMAIN\USERNAME* format.

The screenshot shows the 'New VMware Server' wizard with the 'Credentials' step selected. A modal dialog titled 'Credentials' is open, allowing the user to enter the following information:

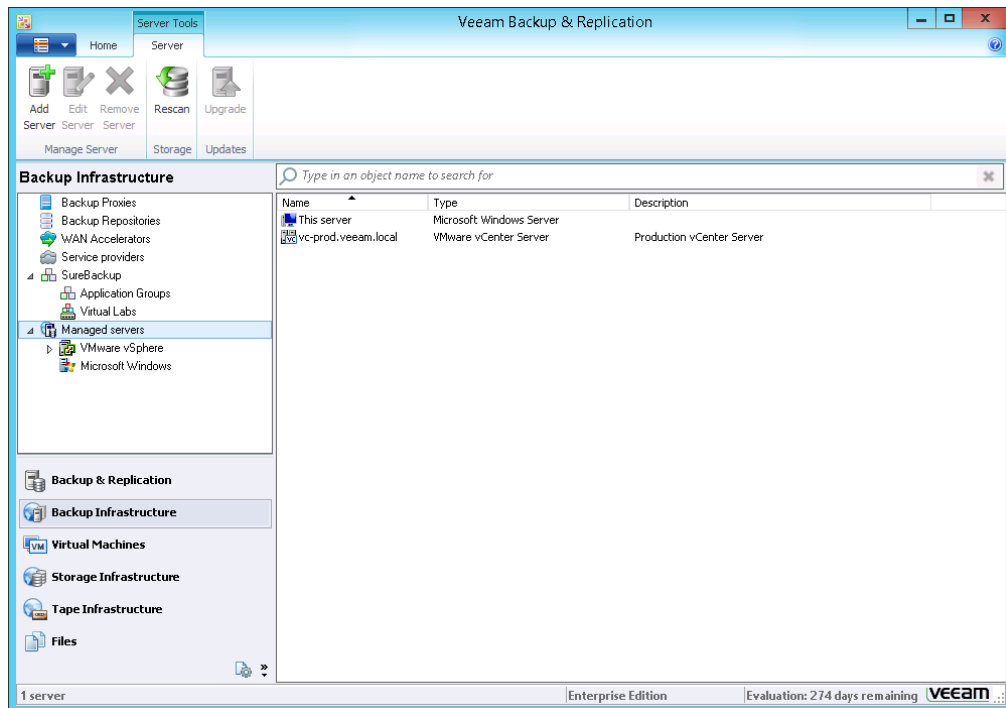
- Username:** VEEAM\Administrator
- Password:** [Masked with dots]
- Description:** VEEAM\Administrator

Buttons for 'Browse...', 'Add...', 'OK', and 'Cancel' are visible. Below the modal, the 'Port' is set to 443. At the bottom of the wizard, there are buttons for '< Previous', 'Next >', 'Finish', and 'Cancel'.

7. Click **Next**, then click **Finish**.
8. Repeat the procedure for all servers that you want to add.

Validation

1. Open the **Backup Infrastructure** view.
2. Click the **Managed servers** node in the inventory pane.
3. Make sure that the added vCenter Server or ESX(i) host is available in the working area.



Configuring a Backup Proxy

Insight into Backup Proxy

In the backup infrastructure, a backup proxy is a “data mover” component responsible for handling jobs and transferring VM data. During backup, replication or VM copy activities, the backup proxy retrieves VM data from the source datastore, processes it and transfers to the destination storage. The backup proxy is also used to write data back to the source datastore during full VM restore and VM disk restore.

The role of a backup proxy should be assigned to a Microsoft Windows machine (physical or virtual) that meets the system requirements.

To retrieve VM data from the source datastore and write it back to the datastore during restore, the backup proxy can use one of the following modes:

- **Direct SAN Access** — in this mode, the backup proxy copies VM data directly through the SAN, bypassing the LAN.
- **Virtual Appliance** — in this mode, the backup proxy retrieves VM data directly from storage through the ESX(i) I/O stack, instead of transporting it over the network.
- **Network** — in this mode, the backup proxy retrieves VM data from the ESX(i) host through the LAN.

You can explicitly select the transport mode that a backup proxy will use or let Veeam Backup & Replication automatically choose the most appropriate mode.

The backup proxy should provide an optimal route for VM data traffic. When configuring a backup proxy, you need to analyze the connection between the source datastore and a backup proxy. Consider the following recommendations:

- If you use FC SAN, assign the role of a backup proxy to a physical server with direct FC access to the SAN to enable LAN-free data retrieval.
- Otherwise, assign the role of a backup proxy to a VM on an ESX(i) host connected the source storage. This type of a backup proxy enables LAN-free data retrieval during backup and replication and LAN-free writing of data back to the storage during full VM restore. Moreover, it does not require you to provision a dedicated physical server.

Evaluation Case

In this exercise, you will configure a backup proxy. Configuration of a backup proxy is performed in two stages:

1. You should add to the Veeam Backup & Replication console a server that will perform the role of a backup proxy.
2. You should assign the role of a backup proxy to the added server.

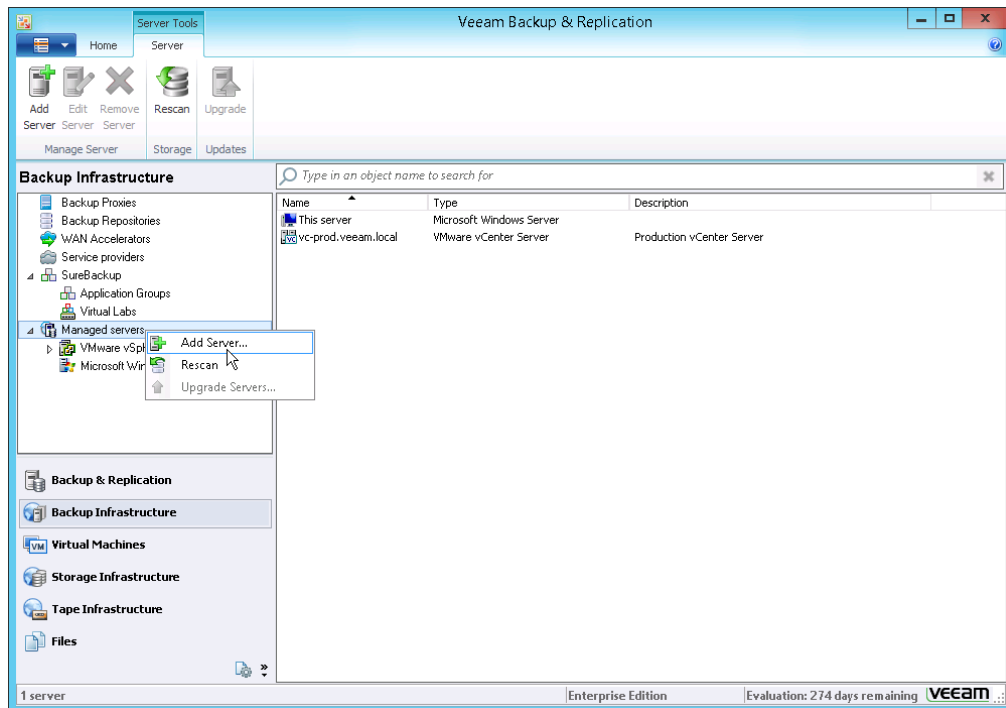
Prerequisites

- You should have a dedicated Microsoft Windows physical or virtual machine that will perform the role of a backup proxy.
- The machine must have access to the Veeam backup server, source datastore and backup repository (for the backup and restore scenarios) and additionally access to the target ESX(i) host (for the replication scenario).
- Make sure that all necessary ports are open.

Procedure

To configure a backup proxy:

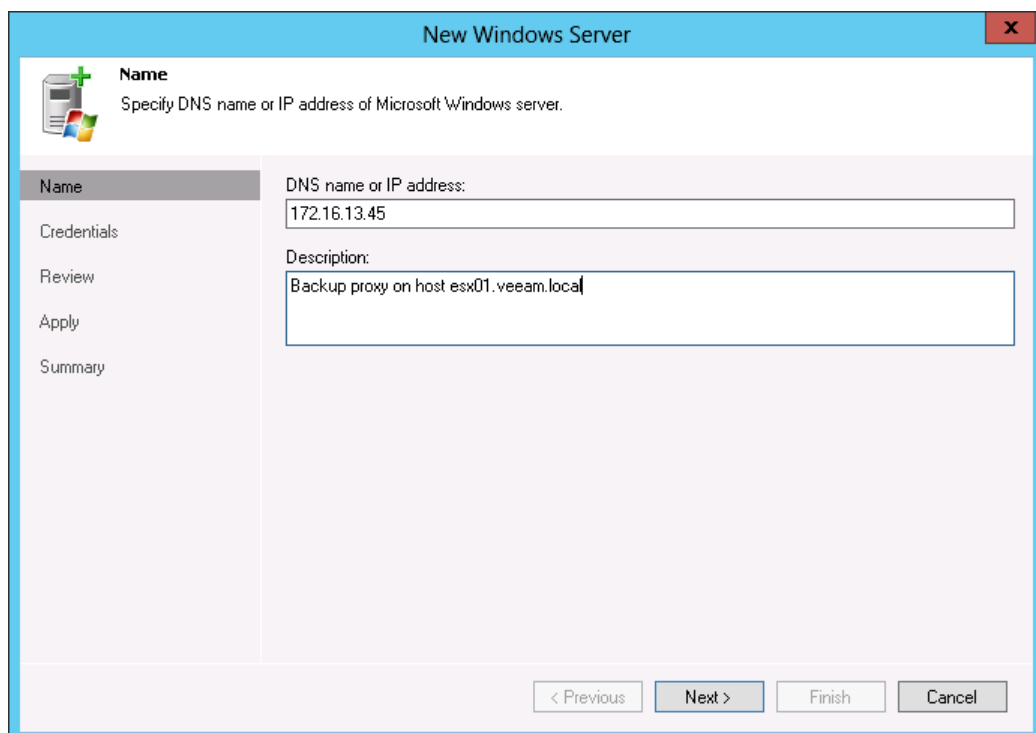
1. Open the **Backup Infrastructure** view.
2. Right-click the **Managed servers** node in the inventory pane and select **Add Server**.



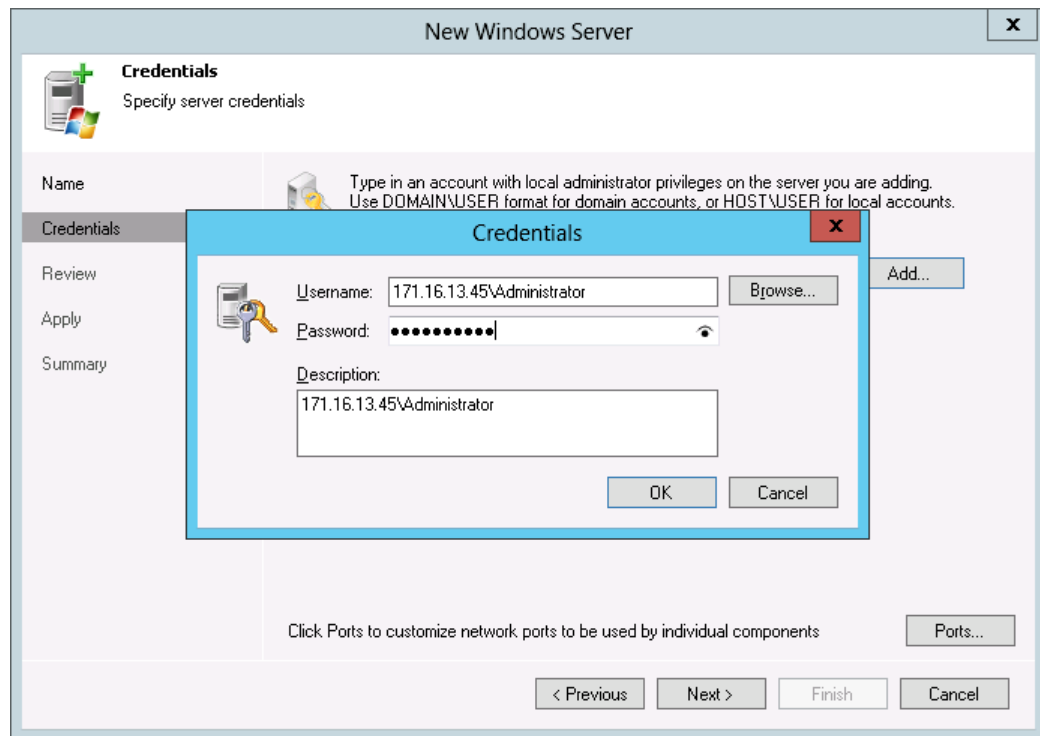
3. In the **Add Server** window, select **Microsoft Windows**.



4. Specify a DNS name or the IP address of a Microsoft Windows server.



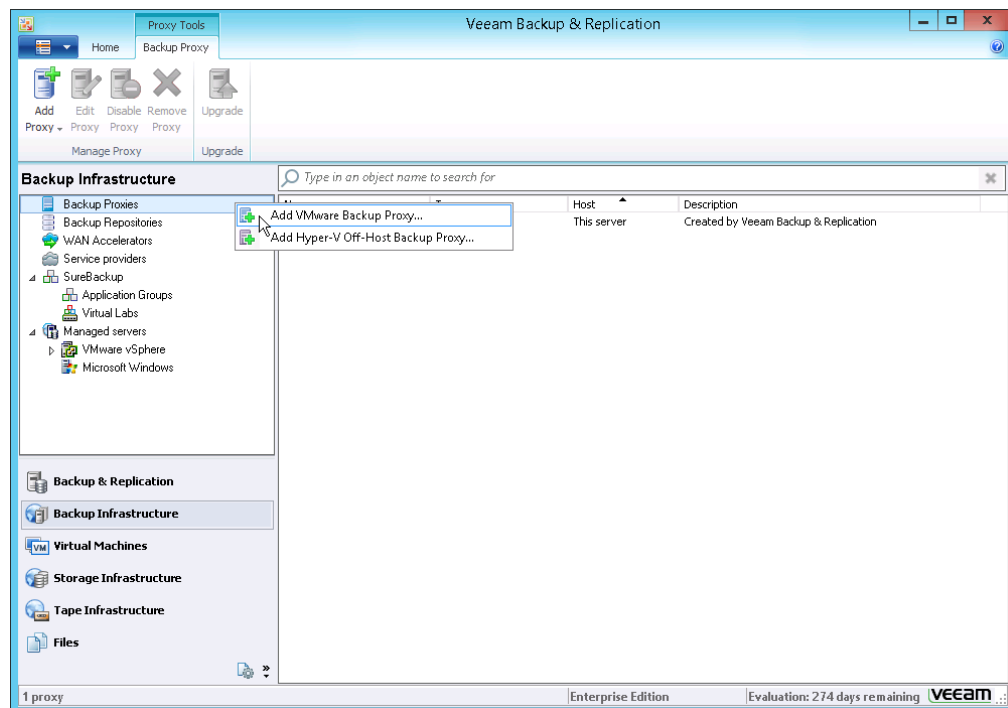
5. Enter credentials for the user account with administrator access permissions to the added server: click **Add** on the right of the **Credentials** field and specify the user name and password to connect to the added server. If you have specified credentials before, you can simply select them from the **Credentials** list. Note that credentials must be specified in the *DOMAIN\USERNAME* format.



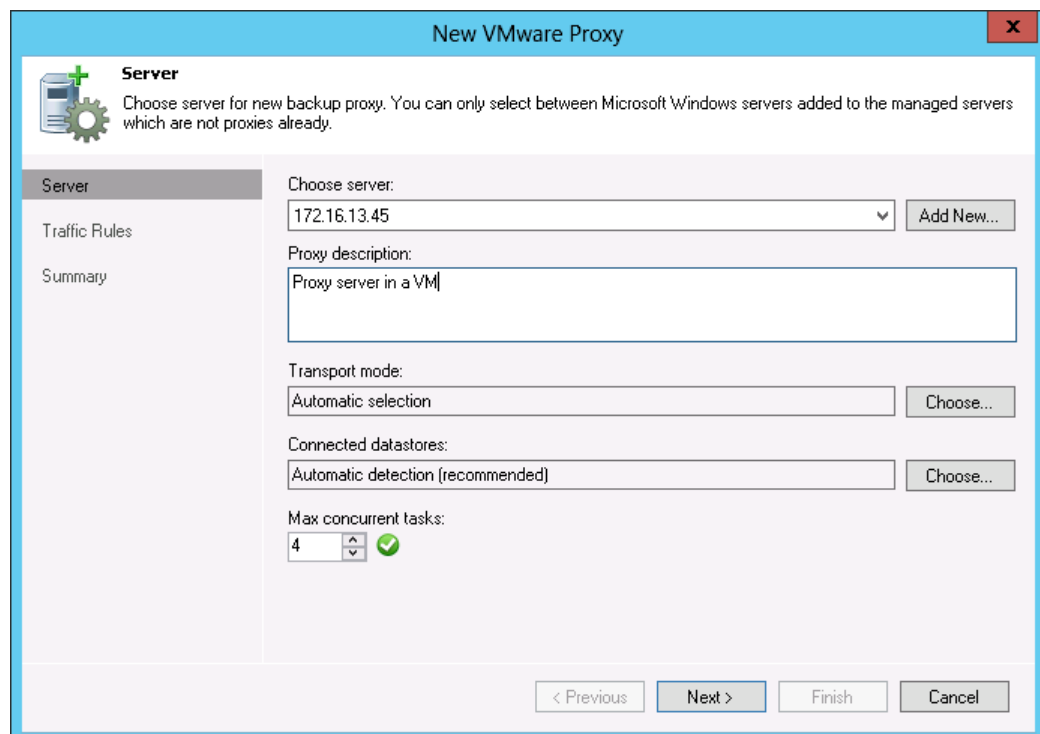
6. Follow the next steps of the wizard. At the last step of the wizard, click **Finish** to add the server.

Now you should assign the role of a backup proxy to the added server.

1. Open the **Backup Infrastructure** view.
2. Right-click the **Backup Proxies** node in the inventory pane and select **Add VMware Backup Proxy**.



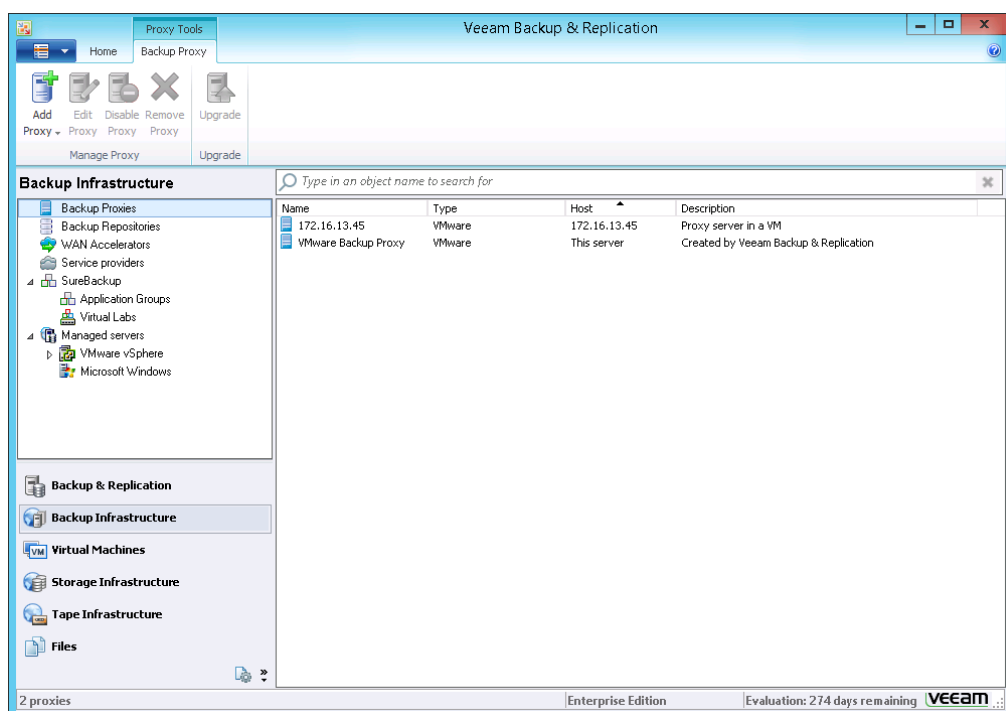
3. From the **Choose server** list, select the Microsoft Windows server that you have added.



4. 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 transport mode depending on the type of connection between the backup proxy and the source datastore.
5. In the **Connected datastores** field, leave the **Automatic detection** option selected. If you use SAN storage, Veeam Backup & Replication will automatically detect all datastores that the backup proxy can access via the *Direct SAN Access* mode.
6. Go through the next steps of the wizard without changing default settings. At the last step of the wizard, click **Finish** to finalize configuration of the backup proxy.

Validation

1. Open the **Backup Infrastructure** view.
2. In the inventory pane click the **Backup Proxies** node.
3. Make sure that the added backup proxy is available in the working area.



Configuring a Backup Repository

Insight into Backup Repository

Backup repository is a location to which created backup files, VM copies and auxiliary replica files are stored. You can use the following types of backup repositories:

- Microsoft Windows-based server with local or directly attached storage
- Linux-based server with local, directly attached or mounted NFS storage
- CIFS share
- Deduplication storage appliance*

* This guide describes backup repository configuration on Microsoft Windows, Linux servers and SMB shares. To learn about using deduplication storage appliances with Veeam Backup & Replication, see Veeam Backup & Replication User Guide at <http://www.veeam.com/documentation-guides-datasheets.html>.

Evaluation Case

In this exercise, you will configure a backup repository in which backup files, VM copies and auxiliary replica files will be stored.

Prerequisites

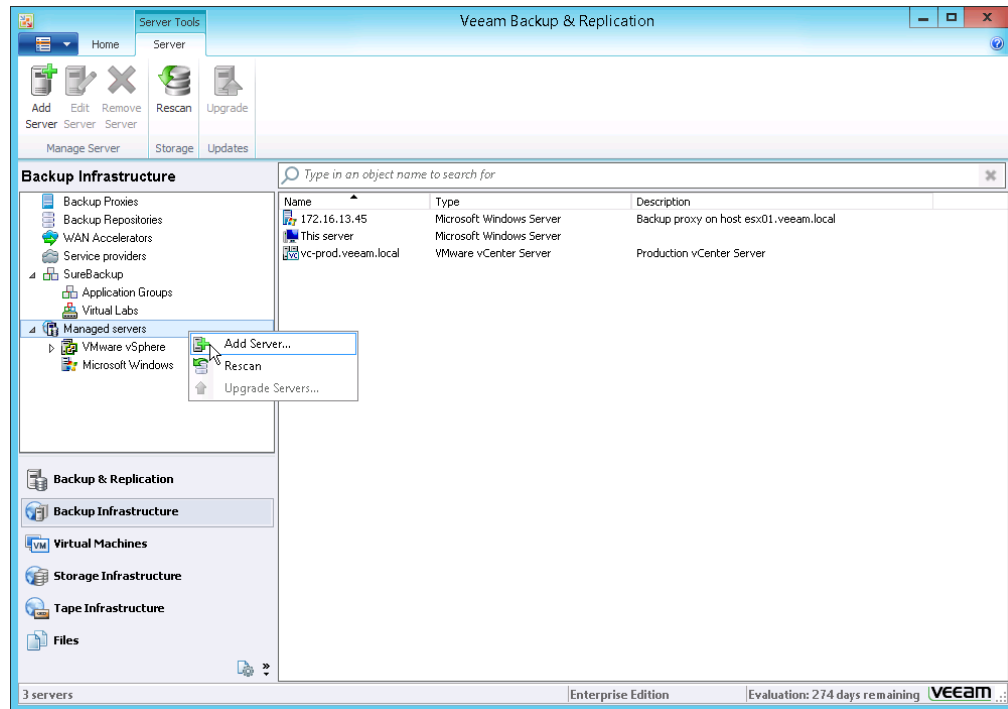
- You should have a Microsoft Windows server with a local or directly attached storage, Linux-based server with a local, directly attached or mounted NFS storage, CIFS share or deduplicating storage appliance that will perform the role of a backup repository.
- If a Microsoft Windows- or Linux-based server is used as a backup repository, these servers must have access to the Veeam backup server and the server that will perform a role of the backup proxy. In case of a CIFS share, make sure that you have an account with the *Full Control* permissions on the share to be able to connect to it.
- Make sure that all necessary ports are open.

Procedure

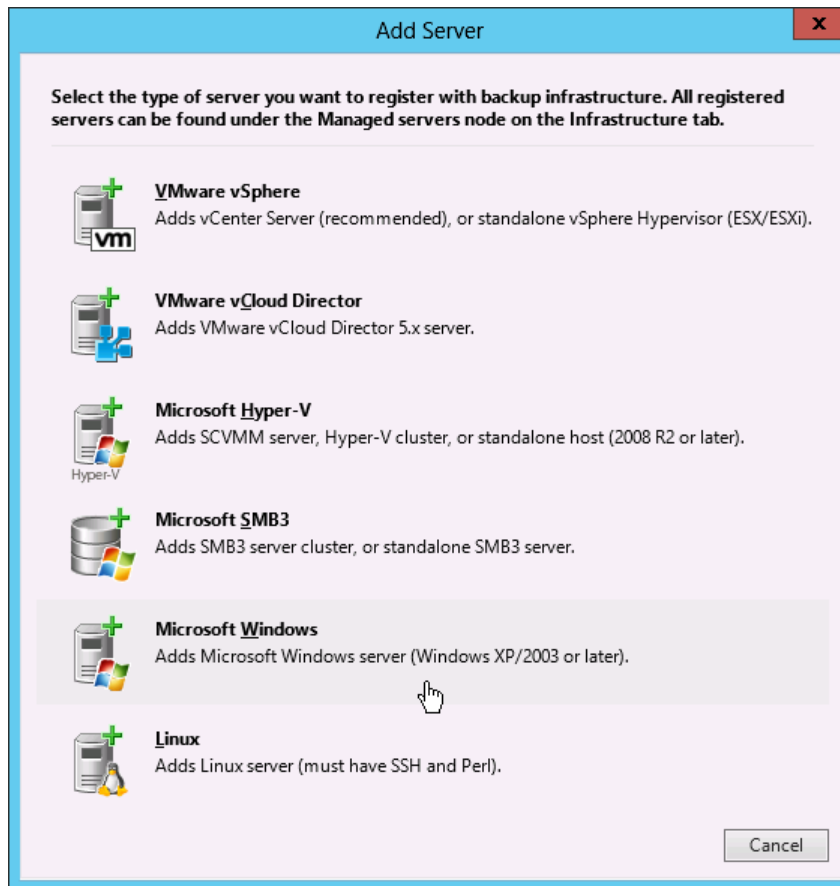
To configure a backup repository:

1. Open the **Backup Infrastructure** view.
2. If you plan to use a CIFS share as a backup repository, go to **step 7** of this procedure.

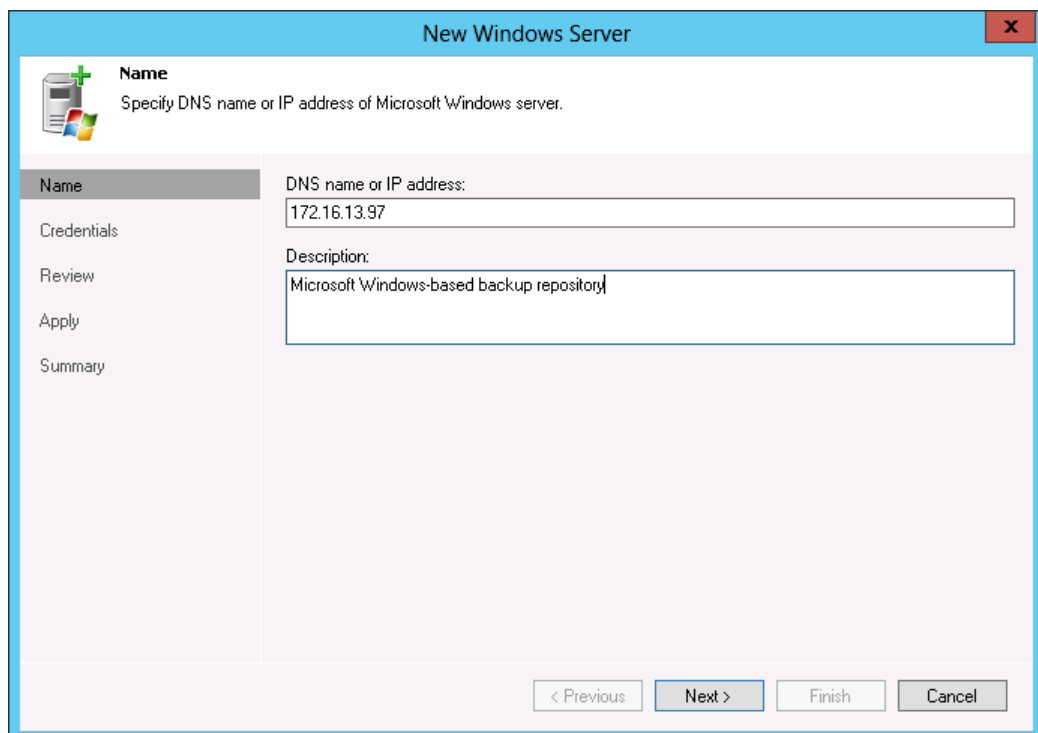
If you plan to use a Microsoft Windows-based or Linux-based server as a backup repository, right-click the **Managed servers** node in the inventory pane and select **Add Server**.



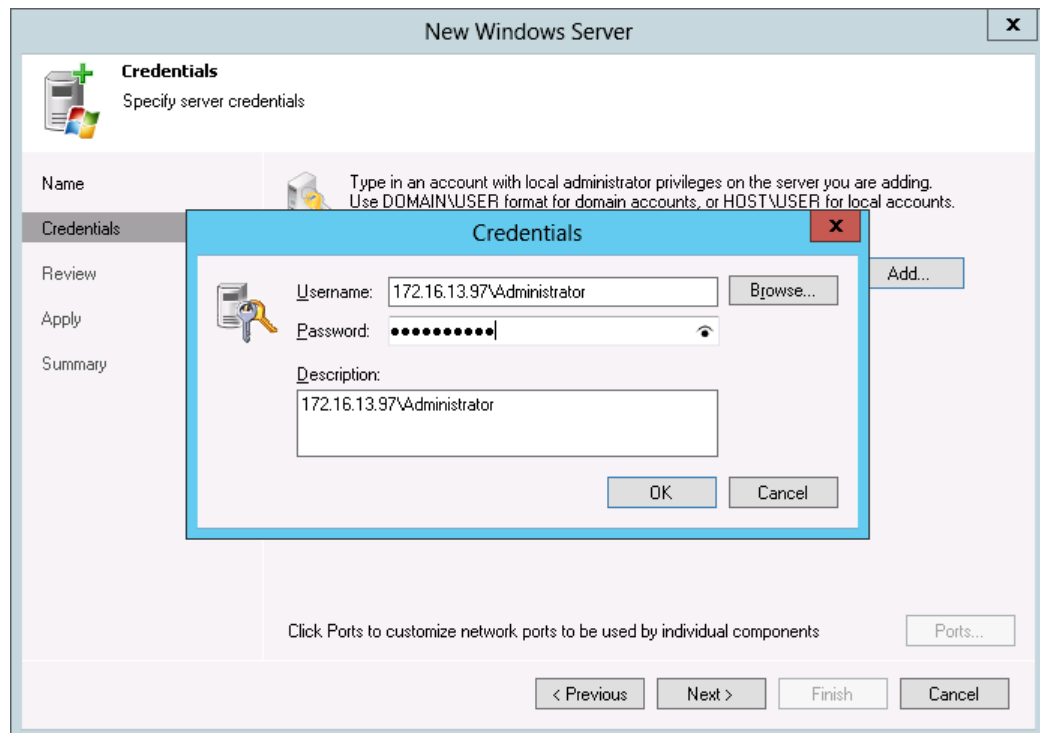
3. In the **Add Server** window, select the server type: **Microsoft Windows** or **Linux**.



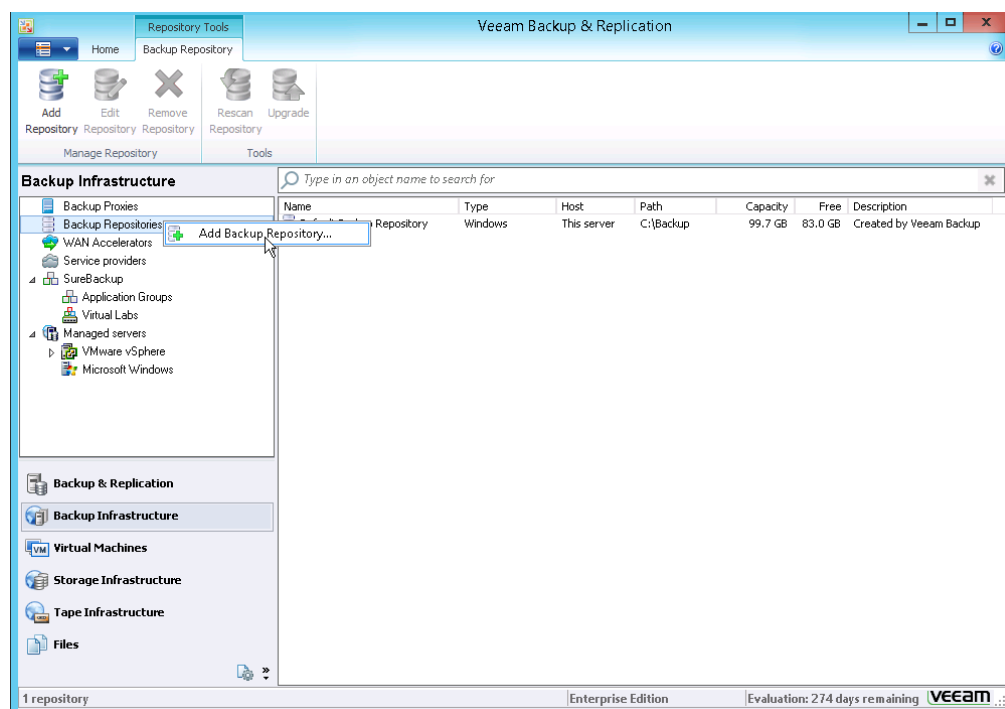
4. Specify a DNS name or the IP address of the Microsoft Windows- or Linux-based server that you want to use as a backup repository.



5. Enter credentials for the user account with administrator access permissions to the added server: click **Add** on the right of the **Credentials** field and specify the user name and password to connect to the added server. If you have specified credentials before, you can simply select them from the **Credentials** list. Note that credentials must be specified in the *DOMAIN\USERNAME* format.



6. Follow the next steps of the wizard. At the last step of the wizard, click **Finish** to add the server.
7. Right-click the **Backup Repositories** node in the inventory pane and select **Add Backup Repository**.



8. Specify a name for the added backup repository.

New Backup Repository

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

Name:
Backup Volume 01

Description:
Microsoft Windows-based repository

< Previous Next > Finish Cancel

9. Select the type of backup repository: **Microsoft Windows server**, **Linux server** or **Shared folder**.

New Backup Repository

Type
Choose type of backup repository you want to create.

Microsoft Windows server (recommended)
Microsoft Windows server with internal or directly attached storage. Data mover process running directly on the server allows for improved backup efficiency, especially over slow links.

Linux server (recommended)
Linux server with internal, directly attached, or mounted NFS storage. Data mover process running directly on the server allows for more efficient backups, especially over slow links.

Shared folder
CIFS (SMB) share. When backing up over slow links, we recommend that you specify a gateway server located in the same site with the shared folder.

Deduplicating storage appliance
Advanced integration with EMC Data Domain, ExaGrid and HP StoreOnce. For basic integration, use the Shared folder option above.

< Previous Next > Finish Cancel

- For a Microsoft Windows- or Linux-based repository, select the server that you have added.

For a CIFS share, specify a UNC path to the shared folder that will be used as a backup repository and enter credentials of an account with administrative privileges on the share.

The screenshot shows the 'New Backup Repository' wizard at the 'Server' step. The left sidebar has 'Server' selected. The main area shows a 'Repository server' dropdown set to '172.16.13.97 (Microsoft Windows-based backup repository)'. Below it is a table with columns 'Path', 'Capacity', and 'Free'. The first row shows 'C:\' with a capacity of '99.7 GB' and free space of '83.0 GB'. A 'Populate' button is to the right of the table. At the bottom are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Path	Capacity	Free
C:\	99.7 GB	83.0 GB

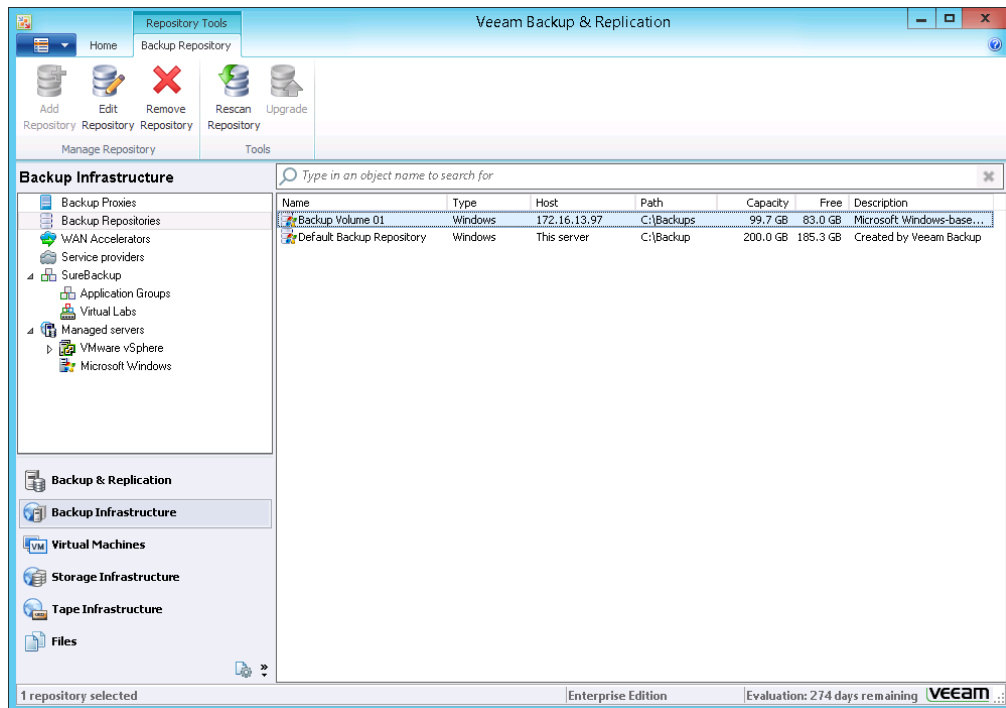
- Specify a path to the folder to which backups and auxiliary replica files will be stored.
- Click **Populate** to see how much space is available on your backup repository.

The screenshot shows the 'New Backup Repository' wizard at the 'Repository' step. The left sidebar has 'Repository' selected. The main area is titled 'Location' and shows 'Path to folder:' set to 'C:\Backups'. Below this, it displays 'Capacity: 99.7 GB' and 'Free space: 83.0 GB'. A 'Populate' button is next to the free space. The 'Load control' section has a warning about concurrent jobs and two options: 'Limit maximum concurrent tasks to:' (checked, set to 4) and 'Limit combined data rate to:' (unchecked). At the bottom right is an 'Advanced...' button. Navigation buttons at the bottom are '< Previous', 'Next >', 'Finish', and 'Cancel'.

- Go through the next steps of the wizard without changing default settings. At the last step of the wizard, click **Finish**.

Validation

1. Open the **Backup Infrastructure** view.
2. Click the **Backup Repositories** node in the inventory pane.
3. Make sure that the added backup repository is available in the working area.



DATA PROTECTION AND DISASTER RECOVERY TASKS

This section describes a set of exercises that you can perform to get to know the product functionality. Each exercise covers one of the most typical data protection or disaster recovery tasks and enables you to evaluate primary features of Veeam Backup & Replication.

Exercise List

To evaluate the key possibilities of Veeam Backup & Replication, perform the following exercises:

Exercise	Description	Time Estimates
Performing backup	Configure a backup job.	5-10 minutes
	Run the backup job to create a full image-level backup.	Varies*
	Run the backup job again to create an incremental backup.	Varies*
Backing up and restoring Microsoft SQL Server databases	Back up a virtualized Microsoft SQL Server and restore a database to a specific transaction.	Varies*
Restoring a full VM	Restore a full VM from the image-level backup.	Varies*
Restoring guest OS files	Restore specific guest OS files from the created image-level backup.	5-7 minutes
Restoring VM disks	Restore a VM drive from the image-level backup, and attach it to another VM.	Varies*
Restoring VM files	Restore a VM configuration file (VMX) from the image-level backup.	5-7 minutes
Creating a backup copy	Create a copy of a backup file and store it on the secondary backup repository.	Varies*
Performing replication	Configure a replication job.	5-10 minutes
	Run the replication job to create a VM replica on the target host.	Varies*
	Run the replication job once again to create a restore point for a VM replica.	Varies*
Failing over a VM replica	Fail over to the VM replica from the original VM.	3-5 minutes
Failing back to a primary VM	Fail back to the original VM from the VM replica.	Varies*

* The actual time required to perform this exercise depends on your test lab configuration, the hardware and software being used and the size of processed VMs.

Performing Backup

Insight into Veeam Backup

Designed specifically for virtual environments, Veeam Backup & Replication performs image-level backup of VMs.

Veeam Backup & Replication backs up a VM image as a whole: it copies VM data at a block level unlike traditional backup tools that process guest OS files separately. Veeam Backup & Replication retrieves VM data from the source storage, compresses and deduplicates it and writes to the backup repository in Veeam's proprietary format. You can use the image-level backup for all types of data restore scenarios: perform Instant VM Recovery*, restore a full VM, separate guest OS files, VM files and VM virtual disks from the same backup file.

To produce a backup, Veeam Backup & Replication leverages VMware snapshot capabilities. When you need to perform backup, Veeam Backup & Replication triggers VMware vSphere to create a snapshot, a cohesive point-in-time copy of a VM that Veeam Backup & Replication can access to retrieve VM data. The VMware snapshot technology lets you back up VMs without suspending them: this is also known as online hot backup.

In Veeam Backup & Replication, backup is job-driven: to perform backup, you need to configure a backup job. A backup job defines when, what, how and where to back up. One backup job can be used to process one or several VMs.

Veeam Backup & Replication conducts both full and incremental backup. During the first run of a backup job, Veeam Backup & Replication creates a full VM backup (VKB). All subsequent job cycles produce incremental backups:

- VIB if forward incremental or forever forward incremental backup is used.
- VRB if reversed incremental backup is used.

The number of increments kept on disk depends on retention policy settings.

The backup technology is a great choice for VMs with lower RTOs: typically, these are VMs running tier 2 applications. When the primary VM fails, you need some time to restore VM data from a compressed and deduplicated backup file. With Veeam's Instant VM Recovery*, however, this time is reduced to the minimum. At the same time, due to compression and deduplication, backup files require less disk space and can be saved to inexpensive storage.

* To learn about the Instant VM Recovery scenario, see [Veeam Backup & Replication User Guide](#).

Evaluation Case

In this exercise, you will create backup of VM(s). You will configure a backup job and run it twice to create full and incremental backups.

It is recommended that you create a backup job for at least two VMs: one Microsoft Windows-based VM and one running OS other than Microsoft Windows, for example, Linux, Unix, BSD or MacOS. This will let you follow the two guest OS recovery scenarios afterwards: [restoring guest OS files from a Windows-based VM](#) and [restoring guest OS files with multi-OS restore wizard](#).

Prerequisites

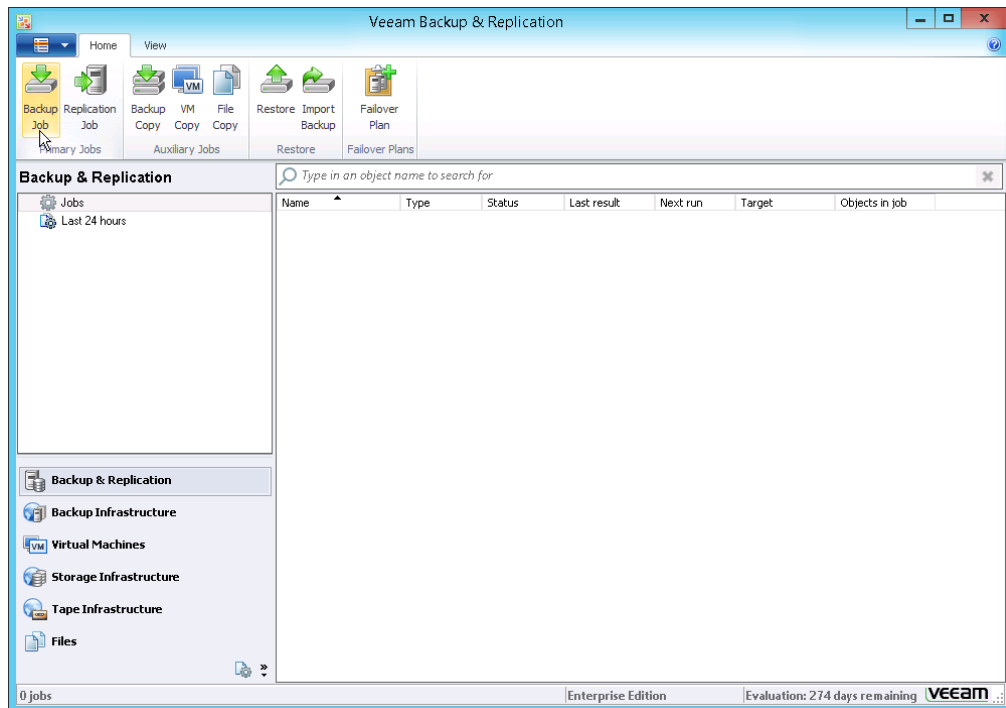
- All backup infrastructure components that will take part in the backup process should be added to the Veeam Backup & Replication console. These include an ESX(i) host on which VM(s) reside or vCenter Server to which this ESX(i) host is connected. If you plan to use a distributed architecture scenario, you should also have a backup proxy and a backup repository.
- [Optional] To receive an email notification when a backup job completes, specify global email notification settings. To do that, select **Options** from the main menu of Veeam Backup & Replication and specify necessary settings on the **Email Settings** tab.
- [Optional] To evaluate the application-aware image processing feature and the file indexing feature, make sure that at least one of backed up VMs runs the following OS'es:
 - Microsoft Windows Server 2003
 - Microsoft Windows Server 2008
 - Microsoft Windows Server 2008 R2
 - Microsoft Windows 7
 - Microsoft Windows 8
 - Microsoft Windows 8.1
 - Microsoft Windows 2012 Server
 - Microsoft Windows 2012 R2 Server

Procedure

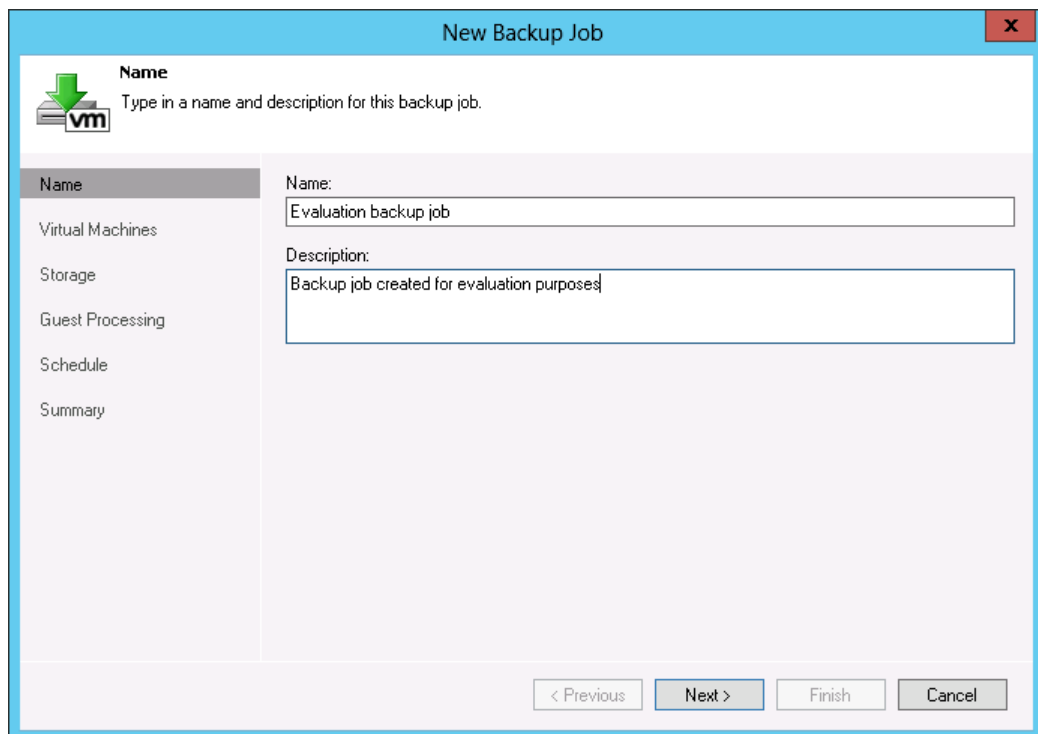
To perform backup of VMs, follow the next steps.

Step 1. Create a backup job

1. On the **Home** tab, click the **Backup Job** button.



2. Specify a name for the created backup job.

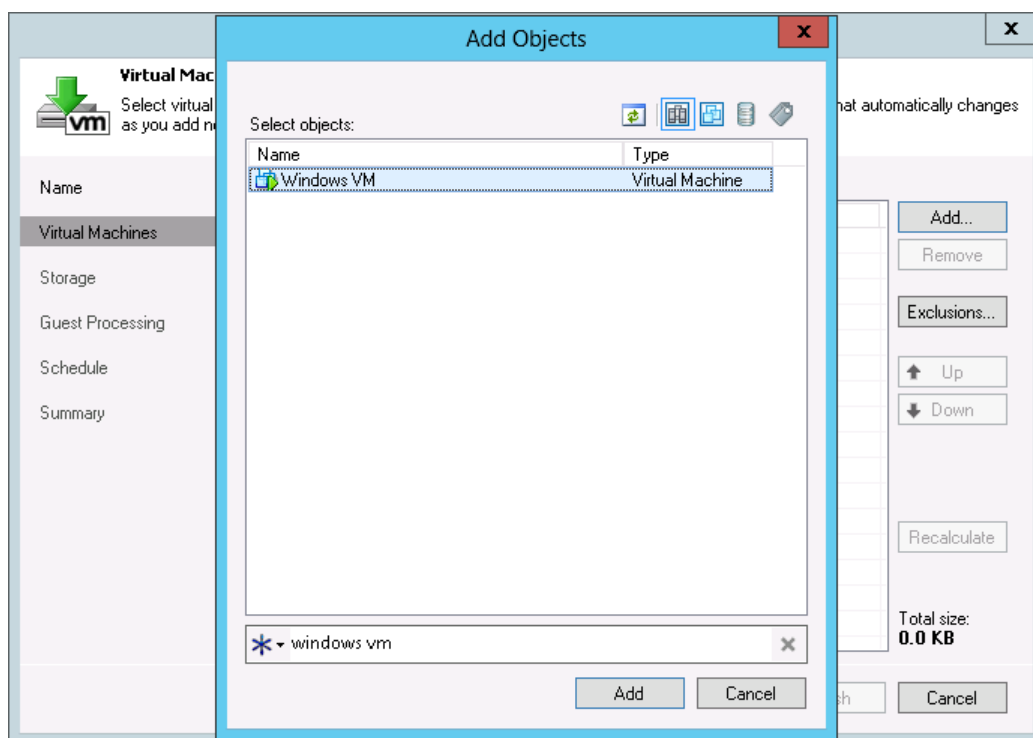


Step 2. Add VMs to the backup job

You can back up individual VMs or VM containers: folders, resource pools, clusters, vApps, datastores and so on. Jobs with VM containers are dynamic in nature: if a new VM is added to the container after the backup job is created, the job is automatically updated to include the new VM.

If you have connected vCenter Server rather than a standalone ESX(i) host to the Veeam backup server, VMs added to the job will be backed up even if they are vMotioned to another host.

1. At the **Virtual Machines** step of the wizard, click **Add**.
2. To quickly find a VM or VM container, enter the name of the object that you want to find in the search field and click the **Start search** button on the right. Select a VM or VM container in the displayed list and click **Add**.

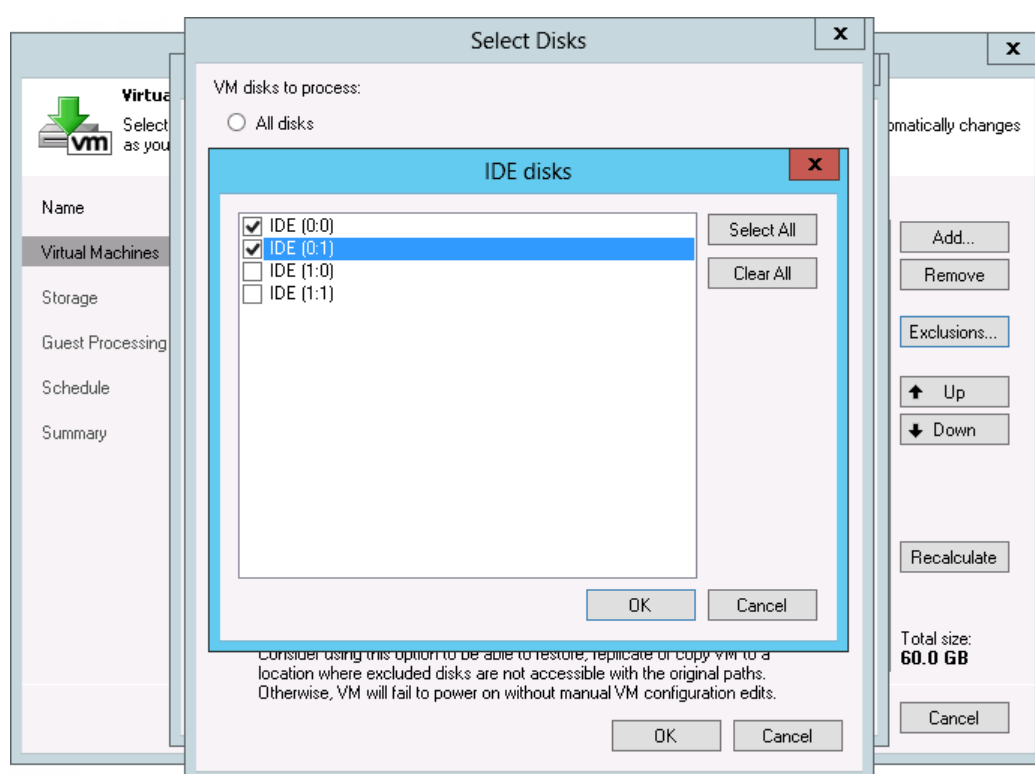


Step 3. Exclude VMs and VM disks

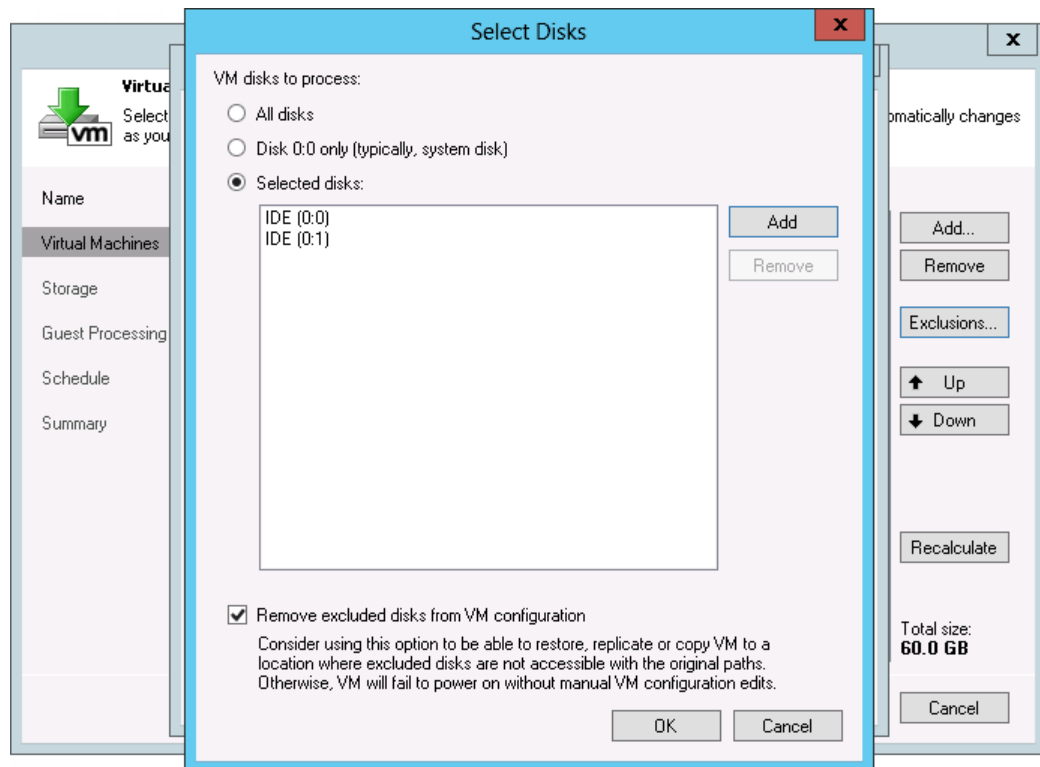
If you create a backup job for a VM container, you can exclude specific VMs or VM containers from the backup job. You can also select which VM disks to back up.

Note that Veeam Backup & Replication automatically excludes VM log files from backups to make the backup process faster and reduce the size of the backup file.

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. Use the **VMs** tab to exclude VMs or VM container from a backup job. Click **Add** and select VMs and VM containers that must be excluded. To quickly find a VM or container, enter the name of the object you want to find in the search field and click the **Start search** button on the right.
3. Click the **Disks** tab. Select a VM and click **Edit** to select disks that must be backed up. This functionality is useful, for example, if you want to back up only VM system drives.
4. To exclude disks of a VM added as part of a container, click **Add** on the right to include the VM in the list as a standalone instance.



5. If you exclude some VM disks from the backup, select the **Remove excluded disks from VM configuration** check box to automatically modify the VMX file with regard to the selected VM disks. With this option selected, you will be able to power on a restored VM without having to edit its configuration.

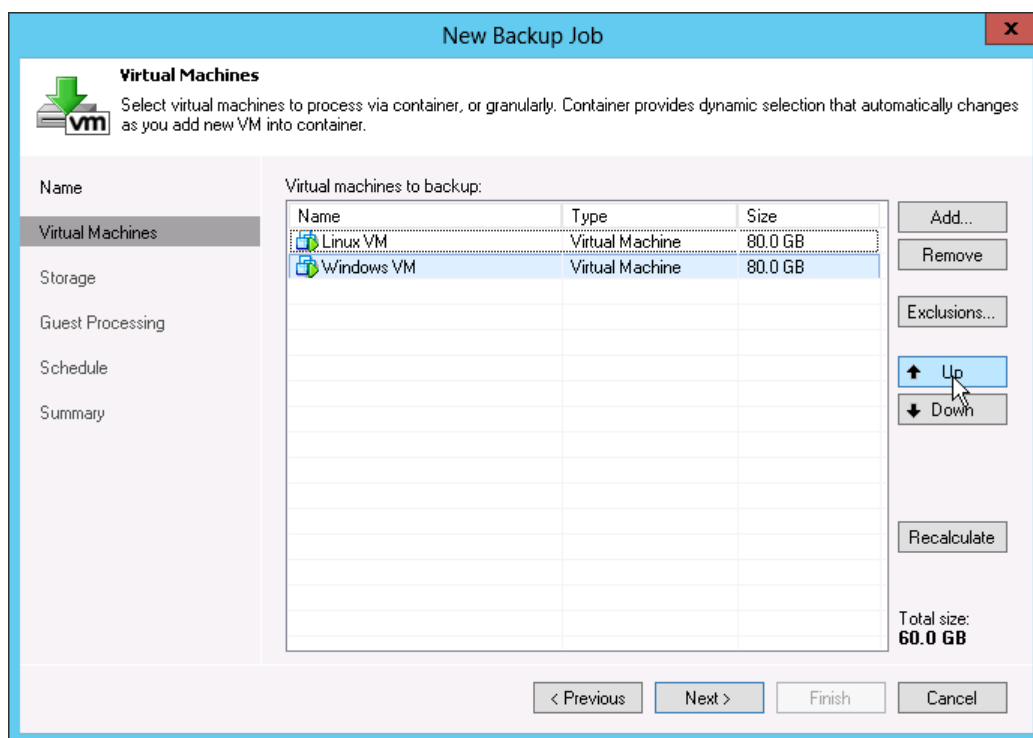


6. Click **OK**.
7. Click **Recalculate** to see the total size of selected objects.

Step 4. Define the VM backup order

If you have included a number of VMs or VM containers to the backup job, you can specify the order in which VMs should be processed. This will help you make sure that the most important VMs in the job are processed first – for example, if you must fit into the backup window and you are unsure how much time VM processing will take.

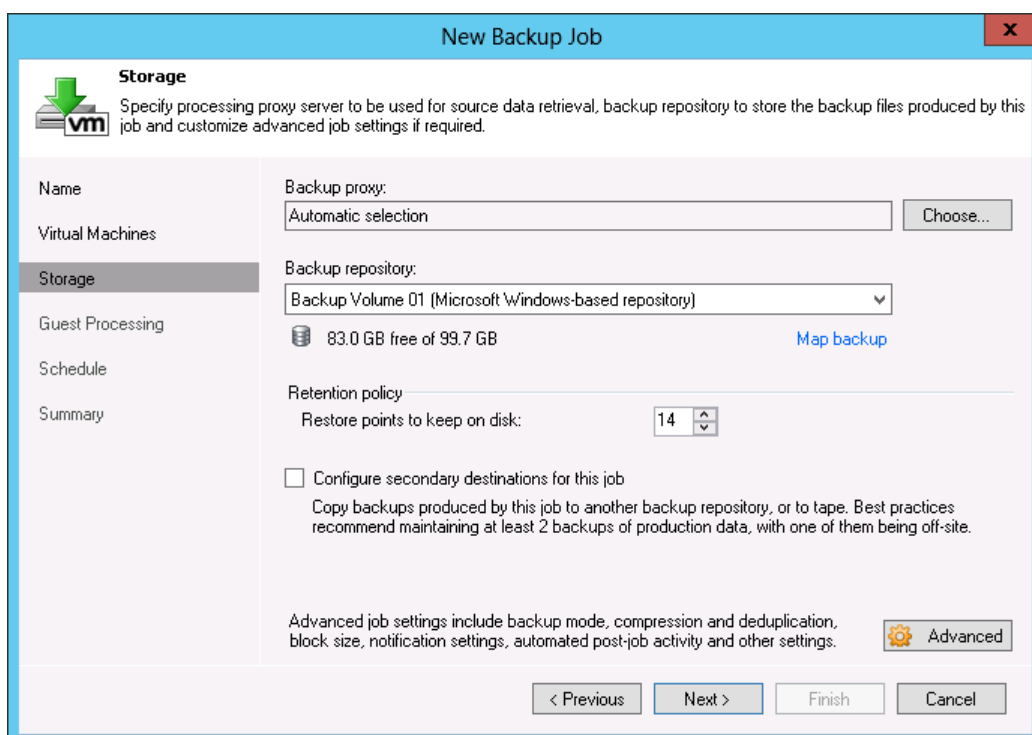
1. At the **Virtual Machines** step of the wizard, select the added VM in the list.
2. Use the **Up** and **Down** buttons on the right to move the VM higher or lower in the list. The higher is VM in the list, the higher its priority. If you added a VM container as a single instance, VMs inside the container will be processed at random.



Step 5. Select a backup proxy and backup repository

If you use a distributed deployment scenario, you must point the created job to the backup proxy that will process VM data and to the backup repository to which VM backups must be stored.

1. At the **Storage** step of the wizard, view the **Backup proxy** field and make sure that **Automatic selection** is specified there. Veeam Backup & Replication will check the settings of available backup proxies and select the most appropriate one for the job: the backup proxy that will enable the most efficient data retrieval from the source datastore. Veeam Backup & Replication first attempts to choose a backup proxy that uses the *Direct SAN Access* mode, then the backup proxy that uses the *Virtual Appliance* mode. If such backup proxies are not available, Veeam Backup & Replication selects the least loaded backup proxy that uses the *Network* mode.
2. From the **Backup repository** list, select the backup repository that you have configured. Backup files created by the job will be written to this backup repository.
3. Define the number of restore points that should be kept. By default, Veeam Backup & Replication keeps 14 restore points.



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 tabs: Name, Virtual Machines, Storage (selected), Guest Processing, Schedule, and Summary. The main area contains the following fields and options:

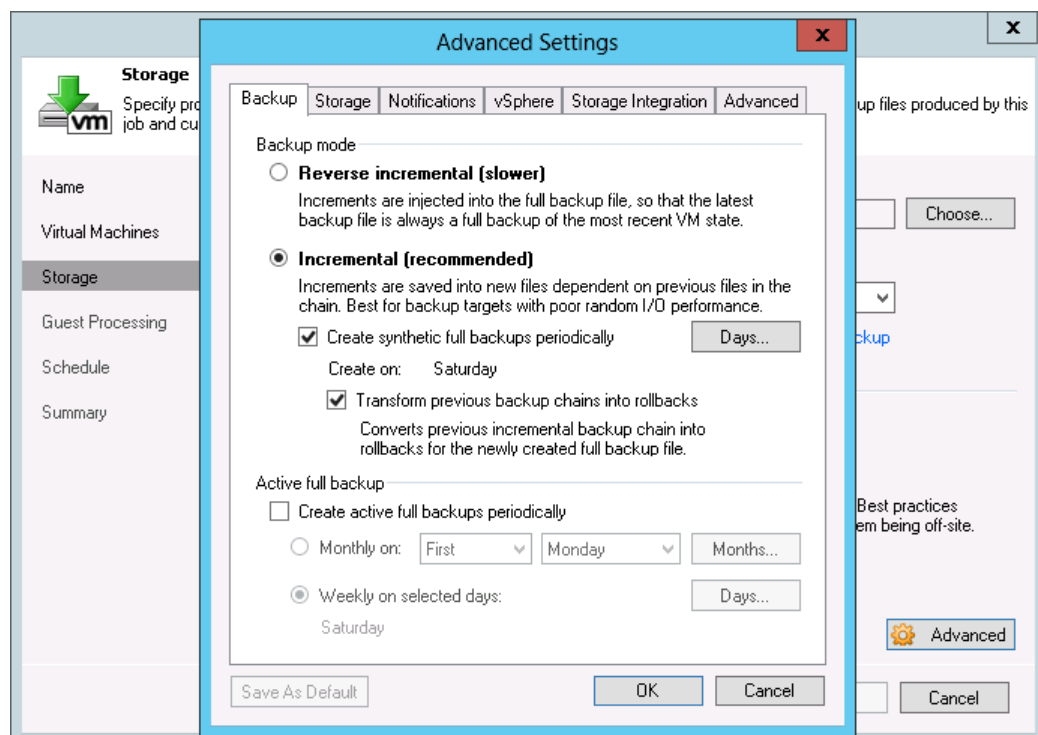
- Backup proxy:** A text box containing 'Automatic selection' and a 'Choose...' button.
- Backup repository:** A dropdown menu showing 'Backup Volume 01 (Microsoft Windows-based repository)' with a 'Map backup' link below it.
- Retention policy:** A section with a 'Restore points to keep on disk:' label and a spinner box set to '14'.
- Configure secondary destinations:** An unchecked checkbox with a description: '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 settings:** A link labeled 'Advanced' with a gear icon.

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

Step 6. Specify advanced backup settings

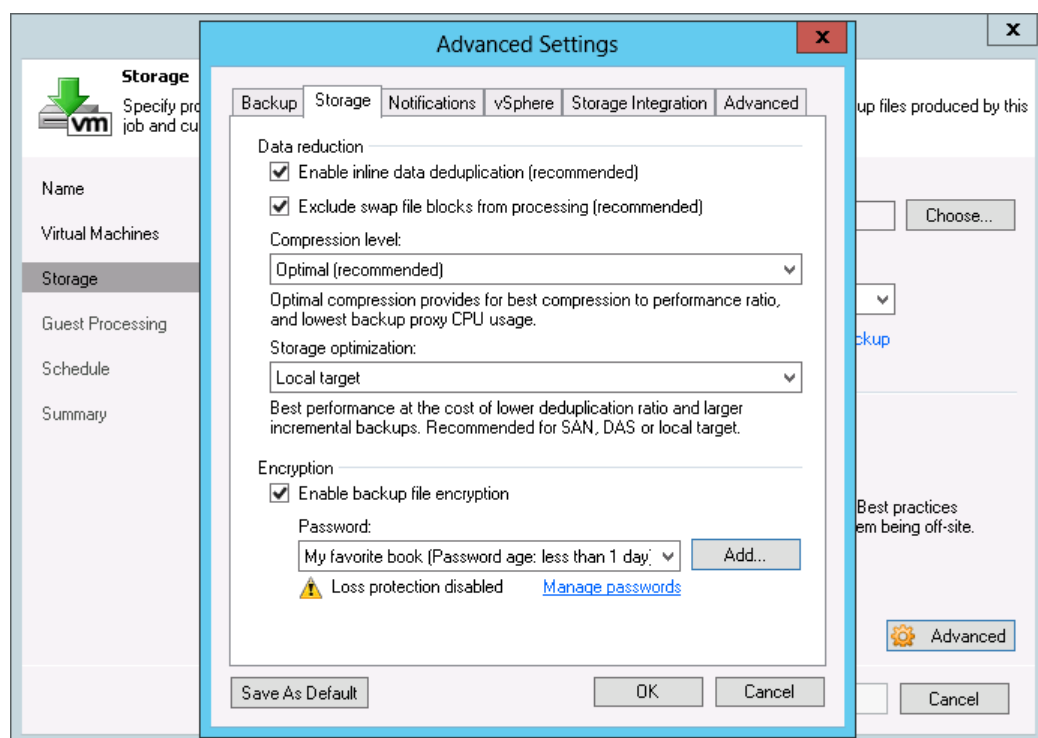
At the **Storage** step of the wizard, click **Advanced** to specify additional backup settings.

1. On the **Backup** tab, you can select the mode in which you want to perform backup: forever forward incremental, forward incremental or reverse incremental.
 - The forever forward incremental backup method produces a backup chain that consists of the first full backup and a set of forward incremental backups following it. To use this backup method, select **Incremental** and do not enable active full and/or synthetic full backup.
 - The forward incremental backup method produces a backup chain that consists of the first full backup and a set of forward incremental backups following it. Additionally, the forward incremental backup chain contains synthetic full and/or active full backups that “split” the backup chain into shorter series. To use this backup method, select **Incremental** and enable active full and/or synthetic full backup.
 - The reverse incremental backup method produces a backup chain that consists of the last full backup and a set of reverse incremental backups preceding it. To use this backup method, select **Reverse incremental**.
2. If you have selected to use forward incremental backup, select **Enable synthetic fulls** and specify a day on which a new full backup file should be created. This option is very useful if your corporate policies require you to periodically create full backups. To create a synthetic full backup, Veeam Backup & Replication uses full and incremental backup files that already reside on the backup repository (instead of retrieving VM data from the production storage). Synthetic full backups do not impact your virtual infrastructure or primary storage.
3. With synthetic backup scheduled, you will have a number of full backups on disk — a full backup created at the first run of the backup job and those create according to the synthetic backup schedule. To save the disk space, you can select the **Transform previous full backup chains into rollbacks** check box. In this case, Veeam Backup & Replication will transform all previous full backup chains to a reversed incremental backup sequence. This option allows you to keep only one full backup image on disk and so reduce the amount of space required to store backups.

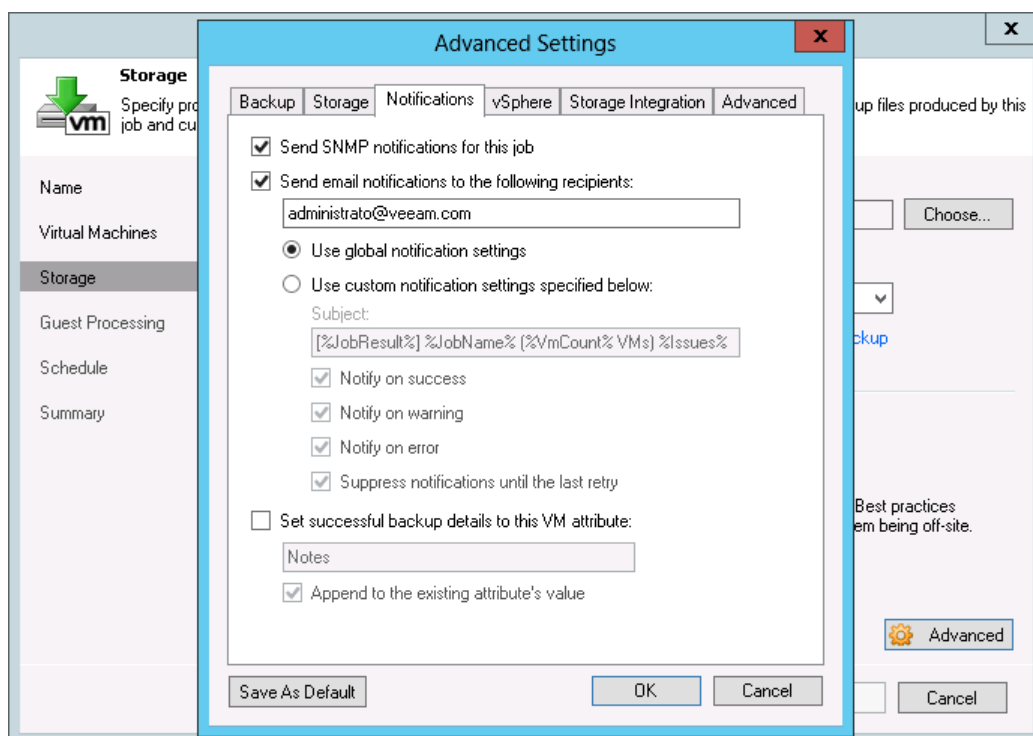


2. Click the **Storage** tab. Leave the **Enable inline data deduplication** check box selected. Veeam Backup & Replication deduplicates identical blocks of data when backing up multiple VMs in one job and eliminates empty space on logical disks of VMs. Use of deduplication dramatically reduces storage costs: you can reduce the backup size up to 90% when backing up VMs created from the same template.
3. Make sure the **Exclude swap file blocks from processing** check box is selected. Swap files are dynamic and change intensively between runs of a backup job. Veeam Backup & Replication will identify data blocks of the Microsoft Windows pagefile in the VM guest OS and exclude them from processing, which will result in increased performance and smaller increments.
4. To reduce the size of a backup file, Veeam Backup & Replication offers 5 compression levels: *None*, *Dedupe-friendly*, *Optimal*, *High* and *Extreme*, that provide different compression ratios to meet the needs of your environment. For this evaluation exercise, make sure that the **Optimal compression** level is selected.
5. Veeam Backup & Replication lets you encrypt backup files and restore data from encrypted backups even if you have lost a password. To enable backup file encryption, select the **Enable backup file encryption** check box and click **Add** on the right to choose the necessary password.

Veeam Backup & Replication lets you decrypt backup files even if you have lost or forgotten the password. To learn more, see the [Restoring Data from Encrypted Backup File Without Password](#) scenario.

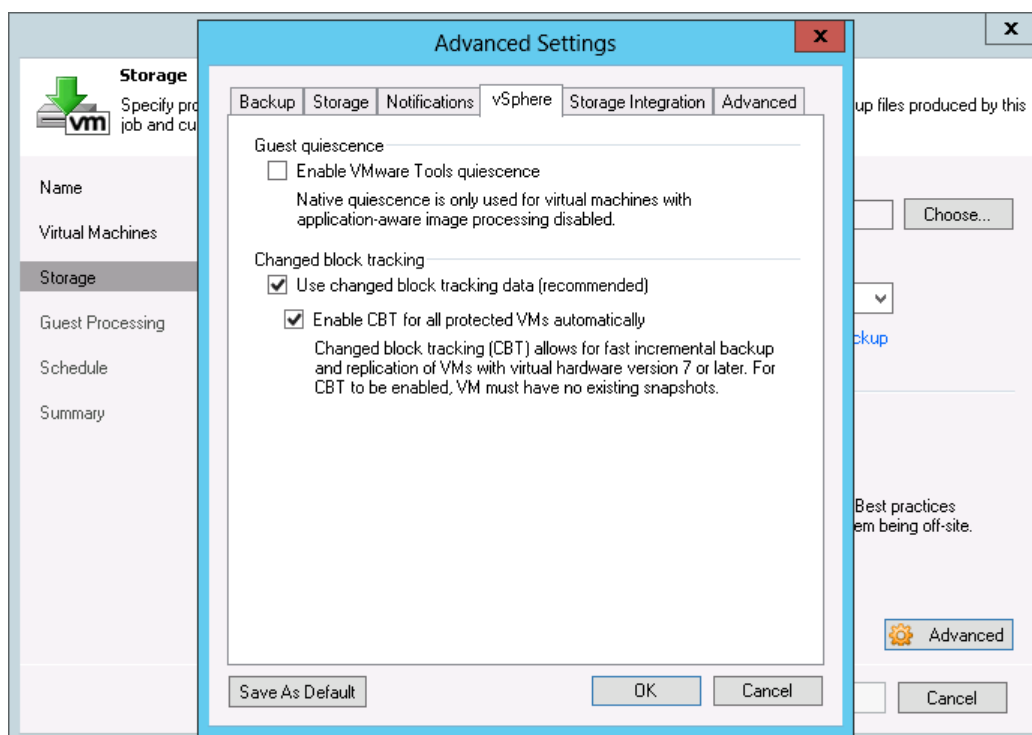


- Click the **Notifications** tab. Select the **Send email notifications to the following recipients** check box. When the job completes, you will receive a notification with details on the job performance. You will be able to receive email notifications only if you specify global email settings. To specify email settings, select **Options** from the main menu.



Note: Make sure that you specify your email address once: either in the **To** field in general notification settings or in job notification settings. If you specify both, you will receive two identical notifications when the job is completed.

- Click the **vSphere** tab.
- If you back up a VM that is not VSS-aware, for example, a Linux-based VM, make sure that the **Enable VMware tools quiescence** check box is selected. This option helps create transactionally consistent backups of such VMs.
- Make sure that the **Use changed block tracking data** check box is selected. For VMware VMs with hardware version 7 or later, Veeam Backup & Replication employs VMware vSphere Changed Block Tracking (CBT). Instead of scanning VMFS to know which data blocks have changed since the previous job run, Veeam Backup & Replication queries the CBT module to get the list of changed blocks. Use of CBT increases the speed and efficiency of block-level incremental backups. For example, if only 5% of a VM changed since the last backup, incremental backup will be performed 20 times faster.
- Make sure that the **Enable changed block tracking for all processed VMs** check box is selected. This option forces use of CBT even in case it is switched off at the level of the ESX(i) host.



Step 7. Specify additional guest OS processing options

At the **Guest Processing** step of the wizard, you can enable additional options for VM guest OS processing:

- Application-aware image processing, which will ensure proper restore of VSS-aware applications. To back up running VMs with VSS-aware applications, Veeam Backup & Replication uses application-aware image processing based on Microsoft VSS. Jobs with application-aware image processing produce transactionally consistent backups, that, unlike crash consistent backups, ensure proper recovery of virtualized applications without any data loss.
- Guest OS file indexing, which will enable you to search for guest OS files in backed up VMs and restore files in 1 click. With this option selected, Veeam Backup & Replication creates a catalog (or index) of VM guest OS files. To learn about the 1-click restore scenario, see the [Searching for Guest OS Files and Performing 1-Click Restore](#) section.

To enable application-aware image processing and indexing:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware image processing** check box.
2. Select the **Enable guest file system indexing** check box.
3. Specify credentials for the user account with Local Administrator privileges on all VMs included into the job: click **Add** on the right of the **Credentials** field and specify the user name and password. If you have specified credentials before, you can simply select them from the **Credentials** list. OS credentials are required to install, start and remove Veeam's runtime process that coordinates indexing and VSS activities inside the VM.

New Backup Job

Guest Processing
Choose additional processing options available for Microsoft Windows guests.

Enable application-aware processing
Quiesces applications using Microsoft VSS to ensure transactional consistency, performs transaction logs processing, and prepares application-specific VSS restore procedure.
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

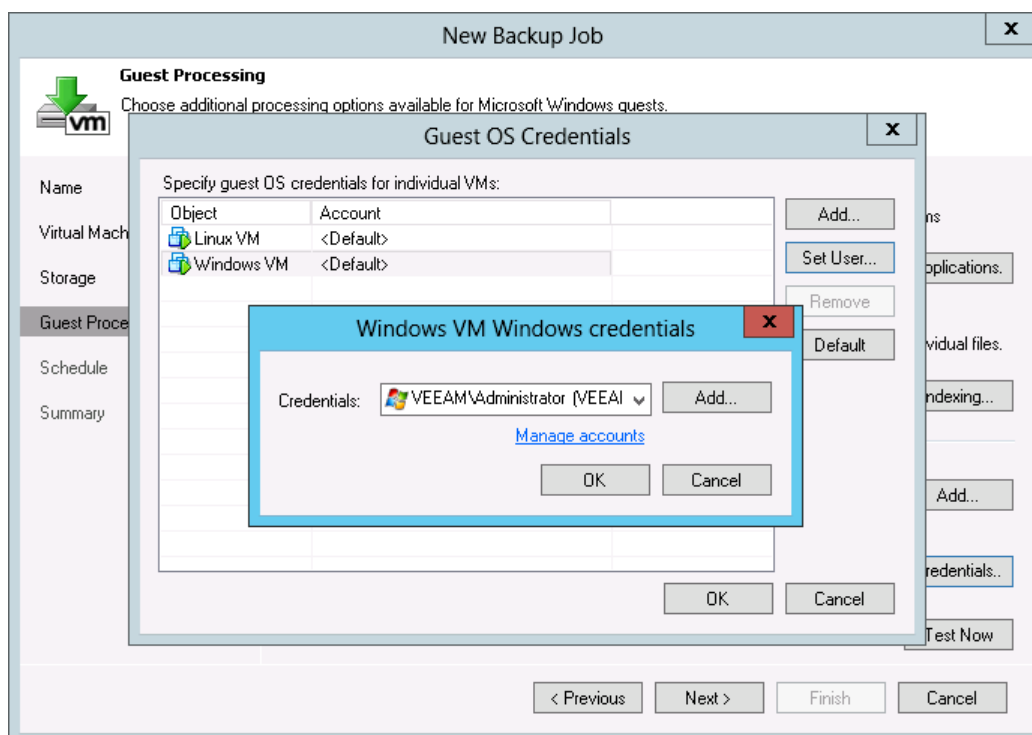
Credentials: VEEAM\Administrator (VEEAM\Administrator, last edited:) **Add...**
[Manage accounts](#)

Customize guest OS credentials for individual VMs and operating systems **Credentials...**

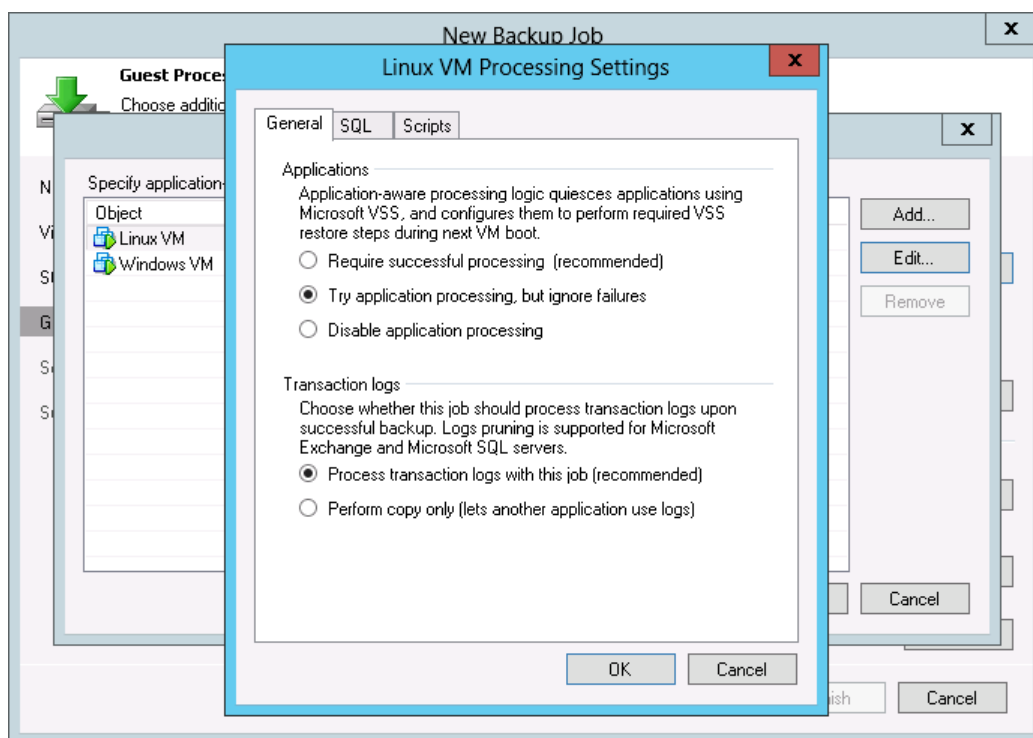
Test Now

< Previous **Next >** **Finish** **Cancel**

By default, the specified guest OS credentials are used for all VMs processed by the backup job. If you back up several VMs that use different guest OS credentials, click **Credentials**. Select a VM in the list and click **Set User > Windows credentials/Linux credentials**. Then enter guest OS credentials with Local Administrator privileges for this specific VM. Repeat the procedure for all VMs in the job.

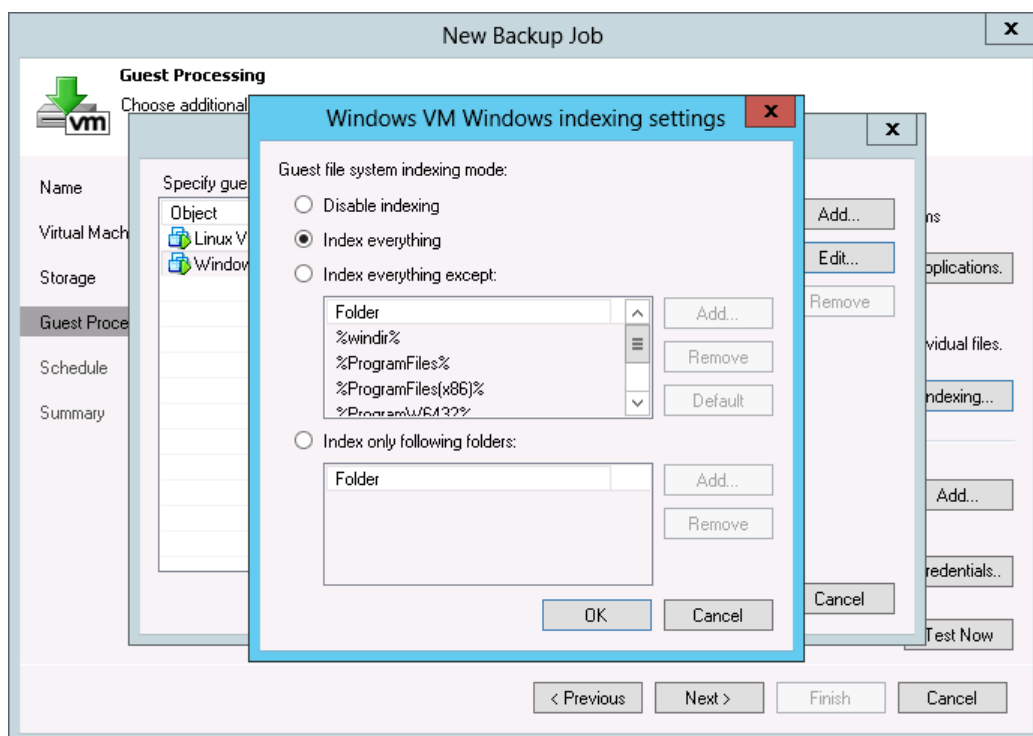


4. To specify advanced options for VSS processing, click **Applications**. Select a VM in the list and click **Edit**.
5. On the **General** tab, select the **Try application processing, but ignore failures** option to continue the backup job even if VSS errors occur. If application-aware image processing fails during the job, Veeam Backup & Replication will try to create a transactionally consistent backup with VMware Tools quiescence.
6. Make sure that the **Process transaction logs with this job** option is selected to correctly handle transaction logs after the backup job is completed. In this case, if the backup job finishes successfully, Veeam Backup & Replication will truncate transaction logs so that they do not overflow storage space. If together with Veeam Backup & Replication you use a third-party backup tool that maintains consistency of transaction logs, select the **Perform copy only** option to prevent possible conflicts.



Note: If you add a virtualized Microsoft SQL Server to the backup job, you can configure the job to copy transaction logs. To learn more, see [Backing up and Restoring Microsoft SQL Server Databases](#).

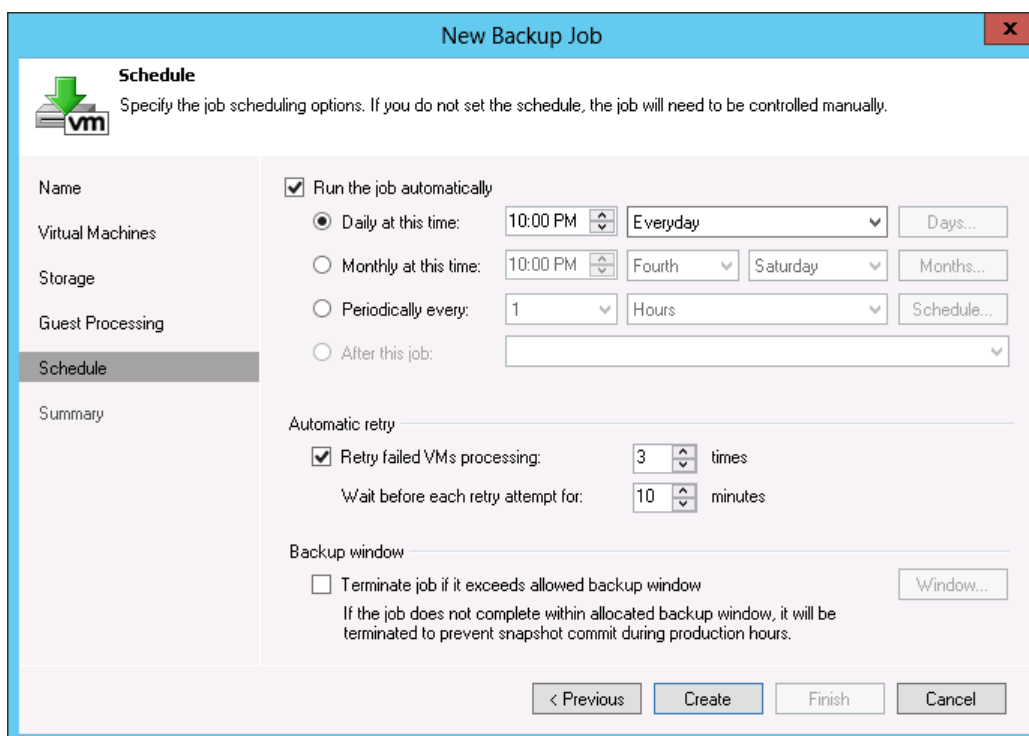
7. Click **Indexing**, select the necessary VM in the list and click **Edit > Windows indexing** or **Linux indexing**. Select **Index everything** to perform indexing of the entire guest file system.



Step 8. Specify job scheduling settings

A backup job can be scheduled or run manually. To schedule a backup job:

1. At the **Schedule** step of the wizard, select the **Run the job automatically** check box. If you do not select this check box, the job will be saved and you will have to run it manually.
2. Select the schedule type: daily, monthly, periodically or continuously. You can also chain the jobs so that they run one after another.
3. Make sure the **Retry failed VM processing** check box is selected. During the retry cycle, only VMs that failed during the main backup cycle will be processed.
4. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**. Define the backup window for your environment. In case the created job overlaps the specified window, it will be automatically terminated not to produce additional overhead on your virtual environment.
5. Click **Create**.



The screenshot shows the 'New Backup Job' wizard with the 'Schedule' step selected. The interface includes a sidebar with navigation options: Name, Virtual Machines, Storage, Guest Processing, Schedule (selected), and Summary. The main area is titled 'Schedule' and contains the following settings:

- Run the job automatically:** Checked. Options include:
 - Daily at this time:** 10:00 PM, Everyday (Days... button)
 - Monthly at this time:** 10:00 PM, Fourth, Saturday (Months... button)
 - Periodically every:** 1, Hours (Schedule... button)
 - After this job:** (Dropdown menu)
- Automatic retry:**
 - Retry failed VMs processing:** Checked, 3 times
 - Wait before each retry attempt for:** 10 minutes
- Backup window:**
 - Terminate job if it exceeds allowed backup window:** Unchecked. A 'Window...' button is available.
 - If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

At the bottom, there are four buttons: '< Previous', 'Create', 'Finish', and 'Cancel'.

Step 9. Review job settings and start the job

1. Review the summary of backup job settings.
2. Select the **Run the job when I click Finish** check box and click **Finish**. The job will start.

The image displays two screenshots of the 'New Backup Job' wizard's Summary page. Both screenshots show the 'Summary' tab selected in the left-hand navigation pane. The top screenshot is for a VMware backup, showing source items as 'Linux VM (vc-prod.veeam.local)' and 'Windows VM (vc-prod.veeam.local)'. The bottom screenshot is for a Hyper-V backup, showing source items as 'Linux VM (172.16.13.45)' and 'Windows VM (172.16.13.45)'. Both screenshots show the 'Run the job when I click Finish' checkbox checked. The 'Finish' button is highlighted in blue in both screenshots.

New Backup Job

Summary
The job's settings have been saved successfully. Click Finish to exit the wizard.

Name
Virtual Machines
Storage
Guest Processing
Schedule
Summary

Summary:
Name: Evaluation backup job
Target Path: C:\Backups
Type: VMware Backup
Enable application-aware processing
Enable guest file system indexing
Source items:
Linux VM (vc-prod.veeam.local)
Windows VM (vc-prod.veeam.local)
Command line: "C:\Program Files\Veem\Backup and Replication\Backup
\Veem.Backup.Manager.exe" backup 75f0e088-2fb5-4513-bfe8-7a58c4087ef8

☒ Run the job when I click Finish

< Previous Next > **Finish** Cancel

New Backup Job

Summary
The job's settings have been saved successfully. Click Finish to exit the wizard.

Name
Virtual Machines
Storage
Guest Processing
Schedule
Summary

Summary:
Name: Evaluation backup job
Target Path: C:\Backups
Type: Hyper-V Backup
Enable application-aware processing
Enable guest file system indexing
Source items:
Linux VM (172.16.13.45)
Windows VM (172.16.13.45)
Target repository: Backup Volume 01
Target repository host: 172.16.13.97
Target repository path: C:\Backups
Command line: "C:\Program Files\Veem\Backup and Replication\Backup
\Veem.Backup.Manager.exe" backup 3b7842bb-cae5-48ad-88c3-5df78149bbf2

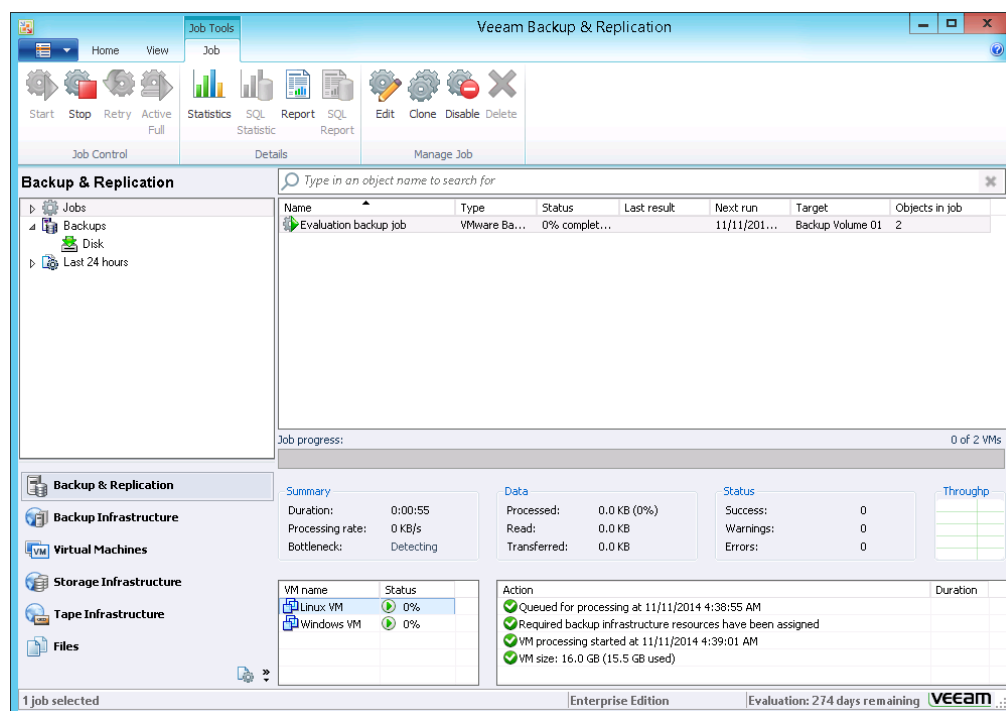
☒ Run the job when I click Finish

< Previous Next > **Finish** Cancel

Step 10. Monitor job performance in real time

When the job is running, you can view job statistics in the real-time mode. Job statistics provide detailed data on the job: job progress, duration, processing rate, performance bottlenecks, the amount of data processed, read and transferred and other details of the job performance. Beside general job statistics, you can view detailed data for each VM or VM container processed by the job.

1. Open the **Backup & Replication** view.
2. Select the **Jobs** node in the inventory pane.
3. Click the job in the working area to open the lower pane with job statistics. Now you can track the job performance as the job runs.
4. Select the name of a specific VM or a VM container to view detailed statistics for this specific object only. Note that Veeam Backup & Replication processes all VMs and VM containers in the job in parallel.

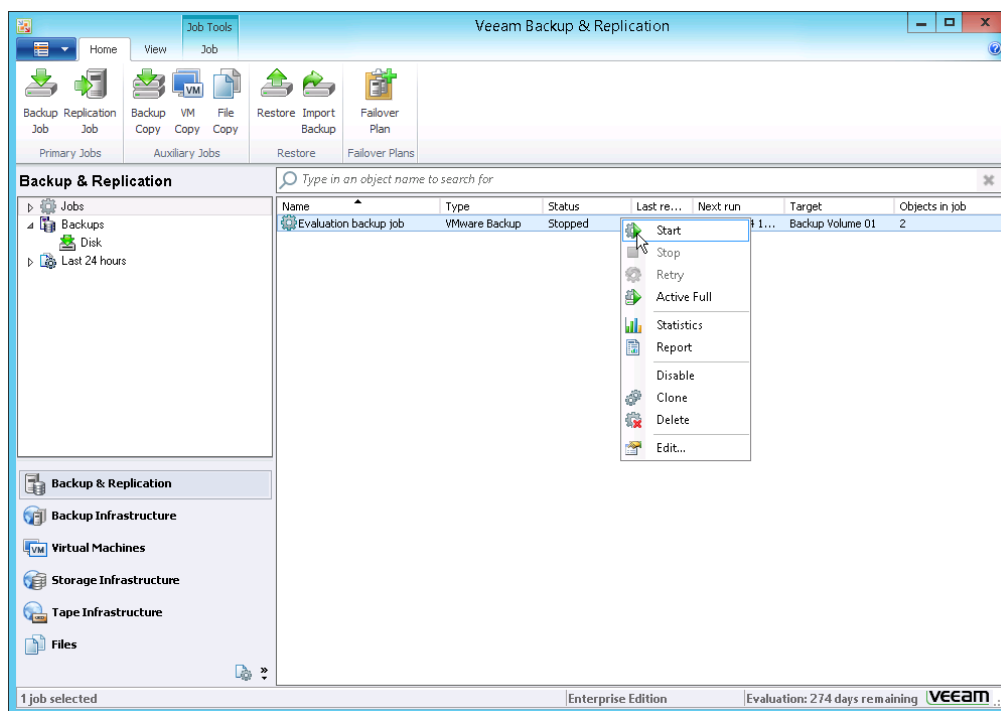


5. Wait for the job to complete. Note that the job must complete with the *Success* or at least the *Warning* status. If the job completes with the *Failed* status, the backup file will not be created and you will not be able to perform restore operations.

Step 11. Perform incremental backup

To perform incremental backup of a VM, do the following:

1. Open the **Backup & Replication** view.
2. Select the **Jobs** node in the inventory pane.
3. Right-click the job in the working area and select **Start**. Wait for the job to complete. Note that the job must complete with the *Success* or at least the *Warning* status.



Validation

After a backup job completes, the resulting backup file is stored to the backup repository that you have selected as a backup target. Veeam Backup & Replication creates a full backup file, VBK, during the first run of a backup job. During every next job run, it copies changes that were made to the VM since the last backup, whether full or incremental.

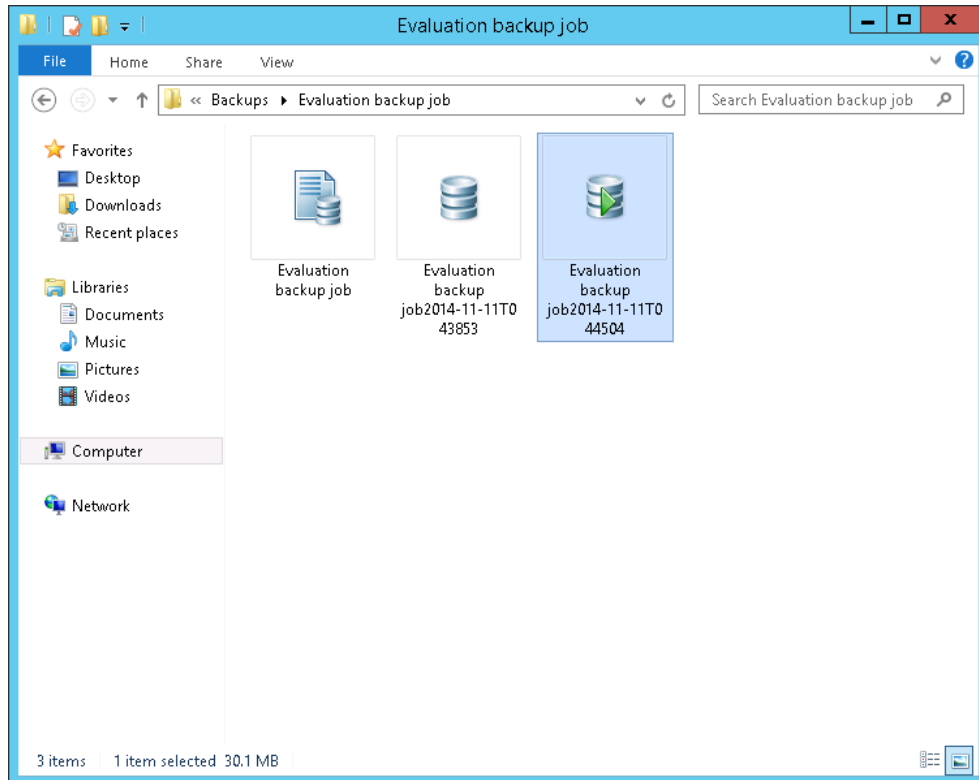
Depending on the backup mode you selected, Veeam Backup & Replication handles incremental changes differently:

- If you use the incremental backup mode (forever forward incremental or forward incremental), Veeam Backup & Replication saves incremental changes to the VIB file next to the VBK file on the backup repository.
- If you use the reversed incremental backup mode, Veeam Backup & Replication injects copied changes to the full backup file and saves replaced blocks of data as a reversed increment file, VRB, next to the VBK file on the backup repository.

Additionally, next to backup files, Veeam Backup & Replication creates a backup 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. This metadata file facilitates import of backups and mapping of backup jobs to existing backups.

To check backup results:

1. Open the **Files** view.
2. In the inventory pane, expand the backup repository file tree.
3. Open the target folder on the backup repository. In this folder, find the subfolder with the backup job name and open it. Make sure that it contains a VBK, .VIB/VRB and VBM files.



4. Open the **Backup & Replication** view.
5. Select the **Jobs** node in the inventory pane. Double-click the backup job in the working area. Check the job results.
6. If you have selected to receive an email message once the job is completed, open your email client and check the *Inbox* folder. Make sure that you have two incoming emails with job results: one for the full backup and one for the incremental backup.

Backing up and Restoring Microsoft SQL Server Databases

To protect a virtualized Microsoft SQL Server, you can configure a backup job that will create image-level VM backups and, in addition, copy database transaction logs. Image-level backups will capture the VM state at specific points in time. Transaction logs will keep records of all transactions performed against protected databases since the moment of the last backup. If a Microsoft SQL Server VM fails, you can recover the VM from the necessary restore point and then apply transaction logs to get databases on the Microsoft SQL Server to the required state between backups.

To configure a backup job that copies transaction logs, you must create a regular backup job and specify advanced settings for transaction logs shipping. In these settings, you define:

- How often you want to back up transaction logs
- How transaction logs must be shipped to the destination
- How long transaction logs must be retained

With these settings enabled, Veeam Backup & Replication actually creates two jobs linked with each other:

- A regular backup job responsible for creating image-level backups
- An auxiliary job responsible for shipping database transaction logs

The regular backup job runs by the defined job schedule. It creates image-level backups and saves them on the backup repository. After the image-level backup has been successfully created, Veeam Backup & Replication truncates transaction logs on the virtualized Microsoft SQL Server.

The auxiliary backup job runs continuously. The job copies transaction logs accumulated between VM restore points at scheduled intervals, for example, every 15 minutes. As a result, on the backup repository you will have a chain of restore points and a set of transaction logs that cover intervals between these restore points.

Veeam Backup & Replication ships transaction logs to the backup repository and saves them in files of VBL format next to VM image-level backups. To ship transaction logs from the virtualized Microsoft SQL Server to the backup repository, Veeam Backup & Replication uses log shipping servers — Microsoft Windows machines added to the backup infrastructure. You can select explicitly what log shipping servers you want to use, or let Veeam Backup & Replication assign log shipping servers automatically.

For restore operations, Veeam Backup & Replication offers a special tool — Veeam Explorer for Microsoft SQL Server. Veeam Explorer for Microsoft SQL Server is integrated with Veeam Backup & Replication. The explorer is installed automatically when you deploy Veeam Backup & Replication.

Veeam Explorer for Microsoft SQL Server supports a number of restore scenarios:

- You can restore a Microsoft SQL Server to a specific point in time or to a specific transaction.
- You can restore the database to a specific point in time or transaction and export it to the necessary location.

Evaluation Case

In this exercise, you will back up a Microsoft SQL VM and restore a database on the Microsoft SQL Server to a specific transaction using the created backup. To do this, you will perform the following actions:

- Configure a backup job that will create a Microsoft SQL VM backup and copy transaction logs.
- Use the created VM image-level backup and transaction logs to recover a database on the Microsoft SQL Server to a specific transaction with Veeam Explorer for Microsoft SQL Server.

It is strongly recommended that you use a non-production Microsoft SQL Server with a sample database for this exercise.

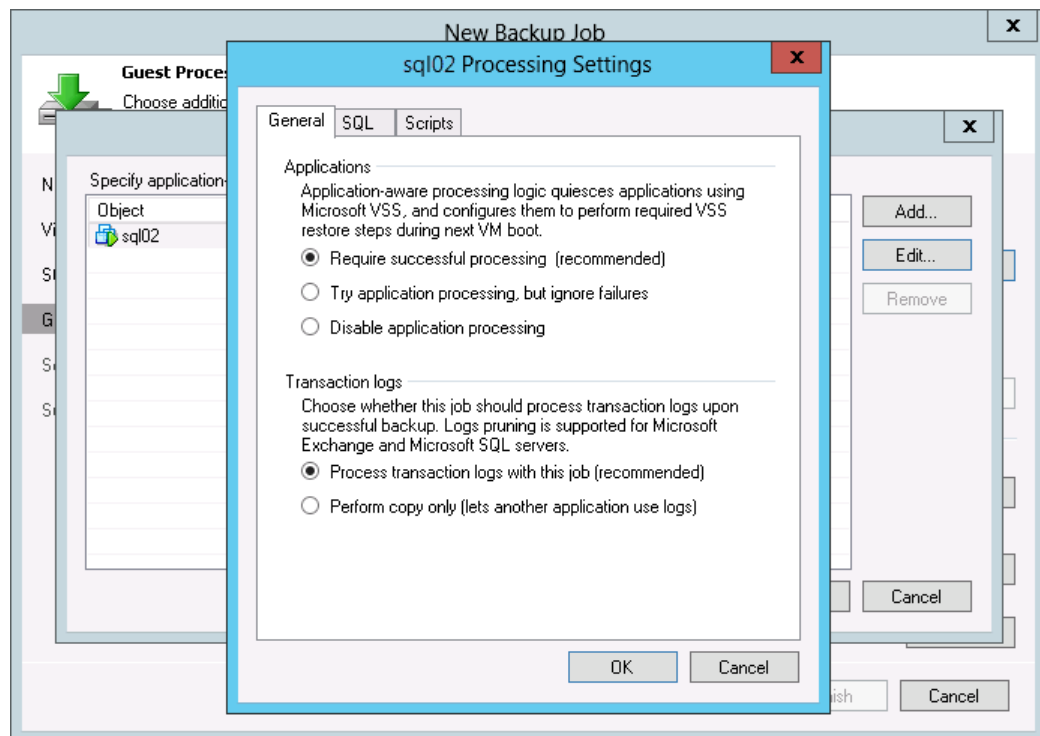
Prerequisites

Make sure that the *Full* or *Bulk-logged* recovery model is enabled for the database on the Microsoft SQL Server that you plan to back up. If the recovery model is set to *Simple*, Veeam Backup & Replication will not detect and process transaction logs.

Procedure

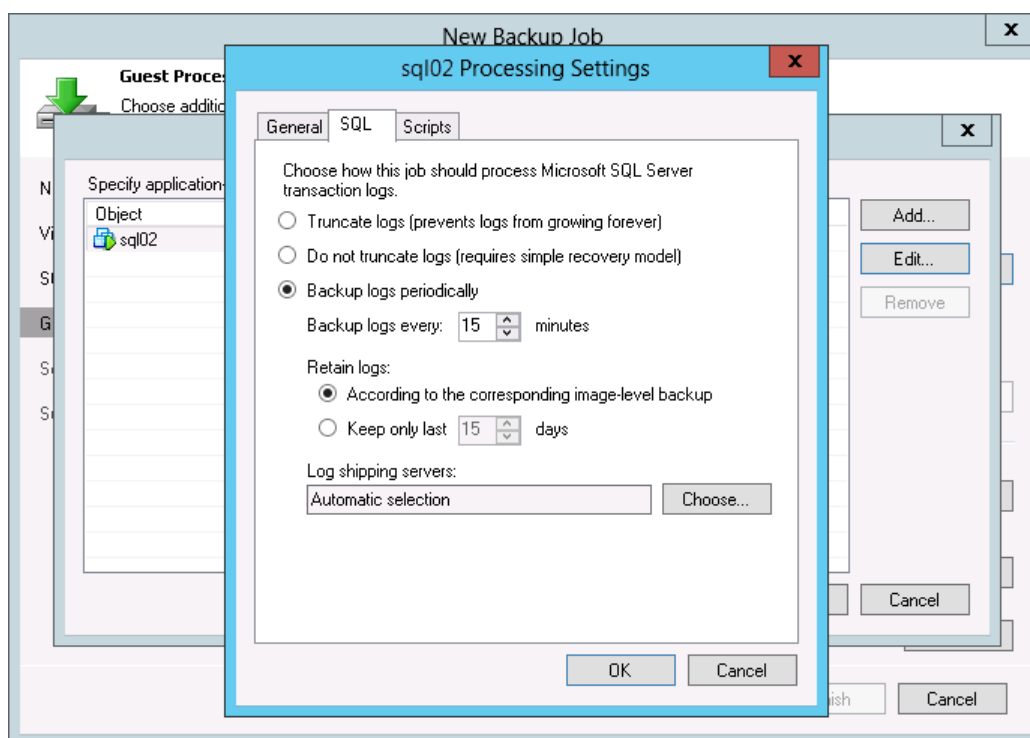
Step 1. Configure the backup job for a Microsoft SQL VM

1. In Veeam Backup & Replication, configure a backup job that processes a Microsoft SQL Server VM (see the [Performing Backup](#) exercise).
2. At the **Guest Processing** step of the **New Backup Job** wizard, select the **Enable application-aware processing** check box. In the **VM Guest OS credentials** section, specify a user account to connect to the VM guest OS. The user account must have the sysadmin privileges on the Microsoft SQL Server. In the opposite case, Veeam Explorer for Microsoft SQL Server will fail to automatically identify Microsoft SQL databases in the created VM backup.
3. Click **Applications**.
4. Select the Microsoft SQL Server VM in the list and click **Edit**.
5. In the **Transaction logs** section, make sure that the **Process transaction logs with this job** option is selected.

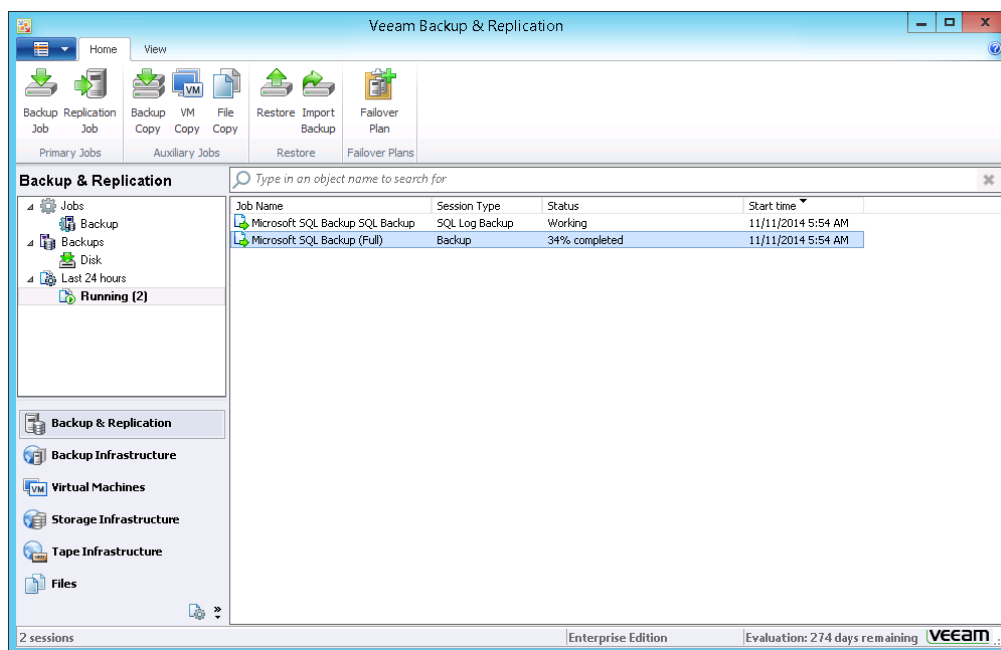


6. Click the **SQL** tab.
7. Select the **Backup logs periodically** option.
8. In the **Backup logs every <N> minutes** field, specify how often you want to ship transaction logs from the Microsoft SQL Server VM to the backup repository. By default, Veeam Backup & Replication ships transaction logs every 15 minutes.
9. In the **Retain logs** section, specify the retention policy for transaction logs. You can store them for a specific number of days or until a preceding image-level backup is removed from the backup chain.

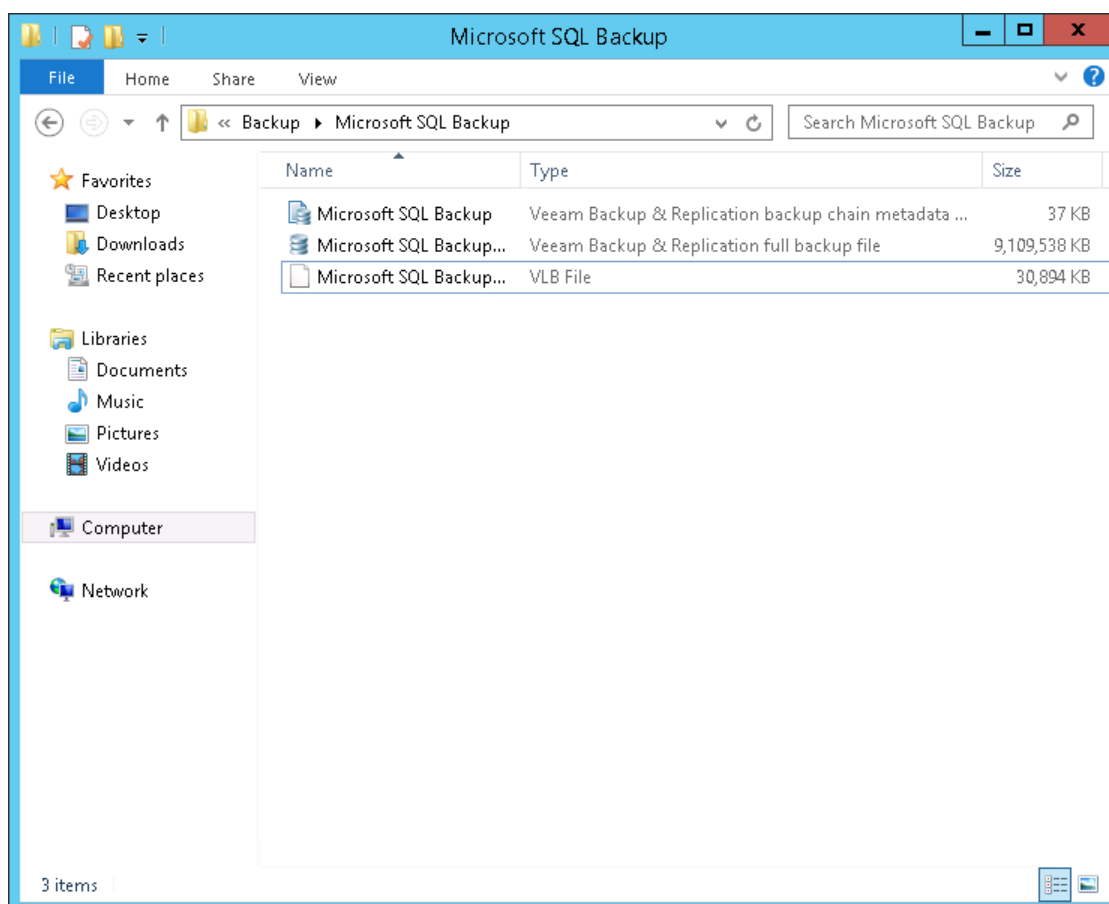
10. In the **Log shipping servers** section, leave the **Automatic selection** option selected. Veeam Backup & Replication will automatically identify the least loaded Microsoft Windows server in your backup infrastructure and use this server to ship transaction logs to the backup repository.



11. At the **Schedule** step of the wizard, define scheduling settings for the job. If you do not specify that the job must run automatically by the defined schedule, the backup job will be unable to ship transaction logs to the backup repository.
12. Finish working with the wizard and run the job to produce an image-level backup of the Microsoft SQL Server VM.
13. When you create a backup job that processes a Microsoft SQL Server VM and enable transaction log shipping, Veeam Backup & Replication creates two jobs: one processing the Microsoft SQL Server VM and the other one shipping transaction logs. In the inventory pane of the **Backup & Replication** view, expand the **Last 24 hours** node to see the created jobs.

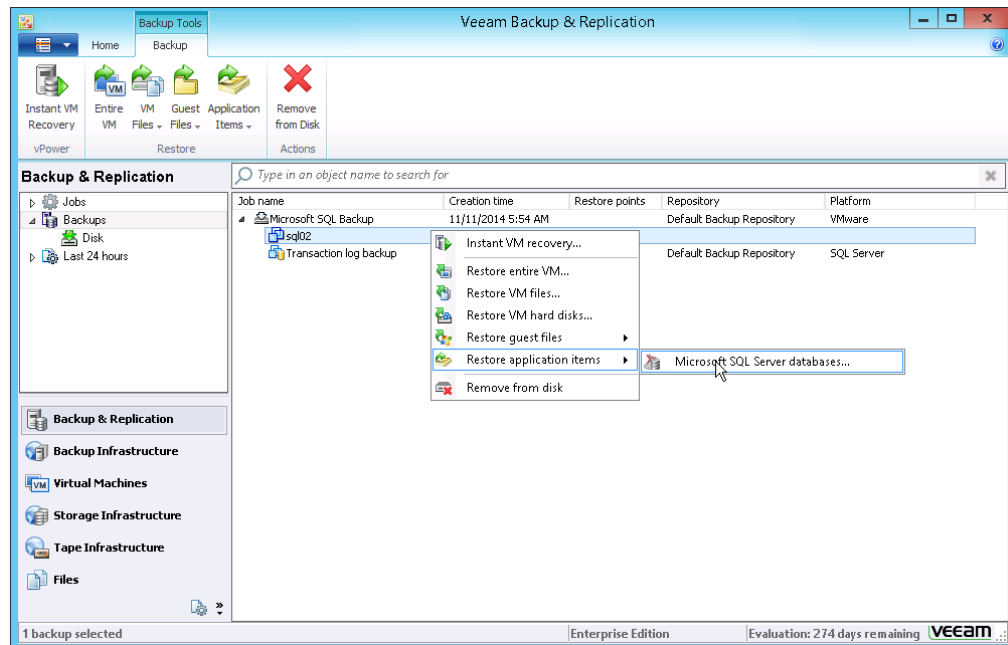


14. After the image-level backup has been created, perform some transaction on the database on the Microsoft SQL Server VM that you have backed up. For example, if you use a test database, you can manually run a simple Microsoft SQL script to insert a record into the database or drop a record.
15. Wait for the period of time that you have defined in the **Backup logs every <N> minutes** field. After this period has expired, Veeam Backup & Replication will ship transaction logs and store them in a file of VBL format on the target backup repository, next to the image-level backup of the Microsoft SQL Server VM.

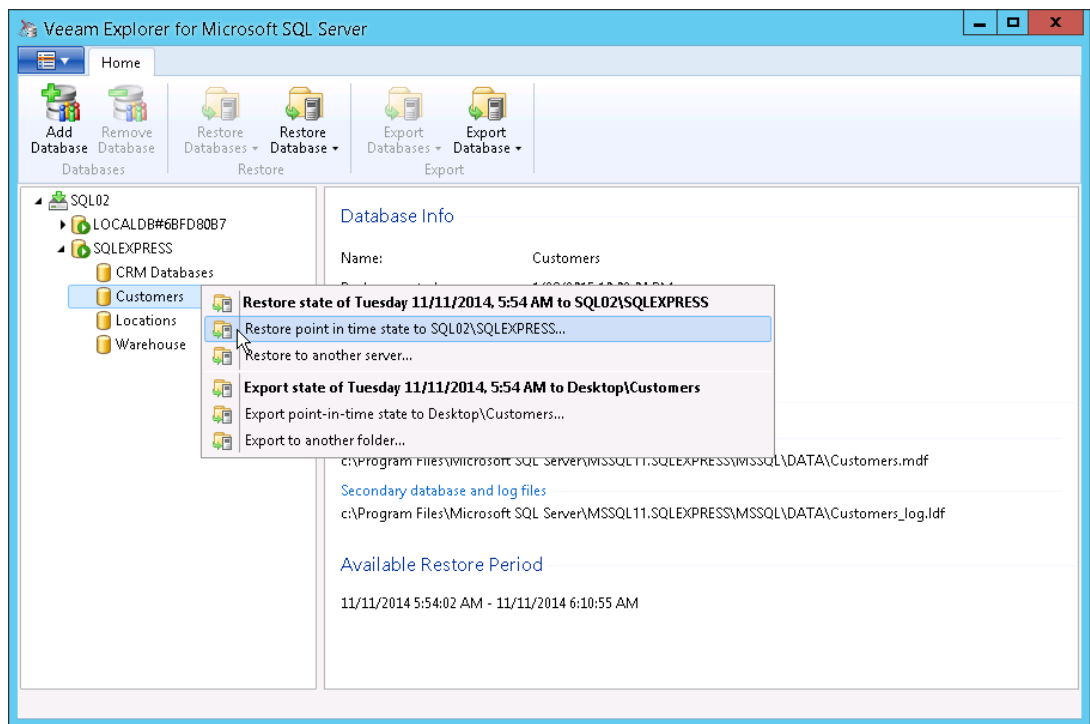


Step 2. Recover a database to a specific transaction

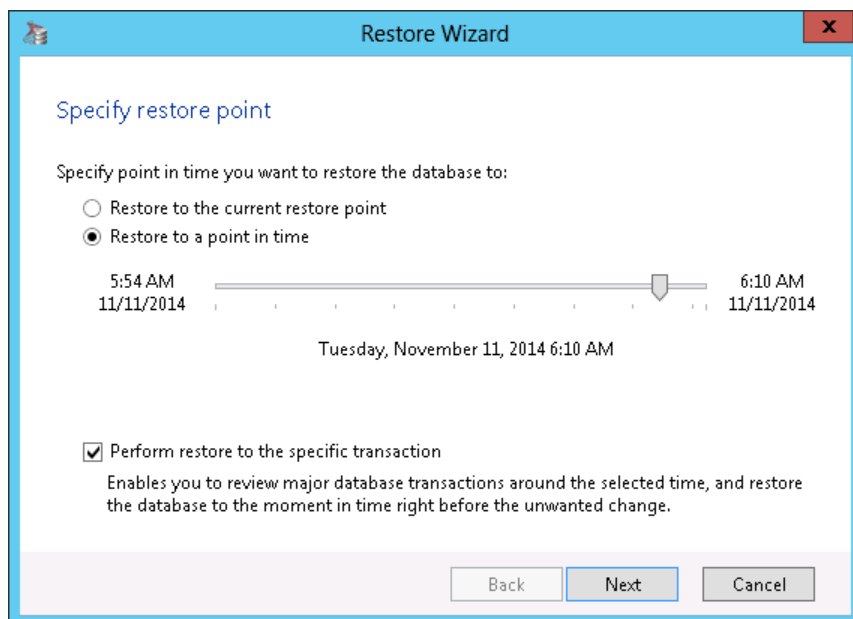
1. Open the **Backup & Replication** view.
2. In the inventory pane, click the **Backups** node.
3. In the working area, expand the backup job processing the Microsoft SQL Server VM, right-click the Microsoft SQL Server VM and select **Restore application items > Microsoft SQL Server databases**.



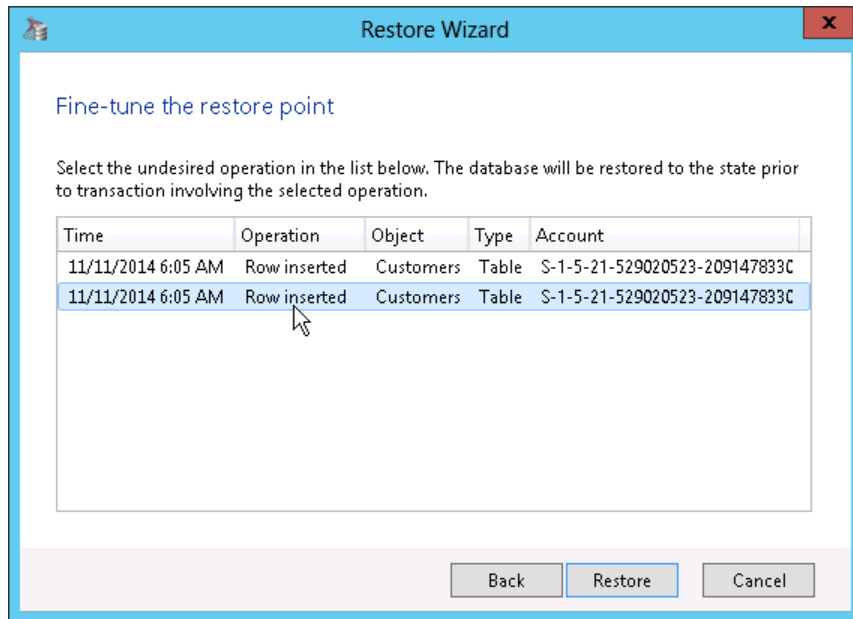
4. Pass through the steps of the **Microsoft SQL Server Database Restore** wizard: select a restore point and specify a restore reason. At the last step of the wizard, click **Finish** to start the recovery process. Veeam Backup & Replication will automatically mount the Microsoft SQL Server VM file system to the Veeam backup server, locate the Microsoft SQL database and attach it to a staging Microsoft SQL Server — Microsoft SQL Server on which the Veeam Backup & Replication database is deployed. After that, Veeam Backup & Replication will automatically open the database in Veeam Explorer for Microsoft SQL Server.
5. In the left pane of Veeam Explorer for Microsoft SQL Server, right-click the necessary database and select **Restore point-in-time state to <Microsoft SQL Server\Instance Name>**.



6. Veeam Backup & Replication will launch the **Restore** wizard. At the **Specify restore point** step of the wizard, select **Restore to a point in time**. Use the slider below to define the exact point in time to which you want to restore the database.
7. Select the **Perform restore to the specific transaction** check box and click **Next**.



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



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

Validation

Check the state of the restored database on the Microsoft SQL Server VM and make sure it has been restored to the necessary state.

Performing Full VM Restore

Insight into Full VM Restore

If a production VM has failed and you need to recover it from the backup, you can use one of the two options that Veeam Backup & Replication offers:

- Instant VM Recovery*, which uses the vPower technology to start a VM directly from a compressed and deduplicated backup file. Instant VM Recovery creates a "temporary spare" of the original VM and provides the minimum restore time possible (several seconds to several minutes).
- Full VM recovery, which actually recovers a full VM from the backup file and registers it on the target host. Though full VM recovery takes more time than Instant VM Recovery as you have to extract the VM image to the production storage, it actually recovers a failed VM on the production storage and provides full disk I/O performance.

You can restore a single VM or a multiple VMs at once, both to the original location or to a new location. VM(s) can be recovered to the latest state or any valid point in time.

When Veeam Backup & Replication creates a backup of a VM, it additionally stores information about the initial VM location to the backup file. As all initial VM settings are available, restore of a VM to the original location is extremely fast: you can do it basically in one click. Restore to the original location mitigates the risk of operator's errors: you do not have to provide any VM data during the restore process, so the chance to specify wrong settings is minimal.

* To learn about the Instant VM Recovery scenario, see [Veeam Backup & Replication User Guide](#).

Evaluation Case

In this exercise, you will restore a VM from the backup to its original location and power it on on the ESX(i) host.

Prerequisites

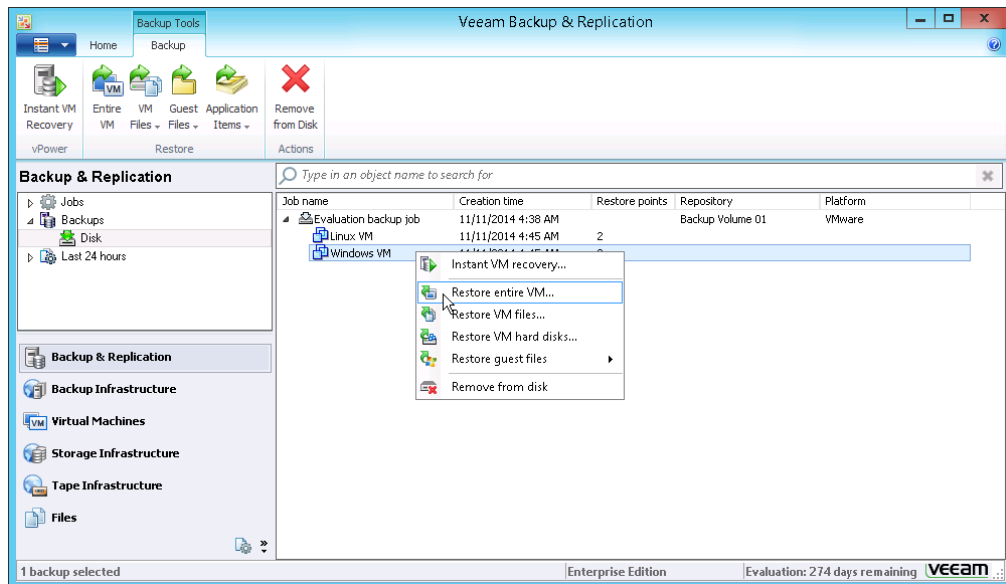
- You can restore a VM from any backup that has been successfully run at least once. Open the **Backup & Replication** view, select the **Backups** node in the inventory pane. Then expand the backup job and verify that there is at least one restore point available for the VM.
- VMs with Virtual Hardware version 8 are supported only by ESX(i) 5.x or later.

Important! When you restore a VM to the original location, Veeam Backup & Replication automatically deletes the initial VM. For safety's sake, make sure that you use a non-production VM for this exercise.


Procedure

To restore a full VM to its original location, do the following:

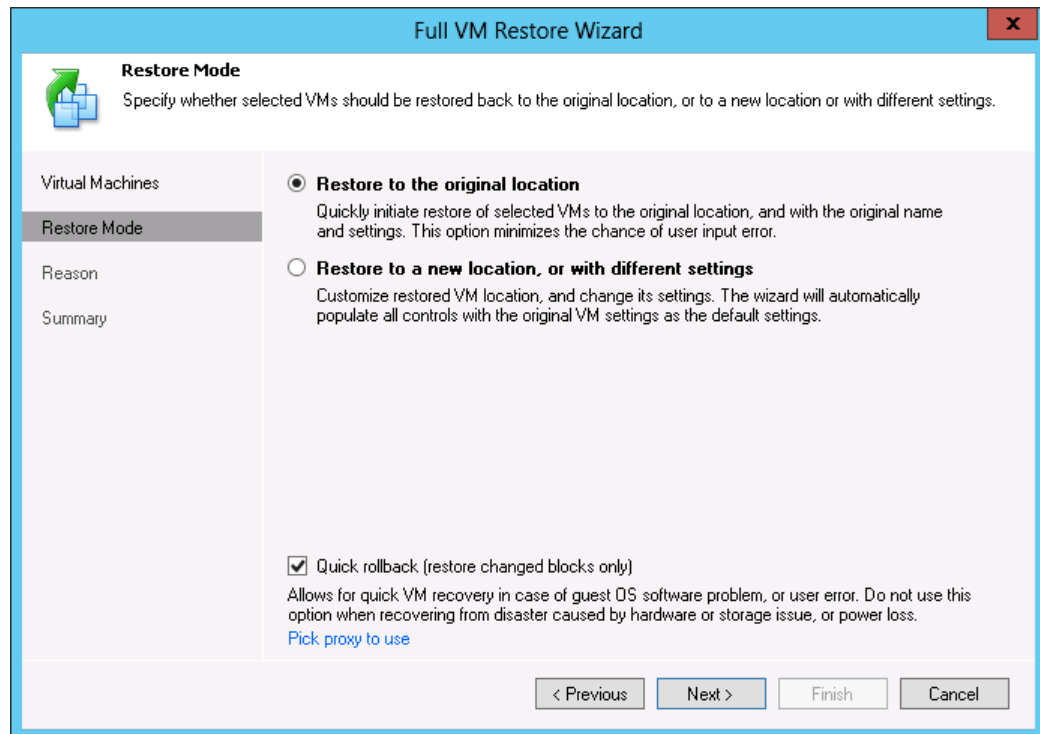
1. Open the **Backup & Replication** view.
2. Select the **Backups > Disk** node in the inventory pane. Expand the backup job in the working area, right-click the necessary VM in the corresponding backup job and select **Restore entire VM**.



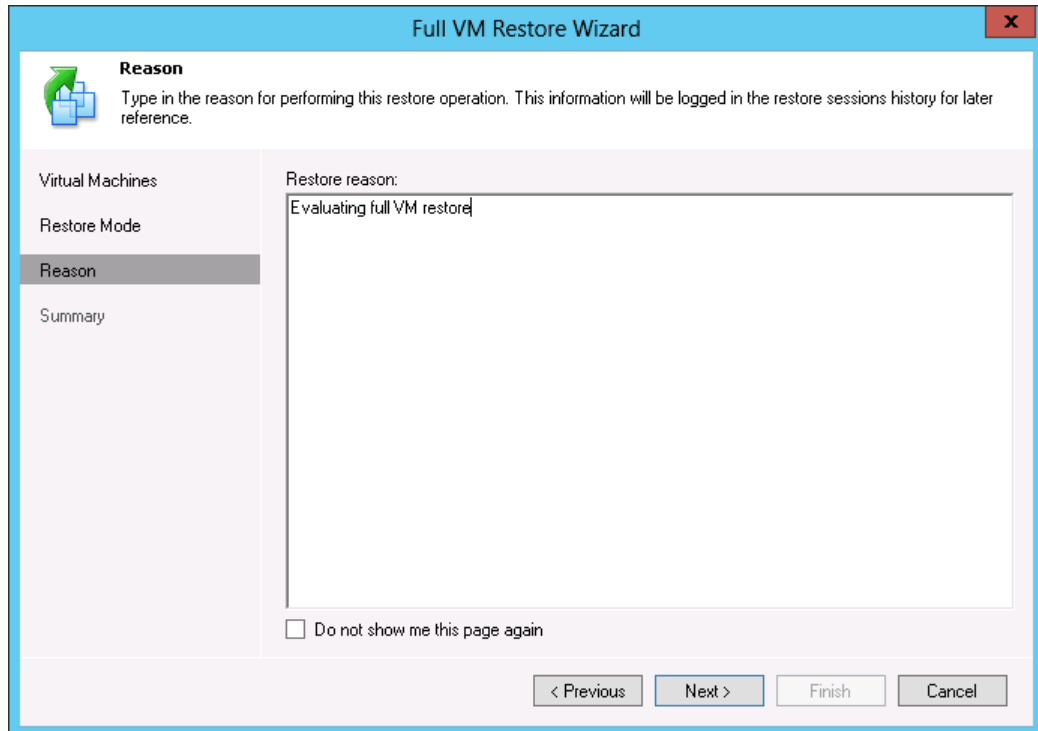
3. At the **Virtual Machines** step of the wizard, select the VM in the list, click **Point** on the right and choose the necessary restore point. If you select an incremental restore point, Veeam Backup & Replication will automatically restore data blocks from the full backup file and the chain of incremental backup files.



4. At the **Restore mode** step of the wizard, select **Restore to the original location**.
5. Click the **Pick proxy to use** link and make sure that **Automatic selection** is specified there. Veeam Backup & Replication will check settings of available proxies and select the most appropriate one for the job — the backup proxy that will enable the most efficient data transport to the source datastore. Veeam Backup & Replication first attempts to choose a backup proxy that uses the *Direct SAN Access* mode, then the backup proxy that uses the *Virtual Appliance* mode. If such proxies are not available, Veeam Backup & Replication selects the least loaded backup proxy that uses the *Network* mode.
6. Select **Quick restore**. Veeam Backup & Replication will perform incremental VM restore — retrieve from the backup file only those data blocks that have changes since the selected restore point was created and write these data blocks to the source datastore. This option lets you significantly speed up the restore process.



- At the **Reason** step of the wizard, specify the reason for restoring the VM.

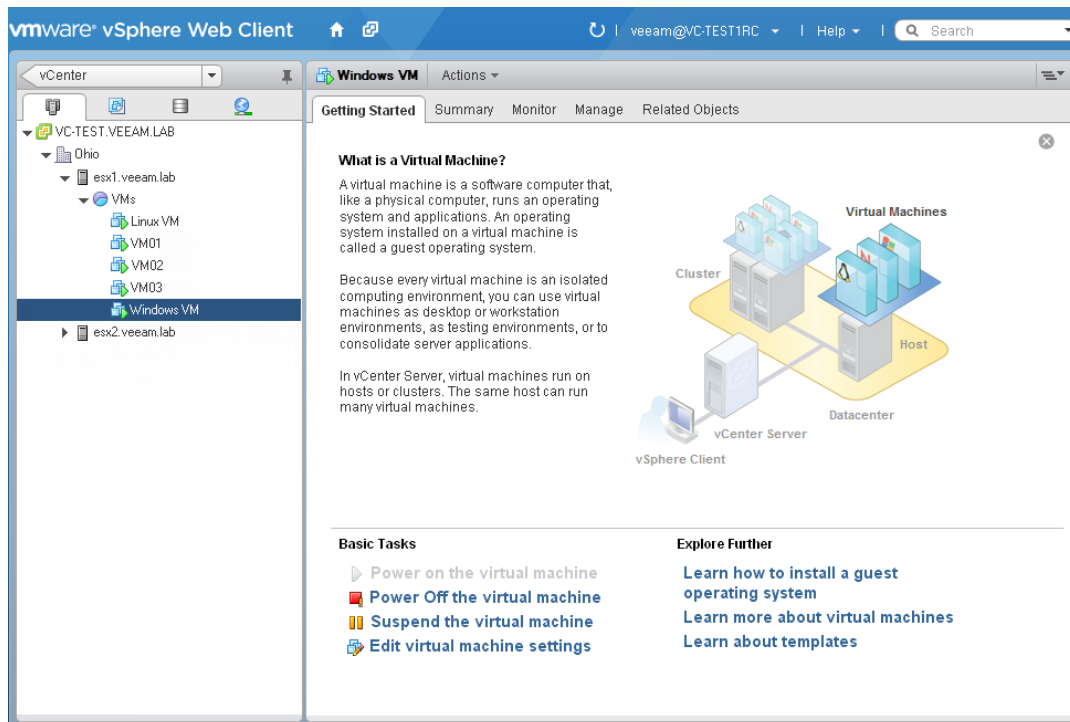


The screenshot shows the 'Full VM Restore Wizard' window, specifically the 'Reason' step. The window has a blue title bar and a sidebar on the left with four tabs: 'Virtual Machines', 'Restore Mode', 'Reason' (which is selected and highlighted), and 'Summary'. The main area contains a text box labeled 'Restore reason:' with the text 'Evaluating full VM restore' entered. Below the text box is a checkbox labeled 'Do not show me this page again'. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

- At the last step of the wizard, select the **Power on VM after restoring** check box; then click **Finish**.

Validation

Open the vSphere Client and make sure that the restored VM is running on on the original host.



The screenshot shows the vSphere Web Client interface. The left sidebar displays a tree view of the vCenter environment, including 'VC-TEST.VEEAM.LAB', 'Ohio', 'esx1.veeam.lab', and 'VMs'. The 'Windows VM' is selected. The main pane shows the 'Getting Started' tab, which includes a 'What is a Virtual Machine?' section with a diagram of a virtual machine architecture. The diagram shows a 'vCenter Server' connected to a 'vSphere Client', which manages a 'Cluster' of 'Hosts' running 'Virtual Machines'. Below the diagram, there are 'Basic Tasks' and 'Explore Further' sections with links to various actions and resources.

Restoring Guest OS Files

Insight into Guest OS Files Restore

Together with full VM restore, Veeam Backup & Replication offers Instant File-Level Restore (IFLR) that lets you recover guest OS files and folders from the VM backup. In contrast to competitive solutions, IFLR does not require a specific file-level backup. Restore of guest OS files, as well as all other restore options, is available from the same image-level backup. Neither should you extract the VM image from the backup prior to restore: VM guest OS files can be recovered on-the-fly, directly from a regular backup or from a VM replica. This makes the restore process extremely fast and does not require you to provision additional storage resources.

Veeam Backup & Replication provides three options for guest OS files recovery:

- Recovery from Windows-based file systems (FAT, NTFS and ReFS)
- Recovery from 15 most used file systems (including Linux, Unix, BSD and soon) with the multi-OS restore wizard
- Universal File-Level Recovery* for any OS or file system through leveraging the Instant VM Recovery technology*

You can recover VM guest OS files to the latest state or any valid point in time.

* To learn about the Instant VM Recovery scenario, see [Veeam Backup & Replication User Guide](#).

Evaluation Case

In this exercise, you will recover guest OS files and folders from the image-level VM backup and save them to the `C:\backup\restored` folder on the Veeam backup server.

The exercise covers two evaluation scenarios:

- **Scenario 1:** Recovering guest OS files from the backup of a Windows-based VM
- **Scenario 2:** Recovering guest OS files from the backup of a Linux-based VM or other non-Windows VM with the help of the multi-OS recovery wizard

Scenario 1. Restoring VM Guest OS Files (FAT, NTFS, ReFS)

When you recover guest OS files from Windows-based VMs, Veeam Backup & Replication mounts the content of a backup file directly to the Veeam backup server and displays the file tree in the built-in file browser. You can copy the files you need and save them locally or anywhere on the network. You can even point applications to restored files and work with these files as usual.

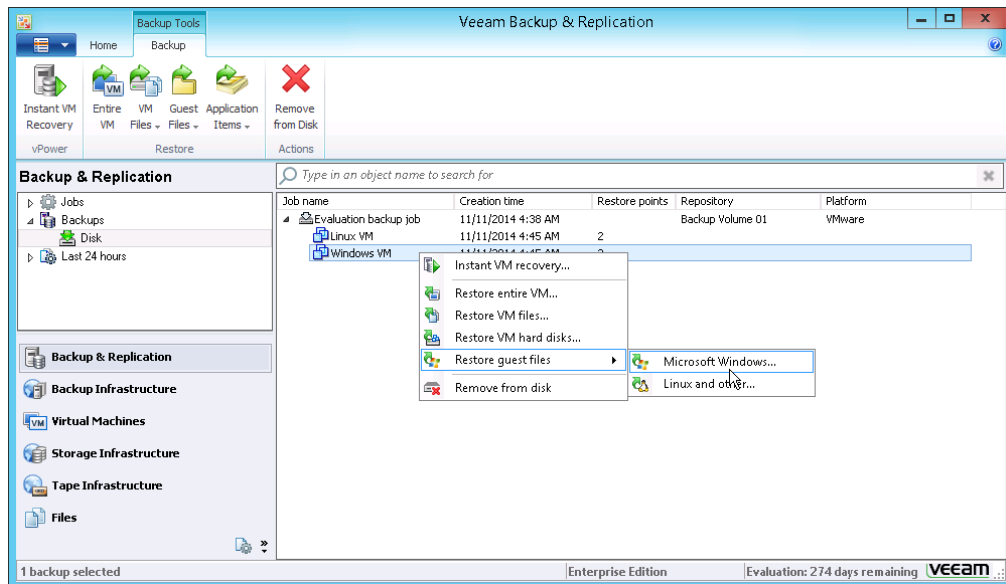
Prerequisites

- You can restore VM guest OS files from any backup that has been successfully run at least once. Open the **Backup & Replication** view, select the **Backups** node in the inventory pane. Then expand the backup job and check if there is at least one restore point available for the VM.
- Windows file restore mode supports Microsoft Windows file systems only (FAT, NTFS and ReFS). To restore files from VMs running other file systems, use the multi-OS file restore wizard.
- You cannot restore files from a VM if it is currently being backed up or replicated.

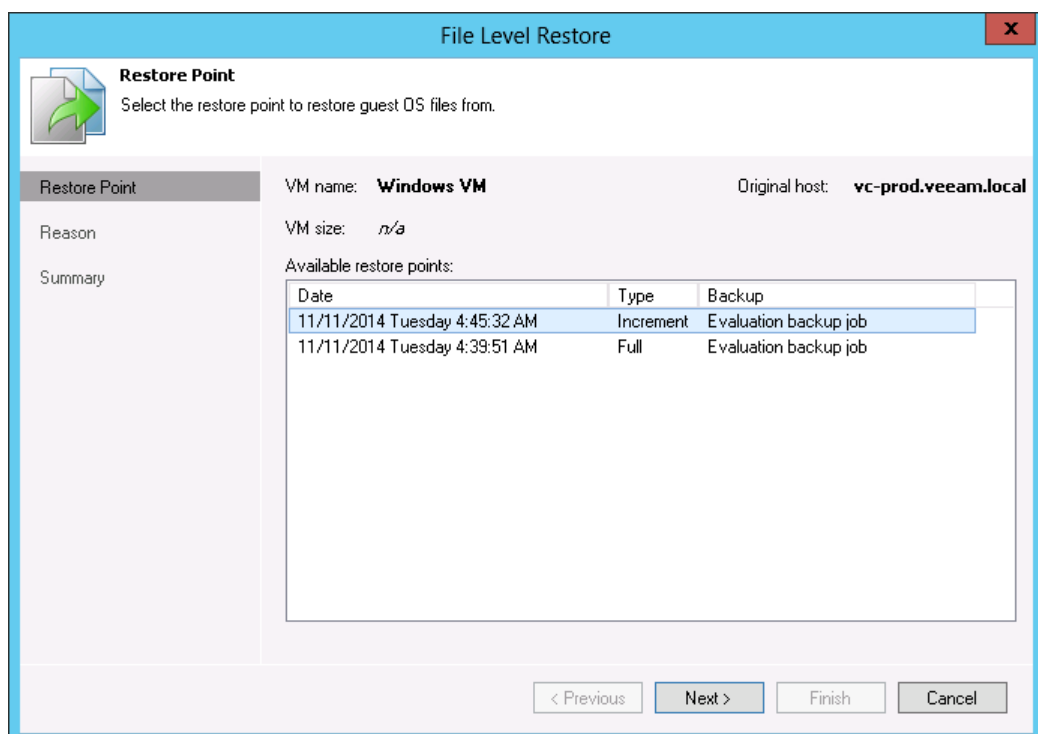
Procedure

To restore guest OS files from a Microsoft Windows VM:

1. Open the **Backup & Replication** view.
2. Click the **Backups > Disk** node in the inventory pane. Expand the backup job in the working area, right-click a necessary VM in the corresponding backup job and choose **Restore guest files > Microsoft Windows**.

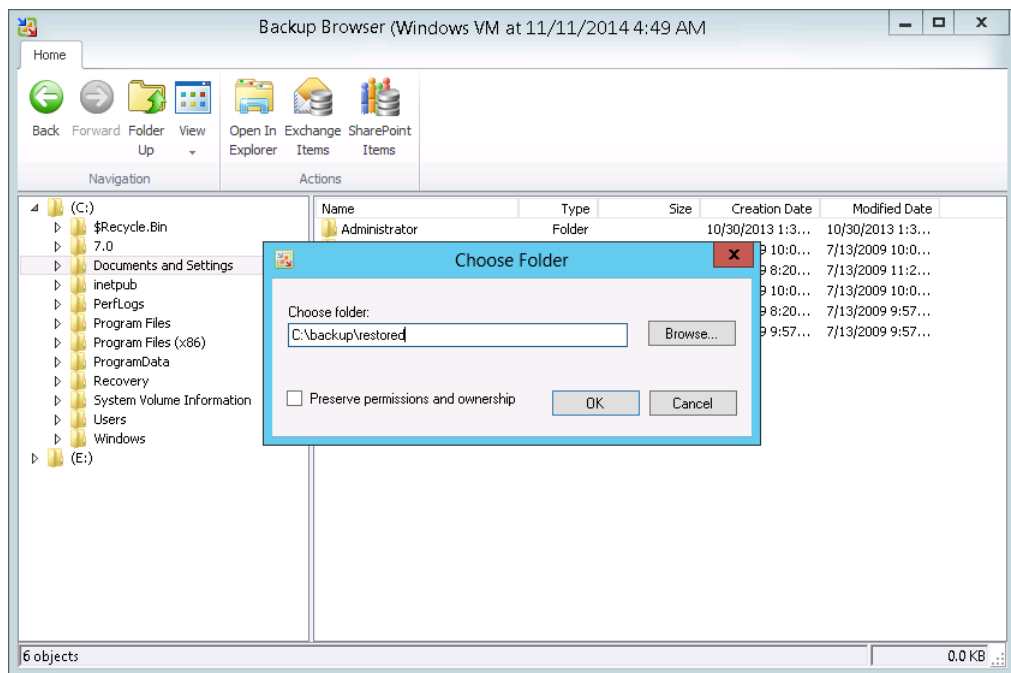


3. At the **Restore Point** step of the wizard, select the necessary restore point.



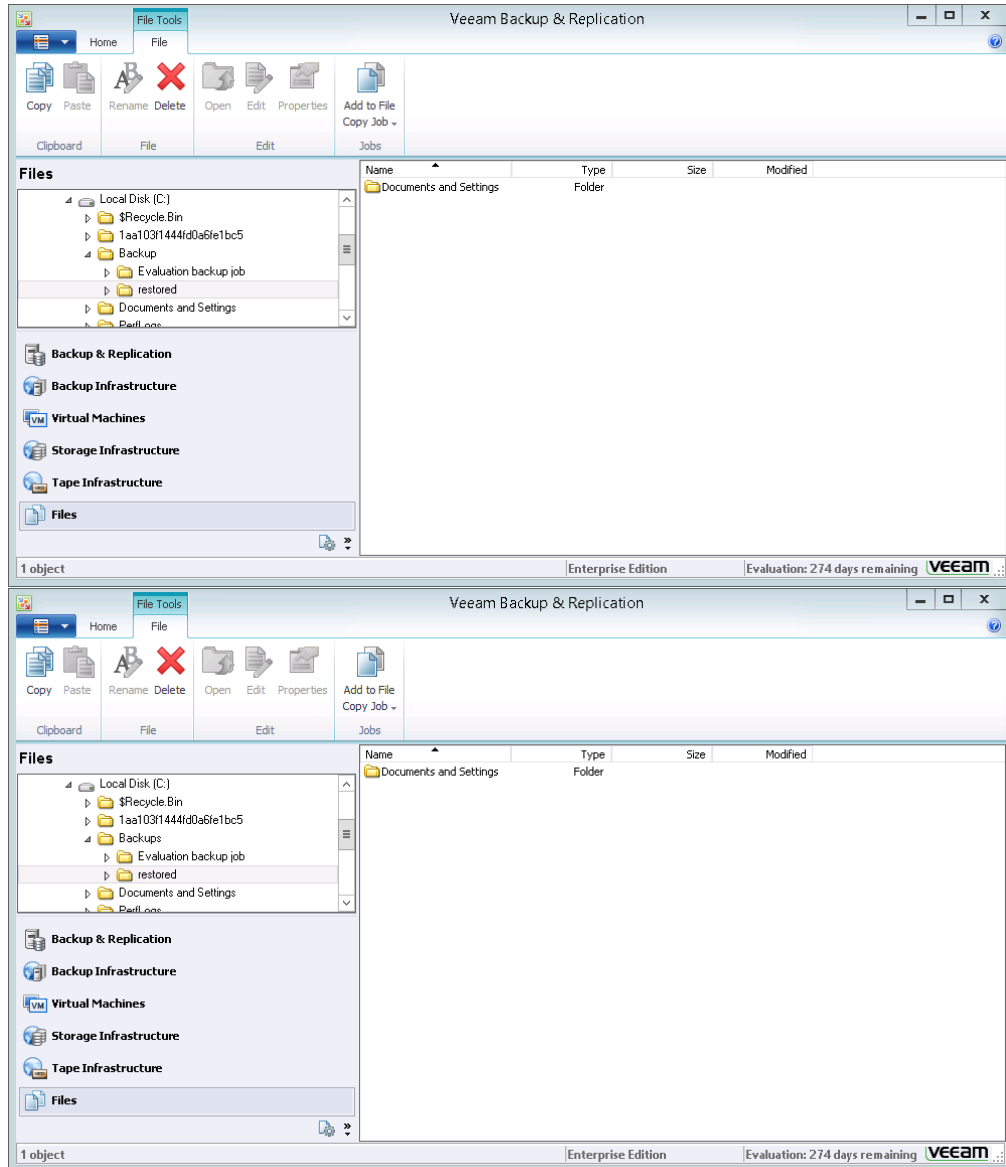
4. At the **Reason** step of the wizard, specify the reason for future reference.
5. Click **Next**. Then click **Finish**.

6. Veeam Backup & Replication will display a file browser with the file system tree of the VM. Right-click the necessary file or folder, select **Copy To** and enter C:\backup\restored in the **Choose folder** field.



Validation

1. Open the **Files** view.
2. In the inventory pane, expand the file tree under **This server**.
3. Select the `C:\backup\restored` folder and make sure that the restored files or folders are available there.



Scenario 2. Restoring VM Guest Files (Linux, Unix and so on)

While most of backup tools offer file-level recovery for a very small number of file systems, typically, Windows and rarely Linux, Veeam's multi-OS restore wizard enables you to recover guest OS files from 15 file systems such as Linux, Unix, BSD, MacOS and many others.

For file-level recovery, Veeam Backup & Replication uses a special FLR helper — a small virtual appliance based on the stripped-down Linux kernel. Whenever you perform file-level restore, Veeam Backup & Replication automatically starts the appliance and mounts VM disks to the FLR appliance as virtual hard drives. Virtual disks are mounted directly from backup files, without prior extraction of the backup content, which makes the restore process much faster in comparison with competitive solutions. You can then copy individual files and folders from VM disks to your local machine drive or a network share.

You can also let the FLR appliance function as an FTP server. In this case, users will be able to access restored files on the FTP using the IP address of the FLR appliance. Bear in mind that users will be able to access restored files on the FTP only while the file browser with restored files remains open on the Veeam backup server. Once the file browser is closed, the FLR appliance will be powered off.

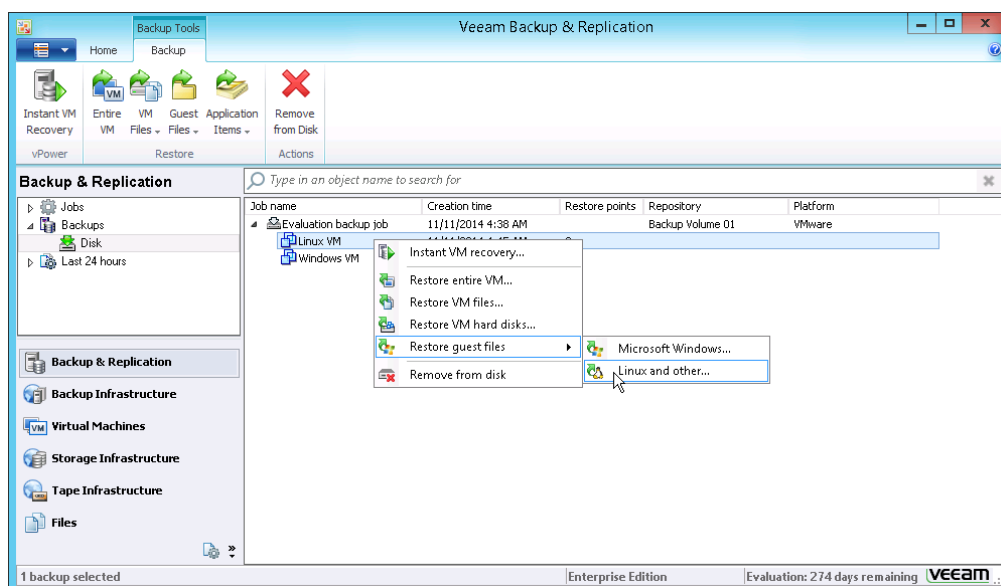
Prerequisites

- You can restore VM guest files from any backup that has been successfully run at least once. Open the **Backup & Replication** view, select the **Backups** node in the inventory pane. Then expand the backup job and check if there is at least one restore point available for the VM.
- Encrypted LVM volumes are not supported.


Procedure

To restore guest OS files from the backup of a Linux-based VM or other non-Windows VM:

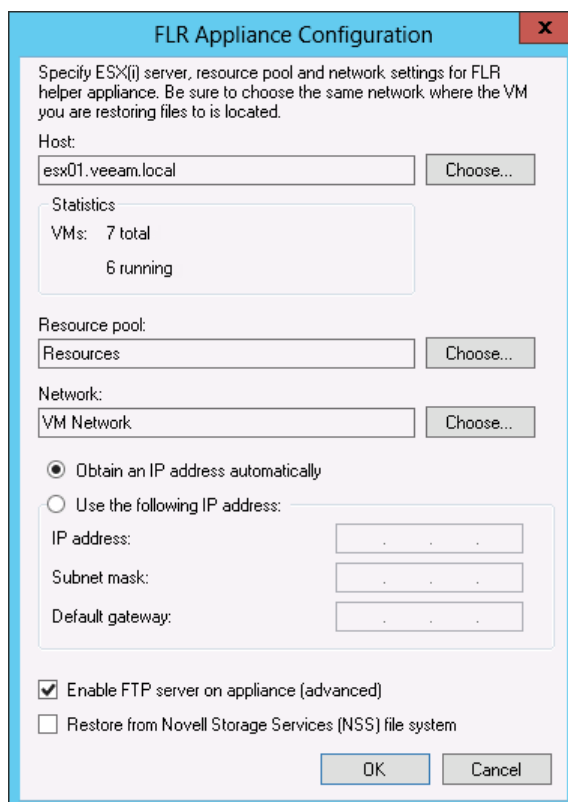
1. Open the **Backup & Replication** view.
2. Select the **Backups > Disk** node in the inventory pane. Expand the backup job in the working area, right-click the necessary VM in the corresponding backup job and choose **Restore guest files > Linux and other**.



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

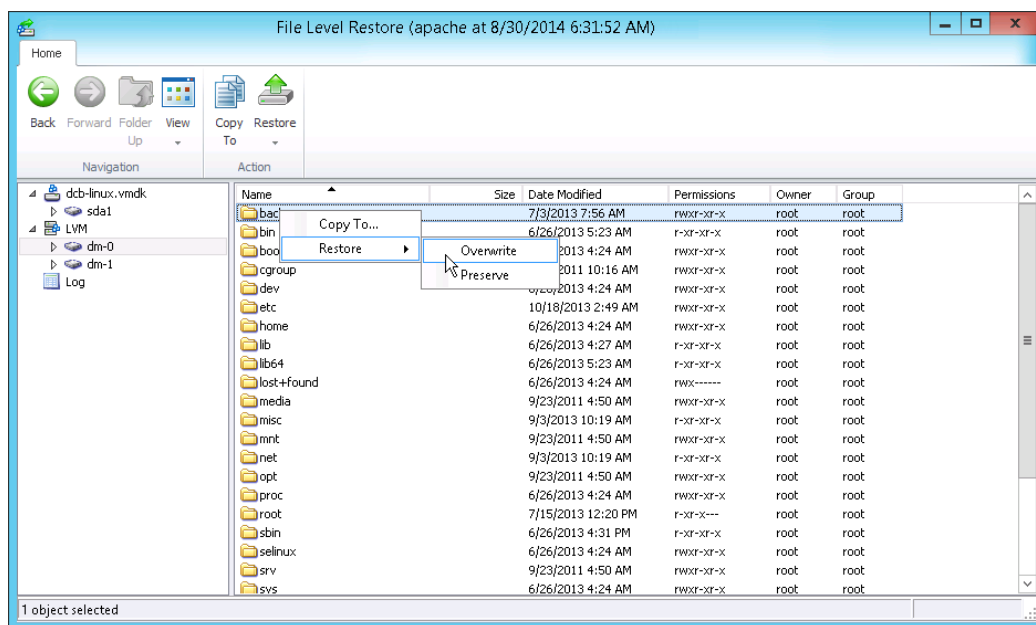


4. At the **Reason** step of the wizard, specify the reason for future reference.
5. At the last step of the wizard, click **Customize** to specify settings for the FLR appliance. Select an ESX(i) host, resource pool and the network on which the FLR appliance will run.
6. If you are restoring files from the NSS file system, select the **Restore from Novell Storage Services (NSS) file system** check box.



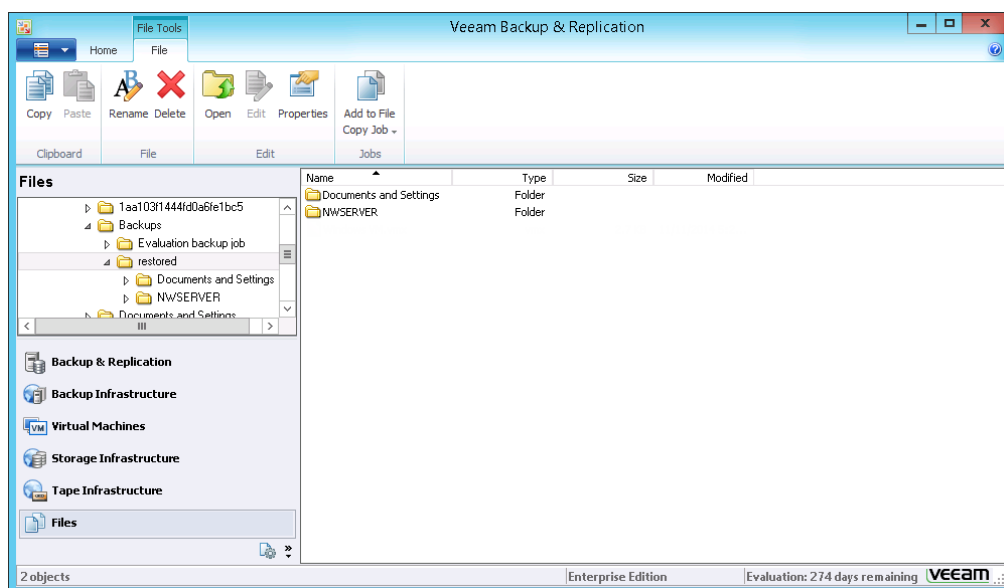
7. Click **Finish**. Note that the FLR appliance may take 30–40 seconds to boot.

8. Veeam Backup & Replication will display a file browser with the file system tree of the VM. Right-click a file or folder, select **Copy To** and enter C:\backup\restored in the **Choose folder** field.
9. If you are restoring files to the original Linux host, select the **Preserve permissions and ownership** check box so that all user access permissions for the file are preserved in the restored copy.
10. Click **Restore**.



Validation

1. Open the **Files** view.
2. In the inventory pane, expand the file tree under **This server**.
3. Select the C:\Backups\restored folder and make sure that the restored files or folders are available there.



Restoring VM Virtual Disk

Insight into VM Virtual Disk Restore

Veeam Backup & Replication enables 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 for some reason, for example, with a virus.

A VM virtual disk can be recovered to the latest state or any valid point in time. You can preserve the format of a recovered virtual disk, or convert it to thin or thick on the fly.

Evaluation Case

In this exercise, you will recover a corrupted virtual disk of a VM and attach it to another VM as a new drive.

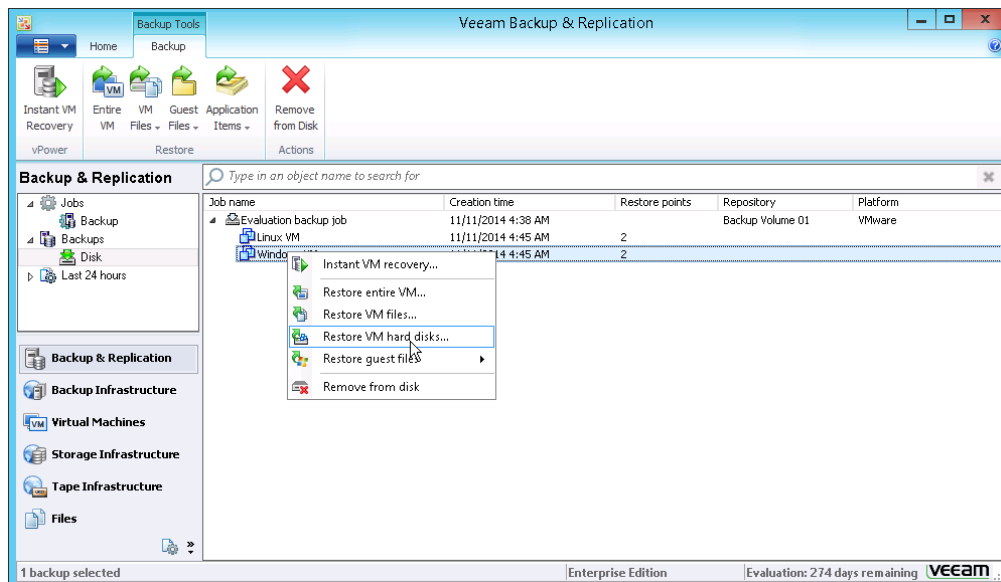
Prerequisites

- You can restore VM virtual disks from any backup that has been successfully run at least once. Open the **Backup & Replication** view, select the **Backups** node in the inventory pane. Then expand the backup job and check if 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. For this reason, it is recommended to stop all active processes on the VM for the restore period.

Procedure

To restore a VM virtual disk:

1. Open the **Backup & Replication** view.
2. Select the **Backups > Disk** node in the inventory pane. Expand the backup job in the working area, right-click a necessary VM in the corresponding backup job and choose **Restore VM hard disks**.



- At the **Restore Point** step of the wizard, select the necessary restore point. If you select an incremental restore point, Veeam Backup & Replication will automatically restore data blocks from the full backup file and the chain of incremental backup files.

The screenshot shows the 'Hard Disk Restore Wizard' window at the 'Restore Point' step. The left sidebar has a tree view with 'Virtual Machine', 'Restore Point' (selected), 'Disk Mapping', 'Reason', and 'Summary'. The main area displays VM details: 'VM name: Windows VM', 'Original host: vc-prod.veeam.local', and 'VM size: 40.0 GB'. Below this, a table titled 'Available restore points:' lists two points:

Date	Type
11/11/2014 Tuesday 4:45:32 AM	Increment
11/11/2014 Tuesday 4:39:51 AM	Full

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

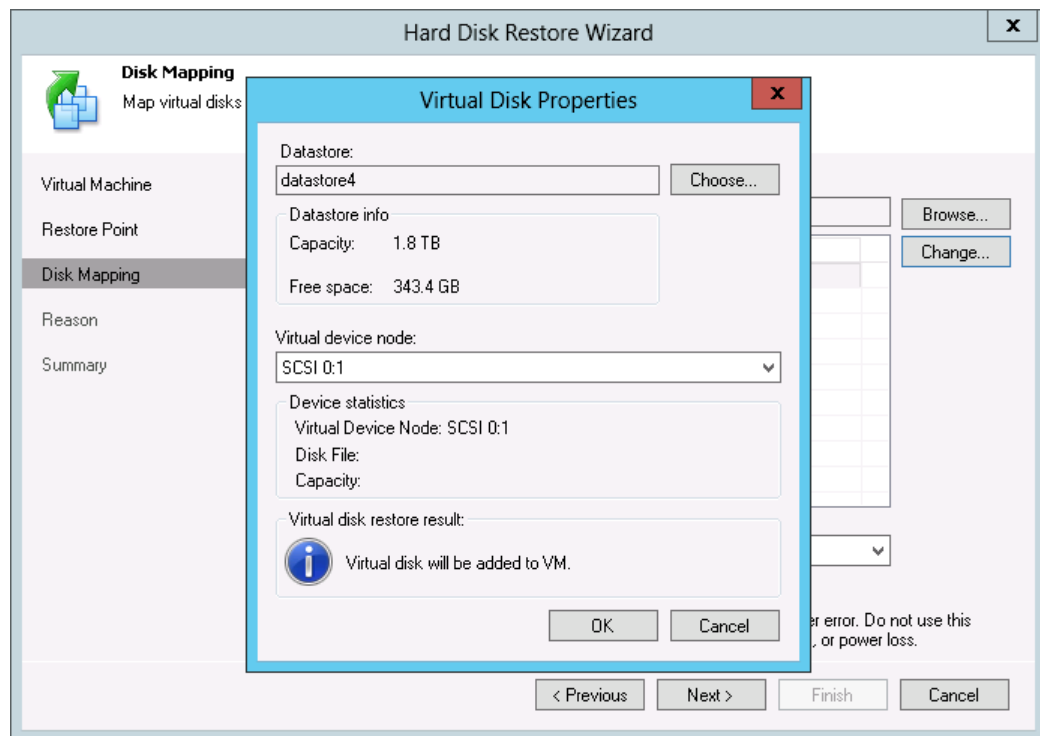
- At the **Disk Mapping** step of the wizard, click **Browse** and select the VM to which the restored hard disk should be attached.
- Select a check box next to the virtual hard disk that you want to restore.
- To change the disk format, select the required option from the **Restore disks** list: same as on the original VM, force thin or force thick.

The screenshot shows the 'Hard Disk Restore Wizard' window at the 'Disk Mapping' step. The left sidebar has a tree view with 'Virtual Machine', 'Restore Point', 'Disk Mapping' (selected), 'Reason', and 'Summary'. The main area displays 'Virtual machine name: Windows VM' with a 'Browse...' button. Below is a table for mapping virtual disks:

Virtual disk	Virtual Device Node	Datastore
<input checked="" type="checkbox"/> Windows VM.vmdk	SCSI 0:0	datastore4

Below the table, there is a 'Restore disks:' dropdown menu set to 'As on original VM (recommended)'. There is also an unchecked checkbox for 'Quick rollback (restore changed blocks only)' with a descriptive text: '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.' At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

7. 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**.



8. At the **Reason** step of the wizard, specify the reason for future reference.
9. At the last step of the wizard, select the **Power on VM after restoring** check box. Then click **Finish**.

Validation

Open the vSphere Client and make sure that the target VM is powered on and a new hard disk is attached to it.

Restoring VM Files

Insight into VM Files Restore

Veeam Backup & Replication lets you restore the VM as the whole or restore specific VM files: VMDK, VMX and others. The latter recovery scenario can be used, for example, if one of your VM files is missing or is corrupted and you need to bring it back. With Veeam Backup & Replication, you can restore the required VM file directly from the image-level backup, without prior de-staging of the VM image from the backup file.

VM files can be recovered to the latest state or to any valid point in time; you can bring them to the original location or to a new location.

Evaluation Case

In this exercise, you will recover a VM configuration file, VMX, and store to the C:\backup\restored folder on the Veeam backup server.

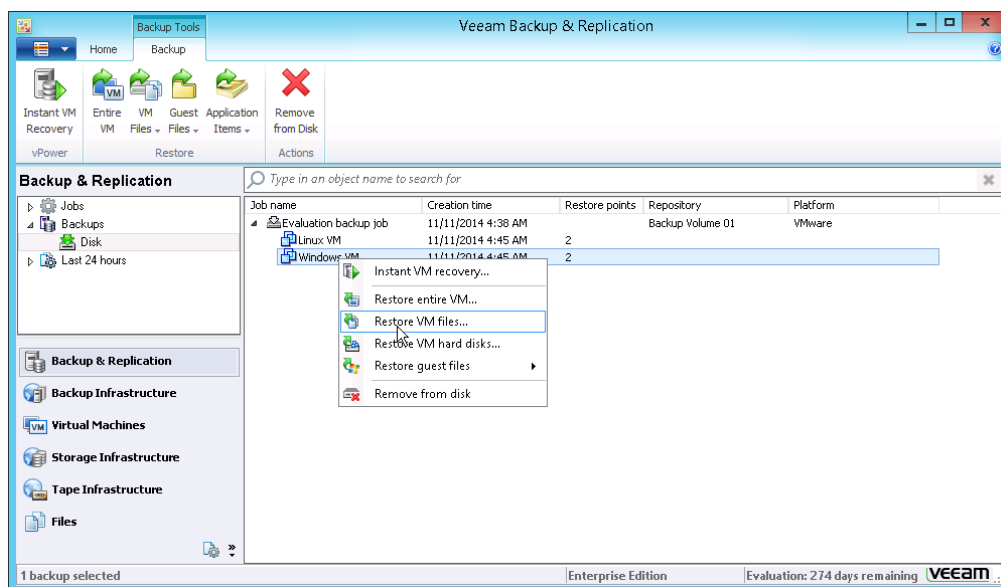
Prerequisites

You can restore VM files from any backup that has been successfully run at least once. Open the **Backup & Replication** view, select the **Backups** node in the inventory pane. Then expand the backup job and check if there is at least one restore point available for the VM.

Procedure

To restore a VM configuration file, do the following:

1. Open the **Backup & Replication** view.
2. Select the **Backups > Disk** node in the inventory pane. Expand the backup job in the working area, right-click a necessary VM in the corresponding backup job and choose **Restore VM files**.



- At the **Restore Point** step of the wizard, select the necessary restore point. If you select an incremental restore point, Veeam Backup & Replication will automatically restore data blocks from the full backup file and the chain of incremental backup files.

The screenshot shows the 'Restore Wizard' window at the 'Restore Point' step. The title bar says 'Restore Wizard'. Below the title bar, there's a sub-header 'Restore Point' and a description 'Select the restore point to restore VM from.' To the right of the description is a green arrow icon. Below this, there are fields for 'VM name: Windows VM' and 'Original host: vc-prod.veeam.local'. Below these is 'VM size: 40.0 GB'. Then, 'Available restore points:' is followed by a table with two columns: 'Date' and 'Type'. The table has two rows: '11/11/2014 Tuesday 4:45:32 AM' with 'Increment' type, and '11/11/2014 Tuesday 4:39:51 AM' with 'Full' type. At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

Date	Type
11/11/2014 Tuesday 4:45:32 AM	Increment
11/11/2014 Tuesday 4:39:51 AM	Full

- At the **Restore Destination** step of the wizard, select **This server**. Use the **Host Summary** button to view information on storage resources.
- Specify a path to the folder on the selected host where files should be restored, for example: `C:\backup\restored`.
- In the **VM files to restore** section, select a check box next to the VMX file.


The screenshot shows the 'Restore Wizard' window at the 'Restore Destination' step. The title bar says 'Restore Wizard'. Below the title bar, there's a sub-header 'Restore Destination' and a description 'Choose server and folder where VM files should be restored, and pick files to restore.' To the right of the description is a green arrow icon. Below this, there's a 'Destination:' section with a dropdown menu showing 'This server (Microsoft Windows Server)' and a 'Host Summary...' button. Below that is a 'Path to folder:' section with a text box containing 'C:\Backup\restored' and a 'Browse...' button. Then, there's a 'VM files to restore:' section with a table listing files and their sizes. The table has two columns: 'Name' and 'Size'. The files listed are 'Windows VM.vmx' (2.7 KB), 'Windows VM.vmx' (0.3 KB), 'Windows VM.nvram' (8.5 KB), 'Windows VM.vmdk' (0.6 KB), and 'Windows VM-flat.vmdk' (40.0 GB). The first file is checked. To the right of the table are 'Select All' and 'Clear All' buttons. At the bottom are three buttons: '< Back', 'Next >', and 'Cancel'.

Name	Size
<input checked="" type="checkbox"/> Windows VM.vmx	2.7 KB
<input type="checkbox"/> Windows VM.vmx	0.3 KB
<input type="checkbox"/> Windows VM.nvram	8.5 KB
<input type="checkbox"/> Windows VM.vmdk	0.6 KB
<input type="checkbox"/> Windows VM-flat.vmdk	40.0 GB

- At the **Reason** step of the wizard, specify the reason for future reference and click **Next**.
- Click **Finish** to restore the VM file.

Validation

1. Open the **Files** view.
2. In the inventory pane, expand the file tree under **This server**.
3. Select the `C:\backup\restored` folder and make sure that the restored VMX file is available there.



Creating Backup Copy

According to the 3-2-1 backup rule, you must adhere to the following requirements when building your backup plan:

- You must have at least three copies of data in different locations: production data, backup and its copy.
- You must use two different types of media to store your data, for example, disk storage and tape.
- You must keep at least one copy of your data offsite, for example, in the cloud or in the remote site.

With the backup copying capability in Veeam Backup & Replication, you can create several instances of the same backup file and copy them to secondary backup repositories that perform the role of the long-term storage. Secondary backup repositories can be located in the same site as the source one or can be created offsite. The backup copy file has the same format as the primary backup so you can restore necessary data directly from it in case a disaster strikes.

During the backup copying process, Veeam Backup & Replication does not simply copy a backup file from one backup repository to another. Instead, Veeam Backup & Replication retrieves data blocks necessary to create a restore point as of the latest point in time and copies this data to the target backup repository. The backup chain produced on the target backup repository is forever forward incremental: the first restore point in the chain is a full backup and all subsequent restore points are incremental backups.

The backup copy process is job-driven. When you create a backup copy job, you define what backup file you want to copy, the target repository for storing the copy, retention policy and other settings for the copying process. The backup copy job supports the GFS retention scheme, allowing you to design a long-term archiving plan.

Unlike the backup job that typically runs on a specific schedule, the backup copy job runs continuously, in cycles. By default, a new backup copy cycle begins every day; however, you can specify any time interval needed. At the beginning of every backup copy interval, Veeam Backup & Replication checks the source backup repository: if a new restore point has been added to the primary backup chain, Veeam Backup & Replication automatically copies it to the target backup repository. After that, the backup copy job is put on hold until a new backup copy interval begins and a new point appears on the source backup repository.

Evaluation Case

In this exercise, you will create a copy of the backup file that has been created in the [Performing Backup](#) exercise with the evaluation backup job. The backup file will be copied from the source backup repository and written to the target backup repository created onsite.

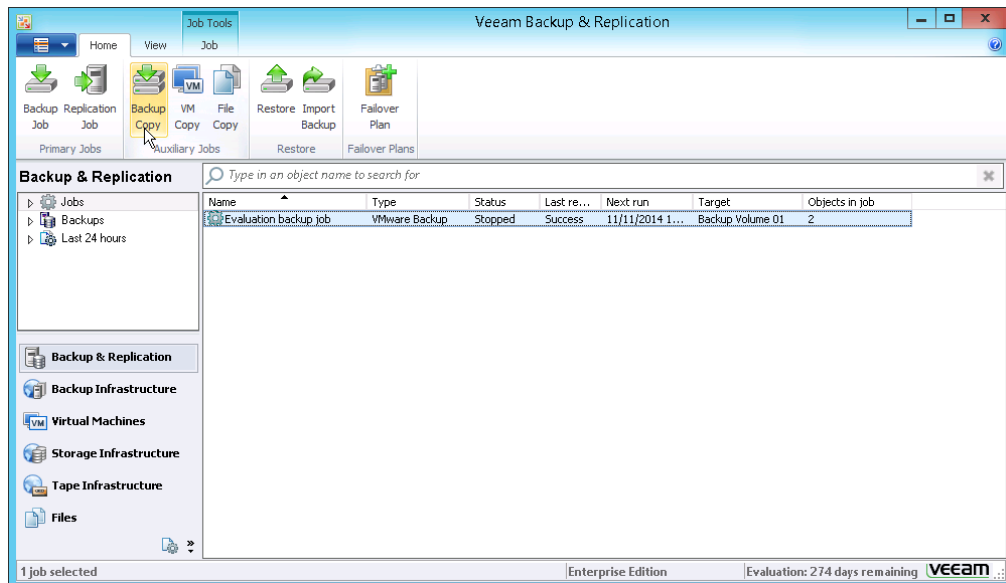
Prerequisites

- The source and target backup repositories that will take part in the backup copy process should be added to the Veeam Backup & Replication console.
- The primary backup job must be completed with the *Success* or *Warning* status and the resulting backup file that you plan to copy must reside on the source backup repository.

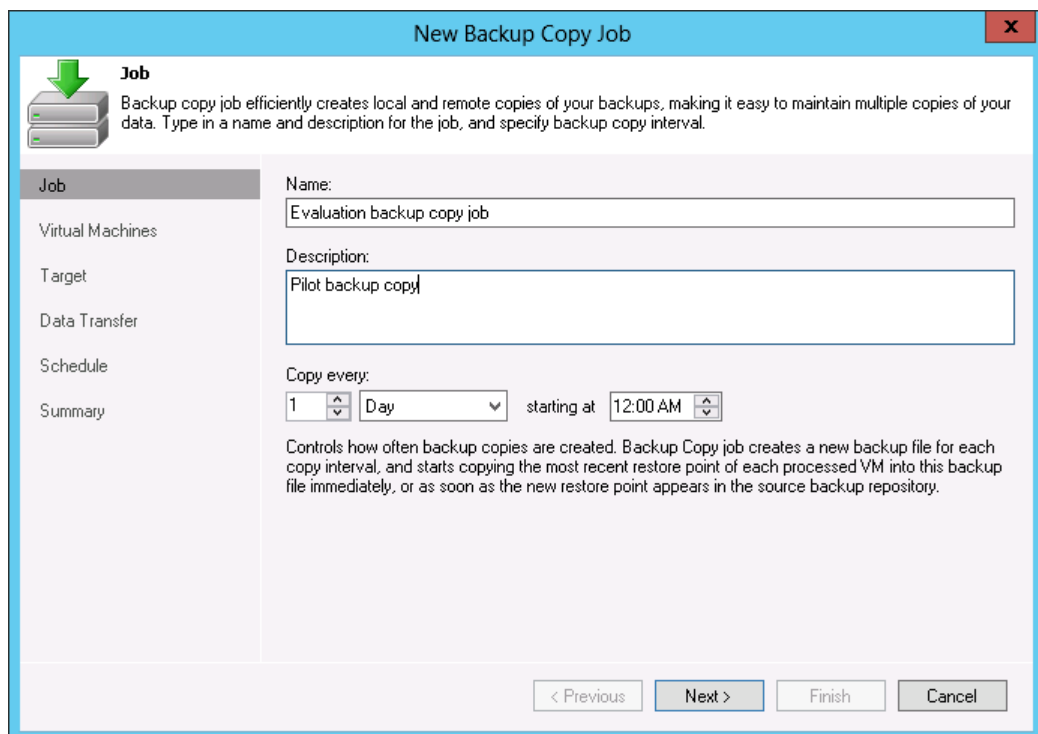
Procedure

To create a copy of the backup file:

1. Open the **Backup & Replication** view.
2. On the **Home** tab, click the **Backup Copy Job** button.

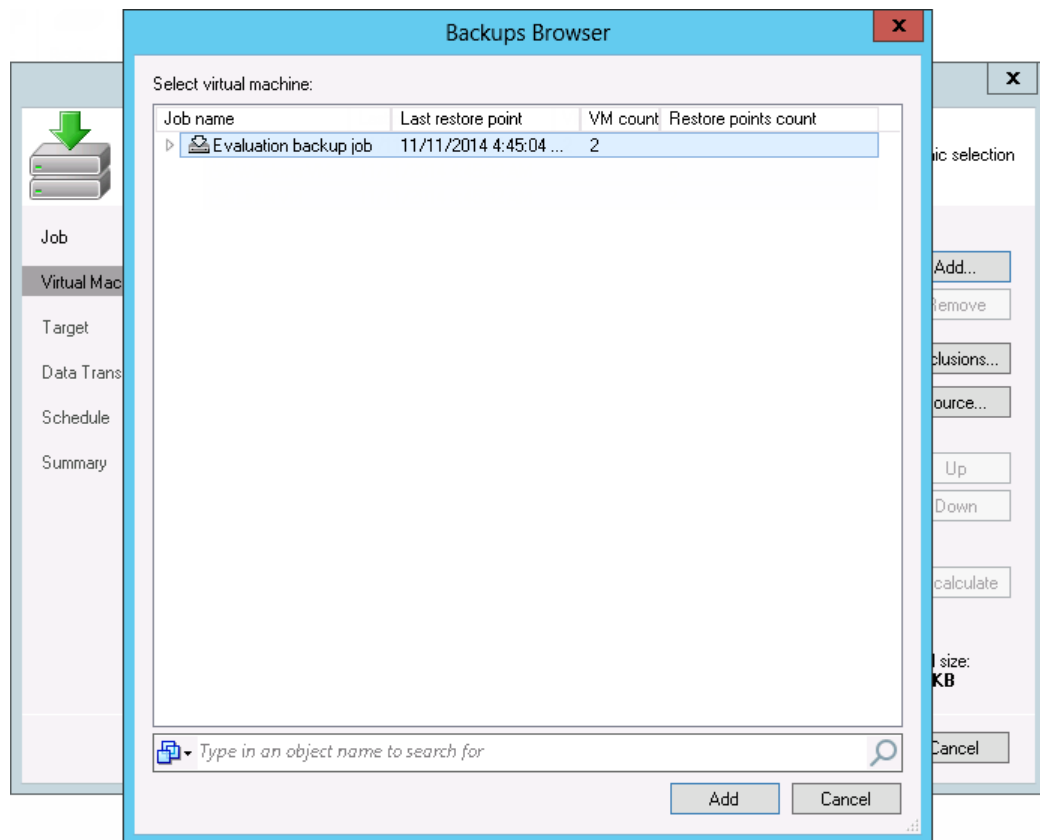


3. Specify a name for the created backup copy job.
4. In the **Copy every** field, specify a synchronization time interval, or the length for the backup copy cycle. At the beginning of every synchronization interval, Veeam Backup & Replication checks the source backup repository. If a new restore point has appeared, Veeam Backup & Replication copies this restore point to the target backup repository.



5. At the **Virtual Machines** step of the wizard, click **Add** and select **From Jobs**.

6. In the list of jobs, select the backup job that is used to create the backup you want to copy. Veeam Backup & Replication will monitor the source backup repository on which the backup job stores the resulting backup file. If Veeam Backup & Replication detects a new restore point on the source backup repository at the beginning of a new synchronization interval, Veeam Backup & Replication will automatically trigger a new backup copy cycle and copy a new restore point to the target backup repository.



7. At the **Target** step of the wizard, select the target backup repository from the **Backup repository** list.
8. In the **Restore points to keep** field, specify the number of restore points you want to retain. By default, Veeam Backup & Replication keeps 7 restore points on the target backup repository.

New Backup Copy Job

Target
Specify the target backup repository, amount of most recent restore points to keep, and retention policy for full backups. You can use map backup functionality to seed the backup files.

Job
Virtual Machines
Target
Data Transfer
Schedule
Summary

Backup repository:
Default Backup Repository (Created by Veeam Backup) Map backup

82.0 GB free of 99.7 GB

Restore points to keep: 7

☐ Keep the following restore points for archival purposes

Weekly backup: 4 Sunday 22:00 Schedule...

Monthly backup: 0 First Sunday of the month

Quarterly backup: 0 First Sunday of the quarter

Yearly backup: 0 First Sunday of the year

Advanced settings include health check and compact schedule, notifications settings, and automated post-job activity options. Advanced

< Previous Next > Finish Cancel

- During the backup copying process, backup data can be transferred directly, from the source backup repository to the target backup repository, or via a pair of WAN accelerators. The latter scenario is recommended for copying backups offsite or over slow network connections. To learn more about WAN acceleration, see [Veeam Backup & Replication User Guide](#).

In this exercise, both backup repositories are located onsite: for this reason, you will use the direct data transfer path. Leave the **Direct** option selected.

New Backup Copy Job

Data Transfer
Choose how VM data should be transferred from source to target backup repository.

Job
Virtual Machines
Target
Data Transfer
Schedule
Summary

☒ **Direct**
VM data will be sent directly from source to target repository. This mode is recommended for copying backups on-site, and off-site over a fast connection.

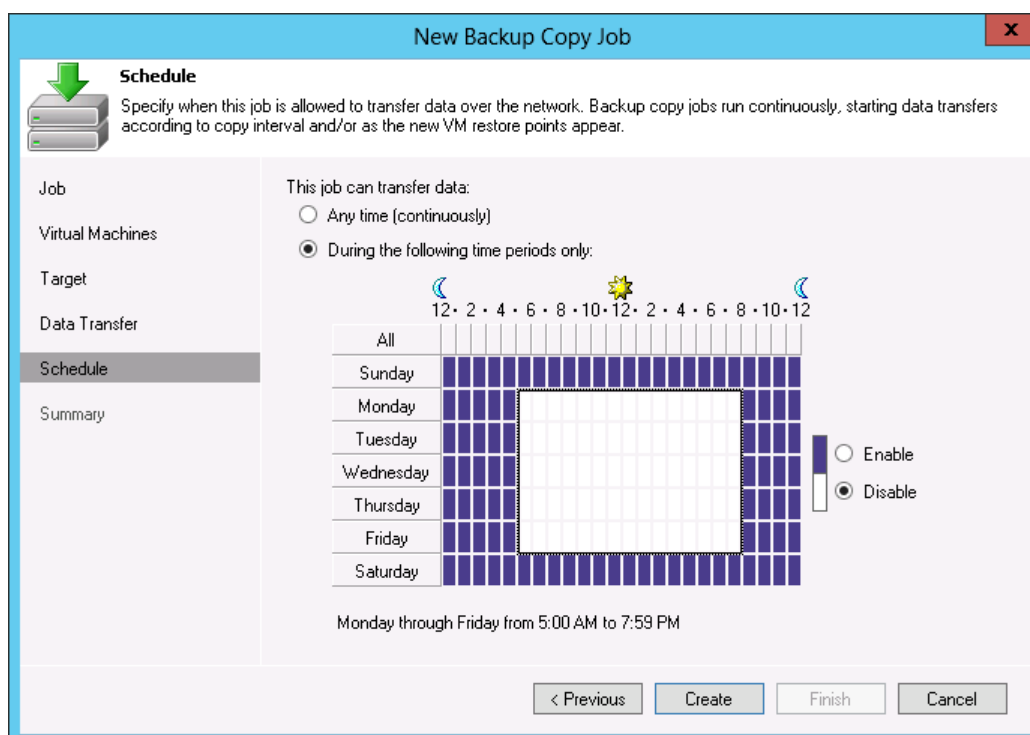
☐ **Through built-in WAN accelerators**
VM data will be sent to target repository through WAN accelerators that must be deployed in both source and target sites. This mode provides for significant bandwidth savings.

Source WAN accelerator:

Target WAN accelerator:

< Previous Next > Finish Cancel

- 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. The backup copy window can be helpful if you do not want the backup copy job to produce unwanted overhead for the production environment or do not want the job to overlap the production hours.



- Click **Create**.
- 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.
- To monitor the job performance, click the **Backup Copy** node under Jobs in the inventory pane. Double-click the created backup copy job in the working area: now you can monitor the job performance in the real-time mode.

Validation

- Open the **Files** view.
- Expand the file tree of the target backup repository.
- Find a subfolder with the backup copy job name and make sure a full backup file is available in it.
- Open the **Backup & Replication** view.
- Click the **Backup Copy** node under **Jobs** in the inventory pane. Double-click the backup copy job in the working area to view its statistics.

Performing Replication

Insight into Veeam Replication

With Veeam, you can not only back up, but also replicate your VMs. When you replicate a VM, Veeam Backup & Replication creates an exact copy of a production VM in the native VMware format on a spare ESX(i) host and maintains this copy in sync with the original VM.

Replication provides the best RTOs and RPOs as you actually have a copy of your VM in a ready-to-start state. If the primary VM goes down for some reason, you can immediately fail over to the VM replica and restore critical services with minimum downtime. Subsequently, replication is most commonly used for VMs running tier 1 applications.

Replication is a job-driven process. During the first run of a replication job, Veeam Backup & Replication copies the whole VM image and registers a replicated VM on the target ESX(i) host. During next cycles of a job, Veeam Backup & Replication copies only incremental changes and creates restore points for a VM replica so you can recover your VM to the necessary state. Every restore point is in fact a usual VMware snapshot. When you perform incremental replication, data blocks that have changed since the last replication cycle are written to the snapshot delta file next to a full VM replica. The number of restore points in the chain depends on your retention policy settings.

To provide extremely fast incremental replication, Veeam Backup & Replication uses the vSphere functionality, ESX Changed Block Tracking (or CBT). With CBT, you can replicate much faster and can schedule replication jobs as often as every few minutes. So you get near-CDP at only a fraction of the cost of traditional CDP solutions.

With Veeam Backup & Replication, you can perform both onsite replication for HA and offsite replication for DR scenarios. For replication over WAN or slow links, Veeam Backup & Replication provides a number of means to optimize data transmission: it performs inline deduplication and compresses replica traffic. You can also configure network throttling rules to prevent replication jobs from consuming the entire bandwidth of your environment and perform replica seeding.

Evaluation Case

In this exercise, you will create a replica of a VM on the target host, and create one restore point next to a full VM replica.

Please note that this guide describes replication to a local target host located in the same network. To learn about replicating offsite, see [Veeam Backup & Replication User Guide](#).

Prerequisites

- All backup infrastructure components that will take part in the replication process should be added to the Veeam Backup & Replication console. These include a source and target ESX(i) hosts, a backup proxy (used as a data mover) and a backup repository (used for storing auxiliary replica files). The latter two components are required for a distributed architecture scenario only.
- [Optional] To receive an email notification when a replication job completes, specify global email notification settings. To do that, select **Options** from the main menu of Veeam Backup & Replication and specify necessary settings on the **Email Settings** tab.

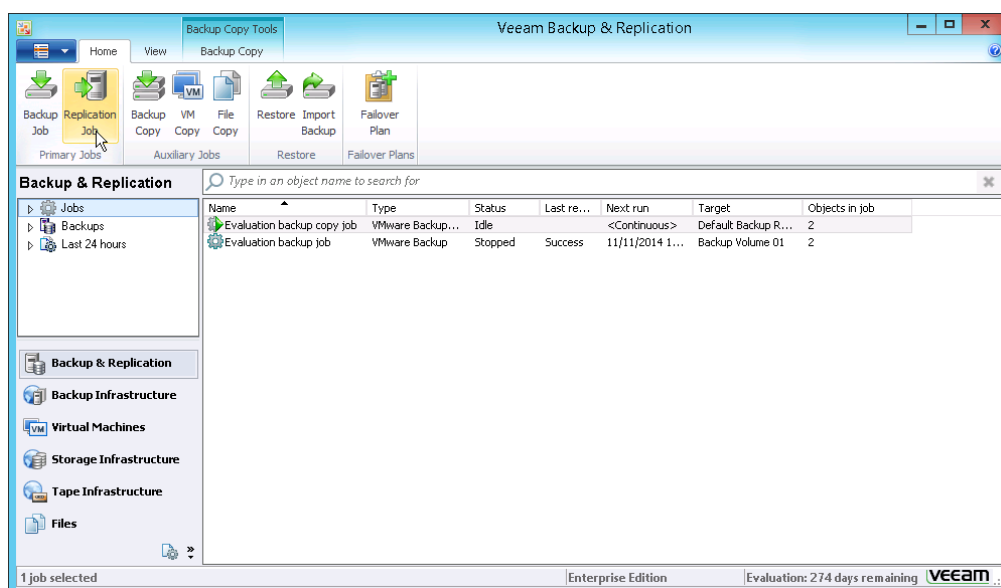
- [Optional] To evaluate the application-aware image processing feature, make sure that at least one of replicated VM runs the following OS'es:
 - Microsoft Windows Server 2003
 - Microsoft Windows Server 2008
 - Microsoft Windows Server 2008 R2
 - Microsoft Windows 7
 - Microsoft Windows 8
 - Microsoft Windows 8.1
 - Microsoft Windows 2012 Server
 - Microsoft Windows 2012 R2 Server

Procedure

To replicate a VM:

Step 1. Create a replication job

1. On the **Home** tab, click the **Replication Job** button.



2. Specify a name for the created replication job.

New Replication Job

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

Name:
Evaluation replication job

Description:
Pilot replication

Describe your DR site:

- ☐ Low connection bandwidth (enable replica seeding)
- ☐ Separate virtual networks (enable network remapping)
- ☐ Different IP addressing scheme (enable re-IP)

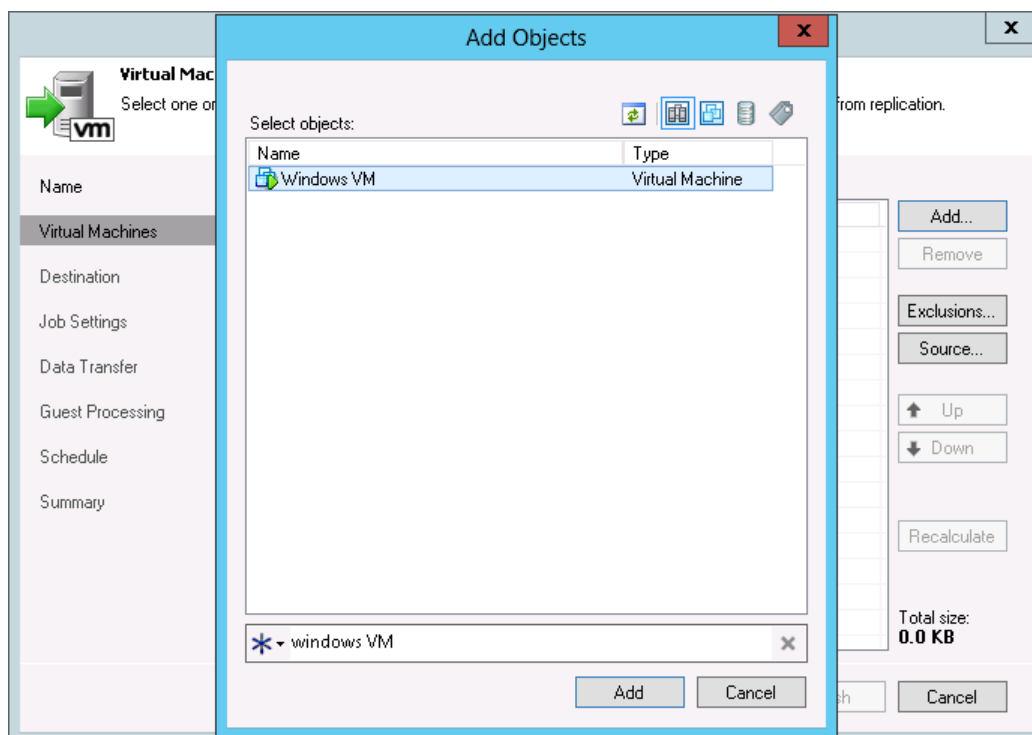
< Previous Next > Finish Cancel

Step 2. Add VMs to the replication job

You can replicate individual VMs or VM containers: folders, resource pools, clusters, vApps, datastores and so on. Jobs with VM containers are dynamic in nature: if a new VM is added to the container after the replication job is created, the job is automatically updated to include the new VM.

If you have connected vCenter Server rather than a standalone ESX(i) host to the Veeam backup server, VMs added to the job will be replicated even if they are vMotioned to another host.

1. At the **Virtual Machines** step of the wizard, click **Add**.
2. To quickly find a VM or VM container, enter the name of the object that you want to find in the search field and click the **Start search** button on the right. Select a VM or VM container in the displayed list and click **Add**.



The object will appear in the **Virtual machines to replicate** list.

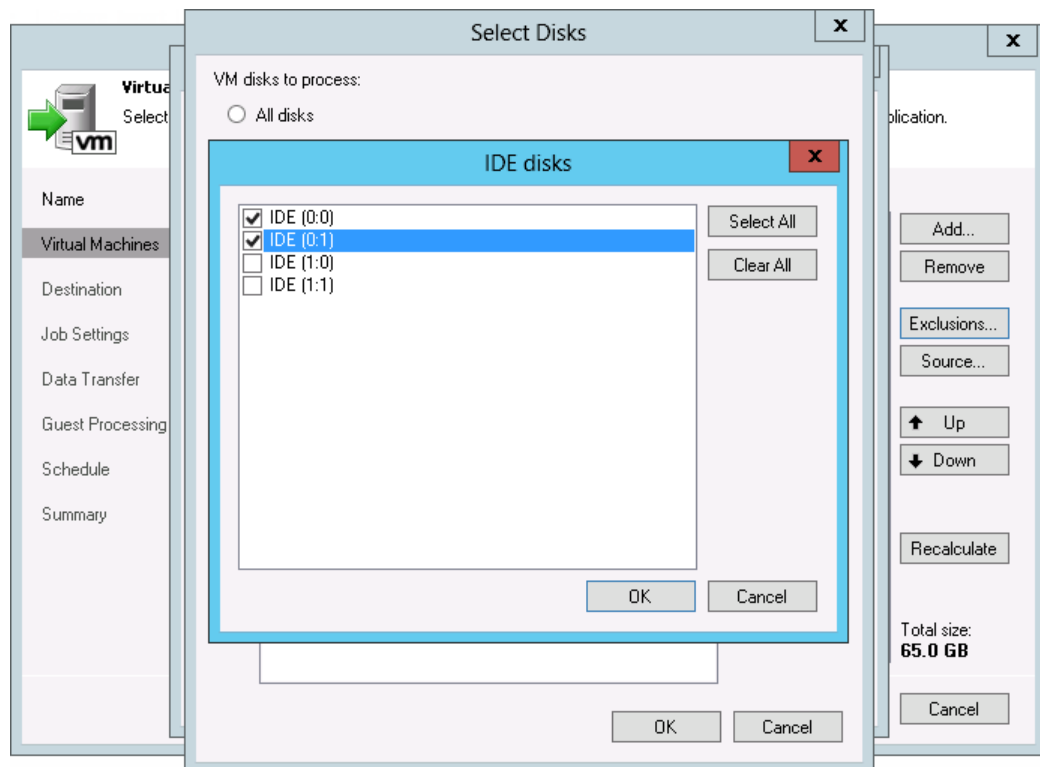
Step 3. Exclude VMs and VM disks

If you create a replication job for a VM container, you can exclude specific VMs or VM containers from the replication job. You can also select which VM disks to replicate.

Veeam Backup & Replication automatically excludes VM log files from replicas to make the replication process faster and reduce the size of the replica.

1. At the **Virtual Machines** step of the wizard, click **Exclusions**.
2. On the **VMs** tab, click **Add**. Select VMs that should be excluded. To quickly find a VM or VM container, enter the name of the object you want to find in the search field and click the **Start search** button on the right.
3. On the **Disks** tab, select a VM and click **Edit** to select disks that should be replicated. This functionality is useful, for example, if you only want to replicate VM system drives.

If you want to exclude disks of a VM added as part of a container, click **Add** on the right to include the VM in the list as a standalone instance.

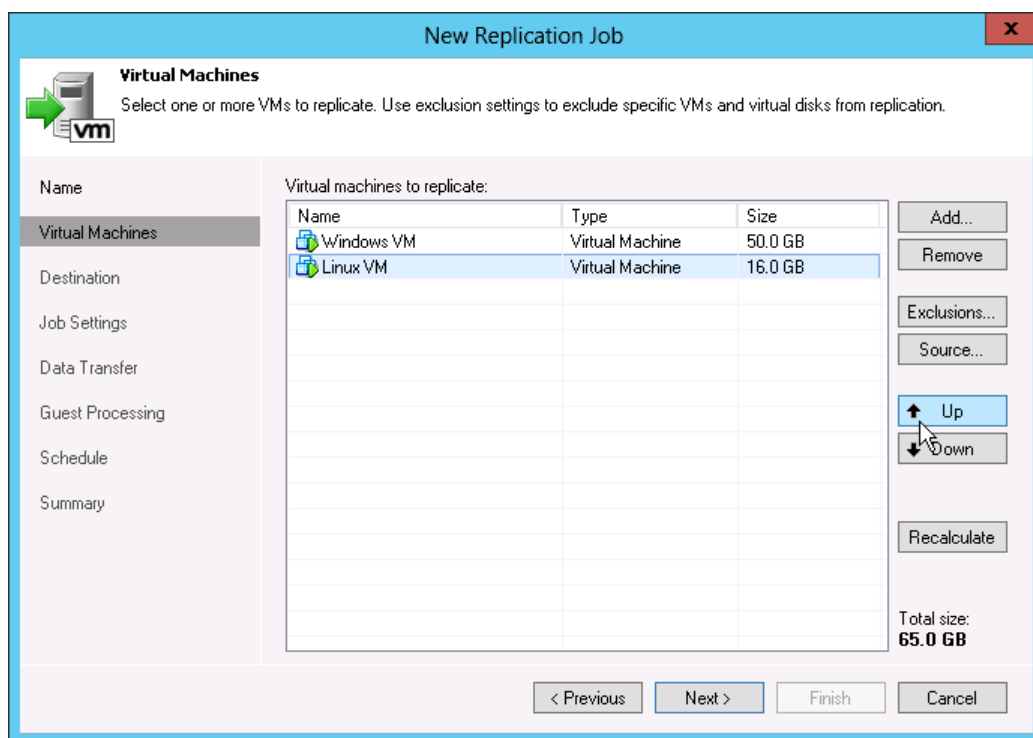


4. Click **OK**.
5. Click **Recalculate** to see the total size of replicated objects.

Step 4. Define the VM replication order

If you have included a number of VMs or VM containers in the replication job, you can specify the order in which VMs should be processed. This will help you make sure that the most important VMs in the job are processed first, for example, if you must fit into the backup window and you are not sure how much time VM processing will take.

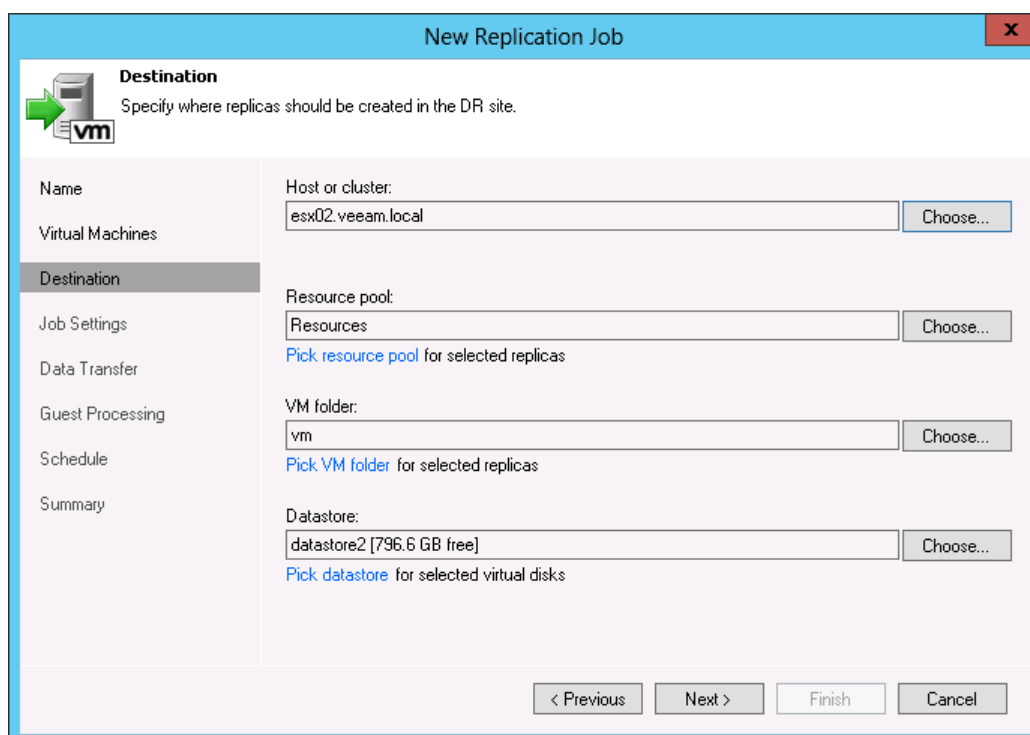
1. At the **Virtual Machines** step of the wizard, click the added VM in the list.
2. Use the **Up** and **Down** buttons on the right to move the VM higher or lower in the list. The higher is the VM in the list, the higher its priority. If you added a VM container as a single instance, VMs inside the container will be processed at random.



Step 5. Select replica destination

At the **Destination** step of the wizard, specify the target for a replicated VM.

1. Click **Choose** next to the **Host** or **Cluster** field and select a host on which the VM replica should be registered.
2. Click **Choose** next to the **Resource pool** field and select the destination resource pool.
3. Click **Choose** next to the VM folder field and select a folder to which a replicated VM should be placed.
4. Click **Choose** next to the **Datastore** field and select a datastore where VM replica files should be stored.



The screenshot shows the 'New Replication Job' wizard window, specifically the 'Destination' step. The window has a blue title bar and a sidebar on the left with icons for 'Virtual Machines', 'Destination' (selected), 'Job Settings', 'Data Transfer', 'Guest Processing', 'Schedule', and 'Summary'. The main area is titled 'Destination' with a subtitle 'Specify where replicas should be created in the DR site.' It contains four fields: 'Host or cluster:' with the value 'esx02.veeam.local', 'Resource pool:' with the value 'Resources', 'VM folder:' with the value 'vm', and 'Datastore:' with the value 'datastore2 [796.6 GB free]'. Each field has a 'Choose...' button to its right. Below the fields are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 6. Specify general settings for the replication job

A replication job must be pointed to a backup repository. The backup repository stores replica metadata (checksums of read data blocks) required to streamline incremental passes of the job.

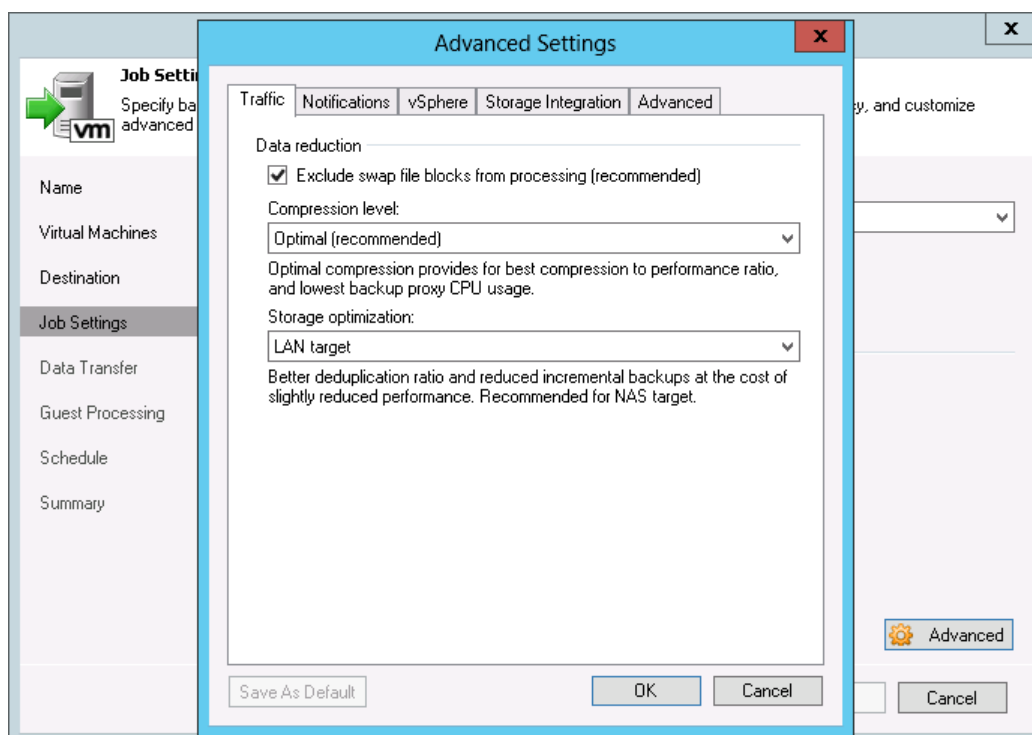
1. At the **Job Settings** step of the wizard, from the **Repository for replica metadata** list, select the backup repository that you have configured. A metadata file created by the job will be stored to this backup repository.
2. In the **Replica's name suffix** field, enter a suffix to append to the name of the replica. Veeam Backup & Replication will use the name of the primary VM with the suffix appended to register a VM replica on the target host.
3. Select the maximum number of restore points that you want to keep on disk. By default, Veeam Backup & Replication keeps 14 restore points.

Step 7. Specify advanced replica settings

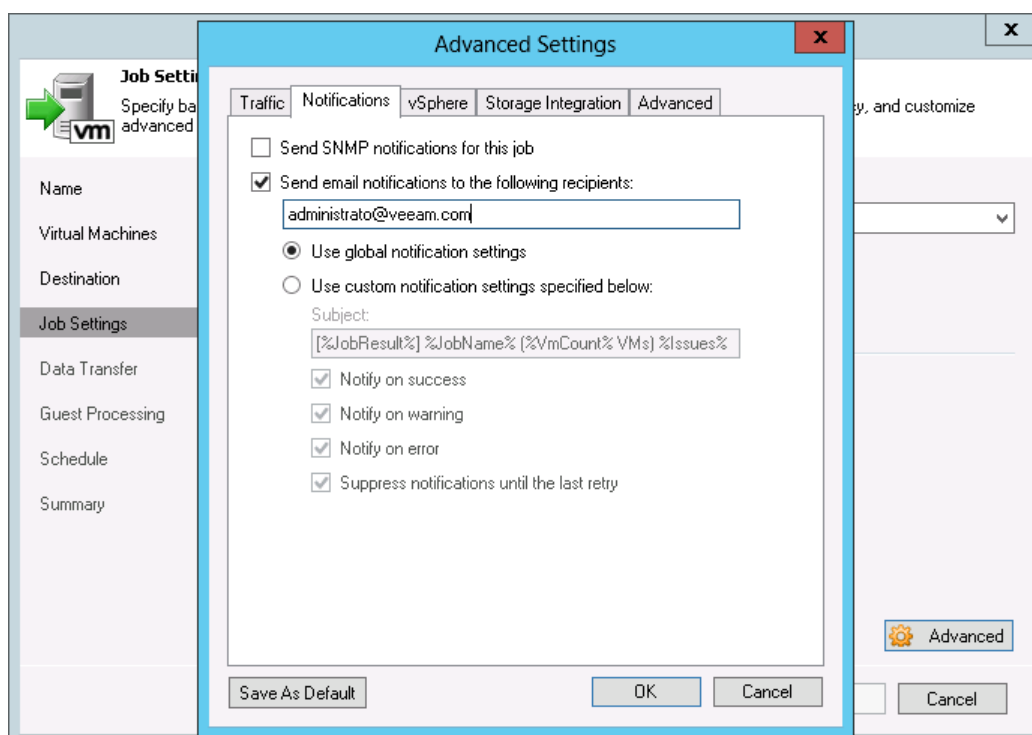
At the **Job Settings** step of the wizard, click **Advanced** to specify additional settings for the replication job.

1. Make sure the **Exclude swap file blocks from processing** check box is selected. Swap files are dynamic and change intensively between runs of a replication job. With this option selected, Veeam Backup & Replication will identify data blocks of the Microsoft Windows pagefile in the guest OS and exclude them from processing, which will result in increased performance and a smaller size of increments.
2. When VM data is transferred between two backup proxies, Veeam Backup & Replication compresses VM data to reduce load on the network. Veeam Backup & Replication offers 5 compression levels that provide different compression ratios to meet the needs of your environment: *None*, *Dedupe-friendly*, *Optimal*, *High* and *Extreme*.

In this exercise, one backup proxy is used as a source and target proxy. For this reason, no compression will be applied.

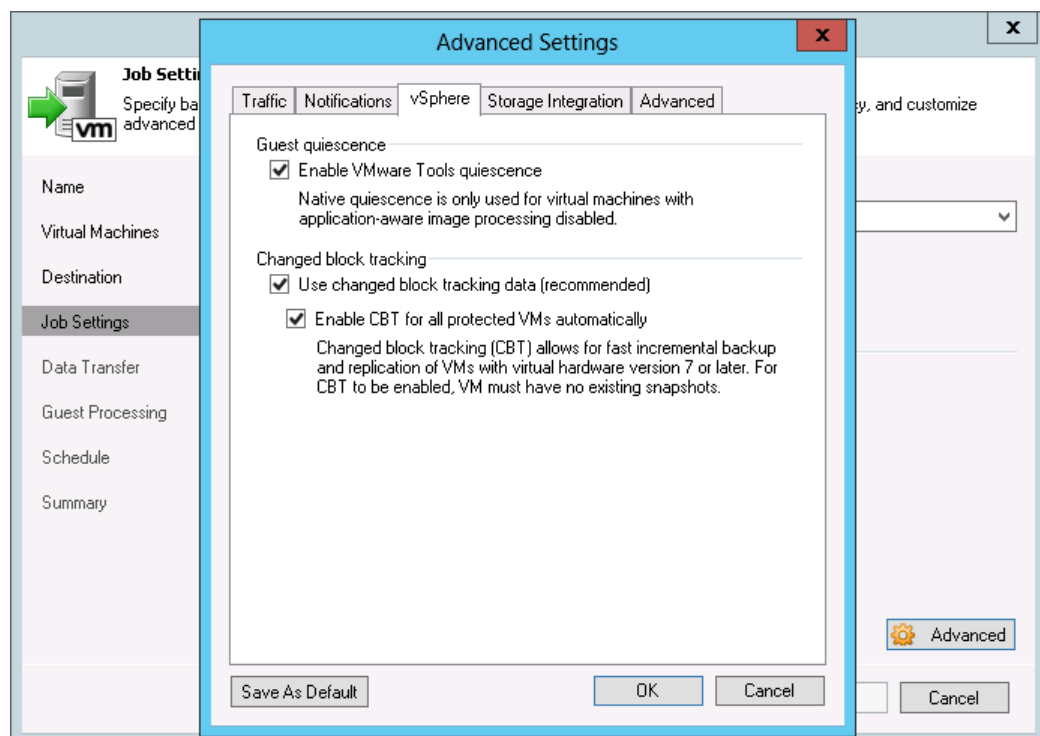


3. Click the **Notifications** tab. Select the **Send email notifications to the following recipients** check box and specify the email address. When the job completes, you will receive an email notification with details on job performance. Note that you will be able to receive an email notification only if you specify global email settings. To specify email settings, select **Options** from the main menu.



Note: Make sure that you specify your email address once: either in the **To** field in general notification settings, or in job notification settings. If you specify both, you will receive two identical notifications when the job is completed.

4. Click the **vSphere** tab.
5. If you replicate a non VSS-aware VM, for example, Linux-based VMs, make sure that the **Enable VMware tools quiescence** check box is selected. This option helps create transactionally consistent replicas of such VMs.
6. Make sure the **Use changed block tracking data** check box is selected. For VMware VMs with hardware version 7 or later, Veeam Backup & Replication employs VMware vSphere Changed Block Tracking (or CBT). Instead of scanning VMFS to know which data blocks have changed since the previous job run, Veeam Backup & Replication queries the CBT module to get the list of changed blocks. Use of CBT increases the speed and efficiency of block-level incremental replication. For example, if only 5% of a VM changed since the last replication, incremental replication will be performed 20 times faster.
7. Make sure the **Enable changed block tracking for all processed VMs** check box is selected, too. This option forces use of CBT even in case it is switched off at the level of the ESX host.

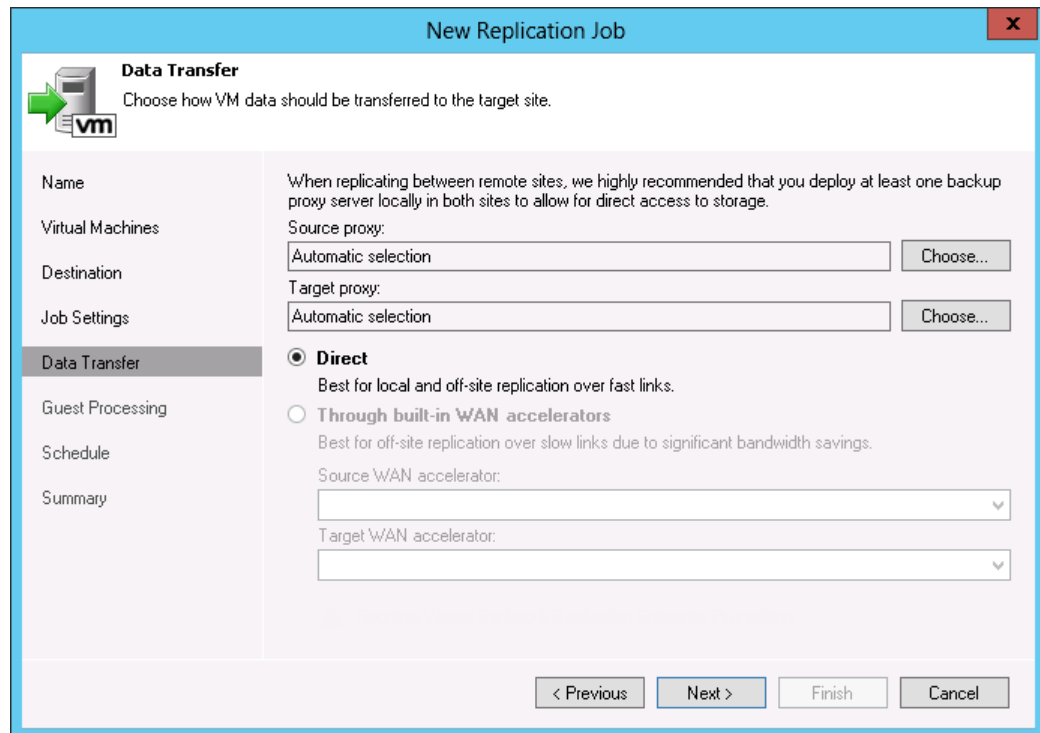


Step 8. Specify replication data path

You must select backup infrastructure components that must be used for the replication process and choose a path for VM data transfer.

1. In the **Source proxy** and **Target proxy** fields, select the backup proxy that must be used for VM replication. You can assign backup proxies explicitly or leave the **Automatic** option selected. In the latter case, Veeam Backup & Replication checks settings of available backup proxies and select the most appropriate one for the job — the backup proxy that will enable the most efficient data retrieval from the source datastore. Veeam Backup & Replication first attempts to choose a backup proxy that uses the *Direct SAN Access* mode, then the backup proxy that uses the *Virtual Appliance* mode. If such proxies are not available, Veeam Backup & Replication selects the least loaded backup proxy that uses the *Network* mode.
2. During the replication process, VM data can be transferred directly via backup proxy(ies) to the target datastore or via a pair of WAN accelerators. The latter scenario is recommended for replicating VM offsite or over slow network connections. To learn more about WAN acceleration, see [Veeam Backup & Replication User Guide](#).

In this exercise, source and target hosts are located onsite. For this reason, leave the **Direct** option selected.



The screenshot shows the 'New Replication Job' wizard in Veeam Backup & Replication, specifically the 'Data Transfer' step. The left sidebar contains a list of steps: Name, Virtual Machines, Destination, Job Settings, Data Transfer (highlighted), Guest Processing, Schedule, and Summary. The main area is titled 'Data Transfer' with a green arrow icon and a 'vm' icon. Below the title, it says 'Choose how VM data should be transferred to the target site.' The main content area has a light blue background and contains the following text: 'When replicating between remote sites, we highly recommended that you deploy at least one backup proxy server locally in both sites to allow for direct access to storage.' Below this text are two sections for proxy selection. The first section is for 'Source proxy:' with a text box containing 'Automatic selection' and a 'Choose...' button. The second section is for 'Target proxy:' with a text box containing 'Automatic selection' and a 'Choose...' button. Below these are two radio button options. The first option is 'Direct', which is selected (indicated by a filled radio button). Below it is the text 'Best for local and off-site replication over fast links.' The second option is 'Through built-in WAN accelerators', which is not selected (indicated by an empty radio button). Below it is the text 'Best for off-site replication over slow links due to significant bandwidth savings.' Below the second option are two text boxes for 'Source WAN accelerator:' and 'Target WAN accelerator:', both with dropdown arrows. At the bottom of the wizard are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

New Replication Job

Data Transfer
Choose how VM data should be transferred to the target site.

When replicating between remote sites, we highly recommended that you deploy at least one backup proxy server locally in both sites to allow for direct access to storage.

Source proxy: Automatic selection Choose...

Target proxy: Automatic selection Choose...

☒ **Direct**
Best for local and off-site replication over fast links.

☐ **Through built-in WAN accelerators**
Best for off-site replication over slow links due to significant bandwidth savings.

Source WAN accelerator: [dropdown]

Target WAN accelerator: [dropdown]

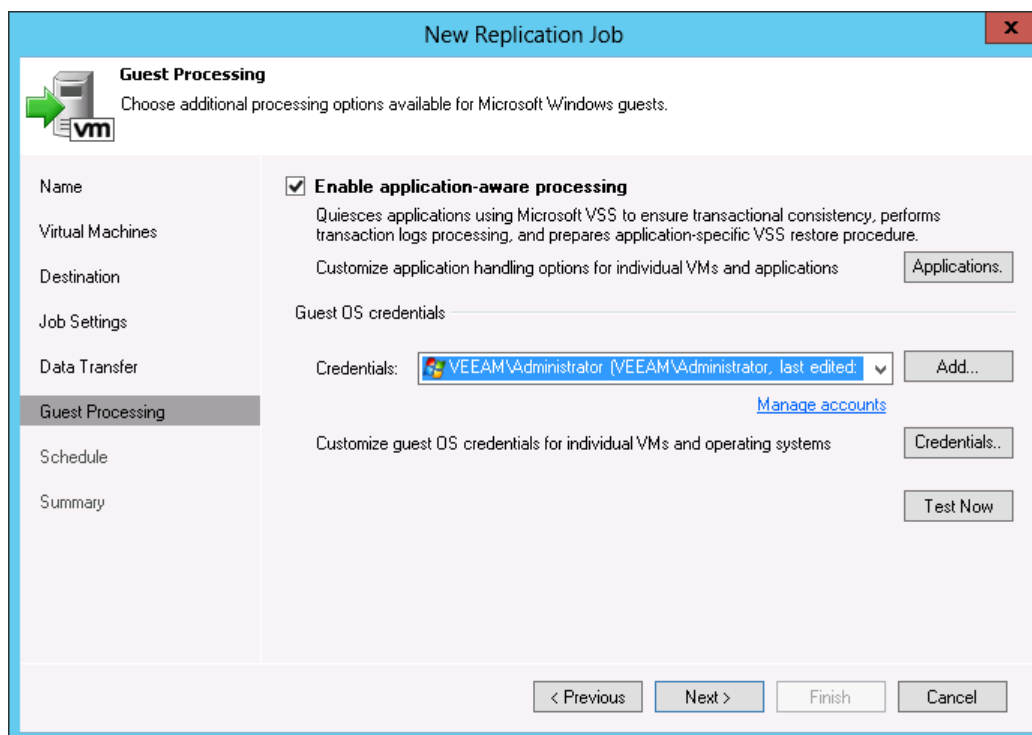
< Previous Next > Finish Cancel

Step 9. Specify additional guest OS processing options

To replicate VMs running VSS-aware applications, Veeam Backup & Replication uses application-aware image processing based on Microsoft VSS. Jobs with application-aware image processing produce transactionally consistent replicas, that, unlike crash consistent replicas, ensure proper recovery of virtualized applications without any data loss.

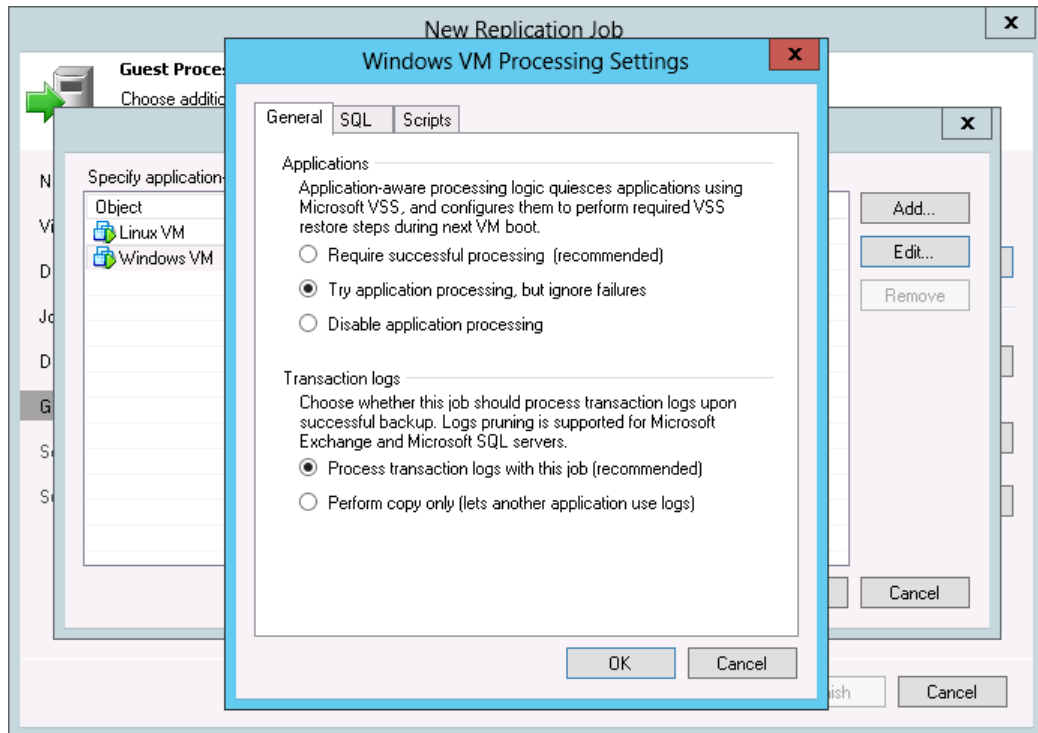
To enable application-aware image processing:

1. At the **Guest Processing** step of the wizard, select the **Enable application-aware image processing** check box.
2. Specify guest OS credentials (username and password) with Local Administrator privileges on all VMs included in the job. OS credentials are required to install, start and remove a runtime process that coordinates VSS activities inside the VM.



3. By default, the specified guest OS credentials are used for all VMs processed by the replication job. If you replicate several VMs that use different guest OS credentials, click **Advanced**. Select a VM in the list and click **Set User**. Then enter guest OS credentials with Local Administrator privileges for this specific VM. Repeat the procedure for all VMs in the job.
4. To specify advanced options for VSS processing, click **Applications**. Select a VM in the list and click **Edit**.
5. On the **General** tab, select **Try application processing, but ignore failures** to continue the replication job even if VSS errors occur. If VSS processing fails during the job, Veeam Backup & Replication will use VMware Tools quiescence to create a transactionally consistent VM replica.

6. Make sure that the **Process transaction logs with this job** option is selected to correctly handle transaction logs after the replication job is complete. In this case, if the replication job finishes successfully, Veeam Backup & Replication will truncate transaction logs so that they do not overflow storage space. If you use a third-party backup tool that maintains the consistency of transaction logs, select the **Perform copy only** option to prevent possible conflicts.



Step 10. Specify job scheduling settings

A replication job can be scheduled or run manually. To schedule a replication job:

1. At the **Schedule** step of the wizard, select the **Run the job automatically** check box. If you do not select this check box, the job will be saved and you will have to run it manually.
2. Select the schedule type: daily, monthly, periodically or continuously. You can also chain the jobs so that they run one after another.
3. Select the **Retry failed VM processing** check box. During the retry cycle, only VMs that have failed during the main replication cycle will be processed.
4. Select the **Terminate job if it exceeds allowed backup window** check box and click **Window**. Define the backup window for your environment. In case the created job overlaps the specified window, it will be automatically terminated not to produce additional overhead on your virtual environment.
5. Click **Create**.

New Replication Job

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

Name

Virtual Machines

Destination

Job Settings

Data Transfer

Guest Processing

Schedule

Summary

☒ Run the job automatically

☒ Daily at this time: 10:00 PM Everyday Days...

☐ Monthly at this time: 10:00 PM Fourth Saturday Months...

☐ Periodically every: 1 Hours Schedule...

☐ After this job: Evaluation backup job (Backup job created for evaluation purpose)

Automatic retry

☒ Retry failed VMs processing: 3 times

Wait before each retry attempt for: 10 minutes

Backup window

☐ Terminate job if it exceeds allowed backup window Window...

If the job does not complete within allocated backup window, it will be terminated to prevent snapshot commit during production hours.

< Previous Create Finish Cancel

Step 11. Review job settings and start the job

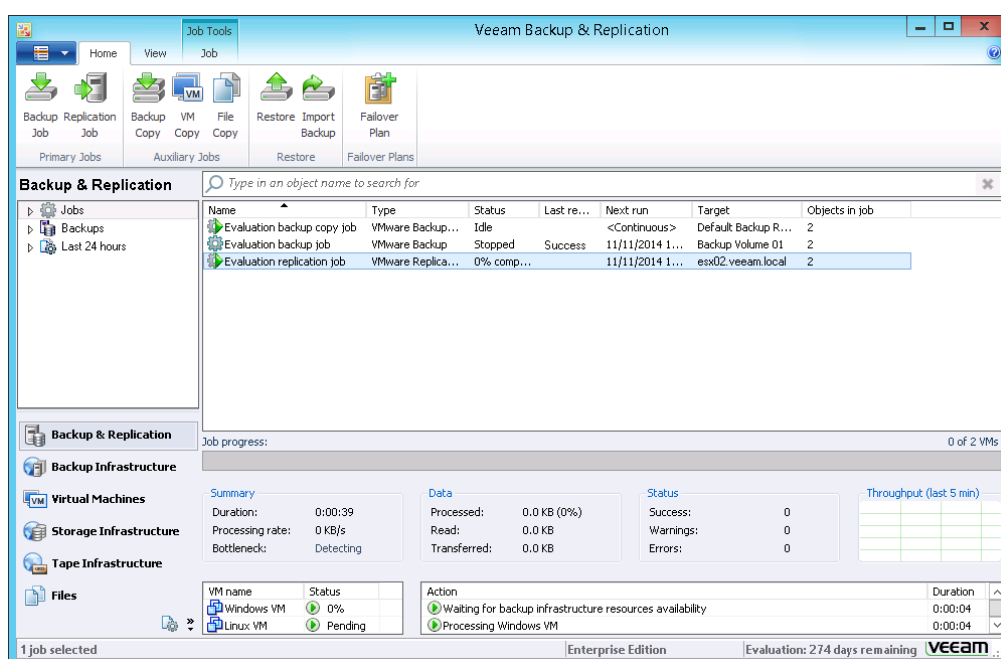
1. Review the summary of replication job settings.
2. Select the **Run the job when I click Finish** check box and click **Finish**. The job will start.

Step 12. Monitor job performance in real time

When a replication job is running, you can view job statistics in the real-time mode. Job statistics provide detailed data on the job: job progress, duration, processing rate, performance bottlenecks, the amount of data processed, read and transferred, and other details of the job performance.

Beside general job statistics, you can view detailed data for each VM or VM container processed by the job.

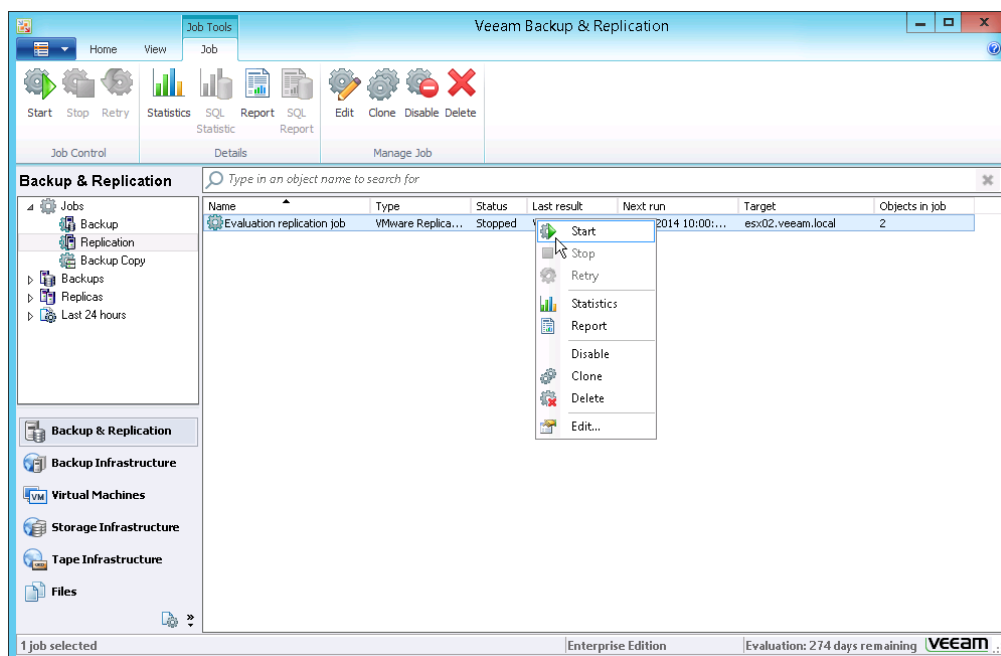
1. Open the **Backup & Replication** view.
2. Select the **Replication** node under **Jobs** in the inventory pane.
3. Click the job in the list. Now you can track the entire job performance as it runs.
4. Select the name of a specific VM or a VM container to view detailed statistics for this specific object only.
5. Wait for the job to complete. Note that the job must complete with the *Success* or at least the *Warning* status. If the job completes with the *Failed* status, the VM replica will not be created, and you will not be able to perform failover and fallback operations.



Step 13. Perform incremental replication

To perform incremental replication of a VM, do the following:

1. Open the **Backup & Replication** view.
2. Select the **Jobs** node in the inventory pane.
3. Right-click the job in the list and select **Start**. Wait for the job to complete. Note that the job must complete with the *Success* or at least the *Warning* status.



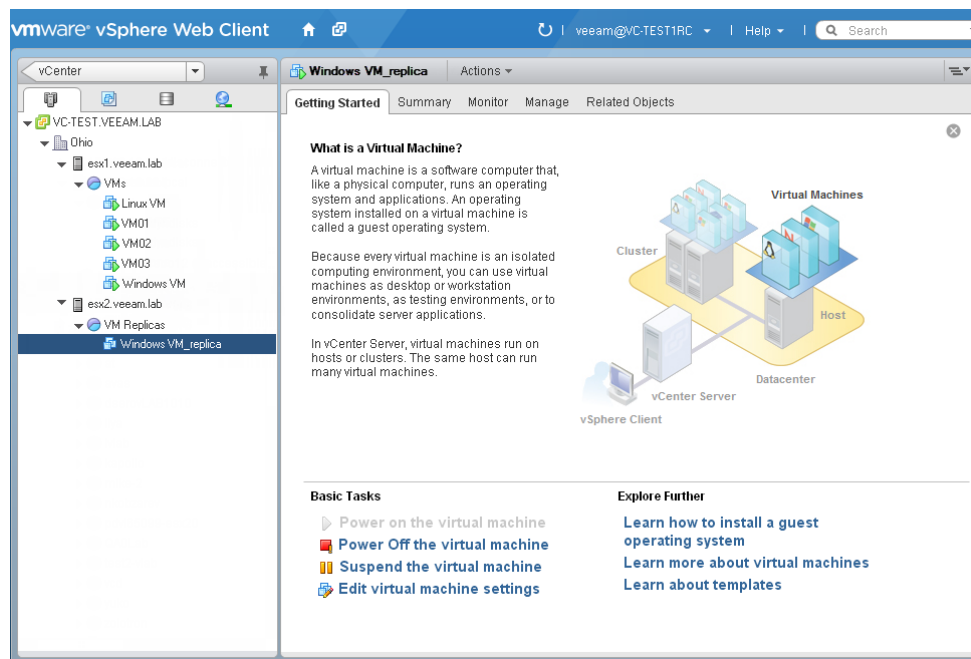
Validation

During replication cycles, Veeam Backup & Replication creates the following files for a VM replica:

- A full VM replica (a set of VM configuration files and virtual disks). During the first replication cycle, Veeam Backup & Replication puts these files to the selected datastore to the *ReplicaName* folder and registers a VM replica on the target host.
- Replica restore points (snapshot delta files). During incremental replication, Veeam Backup & Replication creates a snapshot delta file in the same folder, next to a full VM replica.
- Replica metadata (VBK) used to store replica checksums. Veeam Backup & Replication uses this file to quickly detect changed blocks of data between two replica states. A metadata file is written to the backup repository.

To check replication results:

1. Open the vSphere Client and make sure that a VM replica has been created.



2. In **Veeam Backup & Replication**, open the **Files** view.
3. In the inventory pane, expand the datastore to which the replica has been stored. Browse to the *ReplicaName* folder and make sure that files of the VM replica are available there.
4. Open the **History** view and select the **Jobs** node in the inventory pane. Double-click the replication job in the list. Check the properties of the created replica.
5. If you have configured to receive an email message once the job completes, open your email client and check the *Inbox* folder. Make sure that you have two incoming emails with job results – one for the full replication, and another one for the incremental replication.

Failing Over a VM Replica

Insight into Failover

If a primary 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 over the role of the original VM: you switch from the production VM to its replica and shift your I/O and processes from the production host to a secondary host. As a result, you have your VM up and running within a couple of minutes, and your users can access services and applications they need with minimal disruption.

In Veeam Backup & Replication, you can fail over to the latest state of a replica or to any of its valid restore points.

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

- **Perform permanent failover.** When you perform permanent failover, you “commit failover”. That is, you permanently switch from the original VM to a VM replica and use this replica as the original VM. This scenario is acceptable if your original VM and a VM replica are located in the same site and are nearly equal in terms of resources. In this case, your users will not experience any latency in ongoing operations.
- **Perform failback.** When you perform failback, you “return” from a VM replica to the original VM after the problem in the production site is eliminated. All changes that took place while the VM replica was running are transferred to the original VM. You can follow this scenario if your VM replica is located in a DR site and is running on a lower tier host and storage: that is, it is not intended for continuous operations.
- **Undo failover.** When you undo failover, you switch back to the original VM and work with it in the normal operation mode. All changes made to the VM replica are discarded. You can follow this scenario if you plan to perform some testing and troubleshooting of your VM replica and do not want to affect your production environment in any way.

Veeam Backup & Replication supports failover and failback operations for one VM and for a number of VMs. This way, if you have a problem with an ESX(i) host, you can restore its work with minimum downtime.

Evaluation Case

In this exercise, you will fail over from the original VM to the VM replica you created in the previous exercise, and undo failover to get back to the production VM.

When you fail over from a running original VM to a VM replica, Veeam Backup & Replication does not power off the original VM, it simply powers on the VM replica. If the original VM is running when you perform failover, you will see a notification warning.

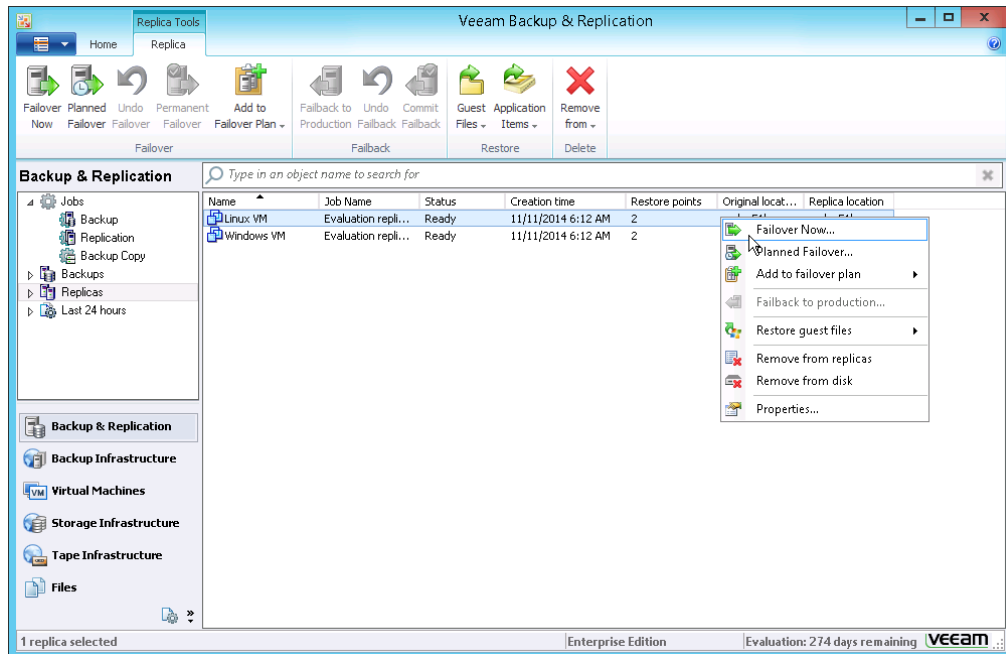
Prerequisites

The failover option can be used with any VM replica that was successfully created at least once. Open the **Backup & Replication** view, click the **Replicas** node in the inventory pane. Then expand the replication job and check if there is at least one restore point available for the replicated VM.


Procedure

To fail over to a VM replica:

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Right-click the replicated VM and select **Failover Now**.



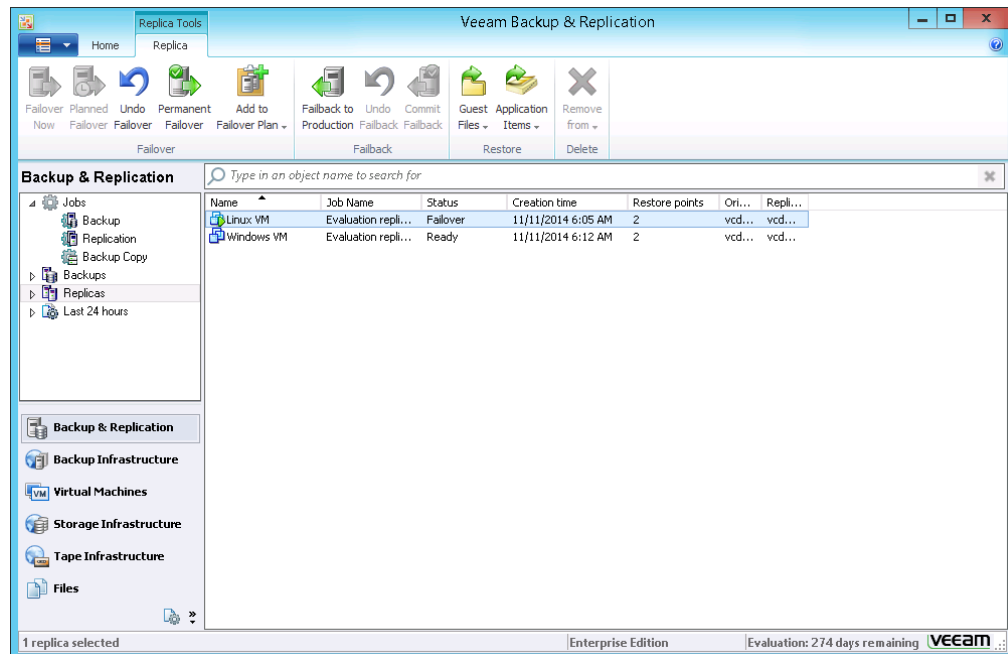
4. If you want to fail over to a specific restore point, select the VM in the list, click **Point** on the right and choose the restore point to which you want to fail over.



5. At the **Reason** step of the wizard, specify the reason for future reference and click **Next**.
6. Click **Finish** to fail over to a VM replica.

Validation

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Make sure that the replicated VM is put in the *Failover* state.



4. Open the vSphere Client and make sure that the VM replica is powered on.

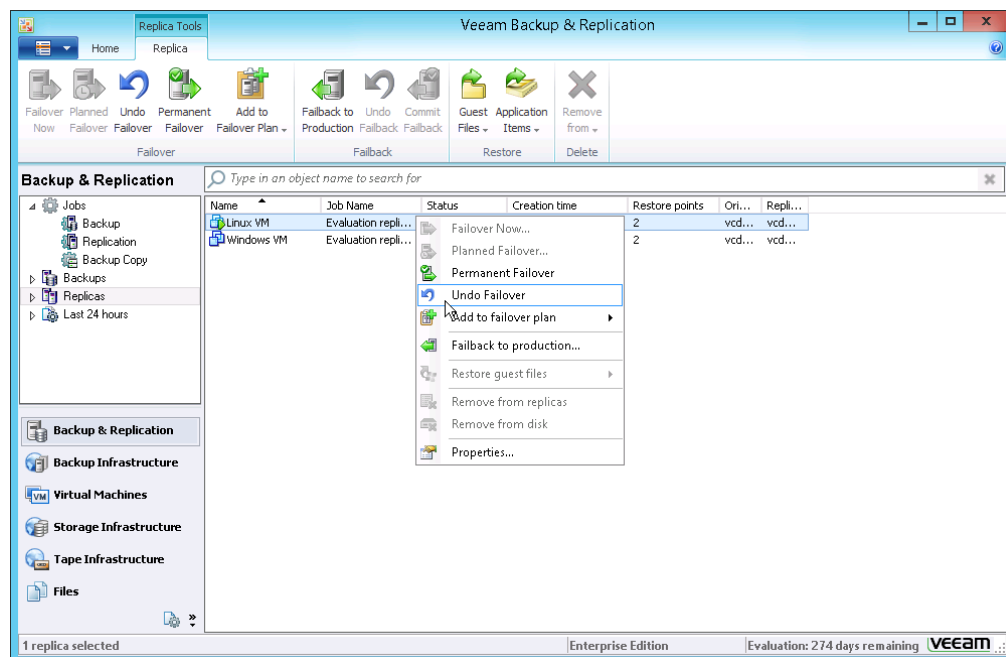
Undoing Failover

When you undo failover, Veeam Backup & Replication reverts the VM replica to its pre-failover state, powers it off, and deletes all changes that have taken place since the VM replica was powered on. You can use the production VM and perform its replication as usual.

Procedure

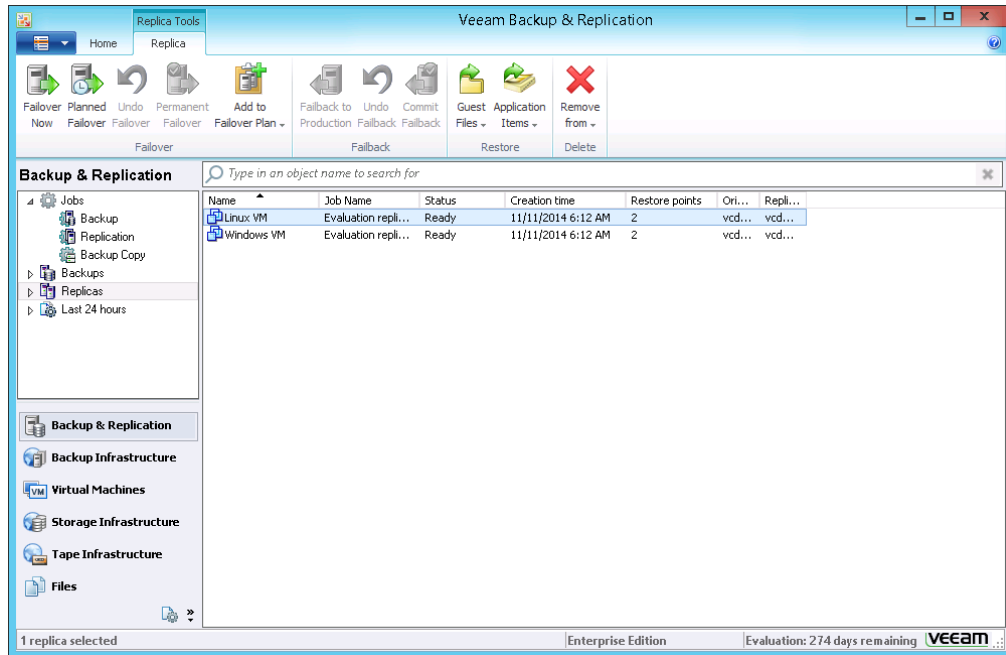
To undo failover:

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Right-click the VM and select **Undo failover**.



Validation

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Make sure that the replicated VM is put in the *Ready* state.



4. Open the vSphere Client and make sure that the VM replica is powered off and the production VM is powered on.

Failing Back to the Primary VM

Insight into Failback

If you want to resume operation of a production VM, you can fail back to it from a VM replica. When you perform failback, you get back from the VM replica to the original VM, shift your I/O and processes from the secondary host to the production host and return to the normal operation mode.

When you perform failback, Veeam Backup & Replication triggers a failback protective snapshot of a VM replica. This snapshot acts as a restore point and saves the pre-failback state of a replica to which you can return afterwards. Veeam Backup & Replication synchronizes the original VM with the VM replica to make sure you get back to the most recent state of a VM. The original VM is then powered on; all replication activities are put on hold.

You can fail back to the VM in the original location or in the new location:

- If you have managed to restore operation of the source host, you can fail back to the original VM on the source host.
- If the source host is not available, you can restore the original VM to a new location and fail back to it. You can also fail back to an entirely new location without restoring the original VM beforehand: in this case, VM replica files will be simply transferred to the necessary destination.

In Veeam Backup & Replication, failback itself is an intermediate action that needs to be finalized. Depending on a disaster recovery scenario, you can do one of the following:

- **Commit failback.** When you commit failback, you finalize recovery of the original VM in the production site. The original VM in the production site or at a new location becomes the primary VM, a VM replica is returned to the normal state and Veeam Backup & Replication resumes replication activities.
- **Undo failback.** If the production VM is not working as expected, you can undo failback and get back to a VM replica. In this case, the VM replica returns to the *Failover* state.

Evaluation Case

In this exercise, you will fail back from a VM replica to the primary VM on the source host, and then commit failback.

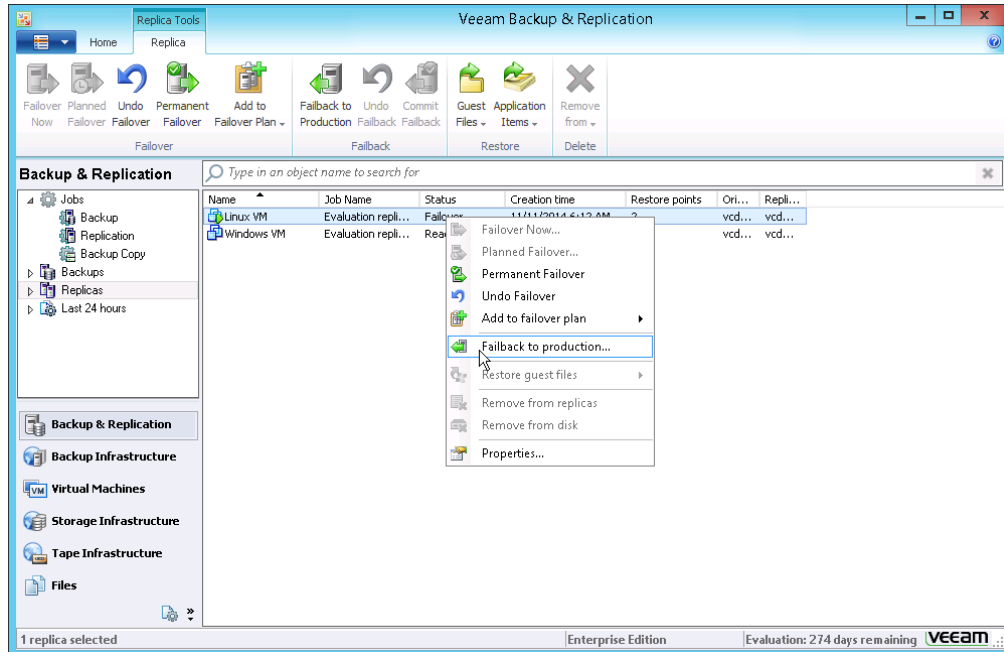
Prerequisites

You can perform failback for a VM replica in the *Failover* state. To put the VM replica to this state, make sure you have performed the [Failing Over VM Replica](#) exercise. If you have already undone failover, repeat the [Failing Over VM Replica](#) exercise for your VM.

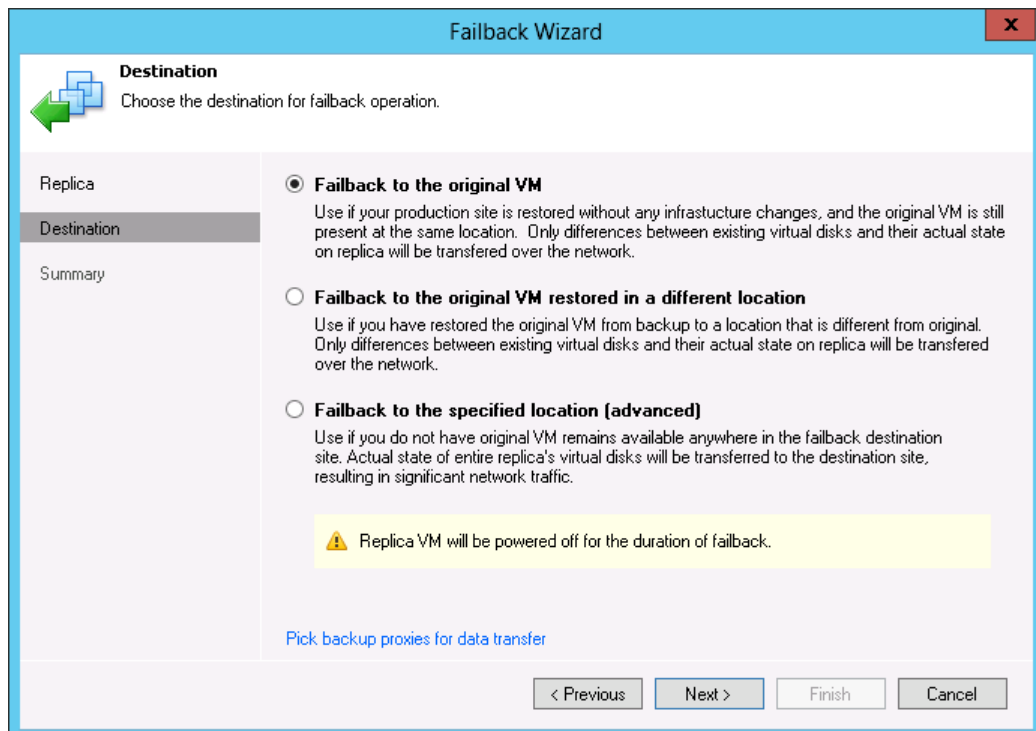
Procedure

To fail back to the production VM:

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Right-click the VM and select **Failback to production**.



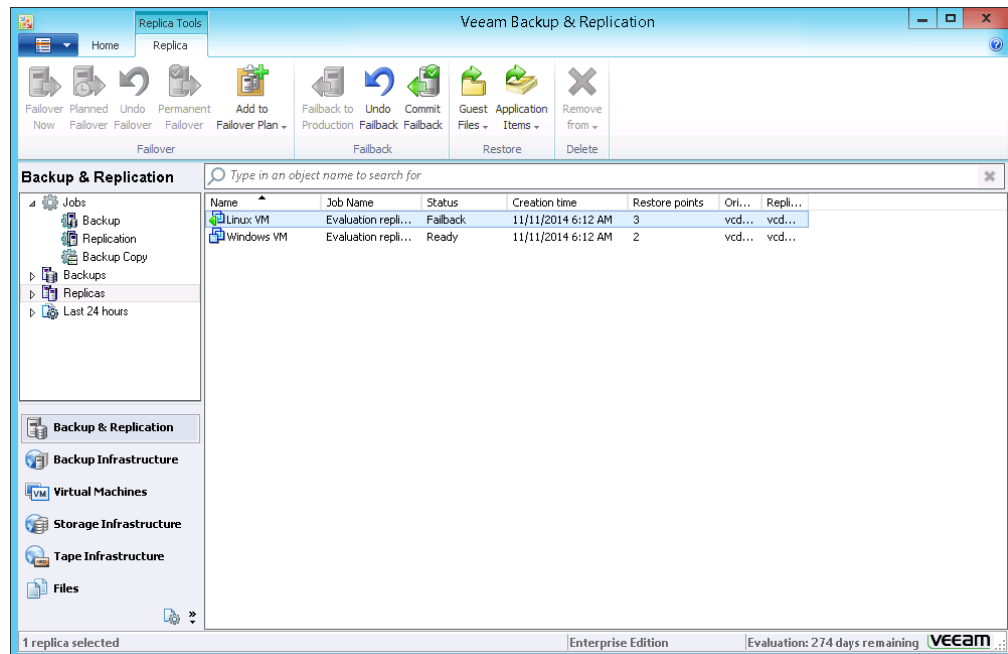
4. At the **Destination** step of the wizard, select **Failback to the original VM**.



5. At the summary step of the wizard, select the **Power on VM after restoring** check box and click **Finish**.

Validation

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Make sure that the replicated VM is put in the *Failback* state.



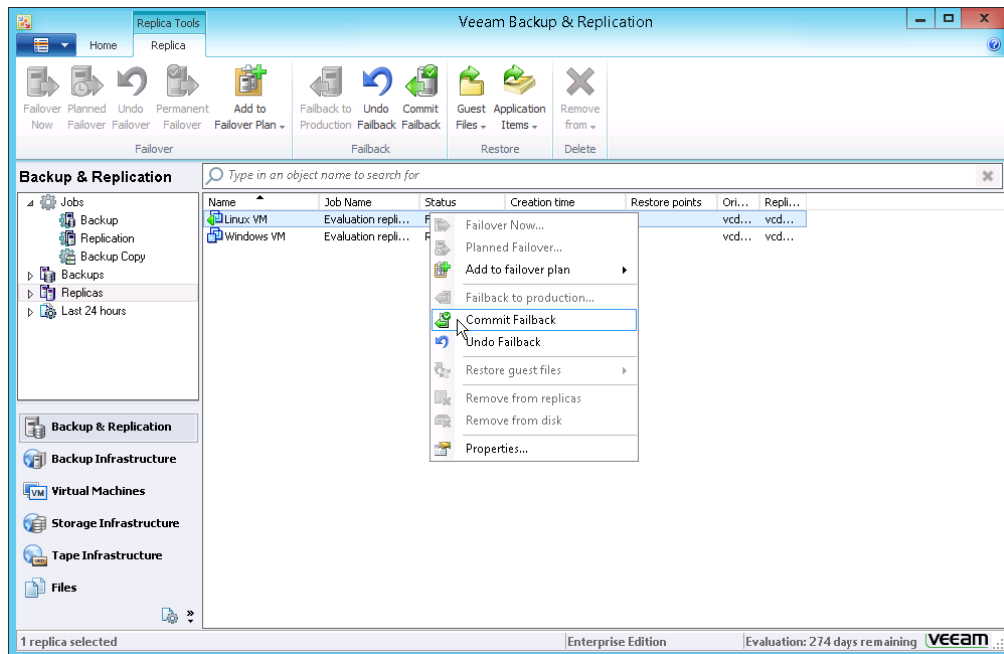
Committing Failback

When you commit failback, you confirm that you want to get back to the production VM. Veeam Backup & Replication resumes replication activities for the VM to which you failed back.

The failback protective snapshot that saves the pre-failback state of a VM replica is not deleted — Veeam Backup & Replication uses this snapshot as an additional restore point for VM replica. With the pre-failback snapshot, Veeam Backup & Replication needs to transfer less changes and therefore puts less load on the network when replication activities are resumed.

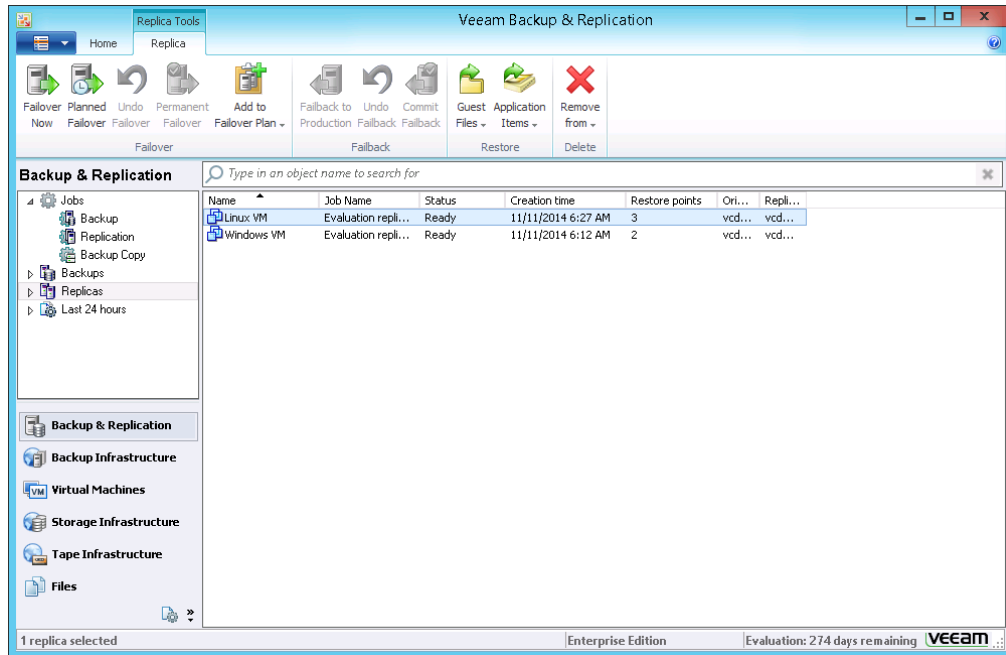
Procedure

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Right-click the VM and select **Commit Failback**.



Validation

1. Open the **Backup & Replication** view.
2. Select the **Replicas** node in the inventory pane.
3. Make sure that the replicated VM is put in the *Ready* state.



4. Open the vSphere Client and make sure the VM replica is powered off and the production VM is powered on.

DISTRIBUTED BACKUP INFRASTRUCTURE MANAGEMENT

For large geographically dispersed virtual environments with multiple Veeam backup servers, it is recommended to use Veeam Backup Enterprise Manager.

Veeam Backup Enterprise Manager is an optional component in the backup infrastructure that federates Veeam backup servers and offers a consolidated view of these servers through a web console. You can centrally control and manage jobs configured on different Veeam backup servers through a single pane of glass, edit and clone jobs, monitor jobs state and get reporting data across all backup servers. Veeam Backup Enterprise Manager also enables you to search for Windows guest OS files in all current and archived backups across your backup infrastructure and restore these files in 1 click.

This section describes a set of exercises that you should perform to learn how to manage a distributed backup infrastructure using Veeam Backup Enterprise Manager.

Exercise List

To evaluate the key features of Veeam Backup Enterprise Manager, perform the following exercises:

Exercise	Description	Time Estimates
Installing and configuring Veeam Backup Enterprise Manager	Install Veeam Backup & Replication on a physical or virtual machine. Add the Veeam backup server to Veeam Backup Enterprise Manager and collect data from the added Veeam backup server.	5-7 minutes
Searching for guest OS files and restoring them in 1 click	Find a guest OS file in a backup of a Windows-based VM and restore this file directly from Veeam Backup Enterprise Manager to the original VM in 1 click.	1-2 minutes
Configuring self-service restore	Delegate the file-level restore task to another team member.	2-5 minutes.
Cloning and editing jobs	Clone a created backup job, save it under a new name and edit its settings directly in Veeam Backup Enterprise Manager.	3-5 minutes
Restoring data from encrypted backups without a password	Create an encrypted backup file and unlock it without a password.	5-7 minutes

Installing and Configuring Veeam Backup Enterprise Manager

Evaluation Case

In this exercise, you will install Veeam Backup Enterprise Manager, connect the Veeam backup server to it and collect data from the connected Veeam backup server. In the production environment, Veeam Backup Enterprise Manager typically federates a number of Veeam backup servers. For evaluation purposes, however, you can connect one Veeam backup server to Veeam Backup Enterprise Manager.

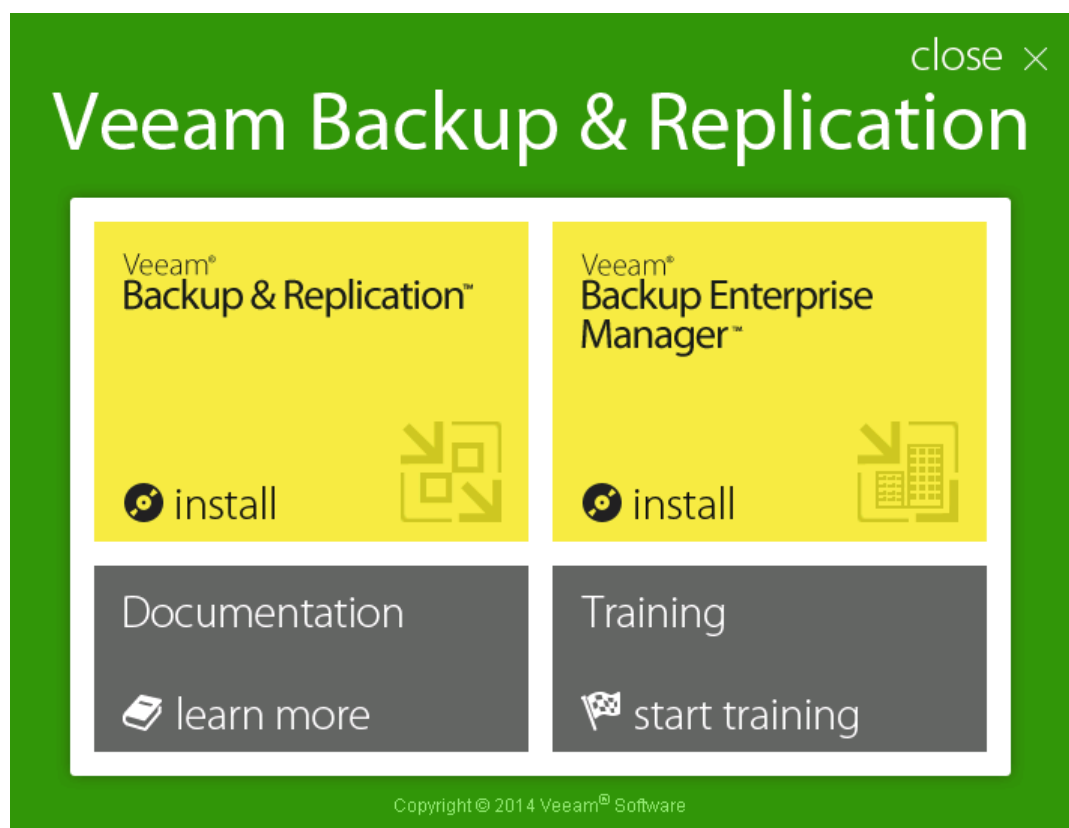
Prerequisites

- The machine on which you plan to install Veeam Backup Enterprise Manager must meet the system requirements.
- You should have at least one instance of the Veeam backup server installed. On the Veeam backup server, you should have at least one backup or replication job that has been successfully performed.
- Make sure that all necessary ports are open.

Procedure

To install Veeam Backup Enterprise Manager, do the following:

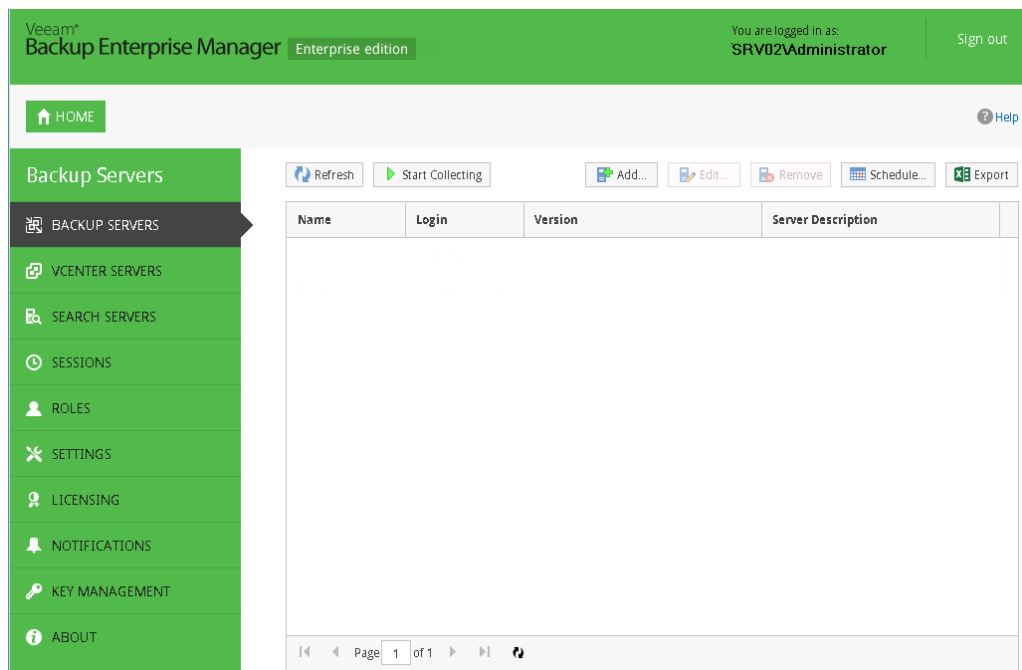
1. Download the latest version of Veeam Backup & Replication from www.veeam.com/downloads.html.
2. Start the autorun from the mounted ISO file and select to install Veeam Backup Enterprise Manager.



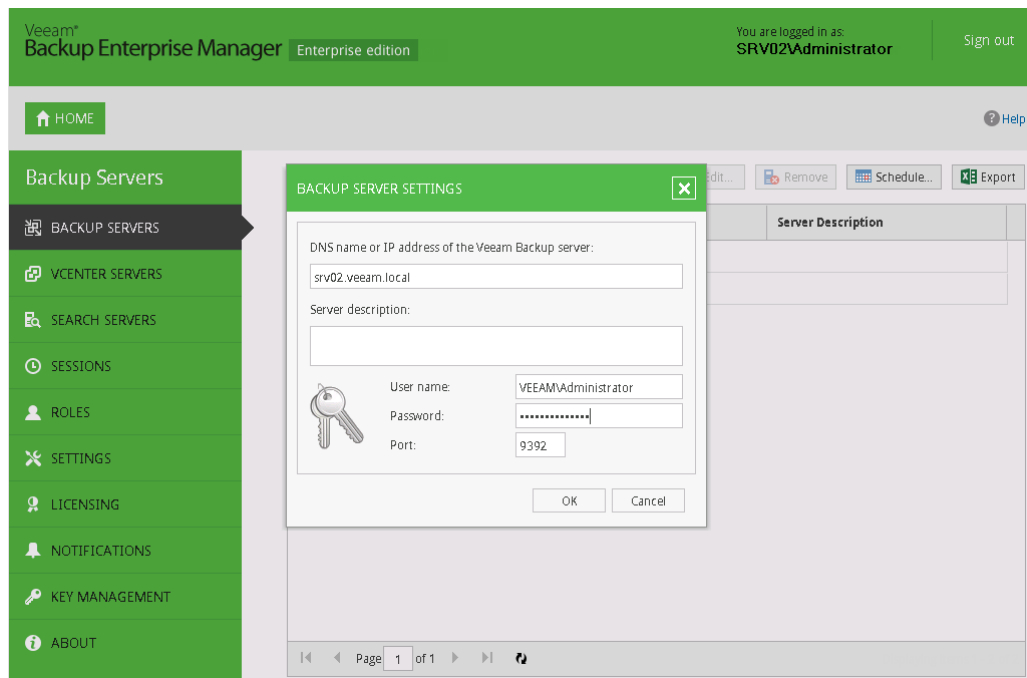
3. Follow the steps in the installation wizard.
4. At the **Service Account** step of the wizard, select to use the Local System account or specify settings of the account that has Administrator permissions on the Veeam Backup Enterprise Manager server.
5. At the **SQL Server Instance** step of the wizard, choose to install a new instance of Microsoft SQL Server or use an existing one (either local or remote). If you select to use an existing Microsoft SQL Server instance, for example, if you are installing Veeam Backup Enterprise Manager on the same machine as Veeam Backup & Replication, provide credentials of a user with database owner rights.

Once the installation process is complete, you can add the Veeam backup server(s) to Veeam Backup Enterprise Manager.

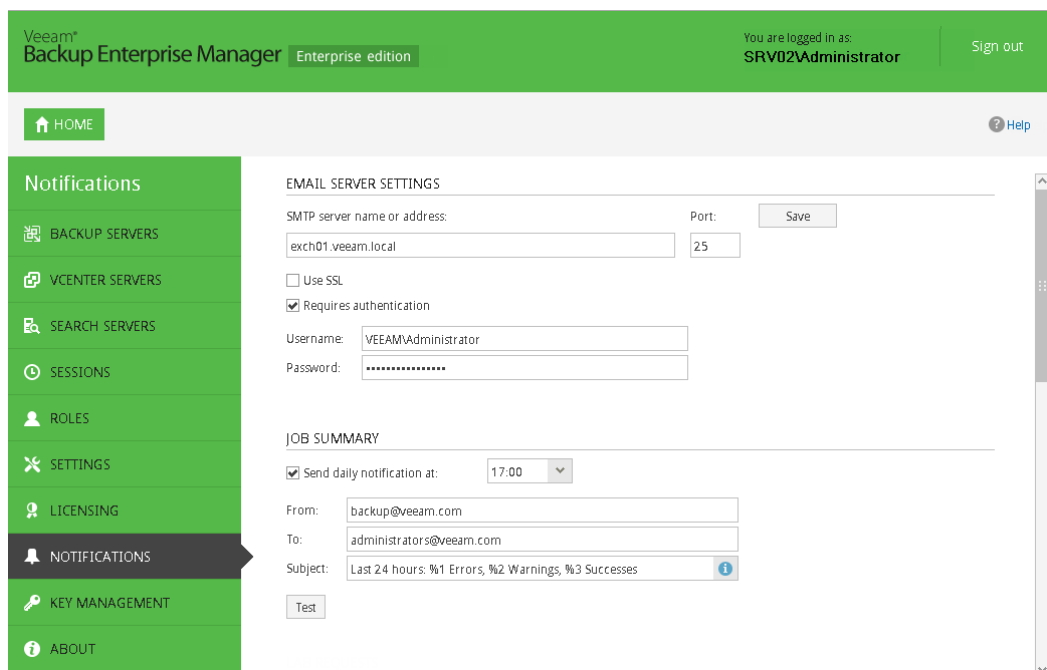
1. Start Veeam Backup Enterprise Manager by choosing **Programs > Veeam > Veeam Backup Enterprise Manager** from the **Start** menu. If you plan to use the web UI remotely, to access Veeam Backup Enterprise Manager, use the following HTTPS address: `https://host-name:9443`.
2. Enter credentials of a user with Local Administrator rights or the user who installed Veeam Backup Enterprise Manager and click **Login**.
3. Click **Configuration** at the top right corner of the window.
4. Click **Add** at the top of the window.



5. Enter a DNS name or the IP address of the Veeam backup server you want to add.
6. Provide a name and password of the user account with administrative rights on the added Veeam backup server. Make sure that this user is included in the *Veeam Backup Administrators* group on the Veeam backup server.

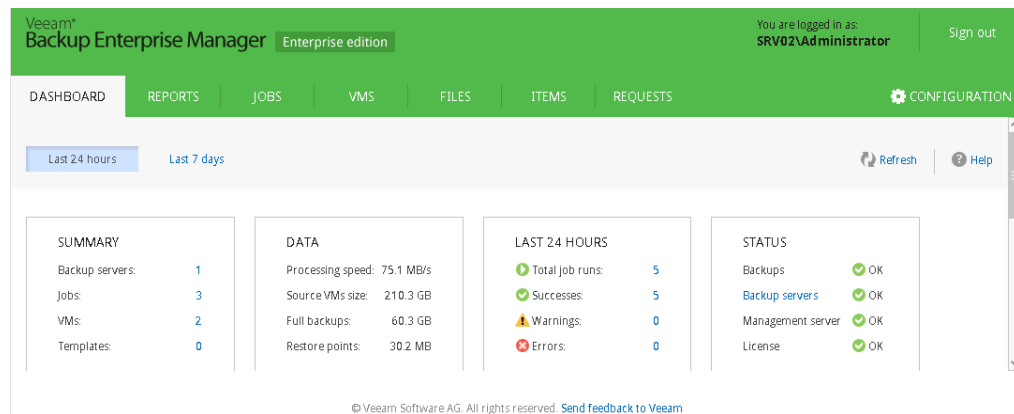


7. Once you click **OK** and add the Veeam backup server, Veeam Backup Enterprise Manager will automatically start collecting data about all backup and replication jobs from added Veeam backup server.
8. To receive consolidated email notifications about the status and summary of all jobs performed in your backup infrastructure, click **Notification** on the left of the **Configuration** view and specify email notification settings.



Validation

1. Click the **Last 24 hours** and **Last 7 days** tabs in Veeam Backup Enterprise Manager and make sure that these tabs provide information about backup and replication jobs performed on the Veeam backup server you have just added.



2. If you have configured email notification settings, open your email client and check the *Inbox* folder. Make sure that you have an incoming email about all performed jobs.

Searching for Guest OS Files and Performing 1-Click Restore

Veeam Backup Enterprise Manager allows you to search for Microsoft Windows and Linux guest files in backed up VMs. This can be useful, for example, if you need a file that has been deleted on the VM and you want to restore it from a backup. Once you find the file, you can immediately restore the file directly from Veeam Backup Enterprise Manager with 1 click. The file can be restored to its original location or saved to a local machine.

Evaluation Case

In this exercise, you will delete a file on the original VM, search for it in the backed up VM image and restore it to its original location.

Prerequisites

- Make sure that you have successfully created a backup of a Microsoft Windows or Linux VM with the guest file indexing option enabled.
- Make sure that you have successfully connected the Veeam backup server to Veeam Backup Enterprise Manager and collected data from it.

Procedure

To find a restore a guest OS file, do the following:

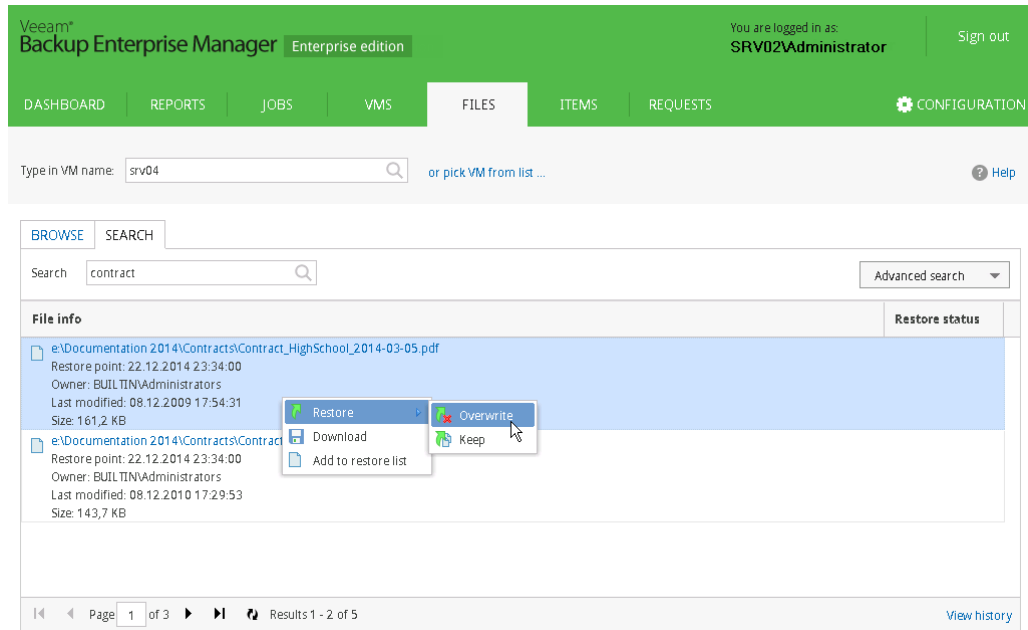
1. Delete a file you want to restore on the original VM.
2. In Veeam Backup Enterprise Manager, click **Home** at the top left corner of the window to get back to the main view of Veeam Backup Enterprise Manager.
3. In the main view of Veeam Backup Enterprise Manager, click the **Files** tab.
4. In the **Type in VM name** field, enter the name of a backed up VM whose file system you want to browse.
5. In the **Restore point** field, select a restore point from which you want to restore the deleted file.
6. In the **Search** field on the right, type in the name of the necessary file or a part of it and click the search icon on the right.

The screenshot shows the Veeam Backup Enterprise Manager interface. At the top, the header includes the Veeam logo, 'Backup Enterprise Manager Enterprise edition', and the user 'SRV02Administrator'. Below the header is a navigation bar with tabs: DASHBOARD, REPORTS, JOBS, VMS, FILES (selected), ITEMS, REQUESTS, and CONFIGURATION. The main area is divided into a left sidebar and a right pane. The left sidebar shows a file tree with folders like C:\, E:\, \$RECYCLE.BIN, Documentation 2011-2014, Contracts (selected), Invoices, Offers, Tax Docs, and System Volume Information. The right pane displays a table of search results for the file 'contract'.

Name	Date Modified	Size	Owner
Contract_Amber_2014-0...	17.11.2009 19:35:19	72,0 KB	Administrators
Contract_DFC_2014-01-1...	25.12.2009 17:22:34	93,3 KB	Administrators
Contract_HighSchool_20...	08.12.2009 17:54:31	161,2 KB	Administrators
Contract_Hosp_2014-05-...	26.01.2011 15:24:43	94,8 KB	Administrators
Contract_Uni_2014-04-1...	08.12.2010 17:29:53	143,7 KB	Administrators

At the bottom of the interface, there is a footer showing 'Page 1 of 1' and 'Displaying 1 - 5 of 5'.

7. Right-click the file in the list and select **Restore > Overwrite**. Click **Yes** to confirm the operation.



Validation

Open the primary VM and make sure that the found file has been successfully restored.

Performing Self-Restore of VM Guest OS Files

You can delegate a task of VM guest OS files restore to users having administrator rights on VMs, for example, to application owners. Delegated file-level restore simplifies the data restore process. Users do not have to wait for backup administrators to recover deleted or modified files and folders.

For delegated file-level restore, Veeam Backup & Replication offers the Veeam Self Service Restore portal. Authorized users can log on to the portal, browse the content of VM backups and restore the necessary VM guest OS files and folders to the original location or download restored files to the local computer drive. The restore process does not require any intervention from the backup administrator side: users can perform search and restore operations on their own, just like the administrator working with Veeam Backup Enterprise Manager.

Only authorized users can browse and restore files in the Veeam Self Service Restore portal. To be able to see the content of VM backups, the user must be added to the Local Administrators group in the VM guest OS. When the user logs on to the portal, the portal displays only those VMs and restore points that the user is permitted to access. Other VMs are not visible to the user.

Evaluation Case

In this exercise, you will perform self-service restore of VM guest OS files from the backup file using the Veeam Self Service File Restore portal.

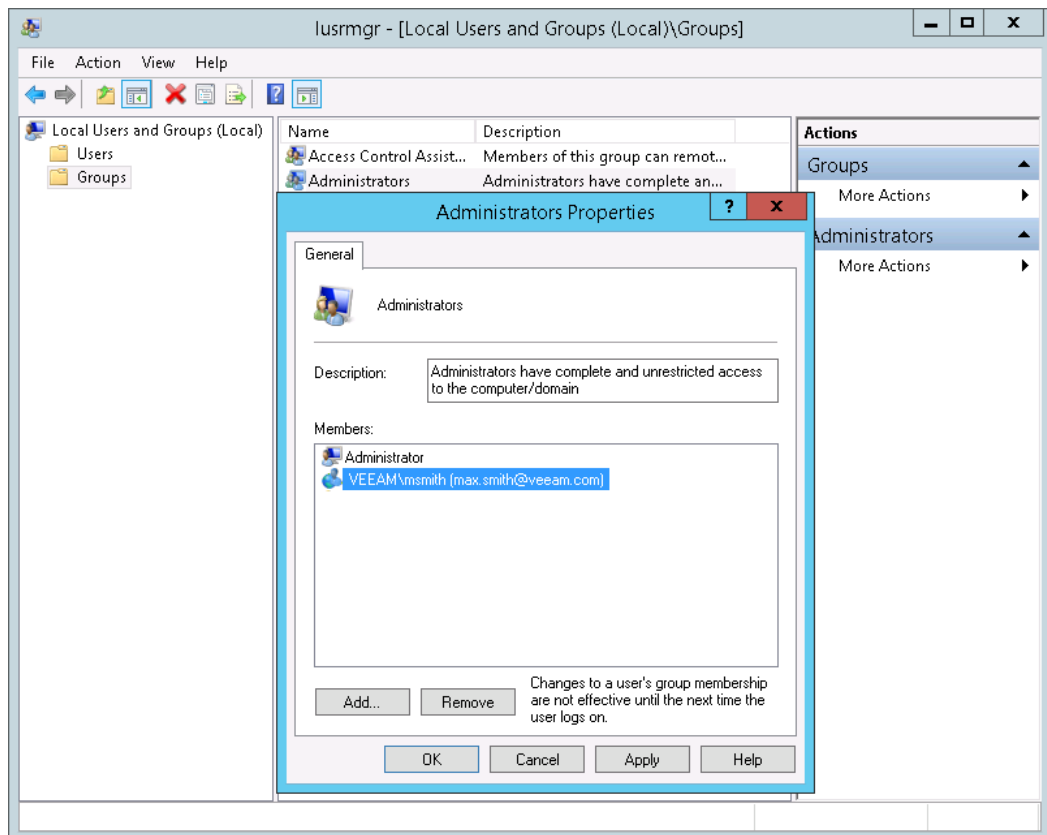
Prerequisites

- Make sure that the Veeam backup server is connected to Veeam Backup Enterprise Manager.
- Make sure that the Enterprise Plus license is installed on the Veeam Backup Enterprise Manager server. You can use a valid trial license or paid license.
- Make sure that 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 Veeam Self Service File Restore portal.

Procedure


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

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



2. In Veeam Backup & Replication, configure a backup job that processes this VM (see the [Performing Backup](#) exercise). At the **Guest Processing** step of the **New Backup Job** wizard, select the **Enable guest file system indexing** check box and specify a user account for file indexing.

New Backup Job



Guest Processing

Choose additional processing options available for Microsoft Windows guests.

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

☐ **Enable application-aware processing**

Quiesces applications using Microsoft VSS to ensure transactional consistency, performs transaction logs processing, and prepares application-specific VSS restore procedure.

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

Credentials: VEEAM\msmith (VEEAM\msmith, last edited: 11/11/ 2014) Add...

[Manage accounts](#)

Customize guest OS credentials for individual VMs and operating systems Credentials...

Test Now

< Previous
Next >
Finish
Cancel

- Run the backup job to produce a backup file.

Self Restore (Full)

Job progress:

0 of 1 VMs

Summary

Duration: 0:05:38

Processing rate: 15 MB/s

Bottleneck: Source

Data

Processed: 3.2 GB (38%)

Read: 3.2 GB

Transferred: 1.8 GB (1.8x)

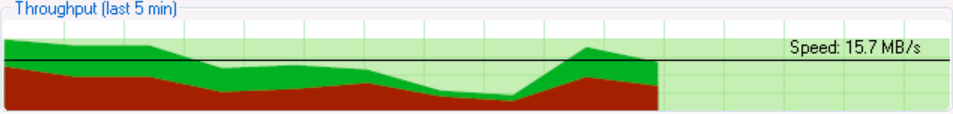
Status

Success: 0

Warnings: 0

Errors: 0

Throughput (last 5 min)



Speed: 15.7 MB/s

VM name	Status
srv04	38%

Action	Duration
Queued for processing at 11/11/2014 6:15:46 AM	
Required backup infrastructure resources have been assigned	
VM processing started at 11/11/2014 6:16:01 AM	
VM size: 60.0 GB (22.0 GB used)	
Inventorying guest system	0:00:04
Creating VM snapshot	0:00:02
Indexing guest file system	0:00:51
Saving [datastore4] srv04/srv04.vmx	
Saving [datastore4] srv04/srv04.vmx	
Saving [datastore4] srv04/srv04.nvram	
Using backup proxy VMware Backup Proxy for disk Hard disk 1 ...	
Hard disk 1 (40.0 GB) 3.1 GB read at 15 MB/s [CBT]	0:03:44
Using backup proxy VMware Backup Proxy for disk Hard disk 2 ...	
Hard disk 2 (20.0 GB) 98.0 MB read at 14 MB/s [CBT]	0:00:11
Getting list of guest file system local users	

Hide Details
OK

- Log on to Veeam Backup Enterprise Manager using the Administrator account.

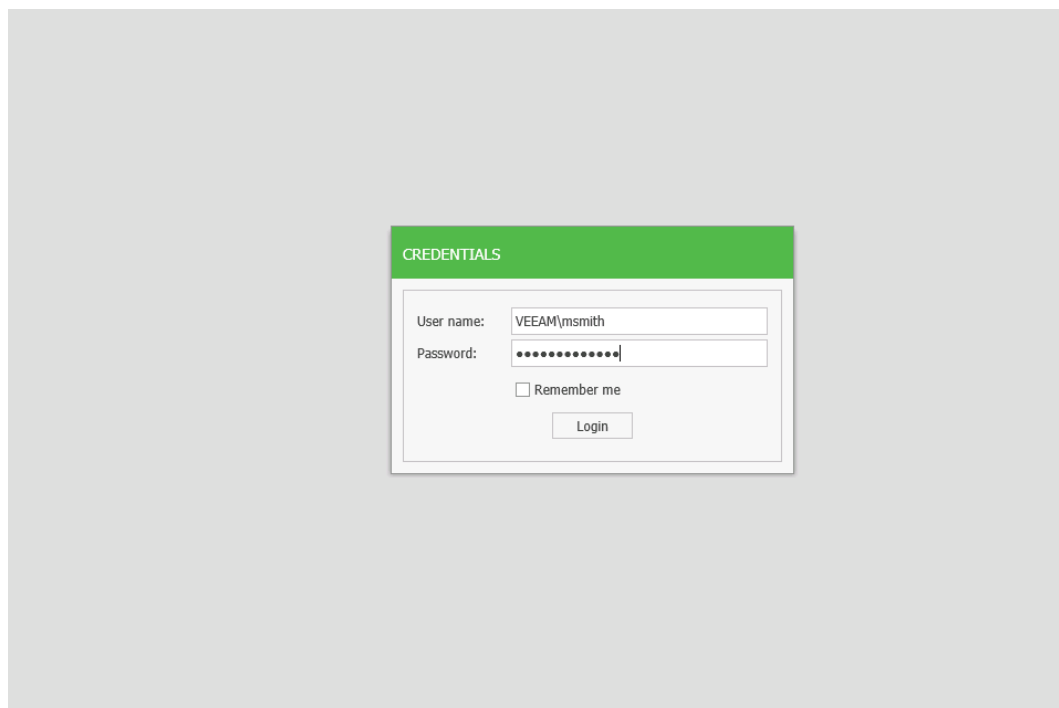
5. Open the **Configuration** view. In the **Backup Servers** section, click **Start Collecting** to collect the latest data about the performed backup job from the Veeam backup server. To check if data has been collected, click **Sessions** on the left and make sure that the data collection session has completed with the *Success* status.

The screenshot shows the Veeam Backup Enterprise Manager web interface. The top header is green with the Veeam logo and 'Enterprise edition' on the left, and 'You are logged in as: SRV02\Administrator' with a 'Sign out' link on the right. Below the header is a navigation bar with a 'HOME' button and a 'Help' link. The left sidebar is green and contains a list of menu items: 'Backup Servers' (selected), 'VCENTER SERVERS', 'SEARCH SERVERS', 'SESSIONS', 'ROLES', 'SETTINGS', 'LICENSING', 'NOTIFICATIONS', 'KEY MANAGEMENT', and 'ABOUT'. The main content area has a toolbar with buttons: 'Refresh', 'Start Collecting' (highlighted with a mouse cursor), 'Add...', 'Edit...', 'Remove', 'Schedule...', and 'Export'. Below the toolbar is a table with the following data:

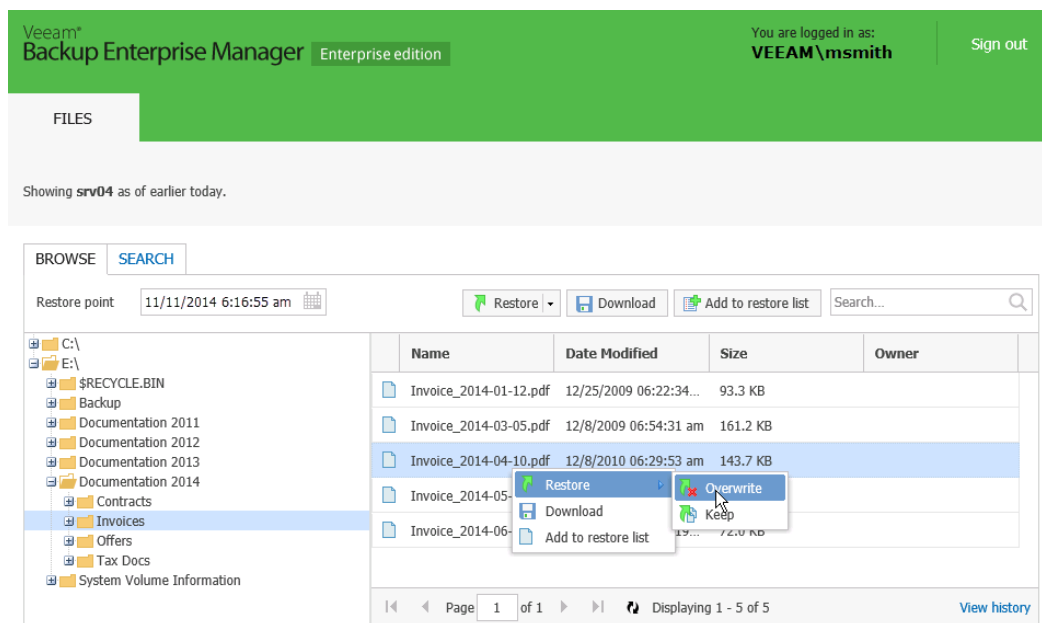
Name	Login	Version	Server Description
localhost	srv02\Administrator	8.0	

At the bottom of the table, there is a pagination bar showing 'Page 1 of 1' and a 'Displaying items 1 - 1 of 1' message.

6. Sign out from Veeam Backup Enterprise Manager.
7. Together with Veeam Backup Enterprise Manager, the Veeam Backup & Replication setup deploys the Veeam Self Service File Restore portal. Open the portal web console using the following link:
https://enterprise_manager_host:9443/selfrestore
8. Log on to the Veeam Self Service File Restore portal under the user account that you plan to use to restore VM guest OS files.



7. The portal will display only one tab — **Files**, with the file tree of the VM guest OS file system. To restore a specific file or folder, right-click it and select the necessary restore option from the list.



© Veeam Software AG. All rights reserved. [Send feedback to Veeam](#)

Validation

Open the file or folder restore location and make sure that the file or folder has been successfully restored.

Cloning and Editing Jobs

Veeam Backup Enterprise Manager lets you clone jobs configured on Veeam backup servers. When you clone a job, you create an exact copy of any backup or replication job available in the job list. Configuration details of a created job copy are written to the same SQL database where details of the original job are stored. You can work with the created job both via Veeam Backup Enterprise Manager and via the Veeam Backup & Replication console on a corresponding Veeam backup server.

The recommended practice is to configure a set of “job templates” in advance, using the Veeam Backup & Replication console on every managed Veeam backup server. Administrators working Veeam Backup Enterprise Manager can then use these “job templates” for cloning and further editing.

Evaluation Case

In this exercise, you will create a copy of the backup job and edit its settings.

Prerequisites

- Make sure that you have at least one backup job that has been successfully performed on the Veeam backup server.
- Make sure that you have successfully connected the Veeam backup server to Veeam Backup Enterprise Manager and collected data from it.

Procedure

To clone a job, do the following:

1. Click the **Jobs** tab.



Name	Type	Backup Server	Platfo...	Current ...	Latest Run	Next Run	Description
✓ Evaluation backup copy j...	Backup...	srv02	VMware	Success	11.11.2014 6:3...	Disabled	Pilot backup copy
✓ Evaluation replication job	Replica	srv02	VMware	Success	11.11.2014 6:1...	11.11.2014 22:...	Pilot replication
✓ Evaluation backup job	Backup	srv02	VMware	Success	11.11.2014 4:4...	11.11.2014 22:...	Backup job created for evaluation purposes

2. Select the necessary job in the list, click **Job actions** on the toolbar and select **Clone** job.
3. In the displayed window, specify the name of the job copy and select a target for the job.
4. Click **Clone**.
5. Select the cloned job in the list.
6. On the toolbar, click **Job actions** and select **Edit job**. Go through the steps of the wizard and edit the job settings as required.

EDIT REPLICATION JOB

Virtual machines

Select virtual machines to process:

Name	Type
 Windows VM	VirtualMachine
 Linux VM	VirtualMachine

Add

Remove

Exclusions

Up

Down

< Previous

Next >

Cancel

Validation

1. In Veeam Backup & Replication, open the **Backup & Replication** view.
2. Select the **Jobs** node in the inventory pane.
3. Make sure that the cloned job is available in the working area.

121 | Veeam Backup & Replication for VMware | EVALUATOR'S GUIDE | REV 2

Restoring Data from Encrypted Backup File Without Password

When you import an encrypted backup file on another Veeam backup server and try to restore VM data from it, Veeam Backup & Replication requires a password to unlock the backup file content. In most products, if you do not provide a password, the file content will remain locked, and the backup will be of no use. Veeam Backup & Replication lets you decrypt encrypted backups even if you have lost a password, or the person who knows the password has left your organization.

To enable data decryption without a password, Veeam Backup & Replication stores the actual encryption key in the backup file twice:

- First, Veeam Backup & Replication encrypts the key with a password that you set for the job.
- Second, Veeam Backup & Replication encrypts the key with a public key from Veeam Backup Enterprise Manager.

The public key from Veeam Backup Enterprise Manager plays the role of a 'duplicate key'. To decrypt the backup file, you can either provide a password or submit a request to Veeam Backup Enterprise Manager. Veeam Backup Enterprise Manager will process your request: it will apply a private key matching the public key that was used for backup file encryption. As a result, you will be able to access backup data in the Veeam Backup & Replication console.

Evaluation Case

In this exercise, you will create an encrypted backup and restore data from it without a password. To do this, you will perform the following steps:

1. Create an encrypted backup with a backup job.
2. To emulate a situation of data decryption on another Veeam backup server, remove the created backup from the Veeam Backup & Replication console and re-import the created backup back to the Veeam Backup & Replication console.
3. Decrypt the backup file without a password.

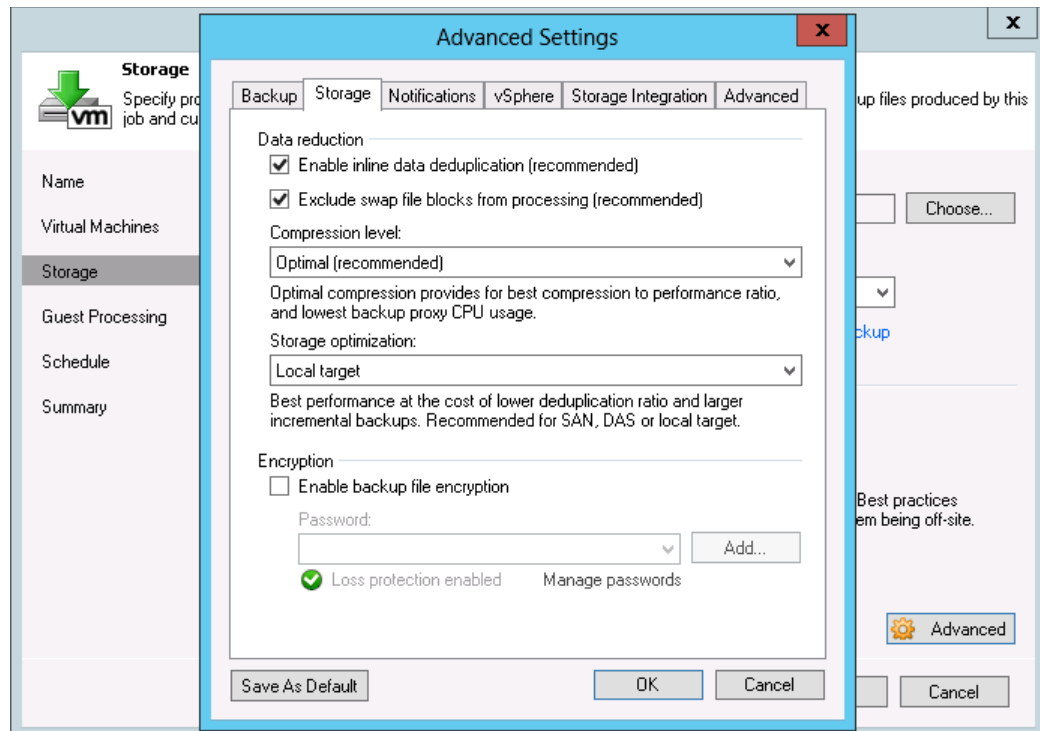
Prerequisites

- Make sure that the Veeam backup server is connected to Veeam Backup Enterprise Manager.
- Make sure that Enterprise or Enterprise Plus license is installed on the Veeam backup server. You can use a valid trial license or paid license.

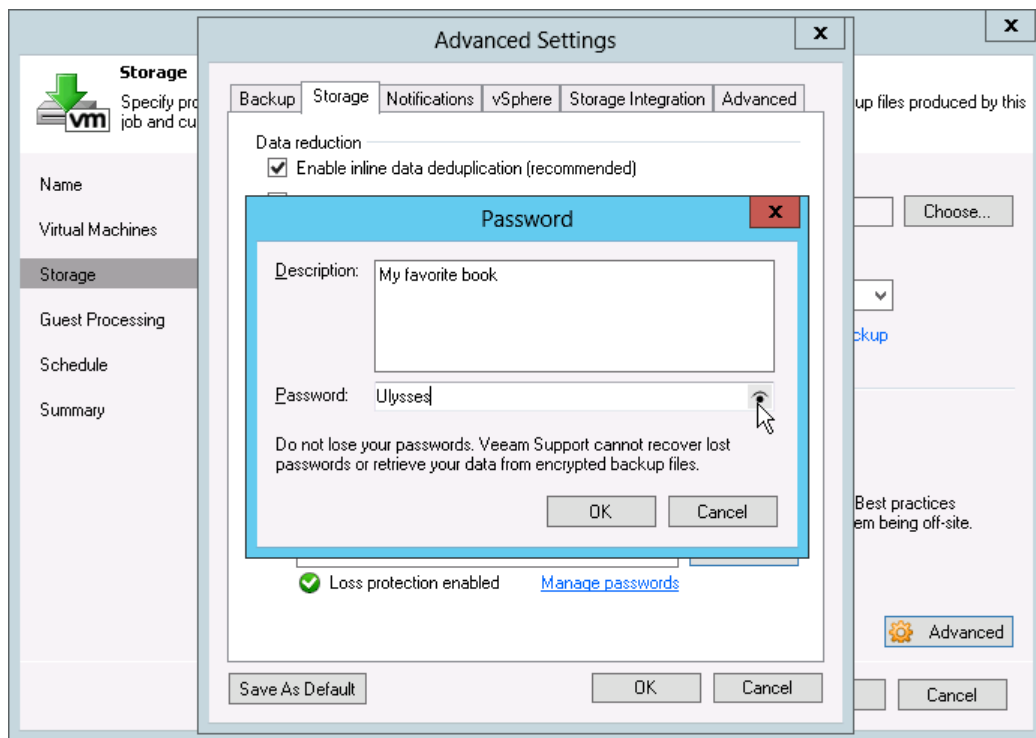
Procedure

Step 1. Create an encrypted backup

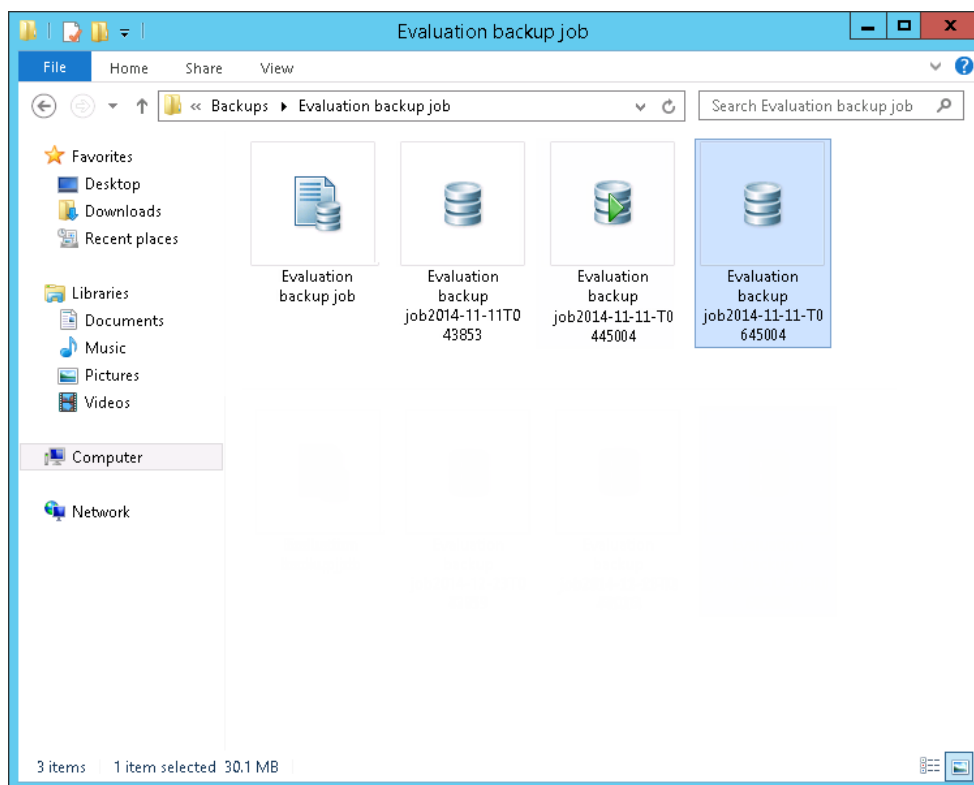
1. Open properties of a backup job that you have configured in the [Performing Backup](#) exercise.
2. Pass to the **Storage** step of the wizard and click **Advanced**.
3. In the **Advanced Settings** window, click the **Storage** tab.



4. In the **Encryption** section, select the **Enable backup file encryption** check box and click **Add** on the right.
5. In the **Password** field, enter a password that you want to use for the backup file encryption. To view the entered password, click and hold the eye icon on the right of the field.
6. In the **Description** field, enter a hint for the password.
7. Make sure that the **Loss protection enabled** label is displayed under the **Password** field. In the opposite case, you will not be able to restore data from the encrypted backup without a password.



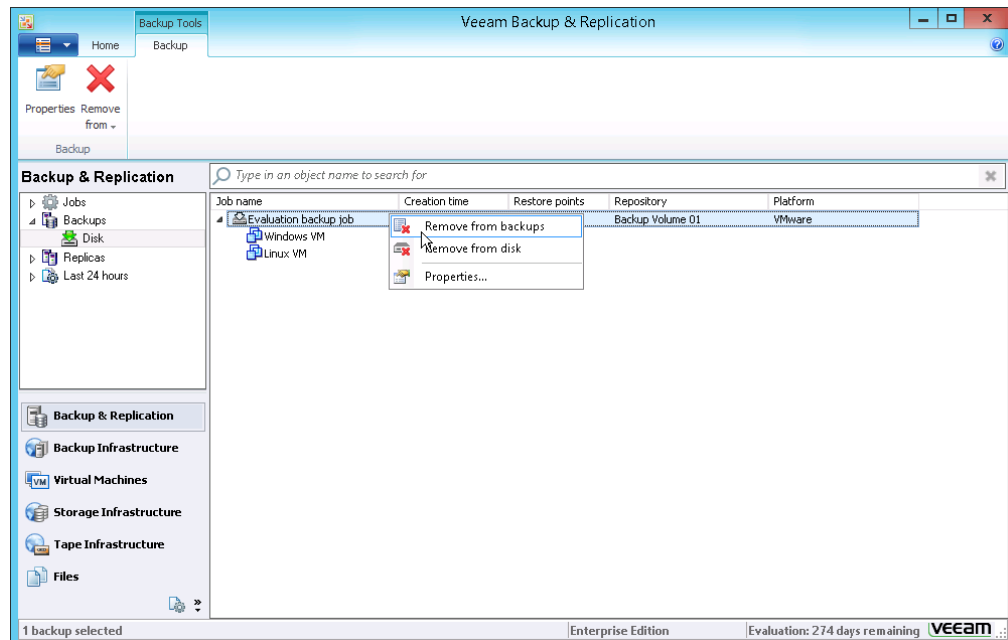
8. Save the new job settings and run the backup job once again to produce an encrypted backup file.
9. When you enable encryption for an already existing backup job, Veeam Backup & Replication restarts the backup chain — it produces a new full backup. To make sure that the encrypted backup has been created, open the target folder on the backup repository, find a subfolder with the backup job name and make sure that a new VBK file is added to the backup chain.



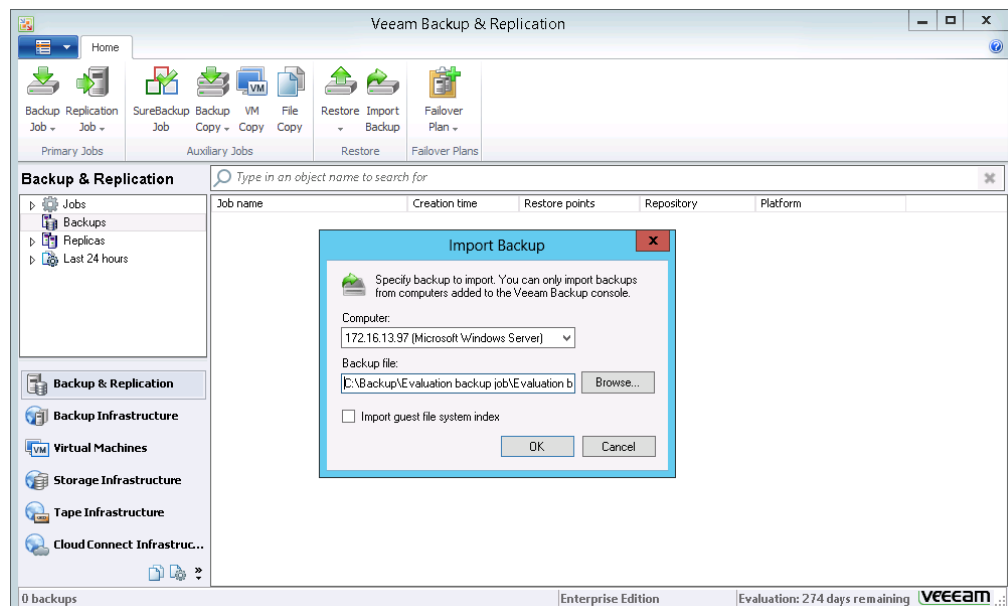
Step 2. Remove the backup from the console and re-import it

1. In Veeam Backup & Replication, open the **Backup & Replication** view.
2. In the inventory pane, select **Backups > Disk**.
3. In the working area, right-click the backup job and select **Remove from backups**.

Veeam Backup & Replication will remove records about the created backup and encryption keys from the Veeam Backup & Replication database. The actual backup files will remain on the backup repository.



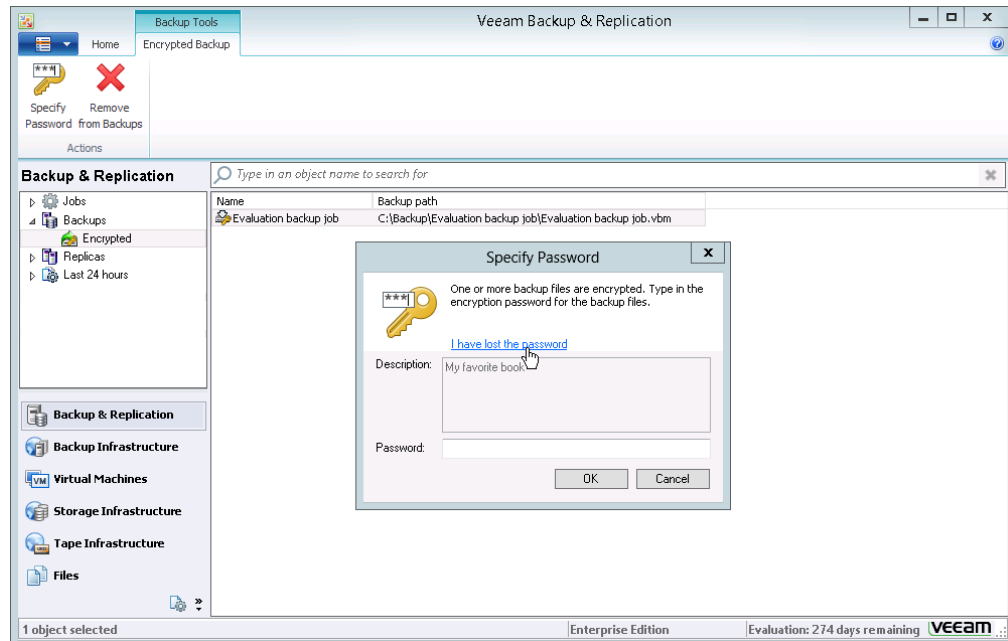
4. On the **Home** tab of the ribbon, click **Import Backup**.
5. From the **Computer** list, select a backup repository where backup files are located.
6. In the **Backup file** field, specify a path to the VBM backup file on the backup repository.



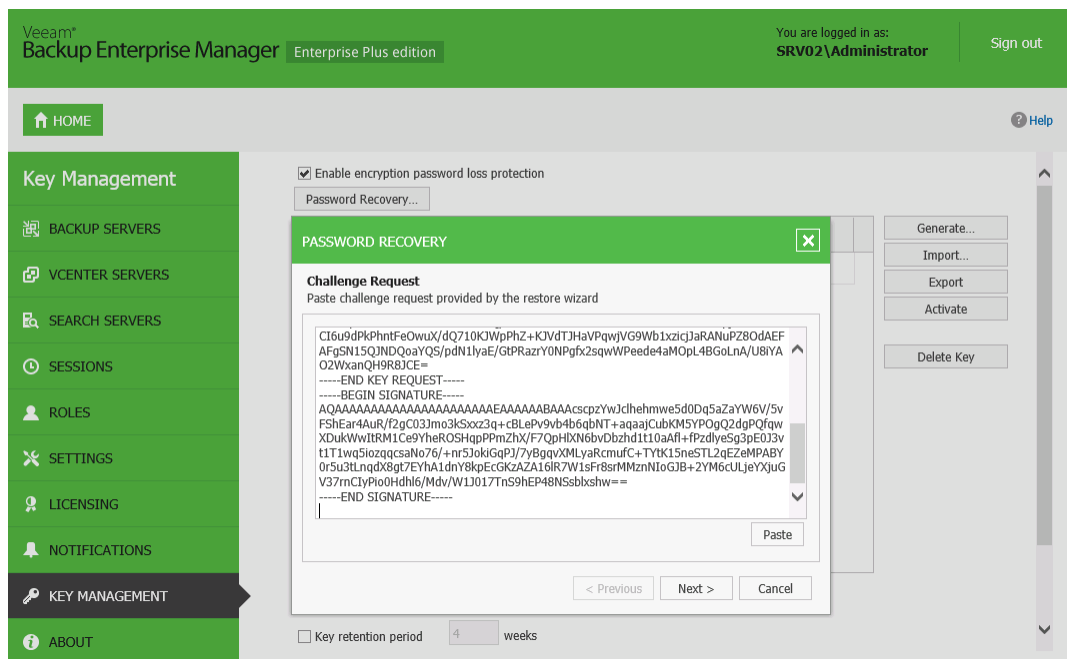
7. Click **OK**. Veeam Backup & Replication will import the backup and place it under the **Backups** > **Encrypted** node.
8. Additionally, Veeam Backup & Replication will display a warning that the backup file you import is encrypted. Click **OK** in the message window to close it.

Step 3. Decrypt the backup file without a password

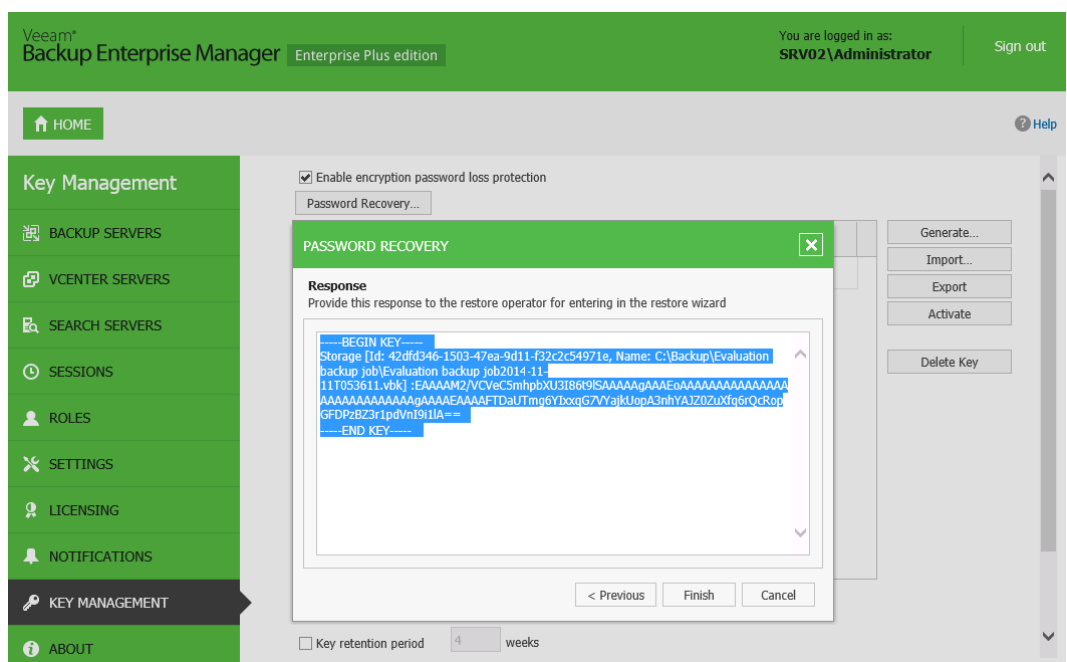
1. In the inventory pane, click the **Encrypted** node under **Backups**.
2. In the working area, right-click the imported job and select **Specify password**.
3. In the **Specify Password** window, click the **I have lost the password** link.



4. In the **Encryption Key Restore** wizard, click **Copy to clipboard** to copy the displayed request for data decryption.
5. Open Veeam Backup Enterprise Manager.
6. In the top right corner of the window, click **Configuration**.
7. Click **Key Management** on the left.
8. Click **Password Recovery** at the top of the view.
9. In the **Challenge Request** window, enter the copied text of the request.



- Pass through the next steps of the wizard. At the **Response** step of the wizard, copy the displayed text to the clipboard.



- Get back to the Veeam Backup & Replication console; in the **Encryption Key Restore** wizard, click **Next**.
- At the **Response** step of the wizard, enter the copied response to the text field and click **Next**. Veeam Backup & Replication will decrypt the backup file and move the imported backup to the **Backups > Disk (imported)** node.

Validation

1. Open the **Backup & Replication** view.
2. Select the **Disk (imported)** node in the inventory pane.
3. Make sure that the imported backup is available in the working area.

